



<![CDATA[Regular article]]>

Ranking institutions within a discipline: The steep mountain of academic excellence



Balázs R. Sziklai^{a,b}

^a Institute of Economics, Centre for Economic and Regional Studies, Hungary

^b Department of Operations Research and Actuarial Sciences, Corvinus University of Budapest, Hungary

ARTICLE INFO

Article history:

Received 30 October 2020

Received in revised form

23 December 2020

Accepted 8 January 2021

Keywords:

University departments

Ranking

Weighted Top Candidate method

Research discipline

ABSTRACT

We present a novel algorithm to rank smaller academic entities such as university departments or research groups within a research discipline. The Weighted Top Candidate (WTC) algorithm is a generalisation of an expert identification method. The axiomatic characterisation of WTC shows why it is especially suitable for scientometric purposes. The key axiom is stability – the selected institutions support each other's membership. The WTC algorithm, upon receiving an institution citation matrix, produces a list of institutions that can be deemed experts of the field. With a parameter we can adjust how exclusive our list should be. By completely relaxing the parameter, we obtain the largest stable set – academic entities that can qualify as experts under the mildest conditions. With a strict setup, we obtain a short list of the absolute elite. We demonstrate the algorithm on a citation database compiled from game theoretic literature published between 2008–2017. By plotting the size of the stable sets with respect to exclusiveness, we can obtain an overview of the competitiveness of the field. The diagram hints at how difficult it is for an institution to improve its position.

© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

According to a popular idiom, half the money spent on advertising is wasted; the issue is that we don't know which half. Likewise, it is difficult to quantify the efficiency of expenditure on research and education. Despite the well-known shortcomings of oversimplification and the skewness toward top institutions, university rankings have become increasingly important in the past decades as they offer a way, however inaccurate, to measure performance.

Institutions are seldom consistent in academic quality. Some departments or research groups have more talented staff whose work gets more recognition than others. Top-performing departments play an important role in an institution's success as they can provide competitive advantage over similar teaching or research programs due to their reputation. Such departments are often able to lobby for additional funding that's disproportionate to their size. The reallocation of resources necessitates performance measurement – the management needs to justify the money spent this way. There might be other reasons why a department or research group needs to be evaluated, such as to provide realistic goals for the staff, justify the existence of a newly formed group or assess the additional effort needed to improve its ranking.

University rankings are a measure of performance at the institutional level and are, therefore, unsuitable for measuring the success of a particular department. One way to resolve this is by evaluating departments based on the performance of the

E-mail address: sziklai.balazs@krtk.mta.hu

<https://doi.org/10.1016/j.joi.2021.101133>

1751-1577/© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

individual researchers. This may be problematic as different departments might not be comparable. For instance, a university of economic studies may accommodate both health economics and sociology departments. Researchers of the former publish in completely different journals than the latter, and these outlets usually have a higher impact factor and greater review speed than those of sociology. Some differences cannot be resolved by normalisation as certain disciplines prefer conference proceedings or monographs over journal publications. Another problematic issue in individual performance measurement is the separation of individual contributions from teamwork. Recently, Flores-Szwagrzak and Treibich (2020) proposed a method to disentangle individual productivity from the effect of coauthors.

Another approach is to evaluate departments based on their success in their field of research. For instance, if a computer science department mainly conducts research in machine learning, one can look at how the institution performs in that field. This approach also has some limitations. If a department has fragmented research interests, then the field that encompasses all the research areas might be too general, and the resulting ranking will not be significantly different from a university ranking. On the other hand, it may happen that there are more than one group or department that conduct research in the same approximate discipline. We can amend this by either looking at specific research areas or by evaluating at a higher level, although the latter decreases the ranking's consistency.

In this paper, we rank academic entities (henceforward *institutions*) based on the field of game theory. The emphasis is on the methodology rather than on the chosen field. We use this case study to introduce the Weighted Top Candidate (WTC) algorithm – a generalisation of an expert selection method (Sziklai, 2018). This algorithm has several advantages such as a sound axiomatic foundation with relevant properties,¹ visually expressive results that lend itself for interpretation and low computational complexity.

The Top Candidate (TC) method originates from the group identification literature and was designed to find expert groups on recommendation networks. The main component of TC is the so-called stability axiom. Stability asserts that experts are the most competent individuals to identify other experts. Remarkably, this circular logic can be implemented; in fact, spectral measures such as PageRank (Page, Brin, Motwani, & Winograd, 1999) and Eigenvector centrality (Bonacich, 1972) are based on similar concepts.

Stability can be decomposed into two elements: recognition (received citations) and recommendation (given citations). Only those institutions who are recognised by expert institutions can be deemed experts themselves. On the other hand, if an institution qualifies as expert, it can nominate (choose from the set of institutions they recommend) other experts. Depending on a parameter, some or all of the recommended institution must belong to the expert group. Parametrisation allows us to modulate how strict we want to be in our expert selection. Note, that unlike centrality measures group identification algorithms do not output a vector of real numbers signifying the importance of the institutions, but a list of institutions that are deemed important – technically, a vector of zeros (non-experts) and ones (experts). In Section 4, we provide a more detailed description of our algorithm and its characterisation.

WTC allows us to uncover the underlying structure of a citation network. If we set the parameter to be as inclusive as possible, we obtain the largest stable set, i.e. those institutions that can be considered experts under the mildest conditions. The other extreme of fixing the parameter to be as exclusive as possible reveals the core institutions that comprise the top of the field.

We constructed the citation network from Web of Science (WoS) data and use direct citations to identify the relevant papers. The principle behind direct citation is that two papers are related if either of them cites the other. In contrast, bibliographic coupling and co-citation analysis cluster papers by common documents they refer to or are referred by. We thus started from a core set of papers that form the main stream of game theoretic literature and looked for related papers. We required two citations in either direction to establish a connection. In addition, we also limited the scope of journals to exclude peripheral branches of the literature.

2. Literature overview

Ranking institutions based on their performance in a research discipline has been studied by a number of papers, although usually from a slightly different perspective.

Abramo and D'Angelo (2015) compared two bibliometric methods to measure the performance of universities – one based on the performance of the individual scientists and the other on that of the scientific fields present in the institution. For the latter, they analysed Italian universities active in nuclear and sub-nuclear physics. Shibata, Kajikawa, Takeda, and Matsushima (2009) investigated the performance of various citation analysis methods for detecting emerging research fronts. They considered three research domains, gallium nitride, complex networks and carbon nanotubes. They found that the direct citation method can detect large emerging clusters earlier than bibliographic coupling or co-citation analysis. In contrast, Boyack and Klavans (2010) compared the accuracy of cluster solutions using biomedical literature and found direct citation to be the least accurate mapping method. In this paper, we do not take sides in the debate; rather, we note that the Leiden University ranking also uses direct citations to define micro-fields (Waltman & van Eck, 2012). Dusansky and Vernon (1998) evaluate eight rankings of U.S. economics departments using four differing methodologies. They assess the strengths

¹ There are a few papers that offer axiomatic characterisations of centrality measures but the underlying properties are not always meaningful from the point of view of Scientometrics, cf. (Boldi & Vigna, 2014).

and weaknesses of the various approaches of which two is based on publications by faculty and two on faculty surveys. [Alma, Coşkun, & Övendireli, 2016](#) propose a field based ranking framework for Turkish Universities. The study's goal is to develop a set of indicators by integrating different perspectives on performance. They argue that general ranking lists might not paint a realistic picture, thus they focus on country specific circumstances.

Perhaps the closest to our work are that of [Zachos \(1991\)](#), [Lazaridis \(2010\)](#) and [Laengle, Merigó, Modak, Modak, and Yang \(2020\)](#). [Zachos \(1991\)](#) compared the mathematics department of two Greek universities. He also considered a 10 year time period, but used only basic scientometric indicators. [Lazaridis \(2010\)](#) used the h-index to rank university departments. He also argued that ranking departments gives a higher resolution picture of the distribution of quality within each university and could provide a strong motive for meritocratic hiring practices. [Laengle et al. \(2020\)](#) identify the most productive and influential research institutions in Operations Research and Management Science by taking into account the most influential journals.

University rankings often use complex score systems in which research excellence is but one factor. Centrality measures such as PageRank ([Page et al., 1999](#)) were suggested for citation analysis purposes ([Ma, Guan, & Zhao, 2008](#)). The Scimago Journal Rank (SJR) indicator is also a variant of PageRank ([Guerrero-Bote & Moya-Anegón, 2012](#)). The Eigenfactor project,² developed by [West, Bergstrom, and Bergstrom \(2010\)](#), is based upon another spectral centrality measure, the Eigenvector centrality ([Bonacich, 1972](#)). In this paper, we employ an expert selection method developed by [Sziklai \(2018\)](#). The idea of the algorithm originates from the group identification problem established in the seminal paper of [Kasher and Rubinstein \(1997\)](#).

Further, there are a few papers that focus on axiomatic study of bibliometric indices and methods. [Marchant \(2009\)](#) have presented an axiomatic characterisation of popular rankings, including the h-index. [Altman and Tennenholz \(2010\)](#) studied personalised ranking systems and trust systems. They adapt several axioms from the literature on global ranking systems and fully classify the set of systems that satisfy all of these axioms. [Bouyssou and Marchant \(2016\)](#) characterised fractionally counting citations that were suggested as a possible way to normalise citation counts between fields of research. [Wąs and Skibski \(2018a, 2018b\)](#) characterized the most popular spectral measures including PageRank and Eigenvector centrality.

3. Data

We constructed a citation network by identifying relevant papers of the research field and then extracting the institution cross-referencing data from this set. We looked at a 10-year period, from 2008 to 2017.

Game theory is a highly diverse field. Research directions extend to microeconomics, social choice and mechanism design, among others. Various branches of the discipline, such as evolutionary and combinatorial game theory, address completely different questions, and the related papers published on various platforms share no common reader base. Consequently, when assessing game theoretical research, we may commit two types of errors – either we include too much such as journals and papers that fall outside the main stream of game theoretical research or we include too little and miss key parts of the literature. There is no single best way to construct a citation network of a scientific discipline. What we opt for in this paper is a pragmatic approach. First, we identified two journals that play a central role in this field: Games and Economic Behavior (GEB) and International Journal of Game Theory (IJGT). These two journals compose *the core* of our citation database. Both journals are well-known and highly respected outlets among researchers of the field. There are others that focus mainly on game theoretical research, including International Game Theory Review, Dynamic Games and Applications, Mathematical Social Sciences, Social Choice and Welfare and MDPI Games. The reasons why we kept the core relatively small is that the two chosen journals

³ are the (only) official journals of the Game Theory Society.³

⁴ publish solely game theoretical research.

The core contains only a fraction of the game theoretical literature. The other parts are scattered in numerous journals that belong to different fields. Adding more journals to the core of our database would undermine its integrity as these journals' scope is typically not restricted to game theory. We decided to define the missing parts in relation with the core. We constructed two sets: the *Ancestor* and *Descendant* set which contain papers from relevant journals.

A journal is considered *relevant* if it cites *and* is cited by the core at least 10 times. In other words, the set of articles published by a journal during the 10-year time frame should refer the papers in the core at least 10 times in total. In addition, the papers in the core should refer the papers published by the journal at least 10 times in total. The logical 'and' guarantees that the journal is part of the discourse and that it is not one-sided communication. There were also exceptions; for instance, International Game Theory Review was not cited enough times by GEB and IJGT, but since the papers published there are relevant, we decided to include the journal in the analysis. Similarly, some journals of mathematical nature were added to the relevant set, although they didn't qualify by the threshold.

² eigenfactor.org

³ The Society as the organiser of the World Congress, is a central institute of game theoretical research.

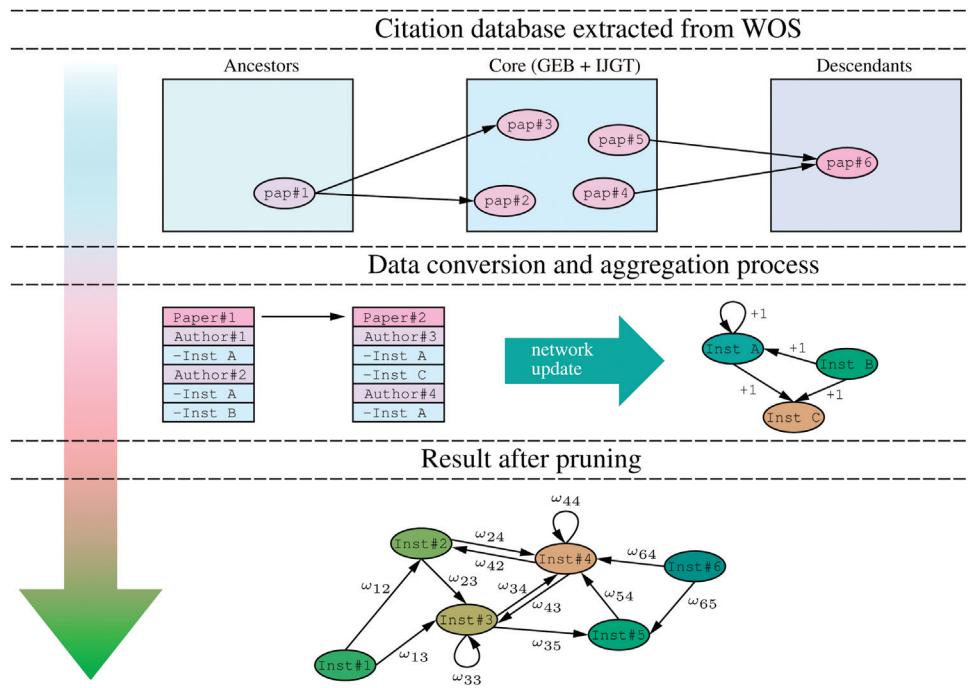


Fig. 1. Data aggregation, procession and network design.

We found 421 journals that cited or were cited by the core at least twice. Although the 10 papers inclusion criterion seems to be relatively lax, only 33 among the 421 journals satisfied it. Based on the journals' scope, we identified 24 additional journals that were potentially relevant. We added 19 of them based on their content – mainly by the published papers' titles and abstracts. In total, 3804 articles published in 52 different journals were considered.

A paper qualified for the Ancestor set if it was obtained from a relevant journal and cited *at least two* distinct papers in the core. Similarly, a paper qualified for the Descendant set if it was from a relevant journal and is cited by *at least two* distinct papers from the core. The top part of Fig. 1 describes the process. The Ancestor, Core and Descendant sets contained 1185, 1730 and 889 papers respectively.

Having identified the admissible papers, we were now ready to construct the institution citation network. Our aim was to create a weighted directed graph, wherein the nodes represent institutions and arc weights represent the number of times the source institution cited the sink institution. There is more than one sensible way for how this can be done. Suppose a paper written by two authors from Institute A cites a paper written by an author affiliated to Institute B. Should this be counted as one or two references? A similar question arises when an author from Institute A cites a paper written by two authors affiliated with Institute B.

We decided that irrespective of the number of authors from the same institution on either of the citing or the cited paper, one citation should increase an edge weight only by 1. However, one citation can increment more than one edge weights if there were more than one institutions involved in one of the papers. The middle part of Fig. 1 demonstrates the calculation. Note that only those citations were considered where both the citing and the cited paper was a member of either the Ancestor, the Descendant or the Core set. References citing non-game theoretical papers were discarded.

As a final step, we had to decide whether to keep self-citations or not. Researchers commonly cite their own works, so larger departments will tend to produce more self-citations entirely by their size. On the other hand, removing self-citations would harm more productive researchers. In addition, self-citation on an institution level does not imply that there is a self-citing author as the citation can come from a colleague of the same institute. Although we opted to keep self-citations, we note that removing them would have also been a valid choice. The drawback of keeping them is that we had to prune the data: an institutions that only cites itself forms a stable component, thus WTC identifies it as an expert (cf. Section 4). Fortunately, such anomalies are rare, and in our citation network we only found one such institution at the periphery of the network. While pruning, we removed its self-citing edge.

The final network constituted 1,002 nodes (institutions) and 23,725 directed edges (references) with a total weight of 41,919.

We had to make a few choices in determining the relevant set of papers. Arguably, the design could be improved. There was a trade-off between the sophistication of the clustering mechanism and efficiency. Direct citation could be replaced by a hybrid method involving bibliographic coupling and keyword analysis. It would be interesting to compare the obtained set

to the micro-field no. #1111 of the Leiden ranking.⁴ Both the journals (GEB, IJGT, Math. Soc. Sci., Soc. Choice and Welf., Lect. Notes in Comp. Sci.) and the keywords (core, shapley value, strategy proofness, cooperative game, judgement aggregation) listed to this micro-field seem to be highly relevant to our targeted set of papers. In comparison, micro-fields no. #1716 and #2833, which are centered around auctions and traffic routing respectively, seemed to show some overlap in journals and keywords as game theorists contributed to both fields significantly. As we said earlier, classification is rather fluid and determining what constitutes as mainstream game theoretical research is a matter of taste. Our aim in this paper is to demonstrate the advantages of the Weighted Top Candidate algorithm, thus we content ourselves with the obtained set.

4. The Weighted Top Candidate algorithm

Some rankings can be constructed on an objective criteria. We can organise competitions to determine who is the best chess player. Beauty contests, on the other hand, are highly subjective and the results express trends and people's preferences rather than the objective truth. In between these two extremes are questions that cannot be decided by competitions and are not even entirely subjective. "Who is the best game theorist?" is one such question. In general, expert identification falls into this category.

We aim to use institution cross-referencing data to uncover the most prestigious institutions in the field of game theory. The WTC algorithm allows institutions to nominate one or more institutions into the expert set. In the beginning, every institution is part of the set, and then we iteratively remove those who are not nominated by anyone from *within the set*. The key assumption is that nominations are not equally valuable, and experts are much more efficient in recognising other experts. Analogously, getting cited by a Nobel-prize winning researcher in a top journal is worth more than getting cited by a PhD student in a second-rate journal. Note that this is not a judgement on the student or the journal – they just might be at the beginning of their journey.

In this section, we make use of the following notations. The institution citation network is represented by a directed graph, $G = (V, E)$, where V denotes the set of institutions (nodes) and E denotes the set of references (directed edges) between institutions. Each reference $e = (\mathbf{u}, \mathbf{v})$ has a weight $w_e \in \mathbb{N}$ that represents the number of times Institution \mathbf{u} cited Institution \mathbf{v} . This also implies that there are no parallel edges; however, loops (self-referencing institutions) are possible. Let $R(\mathbf{v})$ denote the set of institutions referred by \mathbf{v} , that is $R(\mathbf{v}) = \{\mathbf{u} \in V \mid \exists (\mathbf{v}, \mathbf{u}) \in E\}$. Let $L(\mathbf{v})$ denote the set of references given by \mathbf{v} . The *reputation* of an institution is the total number of citations it receives, formally $r_{\mathbf{v}} = \sum_{\{e \in E | e = (\mathbf{v}, \mathbf{u})\}} w_e$. The *top candidate* of an institution \mathbf{v} is the most reputable institution it refers to; that is, the institution \mathbf{u} which has the highest $r_{\mathbf{u}}$ value among institutions that are cited by \mathbf{v} . We denote by $\omega_{\mathbf{v}}$ the weight of the top candidate of institution \mathbf{v} , formally $\omega_{\mathbf{v}} = \max_{\{\mathbf{u} \in R(\mathbf{v})\}} r_{\mathbf{u}}$.

Alternatively, we can consider an institution citation matrix $[a_{ij}]_{n \times n}$, where n denotes the number of institutions and a_{ij} denotes the number of times Institution i cites Institution j . In this way, reputation of Institution j is the sum of weights in column j , formally $\sum_{i=1}^n a_{ij}$.

Algorithm 1. Weighted Top Candidate method

```

1:            $N \leftarrow V, Z \leftarrow \emptyset, b = \text{true} // \text{Initialisation}$ 
2:           while  $b$  do
3:             for all  $\mathbf{v} \in N$  do
4:               for all  $(\mathbf{v}, \mathbf{u}) \in L(\mathbf{v})$  do
5:                 if  $\omega_{\mathbf{v}}(1 - \alpha) \leq r_{\mathbf{u}}$  then
6:                    $Z \leftarrow Z \cup \mathbf{u} // \text{if } \mathbf{u} \text{ is nominated, put it in } Z$ 
7:                 end if
8:               end for
9:             end for
10:            if  $N == Z$  then
11:               $b = \text{false} // \text{if everyone in } N \text{ was nominated, we stop}$ 
12:            else
13:               $N = Z, Z = \emptyset // \text{otherwise we continue with the nominated agents}$ 
14:            end if
15:          end while
16:          return  $N$ 

```

4.1. Algorithm

The Weighted Top Candidate algorithm operates on a set N , which in the beginning contains all the institutions. With each iteration, the algorithm discards institutions that are not nominated by anyone in N . Nomination is controlled by a parameter α : Each institution \mathbf{u} nominates (in terms of reputation) the top α fraction of the institutions that it cites (let us call these α -top candidates of \mathbf{u}). The obtained set – which might be the empty set – is stable in the sense that each institution is nominated (vouched for) by some institution in the set. For a formal description see Algorithm 1.

⁴ Information on the micro-fields of the Leiden Ranking is available in an Excel file downloadable from <https://www.leidenranking.com/information/fields>.

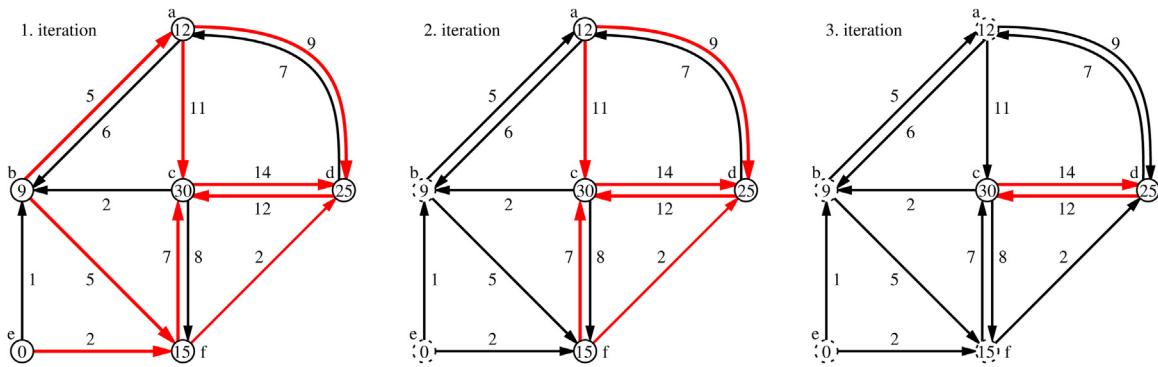


Fig. 2. An example of WTC computation with $\alpha=0.2$. Slightly thicker, red edges represent nominations. Node weights show the reputation of the corresponding institution. Circles with broken lines signify that the institution was not nominated in the previous iteration.

4.2. Example

Fig. 2 shows a simple example of the WTC algorithm with $\alpha=0.8$. Institution **a** nominates **c** and **d**. The former because it is the most reputable cited institution of **a**, the latter because

$$\omega_a \cdot (1 - \alpha) \leq r(d) \quad \text{that is } 30 \cdot 0.8 \leq 25.$$

Similarly, **b** and **f** also nominate two other institutions. Institutions **c**, **d** and **e** have only one candidate. In the second iteration, **b** and **e** are discarded from the set of experts as they are not nominated by anyone. As a consequence, their nominations are cancelled. Note that only the nominations of **b** and **e** are removed from the network, the nodes themselves are not. Thus, the reputations are unaffected and the remaining institutions still nominate the same agents. Since institutions **a** and **f** are not nominated by anyone, they are removed from the expert set in the third iteration. The remaining two institutions **c** and **d** nominate each other, thus they form a stable component. If we increase α to 0.4, **f** also becomes a member of the expert set as **c** nominates it besides **d**. At this level of exclusivity, **e** also nominates **b** but to no avail: **e** cannot become part of any stable component since it is not nominated by anyone at any level of α .

4.2.1. Characterisation

WTC is characterised by three appealing properties: stability, exhaustiveness and decisiveness. Here, we only describe the axioms informatively. Definition and a formal proof of the non-weighted case can be found in (Sziklai, 2018). The characterization carries over to the weighted case in a straightforward manner.

Stability requires that (i) each expert should be nominated by an expert and (ii) the nominees of each expert should belong to the expert set.

Note that different nomination processes produce different stability notions. Institutions in the WTC algorithm nominate their α -top candidates, i.e. α fraction of the most prestigious institutions among the ones they cite. Thus, WTC is stable with respect to the α -top candidate relation.

Exhaustiveness implies that the algorithm identifies every relevant institution. Suppose some members of the network cannot be elected, e.g., due to conflict of interest. An exhaustive algorithm cannot find new experts when the previously selected institutions are marked as non-elective and we re-run the algorithm.

An algorithm that always returns the empty set is both stable and exhaustive. Thus, we need some kind of existence axiom to ensure that the algorithm selects somebody when there are reasonable candidates.

Decisiveness is a guarantee that the algorithm selects somebody when there exists at least one elective institution of a set that is stable with respect to the nomination process.

A group identification method takes a weighted directed network (or a citation matrix) as input and outputs a list of nodes (rows/columns).

Theorem 1. A group identification method is stable with respect to the α -top candidate relation, exhaustive and decisive with respect to the α -top candidate relation, if and only if it is the Weighted Top Candidate method.

A proof follows word by word the non-weighted case, see (Sziklai, 2018). One way to see this is to reduce the problem to the non-weighted case by representing the weighted edges with parallel edges. This is always possible since weights in our model are positive integers.

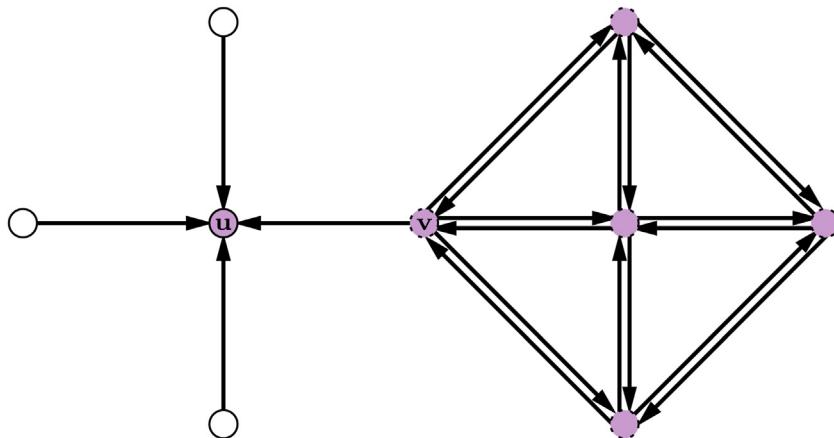


Fig. 3. The difference between k -core and WTC. Each edge weight is set to 1. Colored circles represent the institutions chosen by WTC, while circles with broken lines represent the institutions that belong to the k -core for $k=3$.

5. Benchmark methods

We calculated four alternative measures to serve as comparative benchmarks. Despite its simplicity, **degree** centrality has proven to be a very good indicator of performance. Here, we calculated the weighted in-degree of the nodes, that is, the reputation of the corresponding institutions.

Harmonic centrality was proposed by Marchiori and Latora (2000) to overcome the limitations of Closeness centrality. Harmonic centrality of a node, v is the sum of the reciprocal of distances between v and every other node in the network. For disconnected node pairs, the distance is defined as zero. If a node lies on the periphery, then the distances from most of the other nodes will be large. Thus, the reciprocal of the distances will be small, yielding a small centrality value. In directed networks, it is often more meaningful to work with a graph where the direction of the edges is reversed (if there are many paths leading to a node, then it lies in the center irrespective of how many path begins from that node). Consequently, we reversed the edges when we computed Harmonic centrality.

PageRank is a spectral measure that models an infinite random walk. The PageRank scores indicate how likely it is that the walker occupies a certain node. In our setting, this translates to browsing the game theoretical literature, jumping from one paper to another in a random manner and asking what the probability is that the next paper is written by someone from a given institution. However, a simple random walk entails the following problem: what do we do with sink nodes (institutions that do not cite anybody) or when the walk enters an inescapable component of the graph (when we enter into a group of institutions that cite only themselves)? To amend this, PageRank connects sink nodes with every other node through a link and redistributes some value uniformly in each iteration. The latter is parameterised by the so-called ‘damping factor’ $d \in (0, 1)$. The damping factor is most commonly chosen from the interval $(0.7, 0.9)$; here, we opted for $d=0.8$. PageRank is a core element of Google’s search engine, but the algorithm is used in a wide variety of applications. In particular, the Scimago Journal Rank (SJR) indicator is also based on PageRank (Guerrero-Bote & Moya-Anegón, 2012).

k -core, also known as k -shell, exposes the onion-like structure of the network (Kitsak et al., 2010; Seidman, 1983). First, it successively removes institutions that were only cited by one or less institution in the network. These are assigned a k -core value of 1. Then it removes institutions with two or less citing institutions and labels them with a k -core value of 2. The process continues until every node is classified. The definition of k -core and its variants resemble to that of WTC, however, as Fig. 3 highlights they are not the same. Suppose we would like to determine the k -core of this network for $k=3$. The supporters of Institution u – with the exception of Institution v – are removed one by one. Eventually Institution u is also removed. In the WTC computation, no Institution or citation is discarded. Non-expert institutions only lose their ability to nominate. As a result, Institution u is also deemed as an expert. It would be interesting to see whether there is a nomination process under which k -core is stable.⁵

The k -core method was developed for non-weighted networks. There is more than one way how this procedure can be generalised to weighted networks. We follow Garas, Schweitzer, and Havlin (2012)’s suggestion and compute the so-called **Weighted k -shell**, which instead of the number of citing institutions considers the square root of the product of the number of citing institutions and the total number of citations received by the institute. For example, if Institution a is cited once by Institution b and twice by Institution c , then the weight of Institution a is $\sqrt{2 \cdot (1+2)} = \sqrt{6}$. In comparison, if Institution a receives 3 citation from one source it obtains a weight of $\sqrt{3}$. That is, the Weighted k -shell method favours institutions whose citations come from diverse sources.

⁵ Note that, majority voting is not stable for any nomination process (Sziklai, 2018).

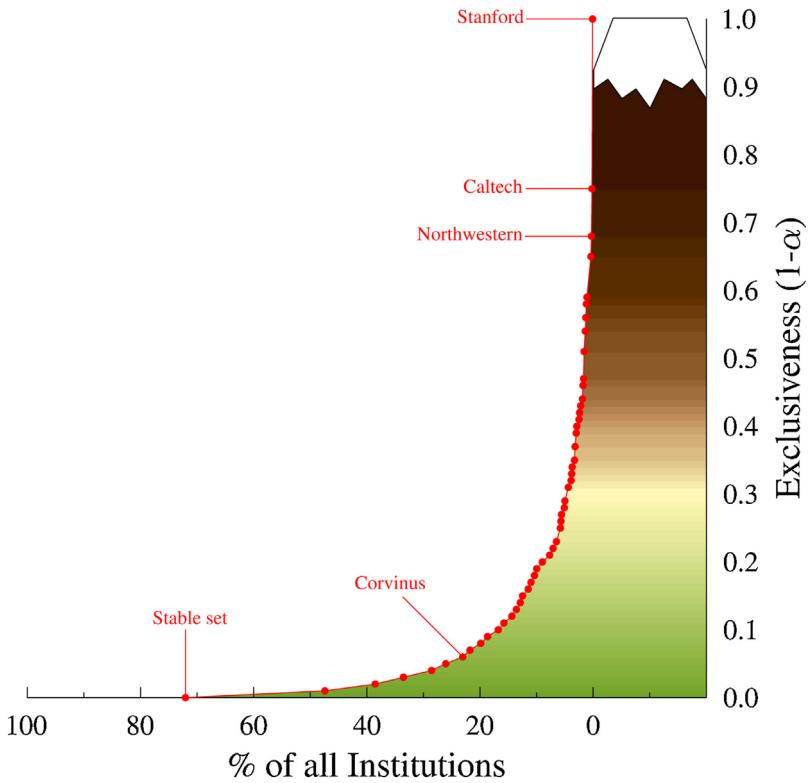


Fig. 4. The steep mountain of academic excellence. As we relax the exclusiveness parameter more and more institutions are selected by the WTC algorithm.

6. Results

The output of the WTC algorithm is a set of experts. This set is stable, thus the membership of each institution is endorsed by someone from the set. By a parameter, we can adjust how exclusive the list of experts should be. To obtain a complete picture, we ran WTC under 101 different parameter setting from 0 to 1 with an increment of 1 percent point.

Under the most relaxed setting ($\alpha = 1$), we obtain the largest stable set – those institutions that can be considered as experts in any sense. Already, 28% of the institutions drop out. We cannot enlarge the expert set in such a way that each member receives a nomination from within the set. The other extreme ($\alpha = 0$) is occupied by one institute. Stanford forms a stable set in itself because it nominates only itself under this parameter setting. Exclusiveness has to drop an astounding 25% to include another institution. Even at ($\alpha = 0.41$), only 1% of the institutions belong to the set of experts. If we plot the fraction of institutions with respect to α , an incredibly steep mountain starts to shape (see Fig. 4).

Corvinus University of Budapest – home affiliation of the author – features a Mathematical and Statistical Institute where game theoretical research is traditionally strong. According to the ranking by the WTC analysis, the university is ranked around the top 25%. With some additional effort and aimed recruitment, it could certainly get into the top 20%. Incidentally, getting into the top 200 institutions in global rankings is a dedicated goal of the recent reforms that were initiated at Corvinus.

However, further improving the university's position seems difficult. The slope of the curve starts to increase dramatically around 15%. The top 5% is like a vertical wall, a tiny advancement would need a significant rise in received citations. It would be interesting to see whether other disciplines have a similar WTC curve.

Tables 1 and 2 compare the WTC ranking to some benchmark rankings induced by well-known centralities. **Table 1** lists the top 11 institutions.⁶ With the exception of Harmonic centrality the top positions of the rankings largely overlap with each other, indicating that there is a consensus between the measures. Considering the differences between the number of received citations, this is hardly surprising.

To obtain a complete picture, we also calculated the distances between the rankings of the expert institutions. **Table 2** displays the normalized sum of ranking differences (nSRD) between the rankings. The nSRD score is the Manhattan distance divided by the maximum possible distance between two rankings, thus the values range from 0 to 1. Here the ranking's size is 722 as this many institutions belong to the stable set. Note that the expected distance between two random rankings of

⁶ For $\alpha = 0.40$ there are only four institution in the expert set. Increasing the parameter by one point adds seven more institutions. This can be interpreted as a seven-way tie for the fifth place. Since WTC outputs a zero-one vector for every parameter setting, ties are common.

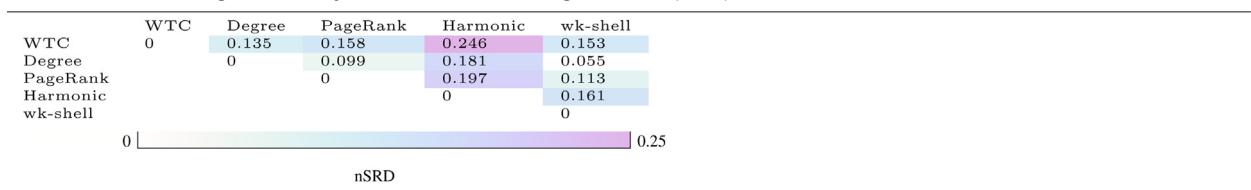
Table 1

Top academic research centers in game theory according to the Weighted Top Candidate algorithm and their ranking according to different centralities. In case of ties, fractional ranking is used, that is, the arithmetic average of the tied ranks is displayed.

	Weighted Top Candidate	Weighted in-degree	PageRank $d=0.8$	Harmonic centrality	Weighted k-shell
Stanford Univ	1	1	1	49	1
CALTECH	2	2	3	7	2
Northwestern Univ	3	3	2	3	3.5
Harvard Univ	4	4	4	41	3.5
Chapman Univ	8	11	32	117	11
Columbia Univ	8	5	6	23	5
NYU	8	10	7	27	9
Univ Autonoma Barcelona &GSE	8	8	9	11	10
Univ Bonn	8	7	8	2	7
Univ Calif San Diego	8	9	5	24	7
Univ Maastricht	8	6	13	1	7

Table 2

Distances between rankings measured by normalised Sum of Ranking Differences (nSRD).



this size follows approximately normal distribution with mean 0.66 and std. deviation of 0.016. Even Harmonic centrality – which seems to be a little bit farther from the others – is very close according to this metric. The smallest distance is displayed between Degree and Weighted k -shell, while WTC and Harmonic centrality are the farthest away from each other. For more details about the SRD statistics the reader is referred to (Héberger, 2010; Héberger & Kollár-Hunek, 2011; Sziklai & Héberger, 2020).

7. Discussion

Let us address a few issues regarding the applied methodology.

We proposed a framework to rank academic entities within a *discipline* without specifying what we mean by the latter. The Journal of Economic Literature developed a classification system (JEL) to categorize scholarly literature in the field of economics.⁷ JEL distinguishes 20 general categories denoted with letters from (A) to (Z). Game Theory and Bargaining Theory (C7) falls under Category (C): Mathematical and Quantitative Methods. Although our database contains JEL codes only sporadically, when it does, the (C) category label almost always accompanied by some other tags among which the most common is (D): Microeconomics, more specifically, (D7): Analysis of Collective Decision-Making. According to the JEL system, game theory is a mid-level category. It is too general to be completely described by some 3 letter tags, but not general enough to be a main category in itself.

The proposed method can be applied at any level of the classification hierarchy, however, there is a trade-off between the accuracy of the ranking and its coverage. For example, it would be much simpler to compile a comprehensive dataset about Cooperative Games (C71). As soon as we aim higher, though, less related papers start to creep in.⁸ A broad view probably incorporates the research of every game theorist but also finds an increasing amount of less related literature. Thus, as we move from specific to general, we lose our descriptive power: the ranking converges to the general university ranking of the field.

Creating a core set of key journals is not always possible, but fortunately isn't necessary either. There are more than one way to identify the papers that belong to a discipline. One advantage of this method (at least for the topic of game theory) is that it makes the choice of the source (WoS vs Scopus) an insignificant matter. Among the relevant set of journals that we extracted from WoS there is only one that is not featured by Scopus: the proceedings series of the Symposium on Algorithmic Game Theory (SAGT) with a total of 16 papers. Although Scopus covers more journals than WoS (Mongeon & Paul-Hus, 2016), we expect that the relevant set of journals are more or less the same in both databases due to the threshold requirement. Torres-Salinas, Lopez-Cózar, and Jiménez-Contreras (2009) find that the works that had published by the University of

⁷ <https://www.aeaweb.org/jel/guide/jel.php>

⁸ According to Microsoft Academic, the most relevant journal paper on the topic of game theory, at the time of writing, is a paper about smart grids (Mohsenian-Rad, Wong, Jatskevich, Schober, & Leon-Garcia, 2010), which might be an important paper but hardly can be categorized as a mainstream game theoretical research.

Navarra in Scopus that are not indexed by Web of Science receive much less citation (nearly 1/5) on average than the works that are indexed in both. Thus, such works are unlikely to survive when we filter for relevant journals. English-language journals are overrepresented in WoS, but for the same reason we do not expect that the language bias affects the results.

In our model each citation carried the same weight. The literature suggests many improvements over the simple citation count. The location and intensity of citations as well as the context (cf. negative citations) matters (Aksnes, Langfeldt, & Wouters, 2019; Catalini, Lacetera, & Oettl, 2015; Maričić, Spaventi, Pavičić, & Pifat-Mržljak, 1998). It would be also interesting to weigh citations based on the quality of the journal the citing paper was published in.

8. Conclusion

University rankings hide the heterogeneity of the faculty. Successful departments and research groups can put a face behind the university's logo and boost their reputation. Measuring the performance of departments has many advantages, but the problem is as difficult as desirable. Ranking departments by evaluating individual researchers can run into the apples-and-oranges fallacy. Successful departments are often organised around a research topic, thus we can measure their performance by looking at how they are ranked within the discipline the topic belongs to.

In this paper, we introduced a novel method for ranking institutions within a research discipline. The Weighted Top Candidate method is a generalisation of an expert selection method. It relies on a simple observation: experts are much more effective in identifying other experts. Consequently, the selected set must be stable: (i) the expert institutions must support each other's membership and (ii) whenever an institution is deemed expert, its recommendation carries weight; that is, it can nominate other experts. WTC has other advantages beside its axiomatic characterisation. Centrality measures output a real vector, while WTC outputs a list of experts. While at first glance this might seem like a restriction, it enables us to point out which institutions *do not* belong to the expert set. The WTC can and does output the empty set if there are no sensible agents that can be called experts. Even if the algorithm finds experts, the largest stable set is usually significantly smaller than the whole set. In contrast, centrality measures such as PageRank will quantify every node, and just by looking at the numbers, we will not be able to notice the quality difference between an expert and a non-expert node.

The output of WTC has an expressive representation. With a parameter, we can adjust how exclusive our list of expert should be. Plotting the size of the stable sets with respect to the exclusiveness parameter reveals the competitiveness of the analysed field. Simple rankings only reveal the current position of an institution, while Fig. 4 also hints at how difficult it is to improve this position.

The difficulty of applying WTC comes from collecting suitable data. It would be interesting to look at more sophisticated databases that describe research disciplines. In particular, using the micro-field classification of the Leiden ranking, we could compare the fierceness of the competition in different scientific disciplines.

Finally, let us note that WTC is a general expert selection method, which is suitable but not limited to ranking institutions. Depending on the underlying data it can rank authors or journals just as well.

Conflict of interest

The authors declare that they have no conflict of interest.

Author contribution

Balázs Róbert Sziklai: Conceptualization, Methodology, Software, Formal analysis, Data Curation, Writing – Original Draft, Writing – Review & Editing, Visualization.

Acknowledgements

The author thanks Sándor Soós for the valuable discussions and Erzsébet Romsics for helping in the collection of WoS data. This research was supported by the Higher Education Institutional Excellence Program 2020 of the Ministry of Innovation and Technology in the framework of the 'Financial and Public Services' research project (TKP2020-IKA-02) at Corvinus University of Budapest and by the Hungarian National Research, Development and Innovation Office, grant number K119930.

References

- Abramo, G., & D'Angelo, C. A. (2015). *Evaluating university research: Same performance indicator, different rankings*. *Journal of Informetrics*, 9, 514–525.
- Aksnes, D. W., Langfeldt, L., & Wouters, P. (2019). *Citations, citation indicators, and research quality: An overview of basic concepts and theories*. *SAGE Open*, 9.
- Alma, B., Coşkun, E., & Övendireli, E. (2016). *University ranking systems and proposal of a theoretical framework for ranking of turkish universities: A case of management departments*. *Procedia – Social and Behavioral Sciences*, 235, 128–138, 12th International Strategic Management Conference, ISM, 2016, 28–30 October 2016, Antalya, Turkey.
- Altman, A., & Tennenholz, M. (2010). *An axiomatic approach to personalized ranking systems*. pp. 57.
- Boldi, P., & Vigna, S. (2014). *Axioms for centrality*. *Internet Mathematics*, 10, 222–262.
- Bonacich, P. (1972). *Factoring and weighting approaches to status scores and clique identification*. *Journal of Mathematical Sociology*, 2, 113–120.
- Bouyssou, D., & Marchant, T. (2016). *Ranking authors using fractional counting of citations: An axiomatic approach*. *Journal of Informetrics*, 10, 183–199.

- Boyack, K. W., & Klavans, R. (2010). Co-citation analysis, bibliographic coupling, and direct citation: Which citation approach represents the research front most accurately? *Journal of the American Society for Information Science and Technology*, 61, 2389–2404.
- Catalini, C., Lacetera, N., & Oettl, A. (2015). The incidence and role of negative citations in science. *Proceedings of the National Academy of Sciences*, 112, 13823–13826.
- Dusansky, R., & Vernon, C. J. (1998). Rankings of U.S. economics departments. *Journal of Economic Perspectives*, 12, 157–170.
- Flores-Szwagrak, K., & Treibich, R. (2020). Teamwork and individual productivity. *Management Science*, 66, 2523–2544.
- Garas, A., Schweitzer, F., & Havlin, S. (2012). A k-shell decomposition method for weighted networks. *New Journal of Physics*, 14, 083030.
- Guerrero-Bote, V. P., & Moya-Anegón, F. (2012). A further step forward in measuring journals' scientific prestige: The sjr2 indicator. *Journal of Informetrics*, 6, 674–688.
- Héberger, Károly. (2010). Sum of ranking differences compares methods or models fairly. *TrAC Trends in Analytical Chemistry*, 29(1), 101–109. <https://doi.org/10.1016/j.trac.2009.09.009>
- Héberger, Károly, & Kollár-Hunek, Klára. (2011). Sum of ranking differences for method discrimination and its validation: comparison of ranks with random numbers. *Journal of Chemometrics*, 25, 151–158. <https://doi.org/10.1002/cem.1320>
- Kasher, A., & Rubinstein, A. (1997). On the question "who is a j?": A social choice approach. *Logique et Analyse*, 160, 385–395.
- Kitsak, M., Gallos, L. K., Havlin, S., Liljeros, F., Muchnik, L., Stanley, H. E., & Makse, H. A. (2010). Identification of influential spreaders in complex networks. *Nature Physics*, 6, 888–893.
- Laengle, S., Merigó, J. M., Modak, N. M., Modak, & Yang, J. B. (2020). Bibliometrics in operations research and management science: A university analysis. *Annals of Operations Research*, 294, 769–813.
- Lazaridis, T. (2010). Ranking university departments using the mean h-index. *Scientometrics*, 82.
- Ma, N., Guan, J., & Zhao, Y. (2008). Bringing PageRank to the citation analysis. *Information Processing & Management*, 44, 800–810.
- Marchant, T. (2009). An axiomatic characterization of the ranking based on the h-index and some other bibliometric rankings of authors. *Scientometrics*, 80.
- Marchiori, M., & Latora, V. (2000). Harmony in the small-world. *Physica A: Statistical Mechanics and its Applications*, 285, 539–546.
- Maričić, S., Spaventi, J., Pavićić, L., & Pifat-Mrzljak, G. (1998). Citation context versus the frequency counts of citation histories. *Journal of the American Society for Information Science*, 49, 530–540.
- Mohsenian-Rad, A., Wong, V. W. S., Jatskevich, J., Schober, R., & Leon-Garcia, A. (2010). Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1, 320–331.
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of web of science and scopus: A comparative analysis. *Scientometrics*, 106, 213–228.
- Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The PageRank Citation Ranking: Bringing Order to the Web, Technical Report 1999–66. Stanford InfoLab.. previous number = SIDL-WP-1999-0120. <http://ilpubs.stanford.edu:8090/422/>
- Seidman, S. B. (1983). Network structure and minimum degree. *Social Networks*, 5, 269–287.
- Shibata, N., Kajikawa, Y., Takeda, Y., & Matsushima, K. (2009). Comparative study on methods of detecting research fronts using different types of citation. *Journal of the American Society for Information Science and Technology*, 60, 571–580.
- Sziklai, B. R. (2018). How to identify experts in community? *International Journal of Game Theory*, 47, 155–173.
- Sziklai, B. R., & Héberger, K. (2020). Apportionment and districting by sum of ranking differences. *PLoS One*, 15, 1–20.
- Torres-Salinas, D., Lopez-Cózar, E. D., & Jiménez-Contreras, E. (2009). Ranking of departments and researchers within a university using two different databases: Web of science versus scopus. *Scientometrics*, 80, 761–774.
- Waltman, L., & van Eck, N. J. (2012). A new methodology for constructing a publication-level classification system of science. *Journal of the American Society for Information Science and Technology*, 63, 2378–2392.
- West, J. D., Bergstrom, T. C., & Bergstrom, C. T. (2010). The eigenfactor metrics $\langle ce : sup \rangle TM \langle /ce : sup \rangle$: A network approach to assessing scholarly journals. *College & Research Libraries*, 71.
- Wąs, T., & Skibski, O. (2018a). An axiomatization of the eigenvector and katz centralities. In S. A. McIlraith, & K. Q. Weinberger (Eds.), *Proceedings of the thirty-second AAAI conference on artificial intelligence, AAAI-18, the 30th innovative applications of artificial intelligence (IAAI-18), and the 8th AAAI symposium on educational advances in artificial intelligence (EAAI-18)* (pp. 1258–1265). AAAI Press.
- Wąs, T., & Skibski, O. (2018b). Axiomatization of the PageRank centrality. In *Proceedings of the twenty-seventh international joint conference on artificial intelligence, IJCAI-18*, pp. 3898–3904. IJCAI.
- Zachos, G. (1991). Research output evaluation of two university departments in Greece with the use of bibliometric indicators. *Scientometrics*, 21, 195–221.