


RESEARCH ARTICLE

# Selecting and ranking leading cases in Brazilian Supreme Court decisions

Jackson José De Souza<sup>1</sup> , Marcelo Finger<sup>1</sup>, Jorge Alberto A. de Araújo<sup>2</sup> and Juliano Maranhão<sup>2</sup>

<sup>1</sup>Department of Computer Science, Institute of Mathematics and Statistics, University of Sao Paulo, São Paulo, Brazil,  
e-mails: [jackson@ime.usp.br](mailto:jackson@ime.usp.br), [mfinger@ime.usp.br](mailto:mfinger@ime.usp.br)

<sup>2</sup>Department of Philosophy and Legal Theory, University of Sao Paulo Law School, São Paulo, Brazil,  
e-mails: [jorge.alberto@ufma.br](mailto:jorge.alberto@ufma.br), [julianomaranhao@usp.br](mailto:julianomaranhao@usp.br)

Received: 7 April 2021; Revised: 1 April 2023; Accepted: 25 April 2023

## Abstract

This work studies quantitative measures for ranking judicial decisions by the Brazilian Supreme Court [Supremo Tribunal Federal (STF)] and selecting leading cases, which are understood as those with broadness of influence on different legal fields. The measures are based on a network built over decisions whose cases were finalized in the Brazilian Supreme Court between 01/2001 and 12/2019, which were obtained by crawling publicly available STF records. Three ranking measures are proposed; two are adaptations of the PageRank algorithm, and one adapts Kleinberg's algorithm. Such measures are compared with respect to agreement on top 100 rankings; we also analyze each robustness measure based on self-agreement under perturbation.

We examine whether the resulting quantitative ranking is congenial to a qualitative intuition of what the legal community usually considers as relevant precedents. We also discuss some possible criteria of relevance in the seek for patterns that suggest how quantitative and qualitative measures would better align. The ranking of leading cases and relevant decisions improved after building decision networks without irrelevant appeals and decisions that overflow the court offers a starting point to discuss the role of STF in the Brazilian judicial system.

In our last work, both versions of PageRank and Kleinberg algorithms produced different rankings and all of them were robust with respect to 10% and 20%-perturbation levels, but none of them retrieved leading cases at the top of these rankings. Then, we took a further step in the studies of the STF decision network and we introduced better filtering of network nodes guided by legal expertise on the works of the Supreme Court. We also introduced more fine-grained perturbation levels to understand the impact of such filters in the STF decision network. We concluded that after filtering low-relevance decision types, the STF decision network is still robust under 10%-perturbation, but it presents higher degradation by increasing perturbation levels. The two versions of PageRank still produce different rankings. Kleinberg's algorithm provides a different ranking, with many relevant criminal cases. Although we improved algorithms rankings filtering decisions from the network, which represents an important methodological step, there is still room for improvement. Given that relevant decisions are well ranked after filtering out a large amount of irrelevant decisions, the results set a starting point to discuss the role of STF in the Brazilian judicial system.

## 1. Introduction

This work investigates the development of network analysis tools based on natural language processing (NLP) aiming at detecting leading cases of the Brazilian Supreme Court and understanding their impact on case law by that same court. This work expands previous results in de Souza and Finger (2020), which did not distinguish between leading cases and regular cases. As before, we must state that there is no widely accepted standard tool either for identifying leading cases or for ranking them, but recent

---

**Cite this article:** J.J. De Souza, M. Finger, J.A. Araújo and J. Maranhão. Selecting and ranking leading cases in Brazilian Supreme Court decisions. *The Knowledge Engineering Review* 38(e7): 1–18. <https://doi.org/10.1017/S0269888923000073>

developments in computer science and artificial intelligence, enhanced with expert analysis of the data, allow us to put forward a few proposals.

This work aims to be another step in the discussion of the desired properties of quantitative feasible measures of relevance of Supreme Court decisions. We believe our research is a relevant contribution to the literature because (1) the Supremo Tribunal Federal (STF) decision network has not been sufficiently explored in the literature to understand its structure and to find a good measure to rank leading cases; (2) we are applying a new measure of network robustness that we have proposed in de Souza and Finger (2020); and (3) we are also applying filters under the guidance of legal experts designed to consider specific aspects of STF functioning.

Selection of leading cases should presumably reflect the influence that a precedent has on the whole system of adjudication, due to the role and higher position of the Supreme Court in the Brazilian judicial system's hierarchy. Applying measures to the decision network may improve the knowledge about Brazilian Supreme Court's functioning in terms of the cohesion of its precedent network, transparency to the legal community with respect to its own rulings, and identification of key precedents. As the STF is a Supreme Court that holds more responsibilities than other well-studied Supreme Courts like the Supreme Court of the United States (SCOTUS) and decides a high volume of cases, it is important to study how different the STF is compared to other Supreme Courts in respect to authority scores and how good these measures are at ranking leading cases on the STF decision network. Also, it is important to investigate if they might reveal or emphasize some of its functions, whether they are particular to the STF or not. We might discover interesting findings for the legal community.

The Brazilian Supreme Court (STF) is of utmost importance to the Brazilian judicial system and to the structure of the democratic republic as it is empowered to enforce the Brazilian Constitution and, therefore, prevent unconstitutional laws from becoming valid or producing effects in the legal system. Since the content of the Brazilian constitution is extensive (not limited to fundamental principles and rights), there is a vast array of subjects that arrive at the court, be it directly or indirectly, through the complex Brazilian system of appeals. This attribution has a negative impact on the performance of the STF, generating a high demand for lawsuits and appeals (Falcão *et al.*, 2014).

Differently from other Supreme Courts that select a few relevant cases for analysis, the Brazilian court is the last instance of appeals on several different subjects, issuing around 80.000 rulings per year.

Several measures were deployed to reduce this enormous amount of cases, such as a new procedural law demanding that appeals to the Supreme Court must have 'general repercussion', which has to be argued for by the appellant, as well as the provision by law of specific rulings by the STF that are binding for the Judiciary. It should restrict access to the Court like the Certiorary Act of 1925<sup>1</sup> did for the SCOTUS. However, the number of appeals that still reach the STF and the variety of subjects is extremely large.

Thus, the court may decide issues of political relevance, such as the eligibility of politicians that were implied in criminal investigations or that were condemned by an inferior court (before a definitive ruling by courts of appeal), or economic impact, such as the constitutionality of a tax collected from industrial activities, and also deliberates on legal theses that become influential in the judicial system and therefore in the legal community as a whole. Particularly important is the Supreme Court's role as the guardian of fundamental rights enshrined in the Brazilian Constitution. This enables the court to rule on cases of social significance and impact, such as the legality of same-sex marriage, the admissibility of abortion of anencephalic fetuses, or the granting of access to the digital content of mobile phones in search and seizure procedures.

Legal scholars are interested in analyzing the influence these STF decisions may have on future cases that deal with similar matters. Decisions with broadness of influence, for example that modify the previous understanding on some significant issue and become the ground of future rulings are usually

<sup>1</sup>The Judiciary Act of 1925, also known as 'Certiorary Act', was designed to reduce the workload of the Supreme Court of the United States (SCOTUS). Since then, any appellant must file petitions for a 'writ of certiorary' with the Court. The acceptance of the petition depends on the discretionary favorable decision of four Justices. Usually, the Court hears 100–150 cases each year.

called ‘leading cases’. An important question raised in this paper is whether such leading cases may be captured by quantitative methods in terms of citations in the decision network.

This work is only possible as the STF makes public and freely available a substantial amount of electronic information on court decisions. We have crawled that material to extract and process decision data in text, written in Portuguese, to build a decision network only with decisions issued by the STF.

Based upon a such complex network and the concept of authority scores, we keep studying the *decision ranking measures* built on PageRank (Page *et al.*, 1999) and Kleinberg’s (1999) ranking algorithms proposed in de Souza and Finger (2020). Our goal is to investigate the level of *agreement between measures*, their *robustness* with respect to the network, and whether most of best-ranked decisions by the measures are *leading cases*. The latter goal is a desirable ranking for legal scholars. For robustness, it is expected that the rankings be preserved under addition/removal of some small random number of nodes; otherwise, the ranking is too unstable to be useful. The proposed measures are an exploratory attempt to provide quantitative support for claims about Brazilian Supreme Court decisions, and should be considered among the first steps toward more complex analyses of decision structure.

To find the leading cases, we also propose filters to remove some decisions from the decision network. This filtering is conducted under the supervision of legal experts that understand the workings of the court and that are capable of proposing criteria to select, among the huge number of decisions issued by the court, which processes are just ‘noise’ in the search for cases of repercussion.

The results achieved are promising, as the filters applied in the decision network helped to retrieve more leading cases and decisions of legal relevance while also putting in evidence a few important characteristics of the STF, some of them hidden, that have the potential to produce social impact. The contrast between the decision relevance rankings of different decision networks shows how much time and importance STF devotes to non-constitutional-related matters such as appeals and public pension cases that overflow the court. These results offer a foundation to discuss the suitability of the functions of STF for the Brazilian judicial system. The present work provides a platform on which more complex NLP may be performed, such as legal argument extraction and evaluation of the impact of specific laws in society.

This paper is structured as follows. We start by analyzing related work on measures and algorithms for quantitative and qualitative ranking measures of legal decisions in Section 2. In Section 3 we provide expert legal analysis that supports the filters our method applies to network data to identify leading cases. Then, we describe how the STF decision data was extracted and preprocessed, how the STF decision network is modeled, and we describe the construction of decision networks studied and compared in this work in Section 4. In the same section, we adapt node ranking algorithms to create measures of decision ranking based on the STF decision network, a measure for the agreement of those rankings that is statistically analyzed, and then we describe a set of robustness tests on those rankings. Results obtained are discussed in Section 5 and we conclude on the compliance of the proposed measures to the desired properties of agreement and robustness and inform main discoveries about ranking leading cases.

## 2. Related work

Algorithms for discovering authoritative nodes in complex networks, which receive a large number of references, and hub nodes, which refer to several nodes, were proposed by Kleinberg (1999). A few networks built on U.S. and European Supreme Courts decisions have been studied in the last decades, such as Agnoloni and Pagallo (2015), Fowler and Jeon (2008), van Opijnen (2012), and Winkels *et al.* (2011). Both Agnoloni and Pagallo (2015) and Fowler and Jeon (2008) have found that Kleinberg’s algorithm leads to scores that usually meet the evaluation of legal experts on relevant decisions, in which relevant *hub decisions* are those that cite many relevant authoritative decisions; similarly, relevant *authoritative decisions* are those cited by many relevant hubs. The analysis of the *in via incidentale* rulings of the Italian Constitutional Court by Agnoloni and Pagallo (2015) identified decisions that are quite debated by legal scholars but are rarely cited by the court due to the definitive resolution of a matter

settled in the ruling; Agnoloni and Pagallo (2015) decision network topology is scale-free according to a power law, and Fowler and Jeon (2008) identified that most cases have a small degree, a few cases are widely cited and other few cases cite a large number of cases, but they did not say that the SCOTUS decision network is scale-free.

Another approach was proposed by van Opijnen (2012), with good results using closeness measures that calculates the distance between decisions, such as *proximity prestige* and *generalized core*, but in their judgment the best results were achieved with ‘Marc in-degree’, a measure that only takes into account incoming citations. On the same line, sink distance measure combined with the single-linkage hierarchical clustering algorithm produced more accurate and more interpretable clusterings in the work done by Bommarito *et al.* (2010). In contrast, van Opijnen (2012) and Winkels *et al.* (2011) did not achieve good or meaningful importance scores with the PageRank algorithm (Page *et al.* 1999), i.e., the relative importance of the referring case does not seem to predict the relevance of the cases it refers to. The results obtained by van Opijnen (2012) using indegree Hyperlink-Induced Topic Search (HITS) showed a low correlation between their results and the number of publications in specialized magazines and the number of citations in literature, too.

There are some studies about the Brazilian Supreme Court, but none of them address the topic of authority scores nor the Supreme Court’s decision citation network. The FGV-Rio Law School<sup>2</sup> has been publishing quantitative analyses in articles and reports since 2010 about the Brazilian Supreme Court. Usually, the reports address questions aiming at fostering debate about the court’s activity through queries that can be answered with basic statistics such as ‘what are the heaviest users of the court?’, ‘which justices take more time to judge preliminary injunctions<sup>3</sup>?’, ‘which types of cases take more time to have a final decision and which ones are judged fastest?’ (Falcão *et al.* 2013). The articles published by FGV-Rio Law School about the Brazilian Supreme Court address more specific issues like ‘Is the time taken to request to view a case shorter in comparison with the U.S. Supreme Court?’ (Hartmann *et al.* 2017) and ‘What affects more the court’s cohesion, the court’s workload, or differences between the justices’ personalities?’ (Almeida *et al.* 2016).

In Section 3 we also cite some works discussing the perception of leading cases by the legal community and ranking lists of cases decided by the SCOTUS, considered as authentic leading cases by the legal community.

With respect to our claims of originality and valuable contribution to the legal community studying the Brazilian Supreme Court, we note that:

- (1) As far as we know, we are the first to study authority scores for the Supreme Court of one of the biggest democracies in the world, with some significant differences from other countries’ Supreme Court analyses, pointing to methodological contributions to the task;
- (2) we apply a new method to evaluate robustness in a network by calculating the intersection of decisions on the top 100 positions in multiple running trials;
- (3) we apply expert-guided filters specific for the STF to filter out decisions considered irrelevant for identifying leading cases in the decision network; and
- (4) due to these filters, we refine robustness levels to get a better understanding of decision network robustness.

The last two items are the techniques we introduced in this work in the pursuit of finding a good authority score for STF leading cases and understand its decision network better.

### 3. Quantitative vs. qualitative conception of relevance

Matching the quantitative evaluation of relevance with the qualitative appreciation by the legal community about what are the leading or most relevant cases in the judicial system, even when this question is

<sup>2</sup><https://portal.fgv.br/en>

<sup>3</sup>In Brazil, preliminary injunctions are petitions that have temporary and urgent features. It aims to avoid likely violations of constitutionally guaranteed rights or irreparable harm in the absence of preliminary relief caused by delay in the legal process.

restricted to the Supreme Court, may prove a difficult or even impossible task. The first difficulty lies in the different criteria of relevance that may be used within the legal community. Actually, relevance is an ‘interpretive concept’ (Dworkin, 1986) that is, it admits different ‘conceptions’ according to the point or goal of its application as stipulated by the interpreter or by a language community.

A case may be considered relevant in terms of its *economic impact*, for instance, in tax law, a conventional case that does not bring any innovation to tax law theory may imply an impact of billions of Brazilian reais on the tax revenue in the country. Unless there is an important and broader constitutional issue at stake, a case in this specific and technical field would hardly be considered relevant in other fields such as criminal, civil, or environmental law.

On the other hand, matters of procedural law, which would certainly have repercussions for many fields, would hardly be considered ‘relevant’ by the legal community, which would mostly be focused on cases with substantive and material issues rather than procedural ones. Thus, a conception of relevance as *broadness of influence on different legal fields* would probably not be congenial to what jurists would call a leading or relevant precedent.

Another interesting conception of relevance would be the *social impact* of a decision modifying previous case law, in terms of its polemic in the social context, and how it has affected not only legal knowledge but the community’s culture or its fundamental values. Such cases are usually associated with the constitutionality of laws or authoritative decisions that hinders the exercise of fundamental rights.

A recent measure by one of the largest repositories of legal papers showed that the most cited cases decided by the SCOTUS are mostly the ones the legal community expected, largely coincident with any list of the most famous or most important in Mattiuzzo (2018). For instance, the first two, *Brown v. Board of Education*<sup>4</sup> and *Roe v. Wade*<sup>5</sup> are prominent in every ranking of landmark decisions, using whichever criteria (Rehnquist, 2002; Irons, 2006; Cushman, 2011; Steinman, 2016; Mattiuzzo, 2018).

Both of them, however, would hardly figure in a citation count examining judicial rulings in the United States. In the first one, the Court decided in 1954 that segregated schools are inherently unequal and determined that schools specifically for black or white students should be integrated. This ruling had a significant social impact in the pre-civil rights era, and even in the 2000s, the problem is still widely regarded as unsolved (Irons, 2002; Patterson, 2004). In the second, the Supreme Court affirmed in 1973 the appellant’s right to an abortion despite legal provisions to the contrary. This ruling was recently overruled by *Dobbs v. Jackson Women’s Health Organization*, No. 19-1392, 597 U.S. \_ (2022)”, which shows that it is still a relevant issue today given that it is discussed every time a new Justice<sup>6</sup> is appointed to the Court (Green, 2020). Despite all social significance, one would be hard-pressed to find an abundance of rulings by the Court citing these two as precedents because there is not much else to debate: segregated schools and state laws forbidding abortions without any qualification are unconstitutional and remain so, with only a handful of reaffirmations of the rulings in the decades that followed.

In Brazil, we may easily think of decisions regarded as significant in terms of social impact by the legal community, which we could hypothesize would fail to appear in a citation count due to the high number of procedural decisions that are cited by hundreds of decisions. In 2012, the STF decided ADPF 54<sup>7</sup>, a much-debated case about abortion rights in the case of an anencephalic fetus (abortion is a criminal offense in Brazil<sup>8</sup>). It was the first time the Court opened the arguments to third parties, and dozens of organizations from the government and civil society debated heatedly before the judges for two weeks. The Court affirmed the right of any pregnant woman with a diagnosis of anencephalic fetus to an abortion without the need for a lawsuit, despite there being no exception of this kind in the Criminal Code.

<sup>4</sup>347 U.S. 483, cited by 23 395 articles as of 2018.

<sup>5</sup>410 U.S. 113, cited by 21 692 articles as of 2018.

<sup>6</sup>Justice is how a Magistrate is called in the United States.

<sup>7</sup>ADPF stands for ‘arguição de descumprimento de preceito fundamental’, which could be translated as a complaint of a violation of a fundamental right. A petition of this kind is presented directly to the STF.

<sup>8</sup>Articles 124–128 of the Brazilian Criminal Code.

A STF judge at the time said that was the most important case ever to be decided by the Court<sup>9</sup>. A year before the Court decided ADPF 132, ruling that same-sex relationships (a ‘stable union’, which is commonly referred to in English as a ‘common-law marriage’ or marriage by habit and repute) should have the same legal status as any other conjugal relationship—despite a specific provision in the Brazilian Constitution of 1988 granting such status only to relationships between a ‘man and a woman’. Two years later, the Brazilian Judicial Council or National Council of Justice (CNJ), an independent body that oversees the Judiciary, used that ruling as a basis to extend formal marriage rights to same-sex couples. Notwithstanding the undeniable social relevance of these rulings, the fact remains that in the aftermath of these decisions, the legal questions settled there will not return to the Court, at least not in the same way. The cases are usually cited as examples of unorthodox decisions whenever another case has to be dealt with in the same or in a similar fashion<sup>10</sup>. The result is that these cases are not cited as much as cases that are recurrent in the court dockets.

Hence, the quantitative measure would probably match a qualitative evaluation of how influential a case is in terms of becoming the ground for several subsequent precedents. This idea captures one aspect of a leading case, which is the fact that the case makes a difference and it usually does because it has changed in some relevant aspect a previous orientation of the case law. It also captures an intuitive and literal notion of a case being ‘influential’ within the very legal community in terms of bringing legal a thesis, which is further reproduced within the Supreme Court and thus, presumably, by inferior courts and by legal doctrine.

However, there are some possible traps that must be dealt with if we are going to use a criterion of relevance as *influence in future case law*. These traps are related to the peculiarities of the Brazilian legal system and the Brazilian Supreme Court.

Prominent among them is the possibility of cases with limited repercussions reaching appellate courts and even the Supreme Court. Cases dealing with a wide range of questions such as social benefits, pensions, and even petitions for *habeas corpus* are common in the Supreme Court’s dockets. At some point, we may have thousands of cases with the same legal question before the Court, and most or all of them end up being decided the same way, in exactly the same terms. Because of that, a precedent might be inserted in a decision that is replicated in thousands of cases—which means that the precedent is cited thousands of times. However, the decision cited as precedent may not even be the ‘technical’ precedent, the leading case that established a legal orientation or changed a previously undisputed understanding about some legal question. We identified this phenomenon appearing in some experiments whose results we discuss in Subsection 5.2.

Another particularity of the proceedings in the Supreme Court is the petition called ‘agravo regimental’ (abbreviated as AgRg or AgR). It is an appeal provided by a Supreme Court internal regulation and can be filed after a ruling by a single Justice that is adverse to the appellant. The petition is then decided by one of the two Panels (of five Justices each) or by the whole Court (of eleven Justices)<sup>11</sup>. These petitions are decided in ‘packages’ and usually in similar terms.

Therefore, considering the high volume of AgRs decided in the STF and using a conception of relevance to detect leading cases as those cases that are most influential in terms of citation count in Section 4.3, we have introduced filters to avoid repetitive issues and recurrent litigants to avoid what we consider spurious leading cases, provided that they do not influence future cases but are actually judged simultaneously or represent subjects of minor relevance that frequently appear at the court with regard to the same litigant.

<sup>9</sup> STF, ADPF 54, 2012.

<sup>10</sup> For instance, the ADPF 132 was cited in 2019, when the Court ruled in the ADO 26 that discrimination by sexual or gender orientation should be regarded as criminal according to legislation that punished racism (ADO stands for ‘ação declaratória de omissão’, which could be translated as a motion for a declaration of legislative omission—in this case, the omission in creating a crime punishing discrimination by sexual or gender orientation).

<sup>11</sup> The President of the Court (which would be the Chief Justice in the United States) is not a member of any of the two chambers, but leads deliberations when the whole Court convenes.

## 4. Decision network modeling and analyses

As data extraction and network modeling were already done in our previous work, we briefly describe those tasks here for the sake of reference. For a detailed description, please refer to de Souza and Finger (2020).

### 4.1. Data extraction

We analyze data from rulings called ‘acórdão’, which are collegiate decisions pronounced by the Plenary and the Panels of the STF. The data was extracted and parsed from case entries found in the STF jurisprudence search engine Supremo Tribunal (Federal, 2022). An entry contains summarized information about one decision organized in sections. Among all the data contained in a case entry, we are interested in the header containing the decision code and petition type, the parties of a case, the ‘Note’ (*Observação*) field, which contains all jurisprudential decisions cited, even ones that do not support the decision, and the “Decisions in the same direction” (*Acórdãos no mesmo sentido*) field that contains decisions that we call *similar decisions*, decisions that share the same matter and content and are decided the same way as the case entry.

### 4.2. Network modeling

We model the decision network as a graph of citations and execute ranking algorithms over it. Our goals in this process are: (1) ranking decisions in the decision network for evaluating the relevance of decisions in the highest 100 positions, identifying leading cases in terms of legal and social impact and the influence of decisions; and (2) evaluating the decision network’s robustness under random error or perturbation to find out if the citations build a stable network, that is, a network that does not change easily structurally speaking.

In the process of building the decision network, we consider nodes as decisions and edges as citations to decisions during judgement. Let  $N$  be the number of decisions, each decision represented by a node  $A_i$ ,  $i = 1, \dots, N$ . If the entry for process  $A_i$  mentions  $n_i$  decisions  $A_{i_1}, \dots, A_{i_{n_i}}$ , we create  $n_i$  edges connecting  $A_i$  to each cited decision. The cited decisions are meant to be precedents for a decision that cites it, but as mentioned in Section 4 that may not always be the case. However, we considered all cited decisions as precedents since we cannot distinguish which citations are part of opinions that followed the majority of magistrates in each case. Another issue concerning equating citation with precedent is that even decisions cited in opinions followed by the majority of magistrates may be overruled due to a change in the court’s position concerning a particular matter. Even under those circumstances, the cited decisions are relevant because they contribute to fostering the decision’s prevailing arguments and can be considered influences.

For each similar decision  $S_k^i$  of  $A_i$ , let us also create edges to connect  $S_k^i$  to  $A_{i_1}, \dots, A_{i_{n_i}}$ , given that each similar decision shares the same matter and content of decision  $A_i$ . Although entries were not created for similar decisions, they were also judged like the one for which the entry was created. We assume similar decisions cite the exact same decisions as entry ones, even if this has not been confirmed by a search in the *complete case files*. But including them in the decision network contributes to consolidating the court’s jurisprudence about the matter at issue in these decisions.

### 4.3. Citation networks construction

In our last work (de Souza & Finger, 2020) we showed the decision network is robust for 10% and 20% of perturbations against the network. However, we failed to find a substantial number of decisions that are leading cases or influential decisions whose legal thesis is reproduced in further decisions. Legal experts took a look at the list of the top 100 best-ranked decisions and found that a high number of decisions are ‘AgRs’ for all algorithms studied in our previous and current work. That happened because, despite

**Table 1.** Number of decisions in each decision network analyzed in this work

Network	Decisions (considering similar and cited decisions)	Decision entries only
$G_{orig}$	280 144	103 168
$G_{no\_agr}$	59 639	27 217
$G_{no\_agr\_inss\_stm}$	50 501	25 739

AgR having little to no relevance by any usual standard since AgR is a procedural appeal, AgRs are filed in large quantities and decided by the judges in the same way, in batches of tens or hundreds of appeals. This practice creates a cycle where AgRs are decided in batches, citing previous decisions in AgRs.

Therefore, this result gave birth to the idea of constructing alternate decision networks by filtering out decisions whose existence in high volume structurally hides, or pollutes, parts of the decision network that would reveal leading cases and decisions with legal relevance. As we are interested in leading cases and decisions of legal relevance that settle legal thesis, which is further reproduced, we need to filter decisions to change the network topology and reveal, that is rank better, these decisions. In the search for such a network, we must first define a graph, which is a network.

A directed graph (Sedgewick & Wayne, 2011) is a pair  $G = (V, E)$  comprising:

- $V$ , a set of vertices (nodes, or the decisions in this work);
- $E \subseteq \{(x, y) \mid (x, y) \in V^2 \text{ and } x \neq y\}$ , a set of edges (also called directed edges, i.e., the decision citations) that are ordered pairs of vertices (i.e., an edge connects an ordered pair of vertices).

In this work we construct three networks: the network with all decisions,  $G_{orig}$ , a network without AgRs and the appeals issued on it,  $G_{no\_agr}$ , and a network,  $G_{no\_agr\_inss\_stm}$ , that does not contain AgR and its appeals and decisions which have the Brazilian Social Security Office<sup>12</sup> or the Superior Military Court (*Superior Tribunal Militar*, STM) as one of the parties in a case. The network  $G_{orig}$  has already been studied in our previous work and will be compared to alternate networks. The network  $G_{no\_agr}$  is motivated by the fact that, as we mentioned earlier, AgRs are not relevant in any way, and their presence in large quantities makes it difficult to rank relevant decisions, as we can see in Table 1. After filtering AgRs' case entries, the network size is reduced by almost 80% considering similars and citations. The idea for the network  $G_{no\_agr\_inss\_stm}$  came out after analyzing  $G_{no\_agr}$  and identifying that many decisions were issued about the same matter related to the INSS because it is among the greatest litigators in the Judiciary—actually, it is the greatest litigant after state institutions—and many decisions related to pensions are decided in batches and some are cited hundreds of times by future decisions. Decisions that have STM as a party because they are neither leading cases nor decisions with significant influence over other decisions in the STF.

An important reminder is that the process of building the network is done by retrieving case entries and their cited and similar decisions. As we may not retrieve parties for some cited decisions and similar decisions, some decisions that have INSS or STM as parties may not be filtered from  $G_{no\_agr\_inss\_stm}$  network.

We analyze and compare in this work all three networks with regard to (1) network robustness, (2) network structure, and (3) ranking measures for each network.

#### 4.4. Network node ranking algorithms

We applied three algorithms for ranking decisions in the decision network and evaluated its robustness as we did in de Souza and Finger (2020):  $PR_1$ ,  $PR_2$ , and Kleinberg's. We call in this work as algorithms the  $PR_1$  and  $PR_2$  models that are different equations to calculate the PageRank value of each decision using

<sup>12</sup>In Portuguese, it is called Instituto Nacional do Seguro Social (INSS), which would translate literally to National Institute of Social Security. We decided to use 'office' instead to better convey the nature of the institution as part of the Executive branch.



the PageRank algorithm (Brin & Page, 1998) and we call a measure the result of a ranking algorithm. PageRank is an algorithm created for ranking website pages in a network using a relevance measure for such pages. It outputs a probability distribution, that is, a probability of reaching each page (node) in the network, which represents the likelihood of a person randomly clicking on links to get to a specific page. PageRank works on the idea that a page pointed to by many other pages may be more relevant than those pointed to by only a few pages; furthermore, the relevance of a page increases if the pages pointing to it are also relevant. The  $PR_1$  model is basically the original PageRank; it calculates the  $PR_1(a)$  of a decision  $a$  as a summation of the  $PR_1(b)$  of each decision  $b$  that points to  $a$  divided by the number of citations  $N_b$  done by  $b$ . It is like decision  $b$  is distributing equally its  $PR_1$  for each decision it cites. The  $PR_2$  equation does the same summation, except there is no factor  $\frac{1}{N_b}$  that multiplies the  $PR_2(b)$  of each decision  $b$ . This change compared to  $PR_1(b)$  is based on legal experts' assumption that the importance of a decision should not be reduced by the number of decisions it cites. The process of calculating  $PR_1$  and  $PR_2$  is iterative; the algorithm initially assigns the same value to each node in the decision network and iteratively updates the value of each node  $a$ , calculating the value  $PR_1(a)$  in each iteration until the difference between the Euclidean distance between one iteration,  $PR_1^m$ , and the previous one,  $PR_1^{m-1}$ , is less than a precision  $\epsilon$ , which for this algorithm in this work is  $10^{-8}$  and the algorithm stops.

Kleinberg's algorithm, also called HITS, was designed to find the most relevant pages as an answer to broad search topics in the context of the Web. It plays on two types of pages: authoritative pages, those that are most relevant to the initial query, usually have a large number of incoming links and there is considerable overlap in the set of pages retrieved in the search that point to them; hub pages, those among the retrieved pages in the initial query that point to, that is, have links to authoritative pages, and there is also an overlap of retrieved pages that are pointed to by them. A good hub is a page that points to many good authorities; a good authority is a page that is pointed to by many good hubs. In this work, the pages are the decisions, and the retrieved pages are the decisions present in the decision network. HITS initially assigns the same value to each decision and iteratively updates the authority weight and the hub weight of decision  $a$  in each iteration. When the difference between the sum of the absolute difference of hub weights between iteration  $m$  and the previous one,  $m - 1$ , is less than  $10^{-8}$ , the algorithm stops.

To build the decision network and run  $PR_1$  and Kleinberg's algorithms, we used the NetworkX Python module (Hagberg *et al.*, 2008). To do the same for  $PR_2$  and analyze experiments' results, we wrote the source code<sup>13</sup>. For a more rigorous mathematical definition of each algorithm and a complete description of the calculation of  $PR_1$ ,  $PR_2$ , and Kleinberg's process, see de Souza and Finger (2020).

#### 4.5. Computing measures of decisions relevance and their robustness

There exist a few measures for specific network topologies to evaluate robustness, like those in Albert *et al.* (2000), Schneider *et al.* (2011), but in this work we propose another one to evaluate if a network is robust or not.

To measure the robustness of a decision network, we calculated the *Top100Decisions perturbation* all measure for all algorithms on all three networks analyzed in this work. The idea is to run a chosen algorithm against a decision network and some altered copies of it and find a degree of similarity in their structures. We do this process in 10 trials. In the first trial, we retrieve the set of decisions, build the decision network with them, and run the algorithm against it. The difference between the first and the other 9 trials is that after retrieving the set of decisions, we randomly sample a percentage of decisions, for example, 10%, remove them, and build the decision network with the remaining decisions in the set<sup>14</sup>. The set of decisions removed in each of the 9 trials is supposed to be different from each other, so we can have 10 different decision networks.

<sup>13</sup>The source code for the scraper, the experiments run and the analysis of the results are available at <https://github.com/jacksonjos/analise-juridica>

<sup>14</sup>As the removal of decisions is done before building the network, if a decision entry is removed, all similar decisions related to this decision are also removed.

The removal of decisions from the network we call perturbation on the network, which can also be understood as a considerable amount of random errors on it. We call *Top100Decisions* the list of 100 best-ranked decisions obtained as a result of running the ranking algorithm against the network. We then consider a decision as ‘agreed upon’ if it reaches 80%-threshold, that is, the decision is present in at least 8 of 10 running trials. The number of decisions present in the *Top100Decisions* for at least 8 of 10 running trials is the *Top100Decisions perturbation all* as a result of an algorithm on a network.

The idea is to obtain a ‘dynamic’ view of robustness, that is, if the list of *Top100Decisions* among multiple trials with random perturbations on the same decision network does not change dramatically, the network is found to be robust. As we mean to do a fair comparison of *Top100Decisions perturbation all* between all measures and perturbation levels, we remove the exact same decisions from the network for each algorithm; that is, the exact same perturbation is applied for each algorithm in the first trial, in the second trial, and so on. In our previous work, we ran each algorithm for 10%, 20%, and 30% levels of perturbation for network  $G_{orig}$ , and we kept it as it is. In this work, we also run 5%, 10%, 15%, 20%, 25%, and 30% levels of perturbation for networks  $G_{no\_agr}$  and  $G_{no\_agr\_inss\_stm}$ .

As we are interested in comparing if  $PR_1$  and  $PR_2$  measures are different, we calculated the Chi-squared hypothesis over the *Top100Decisions perturbation all* results obtained by  $PR_1$  and  $PR_2$ , and we analyzed the intersection of decisions contained in the *Top100Decisions perturbation all* between both measures. A high degree of agreement among those measures indicates that they are capturing, in the higher levels of relevance, a similar notion of decision authority.

Considering the centrality of *Top100Decisions perturbation* measure and that it can take values between 0 and 100, we fix the threshold of 50 as a criterion to determine if the network is robust or not. So, when a *Top100Decisions perturbation* measure is above 50 the network is robust; otherwise, it is weak.

## 5. Results and discussions

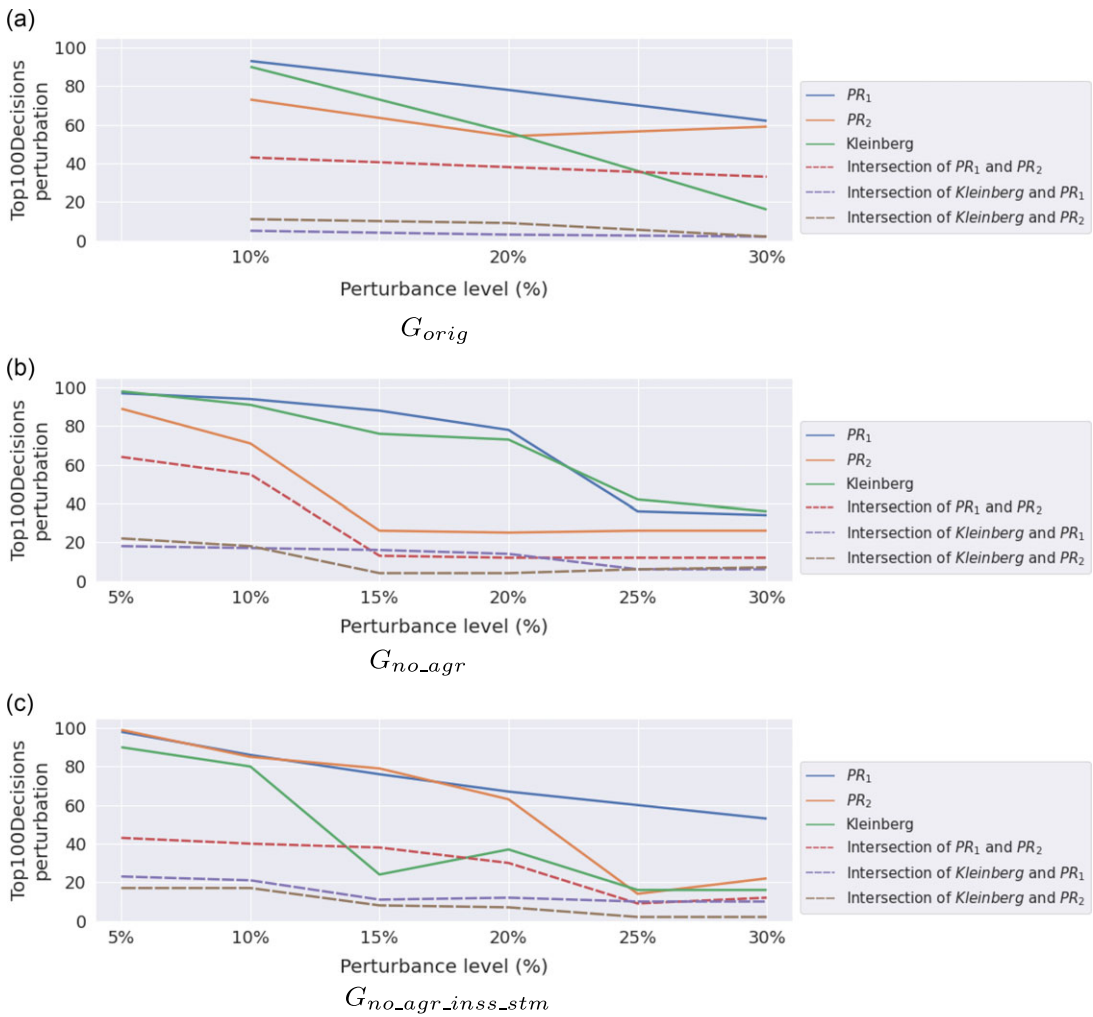
Motivated by the search for a decision ranking that matches the criteria presented by legal scholars, we performed experiments on three decision networks. Two of them were built on semantic filters that removed decisions belonging to a specific class of procedural decisions, for example, AgR, and removed decisions that have INSS and STM as litigants. This is justified by the fact that the removed decisions are neither leading cases nor relevant decisions and do not help to rank them well. After their removal, the network helped rank measures concerning the top leading cases and relevant decisions.

This section is split into quantitative results 5.1 and qualitative results 5.2, aiming to present results regarding network robustness, network topology, and ranking measures for each decision network studied in this work. We also opted to communicate, in this section, the most quantitative results found in this work by means of plots instead of tables, as plots are more concise and clear.

### 5.1. Quantitative results

To obtain a fine-grained view of decision network degradation under perturbation for the smaller networks  $G_{no\_agr}$  and  $G_{no\_agr\_inss\_stm}$ , we adopted the perturbation levels 5%, 10%, 15%, 20%, 25%, and 30%. Analyzing the results, we can see that perturbation levels above 20% reduced by half the values of *Top100Decisions perturbation all* for almost all decision networks, as shown in Figure 1(b) and (c). Thus, we decided to ignore perturbation levels above 20%, as comparisons of network robustness become unfeasible with that level of perturbation. This observation was reached only after filtering out irrelevant decisions from  $G_{no\_agr}$  and  $G_{no\_agr\_inss\_stm}$  decision networks, which supports legal experts’ knowledge that, for the purposes of our work, they distorted the decision network  $G_{orig}$  studied in our previous work, leading to mistaken conclusions.

The results show all three decision networks are robust with respect to the *Top100Decisions perturbation all* measure for perturbation levels 5% and 10% for all ranking measures since the *Top100Decisions*

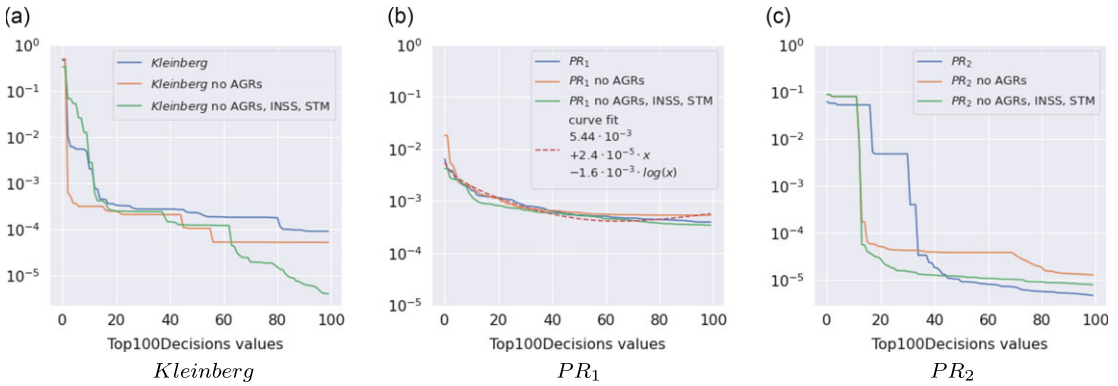


**Figure 1.** Top100Decisions perturbation all results for each perturbation level for  $PR_1$ ,  $PR_2$ , and Kleinberg’s ranking measures and the intersection of each pair of measures for each network

perturbation all is closer to 100 than to zero, as we can see in Figure 1(a). Also, ranking measure  $PR_1$  is robust for perturbation levels of 15% and 20% for all decision networks.

We also analyzed Top100Decisions ranking values for each ranking measure and each decision network without perturbation to analyze how far decisions are from each other and how the difference in ranking values changes as ranking positions decrease. As we can see in Figure 2(b), not only do  $PR_1$  ranking values decrease smoothly from position 1 to 100, but also the ranking values of all decision networks almost overlap. The  $PR_2$  and Kleinberg ranking measures produced ranking values that resembled a step function. We can also see that the ranking values for the  $PR_2$  measure are less stable, changing more abruptly, and the ranking values diverge between decision networks for the  $PR_2$  and Kleinberg measures, in contrast to the  $PR_1$  results.

The Top100Decisions perturbation all results of each decision network studied in this work indicate that the decision network is robust, which can be explained by the fact that the decision networks studied in this work have a scale-free topology, which is also the case for citation networks (Barabási & Pósfai, 2016). The name scale-free comes from the fact that there is not internal scale with respect to the degrees of nodes in the network because some hubs, nodes with a large number of links, coexist with a huge number of small-degree nodes, resulting in a fat-tailed degree distribution. We can see that



**Figure 2.** Ranking values of *Top100Decisions* for each ranking measure and each decision network without perturbation

in the decision network  $G_{orig}$  in Figure 3(a), in which there are a lot of nodes with small degrees and a node with a degree over 5000.

Furthermore, the size of hubs grows as a scale-free network has more nodes. We can see that in the opposite direction when we look at decision networks  $G_{no\_agr}$  and  $G_{no\_agr\_inss\_stm}$  which contain a lower number of nodes than  $G_{orig}$  as we can see in Table 1. As a result, the highest-degree node in both decision networks is smaller than that in  $G_{orig}$  as we can see in Figure 3.

In scale-free networks, degree distribution follows a power law  $P_{deg}(k) = \alpha k^{-\gamma}$ , in which the probability of node degree, that is, number of edges per node, decays as the node degrees  $\langle k \rangle$  increase. However, in real networks, many phenomena change the nature of the degree distribution, resulting in a deviation from a pure power law, which can be observed in the decision networks studied in this work.

To deal with a degree distribution that deviates from a pure power law, we fitted a power law to the decision network degree distributions by applying the gamma method estimation developed by Clauset et al. (2009) and Klaus et al. (2011). We can see in Figure 4 for each decision network that  $P_{deg}(k)$  vs.  $\langle k \rangle$  plots almost form a straight line as in a pure power law. The data is scaled to a log-log plot with logarithmic binning for readability.

Analyzing the results quantitative and qualitatively, we reached the conclusion that the decision networks are scale-free due to the functioning of STF, in which some issues are decided in batches; that is, some issues have a large number of similar decisions, and these decisions share the same citations, increasing substantially the relevance of a certain issue in the network. We showed plots of in-degree decision networks because decision networks are directed, and in this case, we have to study networks of in and out degrees separately (Barabási & Pósfai, 2016). For the sake of conciseness, we decided to show results of in-degree decision networks to focus on cited decisions, but all out-degree decision networks are scale-free, too.

We keep interested in comparing the *Top100Decisions perturbation all* measure between  $PR_1$  and  $PR_2$  to find out if they display different levels of dynamic robustness. We perform a hypothesis test to compare if two distributions with some properties in common are statistically different (de Souza & Finger, 2020). We use the Chi-squared hypothesis test,  $\chi^2$ -test, to compare the dynamic robustness between  $PR_1$  and  $PR_2$  *Top100Decisions perturbation all* results in Figure 1 since the set of perturbation levels for the  $PR_1, PR_2$  pair follows the  $\chi^2$ -distribution. This hypothesis test is done considering the 10%, 20%, and 30% perturbation levels for decision network  $G_{orig}$  and the 5%, 10%, 15%, and 20% perturbation levels for decision networks  $G_{no\_agr}$  and  $G_{no\_agr\_inss\_stm}$ .

The p-value adopted to reject the null hypothesis is 0.05 for  $PR_1$  and  $PR_2$  *Top100Decisions perturbation all* measures. The hypothesis test for the *Top100Decisions perturbation all* measure obtained a p-value of 0.45 for  $G_{orig}$ , as already informed by de Souza and Finger (2020), and p-values of  $2.32 \times 10^{-6}$  and 0.98 for  $G_{no\_agr}$  and  $G_{no\_agr\_inss\_stm}$ , respectively. Therefore, we cannot reject the hypothesis that both PageRank model versions,  $PR_1$  and  $PR_2$ , retrieve the same decision ranking for  $G_{orig}$  and  $G_{no\_agr\_inss\_stm}$

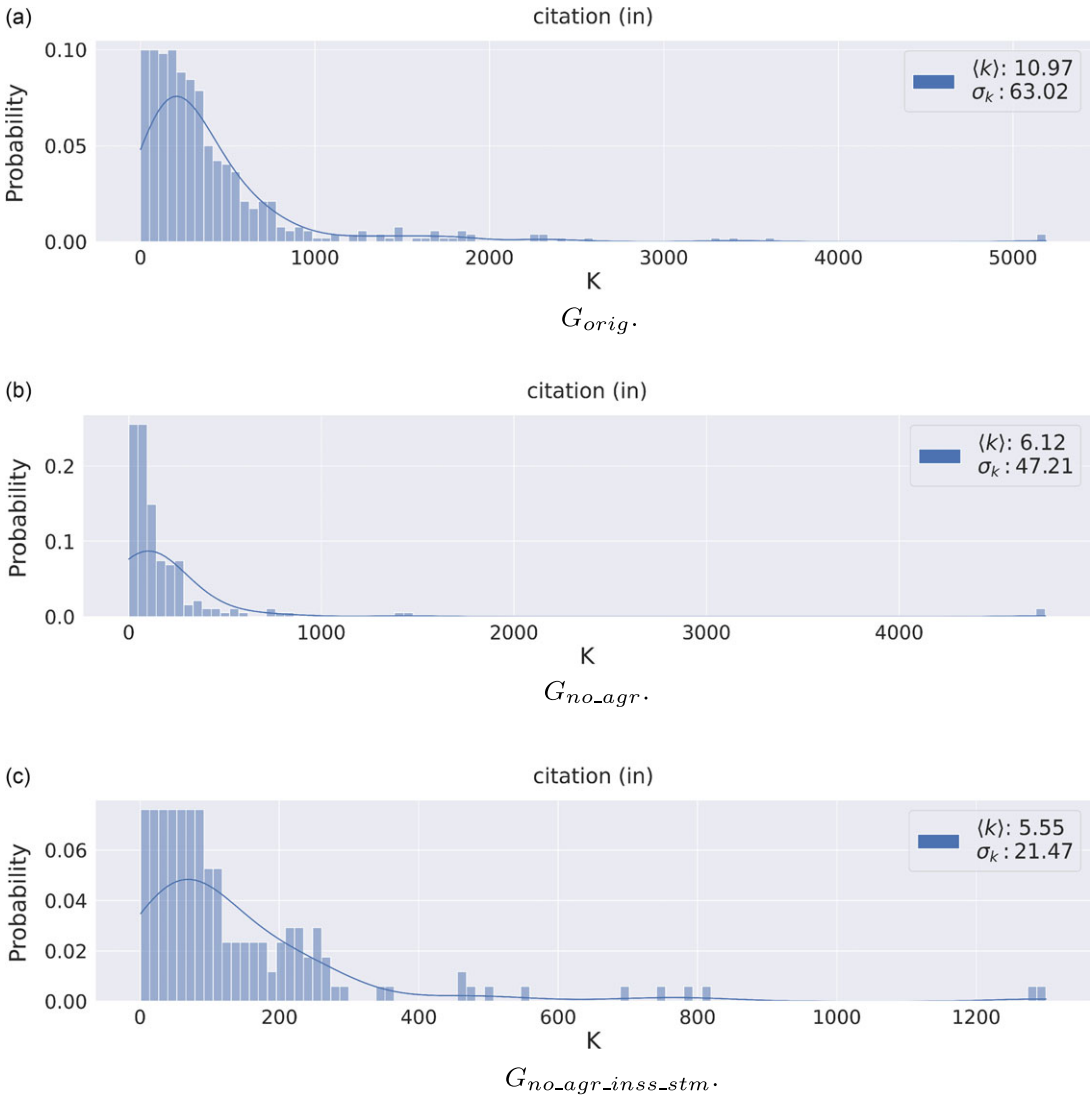


Figure 3. Distribution of node degrees in each decision network

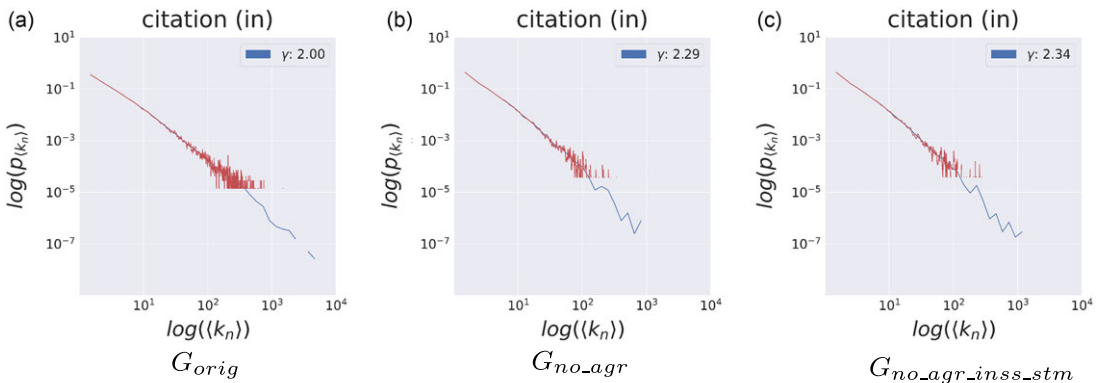


Figure 4. Probability of node degree vs. normalized frequency of node degree (log scale, both) for each network

networks. But we can do so for the  $G_{no\_agr}$  network, which makes sense because  $PR_1$  is robust for all compared perturbation levels while  $PR_2$  is not for perturbation levels 15% and 20%, as we can see in Figure 1(b).

## 5.2. Qualitative results

As discussed in Section 3, leading cases are decisions that modify the previous understanding on some significant subject affecting not only legal knowledge but also society's culture or its fundamental values. Leading cases have such potential for influence because they usually become the ground for future rulings. And this is a result of their association with the constitutionality of laws or authoritative decisions, which hinders the exercise of fundamental rights. Also, there is another category of decisions that have legal impact because they settle relevant theses that are used as precedents by future decisions, but they do not spark the same social impact because they are not necessarily associated with the constitutionality of laws as leading cases. We refer to these decisions as having legal relevance.

The search for meaningful authority scores is our main motivation for creating alternate network versions as defined in Subsection 4.3. With the help of legal experts who are coauthors of this work, we examined the results of ranking measures studied in this work to identify leading cases and decisions of legal relevance. We did that for the *Top100Decisions* lists obtained running every algorithm in the first of 10 trials against each citations network,  $G_{orig}$ ,  $G_{no\_agr}$ , and  $G_{no\_agr\_inss\_stm}$ , without perturbation. This evaluation by a legal expert was essential for doing a quality assessment of results. We will discuss the findings for each network in the subsections below.

### 5.2.1. Decision network $G_{orig}$

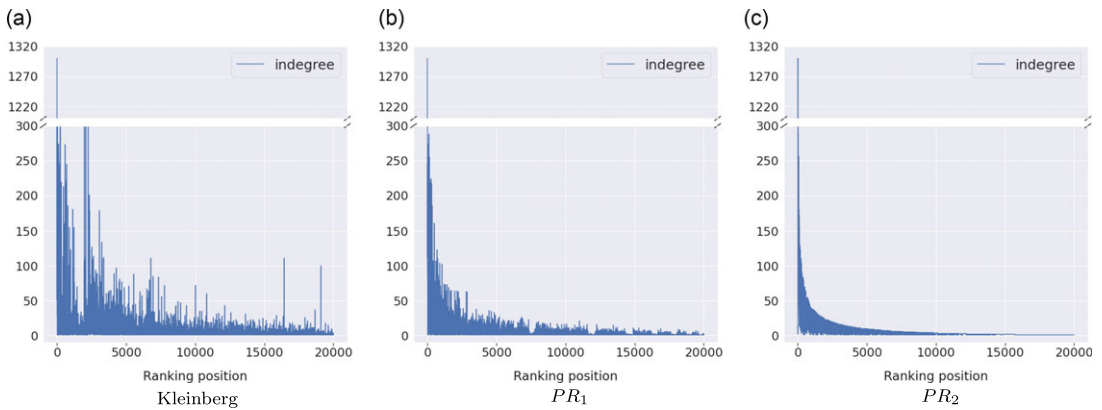
In our previous work, we have found that a high number of decisions are 'AgRs' for all algorithms studied in network  $G_{orig}$  which is undesirable because this category of decisions is a procedural appeal of little relevance. That happened because, as mentioned in Section 3, AgRs are filed in large quantities and decided by the judges in the same way, in batches of tens or hundreds of appeals and some of them share many citations, inflating their respective relevance. Kleinberg's algorithm retrieves just a few leading cases related to social security matters that reach the STF very often. However, these leading cases in particular are not so important and they are overrepresented because the same matter is discussed in many cases that reach STF, and their importance is boosted maybe because they are in the same cluster.

### 5.2.2. Decision network $G_{no\_agr}$

We have found that  $PR_2$  retrieved many leading cases and some decisions of legal relevance. However, there is a leading case in the top 20 positions related to the recalculation of pension that cites other five similar decisions that share the same thesis and are decided in the same way, boosting their ranking to the next positions below the leading case. Other important theses that suffer from the same problem; they are decided by a leading case and replicated in other decisions in *Top100Decisions* list. This shows that the STF decides the same legal thesis repeated times, even for relevant matters, unnecessarily augmenting the decision network and unbalancing the relevance decisions have in the network.

Other leading cases appear in the *Top100Decisions* list and they do not bring other similar decisions with them, which means this leading case is more distinguishable in the list compared, for example, to the leading case that is pension-related. In the *Top100Decisions* list, we also found some decisions about economic plans created during the hyperinflation era in Brazil (+ 6700% annually)<sup>15</sup> from 1988 to 1991. As the intersection of decisions between  $PR_1$  and  $PR_2$  is above 60%, as we can see in Sub Figure 1(b), the same conclusions apply for the  $PR_1$  measure.

<sup>15</sup>For data about IPCA, the inflation index adopted in Brazil, look at: <https://clubedospoupadores.com/ipca-anual>



**Figure 5.** Indegree of best 20 000 ranked decisions by each algorithm for network  $G_{no\_agr\_inss\_stm}$

### 5.2.3. Decision network $G_{no\_agr\_inss\_stm}$

Although we found that around 14 decisions repeat the same legal thesis (billing of enrollment at public universities) in  $PR_2$  *Top100Decisions* list, we also made an interesting and counter-intuitive discovery. About 35 of the 100 decisions in  $PR_2$  *Top100Decisions* list are related to different criminal cases, mostly *habeas corpus*, that settle relevant theses adopted as precedents in future cases. In other words, these criminal cases have legal relevance. This result raises the hypothesis that the STF devotes much more time than other Supreme Courts to criminal matters. Therefore, such an outcome may suggest that Brazilian criminal procedure may be distorted with respect to what is expected of it.

Kleinberg's measure *Top100Decisions* list retrieved just a few leading cases and ranked a big quantity of decisions (35) with general repercussion ('RG'), those that declare the legal thesis in dispute is relevant and when decided should apply to similar cases. Then, the STF selects one or a few cases that stay in the Court as representatives of the legal controversy and sends the others to the lower Courts (or the lower Courts withhold the cases themselves). As soon as the STF decides the matter in hand, the lower Courts must apply the same rationale to the hundreds or thousands of cases waiting for the matter to be settled. That's why the 'RGs' end up being high in the ranking of most cited cases, although it is not the actual decision about the matter in hand, but just a determination of relevance to the court of such a matter.

Also, analyzing the data of the first 10 decisions in the *Top100Decisions* list, most of which are 'RG', we found that they were boosted by a huge amount of decisions called 'ED' ('clarification motion')<sup>16</sup>. In effect, appellants resubmit these cases as 'ED', even when the cases in question do not fulfill specific requirements which must be met for accepting a case as an ED. They do that as an attempt to revert a decision made by a single judge by making a collegiate decision-making body of STF to reanalyze these cases. Then, after analyzing these appeals, the court converts them to AgRs, which is the correct category for such appeals.

To validate some characteristics of the  $PR_1$ ,  $PR_2$ , and Kleinberg measures, we have chosen some of the most famous leading cases decided over the last years. They are ADPF 54 (abortion rights of anencephalic fetus), ADPF 132/ADI 422 (legality of same-sex marriage), and MI 670 (strike of public servants). In  $PR_1$  and  $PR_2$ , they were not in the *Top100Decisions* list because they received at most a few dozen citations, and the ranking of a decision is directly proportional to its indegree measure, the number of citations a decision receives, as can be seen in Figure 5. In Kleinberg's, although there is no direct relationship between indegree and authority ranking, these decisions did not achieve a good ranking because they were cited by decisions that did not have good hub rankings.

<sup>16</sup>ED, or 'embargos de declaração'—which could be translated as a 'clarification motion'—are appropriate when the decision has some ambiguity, obscurity or contradiction.

We have a few hypotheses for why some leading cases do not reach the top of the decision rankings. First, in a qualitative view, several leading cases may be seen as definitive decisions, for example, the right to same-sex marriage. So, such decisions become part of the accepted legal culture and do not promote litigation or opposition. Along similar lines, some relevant decisions may be cited only by low-rank decisions. As a result, they do not ‘inherit’ enough mass to promote their own rank, which affects especially Kleinberg’s method. A different hypothesis explores the fact that ranking measures actually rank poorly relevant decisions but rank well irrelevant decisions that share the same subject because they occupy most of the court’s routine.

In summary, most recent experiments employing legally inspired filtering captured a larger number of leading cases, in the view of legal experts. But, there are more experiments that can be done by using specific knowledge-guided data manipulation to find good relevant case rankings.

## 6. Contributions and future work

We have found that the idea of creating legal knowledge-based filters is key to achieving better results. They allowed the removal of decisions considered noisy, which influenced a mismatch between ranking measures and expert opinion.

Experiments made with decision networks built with these filters show that, besides the decision network being robust under perturbation, they also retrieved a larger number of leading cases, unveiling a strong presence of criminal cases in the STF.

There is now greater intersection between both PageRank-based measure variations, reaching above 40% agreement in the top 100 best-ranked decisions. However, under robustness perturbation tests, these measures tend to diverge when the decision network ceases to be robust.

We conclude that the algorithms employed achieved a good quantitative measure, and they improved substantially in ranking better decisions of legal relevance and leading cases, but they still require further refinement to retrieve a more stable and consistent list of leading cases. Some claims of Vojvodic (2012) encourage refinement in the filters to find more noisy decisions and filter them out of the decision network. This work found that relevant decisions responsible for improving and developing the law and that are of most interest to STF experts, like those concerning constitutional matters, are not in dense regions of the decision network. This happens because they are not used as binding precedents and, therefore, are usually cited by just a few other decisions. The results also show that the decision relevance studied in this work underscores an inherent relevance of a decision in the STF, which much different from that expected by legal scholars like the one found in Fowler and Jeon (2008). Therefore, these results show the need to discuss the role of STF in the Brazilian judicial system.

Possible improvements could add polarity information to the edges of the network on whether the citation supports or opposes a given decision, which could lead to more precise relevance measures. Such capability, however, would require an analysis of the context in which the citation occurred, involving a degree of automated understanding not available at the present development of NLP. A less demanding improvement may come from more legal knowledge-specific filtering, involving the analysis of how specific areas of law are dealt with by the workings of the STF; one possible step in this direction would be to identify EDs converted to AgRs, as cited in the decision text, but that retain an ED identifier, evading the filtering process.

Further improvements to this work include the development of tools to analyze the full content of the decision and identify arguments in it in favor and against it. Also, it is desirable to work alongside STF experts to create separate decision networks per legal area and obtain authority scores for each network to retrieve leading cases and other relevant decisions in such an area.

**Acknowledgements.** We thank Mayara C. Melo and Alessandro Calò for developing previous work that made it possible to push this research forward; Dr Luís Matricardi for helping a lot in understanding concepts of law and properties of STF decisions and suggesting readings regarding these subjects; Felipe Farias and the STF jurisprudence sector staff for helping with questions about the data in case entries.



This work was carried out at the Center for Artificial Intelligence (C4AI-USP), with support from the São Paulo Research Foundation (FAPESP grant 2019/07665-4) and from the IBM Corporation. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001. M. Finger was partly supported by Fapesp, processes 20/06443-5 (SPIRA) and 14/12236-1 (Animals), and CNPq grant 303609/2018-4 (PQ).

## References

- Agnoloni, T. & Pagallo, U. 2015. The case law of the Italian constitutional court, its power laws, and the Web of scholarly opinions. (Proceedings of the 15th International Conference on Artificial Intelligence and Law, Association for Computing Machinery, New York, NY, USA, pp. 151–155. URL: <http://doi.acm.org/10.1145/2746090.2746108>)
- Albert, R., Jeong, H. & Barabási, A.-L. 2000. Error and attack tolerance of complex networks. *Nature* **406**(6794), 378–382. DOI [10.1038/35019019](https://doi.org/10.1038/35019019).
- Almeida, G., Nunes, J. L. & Chaves, L. 2016. Explicando o dissenso: Uma análise empírica do comportamento judicial do supremo tribunal federal e da suprema corte dos Estados Unidos (Explaining dissent. In *An Empirical Analysis of Judicial Behaviour of the Brazilian Supreme Court and Supreme Court of the United States*), Available at SSRN: <https://ssrn.com/abstract=2580904>, or [10.2139/ssrn.2580904](https://ssrn.com/abstract=2580904).
- Barabási, A.-L. & Pósfai, M. 2016. Network science. Cambridge: Cambridge University Press, URL: Available at: <http://barabasi.com/networksciencebook/>.
- Bommarito, M. J., II, Katz, D. M., Zelner, J. L. & Fowler, J. H. 2010. Distance measures for dynamic citation networks. *Physica A: Statistical Mechanics and Its Applications* **389**(19), 4201–4208. URL: <https://www.sciencedirect.com/science/article/pii/S0378437110004930>.
- Brin, S. & Page, L. 1998. The anatomy of a large-scale hypertextual Web search engine. *Comput Netw. ISDN Syst* **30**(1-7), 107–117. DOI [10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X).
- Clauset, A., Shalizi, C. R. & Newman, M. E. J. 2009. Power-law distributions in empirical data. *SIAM Review* **51**(4), 661–703. DOI [10.1137/072F070710111](https://doi.org/10.1137/072F070710111).
- Cushman, C. 2011. *Courtwatchers: Eyewitness Accounts in Supreme Court History*. Rowman & Littlefield Publishers.
- de Souza, J. J. & Finger, M. 2020. Robust ranking of Brazilian supreme court decisions. In *Intelligent Systems. BRACIS 2020. Lecture Notes in Computer Science*, **12319**, Cerri, R., & Prati, R. C., (eds). Springer International Publishing, Cham, Switzerland, 581–594. URL: DOI [10.1007/978-3-030-61377-8\\_40](https://doi.org/10.1007/978-3-030-61377-8_40).
- Dworkin, R. 1986. *Law's Empire*. Harvard University Press, 45–86, chapter 2.
- Falcão, J., Cerdeira, P. & Arguelhes, D. 2013. I relatório do supremo em números-o múltiplo supremo. *Revista de Direito Administrativo* **262**, 399–452.
- Falcão, J., Hartmann, I. A. & Chaves, V. P. 2014. III relatório supremo em números: O supremo e o tempo. Available at: URL: <http://hdl.handle.net/10438/12055>.
- Federal, Supremo Tribunal 2022. Pesquisa de jurisprudência - STF. Accessed 17-August-2022. Available at: URL: <https://jurisprudencia.stf.jus.br/pages/search>.
- Fowler, J. H. & Jeon, S. 2008. The authority of supreme court precedent. *Social Networks* **30**(1), 16–30. URL: <https://www.sciencedirect.com/science/article/pii/S0378873307000378>.
- Green, E. 2020. No one likes amy coney barretts abortion answer. *The Atlantic*, 2020-10-13. URL: <https://www.theatlantic.com/politics/archive/2020/10/amy-coney-barrett-roe-v-wade/616702/>.
- Hagberg, A. A., Schult, D. A. & Swart, P. J. 2008. Exploring network structure, dynamics, and function using networkx. (Proceedings of the 7th Python in Science Conference, Pasadena, CA USA, pp. 11–15, URL: [http://conference.scipy.org/proceedings/SciPy2008/paper\\_2/](http://conference.scipy.org/proceedings/SciPy2008/paper_2/)).
- Hartmann, I. A., dos Santos Junior, F. A. C., Silva, F. A., Appel, O., et al. 2017. Pedidos de vista no tribunal superior eleitoral. *REI - Revista Estudos Institucionais* **3**(2), 1074–1111. DOI [10.21783/rei.v3i2.180](https://doi.org/10.21783/rei.v3i2.180).
- Irons, P. 2002. *Jim Crow's Children: The Broken Promise of the Brown Decision*. Penguin.
- Irons, P. 2006. *A People's History of the Supreme Court: The Men and Women Whose Cases and Decisions Have Shaped Our Constitution*. Revised ed. Penguin.
- Klaus, A., Yu, S. & Plenz, D. 2011. Statistical analyses support power law distributions found in neuronal avalanches. *PLOS ONE* **6**(5), 1–12. DOI [10.1371/journal.pone.0019779](https://doi.org/10.1371/journal.pone.0019779).
- Kleinberg, J. M. 1999. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)* **46**(5), 604–632. DOI [http://doi.acm.org/10.1145/324133.324140](https://doi.org/10.1145/324133.324140).
- Mattiuzzo, L. 2018. Most-cited u.s. supreme court cases in heinonline: Part iii. Accessed on March 29 2021. Available at: <https://home.heinonline.org/blog/2018/09/most-cited-u-s-supreme-court-cases-in-heinonline-part-iii/>.
- Page, L., Brin, S., Motwani, R. & Winograd, T. 1999. The PageRank citation ranking: Bringing order to the Web. Technical Report 1999-66. Available at: <http://ilpubs.stanford.edu>. Stanford InfoLab. Previous number = SIDL-WP-1999-0120.
- Patterson, J. T. 2004. Legacies and lessons. In *The Landmark School Desegregation Case in Retrospect, Supreme Court Historical Society*, Cushman, C., Urofsky, M. I., eds, Black, Hite, W., & Rowen, B., 297–318.
- Rehnquist, W. H. 2002. *The Supreme Court*. Vintage.
- Schneider, C. M., Moreira, A. A., Andrade, J. S., Havlin, S. & Herrmann, H. J. 2011. Mitigation of malicious attacks on networks. *Proceedings of the National Academy of Sciences* **108**(10), 3838–3841. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1009440108>.
- Sedgewick, R. & Wayne, K. 2011. *Algorithms*. 4th ed. Addison-Wesley.

- Steinman, A. N. 2016. The rise and fall of plausibility pleading? *Vanderbilt Law Review* **69**(2753142), 333–401. Available at SSRN: <https://ssrn.com/abstract=2753142>.
- van Opijnen, M. 2012. Citation analysis and beyond: In search of indicators measuring case law importance. In *Legal Knowledge and Information Systems - JURIX 2012*, **250**, Schäfer, B., (eds). IOS Press BV, Amsterdam, The Netherlands, 95–104. DOI [10.3233/978-1-61499-167-0-95](https://doi.org/10.3233/978-1-61499-167-0-95).
- Vojvodic, A.d M. 2012. Precedentes e argumentação no supremo tribunal federal: Entre a vinculação ao passado e a sinalização para o futuro, URL: PhD thesis. Universidade de São Paulo Law School, <https://teses.usp.br/teses/disponiveis/2/2134/tde-27092012-094000/pt-br.php>.
- Winkels, R., de Ruyter, J. & Kroese, H. 2011. Determining authority of dutch case law. In *Legal Knowledge and Information Systems - JURIX 2011: The Twenty-Fourth Annual Conference*, **235**, Atkinson, K., (eds). University of Vienna, Austria, 103–112. *Frontiers in Artificial Intelligence and Applications*, IOS Press. DOI [10.3233/978-1-60750-981-3-103](https://doi.org/10.3233/978-1-60750-981-3-103). URL: 2011-12.