Computer Law & Security Review 59 (2025) 106210

Contents lists available at ScienceDirect

Computer Law & Security Review: The International Journal of Technology Law and Practice

journal homepage: www.elsevier.com/locate/clsr

Process Mining for legal Courts: Visualising, analysing and comparing Italian divorce proceedings ☆

Vittoria Caponecchia, Bernardo D'Agostino, Sima Sarv Ahrabi, Giovanni Comandè, Daniele Licari, Andrea Vandin *

Sant'Anna School of Advanced Studies, Piazza Martiri della Libertà 33, Pisa, 56127, Italy

ARTICLE INFO

Keywords:
Legal process analysis
Judicial performance
Process Mining
Process discovery
Conformance checking

ABSTRACT

Process Mining (PM) is a family of data-driven techniques that use data to study the underlying processes generating the data, i.e., the data-generating process. Despite being initially tailored for the engineering and industrial domain, it is becoming popular also in more human-centric domains like the legal and healthcare ones. We present a PM methodology using the **fuzzy miner technique** aimed at analysing and optimising the complex processes underlying decision making by legal Courts. We consider specifically the domain of civil proceedings, with a focus on divorces. In PM terms, we see a legal proceeding as a process instance, and the different internal phases in which a legal proceeding transits as activities. The studied process is, therefore, the internal process followed by a Court, possibly varying over the years, to handle specific types of proceedings. By leveraging PM techniques, this article compares consensual divorce proceedings within a Court across time, and across Courts. As a case study we take two Courts in Northern Italy. Our PM analysis identifies key performance indicators and uncovers hidden process efficiencies and inefficiencies. The findings highlight the ability of PM to reveal critical process patterns, enabling organisations to make data-driven decisions and implement targeted process improvements.

1. Introduction

The Italian Code of Civil Procedure (here in after: CPC) does not impose formal time bounds within which Courts shall decide on a proceeding.[1] Thislack of strict deadlines contributes to why proceedings often last a long time, even though the principle of *reasonable duration of proceedings* shall be respected, as dictated by Article 111 of the Constitution, already recognised by Article 6 ECHR and Article 47 of the EU Charter of Fundamental Rights, and valid for any type of trial. The commitment to reasonable duration is not only a right of the person involved in the trial, but also an objective guarantee of the proper functioning of justice.

The adoption of data science tools, such as Process Mining (PM) offers a promising approach to address such delays by identifying inefficiencies in procedural workflows and supporting targeted improvements. PM is a family of data-driven techniques that enables the monitoring, enhancement, and validation of the actual process underlying a system [1].

This article analyses and compares the internal processes of two Courts from northern Italy in handling consensual divorce proceedings originating from 2013 to 2021. In particular, we study three general research questions regarding the ability of PM in: (*RQ1*) identifying differences in the management of these proceedings inefficiencies across time, (*RQ2*) identifying differences across different Courts; and (*RQ3*) identifying virtuous practices that could be adopted on a wider scale. Through this analysis, we demonstrate the usefulness of applying PM to study the process underlying Italian legal proceedings.

We consider a new dataset that we have built by downloading data of proceedings from the "Portale dei Servizi Telematici del Ministero della Giustizia" (portal of telematic services of the Italian Ministry of Justice).[2] By leveraging the **fuzzy miner technique** as implemented

☆ This article is part of a Special issue entitled: 'Legal Systems' published in Computer Law & Security Review: The International Journal of Technology Law and Practice.

* Corresponding author.
E-mail address: andrea.vandin@santannapisa.it (A. Vandin).

[1] Except for the labour rite, which, however, does not apply to divorce proceedings, studied in this paper.

[2] We obtained the data from https://pst.giustizia.it/PST/it/services. The copyright does not allow us to provide the dataset or refinements of it. In particular, the data is available at https://servizipst.giustizia.it/PST/it/pst_2_6_1.wp. Each case has its own page with events clearly given in text. From the above link, one can pick the Italian region, and the 'office' (i.e., the Court or similar). Then, in 'Ruolo generale', one has to write the year of interest and the case number (our case ids, where cases are numbered simply by counting them, resetting the counter to 0 at the end of each year).

https://doi.org/10.1016/j.clsr.2025.106210

in Disco [2], we visualise and analyse the processes underlying these proceedings, and graphically compare proceedings across time as well as across Courts. By utilising the open-source pm4py library [3], we delve deeper into the structure of such processes, and compare them using more advanced techniques. In particular, we formally study and quantify *differences* among the processes followed by different legal Courts.

We view legal proceedings as processes , and as such, we consider Courts analogous to business organisations, where project management has consistently demonstrated its effectiveness. We believe that the use of PM in this field could help addressing challenges related to resource management within Courts, often arising from a shortage of personnel. Without compromising the quality of justice or the right to a fair trial merely to reduce processing times, PM can suggest how to streamline procedures by identifying inefficiencies in legal workflows, enabling a more effective analysis and resolution of issues. Our newly built dataset made it possible to reconstruct the activities and processes of these proceedings. By comparing proceedings from different years, we identified procedural enhancements that solved important bottlenecks. By comparing proceedings from different Courts, we highlighted procedural differences and examined how they may influence performance. This work demonstrates that PM can shed light on the handling of legal proceedings, enhancing the efficiency of legal processes, and therefore of the judicial system in general.

2. Related work

Process Mining (PM) has emerged as a powerful tool for analysing and improving workflows across various domains, including healthcare and manufacturing, and is increasingly being investigated in the legal domain by analysing event logs [4,4–14]. Its application in the legal domain, however, is relatively recent and presents unique challenges due to the complexity and variability of judicial processes. This section reviews relevant studies that have applied PM techniques to legal workflows, highlighting their contributions and limitations.

Recent research has leveraged PM techniques to analyse judicial data, uncover inefficiencies, optimise procedural workflows, and enhance decision-making processes. Unger et al. [15] applied PM techniques to Brazilian court logs, identifying bottlenecks and inefficiencies in business law proceedings. Their work demonstrated the potential of PM to uncover structural issues in judicial workflows, providing actionable insights for process optimisation. Similarly, Zhang et al. [16] used PM to extract events from Chinese divorce cases, enabling the identification of dispute-related elements and supporting judicial decision-making. In the Italian context, Pernici et al. [17] analysed civil litigation proceedings using PM and machine learning techniques. Their study focused on identifying long-term inefficiencies and structural bottlenecks, offering valuable insights into reducing delays. However, unlike the present study, they did not compare proceedings across different Courts or examine temporal variations in judicial practices. Campi et al. [18] employed predictive PM to understand how different process variants influence trial duration. Their work highlighted the importance of variant analysis in judicial contexts, aligning with the present study's focus on identifying dominant process paths and their impact on efficiency.

The application of PM in the legal domain often involves the use of specialised tools and techniques. For instance, the fuzzy miner technique [19], has been widely used for visualising complex workflows. This method simplifies process maps by focusing on the most relevant structures, making it particularly suitable for analysing judicial processes. For example, the authors of [13] use it to visualise the Italian lawmaking workflow. The present study leverages this technique to compare divorce proceedings across time and Courts. Advanced PM techniques, such as conformance checking, have also been applied to legal workflows. Carmona et al. [20] provide a comprehensive overview of conformance checking methods, highlighting their ability to assess procedural consistency and identify deviations. The present study extends this approach by using Petri nets to perform alignment-based conformance checking, offering a rigorous comparison of process models across Courts. Legal processes are inherently complex, involving multiple stakeholders, procedural variations, and case-specific requirements. Several studies have addressed these challenges by integrating PM with other analytical methods. For example, López et al. [10] proposed a framework for business process compliance using reference models of law, enabling the alignment of legal workflows with normative requirements. Similarly, Racis et al. [7] explored the use of PM for process reengineering in Courts, emphasising its role in public value creation.

Relation with [21]. This work revises and extends the preliminary workshop version recently presented in [21]. In [21], we applied (only) the fuzzy miner technique to analyse divorce proceedings in (only) a single Italian court, focusing on performance indicators and visualising trends over time. This study extends that work in two main ways: (i) it introduces model-based analysis using Petri nets and conformance checking, and (ii) it expands the scope to include two courts, enabling institutional comparison. In addition, we revised the data pre-processing as well as the analyses and the whole text. In doing so, we contribute to the growing interest in legal PM, especially regarding structural and temporal variation in judicial practices.

3. Domain and case study

We overview the expected structure of the *process underlying* the proceedings of consensual divorces in Italy, used as a case study in this paper.

Divorces may be either consensual or judicial. In the former case, there is an agreement among spouses on the conditions to be adopted, and the proceeding will be filed jointly. In the latter case there is no agreement, and the proceeding may be filed by either spouse. In this work, we focus on consensual divorces, which we simply call 'divorces'.[3]

Fig. 1 graphically depicts the main activities (the internal *statuses* in which proceedings evolve) in a consensual divorce. The labels of the boxes contain the names of the activities.

We will discuss the other numerical information in the figure in the next sections. The request for a divorce is filed in the Court by means of a written summon (GENERAL REGISTER ENROLLMENT). The Chancellery then assigns the proceeding to a section of the Court (ASSIGNMENT TO SECTION), and, in particular, to a judge of that section (JUDGE APPOINTMENT). After this, a date will be set for the first hearing of the spouses before the Court in the designated panel (SCHEDULING CHAMBER HEARING).

On the scheduled date, the Court, after hearing the parties and receiving the opinion of the public prosecutor (DOCUMENTS RETURNED BY PROSECUTOR), pronounces the divorce by judgment (CASE UNDER DELIBERATION). The judgement is first prepared with its minutes by the reporting judge (FILING OF DRAFT FINAL JUDGMENT) and submitted to the President of the panel for her/his signature. After this, the Chancellery will take care of the publication (DEPOSIT OF PUBLISHED DECISION).

At any time during the proceeding, it might be necessary to replace a judge (JUDGE REPLACEMENT) for reasons related, for example, to her/his unavailability, resignation, change of duties, etc. This is not shown in Fig. 1 because it is an irregular event which may not happen, or may happen more than once in a proceeding.

Actually, proceedings might have different *variations* of such events, like REPLACEMENT OF JUDGE FOR THE HEARING, and REPLACEMENT OF JUDGE AND SECTION. To simplify the analysis, we unified these variations under a

[3] Technically, 'divorce' is an improper name. Law 898/1970, which introduced it, never uses it. The law refers to 'dissolution of marriage' (for civil marriages) and to 'cessation of the civil effects of marriage' (for religious marriages).

single label: JUDGE REPLACEMENT, as explained in Section 4.2. Notably, in the next sections we will see that this activity might actually be very frequent, and a major factor for the length of proceedings. In particular, we introduced the activity FIRST JUDGE ASSIGNMENT to represent rare cases where the judge replacement and the judge appointment occur on the same day. This artificial activity was defined using domain knowledge and is further discussed in Section 4.2. We also performed a data enhancement step by relabelling part of the JUDGE REPLACEMENT activities as FIRST JUDGE ASSIGNMENT.

4. Dataset description and pre-processing

In this section, we describe our dataset from a computational and procedural point of view (Section 4.1), we detail the domain-specific pre-processing steps we applied (Section 4.2), and we conduct an exploratory data analysis (Section 4.3).

4.1. Dataset description

The dataset has been built by downloading metadata of proceedings from the "Portale dei Servizi Telematici del Ministero della Giustizia" (portal of telematic services of the Italian Ministry of Justice) [22]. We start considering only one specific Court in northern Italy, while starting from Section 6.1 we will consider also a second one. We focus on consensual divorces within the time-frame 2013–2021. To better identify structural differences in the proceedings across different years, we only considered concluded proceedings that started within the first five months of each year. The obtained dataset regards 2 300 distinct divorce proceedings. It contains 28 693 events over 58 unique types of activities like those discussed in Section 3.

4.2. Domain-driven data cleaning

Before analysing this dataset, it is necessary to apply several *domain-specific* pre-processing steps. Essentially, this can be considered a data-cleaning step.

One issue in this dataset is the presence of *spurious* activities occurring after the actual conclusion of proceedings, which are irrelevant to this study. From our domain knowledge, we know that a proceeding can be considered completed after the activity DEPOSIT OF PUBLISHED DECISION for the divorces. Most of the proceedings in our dataset contained additional activities after those final ones, preventing us from making any reasonable analysis of the duration of the proceedings. These are activities such as ANNOTATION, which involve small changes or notes added to the records of the proceedings, or TRANSFER TO ARCHIVE, which indicate that the document or proceeding is being transferred to the archive.

Based on this understanding, we removed all activities occurring after the"final activity" mentioned above. As anticipated in Section 3, another issue regards groups of activities with different names, but similar meanings. We identified such groups of activities and consolidated them under a single one. In particular, activities JUDGE APPOINTMENT and JUDGE DESIGNATION AND SCHEDULING OF INITIAL HEARING were merged as JUDGE APPOINTMENT. Similarly, the entries REPLACEMENT OF JUDGE FOR THE HEARING, JUDGE REPLACEMENT, and REPLACEMENT OF JUDGE AND SECTION were consolidated in JUDGE REPLACEMENT. Finally, the activities HEARING POSTPONEMENT, HEARING POSTPONED EX-OFFICIO and HEARING RESCHEDULED were merged as HEARING POSTPONED.

We then noticed how in many proceedings the activities JUDGE APPOINTMENT and JUDGE REPLACEMENT occur on the same day. We remark that real substitutions are very unlikely to happen so frequently within a day of the original appointment. One explanation could be that, e.g., proceedings have an initial temporary automatic assignment, which is then refined and updated. To adequately distinguish between true and virtual substitutions, we decided to rename all JUDGE REPLACEMENT activities that happened on the same day of a JUDGE APPOINTMENT to FIRST JUDGE ASSIGNMENT. After these data cleaning steps, we obtained 47 unique activities for the 2 300 divorce proceedings. We then refined the dataset by focusing on a subset of *relevant* activities. Specifically, we identified the 30 most frequent activities in the dataset, and then we manually analysed and selected a subset that we deemed representative for the considered proceedings. Overall, we obtained a list of 15 activities for divorces. The list of all such 30 activities is available in Appendix A together with a brief discussion on each of them. Briefly, the 15 activities that we discarded are those that were not considered essential to the reconstruction of divorce proceedings, e.g., CORRECTION (of marriage date, section, object, etc.), FILING OF GENERAL MOTIONS, etc. The reason is that these are all generic and eventual activities, referring mainly to the correction of information related to the case.

Finally, since, as discussed, our dataset only contains the day on which an activity occurred, and not the time and duration, we encountered proceedings where multiple activities took place on the same day, possibly written in a different order proceeding by proceeding. This led to erroneous loops in the mined processes. To address this,

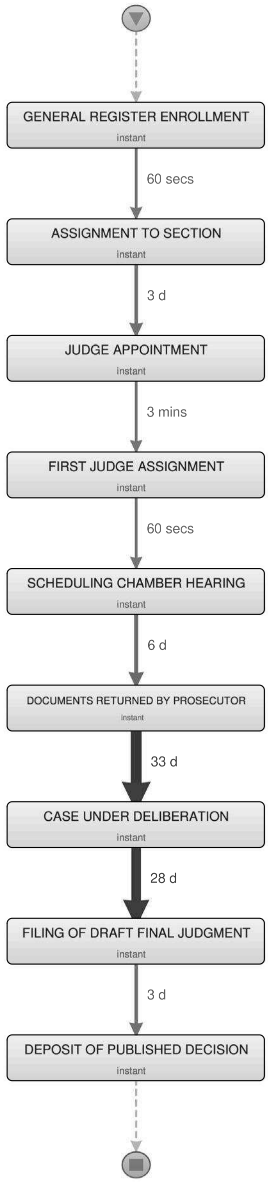

Fig. 1. A consensual divorce proceeding visualised by Fluxicon Disco. Labels, colour and width of edges denote times.

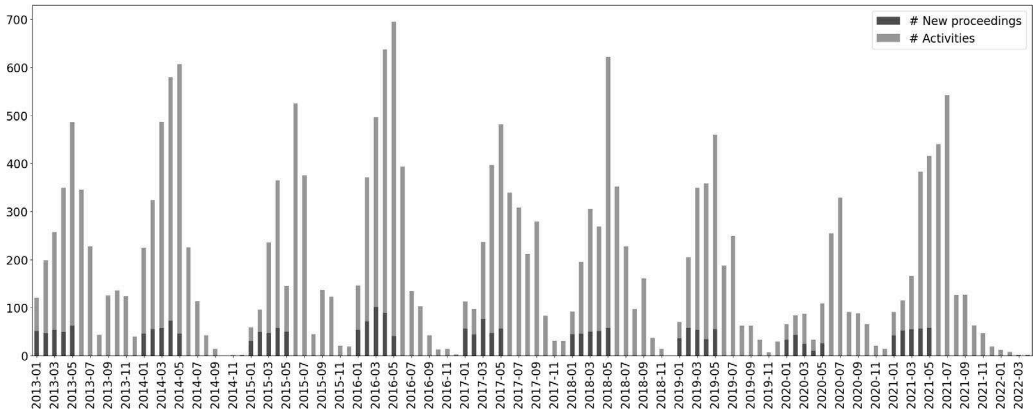

Fig. 2. Frequencies of divorce proceedings and of their activities for each month in the Court. Due to space constraints, we write only every second x-tick label.

Table 1
Duration (in days) and number of divorce proceedings in the Court.

Year	Median length (d)	# divorces
2013	181	263
2014	82	277
2015	159	235
2016	96	357
2017	155	280
2018	120	250
2019	93	238
2020	170	136
2021	140	264

we used domain knowledge to define a fixed order of activities and to consistently sort those occurring on the same day. We perform this by adjusting the beginning time from the default 00:00 by adding as many minutes as the position of the activity in this predefined order. For example, Fig. 1 shows that `GENERAL REGISTER ENROLLMENT` happens 1 min before `ASSIGNMENT TO SECTION` in the considered proceeding. Indeed, we imposed an ordering where the former is first, and the latter is second (therefore, we added 1 min to its beginning time).

After completing all pre-processing and data cleaning steps, we obtained 21 052 events for the 15 selected activities related to 2 300 divorce proceedings.

4.3. Exploratory data analysis

Before delving into the PM capabilities, we present an exploratory data analysis on our curated dataset to provide context.

Table 1 summarises the number of divorce proceedings in the Court for each year along with their median duration. We can see that the median duration fluctuates across the years and that 2020 was the year with the lowest number of divorces and the second-highest median duration. In particular, in Fig. 2 we can see that the number of new proceedings in the first two months of 2020 was in line with that of previous years, while there has been a significant drop in the following three months. This decline is likely attributable to the COVID-19 outbreak and the corresponding restrictions implemented in Italy starting from the 9th of March. In addition, Table 1 also indicates a return to pre-pandemic trends in terms of the number of proceedings and associated waiting times in 2021, suggesting a mitigation of the pandemic's impact.

5. Process mining legal proceedings within a court using process maps

This analysis aims to explore how PM techniques can reveal patterns and performance indicators within legal proceedings from a single Court. We focus on understanding typical process paths and identifying possible inefficiencies. We use process maps discovered through the fuzzy miner technique, as it provides an effective summary of behaviour in complex event logs while preserving the most relevant structures. We begin by presenting two examples of PM scenarios that illustrate and evaluate our data.

We visualise and analyse the considered legal proceedings using the **fuzzy miner technique** [19], as implemented in Disco [2].[4] Indeed, Fig. 1 shows a process map resulting from this method for a specific proceeding. It shows how one consensual divorce proceeding can be visualised using the so-called *performance view*. The performance view depicts the activities appearing in the data (in this case the internal statuses of a proceeding), as well as the waiting times between them (the labels of the edges). Boxes represent the activities and have the label *instant* because we only know the day in which a proceeding changed in the corresponding status. In particular, for each activity of each proceeding, our dataset contains the date when it occurred, the name of the activity, and the identifier of the specific proceeding to which the activity belongs. Feeding the technique with more proceedings enables PM tasks and, in particular, a number of visualisation and analysis features. For example, it computes a number of statistics on the proceedings that can be used as alternative label edges (e.g., min/max/average/median times or counts of proceedings making that status transition). Another example is that the technique can group proceedings in different *variants*, where a variant is a unique sequence of intermediate statuses of a proceeding. The proceeding in Fig. 1 belongs to the most common variant for divorces in our dataset (after applying the data cleaning steps discussed next). Considering the temporal information in the map, we can see, e.g., that it changes within the same day from `GENERAL REGISTER ENROLLMENT` to `ASSIGNMENT TO SECTION`.[5] Instead, it took 3 days to further change into `JUDGE APPOINTMENT`.

In Section 5.1, we compare the proceedings belonging to the variant shown in Fig. 1 with those from a different variant. Then, Section 5.2 offers a comparison of proceedings started in two years of interest.

5.1. Analysis of two variants

As discussed, the method can categorise proceedings into variants. The proceeding in Fig. 1 belongs to the most common one in the dataset. In Fig. 3(a) we show an example of a proceeding belonging to the 10th most frequent variants. We selected this variant to illustrate how PM techniques can uncover meaningful differences in procedural

[4] In Section 7 we will use also the open-source library pm4py [3].

[5] The label should be *0 days*. We will explain later why we modified the dataset adding seconds and minutes during pre-processing.

flows. This proceeding follows the same path as the one from Fig. 1 up to FIRST JUDGE ASSIGNMENT. Then, after 42 days, it progresses to the DOCUMENTS RETURNED BY PROSECUTOR activity. Two days later the hearing is postponed (HEARING POSTPONED), and 2 further days later the judge is replaced (JUDGE REPLACEMENT). After 12 days, the proceeding progresses to the decision phase (CASE UNDER DELIBERATION). Then, after 14 days the draft of the final judgement is submitted (FILING OF DRAFT FINAL JUDGMENT). Finally, on the same day, the proceeding terminates with the publication of the judgment (DEPOSIT OF PUBLISHED DECISION).

In Fig. 3(b) we show an analysis of all the proceedings in our dataset that follow the two mentioned variants. In particular, we apply process discovery to generate a process map[6] for all such proceedings. Overall, the two variants include 389 proceedings, 338 from the variant considered in Fig. 1, and 31 from the newly considered one.

Differently from Figs. 1, 3(b) shows a *frequency view*. Rather than focusing on temporal and performance aspects, it counts the number of times an activity is met by proceedings (numbers in the boxes) and the number of times proceedings change among two statuses (labels of edges). The visualisation of process frequencies provides a quantitative overview of the most common process flows and their relative frequencies. This information can be used to identify critical control points within the process and to assess the impact of potential process changes. By looking at the top edge, we can see that all proceedings start from GENERAL REGISTER ENROLLMENT. Furthermore, the three following down-stream edges confirm that all 389 proceedings share the same initial path. Then, from FIRST JUDGE ASSIGNMENT, 358 proceedings move to SCHEDULING CHAMBER HEARING. Indeed, all 358 proceedings from the variant in Fig. 1 have such a transition, while the 31 from Fig. 3(a) do not have it. Indeed, the 31 proceedings from the 10th variants move directly to the DOCUMENTS RETURNED BY PROSECUTOR activity. Then the comparison of the two variants reveals a significant divergence in the path taken by a minority of proceedings after the DOCUMENTS RETURNED BY PROSECUTOR activity. While the majority of proceedings move directly to CASE UNDER DELIBERATION, a smaller subset undergoes a series of additional steps, first having the hearing postponed (HEARING POSTPONED) and then the judge replaced (JUDGE REPLACEMENT). This atypical path suggests potential complexities or exceptional circumstances affecting these proceedings. The two paths merge back in the activity CASE UNDER DELIBERATION and conclude in the same way. By combining the analysis of process variants with frequency-based insights, domain experts can gain a comprehensive understanding of the underlying process dynamics and identify areas for improvement in the Courts.

5.2. PM-based comparison of proceedings from 2013 and 2019

This section employs PM techniques to perform a comparison of proceedings from 2013 and 2019, representing the earliest and latest available data points prior to the COVID-19 pandemic. These two years were selected as they delineate the temporal boundaries of the available dataset before the pandemic, allowing for a meaningful long-term comparison.

Quantitative comparison. As shown in Table 1, we have similar numbers of proceedings in the two years, 263 in 2013, and 238 in 2019. However, we find a notable difference: the median duration of proceedings in 2019 is about half that from 2013, denoting sensible improvements in the efficiency of the handling of these proceedings. We use PM to shed light on the “organisational and structural reasons" beyond these improvements.

We first analysed the process variants present in the two years. We found 32 variants in 2019 and 43 in 2013. Despite the number of variants being similar, Table 2, which describes the 4 most common variants for both years, highlights an important aspect: 2019 has a much more homogeneous treatment of proceedings. In fact, the most common variant covers 62% of the proceedings, with the others being at most 6%. Instead, there is much less difference among the first three variants from 2013, which together cover a ratio of 63%, slightly above to the top one in 2019 (62%). The comparative analysis of divorce proceedings from 2013 and 2019 reveals a substantial improvement in process efficiency over time. The halving of median processing time in 2019 compared to 2013 underscores a significant transformation in the Court’s operational practices. PM techniques have enabled a deeper understanding of these changes by examining process variations. The emergence of a dominant process variant in 2019, encompassing a significantly larger proportion of proceedings compared to the previous year, indicates a notable increment in process homogeneity. This standardisation is likely a key factor contributing to the observed efficiency gains.

Table 2
Coverage and median time (in days) of different variants for 2013 and 2019.

Variant	2013		2019	
	Coverage	Median length (d)	Coverage	Median length (d)
1	29%	156	62%	88
2	24%	204	6%	110
3	10%	131	4%	112
4	5%	247	4%	98

Procedural comparison. We then proceeded to analyse graphically the proceedings from the two years. In Figs. 4 and 5 we see the resulting performance views for these two sets of proceedings. The performance view was selected utilising the median waiting times as metrics. Notably, the process discovery tool provides a *zooming* functionality which allows focusing on more frequent activities and edges.[7] For the maps in Figs. 4 and 5 we decided to use the settings of 60% for activities and 50% for edges to ignore uncommon edges and activities in this analysis.

The process maps in Figs. 4 and 5 incorporate secondary labels. Indeed, so far, we have focused our performance analyses on the median waiting times among activities. Here we incorporate secondary metrics into the map (secondary labels). In particular, we add frequency information, thereby combining insights from both the performance and frequency views. For example, the first activity for the year 2013 GENERAL REGISTER ENROLLMENT is found in all 263 proceedings, and the edge outgoing from it is taken by all of them.

Considering aspects in common, we can see that, in both years, all edges outgoing from CASE UNDER DELIBERATION take a considerable amount of time, among the highest ones in the process. This means that the activity CASE UNDER DELIBERATION itself takes a considerable amount of time.

Notably, the secondary metrics tell us that these *time consuming transitions* are particularly important in the overall performance of the Court because they are *frequent*: all proceedings in 2019 involve one of such transitions. Instead, e.g., the transition from HEARING POSTPONED to CASE UNDER DELIBERATION has a much lower impact because it involves only 22 proceedings in 2019.

From a legal perspective, apart from what is established for labour proceedings,[8] there are no express prohibitions concerning the postponement of hearings. However, following the requirements of Article 111 of the Constitution and the principles of concentration and efficiency governing the trial, there is a tendency towards not allowing mere postponements except in cases where this is inevitable.[9] This aims

[6] We used Disco to build a directly-follows graph from the event log and filtered out infrequent paths and activities using Disco’s zooming capabilities: 60% for activities, and 50% for edges.

[7] See, e.g., https://fluxicon.com/book/read/mapview/

[8] Article 420 of CPC expressly prohibits mere postponements of hearings, but this does not apply to cognitive proceedings.

[9] E.g., Judgement No. 26935 of October 24, 2018, Supreme Court of Cassation, Civil Section, “Not only in the labour rite... but also in the ordinary rite of cognition mere postponement hearings are not allowed.”

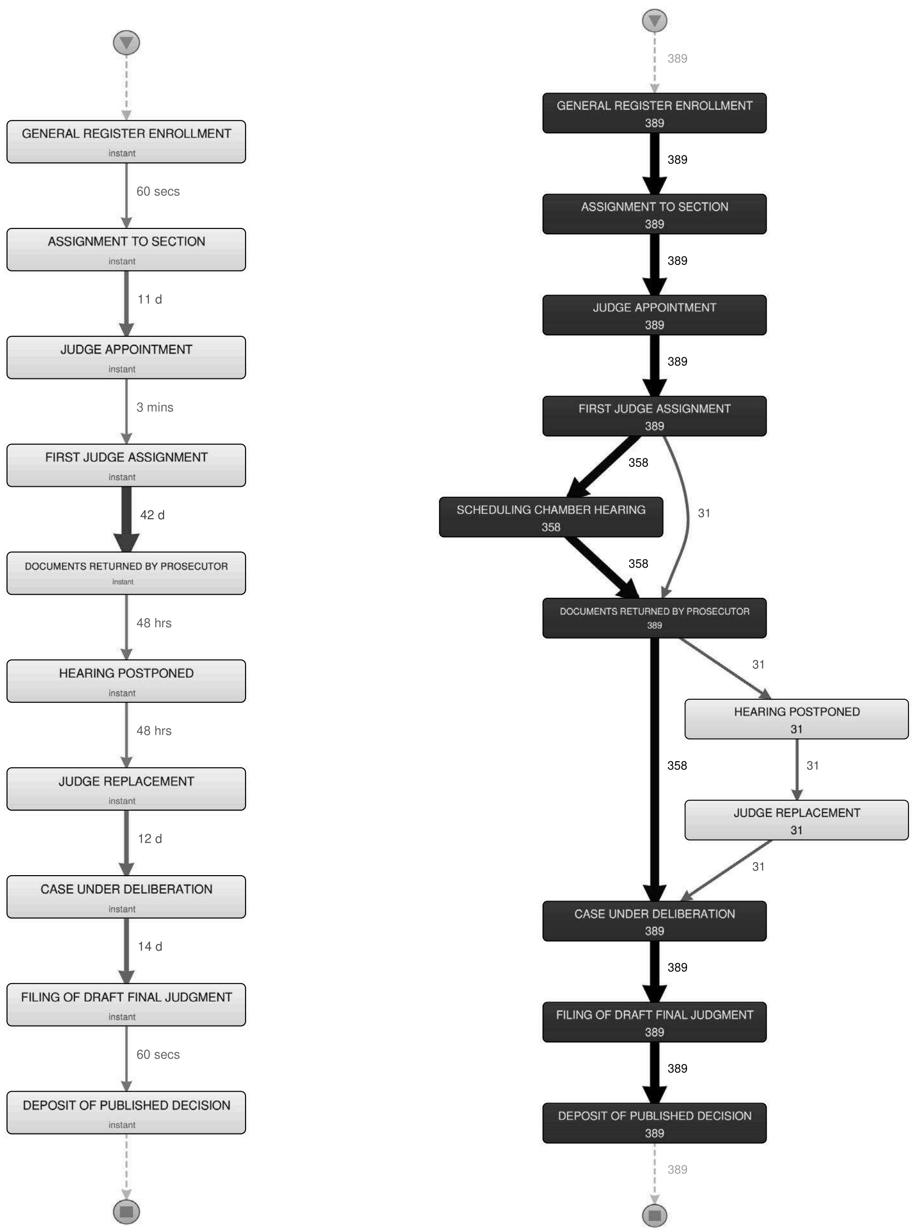

(a) Performance view of one specific variant	(b) Frequency view of the same variant

Fig. 3. Mining of one variant of divorce proceedings using Disco. Fig. 3a shows the performance view, while Fig. 3b shows the frequency view of the same variant.

to avoid lengthy trials with the consequent waste of money. This is probably why, in 2019, far fewer proceedings were postponed to a new hearing and moved more quickly to `CASE UNDER DELIBERATION`.

Two key differences emerge from the process maps. First, appointing a judge to a proceeding in 2019 appears to take nearly twice as long compared to 2013. Second, one of the most common activities in the two years, namely `JUDGE REPLACEMENT`, has a much different role in 2019. In 2013, it was central in the process, with incoming edges with very high times. Instead, in 2019 this activity is much less common with only 81 instances compared to 288 in 2013. Only one of its incoming edges has a high median waiting time of 32 days, and that edge is traversed by just 18 proceedings, resulting in a significant reduction in processing time. These observations suggest an overall improvement in process efficiency, despite the apparent slowdown in the initial judge assignment stage. During the period considered in this study, Law No. 162/2014 came into force and, for this reason, a brief discussion is warranted. It allows spouses seeking a consensual separation or a divorce to either negotiate an agreement with the assistance of a lawyer or, under certain conditions, to formalise the agreement directly before a Civil Registrar. The Law falls within the period covered by our analysis and has likely contributed to a reduction in the number of divorce cases handled by the courts. However, our work specifically focuses on joint divorce proceedings that remain within the court system. Within this subset, we observed structural improvements over time, such as changes in procedures and shorter processing times, that cannot be explained solely by a decrease in caseload.

Finally, we note how in 2019 there are many instances of the activity `FIRST JUDGE ASSIGNMENT` which we have created to denote cases

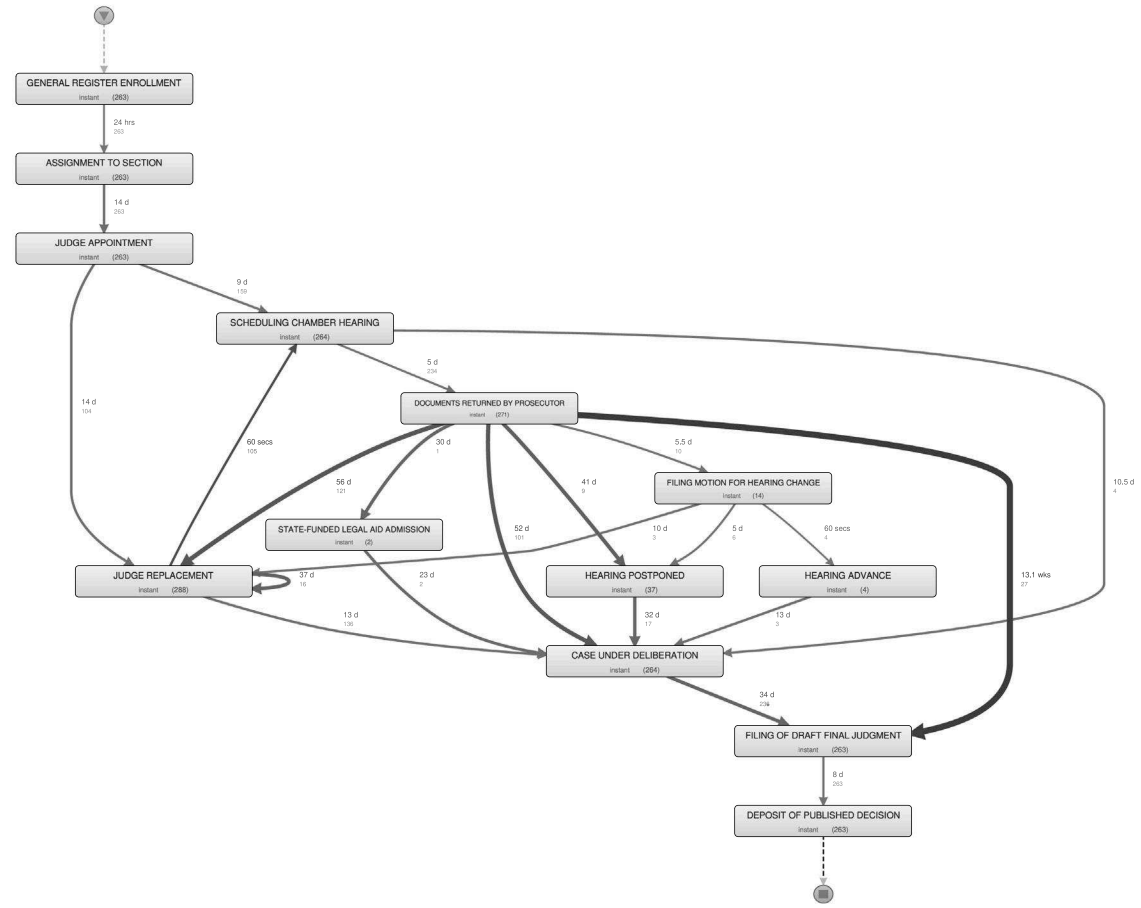

Fig. 4. Median waiting time and frequencies of the proceedings started in 2013.

where a judge substitution happens on the same day of a `JUDGE APPOINTMENT`. This did not happen in 2013 as procedures where instrumented differently.

Key findings from intra-court analysis. Beyond structural and behavioural changes across years, one specific inefficiency in the 2019 process is the 64-day gap between the `FILING OF DRAFT FINAL JUDGMENT` and the `DEPOSIT OF PUBLIC DECISION`. This significant delay suggests a post-deliberation bottleneck, possibly due to human resource limitations, suboptimal internal procedures, or document finalisation issues. Possible improvements may include reinforcing digital tools for managing procedural tasks and transmitting documents, optimising existing systems, or increasing administrative support. Future assessments, such as structured interviews or staff surveys, could help confirm these causes and guide appropriate interventions. This analysis demonstrates the effectiveness of PM in revealing process variations and performance changes over time. By providing a visual and quantitative representation of process execution, the process discovery technique supports the identification of bottlenecks and improvement opportunities, ultimately contributing to more efficient and streamlined legal procedures.

6. Process mining legal proceedings across courts using process maps

We now focus on how differences between Courts are reflected in the structure and performance of their legal proceedings.

6.1. The proceedings from court B

In this section, we consider an additional Court from northern Italy. The intent is to assess the potential of PM in revealing differences across Courts. For the sake of clarity, we refer to the Court analysed so far as Court A, and to the new one as Court B. The dataset for Court B has been built as well by downloading data from the same web portal used for Court A. However, for Court B we restrict only to proceedings started in the first 5 months of 2019. In fact, in the previous sections, we found that 2019 is an interesting case study for Court A. The dataset from Court B contains 1 326 events across 35 unique activity types, belonging to 96 proceedings.

This new dataset was preprocessed following the same steps detailed in Section 4.2. Out of the 15 activities selected for Court A, all but one remained the most relevant also for Court B. The only activity out of the selected subset not present in Court B is `FIRST JUDGE ASSIGNMENT`. In fact, we did not find any `JUDGE REPLACEMENT` activity that could be renamed to `FIRST JUDGE ASSIGNMENT`.

After the necessary pre-processing steps, we obtained 874 events related to 96 divorce proceedings.

6.2. PM-based comparison of proceedings across courts

This section employs PM techniques to perform a comparison of proceedings from Courts A and B during the year 2019. In particular, as in Section 5.2, we conduct a quantitative comparison of process variants, and a graphical analysis of process maps.

Quantitative comparison. From Fig. 6 we can see that in 2019 Court B handles fewer proceedings than Court A, (96 as opposed to 238) and has a significantly worse performance, with a median duration for proceedings of 155 days, about 50% higher than the 93 days of Court A. From Table 3 we can also see that proceedings from Court B are much more fragmented, as the top 3 variants in Court B together account only for 66% of the total proceedings, just slightly more than the

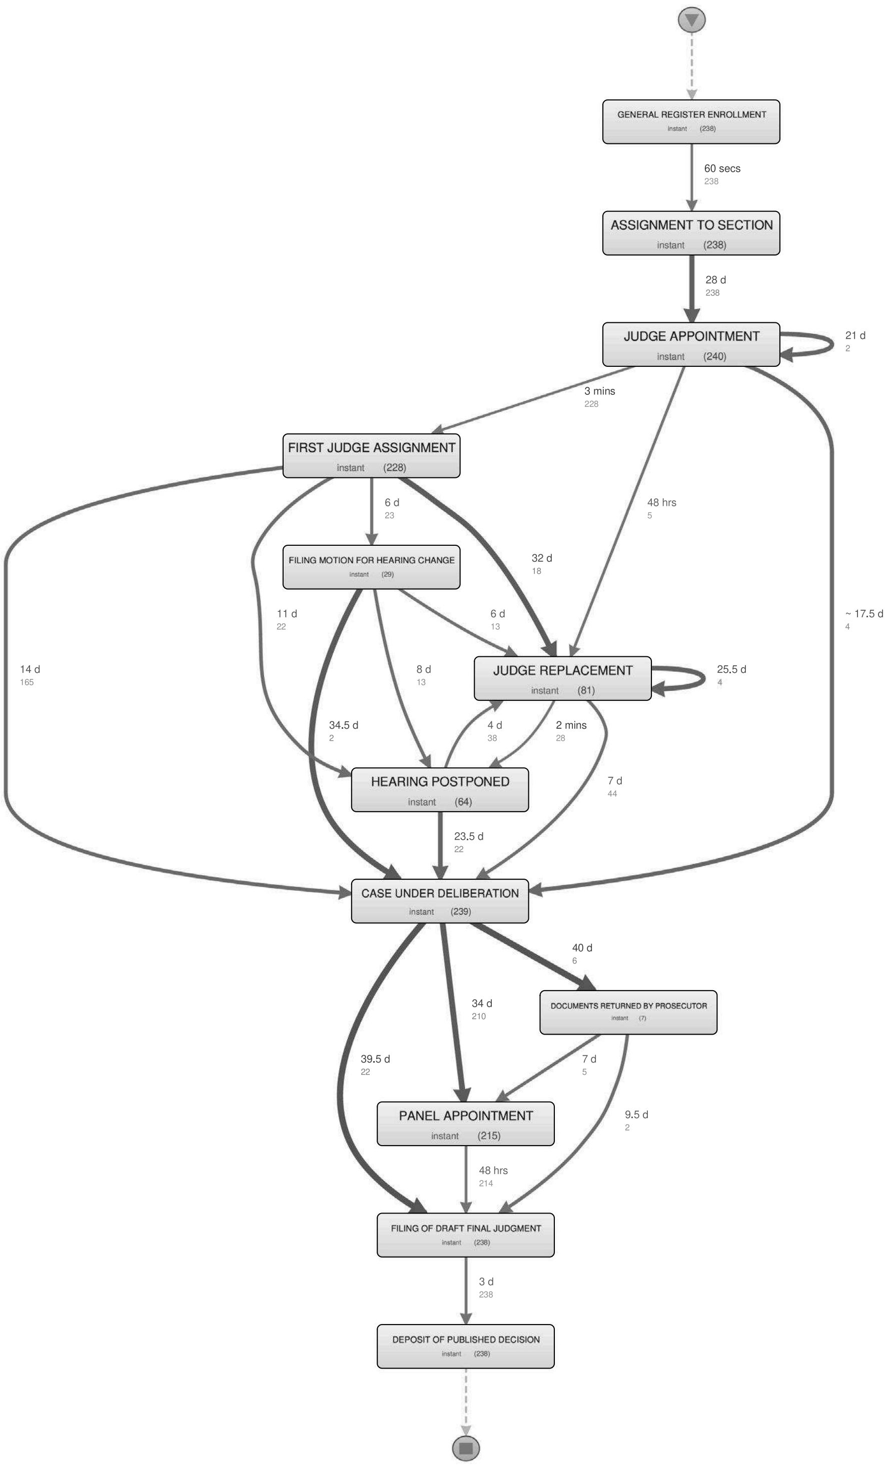

Fig. 5. Median waiting time and frequencies of the proceedings started in 2019.

coverage of the most represented variant from Court A. These findings strengthen our observations from Section 5.2, suggesting that the lower standardisation in proceedings from Court B could be a key factor in its worst performance.

Procedural comparison. We then proceeded to analyse graphically the proceedings from the two Courts in 2019. In Figs. 5 and 7 we see the resulting performance views. As in Section 5.2, the process map in Fig. 7 was generated by selecting median waiting time as the primary metric and absolute frequency as the secondary one. The zooming functionality for activities and edges was also set to 60% and 50%, respectively.

From the two process maps we can see some key differences in how the Courts managed divorce proceedings in 2019. Court B is significantly slower in the initial stages of the proceeding. Specifically, moving from the initial activity `GENERAL REGISTER ENROLLMENT` to `JUDGE APPOINTMENT` in median takes 49 days, while in Court A the same path is completed in 28 days.

An even greater discrepancy emerges when considering the median waiting time of proceedings when moving from the `JUDGE APPOINTMENT`

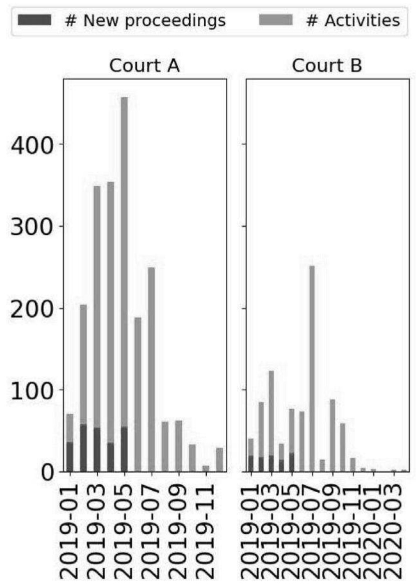

Court	Year	Median length (d)	# divorces
A	2019	93	238
B	2019	155	96

Fig. 6. Divorce proceedings in 2019 for both Courts. (Left) Frequencies of divorce proceedings and their activities for each month (with only every second x-tick label shown for space reasons). (Right) Duration and number of divorce proceedings.

Table 3
Coverage and Median length (in days) of different variants for Court A and Court B in 2019. For the sake of readability, we provide here the numbers for Court A already given in Table 2.

Variant	Court A		Court B	
	Coverage	Median length (d)	Coverage	Median length (d)
1	62%	88	36%	144
2	6%	110	18%	160
3	4%	112	12%	144
4	4%	98	7%	245

activity to CASE UNDER DELIBERATION. Proceedings from Court B need a median time of 90 days to move from one activity to the other, compared to the 16 days needed in Court A. This difference in median time can be attributed to the different paths taken by the two Courts. In Court A, the path is much more standardised with most proceedings moving immediately to FIRST JUDGE ASSIGNMENT and 14 days afterward into CASE UNDER DELIBERATION. The proceedings which follow different paths do so because of judge or hearing changes (represented by the activities JUDGE REPLACEMENT, FILING MOTION FOR HEARING CHANGE and HEARING POSTPONED). On the contrary, in Court B the process structure between JUDGE APPOINTMENT and CASE UNDER DELIBERATION is much less standardised. After the activity, DOCUMENTS RETURNED BY PROSECUTOR we can see two main branching paths followed by most proceedings in Court B. On one path, 46 proceedings move to PANEL APPOINTMENT, while on the other, 38 proceedings move to CASE UNDER DELIBERATION, with both branches exhibiting long median waiting times. This can be attributed to the fact that the activities PANEL APPOINTMENT and DOCUMENTS RETURNED BY PROSECUTOR are executed in this phase of the proceeding rather than after, CASE UNDER DELIBERATION as in Court A.

Interestingly, Court B performed better in the final phase of the proceedings, transitioning from CASE UNDER DELIBERATION to FILING OF DRAFT FINAL JUDGMENT in a much quicker and unified way than Court A. This may be a consequence of Court B handling the PANEL APPOINTMENT and DOCUMENTS RETURNED BY PROSECUTOR activities before deliberation.

We can also notice that the activity JUDGE REPLACEMENT occurs significantly less frequently in Court B. However, this does not lead to a substantial performance improvement with respect to Court A, where these activities are efficiently managed, as emphasised in Section 5.2.

Key findings from inter-court analysis. This analysis reveals that the two Courts follow different strategies in managing proceedings. Court A has a better performance, likely attributable to its more standardised and unified approach to handling proceedings.

7. Conformance checking of legal proceedings across courts using Petri nets

This section aims to assess how well the actual behaviour observed in each Court aligns with the expected process models. This helps identify deviations in execution and assess procedural consistency. To perform this analysis, we use the pm4py library [3] to discover Petri nets, which are richer and more expressive models. Petri nets enable advanced PM techniques, such as alignment-based conformance checking, that are not possible with descriptive process maps. First of all we overview the notion of Petri nets from a legal perspective (Section 7.1). Then, we use pm4py to mine Petri nets for the two courts and compare them by means of conformance checking (Section 7.2).

7.1. Petri nets and their patterns in legal proceedings

Petri nets are a mathematical model used to represent and analyse systems in several domains of science and engineering [23]. Petri nets are one of the most used models in PM to represent processes [1]. They are useful to shed more light on the structure of the process by discovering a number of 'patterns' in which the activities may be structured: *parallel*, *choice* and *optional* that we will discuss later.

A Petri net has three main elements: places, transitions, and arcs (Table 4). Places (shown as circles) represent conditions or states. Transitions (shown as rectangles) represent activities or events. Arcs (shown as arrows) connect places and transitions. Fig. 8 shows three simplified Petri nets that we have synthetically created to for presentation reasons. A token (a dot inside a place) shows in which phase the proceeding is at a given moment. For example, in Fig. 8(a) we see that the token appears in the first place. This means that the proceeding is ready for performing (i.e., to change in status) a GENERAL REGISTER ENROLLMENT. In particular, when the source place (or all the source places) of a transition has a token, then the transition is *enabled*, that is it can be executed, or *fired*. Executing a transition has the effect of consuming the token in all its source places, and of creating one token in each

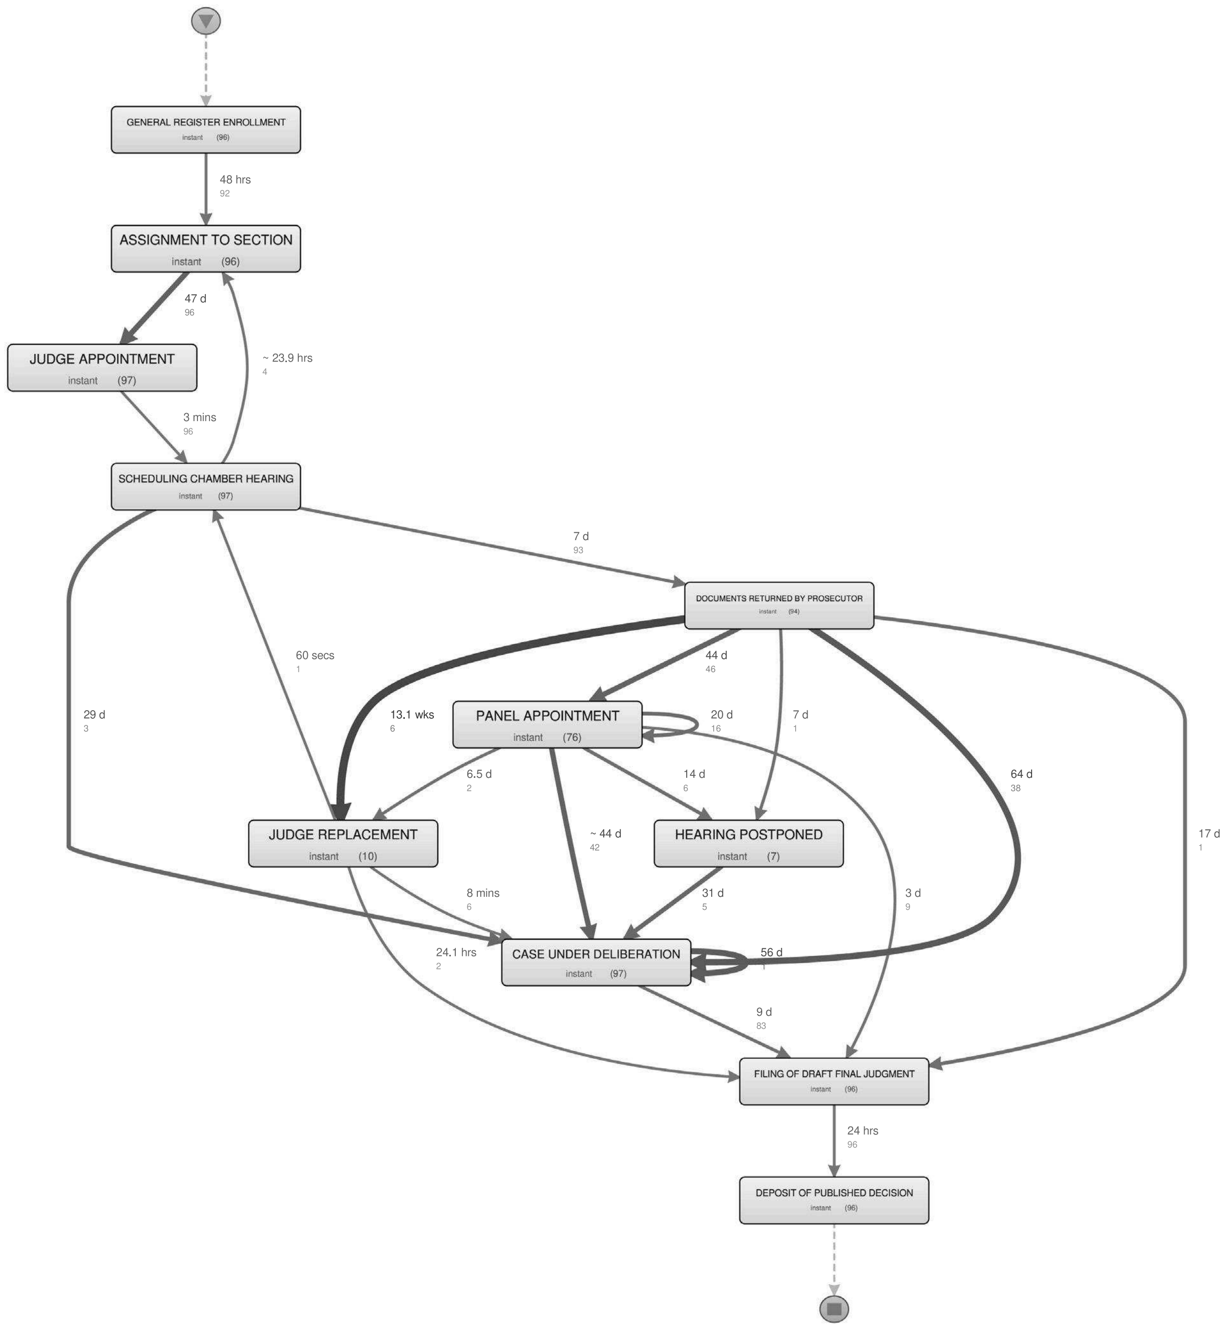

Fig. 7. Median waiting time and frequencies of the Court B proceedings started in 2019.

of its target tokens. The right-most place in Fig. 8 presents 2 circles and a squared dot. That is not a token, but a graphical representation for so-called *terminal places* which denote the end of the process when reached by a token.

Fig. 8 is also useful to discuss the three patterns mentioned above. The figure contains three simple Petri nets, each showing one of such patterns. These are made up simple Petri nets where we use real legal activities to make the logic easier to relate to. These are patterns commonly found in processes from several domains. As we will see in Section 7.2, these appear also in the Courts considered in our dataset. In the rest of this section we provide a legal-oriented discussion of the patterns:

- *Choice pattern*: In some legal procedures, the system must select one path from multiple options. For example, a proceeding might be assigned to either a single judge or to a panel of judges, but not both. The *choice* pattern (Fig. 8(a)) shows this decision point.

Table 4
Petri net elements.

Symbol	Meaning
○ (circle)	Place (condition or state)
□ (rectangle)	Transition (activity or event)
→(arrow)	Arc (flow between place and transition)
•(dot inside a place)	Token (current position in the process)
■ (filled rectangle)	Silent transition (routing only)

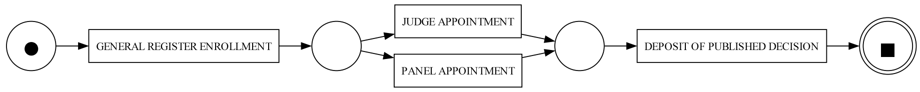

(a) Choice: Either a judge or a panel is appointed.

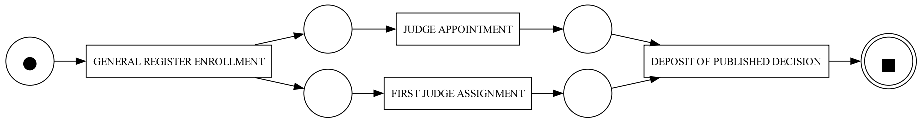

(b) Parallel: A judge appointment and a first judge assignment shall both happen, in any order.

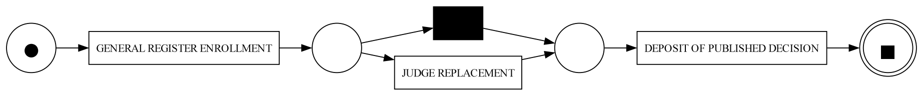

(c) Optional: A judge replacement could occur or not.

Fig. 8. Three simplified Petri nets showing common legal process patterns ; In the choice pattern (a), either a judge or a panel is appointed (not both). In the parallel pattern (b), a judge appointment and a first judge assignment must both happen, but in any order. In the optional pattern (c), judge replacement may or may not occur.

After registration (`GENERAL REGISTER ENROLMENT`), the process splits into two mutually exclusive paths, each leading to a different appointment type. This reflects real-world scenarios where resource allocation requires an either/or decision. The Petri net models this because the two transitions have the same input place. Therefore, once the token will get in such a place, only one of them will succeed in consuming the token and being executed.

- *Parallel pattern*: Certain legal steps can occur in different order, as long as all of them occur. In these cases, the *parallelism* pattern (Fig. 8(b)) allows to compactly capturing all different process executions which differ only for the order in which such independent activities happen. In the figure, we consider two types of appointments which are independent and therefore can happen in any order. In particular, the execution of `GENERAL REGISTER ENROLLMENT` will produce a token in each of its target places. These tokens will enable both `JUDGE APPOINTMENT` and `FIRST JUDGE ASSIGNMENT`. In such a configuration, both activities can be executed, in any order one at a time. After both have been executed, `DEPOSIT OF PUBLISHED DECISION` will have a token in both its input places and therefore will be enabled.
- *Optional pattern*: The optional pattern (Fig. 8(c)) represents activities (or sub-processes) that may or may not occur, such as replacing a judge due to conflicts or workload. After registration, the process can either proceed directly to the next step, or take a detour to handle the replacement. This flexibility mirrors real legal workflows where exceptions (e.g., recusals) occasionally arise. The black transition is a so-called *silent transition* whose execution does not correspond to the execution of any activity. It is just a structural element to enrich the expressive power of Petri nets.

7.2. Process structure and case handling

We now mine Petri nets for our two reference Courts. These models are automatically discovered from real data using the Heuristic Miner algorithm [24] implemented in pm4py. This method builds a model that tries to balance between conformance with the observed data (i.e., high fitness and acceptable precision) and simplicity. Figs. 9 and 10 show the Petri nets discovered for the two Courts. We limited our focus to proceedings started in the two Courts in 2019, and to their ten most frequent activities in 2019. We did not use the same 15 activities as done for the process map analysis to get models that could be shown in the paper. We also tested other process discovery algorithms, including the Alpha Miner and the Inductive Miner. The Alpha Miner failed to generate valid models on our real-world data, while the Inductive Miner produced more complex models with lower precision. For these reasons, we selected the Heuristic Miner for our main analysis, as it produced simpler, more interpretable models with better conformance metrics, especially on the Court A dataset. Detailed results and metrics comparing these algorithms are reported in Appendix B.

As visible in the mined models (Figs. 9 and 10), loops naturally emerge and capture repeated or corrective behaviour in the real process. In the model of Court A, a loop involving `HEARING POSTPONED` suggests that proceedings may return to an earlier stage before reaching deliberation or judgment. This models common delays or the need to reschedule hearings.

In Court B, a different loop connects `JUDGE APPOINTMENT` and `JUDGE REPLACEMENT`, indicating that judges may be reassigned if necessary. Such behaviour reflects real-life administrative or legal adjustments. These loops were automatically discovered by the Heuristic Miner, which encodes frequent re-entries or repetitions present in the log.

While Section 7.1 focused on simple teaching patterns such as choice, parallelism, and optional steps, the complete Petri nets also contain these more advanced constructs. Including loops gives a more realistic and complete view of the dynamic nature of legal processes.

Structural comparison. The two obtained Petri nets are shown in Figs. 9 and 10, while metrics on structural information are provided in Table 5. We observe several differences in how the models are structured, revealing unique features in how the two Courts manage proceedings. In Court A, the model contains fewer transitions than Court B, 15 versus 18, coming from the 10 selected for each Court plus a number of silent transitions. In other words, the process of Court A requires

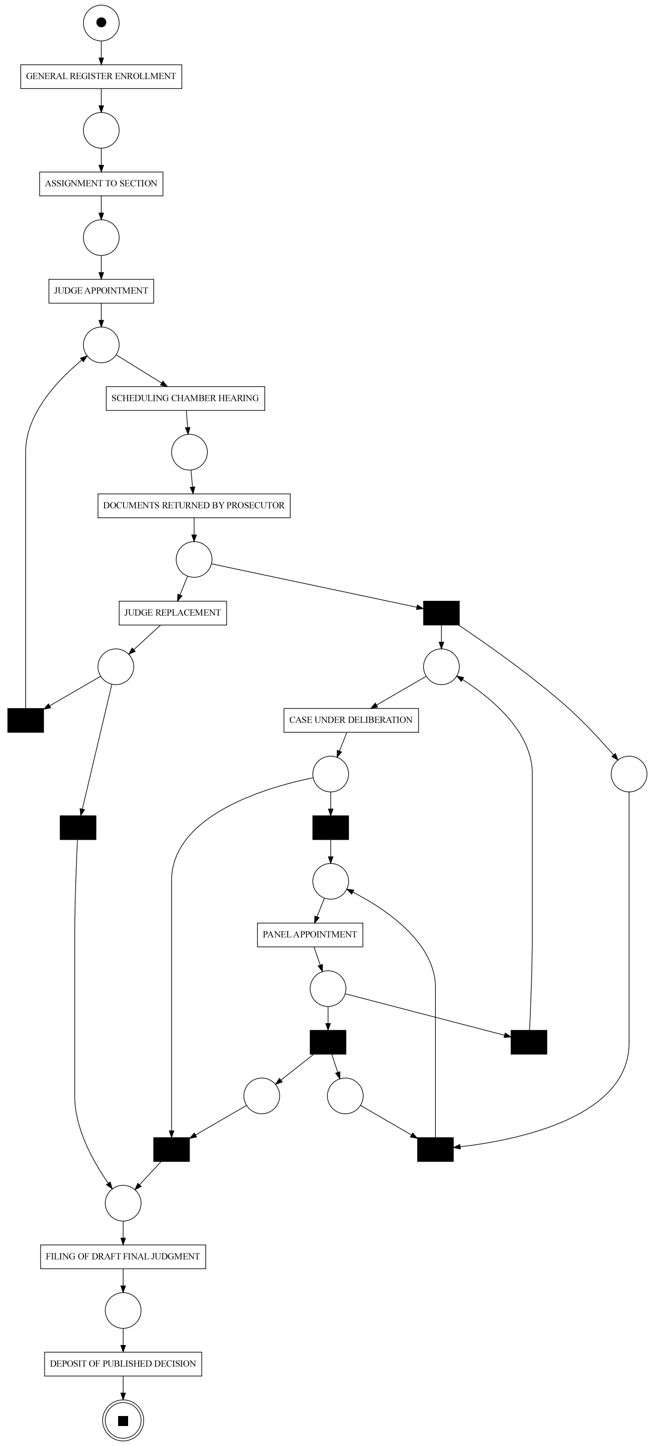

Fig. 9. Petri net for Court A (proceedings starting in 2019).

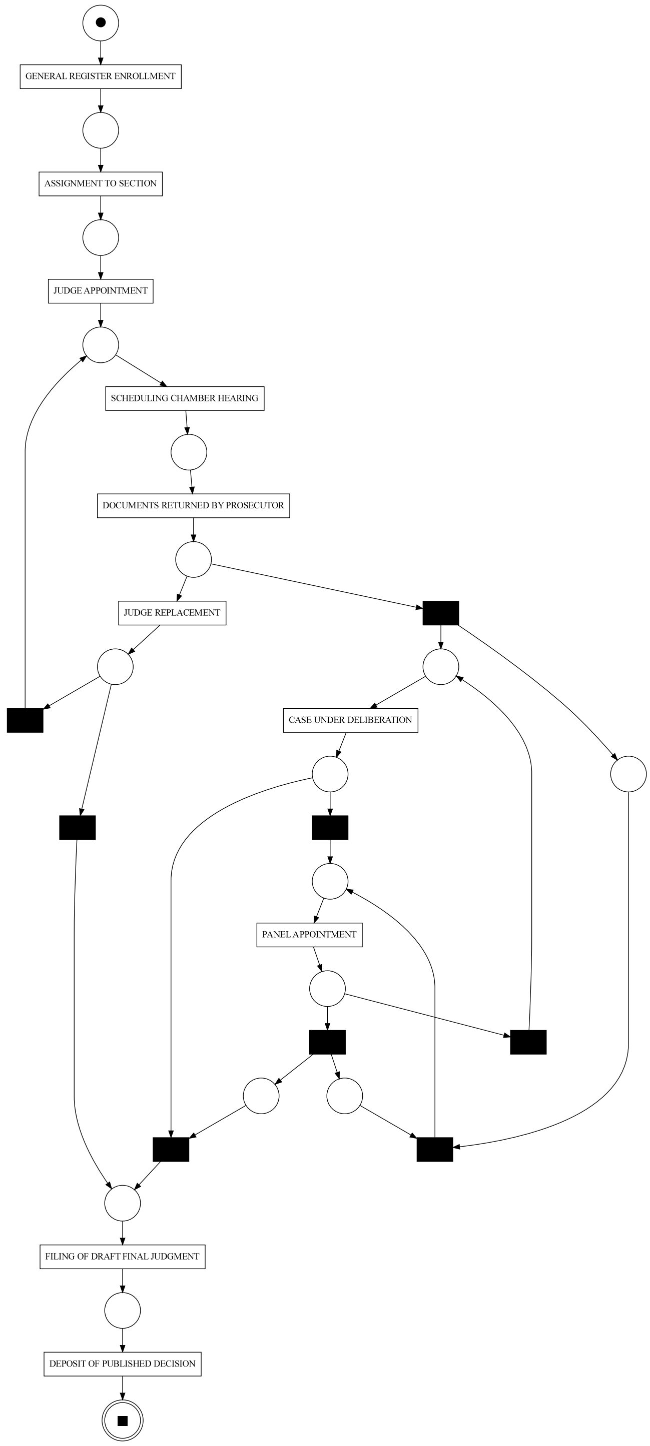

Fig. 10. Petri net for Court B (proceedings starting in 2019).

fewer silent transitions to describe it. As previously discussed, silent transitions are not connected to activities, but are just a structural element to represent more complex processes. The number of places is also slightly lower in Court A (16 compared to 17 in Court B), and the number of arcs is slightly less as well. This may suggest that the process of Court A is structurally simpler. While simplicity does not directly imply efficiency, in this context it may reflect a more streamlined or standardised handling of cases. The visual structure also reflects these differences, as in Court B's model, there is more visible branching and parallelism.

Table 5
Metrics on the structure of the two Petri nets mined for Court A and Court B using the Heuristic Miner.

Metric	Court A	Court B
Number of transitions	15	18
Number of places	16	17
Number of arcs	37	40
Number of silent transitions	5	8

Conformance checking. We now move our attention towards a more rigorous and quantitative comparison of the two processes by means of conformance checking. Conformance checking is a pivotal problem in PM, enabling the identification, analysis and fix of deviations among reference and mined processes [20]. In particular, given a Petri net and an event log (e.g., the legal proceedings from a Court), we can use conformance checking to obtain two metrics of conformance between the model and the data known as *fitness* and *precision*. Both take values in the interval $[0;1]$, where a higher value corresponds to a better conformance. In particular, fitness evaluates how much of the event log can be reproduced by the Petri net. That is, fitness evaluates the extent to which the behaviour observed in the event log can be reproduced by the model.Note that fitness does not require full trace-level matching; rather, it considers how well the control-flow of observed events is supported by the model. Instead, precision tells us how many spurious legal proceedings not present in the dataset may be erroneously generated by the Petri net. As usual in data science, there is a trade off between these two aspects, and in general, it is not possible to optimise both of them.

Table 6 provides such metrics for our dataset from the two Courts. The two top rows consider the Petri net mined on data from Court A, while the two bottom ones consider the Petri net mined on data from Court B. For each model we check its conformance with the dataset of its own Court (within Court, first rows) and of the other Court (across Courts, second rows). We perform the experiments within Courts as a sanity check to evaluate how well the model actually describes the dataset on which it has been mined. Instead, the across Courts experiments are the actual conformance checking tasks to compare the two Courts. In other words, performing conformance checking across Courts allows us to compare how well the process model of one Court can generalise or explain the behaviour observed in the other. This analysis goes beyond simple structural or frequency-based comparisons and enables a more rigorous, behavioural-level assessment. A low conformance across Courts indicates that the processes, although possibly sharing activities, are followed in different orders or combinations, which can reflect different internal workflows, resource policies, or procedural interpretations.

Key findings from conformance checking analysis. The application of conformance checking techniques using Petri nets yielded quantitative insights into the structural and behavioural differences between the legal processes of Court A and Court B. When each Petri net is evaluated on its Court data, it shows high conformance. The model from Court A achieves very high fitness (0.9563) and perfect precision (1.0000), which means that it covers almost all proceedings in the Court, and that it does not generate any spurious proceeding. The model for Court B also performs very well on its data, but with slightly lower metrics. The experiments across Courts show lower results. This is expected, as we check the conformance of a model on proceedings from a different Court. The Petri net from Court A applied to Court B's data drops in fitness (0.7353), meaning that it does not cover all proceedings in Court B. This may mean that Court B presents 'more types of handlings of proceedings', coherently with the analysis of variants in Table 3. However, the model from Court B applied to Court A's data gets even lower fitness. These two results suggest that the two Courts follow different paths or have different handling styles, in line with what was studied in Section 6.2. This hypothesis is further confirmed by looking at precision across courts, equal to 0.75 for both Petri nets. Indeed, each Court has a portion of court-specific handling of proceedings not present in the other Court.

In addition to the behavioural differences identified through conformance checking, the Petri net structure and process map together reveal potential bottlenecks in both Courts. In Court B, the 44-day delay before `CASE UNDER DELIBERATION` suggests a significant internal wait. In Court A, delays between `CASE UNDER DELIBERATION` and `FILING OF DRAFT FINAL JUDGMENT` persist despite a simple process structure. These issues appear to stem from external factors such as workload or coordination, rather than from control-flow complexity.

Table 6
Conformance checking metrics.

Evaluation setting	Fitness	Precision
Petri net discovered on logs in Court A		
On logs in Court A	0.9563	1.0000
On logs in Court B	0.7353	0.7500
Petri net discovered on logs in Court B		
On logs in Court B	0.8731	0.9275
On logs in Court A	0.6826	0.7500

8. Discussion

The excessive duration of civil proceedings has long been a critical issue in the Italian legal system, as repeatedly highlighted by the European Court of Human Rights in cases concerning the right to a fair trial within a reasonable time (Article 6 ECHR, see, e.g. [25]). Reducing delays is not only a matter of efficiency but a constitutional and human rights obligation. This study, responding to these needs, investigates how Process Mining (PM) can help analyse and improve the workflows of Italian legal courts, with a special focus on consensual divorce proceedings. The findings confirm that PM is not only useful for uncovering inefficiencies but also for comparing how different institutions operate and for offering suggestions to support judicial reform.

In the first part of the analysis (RQ1), the internal development of Court A from 2013 to 2021 was examined. Over these years, Court A made visible progress in improving its efficiency. This progress can be seen through reduced median processing times and a higher level of standardisation in its procedures. A particularly notable change occurred in 2019, when a dominant process variant began to emerge, showing that certain procedural paths had become preferred or standardised. In addition, the reduced use of time-consuming steps such as judge replacement suggests that Court A implemented targeted reforms. These likely included better allocation of resources and adjustments to its internal workflows.

The second research question (RQ2) compared Court A and Court B in the year 2019. The results revealed a strong contrast. Court A's process was more homogeneous and efficient, while Court B's was marked by fragmentation and longer durations. Interestingly, Court B performed activities such as panel appointment and documents returned by prosecutor earlier in the process, which may have caused the process to become less standardised and more unpredictable. Conformance checking further confirmed that the two courts followed different procedural logics. In fact, when the process model of one court was tested on the event log of the other, the results showed poor generalisation, which reinforced the conclusion that the structures and practices of the two courts are quite distinct.

The third part of the study (RQ3) explored the relationship between best practices and performance. The evidence highlighted that standardisation is strongly linked with better results. Court A's simpler process structure and its higher conformance metrics clearly show that having a well-defined and consistently applied workflow can lead to improved outcomes. Standardisation can help reduce delays, increase predictability, and support fairness in legal decision-making. This is one of the main goals of PM and business process management [1,26]. The positive developments observed in Court A show that targeted internal reforms, especially when supported by evidence-based tools such as PM, can contribute to the broader objective of ensuring timely justice. Although these improvements predate the implementation of the Italian National Recovery and Resilience Plan (PNRR, see, e.g., [27]), they reflect similar priorities, namely enhancing efficiency and reducing procedural delays. The existence of such local improvements suggests that bottom-up reforms, when guided by data-driven methods, may complement and inform national-level efforts. Consequently, this can increase legal certainty, which is essential for the legitimacy of the justice system.

Methodologically, the study combined both visual and formal PM tools. Process maps offered an intuitive way to see the court workflows and detect where slowdowns or variations occurred. Petri nets provided a formal structure of these processes, which allowed deeper analysis, including conformance checking to detect inconsistencies or deviations. These tools worked together to give both a high-level and detailed view of how the court processes were behaving in practice.

Finally, the findings point to several practical benefits. PM can help court administrators detect inefficiencies and make informed changes. It can also support mutual learning between institutions, as shown by the potential for Court B to adopt some of the successful strategies already used by Court A. These might include reducing procedural complexity and improving workflow consistency. PM also makes legal processes more transparent and accountable, by giving both visual and measurable insights into how justice is carried out. These insights can support legal reform by offering data-driven evidence to identify sources of inefficiency and propose targeted solutions [1,26]. Moreover, the adoption of these instruments by the courts could foster a more uniform and fair application of justice at the national level.

9. Conclusion

This research highlights the value of Process Mining as a promising tool for improving legal workflows. The results suggest that using PM techniques can lead to meaningful efficiency gains without sacrificing legal safeguards. It can help court administrators make better decisions based on real data and evidence.

At the same time, this study also shows how comparisons between courts can lead to positive institutional learning. The clear differences between Court A and Court B underline the opportunity to adopt more standardised procedures that are already working well elsewhere. In this way, PM can support not only internal optimisation but also broader improvements across the judicial system.

However, there are some limitations that need to be addressed in future work. The analysis focused on metadata and did not include qualitative elements such as case complexity or the presence of sensitive situations (e.g., cases involving minors). Including interviews or surveys with court staff in future research could help to provide a richer, more contextual understanding. Also, current conformance checking approaches are qualitative and do not take into account the stochastic (probabilistic) nature of legal processes. Recently, the PM community has started developing stochastic conformance checking algorithms that focus on this aspect [28–31], or domain-specific techniques named Diff models [32–34].

Another suggestion for future development is the use of structural similarity metrics between process models. These could help researchers and administrators better understand how much the structures of different courts differ and where improvements could be made. Finally, the creation of interactive dashboards for real-time monitoring could provide a valuable and practical tool for court managers, allowing for more effective and continuous process improvement.

Finally, while our analysis highlights timing differences between Courts (e.g., delays in judge appointment), we could not investigate the underlying causes, such as resource constraints or personnel availability, due to the limitations of the available data. This is a typical case where the findings of PM analysis should be forwarded to domain experts — such as court administrators or the President of the Court — who can interpret the results in context and take informed action.

In summary, the study confirms that PM has the potential to bring about important changes in judicial administration. It offers a data-driven way to improve efficiency, encourage fairness, and promote transparency in legal proceedings. Moving forward, the integration of qualitative data, stochastic modelling, and user-friendly tools can further strengthen the impact of PM in the legal sector.

Table A.1
30 Most frequent activities from Court A and their corresponding frequencies in Court B. Selected activities are in bold.

Activity Name	Court A	Court B
CASE UNDER DELIBERATION	2310	97
JUDGE APPOINTMENT	2302	97
ASSIGNMENT TO SECTION	2300	96
DEPOSIT OF PUBLISHED DECISION	2300	96
FILING OF DRAFT FINAL JUDGMENT	2300	96
GENERAL REGISTER ENROLLMENT	2300	96
ANNOTATION	2090	135
FIRST JUDGE ASSIGNMENT	1903	0
DOCUMENTS RETURNED BY PROSECUTOR	1660	94
PANEL APPOINTMENT	1350	76
JUDGE REPLACEMENT	927	10
SCHEDULING CHAMBER HEARING	809	97
FILING OF GENERAL MOTIONS	508	0
HEARING POSTPONED	415	7
CORRECTION OF MARRIAGE DATA	263	0
CORRECTION OF CASE DATA	139	4
FILING MOTION FOR HEARING CHANGE	129	3
JUDGE REPLACEMENT FOR THE ROLE	67	9
CORRECTION OF THE OBJECT	48	2
JUDGE CORRECTION	44	3
FILING OF UNCODED DOCUMENT	43	3
HEARING ADVANCE	41	2
UNCODED DOCUMENT	39	0
HEARING DATE CORRECTION	20	0
CORRECTION OF DATE OF DRAFT FILING	19	0
STATE-FUNDED LEGAL AID ADMISSION	17	7
CORRECTION OF THE SECTION	14	0
CORRECTION OF PROCEDURE TYPE	13	1
SCHEDULING PRESIDENTIAL HEARING	8	0
CORRECTION OF CONTRIBUTION AND VALUE	7	1

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The work has been partially supported by project SMaRT COnSTRUCT (CUP J53C24001460006), in the context of FAIR (PE0000013, CUP B53C22003630006) under the National Recovery and Resilience Plan (Mission 4, Component 2, Line of Investment 1.3) funded by the European Union - NextGenerationEU.

Appendix A. The activities selected for the proceedings in the two courts

In this section we provide more information on the activities selected for our Courts as discussed in Section 4. Table A.1 lists the mentioned 30 top activities, showing in bold the 15 actually selected ones. We now briefly discuss each such activity.

Activities considered:

1. Case Under Deliberation: the judge decides the case by issuing the divorce judgment.
2. Judge Appointment: after selecting the appropriate Section, a judge from that Section is appointed to decide the case.
3. Assignment to Section: the case, after being filed with the Chancellery, is assigned to a specific Section of the Court.
4. Deposit of Published Decision: the judgment is officially published. From this point on, the marriage is dissolved and the spouses regain single status.
5. Filing of Draft Final Judgment: the first draft of a judicial ruling, in this case a divorce judgment.

Table B.2
Conformance metrics for Court A and Court B using different process discovery algorithms.

Algorithm	Fitness		Precision		Generalisation		Simplicity	
	Court A	Court B	Court A	Court B	Court A	Court B	Court A	Court B
Inductive Miner	1.0000	1.0000	0.6177	0.5463	0.7909	0.7087	0.6941	0.6774
Heuristic Miner	0.9563	0.8731	1.0000	0.9275	0.8329	0.7634	0.7209	0.7778

6. General Register Enrollment: the parties, in this case the spouses, file the divorce request with the Chancellery.
7. First Judge Assignment: an activity we named this way. It refers to the replacement of a judge that occurs on the same day as the original appointment.
8. Documents Returned by Prosecutor: the Public Prosecutor, after reviewing the divorce terms, gives an opinion on whether or not the case can proceed.
9. Panel Appointment: designation of the judicial panel that will decide the case.
10. Judge Replacement: it occurs when a judge is replaced for any reason (e.g., unavailability, resignation, change of duties, etc.).
11. Scheduling Chamber Hearing: the first hearing before the Court with the designated judicial panel.
12. Hearing Postponed: the hearing is postponed to a later date.
13. Filing Motion for Hearing Change: a motion to change the hearing date, typically due to conflicts or impediments of the parties or lawyers.
14. Hearing Advance: the hearing is anticipated to an earlier date.
15. State-Funded Legal Aid Admission: grants low-income individuals the right to legal representation at the State's expense, for either initiating or defending a lawsuit.

Activities not Considered:

1. Annotation: entry in the margin of the marriage certificate recording the dissolution or termination of the civil effects of the marriage. In this paper, since from the data analysis it appears that it can occur at any point in the proceeding, it is considered as a kind of communication, an addition of necessary information.
2. Filing of General Motions: generic petitions or motions.
3. Correction of Marriage Data: material correction of the marriage date.
4. Correction of Case Data: material correction of certain case information.
5. Judge Replacement for the Role: judge appointed to temporarily replace a colleague, with formal assignment to the hearing schedule.
6. Correction of the Object: material correction of the subject matter of the case.
7. Judge Correction: correction of information regarding the assigned judge.
8. Filing of Uncoded Document: it occurs when, via the telematic system, a document is uploaded, but it doesn't match any predefined document type, so a generic label like "uncoded" is used. Alternatively, it can occurs when the Chancellery manually selects this label when documents of different types are filed together.
9. Uncoded Document: as described above.
10. Hearing Date Correction: correction of an erroneously indicated hearing date.
11. Correction of Date of Draft Filing: material correction of the date on which the draft judgment was filed.
12. Correction of the Section: material correction of the assigned Section.
13. Correction of Procedure Type: correction of the initial procedural classification of the case.
14. Scheduling Presidential Hearing: setting of the hearing before the President of the Court.
15. Correction of Contribution and Value: correction of the case value indicated in the initial petition and the corresponding contribution to be paid.

The above 15 activities were not considered, as previously mentioned in Section 4.2, because they are deemed irrelevant for reconstructing the divorce proceeding. This is due to the fact that they are optional phases, mostly concerning the correction of case data (e.g., value, date, names).

Appendix B. Metrics for additional mining algorithms

In this section we discuss results connected to more Petri nets obtained for the proceedings starting in 2019 in the two Courts. In the main text we presented results obtained for the heuristic miner algorithm. Here we consider also for the Alpha Miner and the Inductive Miner algorithms. The Alpha Miner created models incompatible with the event logs, and we could not calculate valid fitness values. This is a known issue with Alpha Miner when working with real-life, complex data. We also considered the Inductive Miner. Detailed metrics for the Inductive and Heuristic miners are reported in Table B.2. The Heuristic Miner showed better precision and generalisation, especially for the Court A dataset. It also produced simpler models that were easier to interpret. The table provides four conformance metrics are shown for each Court: fitness, precision, generalisation, and simplicity. As mentioned, the Alpha Miner was excluded because it could not produce valid results for the event logs.

The conformance metrics highlight a trade-off between the two process discovery techniques. Inductive Miner produces perfectly fitting models (fitness = 1.0), but at the cost of lower precision, while Heuristic Miner offers higher precision with slightly lower fitness. This pattern is more pronounced in Court B, where lower fitness and precision suggest a more fragmented process. All models were generated using the default settings in PM4Py, without additional tuning of filtering thresholds or dependency parameters.

Data availability

The data that has been used is confidential.

References

[1] Aalst Wil MP van der. Process mining. 2nd ed.. Springer; 2016.
[2] Günther Christian W, Rozinat Anne. Disco: discover your processes. In: Proceedings of demonstration track of the 10th international conference on business process management. Tallinn, Estonia: Springer-Verlag, Berlin, Heidelberg; 2012, p. 40–4, https://research.tue.nl/en/publications/disco-discover-your-processes.
[3] Berti Alessandro, van Zelst Sebastiaan J, van der Aalst Wil MP. Process mining for python (pm4py): Bridging the gap between process- and data science. 2019, CoRR, abs/1905.06169.
[4] Aalst Wil Van der, Gatta Roberto, Llatas Carlos Fernandez, et al. Process mining for healthcare. 2023, https://www.processmining4healthcare.org/, [Access 2023].
[5] Klinkmüller Christopher, Sacher Raphael, Görg Carsten, Mühlen Michael zur. A systematic review of concept drift in process mining. J Decis Syst 2022;31(4):433–55.
[6] Verçosa Luiz Felipe, Silva Vinícius, Cruz Jaqueline, Bastos-Filho Carmelo, Bezerra Byron LD. Investigation of lawsuit process duration using machine learning and process mining. Discov Anal 2024;2(1):9.

[7] Racis Serena, Spano Alessandro, Latti Giorgio. Process reengineering and public value creation: using process mining in courts. Int J Public Sect Manag 2024;38(4):482–506.

[8] Martino Beniamino Di, Cante Luigi Colucci, Esposito Antonio, Lupi Pietro, Orlando Massimo. Temporal outlier analysis of online civil trial cases based on graph and process mining techniques. Int J Big Data Intell 2021;8(1):31–46.

[9] López Hugo A. Challenges in legal process discovery. In: Marrella Andrea, Dupré Daniele Theseider, editors. Proceedings of the 1st Italian forum on business process management co-located with the 19th international conference of business process management. CEUR workshop proceedings, vol. 2952, CEUR-WS.org; 2021, p. 68–73.

[10] López Hugo A, Debois Søren, Slaats Tijs, Hildebrandt Thomas T. Business process compliance using reference models of law. In: Wehrheim Heike, Cabot Jordi, editors. Fundamental approaches to software engineering - 23rd international conference, FASE 2020, held as part of the European joint conferences on theory and practice of software, ETAPS 2020, dublin, Ireland, April (2020) 25-30, proceedings. Lecture notes in computer science, vol. 12076, Springer; 2020, p. 378–99.

[11] López Hugo A, Strømsted Rasmus, Niyodusenga Jean-Marie, Marquard Morten. Declarative process discovery: Linking process and textual views. In: Nurcan Selmin, Korthaus Axel, editors. Intelligent information systems - cAiSE forum 2021, melbourne, VIC, Australia, June 28 - July 2 2021, proceedings. Lecture notes in business information processing, vol. 424, Springer; 2021, p. 109–17.

[12] Debois Søren, López Hugo A, Slaats Tijs, Andaloussi Amine Abbad, Hildebrandt Thomas T. Chain of events: Modular process models for the law. In: Dongol Brijesh, Troubitsyna Elena, editors. Integrated formal methods - 16th international conference, IFM 2020, lugano, Switzerland, November (2020) 16-20, proceedings. Lecture notes in computer science, vol. 12546, Springer; 2020, p. 368–86.

[13] Contestabile Matilde, Ferrara Chiara, Giovannetti Alberto, Parrillo Giovanni, Vandin Andrea. From text to process: Leveraging llms to unveil italian lawmaking. In: Proceedings of the 2nd international workshop on processes, laws and compliance co-located with 23rd conference on business process management. Springer LNBIP Proceedings; 2025, to appear.

[14] Contestabile Matilde, Ferrara Chiara, Giovannetti Alberto, Parrillo Giovanni, Vandin Andrea. The prolific dataset: Leveraging LLMs to unveil the italian lawmaking process. 2025, https://arxiv.org/abs/2509.03528.

[15] Unger Adriana Jacoto, Neto José Francisco dos Santos, Fantinato Marcelo, Peres Sarajane Marques, Trecenti Julio, Hirota Renata. Process mining-enabled jurimetrics: analysis of a brazilian court's judicial performance in the business law processing. In: Proceedings of the eighteenth international conference on artificial intelligence and law. New York, NY, USA: Association for Computing Machinery; 2021, p. 240–4.

[16] Zhang Yuan, Li Chuanyi, Sheng Yu, Ge Jidong, Luo Bin. Judicial intelligent assistant system: Extracting events from chinese divorce cases to detect disputes for the judge. Expert Syst 2024;41(7).

[17] Pernici Barbara, Dilettis Marco, Gerosa Paolo, Campi Alessandro. Why are italian trials taking so long? A process mining approach. In: Falchi Fabrizio, Giannotti Fosca, Monreale Anna, Boldrini Chiara, Rinzivillo Salvatore, Colantonio Sara, editors. Proceedings of the Italia intelligenza artificiale - thematic workshops co-located with the 3rd CINI national lab AIIS conference on artificial intelligence. CEUR workshop proceedings, vol. 3486, CEUR-WS.org; 2023, p. 284–9.

[18] Campi Alessandro, Ceri Stefano, Dilettis Marco, Pernici Barbara. Variants analysis in judicial trials: Challenges and initial results. In: Machine learning and principles and practice of knowledge discovery in databases. Communications in computer and information science, vol. 2133, Cham: Springer; 2025, p. 425–38.

[19] Günther Christian W, Aalst Wil MP van der. Fuzzy mining - adaptive process simplification based on multi-perspective metrics. In: International conference on business process management. Springer; 2007, p. 328–43.

[20] Carmona Josep, Dongen Boudewijn van, Weidlich Matthias. Conformance checking: foundations, milestones and challenges. Cham: Springer International Publishing; 2022, p. 155–90.

[21] Anonimized. Towards visualizing and analysing legal proceedings with process mining. In: Genga Laura, López Hugo A, Sulis Emilio, editors. Proceedings of the 1st international workshop on processes, laws and compliance co-located with 6th international conference on process mining. CEUR workshop proceedings, vol. 3850, CEUR-WS.org; 2024, p. 46–57.

[22] Italian ministry of Justice. Portal of the telematic services of the italian ministry of justice. 2025, https://pst.giustizia.it/PST/it/services.page, [Accessed on April 2025].

[23] Murata Tadao. Petri nets: Properties, analysis and applications. Proc IEEE 1989;77(4):541–80.

[24] Weijters AJMM, Aalst WMP van der, Alves De Medeiros AK. Process mining with the HeuristicsMiner algorithm. In: BETA publicatie : working papers. Technische Universiteit Eindhoven; 2006.

[25] Maurizio de Stefano. La lunghezza della durata dei processi in italia condannata dalla corte europea dei diritti dell'uomo. 2025, http://www.dirittiuomo.it/lunghezza-dei-processi-in-italia, [Access 2025].

[26] Aalst Wil MP van der, et al. Process mining manifesto. In: Florian Daniel Kamel Barkaoui, Dustdar Schahram, editors. Business process management workshops, revised selected papers. Lecture notes in business information processing, vol. 99, Springer; 2011, p. 169–94.

[27] Osservatorio conti pubblici italiani. I tempi della giustizia civile si sono ridotti grazie al pnrr?. 2025, http://bit.ly/41uutf0, [Access August 2025].

[28] Leemans Sander JJ, Li Tian, Montali Marco, Polyvyanyy Artem. Stochastic process discovery: Can it be done optimally? In: Advanced information systems engineering: 36th international conference, cAiSE 2024, limassol, cyprus, June (2024) 3–7, proceedings. Berlin, Heidelberg: Springer-Verlag; 2024, p. 36–52.

[29] Leemans Sander JJ, Maggi Fabrizio Maria, Montali Marco. Reasoning on labelled petri nets and their dynamics in a stochastic setting. In: Business process management: 20th international conference, BPM 2022, münster, Germany, September (2022) 11–16, proceedings. Berlin, Heidelberg: Springer-Verlag; 2022, p. 324–42.

[30] Senderovich Arik, Weidlich Matthias, Yedidsion Liron, Gal Avigdor, Mandelbaum Avishai, Kadish Sarah, Bunnell Craig A. Conformance checking and performance improvement in scheduled processes: A queueing-network perspective. Inf Syst 2016;62:185–206.

[31] Incerto Emilio, Vandin Andrea, Ahrabi Sima Sarv. Stochastic conformance checking based on variable-length markov chains. J Syst Softw Appear 2025.

[32] Casaluce Roberto, Burattin Andrea, Chiaromonte Francesca, Lluch-Lafuente Alberto, Vandin Andrea. White-box validation of quantitative product lines by statistical model checking and process mining. J Syst Softw 2024;210:111983.

[33] Casaluce Roberto, Burratin Andrea, Chiaromonte Francesca, Lluch-Lafuente Alberto, Vandin Andrea. Enhancing threat model validation: A white-box approach based on statistical model checking and process mining. In: Breve Bernardo, Desolda Giuseppe, Deufemia Vincenzo, Spano Lucio Davide, editors. Proceedings of the first international workshop on detection and mitigation of cyber attacks that exploit human vuLnerabilitiES (DAMOCLES 2024) co-located with 17th international conference on advanced visual interfaces. CEUR workshop proceedings, vol. 3713, CEUR-WS.org; 2024, p. 9–20.

[34] Casaluce Roberto, Tschaikowski Max, Vandin Andrea. White-box validation of collective adaptive systems by statistical model checking and process mining. In: Margaria Tiziana, Steffen Bernhard, editors. Leveraging Applications of Formal Methods, Verification and Validation. REoCAS Colloquium in Honor of Rocco de Nicola - 12th International Symposium, ISoLA 2024, Crete, Greece, October (2024) 27-31, Proceedings, Part I. Lecture notes in computer science, vol. 15219, Springer; 2024, p. 204–22.