Machine Learning and Judicial Rulings Traffic Violations

Dušan Kuzmanović
Sector for assuring quality of regulation and services
Republic Secretariat for Public Policies
Belgrade, Republic of Serbia
dusan.kuzmanovic@rsjp.gov.rs

Ana Đilas
eJustice Department
Ministry of Justice
Belgrade, Republic of Serbia
ana.djilas@mpravde.gov.rs

Ivana Jevtić
eJustice Department
Ministry of Justice
Belgrade, Republic of Serbia
Ivana.jevtic@mpravde.gov.rs

***Abstract*— This paper presents a project focused on applying machine learning techniques to develop an assistive tool aimed at facilitating the process of issuing judgments for traffic offenses. Although the project has not yet achieved full implementation through a system-wide rollout for the automatic case management in misdemeanor courts (SIPRES), it has demonstrated significant potential for improving the efficiency and effectiveness of judicial proceedings. The presentation of the project will focus on practical factors and ethics rather than theoretical principles, law, and philosophy. Instead, it will highlight ethical considerations encountered by the authors directly within the model, data, and the proposals of individual decisions.**

Keywords—eJustice, CyberJustice, AI, Judical reform, eGovernment.

I. INTRODUCTION

Machine Learning, a subset of AI, enables computers to learn from data and improve their performance over time without explicit programming. In the legal profession, it can be used to predict case outcomes, identify relevant documents, and automate routine tasks. For example, it can analyze historical case data to estimate the likelihood of success, thereby assisting lawyers in making more informed decisions.

In the Serbian judiciary, in 2019, the Ministry of Justice, in collaboration with external partners, began developing a prototype based on machine learning, aimed at assisting misdemeanor judges in drafting traffic-related rulings. The goal of the prototype was to validate the hypothesis that, using artificial intelligence based on machine learning, a misdemeanor judge can be provided with a ruling that corresponds to a specific criminal act from a finite set of possible decision patterns. This is achieved by suggesting appropriate penalties and quantifying their measures, thereby supporting more consistent and efficient decision-making

In their work, judges faced difficulty due to the lack of available precedent from other courts; for example, a judge rarely had cases involving the determination of a monetary fine for a traffic company. Fines imposed on companies are generally stricter than those on individuals (physical persons), and the fines are prescribed within a broad range. Judges were also concerned about issuing penalties that might be drastically higher or lower than those imposed by other judges in similar situations, for example, speeding 50 km/h over the permitted limit by a commercial vehicle. Therefore, in the initial interviews, they showed interest in an auxiliary tool that would indicate how many times similar situations had resulted in certain penalties imposed by other judges

II. PROCESSING OF MISDEMEANORS WITHIN THE PROTOTYPE FRAMEWORK

The prototype addressed misdemeanors for which the law prescribes a range of penal measures, such as speeding, driving under the influence of alcohol, and failing to yield to pedestrians. Traffic police compile a report on the inspection of traffic participants and vehicles that is ultimately submitted to a magistrate. During this process, the involved party may reach an agreement with the traffic police concerning the implementation of penalties, specifically regarding the relationship between fines and protective measures. Penalty points are imposed in a fixed amount. Upon reaching an agreement, the party forfeits the right to appeal[1].

Alongside the report, the traffic police also prepare and submit a paper-based request to initiate proceedings, which is filed at the court registry and entered the information system. The case then reaches the magistrate as a docket. The case is also available to the magistrate in the information system (SIPRES), where the magistrate issues a decision to initiate proceedings, followed by summonses and hearings, except when an agreement is reached between the party and the traffic police.

The verdict in these proceedings is drafted by a clerk based on the magistrate's dictation, and the details are entered into SIPRES and stored in the system's database. The suggested range of penalties aids the magistrate in deciding consistent with judicial practices.

III. USE CASE SELECTION

The selection of this use case is the result of an analysis showing that due to the obsolescence of misdemeanor prosecution for traffic violations, the state not only misses out on revenue from fines but also incurs legal costs, as in overdue cases, defendants are acquitted and have the right to reimbursement of defense expenses. Although economic objectives should not be the primary focus in this domain, it was considered that the financial benefits would justify the investment and ensure the sustainability of the system.

Additionally, the database of traffic misdemeanors within the information system for case management (SIPRES) contained enough cases, years, and relevant metadata on traffic misdemeanors and imposed penalties, which were of adequate quality (comprehensive and accurate). The model was trained on 700.000 cases (records). Additionally, data from the inception of the system's operation were utilized, specifically records from 2015 onwards. In this regard, a database was established covering a five-year period, which was considered sufficient in terms of historical data. At the same time, the data are not outdated to the extent that they would lose relevance, ensuring that the information used remains pertinent for the development and testing of the model.

Since the Law on Road Traffic Safety was amended during the process of populating this database, specifically with stricter penalties for certain violations, the model was trained to restrict its search to decisions issued under the law in effect at the time of the violation, based on the violation's date. The modeling was correctly implemented because the dataset includes a field with the date of the violation, which is linked to the applicable regulation at that time. Consequently, decisions are associated with the legal provisions valid at the time of the violation. The fact that the legal system generally allows for applying the more favorable law in cases of conflicts regarding temporal validity, that is, applying a law that is more lenient for the defendant was not deemed significant. This is because such instances are statistically negligible given the average duration of proceedings and thus do not substantially impact the model's outcomes.

The records contained information on the type of violation, the relevant legal provisions, as well as data on penalties, including the monetary fine amount, duration of imprisonment, penalty points, duration of community service, and other data that may appear in such cases. Data regarding the specific violations such as the degree of speeding or other precise actions was not directly recorded as a separate metadata attribute. However, the legal structure itself, particularly the articles, sections, and points within the regulations, explicitly specify the range or thresholds related to each offense, such as how much the speed was exceeded. For each individual violation, the precise article, section, and point of the law that applied—serving as the basis for the sanction—are entered into the system. This detailed legal referencing allows for an accurate input of what the offender did, and the corresponding penalty assigned, ensuring clarity and precision in the data used for analysis and model training. Therefore, in this technical sense, there were sufficient assumptions for evaluating that the model developed to assess sanctions would have enough cases for effective training.

The expected results included increasing collection rates and minimizing delays due to prolonged statutory limitations, reducing the duration of misdemeanor proceedings to enhance court efficiency, and harmonizing judicial practice to standardize penal measures.

IV. DESIGN PHASE

During the design phase of this prototype, the following decisions were made. The system should display judges the proposed amounts for each penalty measure, including the average and median values of sanctions assigned for the same type of misdemeanor. Additionally, it should present the ranges of minimal and maximal values of previously imposed penalties for the specific offense. The judge should have the option to accept, modify, or completely reject the suggested values, thus enabling informed decision-making aligned with established judicial practices.

It was unequivocally decided that personal data would not be used; instead, aggregated values would be employed. Additionally, measures for information security were to be implemented, including restricting operations to a closed environment within the Protected Judicial Network. Activity logging and other security protocols were also incorporated. Data transfer from a specially isolated instance was facilitated using web service technology (SOAP), ensuring secure and controlled communication.

In this context, the risk of bias was not identified, as representative data[2] were utilized. Instead of sampling, data encompassing all traffic violations that reached the court were employed. Previously mentioned agreements among involved parties were not processed, which was later recognized as an area for potential improvement. Furthermore, the chosen approach and technology did not pose a significant risk of bias in the model selection process, thereby reducing the likelihood of favoring specific outcomes or groups.

IV. CHALLENGES

The legal and ethical challenges involved issues related to repeat offenders. In judicial practice, the fact that a defendant has previously committed a traffic violation (recidivist) is considered an aggravating circumstance and influences the decision to impose a harsher sentence. If the data on offenders were completely anonymized, it would be impossible for the model to learn that there is a trend of stricter penalties among recidivists. A solution was found in the technique of hashing[1] personal ID data. This approach was implemented in accordance with ethical AI application requirements, whereby the method (key) used for hashing and anonymizing the data was stored in a highly secure environment. As can be seen in Figure 1, in the XML file structure for the web service that populates the model with data for learning, a fixed-size string of characters is used instead of an personal ID. Instead of directly using personal identifiers (LICE_ID is still present, but it's likely an internal ID), the HASH field, a fixed-size string of characters, is used for representing the individual (LICE) in an anonymized way

[1] One-way function that transforms data into a fixed-size string of characters

<ODLUKA_ID>102827</ODLUKA_ID>
<BROJ_ODLUKE>2</BROJ_ODLUKE>
<DATUM_ODLUKE>19.05.2020</DATUM_ODLUKE
<KATEGORIJA_ODLUKE>Осуђујућа
пресуда</KATEGORIJA_ODLUKE>
<TIP_ODLUKE>Осуђујућа
пресуда</TIP_ODLUKE>
<LICA>
<LICE>
<LICE_ID>97744</LICE_ID>
<HASH>7EAD3F70-3A52-45B0-B32F-BEB88EBB0878</HASH>
<TIP_LICA>Физичко лице</TIP_LICA>
<POL>M</POL>
<RADNO_MESTO></RADNO_MESTO>
<ZANIMANJE></ZANIMANJE>
<PRAVNE_KVALIFIKACIJE>
<PRAVNA_KVALIFIKACIJA>

Figure 1. XML file structure with hashed value

Furthermore, another private info is relevant, since it is common in judicial practice to consider it a mitigating circumstance if the offender falls into a specific category of individuals, such as pensioners or military personnel. This information exists as metadata in SIPRES (the case management system) and was used by the model. It is certainly recommended that the list of individual characteristics be expanded, but it also serves as a lesson that databases should be designed to include as much relevant data as possible. The manual data entry process itself can be facilitated either through AI extraction or by achieving interoperability with systems that already possess such data (population register, database of insured persons, etc.).

This is an example of systemic bias, where the bias is consciously built into the system. For instance, the characteristic of being a retiree is considered a mitigating factor, while the fact that an individual has committed an offense again is viewed as an aggravating circumstance. These represent institutionalized biases. The presence of these attributes in the data allowed for their management within the system, ensuring that they do not disproportionately influence the pattern recognition or the model's decisions relative to their actual prevalence in practice. This approach helps maintain a balanced and fairer model by mitigating the impact of embedded biases. It has been observed that bias is present in court rulings and that this is a reality that must not be ignored [3].

The second type of bias we aimed to address in the next phase relates to individual biases[2] among judges, specifically their perceptions regarding the type and severity of sanctions, as well as how frequently they apply certain penalties. In the following paragraph, an overview of the screen interface available to them will be presented. This pertains to personal biases influencing their decision-making processes concerning the appropriate sanctions, and the interface was designed to help mitigate these biases and promote more consistent application of penalties.

Through targeted training on AI models and their operational principles, we seek to dispel misconceptions and biases related to the use of AI tools. The goal is to encourage judges to adopt AI-assisted decisions more confidently and consistently, thereby promoting wider and more uniform application of technology in judicial practice.

V. TESTING

The first evaluation of the system's functionality within the project was conducted during a workshop attended by 8 judges, who had previously received training on how machine learning operates and on the usage of the tool provided for testing. This kind of training for AI is important not only to address the identified skill gap, which has been recognized as a barrier to broader adoption in this field, but also for ethical reasons that necessitate users understanding how the tool operates to utilize it more effectively. In addition to the previously mentioned technical measures, training of users regarding AI systems and the underlying models was implemented as a strategy to mitigate the risk of bias. This educational intervention aimed to enhance user awareness and understanding of potential biases, thereby supporting responsible and informed use of AI technologies.

The proposed type of penalty and amount as example in Table 1. make it easier for the judge to make a decision that is consistent with judicial practice. The system prototype showed the judge the proposed amount for each punitive measure, the most frequently awarded amount for the same type of offense, the mean and median amount for the specific offense, and the ranges of minimum and maximum awarded punitive measures.

TABLE 1. Proposed type of penalty

Penalty	**Value**
Monetary fine	18.000
Penalty points	4
Driving ban	15
Community service (Hours)	0
Imprisonment (Days)	0

The judge had the option to accept, modify, or completely reject the suggested value, as it is presented in Figure 2. The selection of any of these options was also recorded in the model's database, contributing to its ongoing learning process.

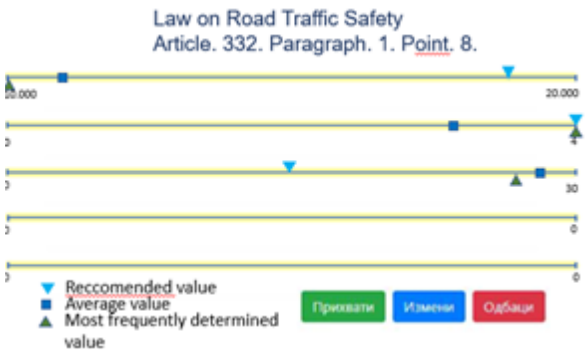

Figure 2. *Graphical user interface (GUI) offers three options, to accept, reject, or modify the proposed decision*

What is particularly sensitive in such a solution is that the judge must not be exposed to suggestions on how to decide in this way. This is considered interference with the judge's work[4]. An acceptable solution is for such a GUI and statistics on decisions to serve as a tool that the judge can call upon if they want to consult existing decisions. This was promptly noted for later integrations with the SIPRES system.

The results achieved by the initial model for suggesting sanctions are presented in the Table 2. 'Conformity' represents the percentage of accurately proposed sanctions, meaning the judges accepted them. The 'Range 20%' column indicates that the judge's proposed amendment was quantitatively within a 20% difference from the model's suggestion. The 'Difference > 20%' column indicates that the judge's proposed amendment differed by more than 20% from what the model suggested. The reasons for rejecting the recommendation were different penalty assessed in 26 cases, acquittal instead of conviction in 4 cases, decision to suspend proceedings in 1 case, plea agreement concluded in 6 cases.

TABLE 2. Results achieved by the initial model testing

Punitive Measure	**Conformity**	**Range 20%**	**Difference>20%**
Monetary fine (Cash)	73%	17%	10%
Penalty points	89%	1%	10%
Driving ban	87%	2%	11%
Community service	100%	0%	0%
Imprisonment	99.75%	0%	0.25%

It has been observed that the solution implemented in the project contributes to the expected effects and that judges' subjective perception of the model's behavior aligns with the statistical presentation of case outcomes

VI. PILOTING

The subsequent step involved piloting the tool at a single misdemeanor court through integration with a case management system that automatically generated rulings based on live case data[1].

During the pilot phase, several significant technical challenges were encountered, including communication issues between servers and time-outs caused by network outages. Additionally, instances were observed where the model suggested measures deviating from the general frameworks prescribed by law. All such issues were thoroughly documented in the event logs and promptly addressed by a dedicated team, ensuring continuous monitoring and resolution throughout the pilot period.

The analysis of these issues led to several conclusions regarding potential system enhancements, including implementing Quality of Service (QoS) on network equipment to improve communication stability. Additional recommendations involve supplementing the model with guidance for agreements, expanding the list of occupations in SIPRES, retraining the model using textual fields, introducing boundary value ranges into the model, and visually customizing the user interface to enhance usability.

The expansion of the occupation list in SIPRES was prompted by the emergence of many individuals registered under the category of unemployed during the pilot phase. This issue may also have resulted from inaccurate data entry aimed at facilitating faster processing, highlighting the need for improved data validation and standardization procedures.

The boundary value ranges integrated into the model would be derived from general provisions of the Law on Offenses, which establish the overarching legal regime applicable to all misdemeanors prescribed by special laws. An example is the Law on Road Traffic Safety, which defines specific violations related to traffic offenses. Incorporating these general legal principles ensures that the model's recommendations are aligned with the statutory framework governing misdemeanors.

This pilot phase lasted for two months. The overall statistics of tool usage after two months of the pilot phase are presented in Table 3.

TABLE 3. Results after two months of pilot phase

Category	**Number of Cases**	**Percentage**
Total cases	234	100%
Proposed suggestions accepted	181	77.4%
Proposed suggestions modified	15	6.4%
Proposed suggestions rejected	38	16.2%

Although the goal was not explicitly defined in a measurable manner regarding the expected accuracy level of responses, the pilot demonstrated a satisfactory degree of success. The sense of optimism was further reinforced by the recognition of significant potential for improvement identified during the pilot phase.

On the ground, the efficiency of procedures and optimization can be improved by enhancing the model to recognize individuals who potentially intend to delay the process to let the statute of limitations expire. This could significantly increase its effectiveness. Such an approach would require the model to utilize data on requests for postponements of hearings and the reasons for delayed hearings [1].

Furthermore, integrating additional data sources, such as assets and income, would enable more accurate assessments and higher quality of recommended measures. This would contribute to better decision-making and increased success of interventions. However, these data have been deemed as excessive processing by the Data Protection Officer, as access to such information is limited to specially authorized personnel designated by the court president.

Due to the emergency during the COVID-19 pandemic and the significantly changed needs and priorities, the project, despite receiving a positive assessment from the Ministry of Justice, did not obtain approval for further rollout to other courts or to additional areas of misdemeanors.

VII. Conclusion

Through the implementation of this pilot project and the use of the new feature initially called "Proposed Sanction" within the SIPRES software, we have concluded that it can significantly contribute, primarily, to standardizing judicial practice, i.e., penal policy, across all misdemeanor courts in the Republic of Serbia. Additionally, the system allowed for the preservation of individualization of the offender such as employment status or prior convictions while also incorporating the factors that judges consider when deciding on monetary penalties.

References

[1] M. Marjanovic, “Application of artificial intelligence in misdemeanor courts,” Belgrade, 2020.

[2] R. Schwartz, A. Vassilev, K. Greene, L. Perine, A. Burt, and P. Hall, “Towards a standard for identifying and managing bias in artificial intelligence,” Mar. 2022. doi: 10.6028/NIST.SP.1270.

[3] D. Kehl, G. Priscilla, and K. Samuel, “Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in Sentencing Citation Link,” 2017. [Online]. Available: https://accessibility.huit.harvard.edu/digital-accessibility-policy

[4] Advisory Board On Artificial Intelligence, “1st AIAB Report on the use of Artificial Intelligence (AI) in the Iudiciary based on the information contained in the Resource Centre on Cyberjustice and AI,” Strasbourg, Feb. 2025. [Online]. Available: https://www.coe.int/en/web/cepej/resource-centre-on-cyberjustice-and-ai