What to Explain and How? The Challenges and Future of Using Large Language Models to Generate Explanations for Lawyers in Autonomous Car Accidents

Yifan Xu[1,†], Xiao Zhan[2,†] and Hazel M Taylor[1,*]

[1] *The University of Manchester, United Kingdom*
[2] *King's College London, United Kingdom*

Abstract
As autonomous cars (ACs) become integrated into society, the legal challenges they introduce, particularly in accident scenarios, become increasingly complex. This systematic review evaluates 31 relevant studies to explore the potential of large language models (LLMs) in generating explanations to assist lawyers in AC-related accident cases. We assess the current capabilities of LLMs within the legal domain, identify the specific explanatory needs of lawyers, and highlight the gaps between what current LLMs can provide and what legal practice demands. Additionally, our review outlines several key challenges and suggests potential directions for future research to better align LLMs capabilities with legal explanation requirements. This paper aims to equip lawyers with improved tools for generating court-usable explanations and insights into handling AC accident cases. Our findings show the growing role of LLMs in the legal domain and suggest ways to advance legal technology in the future.

Keywords
Explanation for Lawyers, Large Language Model, Autonomous Cars

1. Introduction

The introduction of Autonomous Cars (ACs) such as the AMZ Driver-less racecar [1] and the Google Driver-less Car [2] marks a revolution in transportation. This development offers new opportunities but also introduces complex challenges, especially in the legal domain. As ACs become more prevalent in daily life, they engage with human environments in ways that raise important legal issues, especially when accidents occur. Lawyers are at the forefront of navigating these issues, with responsibilities that include litigation, regulatory compliance, policy advocacy, contract management, and addressing concerns related to ethics, privacy, and consumer protection [3, 4].

ECAI'24: Workshop on Multimodal, Affective, and Interactive Explainable AI (MAI-XAI), October 20, 2024, Santiago de Compostela, Spain
*Corresponding author.
†These authors contributed equally.
yifan.xu@manchester.ac.uk (Y. Xu); xiao.zhan@kcl.ac.uk (X. Zhan); hazel.taylor@manchester.ac.uk (H. M. Taylor)
0000-0003-2303-1531 (Y. Xu); 0000-0003-1755-0976 (X. Zhan); 0000-0003-0130-1174 (H. M. Taylor)

Traditional approaches in Explainable Artificial Intelligence (XAI) have been vital in making the decision-making processes of autonomous car systems transparent, understandable, and trustworthy [5]. This transparency is critical as it helps clarify how ACs make real-time decisions, which, in turn, aids lawyers in reconstructing accident scenarios to ascertain liability and fault. However, current XAI implementations primarily focus on the technological aspects of AC systems and do not fully address the wider range of legal requirements. Lawyers require detailed explanations to establish a causal link between the actions of an autonomous system and any resulting harm [6]. These explanations are essential for court claims against manufacturers and users of autonomous cars, aiding in the accurate allocation of responsibility [7]. Therefore, to bridge this gap, a deeper investigation into the specific types of explanations lawyers require is necessary, along with the development of advanced tools and methods. These should be capable of generating detailed, legally-aligned explanations in judicial settings, specialised to meet the expectations and demands of lawyers.

Nowadays, large language models (LLMs) such as recent versions of GPTs [8] represent an advancement in Artificial Intelligence (AI), primarily excelling at natural language understanding and generation [9]. These models have demonstrated substantial value across various fields due to their adeptness at processing and generating human-like text, as well as their capacity to recognise patterns and analyse large datasets. Such capabilities make LLMs particularly effective in providing personalised support and adapting to specific needs in diverse sectors, including healthcare [10, 11], education [12], and finance [13, 14]. The versatility and adaptability of LLMs highlight their potential to transform traditional practices by improving efficiency and accuracy in data-heavy tasks. In this case, LLMs hold promise for enhancing the way lawyers access and use explanations. Their demonstrated success in other text-generation tasks positions them as a potentially transformative tool for legal applications. This suggests a valuable opportunity to explore the capabilities of LLMs further, address existing challenges, and utilise their strengths in creating solutions to the legal domain.

In this paper, we perform a systematic review aimed at identifying the specific XAI needs of lawyers and examining how LLMs can be specialised to support lawyers in managing cases involving autonomous car accidents. Through the analysis, we make the following contributions:

1. Elicit the requirements for explanations of lawyers in AC accidents, including what needs to be explained and what is expected of the explanations.
2. Map out the capabilities of LLMs within the legal domain. Compare these capabilities against the specific XAI requirements of lawyers to identify open challenges.
3. Highlight key directions for future work in this area.

2. Method

In order to assess prior work on how LLMs generate explanations for lawyers to assist with legal issues in AC accidents, we conduct a systemic review by following the PRISMA guidelines[1] proposed by Moher et al. [15]. To achieve this, it involves identifying (1) eligibility criteria and databases searched, (2) parameters for the search and data extraction, and (3) data collection.

[1]We omitted steps from the PRISMA guidelines that are unsuitable for this research, such as summary measures across studies.

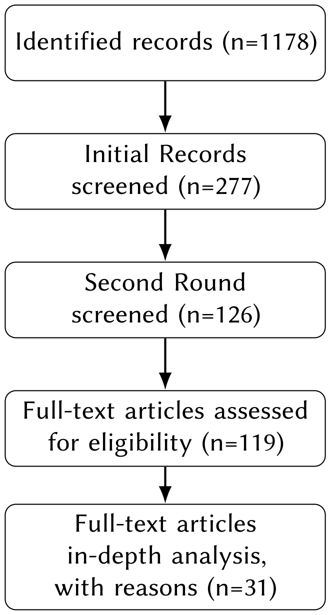

Figure 1: PRISMA flow diagram for the selection process of literature for the systematic review.

For this review, journal articles and conference papers were sourced from reputable databases such as the *ACM Digital Library*, *IEEE Xplore*, and *Web of Science* in accordance with established guidelines for conducting literature reviews [16]. To ensure a comprehensive examination of the relevant literature, *Google Scholar* was also employed.

2.1. Data Extraction

Conducting the search posed a problem: there was no existing literature that precisely reflects the use of LLMs to provide explanations for lawyers in AC accidents. Therefore, we defined the following search parameters to maximise the likelihood of identifying literature relevant to our objectives. These included: *'Explainable AI for Lawyers'*, *'Large Language Model for Lawyers'*, *'Explainable AI and Autonomous Cars'*, *'Autonomous Car and Lawyer'*, *'Autonomous Car and Law'*, *'User Study for Lawyers'*, and, *'Explanation in AI and Law'*.

2.2. Data Collection

The initial keyword search phase returned 1,178 papers. From the 1,178 papers identified, 277 were deemed potentially relevant after abstract review. These were then narrowed down to 126 papers based on their relevance. Out of these 126 articles, only 119 are accessible and downloadable for data analysis.

To identify open challenges in this field and to explore future directions for how LLMs can enhance explanations to assist lawyers in autonomous car accident cases, two of the authors revisited the collection of 119 papers and selected 31 papers which fall into the following categories for more in-depth analysis:

- *Requirements for Legal Explanations in Autonomous Car Cases*: This category encompasses any papers that discuss the specific needs and expectations lawyers have in cases involving ACs, particularly regarding the explanations they require for reference.
- *Capabilities of LLMs in Legal Domain*: Papers that discuss the existing applications of LLMs within legal contexts and how LLMs assist lawyers in their work fall into this category.

A flow diagram of the search process is given in Figure 1. Through this process, papers were scrutinised to ensure alignment with the specific focus of this paper.

Following the manual review process, two authors conducted a collaborative thematic analysis [17]. Initially, each author independently reviewed the selected papers to comprehensively understand their content and identify precise codes relevant to the designated thematic categories. Subsequently, the authors convened to discuss their individual insights and observations, engaging in iterative discussions to refine and consolidate emerging codes. Through this iterative process, a consensus was achieved on the identification of key codes that encapsulate the diverse perspectives and findings within the selected literature.

3. Results

As a result, in this paper, a total of 31 papers were thoroughly analysed (see Table 1). Specifically, 15 papers were categorised under the *'Requirements for Legal Explanations in Autonomous Car Cases'*, and 16 were categorised under the *'Capabilities of LLMs in Legal Contexts'*.

Table 1
Literature Analysed in this paper.

Category	Literature
Requirements for Legal Explanations in AC Cases	[18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32]
Capabilities of LLMs in Legal domain	[33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48]

3.1. Requirements for Legal Explanations in AC Cases

In the legal domain, explanations should be specialised to legal terminology, principles, and requirements to ensure that legal professionals can effectively interpret and act upon the AI's outputs [18, 30]. Traditional legal frameworks are designed to attribute responsibility based on human action or identifiable product flaws, often following a clear chain of causation [49]. However, with autonomous cars, decisions are often made through intricate algorithms and machine learning processes, leading to outcomes that may not align neatly with human understanding or legal precedents. As a result, determining accountability in autonomous car accidents requires a reevaluation of conventional legal concepts to accommodate the unique challenges posed by autonomous cars [32]. In this section, we summarise the key insights derived from these literature.

3.1.1. What to Explain

Explanations specialised for lawyers differ from traditional forms of XAI, such as process visualisation and the deconstruction of 'black-box models'. To ensure the effective representation of the victim's rights in legal proceedings, it is important that the lawyer receives comprehensive explanations regarding the autonomous car system. These explanations should illuminate the operational mechanics of the autonomous car, identify the specific conditions under which the system might malfunction, and highlight any potential design or operational flaws [29, 28]. Additionally, it is crucial to provide detailed insights into the system's behavior, its decision-making processes, and any documented or potential failures that could precipitate accidents or injuries. This information is essential to enable the lawyer to thoroughly understand the issues at hand, thus ensuring that the victim's rights are robustly protected throughout the legal process.

Furthermore, as emphasised in [19], the *case background* described by the parties plays a important role in providing context and understanding the core issues at stake, beyond merely examining the operational aspects of the autonomous car system. The detailed accounts offered by the parties involved, describing the circumstances leading up to and during the incident support the lawyer's judgment and decision-making processes, enabling a more informed and accurate evaluation of the case.

3.1.2. Lawyers Expectation of Explanations

Accuracy and Relevancy It is important that explanations provided by AI systems maintain a high level of accuracy and relevancy [18]. Accuracy ensures that the explanations accurately mirror the decisions made by the system, closely aligning with the algorithms and data underpinning those decisions [18]. Particularly for types of explanations like AC that incorporate reasoning—often based on case backgrounds and legal statutes — it is crucial that the source of information, such as the knowledge base, is meticulously accurate [22]. Relevancy, on the other hand, guarantees that the explanations are suitably adapted to the legal context, utilising terminology and constructs that are familiar to lawyers. This adaptation helps judges, lawyers, and jurors comprehend the AI's decisions in ways that are consistent with legal reasoning and practices [18]. Furthermore, among the 16 papers, the most frequently mentioned capability of explanation tools is their ability to efficiently and accurately process large volumes of legal documents according to lawyers' needs, thereby providing lawyers with filtered and interpreted relevant supporting documents [23, 22, 24, 27]. For such explanations, accuracy and relevancy are crucial.

Effective Information Gathering As mentioned in Section 3.1.1, a comprehensive understanding of the case background and a detailed reconstruction of the events as they unfolded are vital for lawyers to formulate explanations and mount a defense effectively. The process of reconstructing the case largely depends on the narratives provided by the parties involved. However, due to inherent variations in the parties' educational backgrounds and their understanding of legal principles, the reliability and quality of these narratives can be significantly compromised. In many instances, the information provided by the parties may be entirely irrelevant or unhelpful for legal analysis. In response to this challenge, Atkinson et al. [19] has developed an innovative explanation tool. This tool intelligently selects questions to pose to users, thereby collecting relevant information. Subsequently, it synthesises this information to generate coherent and useful explanations, enhancing the efficacy of legal assessments. This means that the ideal tool for generating lawyer-required explanations should intelligently and effectively synthesise rigorously screened information into coherent narratives that enhance legal assessments.

Faithfulness and High Quality To obtain effective explanations to lawyers, faithfulness and overall quality are two other important requirements. Faithfulness refers to how well the explanation reflects the actual reasoning process of the AI model [20]. Quality involves correctness, robustness, and the ability to simulate the AI's decision-making process effectively. Ensuring high-quality explanations is essential for fostering trust and confidence in AI applications within the legal sector, particularly when the stakes are high, and decisions must be justified comprehensively [20].

Unbiased and Ethically Sound In the context of autonomous car accidents, lawyers require explanations that are not only technically sound but also ethically considerate and free of bias [25]. From an ethical standpoint, the explanations need to be transparent, providing clear and understandable reasons behind AI decisions, which allows for accountability and fosters trust among lawyers and the public. Furthermore, it is essential that these explanations are free of any biases that could unjustly influence legal outcomes [50, 26, 31]. This involves ensuring that the AI systems do not perpetuate existing biases in data or algorithms, which could lead to discriminatory practices in legal decision-making [27].

3.2. Capabilities of LLMs in Legal Domain

To gain insights of how LLMs can assist lawyers in handling AC accident cases, we reviewed and summarised the capabilities of LLMs in the legal domain.

Ability for Summarising Extensive Legal Documents Like in other fields, LLMs are primarily utilised for tasks such as multi-document summarisation, demonstrating their potential to assist knowledge workers in managing large document collections [51]. In the legal domain, it is recognised that both lawyers and law students typically spend extensive hours per case analysing multiple relevant documents to produce high-quality summaries of key events and outcomes, thus facilitating document retrieval and use in court settings [38, 37]. LLMs can significantly expedite this process or even undertake it entirely on behalf of legal professionals, as they possess the capability to process and summarise source documents, which often exceed two hundred pages per case, by generating concise explanations.

Recent study [33, 44] has developed tools that utilise LLMs to create coherent and concise summaries that enable lawyers to quickly comprehend case essentials. Furthermore, initiatives such as the Civil Rights Litigation Clearinghouse (CRLC)[2] have been established, which provide

[2]https://clearinghouse.net

information about large-scale civil rights lawsuits to lawyers, scholars, and the general public. Compared to traditional text summarisation tools, LLMs have demonstrated the ability to categorise short text snippets from a variety of legal documents without prior specific training, showcasing a robust capacity for zero-shot learning across diverse types of legal texts [34, 40]. This capability is crucial for tasks such as document drafting and preliminary analysis. Collectively, these advancements illustrate the transformative impact of LLMs in the legal sector by automating complex cognitive tasks that traditionally require substantial human effort and expertise.

Moreover, LLMs are demonstrating an increasing capability to understand, summarize, and even analyze courtroom debates and discussions, providing valuable assistance to lawyers by enhancing their ability to quickly grasp the key points of complex legal arguments [43].

Artificial Lawyer Unlike research focused on how LLM can assist lawyers in their work, there is a group of researchers investigating whether LLMs could potentially replace lawyers [47, 39, 46, 45, 35, 42]. Essentially, they have tested or attempted to use LLMs to interpret legal terms, and court proceedings, and provide legal assistance to non-experts, such as through question answering.

In the role of legal advisory services, LLMs such as GPT models are employed to interpret and apply intricate tax laws [39, 46, 41]. Although these models do not yet match the expertise of human lawyers, they show potential. The use of prompting strategies and retrieval-augmented generation techniques helps enhance the accuracy of the advice provided. While LLMs like ChatGPT are not capable of completely replacing lawyers, they are effective in providing accessible legal information to the general public [39]. Studies comparing LLMs like ChatGPT with expert-based legal systems have shown that LLMs can sometimes match or even exceed the performance of specialised legal tools in terms of user interaction and information provision (Guha et al. [36] also proposed a legal benchmark to evaluate the legal reasoning capabilities of various LLMs.). This is particularly true in non-specialised contexts where high levels of legal expertise are not critical. Furthermore, LLMs assist in legal drafting by helping draft documents that are influential in legal settings, such as jury trials in cryptocurrency securities cases, albeit with some limitations in their ability to perform complex legal reasoning [46, 48].

4. Open Challenges

In Section 3, the evaluation of LLMs against the specific requirements of legal professionals reveals their effectiveness in retrieving legal statutes and facilitating natural language communication. This capability aids in clarifying the contextual background of incidents. However, the application of LLMs in legal practice remains underdeveloped, with existing explanations and standards still insufficient. While LLMs show promise in information gathering, research addressing other critical legal requirements is lacking. We, therefore, delineate the open challenges identified through our comparative analysis of LLM capabilities and legal requirements, and we also summarise the concerns raised by researchers in the existing literature.

4.1. Reduced Quality of Explanations

As discussed in Section 3.1.2 ([18, 22, 23, 22, 24, 27]), the *accuracy* and *faithfulness* of explanations are of great importance. This is especially critical in the legal domain, where issues of right and wrong must be clearly delineated without any ambiguity. Given the importance and permanence of this requirement, we did not find any paper indicating that LLMs have advanced significantly in this area. However, current research has found that LLMs are prone to factual inaccuracies and suffer from what is often referred to as the 'hallucination' problem [52, 53, 48]. This problem leads LLMs to generate responses that appear correct but are in fact fabricated, which is particularly detrimental when providing explanations in legal contexts. Furthermore, LLMs are typically generalist models trained on vast datasets without a dedicated knowledge base [54, 55]. This means that their lacking of domain-specific knowledge significantly impacts their practical applicability in legal settings because they struggle to comprehend and process the complexities involved in intricate legal cases [56]. Moreover, the importance of a continually updated domain knowledge base is critical to ensure that the latest legal regulations and document modifications are incorporated into the model, thus preventing errors. A common problem stemming from a lack of domain knowledge is the variability in legal standards and interpretations across different jurisdictions. LLMs might generate explanations based on general legal principles that do not align perfectly with local laws or recent legal developments [57].

4.2. Verification Challenge

Verifying the accuracy of explanations generated by LLMs presents inherent challenges. LLMs often function as 'black boxes', obscuring the pathways of their reasoning and making it difficult to discern the basis for their outputs [58]. As previously discussed, LLMs are prone to the 'hallucination' problem, wherein they fabricate facts rather than indicate uncertainty or a lack of knowledge about a request that exceeds their current capabilities. This introduces significant complexity: determining when to trust an LLM's explanation becomes a critical question. This issue could lead to a scenario where continual fact-checking is necessary, thereby complicating the process and imposing a substantial burden on lawyers (see [48] for more details). Furthermore, lawyers may struggle to trust the explanations generated by LLMs as they lack sufficient understanding of how explanations were generated.

4.3. Bias and Transparency Considerations

LLMs appear to fall short of fulfilling the requirements set forth in the literature, which advocate that explanations provided to assist lawyers should be devoid of any biases that could unjustly influence legal outcomes [50, 26, 31]. This lack of neutrality can be detrimental, as biased outputs may skew perceptions and decisions in sensitive legal matters. The nature that LLMs are trained on massive amounts of text data from the internet and other sources. If the training data contains biases, stereotypes, or underrepresented of certain groups, LLMs will inevitably learn and reflect those biases in its outputs [59, 60]. For instance, Bender et al. [61] discusses how biases in language models can lead to problematic outputs, emphasising the importance of addressing these issues to ensure fairness and impartiality. Furthermore, the ethical use of LLMs in legal settings requires transparency and accountability, which are often lacking in current AI systems. Although it did not emerge in our findings, it is important to highlight the challenge of privacy and anonymization in regards to data provided to LLMs, which poses potential conflict with General Data Protection Regulations (GDPR) such as the UK's Data Protection Act [62], therefore leaving an open challenge that needs further consideration.

4.4. Integration and Collaboration

Successfully integrating LLMs into the practice of law for autonomous car accidents also requires effective collaboration between lawyers, technologists, and policymakers. Lawyers need a basic understanding of the technology to use LLM outputs effectively, while technologists must ensure that the explanations generated by LLMs are interpretable and actionable for lawyers. This cross-disciplinary collaboration is important for the effective use of LLMs in legal contexts but can be challenging to achieve. The value of interdisciplinary collaboration in legal technology projects emphasises that both legal and technical expertise are necessary for successful integration [63]. Policymakers also play a key role by establishing regulations and standards that guide the use of LLMs in legal settings. Achieving effective collaboration among these diverse groups can be difficult due to differences in language, priorities, and areas of expertise. Bridging these gaps is essential to fully utilise the potential of LLMs in providing reliable and accurate legal explanations for autonomous car accidents.

5. Future Directions

In addition to the open challenges mentioned in the previous section, we outline several potential key areas for tackling the future direction of using LLMs to generate explanations for lawyers in autonomous car accidents.

Development of Specialised Legal LLMs Advancing the capabilities of LLMs for legal applications requires the development of models specifically trained on broad legal documents, such as legislation, court decisions, and contracts. This specialised training is essential for aligning LLMs more closely with the unique demands of the legal field. By fine-tuning these models to handle complex legal tasks, such as classifying legal reasoning, their performance can be enhanced [64]. These specialised LLMs will be better equipped to understand legal terminology and apply legal reasoning to complex cases accurately. To ensure that LLMs can provide precise explanations in autonomous car accidents, it is necessary to train them with a wide variety of data sources. This includes vehicle telemetry, real-world accident data, and simulations. Additionally, integrating legal precedents and regulatory information from multiple jurisdictions will improve the models' ability to generate accurate and relevant legal explanations. Training these models on data that reflects different legal systems will further enhance their adaptability, allowing them to function effectively across various jurisdictions.

Enhancing Contextual Legal Understanding To improve LLMs for helping lawyers with autonomous car accident cases, future models need to get better at understanding complex legal texts and the specific details of legal cases. This means not just handling large amounts of information but also grasping important points like the purpose behind laws, the importance of past cases, and how to interpret different laws. Techniques like few-shot learning, where models learn from a few examples, and reinforcement learning, which improves models through trial and error, are key to this progress. Additionally, using a retrieval module to make responses more reliable is important [65]. This module helps LLMs find and use the most relevant information when explaining legal matters. These improvements will make LLMs more trustworthy helpers in creating legal arguments and giving advice on possible legal outcomes, making them more useful in the legal field.

Robust Testing and Validation Frameworks Before LLMs are widely used in legal practices, they need thorough testing to check their accuracy and reliability. Creating simulated environments that mimic real-life legal scenarios can help assess how these models perform in different situations. This testing should look at how well the models make decisions and handle ambiguous or conflicting information, which often occurs in legal cases. Scenario-based testing can include various factors such as weather conditions, traffic patterns, and driver behaviours. Legal specificity testing ensures the models generate explanations that are both technically accurate and legally relevant, referencing the correct laws, regulations, and precedents. Transparency and accountability are important to evaluate the models' reasoning, ensuring they are free from biases and errors. Ethical considerations ensure LLMs avoid biased or discriminatory language and adhere to legal standards. Continuous learning and improvement, with feedback from lawyers, can help improve the models over time. Implementing these robust testing and validation frameworks can make LLMs reliable and effective for generating explanations in autonomous car accident cases, helping lawyers in their work [66].

Real-Time Explanation Systems Future initiatives should focus on developing real-time LLM systems that can provide instant explanations and legal guidance during the investigation of autonomous car accidents. These systems, integrated with live data feeds from autonomous vehicles and traffic monitoring systems can generate immediate, contextual insights, drastically reducing preliminary investigation times and helping legal professionals make faster, informed decisions. Utilising platforms like LimSim++ [67], which offer multi-modal inputs and real-time decision-making capabilities, these systems can analyse driving scenarios and explain events leading up to an accident [68]. This approach aids in allocating responsibilities by analysing data such as vehicle commands and sensor readings to determine fault, aligning with regulatory frameworks like those from the NHTSA that emphasise transparency and accountability in autonomous driving [69].

Interdisciplinary Collaboration Platforms To further the development and effectiveness of LLMs in this field, it is necessary to create platforms where experts from different fields can work together. These platforms would bring together legal professionals, AI researchers, autonomous vehicle engineers, and ethicists to share knowledge and insights. For example, the Partnership on AI [70] and the AI4People [71] initiative help experts from different fields collaborate on solving problems related to AI. Studies have shown that working together across different fields can make AI systems better and more ethical by including various viewpoints and expertise. By establishing a platform dedicated to the use of LLMs for generating explanations in autonomous car accidents, we can enhance lawyers' understanding of the technology and provide AI researchers with critical insights to fine-tune their models. This collaboration ensures that LLMs are designed to meet the specific needs of legal practice, ultimately leading to more effective and ethically sound applications in the explanation of autonomous vehicle incidents.

6. Conclusion

We systematically reviewed 119 research papers on current LLMs in addressing legal issues related to LLMs and then selected in-depth 31 papers about LLMs assisting lawyers in generating explanations in autonomous car accidents. Our findings show that while LLMs offer promising capabilities in processing and summarising information, their current applications lack the depth and legal specificity needed for effective legal adjudications. We also outline key challenges to be addressed in future work and propose directions for mitigating these issues. These include the development of specialised legal LLMs, enhancing contextual legal understanding, robust testing and validation frameworks, real-time explanation systems, and interdisciplinary collaboration platforms.

Acknowledgments

This project was supported by the University of Manchester and King's College London. We would like to thank the Reviewers for taking the time and effort necessary to review the manuscript.

References

[1] J. Kabzan, M. de la Iglesia Valls, V. Reijgwart, H. F. C. Hendrikx, C. Ehmke, M. Prajapat, A. Bühler, N. B. Gosala, M. Gupta, R. Sivanesan, A. Dhall, E. Chisari, N. Karnchanachari, S. Brits, M. Dangel, I. Sa, R. Dubé, A. Gawel, M. Pfeiffer, A. Liniger, J. Lygeros, R. Y. Siegwart, Amz driverless: The full autonomous racing system, Journal of Field Robotics 37 (2019) 1267 – 1294. URL: https://api.semanticscholar.org/CorpusID:152282677.

[2] S. N. Dethe, V. S. Shevatkar, P. R. P. Bijwe, H. Mandal's, Google driverless car, International journal of scientific research in science, engineering and technology 2 (2016) 133–137. URL: https://api.semanticscholar.org/CorpusID:16886591.

[3] M. H. Shubbak, Self-driving cars: Legal, social, and ethical aspects, Economics of Innovation eJournal (2013). URL: https://api.semanticscholar.org/CorpusID:159857674.

[4] T. H. Pearl, Compensation at the crossroads: Autonomous vehicles and alternative victim compensation schemes, Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (2019). URL: https://api.semanticscholar.org/CorpusID:195878068.

[5] D. D. W. Praveenraj, M. M. Victor, C. Vennila, A. H. Alawadi, P. Diyora, N. Vasudevan, T. Avudaiappan, Exploring explainable artificial intelligence for transparent decision making, E3S Web of Conferences (2023). URL: https://api.semanticscholar.org/CorpusID:259933418.

[6] M. C. Buiten, L. A. Dennis, M. Schwammberger, A vision on what explanations of autonomous systems are of interest to lawyers, 2023 IEEE 31st International Requirements Engineering Conference Workshops (REW) (2023) 332–336. URL: https://api.semanticscholar.org/CorpusID:263231027.

[7] S. Burton, I. Habli, T. Lawton, J. Mcdermid, P. Morgan, Z. Porter, Mind the gaps: Assuring the safety of autonomous systems from an engineering, ethical, and legal perspective, Artif. Intell. 279 (2020). URL: https://api.semanticscholar.org/CorpusID:209932331.

[8] OpenAI, Gpt-4, 2023. URL: https://openai.com/index/gpt-4/, last accessed 26 April 2024.

[9] B. A. y Arcas, Do large language models understand us?, Daedalus 151 (2022) 183–197. URL: https://api.semanticscholar.org/CorpusID:248377874.

[10] A. J. Thirunavukarasu, D. S. J. Ting, K. Elangovan, L. Gutierrez, T. F. Tan, D. S. W. Ting, Large language models in medicine, Nature medicine 29 (2023) 1930–1940.

[11] T. Lai, Y. Shi, Z. Du, J. Wu, K. Fu, Y. Dou, Z. Wang, Supporting the demand on mental health services with ai-based conversational large language models (llms), BioMedInformatics 4 (2023) 8–33.

[12] A. B. Mbakwe, I. Lourentzou, L. A. Celi, O. J. Mechanic, A. Dagan, Chatgpt passing usmle shines a spotlight on the flaws of medical education, 2023.

[13] K. Lakkaraju, S. K. R. Vuruma, V. Pallagani, B. Muppasani, B. Srivastava, Can llms be good financial advisors?: An initial study in personal decision making for optimized outcomes, arXiv preprint arXiv:2307.07422 (2023).

[14] K. Lakkaraju, S. E. Jones, S. K. R. Vuruma, V. Pallagani, B. C. Muppasani, B. Srivastava, Llms for financial advisement: A fairness and efficacy study in personal decision making, in: 4th ACM International Conference on AI in Finance, 2023, pp. 100–107.

[15] D. Moher, A. Liberati, J. Tetzlaff, D. G. Altman, P. Group, et al., Preferred reporting items for systematic reviews and meta-analyses: the prisma statement, International journal of surgery 8 (2010) 336–341.

[16] D. M. Crowther, A clinician's guide to systematic reviews., Nutrition in clinical practice : official publication of the American Society for Parenteral and Enteral Nutrition 28 4 (2013) 459–62. URL: https://api.semanticscholar.org/CorpusID:22213389.

[17] V. Clarke, V. Braun, Thematic analysis, The journal of positive psychology 12 (2017) 297–298.

[18] Ł. Górski, S. Ramakrishna, Explainable artificial intelligence, lawyer's perspective, in: Proceedings of the eighteenth international conference on artificial intelligence and law, 2021, pp. 60–68.

[19] K. Atkinson, J. Collenette, T. Bench-Capon, K. Dzehtsiarou, Practical tools from formal models: the echr as a case study, in: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law, 2021, pp. 170–174.

[20] C. F. Luo, R. Bhambhoria, S. Dahan, X. Zhu, Evaluating explanation correctness in legal decision making., in: AI, 2022.

[21] S. Kalva, F. Geldon, Semantic nlp technologies in information retrieval systems for legal research. adv mach lear art inte, 2 (1): 28-32, 2021.

[22] M. V. Naik, S. Lokhanday, Building a legal expert system for legal reasoning in specific domain-a survey, AIRCC's International Journal of Computer Science and Information Technology (2012) 175–184.

[23] A. Khasianov, I. Alimova, A. Marchenko, G. Nurhambetova, E. Tutubalina, D. Zuev, Lawyer's intellectual tool for analysis of legal documents in russian, in: 2018 International Conference on Artificial Intelligence Applications and Innovations (IC-AIAI), IEEE, 2018, pp. 42–46.
[24] K. Stoykov, S. Chelebieva, Legal data extraction and possible applications, in: IOP Conference Series: Materials Science and Engineering, volume 618, IOP Publishing, 2019, p. 012037.
[25] M. van Otterlo, M. Atzmueller, On requirements and design criteria for explainability in legal ai, in: The 31st International Conference on Legal Knowledge and Information Systems, 2018.
[26] K. Atkinson, T. Bench-Capon, D. Bollegala, Explanation in ai and law: Past, present and future, Artificial Intelligence 289 (2020) 103387.
[27] W. Yu, Z. Sun, J. Xu, Z. Dong, X. Chen, H. Xu, J.-R. Wen, Explainable legal case matching via inverse optimal transport-based rationale extraction, in: Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval, 2022, pp. 657–668.
[28] S. Atakishiyev, M. Salameh, H. Yao, R. Goebel, Explainable artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions, arXiv preprint arXiv:2112.11561 (2021).
[29] Y. Shen, S. Jiang, Y. Chen, K. D. Campbell, To explain or not to explain: A study on the necessity of explanations for autonomous vehicles, arXiv preprint arXiv:2006.11684 (2020).
[30] A. Tabrez, M. B. Luebbers, B. Hayes, Automated failure-mode clustering and labeling for informed car-to-driver handover in autonomous vehicles, arXiv preprint arXiv:2005.04439 (2020).
[31] P. Balboni, A. Botsi, K. Francis, M. T. Barata, Designing connected and automated vehicles around legal and ethical concerns: Data protection as a corporate social responsibility., in: SETN Workshops, 2020, pp. 139–151.
[32] T. Gillespie, S. Hailes, Assignment of legal responsibilities for decisions by autonomous cars using system architectures, IEEE Transactions on Technology and Society 1 (2020) 148–160.
[33] Z. Shen, K. Lo, L. Yu, N. Dahlberg, M. Schlanger, D. Downey, Multi-lexsum: Real-world summaries of civil rights lawsuits at multiple granularities, Advances in Neural Information Processing Systems 35 (2022) 13158–13173.
[34] J. Savelka, Unlocking practical applications in legal domain: Evaluation of gpt for zero-shot semantic annotation of legal texts, in: Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law, 2023, pp. 447–451.
[35] Q. Huang, M. Tao, Z. An, C. Zhang, C. Jiang, Z. Chen, Z. Wu, Y. Feng, Lawyer llama technical report, arXiv preprint arXiv:2305.15062 (2023).
[36] N. Guha, J. Nyarko, D. Ho, C. Ré, A. Chilton, A. Chohlas-Wood, A. Peters, B. Waldon, D. Rockmore, D. Zambrano, et al., Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models, Advances in Neural Information Processing Systems 36 (2024).
[37] A. Dyevre, Text-mining for lawyers: how machine learning techniques can advance our understanding of legal discourse, Erasmus L. Rev. 14 (2021) 7.

[38] J. Savelka, K. D. Ashley, Discovering explanatory sentences in legal case decisions using pre-trained language models, arXiv preprint arXiv:2112.07165 (2021).
[39] J. Tan, H. Westermann, K. Benyekhlef, Chatgpt as an artificial lawyer, Artificial Intelligence for Access to Justice (AI4AJ 2023) (2023).
[40] J. Savelka, K. D. Ashley, The unreasonable effectiveness of large language models in zero-shot semantic annotation of legal texts, Frontiers in Artificial Intelligence 6 (2023).
[41] X. Kang, L. Qu, L.-K. Soon, A. Trakic, T. Y. Zhuo, P. C. Emerton, G. Grant, Can chatgpt perform reasoning using the irac method in analyzing legal scenarios like a lawyer?, arXiv preprint arXiv:2310.14880 (2023).
[42] C. Ryu, S. Lee, S. Pang, C. Choi, H. Choi, M. Min, J.-Y. Sohn, Retrieval-based evaluation for llms: A case study in korean legal qa, in: Proceedings of the Natural Legal Language Processing Workshop 2023, 2023, pp. 132–137.
[43] E. Ash, A. Kesari, S. Naidu, L. Song, D. Stammbach, Enhancing public understanding of court opinions with automated summarizers, arXiv preprint arXiv:2311.06534 (2023).
[44] J. J. Nay, D. Karamardian, S. B. Lawsky, W. Tao, M. Bhat, R. Jain, A. T. Lee, J. H. Choi, J. Kasai, Large language models as tax attorneys: a case study in legal capabilities emergence, Philosophical Transactions of the Royal Society A 382 (2024) 20230159.
[45] J. Savelka, K. D. Ashley, M. A. Gray, H. Westermann, H. Xu, Explaining legal concepts with augmented large language models (gpt-4), arXiv preprint arXiv:2306.09525 (2023).
[46] A. Trozze, T. Davies, B. Kleinberg, Large language models in cryptocurrency securities cases: Can chatgpt replace lawyers?, arXiv preprint arXiv:2308.06032 (2023).
[47] A. Ammar, A. Koubaa, B. Benjdira, O. Nacar, S. Sibaee, Prediction of arabic legal rulings using large language models, Electronics 13 (2024) 764.
[48] M. Dahl, V. Magesh, M. Suzgun, D. E. Ho, Large legal fictions: Profiling legal hallucinations in large language models, arXiv preprint arXiv:2401.01301 (2024).
[49] W. L. Prosser, et al., Handbook of the Law of Torts, volume 4, West Publishing, 1941.
[50] V. Dubljević, Toward implementing the adc model of moral judgment in autonomous vehicles, Science and Engineering Ethics 26 (2020) 2461–2472.
[51] S. Tian, Q. Jin, L. Yeganova, P.-T. Lai, Q. Zhu, X. Chen, Y. Yang, Q. Chen, W. Kim, D. C. Comeau, et al., Opportunities and challenges for chatgpt and large language models in biomedicine and health, Briefings in Bioinformatics 25 (2024) bbad493.
[52] Z. Ji, T. Yu, Y. Xu, N. Lee, E. Ishii, P. Fung, Towards mitigating hallucination in large language models via self-reflection, arXiv preprint arXiv:2310.06271 (2023).
[53] X. Zhan, Y. Xu, S. Sarkadi, Deceptive ai ecosystems: The case of chatgpt, in: Proceedings of the 5th International Conference on Conversational User Interfaces, 2023, pp. 1–6.
[54] H. Naveed, A. U. Khan, S. Qiu, M. Saqib, S. Anwar, M. Usman, N. Barnes, A. Mian, A comprehensive overview of large language models, arXiv preprint arXiv:2307.06435 (2023).
[55] X. Deng, Y. Gu, B. Zheng, S. Chen, S. Stevens, B. Wang, H. Sun, Y. Su, Mind2web: Towards a generalist agent for the web, Advances in Neural Information Processing Systems 36 (2024).
[56] H. Zhong, C. Xiao, C. Tu, T. Zhang, Z. Liu, M. Sun, How does nlp benefit legal system: A summary of legal artificial intelligence, arXiv preprint arXiv:2004.12158 (2020).
[57] M. Chen, A guide to generative ai and llms for in-house counsel, 2024. URL: https://www.streamline.ai/blog/guide-to-generative-ai-and-llms-for-lawyers, last accessed 26 April

2024.
[58] S. Huo, N. Arabzadeh, C. L. Clarke, Retrieving supporting evidence for llms generated answers, arXiv preprint arXiv:2306.13781 (2023).
[59] K. Oroy, A. Nick, Fairness and Bias Detection in Large Language Models: Assessing and Mitigating Unwanted Biases, Technical Report, EasyChair, 2024.
[60] U. P. Liyanage, N. D. Ranaweera, Ethical considerations and potential risks in the deployment of large language models in diverse societal contexts, Journal of Computational Social Dynamics 8 (2023) 15–25.
[61] E. M. Bender, T. Gebru, A. McMillan-Major, S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, in: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, 2021, pp. 610–623.
[62] UK Data Protection Act 2018, 2018. URL: https://www.legislation.gov.uk/ukpga/2018/12/contents/enacted, last accessed 3 September 2024.
[63] D. M. Katz, M. J. Bommarito II, J. Blackman, Predicting the behavior of the supreme court of the united states: A general approach, arXiv preprint arXiv:1407.6333 (2014).
[64] R. E. Thalken, E. H. Stiglitz, D. M. Mimno, M. Wilkens, Modeling legal reasoning: Lm annotation at the edge of human agreement, in: Conference on Empirical Methods in Natural Language Processing, 2023. URL: https://api.semanticscholar.org/CorpusID:264590650.
[65] Q. Huang, M. Tao, Z. An, C. Zhang, C. Jiang, Z. Chen, Z. Wu, Y. Feng, Lawyer llama technical report, ArXiv abs/2305.15062 (2023). URL: https://api.semanticscholar.org/CorpusID:264148648.
[66] F. Doshi-Velez, B. Kim, Towards a rigorous science of interpretable machine learning, arXiv preprint arXiv:1702.08608 (2017).
[67] D. Fu, W. Lei, L. Wen, P. Cai, S. Mao, M. Dou, Y. Qiao, Limsim++: A closed-loop platform for deploying multimodal llms in autonomous driving, arXiv preprint arXiv:2402.01246 (2024).
[68] Z. Yang, X. Jia, H. Li, J. Yan, Llm4drive: A survey of large language models for autonomous driving, arXiv e-prints (2023) arXiv–2311.
[69] M. L. Kubica, Autonomous vehicles and liability law, The American Journal of Comparative Law 70 (2022) i39–i69.
[70] J. Heer, The partnership on ai, AI Matters 4 (2018) 25–26.
[71] L. Floridi, J. Cowls, M. Beltrametti, R. Chatila, P. Chazerand, V. Dignum, C. Luetge, R. Madelin, U. Pagallo, F. Rossi, et al., Ai4people—an ethical framework for a good ai society: opportunities, risks, principles, and recommendations, Minds and machines 28 (2018) 689–707.