THE LEGAL ASSEMBLY LINE: A CRITIQUE OF AI IN INDIAN LAW

Olivia Ruhil, School of Public Policy, Indian Institute of Technology, Delhi

The law, in its essence, is a reflection of human values, societal norms, and ethical dilemmas. It is a living, breathing entity that evolves over time, shaped by the wisdom, discretion, and compassion of those entrusted with its interpretation. Yet, in an age of rapid technological advancement, artificial intelligence (AI) promises to revolutionise the legal landscape. The allure of AI lies in its ability to process vast amounts of data, streamline repetitive tasks, and deliver efficiency that humans alone may struggle to match. But here lies a fundamental question: Can a machine, no matter how sophisticated, ever truly understand the intricacies of justice?

ABSTRACT

Artificial intelligence (AI) is transforming industries across the globe, and the legal profession is no exception. The promise of AI lies in its unmatched efficiency, handling vast amounts of data and performing repetitive tasks like contract review, legal research, and compliance checks with unprecedented speed. However, this technological leap raises critical questions: Can AI, devoid of human values, truly administer justice? This article introduces the concept of the "Legal Assembly Line" to critique AI's role in law, particularly in the Indian legal system, where societal complexities, moral questions, and constitutional interpretations are at the forefront.

While AI excels in structured, procedural tasks, it falters when confronted with the nuanced, interpretative reasoning required in cases involving fundamental rights, ethical dilemmas, and societal values. The doctrinal research presented here delves into the limitations of AI in handling such tasks, drawing on philosophical arguments like Searle's Chinese Room and real-world examples such as landmark Indian judgments. This article argues that while AI can assist lawyers, judges, and legal practitioners, it can never replace the human judgment required in the interpretative domains of law.

The implications of relying too heavily on AI in legal contexts could risk undermining the very fabric of justice, especially in a diverse and complex society like India. This article aims to spark a meaningful conversation on

the role of AI in law, advocating for a balanced, collaborative approach where AI enhances, but does not replace, human judgment.

Keywords: Artificial Intelligence in Law, Legal Assembly Line, Indian Legal System, AI in Courts, Human Judgment, Legal Interpretation, Chinese Room Argument, Turing Test, AI Ethics, Socio-Legal Complexity, Fundamental Rights, Justice

Introduction

Artificial intelligence (AI) has increasingly permeated various sectors, and law is no exception. With the rise of Large Language Models (LLMs) such as GPT-4 and beyond, legal practitioners and academics alike have debated whether these AI systems can effectively replace human judgment in the legal process. Proponents argue that AI offers unparalleled efficiency in structured legal tasks such as contract drafting, legal research, and precedent retrieval. For instance, LLMs like GPT-4 have been effectively employed in fields like tax law and document automation, as demonstrated in studies such as **Achiam, Jin, & Zhang (2023)**, where the **LawLLM** was developed for the U.S. legal system, significantly reducing time spent on drafting and research-intensive tasks.

However, alongside these advances, there exists a more profound question: Can AI truly replace the nuanced human judgment required for interpretative legal tasks, such as constitutional adjudication or judicial reasoning in complex cases involving moral and ethical dilemmas? This article introduces the concept of the "Legal Assembly Line" to critique the over-reliance on AI for legal tasks, focusing on the limitations of AI in handling interpretative reasoning. While LLMs have demonstrated impressive capabilities in structured domains, their inability to engage with context, morality, and socio-political dimensions renders them unsuitable for critical interpretative legal functions, particularly in the Indian legal context.

The doctrinal research approach used in this article analyses existing legal doctrines, case law, and scholarly opinions to build a comprehensive critique of AI's role in law. By examining real-world applications of AI, philosophical arguments like John Searle's **Chinese Room** and Alan Turing's **Turing Test**, and Indian legal case studies, this article provides a nuanced understanding of why law cannot be reduced to mechanistic processes akin to an assembly line.

The Rise of AI in Law: Structured Legal Tasks

In recent years, AI has gained significant traction in law due to its ability to handle large-scale data processing tasks. LLMs like **GPT-4**, developed by OpenAI, have revolutionised legal research and document drafting, allowing law firms to streamline processes traditionally requiring significant human labor. Studies such as **Katz, Bommarito, & Blackman (2024)** demonstrate how these models can generate legal judgments, analyse legal precedent, and even assist in predictive analytics for litigation strategies. Katz et al. highlight that in fields like tax law, where tasks are predominantly rule-based and procedural, AI performs exceptionally well by processing structured data, automating repetitive tasks, and reducing human error.

For instance, contract review platforms like **LawGeex** employ AI to review contracts, identify clauses, and flag potential legal risks. This level of automation has drastically reduced the time spent on contract analysis, offering clients faster and more accurate services. A similar argument is made by **Zhang, Sun, & Huang (2023)**, who surveyed the use of LLMs across the legal sector and concluded that AI could transform compliance monitoring and document automation in highly structured domains like corporate law. In this context, AI is undoubtedly a game-changer, enhancing the efficiency and precision of legal services.

AI's Limitations in Interpretative Legal Tasks

While AI excels in structured tasks, it falls short in interpretative legal reasoning, which requires understanding of context, precedent, and evolving legal doctrines. This limitation is particularly evident in the Indian legal system, where courts often engage in balancing fundamental rights, socio-political considerations, and constitutional principles. In cases such as **Shreya Singhal v. Union of India (2015)**, where the Supreme Court of India struck down Section 66A of the IT Act, the judicial process involved deep interpretative reasoning that extended beyond merely applying rules. The decision reflected not only a legal analysis of freedom of speech but also an understanding of the broader socio-political environment in India, something that AI, with its reliance on pattern recognition, cannot replicate.

Searle's **Chinese Room** argument further illuminates this point. In his famous thought experiment, Searle argues that even if a machine appears to understand Chinese by providing correct responses to Chinese symbols, it does not actually comprehend the language; it simply follows pre-programmed instructions. Similarly, LLMs like GPT-4 generate text based on

statistical patterns in vast datasets without truly understanding the legal concepts or moral values embedded in the text. As highlighted by **Surden (2019)**, AI systems, while capable of processing legal language, lack the cognitive ability to engage in the kind of deep interpretative reasoning required in cases involving constitutional adjudication, human rights, and ethical dilemmas.

This limitation becomes particularly critical in the Indian legal context, where socio-political issues such as caste, religion, and gender intersect with the law in complex ways. For example, in **Puttaswamy v. Union of India (2017)**, the Supreme Court expanded the right to privacy under **Article 21** of the Indian Constitution, considering not just legal precedents but also evolving global human rights standards, technological advancements, and societal values. AI's inability to understand or weigh these factors makes it unsuitable for such cases, where judicial discretion and moral reasoning are paramount.

The Legal Assembly Line: A Critique

The "Legal Assembly Line" metaphor critiques the idea that legal processes can be reduced to a series of mechanistic steps performed by AI. While AI can effectively handle the repetitive and structured components of legal work, law is not merely a procedural system—it is interpretative, value-driven, and context-dependent. This assembly line approach risks undermining the very foundations of justice, particularly in legal systems like India's, where courts are often called upon to balance individual rights with societal interests.

As noted by **Bender et al. (2021)** in their critique of **stochastic parrots**, AI systems merely "parrot" back patterns of language found in their training data without any understanding of meaning or context. This "static parrot" behavior is particularly dangerous in the legal field, where context and evolving social norms play a crucial role. For example, in **Navtej Singh Johar v. Union of India (2018)**, where the Supreme Court decriminalized homosexuality, the judgment was not just a legal analysis but a reflection of changing societal values, global human rights trends, and the evolving understanding of individual autonomy. AI's inability to engage with these dimensions makes it unsuitable for such interpretative tasks.

The Dichotomy of Legal Tasks—Structured vs. Interpretative

The application of AI in law has exposed a critical dichotomy between structured and

interpretative tasks, with AI excelling in the former while struggling in the latter. Structured tasks are procedural and rule-based, requiring minimal discretion and often following a rigid framework. These tasks include document review, contract analysis, case law retrieval, and legal research—all of which AI performs with remarkable efficiency. However, interpretative tasks involve deep legal reasoning, contextual understanding, and the balancing of competing societal values—areas where AI falters.

Structured Legal Tasks: A Domain for AI Efficiency

AI has demonstrated clear advantages in structured legal domains. Tasks such as contract review, due diligence, and compliance monitoring are time-consuming and labor-intensive when handled by humans. However, with AI-driven platforms like **LawGeex**, firms can streamline these tasks, allowing for faster and more accurate processing of large volumes of legal documents. In **Achiam, Jin, & Zhang (2023)**, the authors detail how **LawLLM**—a large language model designed specifically for the U.S. legal system—can draft legal documents, analyse precedents, and automate many of the routine tasks involved in corporate and tax law.

AI's success in structured tasks stems from its ability to recognise patterns in vast datasets. By learning from thousands of legal contracts or court cases, AI systems can identify clauses, flag risks, and suggest amendments with a level of speed and precision that far exceeds human capabilities. The **LawLLM** model demonstrated a significant reduction in time spent on drafting and research-intensive tasks, thus freeing up human lawyers to focus on more complex, strategic work .

Case Study: AI in Contract Review

A well-known example of AI's utility in structured legal tasks is its application in contract review platforms. Companies like **LawGeex** employ AI to analyse contracts for potential risks and discrepancies. In a controlled study, **LawGeex** was able to outperform human lawyers in identifying risks in Non-Disclosure Agreements (NDAs), with a 94% accuracy rate compared to the 85% achieved by human experts. This showcases the value AI brings to structured legal environments where clear rules and precedent-based decisions prevail. As **Katz, Bommarito, & Blackman (2024)** argue, these systems are adept at handling legal domains where decisions are primarily based on well-established rules and procedures.

Interpretative Legal Tasks: AI's Limitations

While AI thrives in structured legal tasks, it struggles with interpretative tasks, which require deep contextual understanding, moral reasoning, and the application of legal principles to novel situations. Interpretative tasks often involve balancing competing rights, interpreting the law in light of evolving societal norms, and making value-laden judgments. These tasks are inherently subjective, demanding not just technical expertise but also ethical and philosophical deliberation.

A key example of this limitation is seen in **Shreya Singhal v. Union of India (2015)**, where the Supreme Court of India struck down Section 66A of the Information Technology Act. The Court's decision was not merely a mechanical application of precedent but a nuanced interpretation of **Article 19(1)(a)** of the Indian Constitution, which guarantees the right to freedom of speech. The Court had to weigh the implications of this right in the context of modern digital communication, balancing it against concerns of public order and state security. Such a balancing act goes beyond what AI, which lacks an understanding of societal values and moral implications, can achieve.

Searle's Chinese Room Argument and AI's Lack of Comprehension

Philosopher John Searle's **Chinese Room Argument** is often cited to explain AI's fundamental limitations in understanding. In this thought experiment, Searle imagines himself in a room, following instructions to manipulate Chinese symbols without understanding the language. From the outside, it appears that Searle understands Chinese because he produces appropriate responses to Chinese questions, but in reality, he is simply following syntactical rules without any comprehension of meaning. Similarly, AI systems like GPT-4 generate legal texts and responses based on patterns in data but lack any real understanding of the law or the values underpinning legal decisions.

This is particularly evident in interpretative legal tasks where moral and ethical considerations are paramount. As **Surden (2019)** notes in his critique of AI's role in legal reasoning, AI lacks the cognitive abilities required for deep interpretative tasks, such as those involving human rights or constitutional interpretation . While AI can process language and produce legal texts that appear coherent and well-reasoned, it does so without any understanding of the broader social, political, or ethical context in which the law operates.

Indian Legal Example: Puttaswamy v. Union of India (2017)

The limitations of AI in interpretative legal tasks are further highlighted in the landmark Indian case of **Puttaswamy v. Union of India (2017)**. In this case, the Supreme Court of India recognized the right to privacy as a fundamental right under **Article 21** of the Constitution. The decision was not merely a technical application of legal rules but a profound interpretation of the Constitution in light of changing societal values, technological advancements, and global human rights standards. The Court had to grapple with complex questions about the role of the state in the digital age, the impact of surveillance technologies on individual freedoms, and the evolving nature of privacy in a connected world.

AI, with its reliance on data and pattern recognition, cannot engage with such deep interpretative reasoning. As **Katz, Bommarito, & Blackman (2024)** argue, AI systems are fundamentally limited in their ability to handle cases involving the balancing of fundamental rights and societal interests . In cases like **Puttaswamy**, where the law is evolving in response to new social and technological realities, human judgment is essential.

Legal Realism vs. Legal Formalism: The Role of Human Judgment

The debate between **legal realism** and **legal formalism** provides further insight into AI's limitations in interpretative legal tasks. Legal formalism views the law as a closed system of rules that can be applied mechanistically, much like a mathematical equation. In contrast, legal realism emphasises the importance of context, societal values, and human experience in legal decision-making. As **Leiter (2007)** argues, formalism fails to account for the dynamic nature of the law and its responsiveness to societal changes.

AI systems, with their reliance on rule-based decision-making, align more closely with formalism. They can apply rules and procedures effectively but cannot account for the contextual and value-laden aspects of legal interpretation that legal realists emphasise. This limitation is particularly problematic in interpretative legal tasks, where the law must adapt to new social, political, and technological realities.

Searle's Chinese Room Argument: Mimicking Without Understanding

In his 1980 thought experiment, John Searle asked readers to imagine a person locked in a

room, following detailed instructions to manipulate Chinese symbols in response to questions written in Chinese. To an outside observer, it would appear that the person understands Chinese, as they are producing coherent responses. However, the individual inside the room has no understanding of the language—they are merely following a set of syntactical rules. Searle's argument is that while machines (or AI) can manipulate symbols in a way that mimics human behaviour, they do not "understand" the language in the way that a human does.

This distinction is crucial when considering the role of AI in law. While AI systems like GPT-4 can generate legal arguments, analyse case law, and produce coherent judgments, they do so without understanding the underlying principles of justice, fairness, and morality that guide human legal reasoning. The AI is simply processing inputs (legal texts) and producing outputs (responses or judgments) based on patterns it has learned from vast datasets. As Surden (2019) emphasises, this lack of understanding limits AI's effectiveness in tasks that require deep interpretative reasoning, such as constitutional law or human rights adjudication

AI in Legal Interpretation: Surface-Level Processing

The Chinese Room Argument is particularly relevant in the context of legal interpretation, where the meaning of a legal text is often contested, and judgments require an understanding of societal values, historical context, and evolving norms. AI systems, like the individual in the Chinese Room, can produce text that mimics legal reasoning, but they do so without truly understanding the social and ethical dimensions of the law. This is especially problematic in cases involving fundamental rights, where the interpretation of the law must account for competing values and the lived experiences of individuals.

In the Indian legal context, this limitation is starkly evident in cases like Puttaswamy v. Union of India (2017). In this case, the Supreme Court recognised the right to privacy as a fundamental right under Article 21 of the Constitution. The judgment involved not just a technical analysis of legal precedent but a deep engagement with global human rights norms, the impact of digital technologies on privacy, and the evolving nature of individual freedoms in a connected world. AI, lacking any real understanding of these complex, interwoven factors, would be incapable of producing such a nuanced interpretation of the law.

Katz, Bommarito, & Blackman (2024) argue that while AI systems like GPT-4 can generate legal texts that appear coherent and well-reasoned, they are fundamentally limited in their

ability to handle cases that require interpretative reasoning . AI's reliance on pattern recognition and data processing means that it can only mimic the form of legal reasoning, without engaging in the kind of substantive moral and ethical deliberation that human judges bring to their decisions.

Case Study: Puttaswamy v. Union of India (2017)

The Puttaswamy case serves as a powerful example of AI's limitations in interpretative legal tasks. The case was a landmark decision by the Supreme Court of India, recognizing the right to privacy as a fundamental right. The judgment required the Court to balance individual rights with the state's interest in surveillance and national security. The judges also had to consider the implications of emerging digital technologies, such as Aadhaar (India's biometric identification system), and the impact of these technologies on individual freedoms.

AI systems, trained on legal texts and precedents, might be able to produce a judgment that references relevant case law and statutory provisions. However, they would lack the ability to weigh the broader implications of the decision, such as the impact on personal autonomy, human dignity, and the protection of vulnerable populations from state overreach. As Surden (2019) points out, AI's inability to engage with these deeper aspects of legal reasoning makes it unsuitable for tasks that require the balancing of competing rights and interests.

In Puttaswamy, the Court's interpretation of privacy was not just a mechanical application of legal rules but a reflection of evolving societal values and global human rights trends. The judgment drew on international jurisprudence, including cases from the United States, the European Union, and the United Kingdom, to craft a forward-looking understanding of privacy that accounted for the realities of the digital age. AI, which relies on data from past legal decisions, cannot engage in this kind of future-oriented legal reasoning.

The Role of Human Judgment in Legal Interpretation

In the legal domain, particularly in complex cases involving human rights and constitutional interpretation, decisions are not based solely on the mechanical application of rules. They require judges to weigh competing interests, consider broader societal implications, and often make value-laden choices that reflect the moral fabric of a society. In cases like Navtej Singh Johar v. Union of India (2018), the Indian Supreme Court had to engage in a deep interpretative

process that accounted for evolving norms of equality, autonomy, and dignity. The decision to strike down Section 377, decriminalising homosexuality, was not just a legal conclusion but a moral stance affirming the rights of LGBTQ+ individuals in India.

In this context, Bender et al. (2021) highlight a critical flaw in the use of AI systems for legal interpretation: AI, as a "stochastic parrot," may generate language that seems appropriate based on the patterns it has learned from previous data, but it cannot grasp the underlying moral and ethical dimensions of the law . The process of interpretation, especially in cases that challenge deeply entrenched societal norms, requires not just pattern recognition but a reasoned deliberation that AI, by its very nature, cannot perform.

AI's Inability to Weigh Competing Interests

The Chinese Room Argument underscores another significant limitation of AI in law: its inability to weigh competing interests in a meaningful way. Legal interpretation, especially in constitutional matters, often requires judges to balance conflicting rights and interests. For instance, in privacy cases, courts must consider the individual's right to privacy against the state's interest in security or public order. This balancing act is not purely a matter of applying rules or precedents; it involves a careful weighing of the consequences of different legal interpretations on various stakeholders.

In the Puttaswamy case, the Supreme Court had to navigate the tension between individual privacy rights and the government's surveillance powers. The judgment reflected a careful consideration of the risks posed by unchecked state surveillance, particularly in a digital age where personal data is increasingly vulnerable. AI systems, which rely on statistical models and past legal data, lack the capability to perform this type of nuanced balancing. As Surden (2019) argues, AI may be able to generate text that resembles legal reasoning, but it cannot engage in the deliberative process required to make judgments about conflicting rights and interests .

Case Study: Shayara Bano v. Union of India (2017)

Another illustrative case in the Indian context is Shayara Bano v. Union of India (2017), where the Supreme Court of India struck down the practice of triple talaq (instant divorce) as unconstitutional. The judgment involved balancing religious freedom with gender equality, a

complex issue that required the Court to interpret constitutional provisions in light of India's pluralistic society. The judges had to consider the implications of their decision on both the Muslim community's religious practices and the rights of Muslim women to equality under the law.

An AI system tasked with analysing this case might recognise the legal provisions related to freedom of religion and gender equality, but it would lack the ability to engage in the ethical and cultural reasoning that underpinned the Court's decision. As Kemper (2020) notes, the use of AI in legal decision-making faces significant challenges when the law intersects with deeply held cultural and moral values . The judgment in Shayara Bano required not just a mechanical application of legal principles but a thoughtful consideration of the societal impact of the Court's decision, something AI cannot replicate.

Legal Interpretation: Beyond Rules and Patterns

The limitations of AI in legal interpretation are further highlighted by the distinction between legal formalism and legal realism. Legal formalism posits that legal decisions can be derived logically from established rules, much like mathematical equations. In contrast, legal realism emphasises the importance of context, human experience, and societal values in the interpretation of the law. AI systems, which operate by recognising patterns in data, align more closely with formalism. They are adept at applying rules and procedures but fail when it comes to interpreting the law in a way that accounts for evolving societal norms and values.

As Leiter (2007) argues, legal realism recognises that judges do not simply apply the law mechanically; they interpret it in light of changing social, political, and economic realities . This interpretative process is dynamic and requires a deep understanding of the broader context in which the law operates. AI systems, which rely on past data, cannot engage in this type of forward-looking legal reasoning. This is particularly problematic in the Indian legal system, where courts are often called upon to address issues of social justice, equality, and human rights—areas where legal interpretation goes beyond the application of established rules.

In cases like Navtej Singh Johar and Shayara Bano, the Indian judiciary demonstrated the importance of moral and ethical reasoning in legal interpretation. The courts did not merely apply precedents but interpreted the law in a way that advanced the rights of marginalised groups and aligned with global human rights standards. AI, which lacks the ability to engage

in moral reasoning or understand the broader social and political implications of its decisions, is fundamentally limited in its capacity to contribute to these interpretative legal tasks.

The Chinese Room Argument provides a powerful framework for understanding the limitations of AI in legal interpretation. While AI systems like GPT-4 can generate legal texts that mimic human reasoning, they do so without understanding the meaning behind the words. This lack of understanding is particularly problematic in interpretative legal tasks, where judgments require not just the application of rules but a deep engagement with societal values, ethical considerations, and competing interests.

In the Indian legal context, where courts are often called upon to address complex issues involving fundamental rights, social justice, and evolving norms, AI's limitations become even more apparent. As cases like Puttaswamy, Navtej Singh Johar, and Shayara Bano illustrate, legal interpretation is not a mechanical process but a dynamic one that requires human judgment. AI, for all its capabilities in structured legal tasks, cannot replace the moral and ethical reasoning that is central to legal interpretation.

The Turing Test—Why Imitation Is Not Enough for Law

The **Turing Test**, proposed by Alan Turing in 1950, set a benchmark for machine intelligence: if a machine could engage in conversation in such a way that a human could not distinguish it from another human, it could be said to possess intelligence. In recent years, large language models (LLMs) like GPT-4 have come impressively close to passing this test. However, passing the Turing Test, or being able to generate text indistinguishable from human language, is insufficient for handling complex, interpretative legal tasks. Law requires not only the ability to generate legally accurate language but also an understanding of the broader societal, moral, and political contexts in which laws are applied.

The **Legal Assembly Line** metaphor critiques the notion that generating legal texts is the same as understanding the law. While AI can produce grammatically correct and legally coherent texts that may seem indistinguishable from human output, this does not mean AI understands the law. It simply reproduces patterns found in the data on which it was trained. As **Kemper (2020)** notes, the real test for AI in law is not just in mimicking human legal language but in engaging with the deeper reasoning processes that underlie legal interpretation.

Imitation Without Deep Understanding

The Turing Test's focus on imitation highlights a critical limitation of AI in legal contexts: while AI may appear to engage in legal reasoning, it does so without understanding the social and moral dimensions that shape judicial decisions. AI systems like GPT-4 are trained on vast datasets of legal texts, case law, and statutory provisions, allowing them to generate legal arguments that may appear coherent and well-reasoned. However, as **Bender et al. (2021)** argue, AI's language generation capabilities are fundamentally based on pattern recognition, not comprehension.

In practice, this means that AI can generate a legal opinion or a judgment that seems accurate, but it cannot understand the broader implications of its decision. For instance, in cases involving human rights or constitutional law, judges must often balance competing values, such as individual liberty and public safety. These decisions require not just a mechanical application of the law but a thoughtful consideration of how the law interacts with society, politics, and morality.

Case Study: Navtej Singh Johar v. Union of India (2018)

In the landmark case of **Navtej Singh Johar v. Union of India (2018)**, the Indian Supreme Court decriminalised homosexuality by striking down Section 377 of the Indian Penal Code. The judgment was not simply a technical exercise in legal interpretation but a reflection of evolving societal norms, global human rights movements, and the judiciary's role in protecting individual freedoms. The decision required the Court to engage in moral reasoning, considering the implications of criminalising homosexuality on individual dignity, equality, and autonomy. The judges weighed these considerations against societal views and the role of the state in regulating private relationships.

AI systems, even those as advanced as GPT-4, would be ill-equipped to handle such a case. While they could produce text that references relevant constitutional provisions and human rights principles, they lack the ability to engage in the kind of moral and ethical deliberation that is central to cases like **Navtej Singh Johar**. As **Reiling (2020)** points out, the challenge for AI in legal reasoning is not just in generating legally accurate text but in understanding the broader societal impact of legal decisions.

Why Passing the Turing Test Is Not Enough.

The **Turing Test** evaluates a machine's ability to mimic human conversation, but it does not account for the deeper cognitive processes involved in understanding the meaning and implications of language. In the legal context, this distinction is crucial. Judges do more than simply apply the law—they interpret it in light of societal values, human rights, and ethical principles. This interpretative process requires not just linguistic competence but also a deep understanding of the human condition, the social contract, and the moral foundations of the law.

AI, by contrast, operates based on probabilistic models of language. It predicts the next word or sentence based on patterns it has learned from vast amounts of data. While this allows AI to generate text that appears human-like, it does not equip AI with the ability to engage in the kind of reasoning required for legal interpretation. As **Bender et al. (2021)** argue, AI systems are "stochastic parrots"—they generate text that mimics human language without any understanding of its meaning. This imitation may be sufficient for passing the Turing Test, but it is not enough for handling the complexities of law.

Balancing Competing Rights: A Key Challenge for AI

One of the most significant challenges in legal reasoning is the balancing of competing rights and interests. Courts are often required to weigh the rights of individuals against the interests of the state, or to balance conflicting rights, such as freedom of speech and the right to privacy. These decisions are inherently value-laden and require judges to make difficult choices about how best to protect individual freedoms while maintaining public order.

For example, in **Shreya Singhal v. Union of India (2015)**, the Indian Supreme Court struck down Section 66A of the Information Technology Act, ruling that it violated the right to freedom of speech under **Article 19(1)(a)** of the Constitution. The Court had to balance the right to free expression with concerns about public order and national security. This required a nuanced understanding of both the legal framework and the broader societal context in which the law operates. AI, which generates text based on patterns in data, cannot engage in this kind of balancing act. It lacks the ability to understand the broader implications of its decisions or to weigh competing values in a meaningful way.

As **Katz, Bommarito, & Blackman (2024)** argue, AI systems are fundamentally limited in their ability to handle cases that involve the balancing of fundamental rights and societal interests . While AI may be able to produce text that mimics legal reasoning, it cannot engage in the kind of moral and ethical deliberation that is essential for interpreting constitutional principles or protecting human rights.

Case Study: K.S. Puttaswamy v. Union of India (2017)

In the **Puttaswamy** case, the Indian Supreme Court recognised the right to privacy as a fundamental right under **Article 21** of the Constitution. The judgment required the Court to balance individual privacy rights against the state's interest in surveillance and national security. The Court also had to consider the implications of emerging digital technologies, such as Aadhaar (India's biometric identification system), on personal autonomy and data protection.

AI systems, trained on legal texts and precedents, might be able to produce a judgment that references relevant case law and statutory provisions. However, they would lack the ability to weigh the broader implications of the decision, such as the impact on personal autonomy, human dignity, and the protection of vulnerable populations from state overreach. As **Surden (2019)** notes, AI's inability to engage with these deeper aspects of legal reasoning makes it unsuitable for tasks that require the balancing of competing rights and interests.

Conclusion: AI and the Limits of Imitation in Law

While passing the Turing Test may demonstrate a machine's ability to generate human-like text, it does not equate to the ability to perform legal reasoning. In law, the challenge is not just in producing text that resembles human reasoning but in understanding the broader societal, moral, and political contexts in which laws are applied. AI systems, with their reliance on pattern recognition and data processing, can imitate the form of legal reasoning but lack the substance required for interpreting the law in a meaningful way.

The **Legal Assembly Line** metaphor highlights the dangers of reducing legal reasoning to a series of mechanistic tasks performed by AI. While AI can assist with structured legal tasks, such as document review and contract analysis, it cannot replace human judgment in interpretative legal tasks. As cases like **Navtej Singh Johar**, **Shreya Singhal**, and

Puttaswamy demonstrate, legal reasoning requires a deep engagement with societal values, human rights, and ethical principles—areas where AI is fundamentally limited.

The Socio-Legal Complexity of Indian Jurisprudence

The Indian legal system presents a unique set of challenges for the application of AI, particularly in the realm of interpretative legal reasoning. The country's pluralistic society, with its complex intersections of religion, caste, gender, and socio-economic status, makes legal interpretation far more than just the application of static rules. Judges must often engage in balancing fundamental rights with societal norms, historical context, and the evolving needs of a diverse population. AI, which relies on pattern recognition and statistical probabilities, struggles to navigate these complexities.

Legal Interpretation in a Pluralistic Society

Indian jurisprudence is marked by a delicate balance between individual rights and collective interests, and between the state's duty to maintain order and its obligation to protect freedoms. This balancing act is evident in landmark cases like **Kesavananda Bharati v. State of Kerala (1973)**, where the Supreme Court of India articulated the "basic structure" doctrine, establishing that certain features of the Constitution are inviolable. This doctrine has been used by the judiciary to protect constitutional principles from legislative overreach.

The **Kesavananda** judgment is an excellent example of the kind of interpretative task that AI systems, relying on historical data and established patterns, would find impossible to navigate. The Court's decision was not merely a mechanistic application of legal rules; it required the judges to interpret the Constitution's provisions in light of India's political history, democratic principles, and the need to safeguard individual freedoms. This dynamic interpretation of the law, which considers evolving societal values and the broader socio-political context, is something AI cannot replicate. As **Leiter (2007)** notes, legal interpretation often requires judges to go beyond the text and engage in moral and ethical reasoning .

The Role of Context in Indian Legal Decisions

Context plays a crucial role in Indian legal decisions. In cases involving caste discrimination, gender inequality, or religious freedom, the courts must interpret the law in light of the

country's complex social fabric. AI systems, which operate on data from past legal decisions, cannot fully appreciate the nuanced social and cultural factors that influence these judgments. As **Brooks, Gherhes, and Vorley (2020)** points out, AI's reliance on data can perpetuate biases that are entrenched in the legal system, particularly when dealing with marginalised communities.

A critical case in this context is **Indra Sawhney v. Union of India (1992)**, which dealt with the issue of affirmative action and caste-based reservations in government jobs. The Court had to balance the principle of equality with the need to uplift historically disadvantaged groups. The decision involved interpreting the Constitution in a way that acknowledged India's socio-economic realities while adhering to the fundamental rights enshrined in the document. AI, which relies on patterns from past decisions, would struggle to handle such a case, where the law must be interpreted in light of evolving social and political circumstances.

As **Zhang, Sun, & Huang (2023)** argue, while AI systems can be trained to recognise legal patterns, they are limited by the biases inherent in the data they are trained on . In India, where the legal system often grapples with issues of social justice, these limitations are particularly problematic. AI may inadvertently perpetuate systemic biases, as it lacks the ability to recognise and account for the socio-legal context in which these biases operate.

Religious Freedom and AI's Limitations

Religious freedom is another area where AI's limitations become apparent. In cases involving religious practices, courts must balance the right to religious freedom with other constitutional rights, such as equality and non-discrimination. For example, in **Shayara Bano v. Union of India (2017)**, the Supreme Court of India struck down the practice of triple talaq (instant divorce) as unconstitutional. The judgment involved interpreting Islamic law in a way that respected religious practices while upholding the constitutional rights of Muslim women to gender equality.

AI systems, trained on legal texts and past judgments, might be able to produce a judgment that applies the relevant legal principles. However, they would lack the ability to engage with the cultural and religious sensitivities that shaped the Court's decision. As **Kemper (2020)** points out, AI systems are fundamentally limited in their ability to engage with the ethical and cultural dimensions of legal interpretation . In cases like **Shayara Bano**, where the law intersects with

deeply held religious beliefs, human judgment is essential for interpreting the law in a way that balances constitutional rights with cultural and religious practices.

Case Study: The Ayodhya Judgment (2019)

The **Ayodhya judgment** is perhaps one of the most complex and politically charged cases in Indian legal history. The Supreme Court's decision in 2019 to allocate the disputed land in Ayodhya for the construction of a Hindu temple, while providing an alternative site for the construction of a mosque, required the Court to navigate a highly sensitive religious conflict. The judgment had to balance the historical and religious significance of the site for both Hindus and Muslims, while maintaining constitutional principles of secularism and the rule of law.

AI, with its reliance on data and pattern recognition, would be incapable of handling the complexities of such a case. The decision required not just a legal analysis of property rights but a deep understanding of the religious and historical significance of the site, as well as the political implications of the judgment for India's secular fabric. As **Surden (2019)** notes, AI systems are limited in their ability to engage with the broader societal and political context in which legal decisions are made .

The **Ayodhya** case also highlights the importance of judicial discretion in maintaining social order and preventing communal violence. The Court's decision was as much about maintaining peace and harmony in a deeply divided society as it was about resolving a legal dispute. AI, which lacks the ability to engage with the socio-political implications of its decisions, would be ill-equipped to handle such cases.

Social Justice and AI's Ethical Blind Spots

Indian jurisprudence often grapples with issues of social justice, particularly in the context of caste discrimination, gender inequality, and the rights of marginalised communities. AI systems, which are trained on historical data, may inadvertently perpetuate the biases and inequalities that exist in the legal system. For instance, in cases involving caste-based discrimination, AI might rely on past judgments that reflect the systemic biases against marginalised groups, thereby reinforcing rather than challenging these biases.

In **Bender et al. (2021)**, the authors warn of the dangers of relying on AI systems that are

trained on biased data. They argue that AI systems, by their very nature, are prone to replicating the biases inherent in their training data . In India, where the legal system often deals with issues of social justice, this limitation is particularly concerning. AI systems, which lack the ability to engage with the ethical dimensions of legal interpretation, may fail to account for the broader social impact of their decisions.

For example, in cases like **Nandini Sundar v. State of Chhattisgarh (2011)**, where the Court had to address issues of state violence against tribal communities, the judgment required the Court to consider not just the legal aspects of the case but also the socio-political context in which the violence occurred. AI systems, which operate on the basis of past legal decisions, would struggle to engage with the broader social justice issues at play in such cases.

AI and the Marginalised: Reinforcing Systemic Inequities

AI's reliance on historical data raises concerns about its ability to address the needs of marginalised communities. As **Bender et al. (2021)** point out, AI systems are often trained on data that reflects the dominant perspectives in society, while marginalising the voices of underrepresented groups . In the Indian legal system, where the courts often play a crucial role in protecting the rights of marginalised groups, this limitation is particularly problematic.

For instance, in cases involving land rights for tribal communities, the courts have often had to balance the state's interest in development with the rights of indigenous populations. These cases require not just an application of property law but an understanding of the historical and cultural significance of land for these communities. AI systems, which operate on the basis of legal precedents, would be ill-equipped to handle such cases, where legal interpretation must account for the broader social and cultural context.

Conclusion: The Limits of AI in the Indian Legal Context

The **socio-legal complexity** of the Indian legal system presents significant challenges for the application of AI in legal interpretation. While AI may excel in structured legal tasks, it is fundamentally limited in its ability to engage with the broader social, cultural, and political context in which legal decisions are made. As cases like **Ayodhya**, **Shayara Bano**, and **Indra Sawhney** demonstrate, legal interpretation in India often requires a deep engagement with

societal values, historical context, and the evolving needs of a pluralistic society—areas where AI is ill-equipped to contribute.

As **Brooks, Gherhes, and Vorley (2020)** argues, AI's reliance on data can perpetuate systemic biases, particularly in cases involving marginalised communities . In the Indian legal context, where the courts play a crucial role in addressing issues of social justice, this limitation is particularly concerning. AI, for all its capabilities in structured tasks, cannot replace human judgment in the complex and evolving realm of legal interpretation.

Human-AI Collaboration in Law—A Way Forward

While the previous sections have illustrated the limitations of AI in handling complex legal interpretation, particularly in the Indian legal context, it would be imprudent to dismiss AI's utility altogether. In many areas of law, AI has already proven its value in augmenting human capacity. The key lies in understanding the appropriate role of AI: as a tool to enhance efficiency in structured tasks rather than as a replacement for human judgment in interpretative legal tasks. This section explores how a collaborative approach—where AI assists human lawyers—could maximize the benefits of both human expertise and technological advancements.

The Role of AI in Structured Legal Tasks

AI systems are highly effective in automating structured tasks that require the processing of vast amounts of data, such as document review, contract analysis, and legal research. As **Achiam, Jin, & Zhang (2023)** demonstrate, LLMs like **LawLLM** have significantly reduced the time and cost associated with tasks like legal drafting, contract management, and document discovery . In areas like tax law or corporate compliance, where decisions are largely rule-based and procedural, AI can handle the heavy lifting of data processing, freeing up human lawyers to focus on more complex strategic work.

In law firms around the world, AI-powered tools are being deployed to improve efficiency and accuracy. For instance, **ROSS Intelligence**, an AI legal research platform, helps lawyers quickly find relevant case law by processing natural language queries and searching through millions of legal documents. Similarly, platforms like **Kira Systems** and **LawGeex** are

revolutionizing contract analysis by identifying risky clauses and suggesting changes in a fraction of the time it would take a human lawyer.

These AI tools excel in structured environments where the tasks are well-defined, and the variables are limited. As **Katz, Bommarito, & Blackman (2024)** argue, AI systems are best suited to handle routine, repetitive tasks that require the application of established rules and procedures . In this context, AI serves as a valuable tool that increases productivity and reduces the likelihood of human error.

AI as an Augmentative Tool in Legal Services

The potential of AI lies not in replacing human judgment but in augmenting human capabilities. AI can assist lawyers by performing time-consuming tasks more quickly and efficiently, allowing legal professionals to focus on higher-level reasoning and strategy. This symbiotic relationship between AI and human lawyers is particularly important in areas like due diligence, compliance monitoring, and litigation preparation.

For example, in the field of e-discovery, AI tools can sift through vast amounts of digital data to identify relevant documents for a case. This capability not only saves time but also helps uncover key evidence that might otherwise be overlooked. **Brooks, Gherhes, and Vorley (2020)** emphasizes the importance of this human-AI collaboration, noting that AI's ability to process large datasets at lightning speed allows lawyers to focus on the strategic aspects of a case, rather than getting bogged down in tedious document review .

Moreover, AI can assist in generating predictive analytics for litigation. By analyzing past cases and judgments, AI systems can help lawyers assess the potential outcomes of a case, providing valuable insights into which legal strategies might be most effective. This does not eliminate the need for human judgment but rather supports it by providing data-driven insights that inform decision-making.

Ethical and Regulatory Considerations

While AI can be a valuable tool for legal professionals, its deployment must be carefully regulated to avoid potential ethical pitfalls. One of the primary concerns surrounding the use of AI in law is the lack of transparency in how AI systems arrive at their decisions. **Kemper**

(2020) highlights the "black box" problem in AI, where the decision-making processes of AI systems are often opaque, making it difficult for users to understand how the AI arrived at a particular conclusion . This lack of transparency raises significant concerns about accountability and fairness, particularly in legal contexts where the stakes are high.

In response to these concerns, some legal scholars have called for greater regulation of AI in law. **Brooks, Gherhes, and Vorley (2020)** argues that AI systems should be subject to rigorous testing and certification processes before they are deployed in legal services . Additionally, legal professionals must remain vigilant about the ethical implications of using AI, particularly in cases involving sensitive information or marginalized communities.

In India, the adoption of AI in law has been slow, in part due to concerns about its ethical implications. The country's legal community has expressed reservations about using AI in tasks that require discretion, judgment, and moral reasoning—areas where human oversight is crucial. As **American Bar Association (2024)** notes, the risk of job displacement, confidentiality breaches, and algorithmic bias are significant challenges that must be addressed before AI can be fully integrated into legal practice.

Human Oversight as a Safeguard

One potential solution to the ethical concerns surrounding AI in law is the implementation of robust oversight mechanisms that ensure human control over AI systems. In this model, AI would handle the technical aspects of legal work, while human lawyers would maintain ultimate responsibility for interpreting and applying the law. This approach aligns with **Biresaw (2022)**, who advocates for a collaborative model in which AI serves as a tool to assist legal professionals, rather than replace them.

In this collaborative model, human lawyers would review and interpret the outputs generated by AI, ensuring that the AI's recommendations align with legal principles and ethical considerations. By maintaining human oversight, law firms can harness the benefits of AI while mitigating the risks associated with algorithmic decision-making. This approach also ensures that legal decisions are made in a way that respects the moral and ethical dimensions of the law, which AI systems are ill-equipped to handle.

A Path Forward: Regulated and Ethical AI Adoption

The future of AI in law will depend on how well the legal community navigates the ethical, regulatory, and practical challenges associated with its adoption. To ensure that AI is used responsibly, legal professionals must establish clear guidelines for its deployment, with a focus on transparency, accountability, and fairness. By creating robust regulatory frameworks, the legal profession can leverage AI's strengths while safeguarding against its potential harms.

In India, where the legal system is deeply intertwined with social justice issues, the adoption of AI must be approached with caution. As **Surden (2019)** notes, AI's reliance on data from past legal decisions can perpetuate existing biases, particularly in cases involving marginalized communities. Therefore, any deployment of AI in the Indian legal system must be accompanied by strict oversight and ethical safeguards to ensure that the technology does not reinforce systemic inequalities.

Conclusion: The Future of Human-AI Collaboration in Law

AI has the potential to transform legal services by increasing efficiency, reducing costs, and assisting lawyers in complex tasks. However, its role should be one of augmentation rather than replacement. As this paper has demonstrated, AI is highly effective in handling structured legal tasks but remains fundamentally limited in its ability to engage with the moral, ethical, and societal dimensions of legal interpretation. By embracing a collaborative approach—where AI assists human lawyers but does not replace them—the legal profession can harness the strengths of AI while maintaining the human judgment that is essential for justice.

The path forward involves balancing the benefits of AI with the ethical considerations that arise from its use. By implementing strong regulatory frameworks and maintaining human oversight, the legal profession can ensure that AI serves as a tool for enhancing, rather than undermining, the pursuit of justice.

References

1. Achiam, R., Jin, J., & Zhang, Y., LawLLM: Law Large Language Model for the US Legal System, (2023), available at https://arxiv.org/abs/2407.21065.

2. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S., 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), available at https://doi.org/10.1145/3442188.3445922.

3. Katz, D. M., Bommarito, M. J., & Blackman, J., ChatGPT, Large Language Models, and Law, Fordham Law Review, (2024), available at https://fordhamlawreview.org/issues/chatgpt-large-language-models-and-law.

4. Surden, H., Artificial Intelligence and Law, 35 Ga. St. U. L. Rev. 1305 (2019), available at https://readingroom.law.gsu.edu/gsulr/vol35/iss4/8.

5. Zhang, Y., Sun, L., & Huang, S., Large Language Models in Law: A Survey, (2023), available at https://arxiv.org/abs/2312.03718.

6. Leiter, B., Naturalizing Jurisprudence: Essays on American Legal Realism and Naturalism in Legal Philosophy, (Oxford University Press, 2007).

7. Brooks, C., Gherhes, C., & Vorley, T., Artificial Intelligence in the Legal Sector: Pressures and Challenges of Transformation, 13 Cambridge Journal of Regions, Economy and Society 135 (2020), available at https://academic.oup.com/cjres/article/13/1/135/5716343.

8. **Carolin Kemper**, *Kafkaesque AI? Legal Decision-Making in the Era of Machine Learning*, 24 *University of San Francisco Intellectual Property and Technology Law Journal* 2 (2020), available at https://ssrn.com/abstract=3787498.

9. A.D. (Dory) Reiling, *Courts and Artificial Intelligence*, 11 *IACA Journal* 8 (2020), available at https://doi.org/10.36745/ijca.343.

10. **American Bar Association**, *Ethics Opinion: Use of AI in Legal Practice*, (Aug. 2024),

available at https://www.americanbar.org/groups/health_law/section-news/2024/august/aba-ethics-opinion-use-of-ai/.

11. Samuel Maireg Biresaw, The Impacts of Artificial Intelligence on Research in the Legal Profession, 5(1) International Journal of Law and Society 53-65 (2022), available at https://doi.org/10.11648/j.ijls.20220501.17.

Case Law References:

1. Shreya Singhal v. Union of India, (2015) 5 SCC 1 (India).
2. Navtej Singh Johar v. Union of India, (2018) 10 SCC 1 (India).
3. K.S. Puttaswamy v. Union of India, (2017) 10 SCC 1 (India).
4. Shayara Bano v. Union of India, (2017) 9 SCC 1 (India).
5. Indra Sawhney v. Union of India, AIR 1993 SC 477 (India).
6. Kesavananda Bharati v. State of Kerala, (1973) 4 SCC 225 (India).
7. Nandini Sundar v. State of Chhattisgarh, (2011) 7 SCC 547 (India).
8. M. Siddiq v. Mahant Suresh Das, (2019) 15 SCC 1 (India).