PATENT DISCLOSURE AND BLACK-BOX AI SYSTEMS: DOCTRINAL BREAKDOWN AND THE LIMITS OF THE ENABLEMENT REQUIREMENT

Amirtha K, Presidency University, Bangalore

ABSTRACT

The rapid advancement of artificial intelligence (AI), particularly through machine learning systems, has begun to generate significant tensions within patent law, especially regarding disclosure requirements. Many of these systems exhibit opacity, unpredictability, and are heavily reliant on data, complicating their explanation in conventional terms. Central to patent disclosure is the enablement requirement, which mandates that inventors articulate their inventions sufficiently for a person possessing ordinary skill in the art (PHOSITA) to make and utilize them without undue experimentation.

The challenge arises with modern AI systems, particularly the so-called "black-box" models, which do not conform to these expectations. Their internal mechanisms are often difficult to interpret, challenging to reproduce, and sometimes not fully comprehensible even to their creators. This brings forth a crucial inquiry: how can one adequately "disclose" that which cannot be fully elucidated?

This paper examines the legal foundations of the enablement requirement and assesses its efficacy when applied to AI-based inventions. It posits that black-box AI reveals both conceptual and practical deficiencies within the current patent framework, thereby destabilizing the traditional quid pro quo of patent law. By analyzing legal developments, case law trends, and policy concerns, the paper contends that the existing approach to enablement may be insufficient for fostering AI innovation. Ultimately, it advocates for a re-evaluation of the framework, potentially by embracing more flexible disclosure models or alternative standards of reproducibility, to ensure that the patent system continues to promote innovation while preserving its core objectives.

Keywords: Artificial Intelligence (AI), Enablement Requirement, Black-Box Models, Patent Disclosure, Reproducibility in Machine Learning

1. Introduction

Patent law is fundamentally predicated on a basic trade-off: in exchange for a limited monopoly, inventors are obligated to disclose their inventions in adequate detail, enabling others to understand, utilize, and further develop them. This principle, often articulated as the quid pro quo of the patent system[1], transcends mere technical regulation; it underpins the system's capacity to foster innovation while simultaneously serving the public interest. The enablement requirement is pivotal in this context, ensuring that patents contribute to public knowledge[2] rather than sequestering it.

However, the emergence of artificial intelligence, particularly in the realm of machine learning, threatens to disrupt this equilibrium. Traditional inventions typically lend themselves to step-by-step explanation, characterized by clear processes and logic. In contrast, AI systems often function through intricate models that learn from vast datasets, producing outputs based on patterns rather than fixed instructions. Consequently, even the developers may lack complete understanding of how specific results are generated, giving rise to the so-called "black-box" problem[3].

This scenario presents a significant challenge for patent law. If an invention cannot be clearly articulated, can it truly be considered properly disclosed? This issue transcends mere technicality; it strikes at the heart of patent law's foundational assumptions regarding knowledge. The system is predicated on the notion that inventions can be described, shared, and comprehended. However, black-box AI systems frequently defy this notion, as their inner workings may be excessively complex or insufficiently transparent to facilitate conventional communication.

This paper addresses this dilemma, asserting that the challenges posed by black-box AI are not trivial concerns amenable to minor adjustments. Rather, they unveil profound limitations in the current understanding of disclosure and enablement within patent law.

2. The Enablement Requirement: Legal Foundations

2.1 Statutory Basis and Objective

[1] William M. Landes & Richard A. Posner, The Economic Structure of Intellectual Property Law 294–95 (2003).
[2] Jeanne C. Fromer, Patent Disclosure, 94 Iowa L. Rev. 539 (2009).
[3] Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (2015).

The enablement requirement constitutes a fundamental tenet of patent law, necessitating that an invention be described in "full, clear, concise, and exact terms" [4] to enable a skilled person to make and utilize it without undue experimentation. While this requirement may appear straightforward in theory, it embodies a broader principle: patent protection is warranted only when the inventor contributes meaningfully to public knowledge.

Historically, this requirement has functioned as a crucial safeguard within the system, preventing inventors from asserting broad claims [5]without adequately elucidating their inventions. In essence, a monopoly is granted only when the inventor agrees to impart knowledge about the workings of their invention, thereby preserving a balance between private advantage and public accessibility.

The challenge arises from the assumption that inventions can be clearly articulated. This assumption falters when we examine AI systems, wherein the relationship between input, process, and output often lacks clarity. Even comprehensive descriptions may fail to capture the nuances of how a system derives its results.

Thus, while the language of enablement remains applicable, its practical utility becomes questionable in the context of AI. The law anticipates clarity and reproducibility, yet AI systems do not consistently provide either in a manner that aligns with traditional disclosure standards.

2.2 The PHOSITA Standard

The PHOSITA standard aims to introduce a degree of flexibility into the enablement requirement. Rather than assessing an invention's comprehensibility in absolute terms, the law evaluates it through the lens of a hypothetical individual possessing ordinary skill in the relevant field.[6] This approach assists courts and patent examiners in determining the sufficiency of disclosure, contingent upon the complexity or predictability of the technology. In principle, this framework is effective; it alleviates undue pressure on inventors while ensuring that others in the field can utilize the invention.

[4] Patents Act, 1970, § 10(4).
[5] Mark A. Lemley, The Myth of the Sole Inventor, 110 Mich. L. Rev. 709 (2012).
[6] John M. Golden, The Supreme Court as “Prime Percolator”: A Prescription for Appellate Review of Questions in Patent Law, 56 UCLA L. Rev. 657 (2009).

However, the application of this standard in the context of AI is far from straightforward. A pivotal challenge lies in delineating who constitutes this "ordinarily skilled person." AI is not a fixed domain[7]; it is continually evolving and intersects with diverse fields such as computer science, statistics, and data science. Consequently, practitioners in AI often possess varying skill sets, rendering it challenging to define what "ordinary" encompasses in this context.

This ambiguity impacts the application of the standard. If we presume that the PHOSITA possesses advanced skills, the requisite level of disclosure diminishes, potentially permitting patents that lack specificity. Conversely, if we envision a less specialized PHOSITA, inventors may be expected to provide explanations that exceed realistic capabilities. In either scenario, the absence of clarity complicates the consistent application of the enablement requirement in AI-related cases.

2.3 Excessive Experimentation and the Wands Factors

The concept of "undue" or excessive experimentation serves as a practical criterion for evaluating whether an invention has been adequately disclosed. Courts generally accept that a degree of trial and error is inherent, particularly in complex domains. However, there are limits; if an individual must engage in excessive conjecture to realize the invention, the disclosure is deemed insufficient. The factors established in ***In re Wands***[8] aid in this assessment by considering aspects such as the extent of experimentation required, the guidance provided, and the predictability of the field.

This framework proves effective in traditional technological contexts. However, its applicability becomes significantly more challenging in the realm of AI. Machine learning systems are influenced by numerous variables such as the data utilized, the model architecture, and parameter adjustments and even minor alterations in these elements can yield markedly different outcomes[9]. Consequently, even a patent that offers detailed exposition may still necessitate considerable effort to accurately reproduce the system.

As a result, the definition of what constitutes "undue" experimentation becomes ambiguous. AI systems, by their very nature, are not entirely predictable. Thus, the pertinent question

[7] Stuart Russell & Peter Norvig, Artificial Intelligence: A Modern Approach (4th ed. 2021).
[8] In re Wands, 858 F.2d 731 (Fed. Cir. 1988).
[9] Ian Goodfellow et al., Deep Learning (MIT Press 2016).

arises: how much experimentation constitutes "too much" in a discipline where some level of uncertainty is inevitable? A stringent standard may preclude numerous legitimate AI inventions, while a lenient one risks undermining the requirement's significance. This illustrates the challenges inherent in applying existing legal tests to emergent technologies like AI.

3. Artificial Intelligence and the Black-Box Dilemma

3.1 Characteristics of AI Systems

AI systems, particularly those grounded in machine learning, do not conform to the traditional conception of an invention. In most conventional technologies, one can elucidate the workings of an invention in a clear, sequential manner. AI systems diverge from this paradigm; they learn from data rather than adhering to fixed protocols, employing methodologies such as supervised learning, unsupervised learning, or reinforcement learning to discern patterns and make determinations.[10]

This complexity renders the explanation of AI systems challenging. The "logic" underpinning these systems is not articulated in a straightforward manner; rather, it resides within layers of computations, parameters, and interactions with data. Even when a system performs effectively, deconstructing the precise manner in which it arrived at a specific outcome can prove difficult.

This raises a fundamental yet crucial inquiry: what precisely constitutes the invention within an AI system? Is it the algorithm itself, the training methodology, the data utilized, or the resultant model that emerges from the process? The absence of a clear answer complicates the consistent application of traditional patent rules regarding disclosure.[11]

3.2 The Black-Box Challenge

The opaque nature of numerous AI systems underscores a fundamental concern. While these systems can yield highly accurate results, the rationale behind their conclusions often remains elusive, even to seasoned experts[12]. In models such as deep neural networks, the decision-

[10] Kevin P. Murphy, Machine Learning: A Probabilistic Perspective (2012).
[11] Ryan Abbott, The Reasonable Robot: Artificial Intelligence and the Law (2020).
[12] Zachary C. Lipton, The Mythos of Model Interpretability, 16 Queue 31 (2018).

making process is obscured by numerous layers of computation, complicating the ability to trace the origins of a specific output.

From the perspective of patent law, this presents a significant challenge. The primary objective of disclosure is to ensure that others can comprehend and replicate the invention. However, if the internal mechanisms of a system cannot be elucidated clearly, fulfilling this requirement becomes increasingly difficult. In essence, while the invention may perform exceptionally well in practice, it cannot be adequately articulated.

This disconnect between performance and explainability also raises broader issues. In critical sectors such as healthcare, finance, and criminal justice, the deployment of black-box AI has ignited discussions surrounding fairness, bias, and transparency. [13]When decisions are derived from systems that lack full explainability, it becomes more challenging to scrutinize or validate those decisions. This emphasizes the importance of disclosure and raises questions regarding the appropriateness of granting patent protection to such systems without clearer elucidation.

3.3 Stochastic and Non-Deterministic Dynamics

Another salient characteristic of AI systems is their inherent unpredictability. Unlike conventional software, which typically produces the same output for identical inputs, machine learning models may exhibit varying behaviors influenced by factors such as training methodologies, data sets, or even random initial conditions. This leads to a degree of variability in their operations.[14]

This variability poses a problem for patent law, where reproducibility is paramount. If an invention cannot be reliably reproduced, it becomes challenging to assert that it has been adequately disclosed. In the context of AI, even minor alterations in training conditions can yield significantly different results, complicating consistent reproduction compared to traditional technologies.

Consequently, AI does not conform neatly to the established expectations of patent law. The framework presupposes that an invention can be recreated by adhering to provided instructions. However, this assumption does not always hold true with AI, prompting a broader inquiry into

[13] Solon Barocas, Moritz Hardt & Arvind Narayanan, Fairness and Machine Learning (2019).
[14] Christopher M. Bishop, Pattern Recognition and Machine Learning (2006).

whether the current interpretation of enablement is sufficiently adaptable to accommodate inherently unpredictable technologies.

4. Doctrinal Disruption: Enablement Meets Black-Box AI

4.1 The Reproducibility Challenge

Reproducibility constitutes a fundamental aspect of the enablement requirement[15], yet it is precisely here that AI systems introduce complications. Unlike traditional inventions, machine learning models are dynamic and not static. They are developed through a complex interplay of data, algorithms, and multiple computational layers. To replicate such a system, one often requires comprehensive access to information regarding the training data, model architecture, parameters, and even the computing environment utilized.

In numerous instances, this information is not fully disclosed. It may be too extensive to incorporate into a patent, kept confidential, or simply too unstable to reproduce accurately. Consequently, patents may articulate the general methodology or concept behind the system, yet fail to provide sufficient detail for others to successfully recreate it.

Even when additional details are supplied, there remains no assurance that identical results can be attained. Due to the inherent variability in training AI systems, outcomes may diverge even when similar procedures are followed. This renders the conventional notion of reproducibility where an outcome is either successful or not less applicable in the AI context. Instead, reproducibility may exist on a spectrum, raising questions regarding the acceptable degree of variation.[16]

This scenario presents a deeper challenge for patent law. If the public cannot realistically utilize what has been disclosed, the foundational exchange underpinning the patent system begins to deteriorate. This suggests that the prevailing understanding of enablement may require reevaluation in light of AI technologies.

4.2 The Explainability Divide

Closely intertwined with reproducibility is the issue of explainability. Patent law has

[15] National Academies of Sciences, Reproducibility and Replicability in Science (2019).
[16] Pamela Samuelson, Trade Secrets vs. Patents for AI, 21 Stan. Tech. L. Rev. 1 (2018).

historically favored clear elucidations of an invention's operation. Traditionally, inventors are expected to delineate the structure or process underlying their invention in a manner that others can replicate. However, this task becomes increasingly challenging with black-box AI systems, where internal mechanisms are not easily articulated.[17]

As a result, inventors often resort to describing the functionality of the system rather than its operational intricacies. While this approach may demonstrate the utility of the invention, it does not adequately facilitate replication by others. Absent a clear exposition of the underlying process, a skilled individual may still need to engage in extensive trial and error to effectively implement the system.

This discrepancy creates a schism between the actual development of AI and the expectations of patent law. More critically, it raises a fundamental question: should exclusive rights be conferred upon systems that are not fully comprehended or explicated? If even the inventor cannot provide a clear description of the invention's workings, the justification for restricting others from engaging in the same domain becomes tenuous.

In this regard, the challenge is not merely technical; it fundamentally questions the rationale for patent grants. It challenges the premise that protection should be predicated on meaningful disclosure and suggests that the existing framework may not be entirely equipped to address the complexities associated with AI technologies.

4.3 Overbreadth and Claim Scope

A significant concern in the realm of artificial intelligence (AI) patents is the issue of claim scope. Inventors often endeavor to draft expansive claims to encompass a multitude of variations of their inventions. From a commercial perspective, this approach is rational, as it enhances the patent's value and broadens its protective reach. However, patent law stipulates that claims must not exceed the disclosures made by the inventor.

Judicial precedents have established that the enablement requirement must be fulfilled throughout the entire scope of the claim. This necessitates that the patent must not merely describe a singular embodiment of the invention; rather, it must furnish [18]sufficient detail to

[17] Dan L. Burk & Mark A. Lemley, The Patent Crisis and How the Courts Can Solve It (2009).
[18] Amgen Inc. v. Sanofi, 598 U.S. 594 (2023).

enable the creation and utilization of all claimed variations.

The complexities introduced by AI exacerbate this issue. AI systems can differ significantly based on the data utilized, the model architecture, and the training methodologies employed. Even minor modifications can result in entirely distinct implementations. Consequently, disclosing every conceivable iteration of an AI-based invention is often impractical. Thus, numerous AI patents face challenges of overbreadth, as the disclosures may not adequately support all claims made.

This situation presents a practical dilemma for inventors. Narrowing claims facilitates compliance with the enablement requirement but limits the scope of protection. Conversely, broader claims may provide enhanced coverage but introduce a heightened risk of invalidation. This inherent tension between claim breadth and adequate disclosure is particularly pronounced in the context of AI.

4.4 The Black-Box as a Legal Barrier

At a fundamental level, the black-box problem transcends technical challenges; it fundamentally interrogates the foundational tenets of patent law. The essence of enablement is predicated on the premise that an invention can be comprehended, articulated, and replicated by others. However, black-box AI systems frequently do not conform to this paradigm.

In numerous instances, the "knowledge" encapsulated within these systems cannot be succinctly documented or systematically elucidated. It is dispersed across various layers of the model and influenced by data in ways that lack full transparency. This opacity complicates the task of describing the invention in a manner that satisfies traditional disclosure obligations.[19]

This leads to a broader inquiry: what if certain categories of AI inventions are inherently incompatible with the patent system? If an inventor is unable to elucidate the operational mechanics of the system, justifying the conferral of exclusive rights becomes increasingly problematic.

Viewed in this light, the black-box issue extends beyond mere drafting challenges; it calls into question the foundational principles upon which patents are granted. It implies that the current

[19] Michael Polanyi, Personal Knowledge: Towards a Post-Critical Philosophy (1958).

legal framework may be inadequately equipped to address certain modern technological advancements. Whether the resolution lies in revising the enablement standard, developing alternative forms of protection, or fundamentally rethinking the role of patents warrants further exploration.

Although it does not directly address enablement**, *Ex parte Allen*** [20]sheds light on a related concern within patent evaluation—the risk of relying on speculative reasoning rather than substantial technical disclosure. The Board rejected an obviousness claim in which the examiner sought to combine prior art based on a problem that had already been resolved. This reasoning is particularly significant in the context of AI, where claims are sometimes supported by broad, outcome-oriented descriptions that lack sufficient technical detail. As with Allen, reliance on assumptions or hypothetical improvements, rather than authentic disclosure, can jeopardize the validity of patent claims.

5. Case Law and Administrative Developments

The evolution of patent law in relation to AI has not solely been dictated by statutes; it is also being shaped by judicial rulings and directives from patent offices. Entities such as the United States Patent and Trademark Office[21] and the European Patent Office [22]have begun to scrutinize the disclosure practices pertaining to AI-related inventions more closely.

Recently, these offices have articulated that a mere reference to "AI" or "machine learning" is insufficient. Applicants are now expected to provide detailed information, including the structure of the model, the training processes, the types of data utilized, and the specific technical contributions of the invention. The objective is to prevent applicants from claiming excessively broad rights based on ambiguous descriptions.

However, this requirement introduces practical challenges. AI systems can be extraordinarily intricate, and disclosing excessive detail in a patent may inadvertently expose sensitive or proprietary information. Conversely, insufficient detail risks non-compliance with the enablement requirement. Thus, applicants often find themselves navigating the delicate balance between clarity and the protection of their proprietary interests.

[20] Ex parte Allen, Appeal No. 2018-008208 (P.T.A.B. July 1, 2020)
[21] USPTO, Public Views on Artificial Intelligence and Intellectual Property Policy (2020).
[22] European Patent Office, Guidelines for Examination (2023).

Courts have also adopted a more stringent stance. There is an increasing emphasis on the necessity for the invention to be enabled across the full scope of the claim. If the disclosure fails to provide adequate guidance for implementation without undue experimentation, the patent may face invalidation.

In the realm of AI, this has resulted in heightened scrutiny of patents that prioritize end results over the methodologies employed to achieve those results. Courts appear cautious about permitting broad claims that could encompass extensive areas of technological development without corresponding levels of disclosure. This raises a pertinent concern: are legal standards crafted for more predictable technologies being applied too rigidly to AI?

This tension illustrates that patent law is in a continuous state of adaptation in response to AI. While there is a discernible shift toward stricter disclosure requirements, an underlying uncertainty persists regarding the suitability of the existing framework for this category of technology.

A significant legal development at the intersection of patent law and artificial intelligence is exemplified by the case of ***Thaler v. Vidal***[23]**.** In this matter, the U.S. Court of Appeals for the Federal Circuit ruled that only natural persons may be recognized as inventors under the Patent Act. This case arose from Stephen Thaler's attempt to designate his AI system, DABUS, as the sole inventor of two patent applications. The court rejected this proposition, concluding that the statutory interpretation of the term "individual" pertains exclusively to human beings. While this ruling does not preclude patents for inventions facilitated by AI, it establishes a definitive boundary by asserting that a discernible human contribution is essential to the inventive process.

This judgment is particularly relevant in the context of AI-driven innovation, reflecting the judiciary's broader reluctance to extend traditional patent concepts to non-human entities. More importantly, it highlights a deeper structural concern: patent law continues to operate under a human-centric model of invention, which is at odds with the increasing autonomy and complexity of contemporary AI systems.

In this respect, ***Thaler v. Vidal***[24] underscores the argument that existing legal frameworks may

[23] Thaler v. Vidal*,* 43 F.4th 1207 (Fed. Cir. 2022), cert. denied, 143 S. Ct. 1783 (2023).
[24] Thaler v. Vidal*,* 43 F.4th 1207 (Fed. Cir. 2022), cert. denied, 143 S. Ct. 1783 (2023).

be ill-equipped to address the challenges posed by AI, particularly those related to disclosure and enablement.

6. The Constraints of the Enablement Requirement

6.1 Knowledge Limitations

At its core, the challenges presented by opaque artificial intelligence (AI) systems signify that the issues at hand extend beyond legal parameters; they also pertain to the nature of knowledge itself. The enablement requirement presupposes that an invention can be understood and subsequently articulated with clarity to others. This presumption typically holds for conventional technologies, where the inventor possesses a clear understanding of the invention's operation and can describe it in a linear manner.[25]

However, AI systems frequently diverge from this model. In many cases, even the developers are unable to fully explain how the system derives a particular outcome. The model "learns" from data, and its internal logic is not always readily articulated. This presents a significant dilemma for patent law: how can an inventor be expected to disclose something that is not entirely comprehended?

This raises a more profound question: should patent law mandate a level of explanation that may not be genuinely attainable within AI systems? If the answer is affirmative, a substantial number of AI innovations may remain ineligible for patent protection. Conversely, if the response is negative, the fundamental purpose of the enablement requirement diminishes, as it fails to ensure meaningful disclosure.[26]

This predicament also alters our understanding of the inventor's role. Patent law typically assumes that the inventor possesses a comprehensive understanding of their creation. In the realm of AI, this assumption begins to wane. Although the system may produce valuable results, the mechanisms underlying those results are not entirely transparent. Consequently, a widening gap exists between the act of invention and complete comprehension, creating significant tension within the current legal framework.

[25] Thomas S. Kuhn, The Structure of Scientific Revolutions (1962).
[26] Harry Collins, Tacit and Explicit Knowledge (2010).

6.2 Practical Constraints

Beyond these theoretical considerations, substantial practical challenges arise in the enforcement of the enablement requirement for AI. One of the primary obstacles pertains to data. Machine learning systems often depend on extensive datasets, which are impractical to include within a patent document. In many instances, these datasets are also confidential or contain sensitive information, further complicating the requirement for full disclosure.

Even when inventors endeavor to describe the data in broad terms, such descriptions may prove insufficient. Minor alterations in the dataset can yield drastically different outcomes, rendering the reproduction of the system unfeasible without the precise data.

Another significant concern involves resource requirements. Training advanced AI models typically necessitates powerful hardware, such as high-performance GPUs or substantial computing systems. Not all qualified individuals in the field will have access to such resources. This raises a critical question: should enablement be assessed based on theoretical possibilities or practical attainability?

If we prioritize practical feasibility, many AI inventions may struggle to meet the criteria, as they cannot be easily replicated. Conversely, if we consider only theoretical potential, the requirement loses its significance by disregarding real-world constraints.

These considerations suggest that the traditional "one-size-fits-all" approach to enablement is ill-suited for AI. The unique characteristics of these systems necessitate an evolution of the law.

6.3 Policy Conflicts

These challenges also give rise to broader policy dilemmas. Patent law is designed to simultaneously achieve two objectives: stimulate innovation and ensure the dissemination of knowledge to the public. In the context of AI, reconciling these objectives becomes markedly more complex.

On one hand, robust patent protection can incentivize companies to invest in AI research, which often requires significant time, capital, and expertise. Without adequate protection, the motivation to develop new technologies may diminish.

On the other hand, the patent system is intended to foster transparency. If patents are issued without sufficient disclosure, they fail to enhance public knowledge meaningfully. Instead, they may obstruct others from participating in the same field without offering substantial value in return.

This creates a challenging predicament. If the enablement requirement is enforced too stringently, companies might opt to forgo patents altogether, relying instead on trade secrets, thereby limiting public access to information. Conversely, if the requirement is relaxed excessively, the system risks granting broad rights without adequate disclosure, undermining its credibility.[27]

Thus, the challenge lies in identifying a balanced approach. The law must protect innovation while also fulfilling its essential mission of knowledge sharing. In the domain of AI, achieving this equilibrium is far from straightforward, and it is evident that the existing framework is under considerable strain.

7. Potential Doctrinal and Policy Reactions

In light of the challenges identified, a range of potential doctrinal and policy responses has been proposed to reconcile the enablement requirement with the characteristics of artificial intelligence (AI) systems. One approach involves enhancing disclosure standards to require more extensive information regarding training datasets, model architectures, and performance metrics. This initiative seeks to ensure that AI inventions are articulated with sufficient clarity to facilitate replication, even if their internal operations remain somewhat opaque[28]. However, this approach may impose significant burdens on inventors and raise concerns regarding the protection of proprietary information.

An alternative response could be to adopt a more flexible interpretation of enablement, allowing for functional descriptions in circumstances where structural details are not feasible. This would represent a departure from traditional doctrine, yet it may be necessary to accommodate the unique nature of AI systems. Nonetheless, such flexibility must be carefully monitored to avoid endorsing overly broad claims that exceed the inventor's actual

[27] Mark A. Lemley, Surviving Patent Law's AI Disruption, 97 Tex. L. Rev. 1 (2019).
[28] WIPO, Revised Issues Paper on AI and IP Policy (2020).

contributions.

A more nuanced solution may involve developing hybrid disclosure models that integrate elements from both strategies. For example, patent applicants could be required to provide partial disclosures of internal processes alongside external validation data, benchmarking results, and reproducibility protocols. This approach would shift the emphasis from complete transparency to practical replicability, thereby aligning the enablement requirement with the realities of AI innovation.

Institutional responses may also be crucial. Patent offices could formulate specialized guidelines and examination procedures tailored to AI-related inventions, incorporating interdisciplinary insights and standardized disclosure formats. Such initiatives could promote consistency and clarity in the application of enablement criteria while adapting to the rapidly evolving landscape of AI technologies.

8. Comparative Viewpoints

A comparative analysis of various jurisdictions reveals a diversity of strategies addressing the challenges posed by AI and patent law, reflecting differing policy priorities and legal traditions. The European Patent Office, for instance, underscores the technical nature of inventions and mandates that AI-related claims demonstrate a tangible technical effect. This position indirectly reinforces the need for detailed disclosure, as applicants must articulate the specific technical contribution of their invention. Conversely, the United States has focused more directly on the legal standards of enablement and written description, with judicial decisions significantly shaping the doctrine through case law.

Emerging jurisdictions, such as India, are also grappling with these issues, often striving to foster technological innovation while balancing broader considerations of access and public welfare. [29]The inconsistency across jurisdictions presents both challenges and opportunities. On one hand, it may lead to forum shopping and disparities in patent protection; on the other hand, it allows for experimentation with varied regulatory frameworks, potentially resulting in the identification of best practices that could inform global standards.

[29] NITI Aayog, National Strategy for Artificial Intelligence (2018).

9. Future Consequences

The trajectory of AI development suggests that the challenges associated with enablement are likely to intensify in the future. As systems become increasingly complex and less interpretable, the gap between technical capacity and legal obligations may widen. This could encourage a gradual shift away from patent protection toward alternative mechanisms such as trade secrets, particularly in fields where disclosure is problematic or undesirable.[30] Such a transition would have significant implications for the transparency and accessibility of technological knowledge.

At the same time, the growing importance of AI in critical sectors such as healthcare, finance, and governance underscores the need for robust accountability frameworks. Patent disclosure, if appropriately calibrated, could play a vital role in promoting transparency and oversight. However, this will require a reevaluation of traditional doctrines to ensure they remain relevant and effective in the context of technological advancements.

The evolution of patent law in response to AI may also have broader implications for the legal system as a whole, prompting a reassessment of how law interacts with complex, data-driven technologies. This may lead to a deeper integration of technical expertise into legal processes and a more dynamic approach to doctrinal evolution.

10. Conclusion

The emergence of black-box AI systems presents a complex and multifaceted challenge to the traditional framework of patent law, particularly concerning the foundational enablement requirement. As this paper has demonstrated, the opacity, unpredictability, and data dependence of contemporary AI systems expose both epistemic and practical limitations within existing disclosure doctrines. The imperative for inventions to be thoroughly understood, described, and reproducible is becoming increasingly difficult to fulfill in a technological landscape where innovation often outpaces comprehension. This challenge transcends mere technical complexity, reflecting a significant misalignment between the assumptions embedded in patent law and the realities of modern scientific and technological practices.

At a broader level, the challenges associated with enablement in artificial intelligence (AI)

[30] Deepa Varadarajan, Trade Secret Fair Use, 83 Fordham L. Rev. 1401 (2015).

transcend mere legal concerns; they signify a fundamental transformation in our comprehension of knowledge itself[31]. The foundational principles of patent law originate from an era when knowledge was perceived as something that could be explicitly documented, systematically articulated, and transmitted in a stable format. This premise is effective for most conventional technologies. However, black-box AI systems often resist this conventional framework.

In numerous AI systems, the "invention" does not manifest as a tidy compilation of instructions or a distinctly delineated mechanism.[32] Rather, it arises from the interplay of data, algorithms, and training processes. Individually, none of these components adequately elucidates the operational mechanics of the system. Consequently, articulating the invention in traditional terms poses significant challenges. This prompts a critical inquiry: can patent law, in its current iteration, adequately accommodate such technology, or must it undergo evolution?

The ramifications of this issue extend beyond legal doctrine and resonate with the overarching objectives of the patent system. Ideally, patents are intended to fulfill two primary functions: to foster innovation and to disseminate knowledge to the public. However, in the context of AI, there exists a risk that the system may falter in achieving either goal. If the enablement requirement is enforced too stringently, innovators may opt to forgo patents entirely in favor of trade secrets, thereby constraining knowledge sharing. Conversely, if disclosure standards are excessively relaxed, patents may be issued without a substantive contribution to public understanding, potentially undermining the credibility of the system and hindering further innovation.

This scenario places patent law in a precarious position, necessitating a balance between safeguarding innovation and ensuring meaningful disclosure. One potential pathway forward involves reexamining the expectations surrounding enablement. Instead of mandating a comprehensive elucidation of an AI system's internal workings, the focus could shift to whether the invention can be practically reproduced. In this context, the inquiry would not center on the completeness of the explanation, but rather on whether a skilled individual can achieve comparable results utilizing the information provided. This approach would acknowledge the inherent opacity of AI while preserving the essential principle of public access. Nevertheless,

[31] Thomas S. Kuhn, The Structure of Scientific Revolutions 52–65 (4th ed. 2012).
[32] Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information 3–15 (2015).

this must be approached with caution to prevent the issuance of overly broad or ambiguous patents.

An alternative strategy may involve looking beyond traditional patent specifications. Given the complexities inherent in AI systems, it may not always be feasible to encapsulate everything within a written document. Innovative methods such as employing standard datasets, model repositories, or independent validation could enhance the reproducibility of inventions without necessitating exhaustive disclosure. Additionally, the exploration of deposit systems, akin to those utilized in biotechnology, where critical materials are stored and made accessible under regulated conditions, could represent a more pragmatic and adaptable approach to disclosure.

Concurrently, institutions such as patent offices and courts will be instrumental in determining how these matters are addressed. It is imperative that decision-makers possess a thorough understanding of the operational dynamics of AI systems. This may necessitate closer collaboration with technical experts and the development of updated examination guidelines that accurately reflect the realities of AI advancement. Absent this understanding, there exists a risk that legal standards will be applied in a manner that fails to align with the technology.

Moreover, a more profound inquiry warrants consideration should all categories of AI inventions be patentable? If certain systems cannot be meaningfully disclosed, it may be prudent to evaluate whether patents represent the most appropriate form of protection for them. This does not suggest that such inventions lack significance; rather, it implies that alternative regulatory frameworks may be more suitable. In sectors where AI exerts considerable societal influence, such as healthcare or decision-making systems, regulations pertaining to transparency, accountability, and fairness may take precedence over exclusive rights.

Lastly, given the global nature of AI development, these issues are not confined to a single jurisdiction. Divergent approaches to disclosure and enablement across different countries can engender confusion and uncertainty for innovators. A degree of international coordination would contribute to the establishment of a more predictable system. However, any harmonization efforts must consider that various countries may possess differing priorities and levels of technological advancement.

Ultimately, the challenges posed by black-box AI systems underscore a critical aspect of the interplay between law and technology: the law must not remain static while the technology it

regulates continues to evolve. The enablement requirement has long been a fundamental principle of patent law; however, it cannot afford to be inflexible if it is to maintain its relevance. Its future viability hinges on its ability to adapt to emerging forms of innovation while preserving its essential function of ensuring meaningful disclosure.[33]

This necessitates a shift away from excessively rigid, mechanical interpretations of the rule, favoring a more flexible, principle-based approach that reflects the operational realities of AI systems. Concurrently, such flexibility must not undermine the protective measures that prevent vague or overly broad patents. Achieving this equilibrium is of paramount importance.

More broadly, the issues delineated in this paper indicate that AI represents not merely another technological advancement; it fundamentally challenges some of the foundational assumptions underlying patent law. Consequently, superficial or incremental modifications to existing regulations are likely insufficient. A comprehensive reevaluation of how concepts such as disclosure and enablement should function in this new paradigm is essential.

If approached with due diligence, this challenge also represents an opportunity. By adapting to the realities of AI, the patent system can become more responsive and better equipped to address future innovations. In this regard, the tension between black-box AI and patent disclosure should not be viewed solely as a problem to be resolved, but rather as a pivotal moment that could influence the evolution of intellectual property law moving forward.

[33] Lawrence Lessig, Code and Other Laws of Cyberspace 6–8 (1999).