PUBLISH, PATENT, OR PERISH: A STRATEGIC FRAMEWORK FOR UNIVERSITY INTELLECTUAL PROPERTY IN THE GENERATIVE AI REVOLUTION

By

Shriya Shah

An Undergraduate Thesis submitted as part of the

WHARTON RESEARCH SCHOLARS

Faculty Advisor:

Steven Weiner

Senior Director, Innovation & Technology – Penn Engineering, Adjunct Lecturer Management Department

THE WHARTON SCHOOL, UNIVERSITY OF PENNSYLVANIA

APRIL 2025

ABSTRACT

Universities are at a crossroads as generative AI (GenAI) transforms the landscape of academic innovation and competition. This research addresses the strategic dilemma facing U.S. universities: how to protect and maximize the value of AI-assisted discoveries while advancing their core missions. Through a qualitative methodology, combining literature review, expert interviews, and scenario-based analysis, the study examines both patenting and non-patenting pathways for GenAI outputs. Findings reveal that while patents can offer exclusivity and commercial potential, universities often benefit from integrating alternative strategies such as industry partnerships, capacity building, and data deployment. The proposed decision framework enables institutions to align intellectual property management with evolving technological, regulatory, and mission-driven priorities. Ultimately, this work provides actionable guidance for universities seeking to navigate the complexities of GenAI innovation, ensuring that academic research remains a driving force in technological progress.

Keywords: Generative AI, university innovation, intellectual property, patent strategy

INTRODUCTION

Innovation has long been recognized as a crucial driver in economic growth and prosperity for nations. Disruptive technologies, dating back to the industrial revolution, have played an instrumental role in shaping the trajectory of developed, knowledge-based economies. Numerous studies have highlighted the indirect impacts on productivity growth, capital utilization, and income per worker across countries. Research by the U.K.'s National Endowment for Science, Technology and the Arts attributed two-thirds of the private-sector productivity growth in the U.K. between 2000 and 2007 to innovation (NESTA 2009). Similarly, Klenow and Rodriguez-Clare's cross-country study found that over 90% of the variation in the growth of income per worker resulted from innovations that changed how capital was utilized (Klenow and Rodriguez-Clare 1997). These findings underscore the profound significance of innovation in economic expansion, crafting employment opportunities, and fostering a virtuous cycle of growth.

However, innovation is rarely driven without the appropriate incentives, thus its history is inextricably linked to the development of intellectual property (IP) protection. The first and second industrial revolutions witnessed sharp increases in patenting activities as countries became more inventive, with global patent systems evolving to follow the archetypal United States model (Billington, Hanna, and Lane 2024). The codification in the Statute of Anne during the eighteenth century marked a pivotal moment, establishing IP rights for intangible property and reshaping the trajectory of global growth (Norman n.d.). The power of IP enables individuals to reap the benefits of their creativity through three primary mechanisms: patents, which provide exclusive rights for inventions; copyright, which predominantly serves literary and artistic industries; and trade secrets, which protect confidential information from unauthorized use.

However, the rapid pace of technological advancement has frequently created tensions in our legal and regulatory frameworks. Nowhere is this more evident than in the realm of artificial intelligence (AI), a field that has witnessed extraordinary growth since its conceptualization in the 1950s. By 2016, the AI domain had generated 1.6 million scientific publications and nearly 240,000 patent applications. Within this sphere, machine learning—specifically deep learning and neural networks—have driven revolutionary advancements in generative AI, with deep learning patent mentions growing annually by 175% from 2013 to 2016, while neural networks experienced 46% annual growth (WIPO 2010). Yet the substantial increase in AI patents reveals a more nuanced disconnect than simply legal frameworks failing to keep pace with innovation. Rather it highlights significant imbalances in who leads this technological revolution: universities have increasingly trailed behind their industry counterparts, while the United States faces growing competition from China's strategic investments in AI research and development. These asymmetries make it all the more crucial to examine how the traditional symbiotic relationship between academic research and commercial application has evolved, particularly with the introduction of Generative AI.

The emergence of Generative AI, exemplified by Google's Bard and OpenAI's ChatGPT, has ushered in a transformative era of innovation. Bloomberg projections indicate that the generative AI market will reach $1.3 trillion over the next decade, achieving a compound annual growth rate of 42% (Castaros and Bloomberg Intelligence 2023). This explosive growth is driven by demand for training infrastructure, inference devices for large language models, digital advertising, specialized software, and services. However, this has marked a significant shift in the traditional innovation landscape - a shift from ivory towers to executive powers.

Historically, universities, especially in the United States, have been the powerhouses of technological innovation, serving as the primary catalysts for breakthrough discoveries and fundamental research that shapes entire industries. Their unique ecosystem of diverse expertise, academic freedom, and long-term research horizons has traditionally positioned them at the forefront of transformative technologies, adhering to the “publish or perish” mentality while also seeking patents for inventions with valuable commercial potential (Singh n.d.). This is corroborated by empirical evidence that companies spun out by research universities have significantly higher success rates than other companies, with university research spin-offs generating more patent applications and radical product innovations. The United States has been a beneficiary to this, with industry giants such as Google, Medtronic and iRobot owing their origins to universities. However, the generative AI revolution has disrupted this established paradigm. As noted by industry experts, "For the first time, universities are behind" in research and development, with firms taking the lead in pushing technological boundaries. This unprecedented reversal of traditional innovation dynamics demands urgent attention, as the potential consequences of universities losing their position as primary innovators could significantly impact long-term technological progress, retention of top talent and overall economic growth.

This risk is already evident in the global patent landscape. China has filed 7,582 generative AI patent applications between January 2016 and February 2023, far surpassing the United States' 2,391 applications. Most tellingly, seven Chinese universities rank among the worldwide top 10 filers for generative AI technology, while only one U.S. university (Northeastern) has achieved this distinction (Sarkar 2023). This disparity has significant real-world implications for global

economies, as seen in the recent launch of DeepSeek's R1 "reasoning" model, which triggered a stock market sell-off with ripple effects from San Francisco to Wall Street (The Outpost 2024).

To address this widening innovation gap and develop effective strategies for university engagement in generative AI development, it is essential to first understand the fundamental structure of the innovation process itself. Through a systematic examination, we can identify specific intervention points where universities might maximize their competitive advantage while navigating the complex landscape of patent eligibility.

The generative AI innovation process can be broadly divided into three stages (as shown in Figure 1): 1) training the model and generating the data, 2) the model itself with its specific architecture, parameters, and configurations, and 3) the outputs or creations generated by the model - of which the third is most interesting for the purposes of this study.

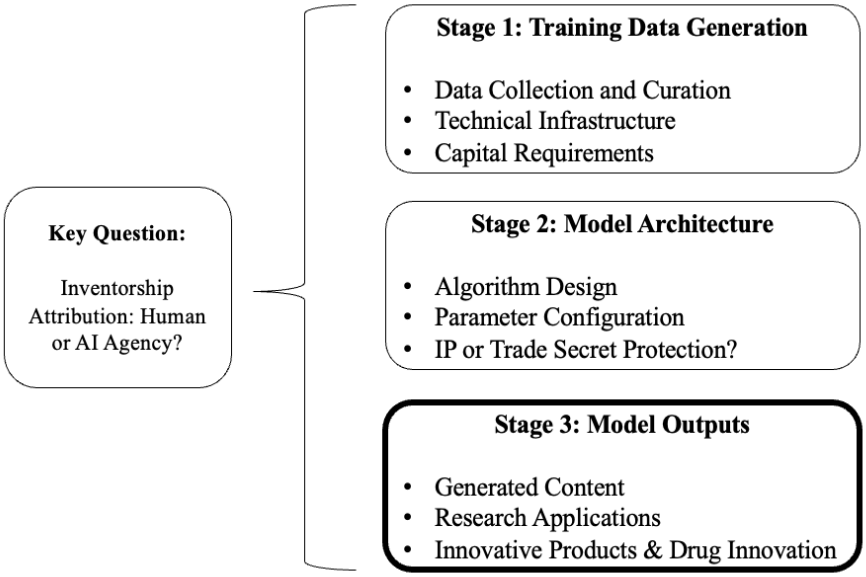

Figure 1; Delineation & Key Considerations of the Generative AI Innovation Process

The inherent nature of the training phase presents barriers primarily characterized by resource intensity rather than novel intellectual property. Current estimates indicate that training a large language model (LLM) can exceed $4 million in direct costs – exemplified by Meta's LLaMa model, which required over 1 million GPU hours at an approximate cost of $2.4 million. OpenAI's reported monthly expenditure of $40 million for ChatGPT prompt processing further illustrates that the competitive advantage at this stage derives predominantly from capital deployment rather than patentable innovation (Vanian 2023). These substantial capital requirements naturally restrict market entry, creating a de facto protective mechanism that diminishes the additional strategic value of patent protection.

The second stage, focusing on model architectures, encounters distinct epistemological and practical barriers to effective patent protection. Contemporary model architectures predominantly build upon an established corpus of academic literature and open-source implementations, creating significant hurdles in satisfying the fundamental patent requirements of novelty and non-obviousness. The field's historical commitment to open science, manifested through extensive documentation in academic publications and collaborative development in open-source repositories, has created a rich body of prior art. Moreover, the generative AI community's deeply embedded culture of knowledge sharing—while instrumental in advancing the field—has established a precedent that complicates attempts to assert exclusive rights over architectural innovations.

In contrast, the third stage—centered on model outputs and their applications—emerges as the most promising domain for university patent strategies, though it requires navigating evolving legal standards. This stage presents unique opportunities precisely because it intersects with universities' traditional strengths in domain-specific research and innovation. Unlike the

resource-intensive training phase or the architecturally constrained model development stage, the output phase allows universities to leverage their diverse expertise in fields such as product design, chip design, drug discovery, and material science. The distinctive nature of generative AI outputs, which transcend traditional numerical processing to create novel content—spanning from text and images to molecular structures and code—introduces new paradigms for innovation that align with academic research capabilities and its collaborative nature. Nevertheless, the path from generative AI output to commercially viable intellectual property is neither immediate nor straightforward. The timeline for realizing tangible results in this stage, such as license agreements, startup formation, or even patent filings, tends to be significantly slower than in earlier stages of AI innovation. This lag is due in part to the need for rigorous validation, refinement, and the integration of generative AI technology into practical applications, before the work of using such technology to actually generate and refine potentially valuable output can begin – processes that often require years of additional R&D. Moreover, we may have difficulty identifying whether a given invention was derived from GenAI output, unless the owner voluntarily discloses that fact. Such disclosure may be less likely to the extent that intellectual property law necessitates emphasizing human contribution to conception and development, and deemphasizing the role played by the AI system itself.

Despite this, the significance of a third stage-centered investigation becomes even more pronounced considering the Federal Circuit's recent decision in Recentive Analytics v. Fox Corp.[1] The decision was explicit in drawing a crucial distinction: it reaffirmed that patent eligibility does not extend to claims that simply use established machine learning methods for

[1] Recentive Analytics, Inc. v. Fox Corp., No. 23-2437 (Fed. Cir. Apr. 18, 2025), https://www.cafc.uscourts.gov/opinions-orders/23-2437.OPINION.4-18-2025_2500790.pdf.

new tasks or fields, absent any concrete improvement to the underlying technology (Winston & Strawn LLP 2025). However, the opinion also recognized that ML and AI remain "burgeoning and increasingly important" fields, and crucially, that patent-eligible improvements are still possible, provided the claimed inventions move beyond abstract applications, what this paper would classify as Stage 2 patents, and demonstrate specific contributions. Thus, this legal clarification directly supports a continued focus on stage 3 for university patent strategies. The output phase aligns with the court's guidance by offering fertile ground for genuine invention. Here, universities can leverage their domain expertise to create and validate new applications that embody inventive steps beyond the mere use of AI as a tool, satisfying the heightened standards articulated in Fox Corp.

Considering the aforementioned analysis, this research takes an institutional perspective, exploring how United States universities can harness generative AI patenting strategies to establish competitive advantages. By developing a novel analytical & decision framework, this research aims to identify and evaluate the critical elements to consider for universities to successfully protect and integrate GenAI into their research enterprises. The findings hold significant potential to inform policy decisions, shape institutional strategies, and foster an environment conducive to university-driven innovation in the generative AI era.

LITERATURE REVIEW

Now that we have narrowed down the scope of the research to patent strategies, centering on outputs and United States universities it's important to delve into the existing literature within this domain, which will show us the gaps that need further exploration - gaps that I hope to address.

Intellectual Property Theoretical Frameworks

The traditional paradigm of patent strategy and patent rights has undergone an evolutionary process, serving multifaceted and interrelated objectives – cultivating incentives for innovation by conferring exclusivity over inventions, enabling technological transfer and commercialization through licensing of protected intellectual property (IP), and signaling the underlying value of companies to attract investment capital (WIPO 2017). This section investigates the foundational theoretical underpinnings that have shaped the contemporary landscape of patent strategy, with the aim of ascertaining whether these established notions remain applicable within the burgeoning realm of Generative AI.

A confluence of key frameworks has emerged, empowering organizations to strategically leverage patents. These encompass the utilization of patents as "isolating mechanisms" to safeguard valuable innovations from imitation, thereby preserving competitive advantages or economic rents derived from the commercialization of the IP (Kennedy et al. 2020). Furthermore, patents can be viewed through the lens of maximizing licensing revenues by monetizing protected inventions. Finally, given the intrinsically crowded nature of modern technological landscapes, characterized by overlapping patent rights, companies must often

devise "defensive" strategies involving the development of their own patent portfolios, a necessary measure to maintain freedom to operate and circumvent potential litigation (Somaya 2002).

Historically, the rationale of "patent rights as incentives to innovate" has been a guiding principle (Lerman 2015). Fundamentally, patents derive their value from the exclusivity they confer upon the patent holder, enabling them to commercially make, use, sell, or import the patented invention for a finite period. As a guiding tenet, "a patent system benefits innovators, and they, in turn, offer their new developments to society" (WIPO 2017). However, further research has unveiled that with the evolution of the secondary market, the "contextual value of patents" and signaling capabilities have emerged as significant drivers influencing bid-offer spreads (IAM 2018). Consequently, "patent arms races" have materialized, characterized by the defensive stockpiling of even low-quality, vague patents, employed exploitatively in litigation (Chien 2010). These non-legal developments underscore that the dynamics of market operations are no longer solely governed by the purview of Congress, Courts, and the USPTO, but rather, patentee behavior plays a critical role.

In the context of university patenting, this concept assumes particular relevance, as for universities and their faculty members, "publishing" or disseminating research findings is a pivotal consideration for all stakeholders involved. Thus, a pertinent inquiry arises: for universities specifically, will the intrinsic value of patents transcend mere licensing considerations, rendering the contextual value (signaling) of even greater significance?

Artificial Intelligence & Generative AI Theoretical Frameworks and Landscape

Turning our attention to the realms of artificial intelligence and generative AI, the prevailing landscape appears to offer limited opportunities for sustainable differentiation through proprietary knowledge or the cultivation of loyal customer bases. This assertion is attributed to the emergent nature of the technology and the multiplicity of players involved. Bob Steinberg, a retired partner of Latham & Watkins and former Global Co-Chair of the Intellectual Property Litigation Practice, characterizes the current stage of the artificial intelligence industry as a first-mover stage, devoid of a dominant provider owning an entire vertical, with the exception of Nvidia in the GPU domain. Consequently, access to scarce data assets emerges as a critical second source of advantage for developing high-performing generative systems.

This revelation raises a thought-provoking inquiry: specifically, given our previous discussions illuminating the substantial costs associated with acquiring adequate data for a model and sustaining its ongoing training, what ramifications arise if institutions like universities lack the requisite funding to undertake such endeavors?

This predicament could potentially engender the bifurcation of two distinct realms: one in which large firms like Microsoft and Google focus their competitive advantage on the data upon which their models are trained, leveraging the efficacy, validity, and sheer volume of their data assets. On the other hand, when considering institutions such as universities or smaller firms, a secondary approach has emerged based on initial discussions within this study – the concept of "prompt engineering" and its critical significance, wherein resources within these institutions will be directed towards this endeavor. According to insights from Marc Sedam, the VP of

Technology, Operations & Ventures at NYU at the AI Series Part 1 by Osage University Partners, prompt engineering represents the research and development required to refine how we pose questions to these models, thereby eliciting more accurate and insightful responses. Notably, at NYU Langone Medical School, it took a concerted effort spanning 30-40 hours to craft a singular prompt that would be effectively perfect.

While adjustments were made to elements such as the model's temperature, which governs its creativity, the development of this "defensive arsenal" of unique prompts tailored to a firm's specific practices and objectives has emerged as an alternative pathway to cultivating a competitive advantage. Whether this approach is scalable, efficient, or feasible within the constraints of other university frameworks remains a question that will be addressed within the scope of this study.

Furthermore, as we revisit the question of patenting through the lens of a hypothetical analogy involving hardware and software, it can be argued – as posited by Robert Plotkin, one of the preeminent AI patent attorneys in Massachusetts – that the models themselves can be viewed as the hardware upon which the prompts (software) operate. This raises the intriguing proposition: why can the prompts themselves not be patented in a manner akin to software? This inquiry represents another avenue that will be further explored throughout the duration of this study.

Underlying all these considerations, we must return to the fundamental rationale that initially prompted this research question. The distinguishing characteristic of Generative AI lies in its unique nature, wherein, in both discriminative and early-stage generative models, while the

training process is automated, human judgment remains crucial for tasks such as architectural design and defining success metrics. However, as Smits & Borghuis posit, once the training model is established, the absence of direct human intervention in creating outputs raises challenges aligned with intellectual property policies aimed at safeguarding works traditionally attributed to human authors (Smits and Borghuis 2022). Furthermore, in instances where no authorship is established, the work falls into the public domain, which Cockburn argues presents a considerable disadvantage as it eliminates tangible incentives for developers to continue advancing the underlying generative systems (Cockburn, Henderson, and Stern 2018).

As a consequence, granting protections also proves complex given that under existing US patent law, qualifying subject matter must meet thresholds related to human ingenuity and inventorship. As explained through USPTO guidance, the current statute interprets inventors as natural person individuals. Though AI-assisted inventions are not categorically unpatentable, analysis should focus on human contributions, as patents intend to reward human creativity. Reinforcing this, the Thaler v. Vidal precedent solidifies that US law requires inventorship to reside with a natural person, not AI systems themselves: "The Patent Act expressly provides that inventors are 'individuals'...the statute does not define 'individual'. However, as the Supreme Court has explained, when used “[a]s a noun, 'individual' ordinarily means a human being, a person" (Thaler v. Vidal 2022). These perspectives make virtually no allowance for directly acknowledging generative models' autonomous outputs through patent protections, despite the levels of creative problem-solving they exhibit. However, an interesting facet is presented when a USPTO guidance provided some allowance for AI-assisted inventions which stated that

“inventions are not categorically unpatentable due to improper inventorship if one or more natural persons significantly contributed to the invention” (USPTO 2024).

In the university context, the patent ownership nuances encompass critical considerations regarding assignment rights and tech transfer mechanisms. Traditional assignment policies mandate that faculty members transfer rights of inventions developed using university resources to their institutions, which raises an interesting fundamental question about attribution determination, and if one can transfer the rights to an invention that they independently did not ideate. Furthermore, a critical inquiry emerges regarding the “significant contribution” threshold: how can universities demonstrate sufficient human intellectual input to satisfy patent eligibility requirements? This question becomes particularly salient when considering that innovations failing to meet these criteria lack protection mechanisms, potentially leaving substantial institutional investments in advanced generative systems vulnerable to imitation or infringement.

The Role of University Patenting in the Generative AI Ecosystem

Examining the integral role that universities play within the generative AI ecosystem unveils a multifaceted tapestry of characteristics, drivers, and implications. Universities are quintessential knowledge houses, serving as incubators for students, professors, and researchers. Historically, these institutions have adhered to the "publish or perish" mentality as discussed earlier, prioritizing the dissemination of research findings. However, in recent years, a pivotal shift has emerged, with universities recognizing the tangible benefits of pursuing patenting strategies. These benefits encompass financial independence and profitability, bolstering institutional rankings, and cultivating an ecosystem conducive to nurturing knowledge-based startups.

This transformation in mindset is further amplified by the global paradigm shift regarding the role universities will assume in AI patenting, as exemplified by the dominance of Chinese universities in this arena. These institutions are intricately intertwined with governmental interests, reflecting the strategic imperatives of their respective nations.

The Bayh-Dole Act, a landmark legislation, has been a catalyst for this transformation. It empowers universities to retain patents for inventions arising from government-funded research, provided they satisfy certain social and safety criteria. Moreover, it enables them to license their innovations to private sector partners, facilitating technology transfer and commercialization (Drexel University n.d.). The impact of this act has been profound, leading to a surge in patenting and licensing activities by universities, as well as the proliferation of startup companies dedicated to developing and commercializing federally funded technologies.

Prior to 1980, the patent rights to nearly all discoveries made with federal aid belonged to the government, resulting in a stagnant pool of untapped potential, with fewer than 1,400 of the 28,000 patents held by federal agencies being licensed, and none encompassing new drugs (Fraser 2024). However, the Bayh-Dole Act has catalyzed a remarkable transformation, contributing over $1 trillion to the US GDP between 1996 and 2020. Notably, the commercialization rates of federally funded research have increased tenfold since the passage of the act between 1980 and 2009 (Ezell 2019).

The technology transfer facilitated by academia-private sector patent licensing across all industries has been a driving force behind this economic impact, bolstering the US GDP by up to

$865 billion and supporting up to 5.9 million person-years of US employment between 1996 and 2017 (Pressman et al. 2015). However, a disparity persists between the level of funding provided by the government for federally funded innovations and the relative position of the USA in terms of business funding of university research and development (R&D) as a share of GDP. In 2018, the USA ranked a mere 24th out of 36 nations, falling short of its OECD peers, despite boasting world-leading research universities that have been pivotal in driving American technological supremacy since World War II (Atkinson and Gawora 2021).

To regain its competitive edge, research indicates that the US would need to invest an additional $17 billion per year to ascend to 15th place in government funding of university research, $31 billion per year to reach 10th place, $51 billion to secure 5th place, or a staggering $90 billion to match Norway, the current leader (Fraser 2024). The advent of generative AI presents a compelling opportunity for the United States to reclaim its competitive advantage in this domain.

However, to justify such substantial investments, it is imperative that universities are capable of patenting their inventions that derive their core value proposition from generative AI. A recent proposal by the Biden administration has introduced a potential hurdle, as it would allow the government to claim patent rights on federally funded, private-sector inventions deemed unreasonably priced. While the Bayh-Dole Act previously permitted the government to "march in" and relicense patents to third parties under well-defined, narrow circumstances typically focused on social or safety concerns, this new provision could deter startups and investors from partnering with universities, knowing that the government could swoop in and claim rights at any point during the development process (Thaler v. Vidal).

In this context, the significance of universities successfully demonstrating the value in licensing out patents for generative AI innovations becomes integral. Such a demonstration is crucial to assuage the concerns of startups and firms, ensuring their willingness to overlook the potential risk posed by the government's ability to claim rights. Achieving this objective hinges on universities' ability to secure the appropriate patenting strategies, a challenge that lies at the heart of this research endeavor.

Moreover, universities present a unique advantage in the generative AI landscape, as they foster a diversity of thought and applications that transcend the boundaries of any single industry. From Stanford University researchers developing generative AI for the creation of potential new drugs to combat antibiotic-resistant bacteria, to the Penn Research in Embedded Computing and Integrated Systems Engineering Center exploring the use of AI in medical practice the breadth of applications is vast (Tompa 2024; Singhi 2023). At MIT's Center for Art, Science and Technology, researchers have harnessed AI to create a non-linear documentary based on video material from six different monuments, effectively fabricating a synthetic memory (Volpe and Yuce 2023).

In a similar vein, Oxford University's Edward Brooks and Corey Crossan are pioneering the development of "ManagementGPTs," envisioned as co-thinkers, well-versed in a variety of management practices, supporting managers across levels, power structures, educational backgrounds, and geographic locations (Brooks and Crossan 2023). Simultaneously, we are witnessing university inventions that aim to counteract the use of AI, a development of Stage 1

inventions, such as the University of Chicago's Professor Ben Zhao and his team's creation of Glaze, a product designed to protect artwork from being mimicked by AI. They are now working on Nightshade, a tool that attacks generative AI by altering the image data it receives (Chen 2024).

Lastly, this study also hypothesizes a potential deviation from traditional strategies towards small language models or private large language models could offer a viable patenting pathway for universities in the generative AI landscape. In contrast to public LLMs which inherently lack data privacy and carry the risk of all inputs being used for further model training, private LLMs built on foundational models like GPT, OpenAI, or Azure can ensure data confidentiality and enable technology transfer offices to quantify returns on investment. Furthermore, SLMs, with their relatively smaller model sizes and domain-specific training data, may prove advantageous for universities seeking to excel in function-specific use cases that require highly specialized and proprietary knowledge (Raza 2024). This study aims to evaluate the merits and feasibility of this hypothesized deviation from traditional strategies as a means for universities to secure robust generative AI patenting strategies.

The multifaceted role of universities in the generative AI ecosystem, the challenges and opportunities surrounding patenting strategies, and the diverging pathways of LLMs and SLMs collectively underscore the pressing need for a comprehensive framework to guide universities in harnessing the potential of this disruptive technology.

METHODOLOGY

In pursuit of a comprehensive understanding, this study employed a synergistic blend of qualitative data gathering and analysis methods. Given the nascent and rapidly evolving nature of the technological domain, the design prioritized both flexibility and depth in data collection while maintaining methodological rigor throughout the process.

The foundational phase of the research involved a comprehensive analysis of the existing knowledge landscape as detailed in the aforementioned literature review. This comprised examining fragmented academic literature, legal analyses, policy documents, and industry reports pertaining to generative AI patenting approaches and considerations. The literature review proved to be particularly crucial given the emergent nature of the field, as it enabled the synthesis of disparate perspectives that currently inform university strategies. Special attention was devoted to guidance from authoritative bodies, notably the United States Patent and Trademark Office (USPTO), to establish a robust understanding of the current regulatory environment. The primary data collection phase centered on expert interviews, which were conducted through various modalities to accommodate participant availability and geographical distribution. The study engaged with subject matter experts, strategically selected to represent diverse perspectives across the innovation and intellectual property landscape. The participant pool included Chief Innovation Officers, and heads of technology transfer offices from leading academic institutions, providing crucial insights into institutional approaches and ongoing conversations regarding GenAI patent strategies. Legal expertise was garnered through interviews with specialized intellectual property attorneys who actively engage with emerging AI

technologies. The research also benefited from conversations with PhD economists and prior expert witnesses who have participated in patent litigation, offering perspectives on the practical implication and potential judicial outcomes of AI patents. To complement these perspectives, the study also incorporated insights from individuals at the forefront of research and development, who rely upon generative AI foundations.

The interview process employed a semi-structured approach, allowing for organic exploration of themes while attempting to maintain consistency across conversations. They were conducted primarily through remote video conferences due to geographical constraints, with supplementary email correspondence when necessary to clarify or expand upon prior discussions. The research process also benefited from informal conversations, following conferences, speaker series and academic lecturers. Data collection methods were carefully documented as well, with interview sessions being recorded with participant consent and subsequently transcribed using artificial intelligence tools, with manual verification performed to ensure accuracy. Detailed written notes were maintained for all interactions, including informal conversations and email exchanges in order to effectively capture the nuanced perspectives and insights that emerged during these discussions.

A notable strength of the research design lay in its incorporation of hypothetical test cases, particularly during interviews with legal experts and expert witnesses where the scenarios served as analytical tools to examine how existing and proposed patenting frameworks might be interpreted and applied in various contexts. This approach proved to be highly valuable in

understanding the uncertainties and implications of current patenting approaches applied to the generative AI domain.

It is important to acknowledge that the study's methodological framework lent itself to certain limitations and contextual factors that shaped the research process. While the sample demonstrated substantial representation from the University of Pennsylvania ecosystem, this was mitigated through deliberate efforts to incorporate diverse perspectives from experts across different organizations and geographic locations to mitigate any institutional homogeneity. Furthermore, temporal and resource constraints inherent to thesis research inevitably influenced the scope of the primary data collection. Although a larger sample of formal interviews would have potentially yielded additional insights, the temporal boundaries of the academic program necessitated a carefully bounded design. These constraints, while noteworthy, did not fundamentally compromise the study's analytical rigor or theoretical contributions.

Overall, this qualitative approach enabled deep exploration of complex themes surrounding generative AI patenting strategies while acknowledging the field's dynamic and ever-changing nature. Through methodological triangulation—synthesizing insights from literature review, expert interviews, and hypothetical test cases—the research provided a nuanced understanding of how organizations approach generative AI patenting as a competitive advantage strategy. The methodology's flexibility allowed for adaptation to emerging insights while maintaining scholarly rigor, though it is important to note that findings should be interpreted within the context of the field's rapid evolution and the study's qualitative nature.

ANALYSIS

Drawing from the aforementioned methodology, the culmination of this study is an analytical framework (Figure 2) designed to navigate the complex landscape of AI innovation protection within academic institutions. This framework, visualized as a hierarchical decision structure, aims to serve as a comprehensive analytical tool for university technology transfer offices, researchers envisioning the future of their innovations, and institutional policy makers, among others. The significance of this framework lies in its ability to address a critical gap in existing literature and practice: the absence of a structured approach that simultaneously considers institutional priorities, technical innovation assessment, and contextual implications specific to AI technologies. Importantly, the framework acknowledges that patent protection may not always represent the optimal strategy, incorporating detailed consideration of alternative protection mechanisms that might better serve institutional and societal interests. The following sections provide a detailed examination of each component within this framework, beginning with institutional strategic evaluation and proceeding through innovation assessment, protection strategy determination, and contextual considerations. Finally, this paper presents a detailed examination of Synthemol as a case study, demonstrating how the theoretical framework developed herein manifests in practical application, thereby bridging conceptual analysis with real-world implementation.

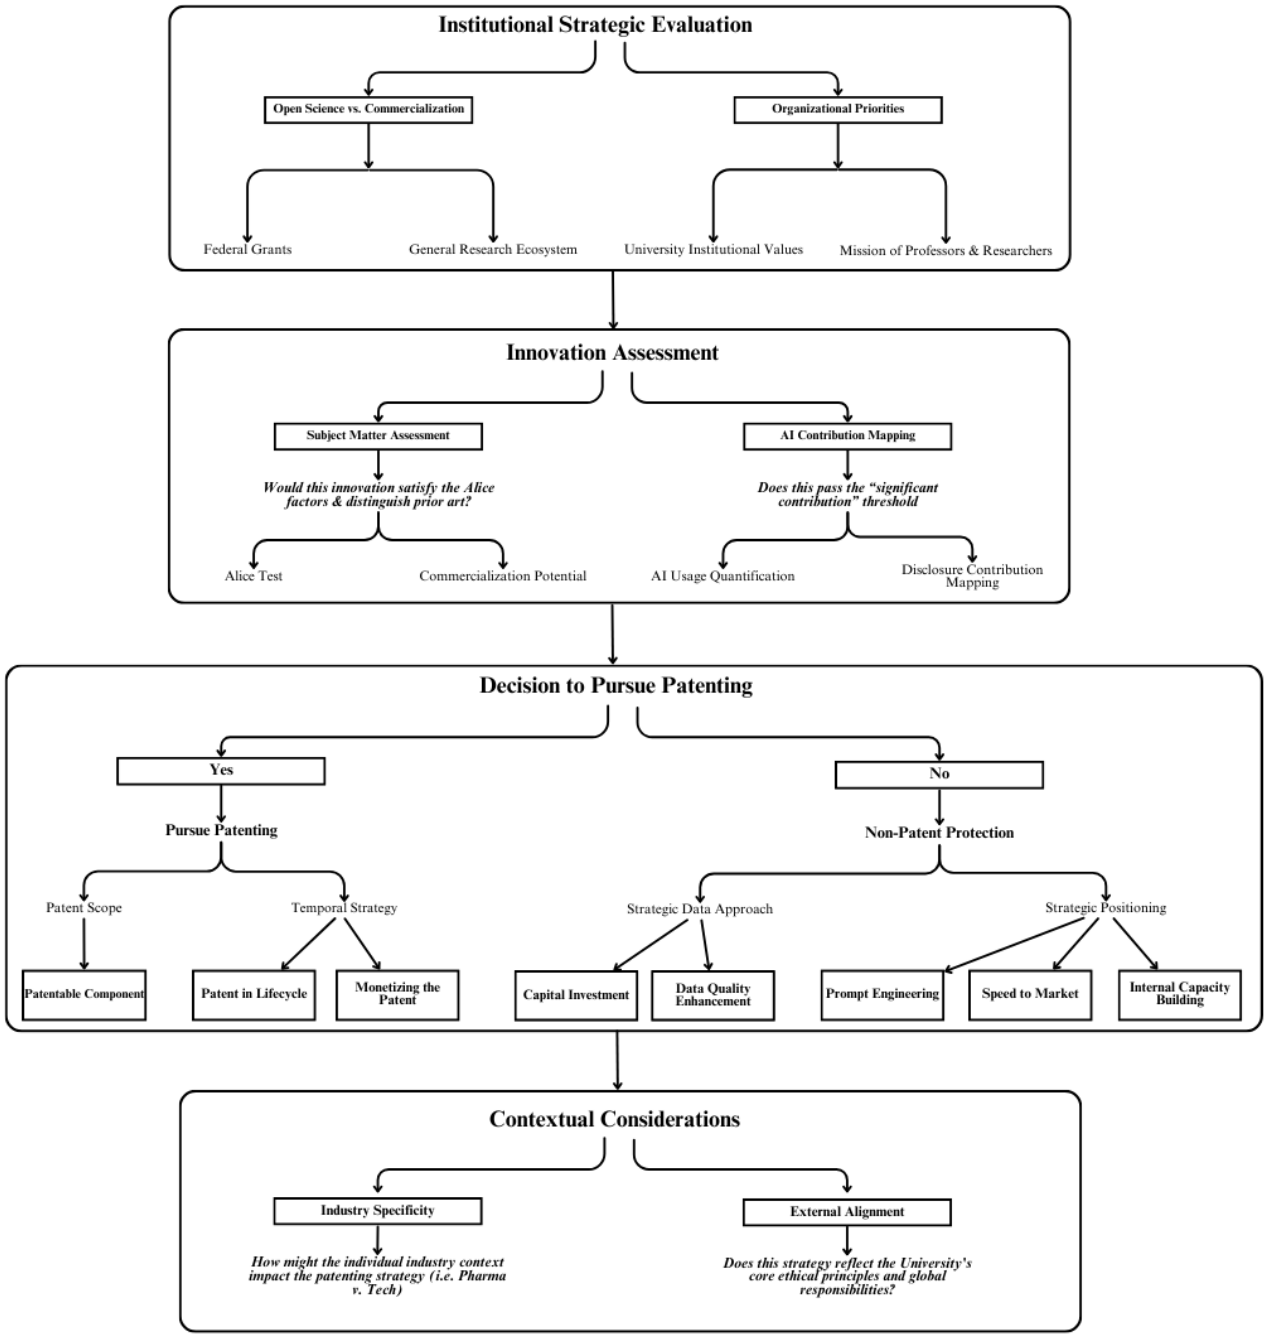

Figure 2: Decision Analysis and Considerations for Generative AI Patenting Strategies

INSTITUTIONAL STRATEGIC EVALUATION

The primacy of institutional strategic evaluation within the proposed framework reflects a foundational understanding that university intellectual property decisions do not occur in isolation, but rather within distinct and at times, bureaucratic organizational systems shaped by institutional values, policies and constraints. Therefore, while the findings pertaining to technical assessments of patentability or market potential will be largely applicable to inventions utilizing GenAI, this section situates the protection strategies within a broader context that will ultimately govern their implementation and impact. This section finds that when considering GenAI innovations, there exists a conflict between supporting the rapid pace of technological advancement and the role of academic institutions in knowledge dissemination. Moreover, the heterogeneity of universities themselves plays a key role in the decision-making process.

Open Science v. Commercialization

Universities occupy a unique position within the innovation ecosystem, one characterized by a fundamental tension between two seemingly contradictory imperatives. This dual mandate, as articulated by Chris Callison-Burch, is one of "balancing the pursuit of open science with a need for commercialization and revenue generation from innovation".[2] This tension is even reflected in international normative frameworks, as evidenced by its codification in Article 27 of the Universal Declaration of Human Rights, Article 15 of the ICESCR, and Article 15 of the UNESCO Universal Declaration on Bioethics and Human Rights, all of which explicitly articulate the obligation to disseminate scientific knowledge broadly while ensuring equitable

[2] Chris Callison-Burch, conversation with the author, Philadelphia, November 14, 2024.

distribution of its benefits (Chapman 2011). The complexity of this balancing act has intensified contemporaneously with the escalating costs of cutting-edge research and increasingly fragmented institutional structures. This results in contradictory external pressures for researchers: “on the one hand, they are told to commercialize their research by patenting, licensing, and forming close partnerships with industry…[as] it is presumed that this will generate maximum wealth…On the other hand, researchers are encouraged to share data and disseminate knowledge…so as to foster scientific progress (Caufield, Harmon, and Joly 2012)”. Caufield, Harmon, and Joly highlight that there has also been a significant historical transformation in the conceptualization of the university’s societal role, evolving from primarily “knowledge generators” to active “economic developers”. This paradigmatic shift received explicit political endorsement through President Obama’s 2011 State of the Union address, which introduced the “Startup America” initiative and called upon academic research communities to function as economic catalysts (White House Archives 2011). However, recent policy developments have further complicated this landscape, particularly through the proposed reductions in medical research funding, which have precipitated systemic repercussions throughout higher education facilities. A particularly significant policy change involves the National Institutes of Health (NIH) implementing a 15% cap on Facilities & Administration (F&A) rates for university research grants. The implications of this policy are substantial: at the University of Pennsylvania alone, the announced measures would reduce annual federal funding by approximately $240 million, potentially reaching $315 million if adopted across all federal agencies (Jameson 2025). This reduction in federal support—historically a cornerstone of university research funding—will likely intensify pressure toward alternative revenue sources,

with commercialization activities emerging as a primary avenue, particularly those with significant revenue potential such as Generative AI applications.

Policy Manifestations of the Dual Mandate:

The dual mandate finds concrete expression in university policies and faculty handbooks across higher education:

Michigan Technology University's Faculty Handbook, in its "Tenure, Promotion, and Reappointment" section, explicitly states: "Innovation and commercialization are consistent with and directly support the education, research and public service goals of Michigan Tech. Innovation refers to the process of creating economic and social value from inventions and discoveries. Commercialization supports innovation and results in the creation of societal and economic value from knowledge, discoveries, and experiences." The document further emphasizes that "promotion and tenure and merit evaluations shall recognize innovation and commercialization as an indicator of commitment and expanding traditional boundaries and to the economic prosperity of society."

Similarly, the University of California's Policy on Innovation Transfer & Entrepreneurship aims to "promote the translation of UC's discoveries into useful products, services, and innovations that not only provide value to individuals and society, but also endeavor to uplift the human condition," while simultaneously seeking to "pursue fair value for our intellectual property so UC can continue to grow its excellence in scholarship, research, and global impact."

Lastly, Duke University's policy on Innovation and Intellectual property further demonstrates this bifurcated approach: "Duke University is committed to disseminating the fruits of its research and scholarship as widely as possible. In keeping with this commitment, Duke adopts the following policy: Duke University and its faculty, staff, and students are committed to ensuring the greatest possible scholarly and public access to the scholarly products that they create." Concurrently, the policy establishes that "Duke University has established and maintains a policy on Intellectual Property Rights, which includes the definition and general principles of ownership," accompanied by detailed financial distributions from commercialization revenues." These institutional policies illustrate how universities are deliberately constructing governance frameworks that accommodate both open science imperatives and commercialization objectives, attempting to resolve the inherent tension between the competing mandates.

Navigating the Balance: Licensing as a Strategic Tool

The dual mandate takes on heightened significance in Generative AI patenting strategies, where the industry's nascent state has simultaneously amplified open science imperatives—evidenced by the formation of entities like Anthropic and OpenAI with explicit commitments to industry development—and intensified commercialization pressures, as securing exclusivity becomes a key differentiator between competing models. Herb Hovenkamp, an American legal scholar and professor at the University of Pennsylvania Carey Law School notes that when looking at companies, "a firm that develops a superior product does not have a duty to share or license their assets to rivals."[3] This industry principle is directly contrasted in the world of universities, with Marc Sedam articulating a fundamental perspective on this tension: "the business model of the

[3] Herb Hovenkamp, Current and Emerging Issues in Law and Technology panel discussion at the University of Pennsylvania Carey Law School, September 24 2024

university [is] the creation of knowledge and the dissemination of it. The commercialization is nice to have, but it's an alternate path of dissemination." This foundational understanding ought to inform university approaches to GenAI patenting, though how this orientation evolves in response to the federal funding landscape and shifting administrative priorities should be researched further.

In practice, universities have begun to experiment with licensing strategies that attempt to reconcile open science and commercialization. One approach involves the use of copyleft licenses such as the General Public License (GPL), as evidenced by policies found at the University of Texas, Univates in Brazil and the International Institute of Information Technology in Hyderabad, India (Stallman 2021). Releasing an innovation under a GPL license allows the university to fulfil its open science mandate by making the technology widely accessible, while the copyleft provisions of the GPL ensure that any derivative works must also remain open, potentially deterring commercial entities from incorporating the work into proprietary products without also open-sourcing their own modifications. In this way, the GPL can serve as a shield for open science and a barrier to certain forms of commercialization, since firms seeking to develop proprietary applications may be dissuaded by the requirement of disclosure. This is not to say though, that this approach is without its drawbacks. If a university invests in patenting a technology but then releases it under a GPL, any attempt to enforce the patent against a commercial infringer may be undermined by the license itself. Moreover, while it may facilitate collaboration among academic and non-profit organizations, it could limit the technology's uptake by industry partners who require greater freedom to commercialize without open-sourcing their enhancements (Tran 2025). This generates a strategic dilemma: maximizing

openness may then come at the expense of commercial impact, while prioritizing commercialization can constrain the reach and collaborative potential of university-generated knowledge. Recognizing this tradeoff, some universities and research consortia have turned to additional nuanced licensing models, such as permissive open-source licenses (i.e. MIT, Apache 2.0) or a new generation of AI-specific licenses, known as OpenRAIL or RAIL variants, which layer ethical, safety, and misuse restrictions onto traditional copyleft frameworks. Such licenses represent a deliberate effort to reconcile openness with institutional responsibility and unlike purely permissive licenses, exemplifying how universities can tailor their strategies to align with both their public mission and the realities of a competitive innovation landscape. By adopting GPL or OpenRAIL-type licenses, institutions mitigate the social costs of uncontrolled proliferation. For example, a university releasing a GenAI model under these licenses could freely share it with academic researchers while restricting commercial entities from using it for surveillance applications, a balance that traditional patenting strategies alone cannot achieve. This addition then transforms intellectual property from a purely economic asset into a governance mechanism, allowing universities to assert their stewardship over how innovations propagate throughout the ecosystem (Synk 2020; The Turing Way 2019; Licenses.ai 2023; HackMD 2022; FOSSA 2025).

Role of Federal Funding: Evolving Constraints and Opportunities

Beyond the strategic considerations posed by the "general research ecosystem", the critical role of federal funding is a valuable driver of institutional patenting strategy. As Swartley notes: "we have received a huge amount of support from federal institutes, and they impose certain obligations on us, which we are constantly having to adapt to." Federal support has constituted the predominant funding source for academic research and development since 1953, accounting

for 55% of total academic R&D expenditures in 2021. Universities such as Johns Hopkins received approximately $1,022,300,000 in NIH research funding during fiscal year 2024 across more than 3,200 active awards, with approximately $281,446,000—representing 27.5% of the total—allocated as indirect cost reimbursement (Hub Staff 2025). The potential reduction or elimination of these funds would significantly impact institutional research capacity; thus, it falls upon the researchers and innovation offices to comply with the extensive regulatory requirements that accompany the federal grants. The NIH grants policy applicable to all grants initiated after October 1, 2023, stipulates that “If a recipient has failed to comply with the terms and conditions of award, NIH may take one or more enforcement actions which include disallowing costs, withholding of further awards, or wholly or partly suspending the grant, pending corrective action" (NIH 2024). Compliance necessitates adherence to the Code of Federal Regulations (CFR), particularly sections devoted to federal grants and contracts such as 2 CFR section 200, with additional requirements potentially imposed on institutions which are flagged for elevated grant management risk. As a result, non-compliance carries significant financial penalties, as demonstrated by Pennsylvania State University’s $1,250,000 settlement resolving False Claims Act violation allegations regarding cybersecurity requirement non-compliance across fifteen Department of Defense and NASA contracts (Insurance Journal 2024).

The dedicated funding and regulatory oversight of federal grants has heightened with regards to artificial intelligence innovations. The National Science Foundation (NSF) has assumed a pivotal role in the funding domain, announcing $140 million in 2024 to establish seven dedicated AI research institutes (University of Maryland 2023). The federal investment is strategically directed towards priority areas: trustworthy AI and ethical frameworks, cybersecurity

applications, climate-smart agricultural initiatives, neural and cognitive foundations, and decision-making processes with societal impact assessments. Simultaneously, the NSF has implemented specialized AI research protocols, including prohibitions against reviewers uploading proposal content to non-approved AI tools (Burgan 2023). The emergent nature of AI technologies has prompted enhanced oversight mechanisms to ensure appropriate fund allocation and utilization.

This intensive monitoring is exemplified by the NSF Foundry's support for the AI-driven RNA Foundry initiative, known as AIRFoundry, at the University of Pennsylvania and University of Puerto Rico through a $18-million, six-year grant (Penn Center for Innovation 2024). While the AIRFoundry is not formally designated as an Engineering Research Center (ERC), it is subject to NSF's rigorous review standards typically reserved for major multi-year awards, such as IoT4Ag at the University of Pennsylvania, including comprehensive annual site visits and a particularly demanding multi-day evaluation at the project's midpoint, designed to ensure continued alignment with NSF's ambitious research and commercialization objectives (U.S. National Science Foundation, n.d.). However, Professor Steven Weiner, who is involved with the project underscores the importance of this funding: "ultimately [the funding] is what is trickling down and incentivizing our research."[4]

Implications for Generative AI Patenting Strategies:

The aforementioned landscape raises significant operational considerations, as Swartley notes: "Being purely academic [has been] a bit of a luxury. It only exists in spaces where somebody is

[4] Steven Weiner, conversation with the author, April 2025

willing to pay for it. And…[we have] been in an environment that has been conducive to this…unrestricted, unconstrained research and scholarship for decades now."

Thus, two critical implications emerge from these funding challenges. First, research priorities and funding allocations will likely exhibit increased volatility, aligned with governmental priorities and development objectives. Through conversations with patent experts such as Bob Steinberg, it becomes evident that the United States' holds a presumptive orientation towards protecting American researchers, thus universities will need to clearly demonstrate the substantive contributions of human researchers to justify continued investment in individuals rather than merely in technology. Still, as GenAI technologies become more deeply integrated into the research processes, it remains unclear how this orientation will translate. If federal agencies continue to favor established human researcher over purely technological solutions, the case for ongoing investment in individuals may persist. Yet, the rise of GenAI could prompt a reevaluation of funding rationales, especially if the locus of innovation shifts from individual expertise to the capabilities of advanced AI systems. Second, institutions will increasingly seek alternative funding sources. This diversification may manifest through either intensified patenting and licensing of university research or expanding reliance on private funding, each of which presents their own distinct challenges. Currently, as Swartley mentions, "commercial sponsored research has been maybe 20% of Penn's annual research enterprise and people seem generally comfortable with 20%. What if it was 50%? That [may] feel really different, and if it was 50%, how do you get [there]. I think you might only get to 50% if you started to compromise some of those academic principles."

An experimental attempt to accommodate sponsor requirements to increase private funding was conducted at Penn State, yielding adverse institutional consequences and

highlighting the tensions inherent in such approaches.[5] As commercial sector engagement intensifies, deviation from core academic principles becomes increasingly probable. Consequently, developing internal, iterative, and independent funding streams through licensing Generative AI innovations will become increasingly vital in the coming years. The commercial viability of such strategies depends heavily on industry attitudes toward AI-driven IP, which vary significantly across sectors. In high-tech and software domains, where GenAI's foundational models (Stage 1 & 2) reside, industry resistance to traditional university licensing models is rather pronounced due to open-source norms and the perception of AI as a broadly applicable tool rather than a patentable invention (IAM 2024). By contrast, biotechnology and life-sciences sectors, where GenAI is applied to design tangible innovations such as materials or therapeutics (Stage 3), exhibit greater receptiveness, even potentially mirroring existing practices for non-GenAI discoveries (Bio.Law 2025).

Organizational Priorities

The second dimension of institutional strategic evaluation encompasses organizational priorities, which can be further divided into university institutional values and the mission of professors and researchers.

University Institutional Values

Unlike corporate patenting strategies, which are primarily motivated by exclusivity benefits—whether through first-mover advantage or revenue generation—university strategies cannot solely be tethered to financial outcomes. The shared value orientation characteristic of

[5] John Swartley, conversation with the author, December 2024

commercial entities exhibits substantial divergence within the university landscape which drives unique patenting strategies.

As John Swartley observes, “just like anything, when it comes to IP management or IP commercialization, you'll find a pretty wide range of approaches between institutions. Some of it just depends on the historical culture, but also the tolerances—there is always going to be a push and pull between pure academia and the possibility of more commercial-looking translation of technology." This can be seen in the variance between institutions such as NYU and Penn, where “academia will almost always win, at least at Penn” (Swartley), whereas “[Universities like] NYU in particular [have] been pretty forward-thinking about these issues that a lot of the rest of the academic community is catching up”. Similarly, Bray reaffirmed Penn’s prioritization: “Penn would typically err on the side of disclosure and if that compromises the commercial opportunities, so be it.”

Mission of Professors and Researchers

The historical foundation of intellectual property rights establishes that inventions are initially owned by their creators—a principle enshrined in both the United States Constitution and federal patent law. The U.S Supreme Court affirmed this principle specifically for university contexts in Board of Trustees of the Leland Stanford Junior Univ. v. Roche Molecular Systems, Inc., 563 U.S. 776 (2011), establishing that federal funding of faculty-led research does not alter this principle: inventors in university settings utilizing federal funds remain the initial owners of their inventions.[6] Universities hosting federally supported research have neither an obligation nor a

[6] Board of Trustees of the Leland Stanford Junior Univ. v. Roche Molecular Systems, Inc., 563 U.S. 776 (2011)

mandate under federal law to assume ownership of faculty-developed inventions. Ownership of patent rights may be transferred through a written instrument signed by the inventor, distinguishing ownership from control (Subcommittee of the Association's Committee A on Academic Freedom and Tenure 2013).

Despite this, some universities have established ownership of faculty patent rights as a general condition of employment, citing university facility utilization or participation in externally funded research as justification. This holds the potential to generate institutional conflicts of interest between governance responsibilities and financial interests in patent exploitation. In May 2013, the University of Pennsylvania established a precedent where employment status supersedes traditional expectations of faculty independence (Jackson, Jr., and Bonnell 2023). While this precedent specifically addressed faculty development of external online courses that might compete with university offerings, the principle could extend to other intellectual property contexts, such as publishing through university-affiliated press. As Chris Callison-Burch observes regarding technology introduction, faculty is typically inclined to "publish everything, especially in a university setting." In his assessment, professors tend to prioritize the open science mandate given their reliance on government grant funding, while students demonstrate greater commercialization orientation, positioning professors in mentorship or advisory capacities. This aligns with research that indicates that over time, universities are increasingly expected to assume central roles in the knowledge-centered economy. However, transformation has been relatively limited and gradual, which is potentially attributable to unchanged incentive structures for key stakeholders, particularly faculty members (Susalka 2019). The predominant

model for academic career advancement remains heavily focused on publication and citation-based metrics (Sandberg et al. 2014).

Universities are also increasingly attentive to the risks associated with AI-generated content in research. For example, the University of Maryland also advises cautious approaches to AI-generated content in research contexts, noting potential third-party intellectual property infringement risks (University of Maryland 2025). AI-generated outputs deriving from preexisting works may necessitate additional scrutiny to avoid legal and ethical implications. For potentially patentable research, while the US Patent Office recognizes AI-assisted inventions where persons have made significant contributions, fundamental questions may sometimes arise regarding whether utilization of others' copyright-protected information and data to train that AI might impact ownership determinations based on their technical contributions.

Implications for GenAI Patenting Strategies

The varying priorities of professors and researchers highlight two key takeaways for universities. Firstly, developing more robust incentive structures tied to patenting activities is essential, particularly in Generative AI fields where technological leadership offers substantial benefits. Providing faculty with meaningful involvement in patenting decisions and applied research incentives can also enhance the likelihood of successful commercialization. However, as universities pursue these advancements, maintaining vigilance over ethical data and AI utilization remains crucial to avoid reputational risks that could undermine institutional value as well as impair patent value.

As Marc Sedam emphasizes, "We do not want to be the Henrietta Lacks of human data... The moral and ethical implications we have of patient records, patient data, how do we de-identify it? How do we make sure that we make it accessible but not available?" Ensuring transparency and responsible handling of sensitive data is essential to sustaining long-term credibility, which in turn bolsters the marketability and legitimacy of patented innovations.

Ultimately, the path forward for universities involves more than simply increasing patent output; it requires a deliberate recalibration of how patenting fits into broader institutional goals. In the unique context of AI, universities have an unprecedented opportunity to shape the norms of responsible development, not only through the inventions they patent, but by what they choose to keep open, defensively publish, or share for the public good. Guided in part by their academic mission rather than purely commercial interests, their success will depend on structuring incentives that reflect evolving funding and faculty motivations, as well as making strategic decision about how to lead by example in fostering an ethical and inclusive AI ecosystem. By integrating these considerations, universities can ensure that their pursuit of technological leadership in AI remains both sustainable and true to their core academic values.

INNOVATION ASSESSMENT

This section represents a critical preliminary analysis that precedes more nuanced discussions of institutional patent strategies. While subsequent sections will explore the intricate landscape of university-specific approaches, the present investigation focuses on a more universal question: whether an innovation can substantively meet the essential thresholds of patentability. The principles articulated herein are intentionally broad and universally applicable, derived from authoritative sources including the United States Patent and Trademark Office (USPTO), federal legislative frameworks, and judicial precedents. By establishing a standardized evaluative mechanism that considers the eligibility of inventions utilizing GenAI, this approach provides a rigorous, objective methodology for assessing fundamental patent eligibility. Should an innovation fail to satisfy this criteria, researchers can strategically redirect their efforts towards alternative protection strategies, ensuring an efficient use of finite institutional resources.

The decision-making process can be structured into two primary assessments: the Subject Matter Assessment and AI Contribution Mapping. These assessments ensure that the innovation not only aligns with existing patent eligibility requirements but also complies with the evolving regulatory requirements to demonstrate sufficient human contribution.

Subject Matter Assessment: Satisfying the Alice Framework and Distinguishing the Prior Art

The first step in evaluating patentability involves assessing whether the innovation satisfies the criteria established in the *Alice* framework, which governs subject matter eligibility for patents. For Stage 1 AI inventions, these challenges are particularly acute as they often involve processes

that either replicate human activities or rely on intangible software. Courts have repeatedly held that speed or mere efficiency alone does not render an invention patentable; instead, they ask whether a human could perform the same process without a computer, regardless of the time or difficulty of execution it would take (Fisher n.d.). As Bob Steinberg succinctly explains, "the invention will not be patent eligible if a computer is simply used to automate a routine task, or simply lends speed or efficiency to the performance of a routine task" (Steinberg 2024).

To survive Alice scrutiny, Stage 1 claims must go beyond generic invocations of AI tools such as neural networks or machine learning and ought to describe specific implementations that humans cannot replicate. For example, patents tied to unique data-reduction techniques that require extreme precision, down to the single bit, have been upheld in cases such as *Blue Spike LLC v. Google.*[7] Similarly, patents fare better when they are tied to specific applications or physical environments, as demonstrated in Ocado Innovation Ltd. v. AutoStore. In this case, the court found that the "control system described in the [patent] could not be performed 'entirely' in the human mind," rejecting analogies to human activities and emphasizing the innovative aspects of the claimed technology. Stage 3 inventions, by contrast, should face fewer (if any) Alice hurdles in applications where the output of a generative AI system is e.g., the design for a new, tangible drug, device, material, or circuit. Furthermore, recent developments suggest that there is a willingness within the U.S. Patent & Trademark Office to increasingly allow AI-related patents to pass subject matter eligibility thresholds. The 2019 Revised Patent Subject Matter Eligibility Guidance significantly reduced uncertainty in patent determinations and led to a 25% decrease in first office action rejections for Alice-affected technologies, including GenAI innovations

[7] Blue Spike, LLC v. Google Inc., No. 16-1054 (Fed. Cir. 2016), https://law.justia.com/cases/federal/appellate-courts/cafc/16-1054/16-1054-2016-10-14.html

(United States Patent and Trademark Office, Commerce 2019; Sosna 2024). Building on this, the USPTO's 2024 Guidance Update – issued in response to Executive Order 14110 provides even greater clarity for AI inventions. The update introduces new illustrative examples specifically addressing AI-related claim, such as anomaly detection, speech separation, and personalized medical treatments, which can be used to guide universities navigating eligibility under 35 U.S.C. § 101 (USPTO 2024). This momentum alongside resources such as the Artificial Intelligence Patent Dataset (AIPD), which analyzes over 15 million U.S. patent documents published since 1976, can be leveraged by universities to identify successful strategies for drafting claims and shaping applications (Office of the Chief Economist USPTO 2021).

Beyond legal considerations, the commercialization potential of an innovation plays a crucial role in determining whether to pursue patent protection. This aspect is particularly relevant when considering the practical value and utility of the patent in the marketplace. Scott Lobel, a testifying expert with over 20 years of experience in economic and financial analysis, articulates this concept eloquently, emphasizing the importance of differentiation and commercial viability: "There is this notion of commercial success, how differentiated the product is... So if what AI is doing, or the use of AI, is coming up with a green pill and not a red or white pill, it's not obvious that will come out as part of the commercial success of the patent."[8] Lobel's analogy underscores a valuable principle: in the context of AI-enabled innovations, simply applying GenAI to produce a slightly different outcome—such as changing the color of a pill—may not be sufficient to warrant patent protection or guarantee commercial success. Universities must therefore carefully evaluate not only the technical novelty of their AI-driven innovations but also their

[8] Scott Lobel, in conversation with the author, July 10 2024

potential to create significant value in the industry. This assessment should consider factors such as the innovation's ability to solve pressing problems, its potential to disrupt existing markets or create new ones, and its capacity to generate substantial revenue or cost savings for potential licensees or end-users.

Moreover, the commercial potential of an innovation can influence its chances of surviving legal challenges. Courts and patent examiners may view commercially successful inventions more favorably, as their market impact can serve as evidence of non-obviousness and inventive step (Xie 2022). Consequently, universities should prioritize innovations that demonstrate clear market differentiation and substantial commercial promise when crafting their AI-related patent strategies.

AI Contribution Mapping: Meeting the "Significant Contribution" Threshold:

The second critical threshold for determining whether an innovation is patentable involves assessing whether sufficient human intellectual input exists to meet inventorship standards as outlined by the USPTO (United States Patent and Trademark Office 2024). This requirement is particularly relevant for Stage 3 AI inventions, such as pharmaceuticals, materials, or medical devices, where human inventors typically contribute substantive, delineable input. The 2024 Inventorship Guidance for AI-Assisted Inventions reaffirms the principle established in Thaler v. Vidal: patents are intended to incentivize human ingenuity and require significant contributions from natural persons during the creation process. However, the guidance also acknowledges the complexity of human contributions in AI-assisted innovations, which can manifest in diverse forms, such as designing, building, or training AI systems tailored to specific problems; selecting

high-quality datasets for foundational training; or incorporating AI-generated outputs into broader systems with substantial modification (Hsu 2024).

For example, reinforcement learning with human feedback (RLHF) introduces variance based on human preferences and serves as evidence of substantial intellectual input. Similarly, selecting model architectures optimized for specific applications, rather than relying on general-purpose models like GPT-4, can differentiate an invention and strengthen its eligibility for patent protection.[9] Courts have long-applied factors such as those outlined in Pannu v. Iolab Corp. to evaluate inventorship claims. To qualify as an inventor, an individual must contribute significantly to either the conception or reduction to practice of the invention and must offer more than routine explanations of well-known concepts or prior art. Additionally, the guidance emphasizes that conception does not occur when there is only an unrecognized accidental creation. The Pannu factors provide a framework for determining whether a human's role meets the threshold of inventorship: contributions must be significant in quality when measured against the full scope of the invention and must go beyond merely explaining existing knowledge[10].

The USPTO has also clarified principles for identifying inventorship in AI-assisted inventions. For instance, while using an AI system in creating an invention does not negate a person's contributions as an inventor, merely recognizing a problem and prompting the AI system to produce a solution does not qualify someone as an inventor. Inventorship requires active engagement—such as modifying AI-generated outputs or designing systems to elicit specific solutions—and not passive reliance on AI tools.

[9] Jonathan Frankle, "How to Train an LLM Model from Scratch," talk, February 2024.
[10] Pannu v. Iolab Corp., 155 F.3d 1344, 1349 (Fed. Cir. 1998)

Despite these guidelines, challenges arise when AI performs most of the work—a phenomenon referred to as "thin wrappers." As Robert Plotkin explains, "Companies where ChatGPT is doing 99% of the work…[are] going to be hard to patent because it's obvious" (Plotkin 2025). Merely training or fine-tuning a model using novel datasets often fails to meet inventorship standards unless these efforts result in substantial improvements beyond what is routine or expected.

For universities, this results in a heightened necessity for careful documentation of human contributions throughout the innovation process to ensure compliance with legal standards and avoid disputes over ownership or liability. As Terry Bray emphasizes, "inventorship is critically important…you want proper legal counsel involved because if it ends up in court you want confidence that inventorship is correct."

Contextual & Global Considerations:

An intriguing nuance emerges when universities, such as Penn, elect to file patents outside the United States, particularly when driven by the interests of innovators, institutional goals, or strategic partnerships. Filing internationally introduces a host of complexities that extend beyond U.S. patent law and require careful consideration of jurisdiction-specific nuances. While U.S. law emphasizes human inventorship and treats AI as a tool rather than an inventor, other countries may interpret these principles differently, potentially impacting patent eligibility and inventorship claims for AI-assisted innovations.

For instance, European patent law mandates that inventors must be "natural persons" while Japan explicitly defines inventions as "products of human activity," excluding non-human creators.

South Korea similarly adheres to strict requirements for human involvement in the inventive process under Article 33(1) (Phelan 2022; Matsumura and Takahasi 2025; Lee and Kim 2024). Ownership is further complicated when AI tools are licensed rather than developed in-house. As John Swartley explains, "[Ownership] is always going to be if the invention is created by an employee of the university…That also includes inventions created by employees through AI tools." However, should it be obtained through a licensing agreement that conveys certain rights to the supplier, universities must navigate these agreements prior to the product or service being developed.

Given this, it may be advantageous for universities to capitalize on the Patent Cooperation Treaty (PCT), which offers a streamlined pathway for filing patents in over 150 member countries through a single application. This allows universities to defer costs associated with national filings while gaining valuable feedback on patentability during the international search phase. Additionally, the extended timeline provided by the PCT—up to 30 months—enables institutions to assess market opportunities and refine their filing strategies based on jurisdiction-specific requirements (USPTO 2018). Furthermore, having region and industry-specific legal representatives who can navigate various jurisdictional laws and practices could optimize filing strategies for maximum efficiency.

Taking this into account, as AI technologies continue to advance, institutions must adapt their strategies to ensure they can effectively protect and commercialize their innovations. Leveraging tools such as the AIPD, engaging in strategic international filings through the PCT, and fostering collaborations with industry partners across jurisdictions will enable them to demonstrate clear

technological advancements, substantial human input, and compelling commercial potential – all necessary in ensuring the university patent strategy succeeds.

DECISION TO PURSUE PATENTING

Building upon the preliminary analysis of patent eligibility thresholds, universities face a critical decision point when managing GenAI innovations. Three distinct pathways emerge from this juncture: when an innovation satisfies patentability requirements, institutions can either pursue formal patent protection or deliberately opt for non-patent protection mechanisms. Conversely, if an innovation fails to meet patentability thresholds, universities must explore non-patent protection strategies. This decision framework has gained particular attention as research has demonstrated, "new forms of non-patent IP are placing [pressure] on the traditional university intellectual property and technology transfer regime" (Marr and Phan 2020). The divergent pathways—patent protection versus non-patent protection—each offer distinct advantages and limitations for academic institutions. On the one hand, patent protection offers exclusive rights but requires substantial investment and disclosure, while on the other non-patent approaches offer faster implementation timeframes and strategic flexibility. Understanding when each approach serves institutional objectives optimally is likely to become a critical competency for university technology transfer offices. Thus, this section analyzes the factors influencing these strategic decisions. For innovations pursuing the patenting pathways, considerations related to patent scope and temporal strategy will be examined and for non-patent approaches, strategic data and positioning will be analyzed.

Relevancy & Determination of the Patent Scope:

The dramatic growth in GenAI patent families—from a mere 733 to an impressive 14,000 by 2023—underscores the field's accelerating development as well as its competitive significance (Alemán et al. 2024). When crafting patent applications, universities must then navigate a delicate balance between ensuring claims are sufficiently broad to defend against competitors while remaining specific enough to withstand validity challenges. The USPTO has provided valuable guidance in this domain, noting that "claims that involve specific hardware components and improvements to practical applications are more likely to be considered patent-eligible" (International Intellectual Property Law Association, n.d.). This direction proves to be particularly salient given that a fundamental challenge in patenting GenAI innovations stems from the inherent "complexity and opacity" of artificial intelligence systems (Cheong 2024).

The first element of the strategy is identifying the specific patentable component—whether it resides in the algorithm itself or in the output the algorithm produces. Current legal frameworks generally render the output considerably easier to patent than the underlying algorithm, from a subject matter eligibility standpoint, but they also necessitate meticulous documentation of human involvement and establishing the essential "human-in-the-loop" element that strengthens patentability claims. Additionally, successful patent strategies typically target specific technical implementations and demonstrable improvements rather than abstract principles, as explored in the forthcoming example of Insilico Medicine. In 2022, Insilico secured a patent, US11403521B2, for their "Mutual Information Adversarial Autoencoder (MIAAE)" technology, in 2018. This patent protects their proprietary system for utilizing GenAI to create novel small molecules suitable for synthesis and pharmaceutical testing. The technology functions as the

algorithmic foundation for their Chemistry42 engine, which has already yielded new treatments for conditions such as fibrosis, cancer, and COVID-19. To understand what made their patent successfully meet the threshold, it is vital to understand the underpinning technology. Insilico's MIAAE is able to incorporate object encoders and decoders that process molecular structures to generate latent representations and subsequently new molecular entities. It includes condition encoders and decoders that process variables such as cell states or protein binding sites, and discriminator networks that ensure generated molecules maintain properties similar to training data while specifying certain conditions (Buvailo 2023). Following an analysis of their patenting approach, the strategies Insilico Medicine adopted to successfully navigate patentability thresholds become evident. First, they patented the specific implementation of generative AI for drug discovery rather than attempting to claim broader algorithmic principles. Second, they maintained significant human scientific involvement throughout the development process—while GenAI generated molecular candidates, human experts directed development, selected targets, and optimized compounds, thereby establishing the critical "human in the loop" element required for patentability – should they attempt to patent drugs discovered by the platform. Alex Aliper, Insilico Medicine's president, succinctly captured their strategic approach: "Everybody patents molecules. For algorithms, it's a choice" (Ratner 2023).

For universities, the Insilico case provides a tactical blueprint for successful patenting GenAI innovations. Insilico secured their patent by focusing on the practical applications of their model, rather than merely on the underlying algorithms, a strategy that aligns with current patent office requirements. Demonstrating tangible, measurable improvements—such as Insilico's achievement of reducing drug discovery timelines to 18 months at just 10% of conventional

costs—can strengthen the case for patent approval and highlight the real-world impact of the innovation (Insilico Medicine 2025). Additionally, documenting human involvement throughout the process remains crucial when it comes to patenting new therapeutics discovered by using systems like MIAAE. Insilico's approach involved scientists curating datasets, training the target discovery engine, and making final selections from AI-generated options for further testing, ensuring compliance with patent eligibility criteria that require significant human contribution. An interesting nuance is Insilico's approach to strategic disclosure management, wherein they balanced the benefits of patent protection with the need to maintain a competitive advantage on their broader Pharma.AI platform. Thus, while patents must provide enough detail to meet legal requirements, universities can withhold proprietary methodologies or datasets that are not essential for understanding the invention and for enabling its implementation in a basic, illustrative way.

If a researcher does wish to pursue algorithm patenting, significant challenges are likely to arise. Namely, patent applications must describe and claim how the invention holistically integrates the algorithm with non-routine technology, and/or concretely improves a technology. Simply stating "apply it with a computer" is woefully insufficient per current USPTO guidance; nor is simply reciting an intended field of use sufficient by itself to overcome *Alice*. Matthew Chun, a patent agent at Fish & Richardson, advises applicants to explicitly describe which aspects of their machine-learning processes cannot be performed by human cognition, frame innovations as technical improvements, tie them to particular machines, and be highly specific in their claims. However, Shelby Newsad, an investor with venture capital firm Compound, points out a core issue with algorithm patenting: "for many, algorithms themselves aren't the sole proposition,

because there's a new state-of-the-art method published every couple of weeks" (Ratner 2023). This rapid evolution creates uncertainty about long-term patent value, particularly in academic settings where publication pressures are high.

The decision context to pursue patenting also depends heavily on the industry and the stage of AI innovation involved. For Stage 1 inventions, as seen in autonomous vehicles algorithms, the AI is in the product itself, early patenting is essential not only as public deployment can preclude later protection, but also because Alice-related eligibility challenges often force narrow claim scope. Patents in this domain risk being limited to highly specific technical implementation, such as sensor fusion architectures, to avoid abstract idea rejections, even if the underlying AI breakthroughs are transformative. On the other hand, Stage 3 inventions, such as pharmaceutical innovations, put greater emphasis on the outputs, which aligns with current patent offices' greater clarity on patenting tangible materials and devices than intangible methods and algorithms. Here, human contribution requirements are procedural hurdles, perhaps resulting in in-house best practices through documentation, but will likely rarely constrain claim scope for the final product. A drug's structural novelty and efficacy remain the primary criteria, regardless of whether AI accelerated its discovery. Despite these differences, it is ultimately a well-defined patent scope that ensures GenAI inventions are effectively protected, balancing specificity to meet legal requirements with strategic breadth to maximize both their commercial and academic potential.

Temporal Strategy: Optimizing Timing and Value Creation

The effective protection of GenAI innovations through patents requires not only careful attention to scope, as previously discussed, but also a sophisticated temporal strategy. It is the timing of patents, where in the life cycle of innovation they are filed, that often substantially determines

their commercial and academic value. The following section analyzes two components of this consideration: the patent lifecycle and approaches to licensing and monetization.

Patent Lifecycle Management:

Strategic timing has become increasingly important in the GenAI patent landscape, which has experienced a compound annual growth rate (CAGR) of 31% in recent years (Azhar 2024). This remarkable growth has created a highly competitive environment where first-mover advantage can be decisive. One significant consideration is the typical 18-month lag between patent filing and publication, which creates a window during which competitors remain unaware of pending inventions (TRP International 2025). This lag creates a window for university researchers, offering both opportunities and challenges. During this period, universities may refine their research and explore commercialization options without external influence. Furthermore, this secrecy can be advantageous for university technology transfer offices to identify potential industry partners or licensees before the invention becomes public knowledge (Awang and Abdullah 2024). However, it is important to note that research shows that post-publication licensing activity often increases significantly after the 18-month mark, particularly for investors with limited previous licensing experience. This is understandable given that "requiring inventions to be published through a credible, standardized, and centralized repository…mitigates information costs for buyers and sellers and, thus, facilitates transactions in the market for ideas" (Hegde and Luo 2018). The delay then enables universities to prepare a robust commercialization strategy during the pre-publication period to maximize the value of their patents should they enter the public domain. Nevertheless, as a result of the pervading academic publication pressures, universities also ought to be careful of premature disclosure through academic journals or conferences as this can jeopardize the novelty requirements.

Therefore, strategies such as filing provisional applications early could be highly advantageous. Provisional applications secure a priority date while allowing researchers a 12-month window to refine their invention—particularly valuable in the rapidly evolving GenAI field. Beyond being cost-effective, provisional applications create an opportunity to assess commercial viability and attract potential investors before committing resources to a full patent application. While some industry perspectives suggest provisional filings might signal uncertainty about the invention, this approach is uniquely beneficial for GenAI innovations given the fluid regulatory landscape (Krajec 2021). As patent laws specific to GenAI continue to evolve through legislative, judicial and regulatory reform, the provisional period provides universities crucial time to adapt their innovations accordingly. This flexibility would allow institutions to optimize finite resources while strengthening the inventions' prospects of meeting increasingly sophisticated patent thresholds in the GenAI domain.

It is important to note that despite careful planning, patent application processes for these innovations still carry inherent uncertainties. As Gregory Rabin, a patent attorney at Young Basile, observes, it is not always clear whether an applicant will receive a patent, as especially when it comes to GenAI, the changing interpretations of patent law by the patent offices will become a key determinant (Ratner 2023). This uncertainty is then compounded by challenges in infringement detection and enforcement. For example, if a company secretly adopts a patented technology for in-house use, infringement can be highly difficult to detect. In a similar vein, the current regulatory framework encourages patenting the implementation of an idea, not the abstract idea or model itself, meaning that competitors may find alternative implementations that achieve similar results without infringement. As Verseon's Ed Ratner states, deciding to patent is

ultimately “a business question— a cost–benefit tradeoff,” one that universities must carefully weigh.

Licensing and Monetizing GenAI Patents:

Once secured, patents require thoughtful monetization and licensing strategies to realize their full potential. This aspect is particularly complex for GenAI innovations due to their versatility and multiple opportunities for applications across different industries. This raises an interesting question: If there are arguably "infinite" possibilities for a generative AI model’s application, how could one adequately determine an appropriate licensing valuation that reflects its true potential? A comprehensive solution to this valuation problem does not yet exist, primarily because the field lacks established licensing precedents. Nevertheless, in speaking with experts, the following analysis presents relevant criteria and considerations which could aid universities in developing forward-looking licensing strategies.

As previously mentioned, patenting valuation for GenAI innovations presents complexities that traditional methods struggle to address adequately. While conventional approaches such as financial valuation techniques or comparative market analysis offer standard frameworks, this study deliberately examines legal frameworks for assessing patent infringement damages as an alternative lens that provides more instructive insights. By analyzing how courts determine monetary compensation when patents are infringed, one can reverse-engineer principles for establishing intrinsic value in these cutting-edge technologies. This approach, rooted in practical jurisprudence rather than theoretical models, offers pragmatic lessons for determining the worth of GenAI-assisted inventions in an otherwise speculative market. Expert consultation with Scott

Lobel highlights the significance of the Georgia-Pacific-Factors in this context.[11] As Lobel explains, “It sets up a negotiation that says, if you imagine a world where the infringer didn't infringe but came to the patent holder, and said, how much do I have to pay for the patent? There'd be a negotiation that took place at that time, between the patent holder, who has a minimum willingness to accept the royalty value, and the potential licensee, who has in his or her mind, a maximum willingness to pay for that time. And you think about all the factors that would go into that negotiation as to whether or not they could come to an agreement before the infringement took place. Because that is, at the end of the day, the fair value of a royalty right to that patent." This valuation approach intersects with the Entire Market Value Rule (EMVR), which permits damages based on an entire product’s value, rather than just the patented feature, in cases where the patented feature serves as the primary driver of consumer demand (Wheelock 2018). This distinction becomes especially relevant when considering innovations that universities may produce across industries. For instance, in the technology sector, royalties must be calculated based on the “smallest saleable unit” incorporating the patented innovation—such as a specific chip within a smartphone rather than the entire device. In contrast, pharmaceutical patents often cover the active ingredient that constitutes the entirety of a drug's value (IAM 2017). These industry differences create varying valuation landscapes for GenAI innovations applied across different sectors.

Regarding the broader question of monetizing GenAI patents, several strategic approaches have emerged for maximizing patent value. Cross-licensing has proven especially valuable “in industries where products rely on multiple interconnected patents," an increasingly common

[11] Georgia-Pacific Corp. v. United States Plywood Corp., 318 F. Supp. 1116 (S.D.N.Y. 1970).

scenario in the GenAI landscape where algorithms, training methodologies, and applications often involve complementary technologies. This approach would enable universities to exchange licenses with industry partners who may possess implementation expertise or computational resources that universities lack, while the industry receives access to cutting-edge academic research and specialized domain knowledge. For instance, university-developed GenAI models in healthcare diagnostics or materials science can be cross-licensed with industry-developed infrastructure, creating mutual benefit while accelerating real-world applications (Chakraborty 2025).

"Carrot" licensing represents another effective strategy, offering "potential licensees a glimpse of the benefits and profits that could be reaped by licensing a particular patent" (PJ Parker n.d.). This approach precisely aligns with the novelty of GenAI research that arises from fundamental research and may not have obvious commercial applications at first glance. By demonstrating how a groundbreaking GenAI algorithm could improve drug discovery process or enhance workflows, universities can break into the industry and thus bridging the gap between academic innovation and implementation. Notably, universities are also pioneering alternative models that bypass traditional IP licensing frameworks altogether. The NSF AIRFoundry exemplifies this shift as instead of pursuing royalties through patent licensing, they have a proposed fee-for-service model, offering external partners access to its AI-optimized RNA design platform and expertise. Partners would retain ownership of their intellectual property while paying for the refinement and optimization of their assets, a strategy particularly effective in biotechnology, where traditional licensing often faces friction due to the fluid, collaborative nature of innovation. For example, a pharmaceutical company developing an RNA therapeutic might pay

AIRFoundry a fixed fee to iteratively optimize delivery vehicles using proprietary algorithms, avoiding protracted royalty negotiations.[12] This represents a hybrid approach tailored to sector-specific realities while exemplifying how universities can develop internal, iterative funding streams by monetizing GenAI without relying on licensing frameworks. This pioneering ability to generate independent revenue will prove vital for sustaining innovation in domains where traditional IP frameworks and pre-existing funding opportunities may falter. The emergence of university innovation showcases and AI-focused technology transfer events, such as the Ohio State Research and Innovation Showcase, AI Infra Summit Innovation Showcase, and Georgetown's University Research and Innovation Showcase further facilitates these demonstrations, allowing academic institutions to present compelling use cases of their gen-AI patents to potential industry licensees, thereby securing more advantageous licensing arrangements (OSU 2024; Georgetown 2025; Kiasco Research 2025).

In the end, the decision to pursue patenting for university GenAI innovations requires a multipronged approach that balances technical, legal, and commercial considerations. By defining a clear and defensible patent scope, universities can ensure their claims are robust enough to protect their innovations while remaining flexible for future developments. Strategic timing in the patent lifecycle—whether through early provisional filings or selecting to delay patenting till the internal benefit has been exercised—provides researchers with critical advantages in securing rights and adapting to the ever-changing regulatory environment. Furthermore, effective monetization strategies, such as cross-licensing, "carrot" licensing, and

[12] Steven Weiner, in conversation with the author, April 2025

service-based models like AIRFoundry's, enable universities to capitalize on the value of their patents.

These two components of the overall framework outlined in this paper have become increasingly critical as universities continue to face challenges in securing federal funding. With persistent reductions in overhead budgets and decreased grant allocations, intellectual property revenue streams are poised to play a pivotal role in sustaining research and innovation efforts. As such, strategic discussions surrounding patent valuation, licensing strategies, and commercialization pathways must become immediate priorities for universities across the United States to ensure long-term financial stability and maintain their leadership in cutting-edge research.

In light of this, pursuing patenting is no longer merely a legal exercise but a strategic decision that aligns intellectual property protection with broader institutional goals and success. For universities advancing GenAI-assisted innovations, this pathway not only safeguards groundbreaking research but also establishes avenues for impactful commercialization and societal contribution. By thoughtfully navigating the complexities of both patent and non-patent strategies as outlined in the framework, universities can strengthen their position as leaders in innovation and technology transfer, while also setting powerful precedents that may redefine the future of discovery and impact generations to come.

DECISION TO PURSUE NON-PATENTING PROTECTION

Despite an invention qualifying for patent protection, universities may still choose to pursue non-patent protection strategies, as either the innovation does not meet the stringent requirements for patent eligibility or because the costs associated with applying for and defending patents simply outweighs the perceived benefits. As Vijay Pande, a general partner at Andreessen Horowitz and a Stanford biologist, aptly notes in the Verseon article, "in Silicon Valley, we talk about software patents as being the sport of kings," highlighting the financial and logistical barriers inherent to the patenting process. For universities operating under constrained budget and competing with well-funded industry players, alternative approaches to securing competitive advantages are not just pragmatic but often essential. Thus, this section examines non-patent protection strategies as means to safeguard GenAI innovations and create value in a rapidly evolving technological landscape. The focus shifts to leveraging model architecture, infrastructure, and other critical elements that underpin GenAI systems. These strategies are particularly relevant in domains where industry leaders currently dominate, offering universities an opportunity to establish their own niche by capitalizing on unique institutional strengths. For example, NYUTron—a large language model (LLM) trained on a decade's worth of clinical notes from NYU Langone Health—demonstrates how universities can harness proprietary datasets and tailored model architectures to achieve breakthroughs in specialized fields such as healthcare (Oermann 2025). The mechanisms by which this model has harnessed such power are to be examined in the forthcoming analysis. By focusing on strategic data approaches and positioning, this section explores how strategies such as prompt engineering, speech-to-market, and data quality enhancement, can not only complement patent strategies but provide standalone pathways for innovation protection and commercialization. In an era where the pace of discovery often

outstrips the ability of legal frameworks to adapt, non-patent protection offers universities a flexible and dynamic toolkit to remain at the forefront of GenAI research while navigating the challenges of resource constraints and industry competition.

Strategic Data Approach

Capital Investment Asymmetry and Alternative Pathways

The GenAI landscape is characterized by a significant capital investment asymmetry between the predominant technology firms and academic institutions. Companies such as Microsoft, Amazon, and Google possess the financial resources to construct large data centers and procure extensive arrays of graphics processing units (GPUs), enabling them to develop and train large-scale models with capabilities that smaller entities cannot match through direct competition. This capital-intensive barrier presents a fundamental challenge for universities, which typically operate with more constrained budgets and infrastructure limitations. In response to this asymmetry, universities must pursue alternative strategic pathways, and two promising approaches have emerged: computational efficiency and specialized small language models (SLMs) tailored to individual research domains.

The computational efficiency approach recognizes that perhaps brute-force computing power may not be the only path to innovation. As research from Carnegie Mellon University highlights, “creating a single image can use as much energy as is needed for a full smartphone charge” (Heikkilä 2023). The environmental and economic cost this poses creates an opportunity for universities to distinguish themselves through more efficient algorithmic designs. This shift is exemplified by DeepSeek, whose DeepSeek-V3 model achieves performance comparable to

leading models such as ChatGPT while using 10-40 times less energy through innovations such as FP8 mixed-precision training, dynamic expert routing architectures, and latent space compression (Varshney 2025; Gupta 2025). DeepSeek's focus on efficiency using only the most relevant sub-networks for each task, demonstrates that smarter, not just larger, models can deliver state-of-the-art results while drastically reducing resource consumption. Penn State researchers have also demonstrated the viability of this approach, developing a "selective learning" methodology that achieved "a 435-fold reduction in compute power in comparison to typical neural networks" (Mcllwain 2025). By focusing on developing computational efficiency, universities can establish intellectual assets that derive value from their elegant design as opposed to their resource requirements.

Aligned with the goal of efficiency, the second strategic pathway involves developing SLMs rather than attempting to compete with industry-scale LLMs. As Marc Sedam observes, "the place where universities are bringing value [are] small language models…bespoke applications of AI…[to find] a solution and driving down to make that solution work." This approach capitalizes upon one of academia's fundamental strengths: depth of domain knowledge within specialized fields. SLMs offer several distinct advantages that align with university priorities and constraints. From a resource perspective, these models require significantly less computational power, making them more cost-effective to develop and maintain within typical university environments. This would also allow for greater experimentation and iteration, thus enabling research teams to test multiple approaches without prohibitive costs. Beyond mere cost considerations, SLMs also provide greater transparency in their reasoning processes due to their more limited parameter space. This interpretability proves invaluable for research applications

where understanding the model's decision-making process is essential for scientific validity and ethical implementation. Perhaps most significantly, SLMs can be narrowly tailored to specific disciplines and trained on highly specialized datasets. Universities can then create models that in some instances would outperform general-purpose systems for specific applications, while adhering to the ethical and data security principles necessary within academic institutions. The focused nature of these models also facilitates clearer documentation of human contributions to the development process, potentially strengthening the patentability claims should the inventors ever choose to pursue that path.

However, the value of SLMs extends beyond theoretical advantages to practical applications, particularly as proof-of-concept tools for novel research directions (Blondeau and Moroney 2024). Particularly in the GenAI field, where universal acceptance of the notion remains incomplete, institutions can utilize SLMs to quantify potential returns on investments and effectively showcase the value proposition to skeptical stakeholders. This approach proves especially valuable in research environments like César de la Fuente's laboratory at the University of Pennsylvania, where artificial intelligence generates novel peptide sequences with potential antimicrobial properties, effectively exploring what researchers term the "microbial dark matter" for new antibiotic candidates. Such pioneering research can begin with relatively modest computational resources using specialized SLMs, with initial promising results justifying subsequent scaling to more resource-intensive approaches (Penn Medicine 2024).

Through strategic development of computational efficiency and specialized SLMs, universities can establish competitive advantages in the GenAI landscape despite resource limitations

compared to industry leaders. These methods leverage core academic strengths—theoretical innovation, domain expertise, and specialized knowledge—while acknowledging the practical constraints of university computing environments. By pursuing these strategic data approaches, universities can develop valuable intellectual assets that complement or substitute for traditional patent protection while advancing their research missions.

Data Quality Enhancement

In the instance that universities are unable to match the capital investment or achieve the same level of efficiency as industry leaders, they can still gain a competitive advantage by focusing on the quality of the data used to train the models by which innovations are generated. Marc Sedam underscores the critical importance of data in determining the success in AI systems, stating that "algorithms are rendered worthless, but highly curated datasets are incredibly valuable." This observation underscores how universities can prioritize the development of high-quality datasets to compensate for resource disparities. Training a language model involves multiple stages of data utilization, each with unique demands and costs. Jonathan Frankle, Chief Scientist at Databricks and Founder of MosaicAI, explains that pre-training alone requires billions of general data examples. Following this, domain-specific pre-training narrows the focus to millions of examples, while supervised fine-tuning (SFT) for instruction following involves tens of thousands of high-quality examples. The costs associated with this process, per his determination, can be significant: a single instruction-response pair costs $30–$80, while multi-turn chat conversations can cost $100–$200 each. Reinforcement Learning with Human Feedback (RLHF) adds further complexity and expense by requiring human evaluators to refine model outputs. Given these constraints, universities must carefully consider their data strategies. From a high level, key decisions include selecting high-quality datasets over low-quality options,

deduplicating data to reduce redundancy, and optimizing sequence lengths to balance memory usage and processing efficiency. Technical considerations such as positional encodings, optimizers, activation functions, tokenizers, and attention weights also play critical roles in ensuring model performance. For instance, the sequence length—the maximum number of tokens a model can process at once—requires careful calibration to avoid excessive memory demands while maintaining accuracy. Similarly, attention weights dictate how much importance is assigned to different input sources, directly influencing the model's effectiveness.

Given the aforementioned discussion of resource constraints, a potential strategy for universities is the formation of data consortia. Collaborative initiatives such as NextGenAI—a consortium comprising 15 leading research institutions—demonstrate how universities can pool resources to gain access to additional AI tools, research funding, and computational powers (OpenAI 2025).[13]The SEC Artificial Intelligence Consortium offers another example of collaboration among member universities by providing shared access to curricular materials and online presentations (SEC n.d.). Similarly, the American Council on Education's proposed Global Data Consortium aims to create a decentralized data architecture (Ashburn 2024). It is imperative to note that while data consortia offer significant advantages, such as cost-effective exploration of personalized learning applications and the creation of learning libraries, they also present substantial challenges. Privacy concerns within an incredibly digital environment remain paramount, and failures in one institution's security could compromise data integrity across the entire consortium. Additionally, technical barriers such as the lack of standard formats or

[13] The 15 leading research institutions are Caltech, the California State University system, Duke University, the University of Georgia, Harvard University, Howard University, Massachusetts Institute of Technology, the University of Michigan, the University of Mississippi, Ohio State University, the University of Oxford, Sciences Po, Texas A&M University, Boston Children's Hospital, and Boston Public Library.

potential misinterpretation of shared data could hinder collaboration. Lastly, sharing proprietary datasets may dilute individual institutions' competitive advantages. For example, NYU Langone's billion-point hospital dataset drives its GenAI healthcare innovations; sharing this dataset within a consortium could reduce its exclusivity while benefiting broader advancements in medical AI. The consortium approach essentially mirrors cross-licensing strategies with respect to data sharing, and collaboration would allow smaller entities to compete against the dominant players. For universities facing the tech giants, whose control over platforms like Instagram and Facebook provide them with an exponentially large repertoire of data scraping possibilities, combining collective power could significantly advance their respective GenAI research agendas.

Strategic Positioning

Having explored the foundational architecture approaches as means for universities to enhance their GenAI capabilities, the focus now shifts to how institutions can position themselves to maximize the impact and utility of their innovations. The discussions surrounding prompt engineering center on empowering researchers within universities to maximize the potential of the AI systems assisting their inventions. In contrast, speed-to-market focuses on the role of technology transfer offices in executing strategies for introducing innovations to the market. Internal capacity building occupies the middle ground, encompassing both the development of appropriate skills across the research community and the cultivation of an innovation-oriented institutional culture that fosters either patentable outcomes or competitive advantages.

Prompt Engineering

Prompt engineering, a process akin to refining search filters to yield more tailored results, involves crafting precise and structured inputs for GenAI systems with the intention of optimizing outputs. This practice has gained traction across industries, with McKinsey's Quantum Black reporting that 7% of organizations adopting AI are already hiring for prompt engineering roles—a figure expected to grow as the field matures (Segel and Hatami 2023). For universities, prompt engineering offers a unique opportunity to create a proprietary arsenal of inputs that maintain control over their intellectual assets. This approach establishes a zone of exclusivity comparable, though not exact, to traditional patenting, as the specialized prompts yield results of a quality and specificity that cannot be readily replicated without access to the carefully engineered input patterns themselves. Universities have begun recognizing the strategic value of this practice, as evidenced by seven proprietary in-house prompts developed by NYU's tech transfer office, enabling their compliance manager, John Keary, to review an entire year's worth of license agreements in merely three days. Similarly, their medical school counterparts created a 720-page prompt over six months with a dedicated team, demonstrating the depth and complexity involved in crafting effective prompts (Keary 2024).

A pivotal question previously addressed in the literature review resurfaces here: could prompts eventually attain patent-eligible status similar to software? Current legal frameworks indicate that prompts, while not patentable, could be protected under copyright. Though the nuances of copyright law exceed this inquiry's scope, this position aligns with Sedam's argument that prompts constitute physical manifestations of written ideas—precisely what copyright law is designed to protect. Regardless of formal intellectual property classification, prompt

engineering's fundamental advantage lies in its capacity to generate institution-specific frameworks that competitors cannot easily replicate. For instance, UC Davis has outlined a "CLEAR" framework for effective prompt engineering: prompts should be concise, logical, explicit, adaptive, and reflective (Lo 2023). This structured approach ensures clarity and precision while allowing for continuous refinement based on feedback—a critical component for achieving actionable results. Although this specific framework may not be universally applicable across all higher education institutions, developing a comparable structure is advantageous as it results in university-wide yet customized methodologies, and at its core promotes a consideration of prompt engineering from the outset. Taking this further, universities may consider creating national standards for prompt engineering practices, as such frameworks would not only enhance consistency but also ensure that human judgement and reflectiveness are embedded within the process. This is particularly important for GenAI-related inventions where ethical considerations and interpretability play significant roles. In a similar manner to SLMs, prompt engineering also offers advantages with regards to scalability. For example, researchers can use refined prompts to generate preliminary outputs or test hypotheses prior to full-scale model development.

Speed-to-Market Considerations

As mentioned above, speed to market represents a complementary strategy focused on the timing and execution of introducing innovations to the public or industry. The decision of when and how to disclose an invention or its underlying model can significantly impact its competitive advantage, intellectual property protection, and broader contribution. The dichotomy between open science principles and commercialization imperatives creates a pronounced tension in this context, presenting universities with two divergent pathways. Universities can choose to release their inventions as open source, thereby establishing themselves as the originators of the

innovation while simultaneously preventing others from patenting the innovation by creating prior art. Alternatively, institutions may opt to retain exclusivity for a period, using the invention internally to refine its capabilities or gain competitive advantage prior to public disclosure. Research on first-mover advantage underscores the importance of timing in innovation as early disclosure can provide technology leadership by making products or services harder for later entrants to replicate. Economists refer to this as an “absolute cost advantage”, where first movers can control resources such as equipment, data, and methodologies that create barriers for competitors (CFI Team n.d.). In university research, this advantage manifests through control over both tangible assets—such as proprietary datasets—and intangible assets like specialized expertise and methodologies. This notion becomes complicated within the GenAI space with the onset of advanced reverse engineering technologies.

The second pathway involving deliberate publication to establish prior art, stems from theoretical grounding in intellectual property law. As Talia Bar, an Associate Professor at the University of Connecticut researched, publishing details of research findings creates prior art that strategically affects the patentability of related innovations (Bar 2006). This approach allows universities to maintain their commitment to open science while simultaneously safeguarding their inventions against competitors seeking patent protection for similar technologies. Alternatively, universities may choose to delay public disclosure and utilize their innovations internally for a period of time prior to releasing. Referring back to NYU's GenAI models, NYU Langone strategically delayed licensing opportunities and immediate disclosure, recognizing the competitive advantage conferred by maintaining exclusivity during the capability refinement process. As evidenced in one internal discussion, the sentiment "do [we] really want to tell everyone what [we] are doing

right now?" reflects this deliberate approach to information management and competitive positioning.[14]

Overall, speed-to-market strategies allow higher education institutions to navigate the contradictions of academia and industry effectively. By integrating these approaches into broader institutional frameworks, technology transfer offices can ensure that university-developed GenAI innovations achieve both impactful knowledge contributions and sustainable commercial success.

Internal Capacity Building

The final pillar of the strategic positioning strategies is internal capacity building, a cornerstone of non-patent strategies aimed at maximizing impact and fostering innovation. As Robert Plotkin aptly describes, “[if] you imagine a company with a thousand people, if two people come up with great prompts, but the company doesn’t have a way of distributing them and training other people on how to use them, it’s going to be those two people sitting at their desk using those prompts over and over…it’s not going to provide widespread value to the company.” This highlights the critical importance of institutionalizing knowledge and expertise rather than allowing it to remain siloed among a few individuals. For universities, this means a system in which GenAI tools and methodologies are broadly accessible across departments and disciplines.

Arizona State University exemplifies effective internal capacity building through its "institution-wide AI Innovation Challenge," which engaged faculty, staff, and students in conceptualizing

[14] Marc Sedam, in conversation with the author, December 2024

novel applications for GenAI technologies. By incorporating diverse stakeholders in the discovery process, ASU expanded GenAI adoption beyond specialized laboratories and expert offices. Dr. Elizabeth Reilley from ASU emphasizes that effective capacity building need not require "large-scale partnerships or a team of highly skilled AI engineers," but rather benefits from empowering faculty and staff to spearhead innovation initiatives (Reilley 2025).

This internal empowerment approach addresses a critical challenge in academic settings: despite substantial annual investments in computational resources and talent development, universities frequently witness their expertise being absorbed by industry through recruitment. By fostering broad-based participation and creating meaningful engagement opportunities with cutting-edge technologies, institutions can develop robust internal capacity that not only cultivates talent but actively preserves it within the academic ecosystem, creating sustainable innovation capabilities rather than merely serving as industry talent pipelines.

The decision to pursue non-patent protection strategies represents a sophisticated approach for universities seeking to safeguard their GenAI innovations amid resource constraints and intensifying competition. Universities can capitalize on their distinctive strengths in theoretical advancement and domain expertise through strategic data initiatives—including computational optimization, purpose-built small language models (SLMs), and enhanced data quality frameworks. These approaches enable academic institutions to develop valuable intellectual assets that serve as effective alternatives or complements to conventional patent protection mechanisms. Strategic positioning is able to amplify these initiatives by utilizing methodologies that optimize the functionality of GenAI tools. Refined prompt engineering techniques empower

researchers to extract maximum value through precision-crafted inputs, while accelerated commercialization pathways enable technology transfer offices to adeptly balance open science principles with market imperatives. Comprehensive internal capacity building transforms isolated expertise into institutional knowledge, facilitating the pervasive integration of GenAI technologies throughout university ecosystems. Together, these non-patent protection strategies provide universities with a versatile and progressive framework for action. By adopting this multifaceted approach, academic institutions can maintain their position at the forefront of GenAI innovation while effectively safeguarding their intellectual contributions.

SYNTHEMOL CASE STUDY

Building upon the comprehensive framework for university GenAI protection strategies, this inquiry now turns to a revelatory case study that demonstrates these principles. The case that served as the impetus for this research project is SyntheMol, a groundbreaking GenAI platform developed in collaboration between Stanford University and McMaster University. Designed to address the urgent global challenges of antibiotic resistance, SyntheMol's model is able to generate novel antibiotic molecules, including six compounds effective against A. *baumannii,* a drug-resistant pathogen identified as a critical threat by the World Health Organization. By leveraging a library of 132,000 molecular fragments and integrating property prediction models, SyntheMol not only designs entirely new molecules but also provides detailed synthetic pathways for their production—an innovation that bridges computational chemistry with practical laboratory applications.

SyntheMol is particularly compelling as a case study as it encapsulates many of the theoretical and practical challenges discussed in this thesis, including the balance between human and AI contributions, intellectual property considerations, and the role of data transparency in upholding the dual mandate of universities. Therefore, the following section draws from direct correspondence with Kyle Swanson, one of Synthemol's lead developers, providing valuable insights into how academic institutions navigate the complex terrain of AI innovation protection.

The key questions surrounding Stanford's SyntheMol research center on the role of human contribution and its implications for IP strategies. A critical consideration was whether the outputs generated by SyntheMol could be patented and, if so, how the human element could be substantiated within an AI-driven workflow. Additionally, this inquiry wished to examine whether Stanford had provided institutional guidance or policies that influenced the team's approach to IP protection.

The SyntheMol system exemplifies the intricate balance between algorithmic creativity and human guidance in generative AI innovation. As Swanson emphasized, "human involvement is crucial, even if the core design process is done by AI." SyntheMol's model process begins with human-curated training data: a set of approximately 13,000 molecules tested against A. *baumannii,* assembled entirely by wet-lab researchers. Critical human judgment was required in establishing activity thresholds for classifying molecules as effective or ineffective against bacteria, decisions Swanson notes were "based both on our prior experience and based on the scientific literature."

The development of SyntheMol's property predictor—a graph neural network—further illustrates the collaborative nature of AI development. While model training was "mostly automated", Swanson acknowledges that researchers made essential decisions pertaining to hyperparameters, including "how many models to train in our ensemble and whether to train a regression or classification model." These choices reflect what Swanson describes as "high-level decisions about how SyntheMol would operate," highlighting the human agency required to steer AI tools towards meaningful outcomes. Researchers made critical choices regarding iteration counts and the balance between "exploiting known good molecular building blocks vs exploring new ones." Additionally, they established computational filtering rules to identify promising candidates for synthesis and validation. Swanson notes that while "the filtering rules were all computational," their design was entirely human driven. Even the chemical space SyntheMol operates within—the Enamine REAL Space—is algorithmically enumerated but rooted in rules devised by human chemists. The final phases of synthesis and validation remained "entirely human endeavors without involving AI." This layered involvement aligns with USPTO guidance requiring "significant human contribution" for patent eligibility, positioning SyntheMol as a tool rather than an autonomous inventor.

When asked how SyntheMol-generated molecules might satisfy federal patent requirements, Swanson reaffirmed the centrality of human agency: "Human expertise is particularly crucial in determining what disease to go after, what training data to use… and how to interpret and select the molecules generated."

He positions Synthemol as "a tool that helps scientists more rapidly identify good drug candidates," comparing it to other computational chemistry tools where "it is still the scientist using Synthemol that is responsible for discovering the drug due to all the decisions they make before and after using Synthemol." This perspective mirrors legal arguments that AI-assisted inventions meet patent thresholds when human decisions ensure novelty, non-obviousness, and enablement. While Swanson acknowledges that Stanford has filed patent applications for SyntheMol's algorithms and drug candidates, he clarifies that institutional policies on AI-generated inventions remain opaque—a reflection of broader ambiguities in GenAI patent strategy.

This uncertainty underscores a key tension in AI-driven academic research: even as tools like SyntheMol enhance discovery efficiency, universities must navigate evolving legal standards that prioritize human inventorship. Swanson's characterization of SyntheMol as a "tool that helps scientists more rapidly identify good drug candidates" reinforces the thesis argument that AI systems function as collaborators rather than replacements for human ingenuity.

Regarding SyntheMol's approach to data transparency, the team adopted a notably accessible approach to data and code sharing—a strategy that contrasts with defensive non-disclosure approaches observed in other university AI innovations. "We are very transparent regarding SyntheMol," Swanson explained, "all of our data and property predictor model weights for our antibiotic application are available... on Zenodo" and "all of our code is available... on GitHub." Swanson described the platform as "disease-agnostic," requiring only disease-specific training data in CSV format—a design choice that prioritizes broad academic and industry adoption. This

open-source strategy aligns with the consortium model discussed earlier, where shared resources accelerate innovation but risk diluting competitive advantages. However, SyntheMol's transparency includes trade-offs. While public accessibility enhances Stanford's reputation as a leader in ethical AI research, it may limit opportunities for exclusive commercialization. In this instance, Swanson's team appears to prioritize societal impact over proprietary control, stating, "we hope that researchers will use SyntheMol for their own drug discovery applications."

Looking forward to future developments, Swanson envisions an evolving landscape of human-AI collaboration in drug discovery. While acknowledging that research is "still guided primarily by humans," he anticipates that "over time, more of those decisions could potentially be made with AI models such as large language models (LLMs)." The team has already begun exploring this frontier, having released a preprint demonstrating "how LLMs can make more of those research decisions, at least in conjunction with human researchers." Nevertheless, Swnason maintains that "there will likely still be a big role for humans to play in drug discovery in terms of decisions such as what diseases to prioritize and what types of models are worth exploring."

Overall, SyntheMol's development crystallizes key themes of this thesis. It reflects an institutional strategic evaluation of the mission of the individual researchers as well as Stanford's values of integrating AI into traditional research workflows without relying on large-scale industry partnerships. The detailed mapping of researcher involvement throughout the development process showcases practical mechanisms universities can adopt to satisfy the "significant human contribution" threshold required for patent eligibility. With respect to the decision to pursue patenting, Stanford appears to have chosen both paths simultaneously:

according to Swanson, who specializes in the computational element of SyntheMol, the team has filed patent applications for the SyntheMol algorithm and some of the drug candidates discovered using it, while also making the data and code for SyntheMol openly available. The ultimate outcome of Stanford's patent applications has yet to come to fruition but will likely set a highly valuable precedent for higher education institutions. Their strategic positioning prioritized speed-to-market by establishing first-mover advantage in the scientific community while building internal capacity through interdisciplinary collaboration. However, making the code open-source may complicate the enforceability of any algorithm-related patents, since users could argue, they are licensed to use the algorithm under the terms of the open-source release. Finally, SyntheMol's approach reflects careful industry-specific adaptation for pharmaceutical development, where drug candidates themselves hold tremendous value, while maintaining external alignment with Stanford's ethical principles by addressing antibiotic resistance—a global health priority—and openly sharing their methodology to advance scientific knowledge.

AREAS FOR FURTHER RESEARCH AND CONTEXTUAL CONSIDERATIONS

This study's theoretical framework offers a starting point for understanding university approaches to GenAI innovation protection, however, several critical contextual dimensions warrant deeper exploration. The final component of the proposed framework—contextual analysis—offers rich territory for subsequent research that could significantly enhance our understanding of how different environments shape patenting decisions and strategies.

Industry-Specific Implications

While universities were deliberately chosen as the focus of this study due to the diverse range of innovations they produce, future research should explore how specific industry contexts influence GenAI patenting decisions. The pharmaceutical and biotechnology sectors have received disproportionate attention in current literature, largely as these fields historically generate substantial licensing revenue for academic institutions. However, university innovations span numerous other domains—from agriculture to renewable energy, transportation to educational technology—each with distinct market dynamics, regulatory environments, and commercialization pathways. For instance, software-based GenAI innovations face substantially different patentability challenges than hardware-integrated systems Research examining how these industry-specific factors influence patent strategy could provide valuable differentiated guidance for technology transfer officers operating across a multitude of fields.

Global Competitive Landscape

This research focused primarily on US universities due to information accessibility and regulatory familiarity, however, Chinese academic institutions have emerged as particularly prolific in GenAI patenting. There exists a high likelihood that their substantially different approaches to intellectual property protection, data sharing, and university-industry collaboration could provide valuable insights for Western universities seeking to maintain competitiveness internationally. Comparative case studies could address several insightful questions: How do varying national intellectual property regimes influence university GenAI patenting strategies? What cultural or institutional factors shape attitudes toward open science v. commercialization?

How do differences in government funding models impact relevant incentives for patent protection?

Judicial Clarification and Precedent Development

Perhaps the most significant limitation of current research is the absence of substantial litigation history specific to GenAI patents. While this thesis provides a framework based on existing legal principles and early patent office guidance, the true boundaries of patentability will be ultimately defined through judicial interpretation. As Swanson's experience with SyntheMol demonstrates, universities are actively pursuing patents for both GenAI systems and their outputs, but the enforceability and scope of these patents remains untested. In a similar vein, longitudinal research tracking key GenAI patent cases through the court system would provide invaluable insights into how judges interpret the notion of human contribution, how prior art is evaluated in ever changing technological domains, and what remedies would be available for infringement. While this inquiry primarily focused on business and strategic considerations to align with the decision-making frameworks employed by firms, incorporating a legal perspective could provide valuable insights. A systematic analysis of court decisions to identify patterns in judicial reasoning would complement the strategic lens and offer a deeper understanding of how patent law evolves in response to generative AI innovations.

CONCLUSION

As GenAI fundamentally reshapes innovation processes across disciplines, universities face an unprecedented strategic challenge: how to protect and leverage AI-assisted discoveries while fulfilling their core missions of knowledge creation and dissemination. This research has

responded to this challenge by developing a comprehensive analytical framework that transcends the limitations of existing approaches, which have typically treated AI innovations through conventional patenting lenses without consideration for their unique characteristics and institutional contexts.

Rather than prescribing universal solutions, this framework introduces a structured decision-making methodology that systematically balances varying priorities and attributes. The hierarchical decision structure presented in this thesis bridges a critical gap in both theoretical literature and practical guidance by recognizing that patent protection, while valuable, may not always represent the optimal strategy for universities. By incorporating alternative mechanisms—from strategic data approaches to capacity building initiatives—this framework enables universities to align their intellectual property strategies with their specific institutional values and objectives.

The SyntheMol case study demonstrates the framework's practical utility, illustrating how human contributions can be methodically identified throughout the innovation process while strategically positioning assets for maximum impact and value creation. This application validates the framework's adaptability to complex real-world scenarios while highlighting its potential to resolve the tensions inherent in university-based AI innovation.

This research is not intended to provide static solutions but rather to catalyze an evolution in how academic institutions approach emerging technologies. As universities incorporate these tools into their research enterprises, they can utilize this framework as a foundation for evidence-based

policy development, institutional strategy formulation, and cross-disciplinary collaboration. The framework's true value will emerge as institutions adapt it to address evolving challenges—particularly as judicial decisions clarify the boundaries of AI patentability and as global innovation leaders shape new approaches to intellectual property protection.

In an era where the boundaries between human and machine contributions are increasingly indistinct, this research ultimately reaffirms the irreplaceable role of universities as architects of responsible technological advancement. By highlighting the distinctly human elements that persist even in AI-assisted discovery, this thesis underscores the importance of academic institutions in balancing innovation protection with their broader societal missions. This work invites universities not merely to react to technological change, but to actively shape how generative AI transforms the creation and dissemination of knowledge—ensuring that these powerful tools serve as catalysts for equitable progress in a landscape of unprecedented technological potential.

REFERENCES

Alemán, Marco, Alejandro R. Campaña, Andrew Czajkowski, and Christopher Harrison. 2024. "Generative Artificial Intelligence: Patent Landscape Report." WIPO. https://www.wipo.int/en/web/patent-analytics/generative-ai.

Atkinson, Robert D., and K. Gawora. 2021. "U.S. University R&D Funding Falls Further Behind OECD Peers." Information Technology and Innovation Foundation, April 12. https://itif.org/publications/2021/04/12/us-university-rd-funding-falls-further-behind-oecd-peers/.

Azhar, Ali. 2024. "New AI Patent Study Shows IBM Leads Google and Microsoft in GenAI Race." AI Wire. https://www.aiwire.net/2024/02/07/new-ai-patent-study-shows-ibm-leads-google-and-microsoft-in-genai-race/.

Bar, T. 2006. "Defensive Publications in an R&D Race." Journal of Economics & Management Strategy 15 (1): 229–254. https://doi.org/10.1111/j.1530-9134.2006.00099.x.

Billington, Stephen, Alan Hanna, and Joe Lane. 2024. "What can we learn about patents and innovation from the past?" Economics Observatory, January 30. https://www.economicsobservatory.com/what-can-we-learn-about-patents-and-innovation-from-the-past.

Blondeau, Harrison, and John Moroney. 2024. “Government AI: Starting Small with Small Language Models (SLMs).” Elder Research. https://www.elderresearch.com/blog/government-ai-starting-small-with-small-language-models-slms/.

Blue Spike, LLC v. Google Inc., No. 16-1054 (Fed. Cir. 2016). https://law.justia.com/cases/federal/appellate-courts/cafc/16-1054/16-1054-2016-10-14.html.

Board of Trustees of the Leland Stanford Junior Univ. v. Roche Molecular Systems, Inc., 563 U.S. 776 (2011).

Brooks, E., & Crossan, C. 2023. “Generative Ai for Strategy and Innovation.” LinkedIn, August 30. https://www.linkedin.com/pulse/managementgpt-empowering-every-manager-co-thinker-paolo-cervini/?trackingId=gYGI5TD%2BSQSWlLKHuc7aKg%3D%3D.

Buvailo, Andrii. 2023. “Insilico Medicine’s Generative AI Patent Provides Advantage in AI Drug Discovery Race.” BioPharmaTrend. https://www.biopharmatrend.com/post/631-insilico-medicines-generative-ai-patent-provides-advantage-in-ai-drug-discovery-race/.

Caulfield, Timothy, Shawn H. Harmon, and Yann Joly. 2012. "Open science versus commercialization: a modern research conflict?" Genome Medicine 4, no. 17. https://doi.org/10.1186/gm316.

Castaros, O., & Bloomberg Intelligence. 2023. “Generative AI to Become a $1.3 Trillion Market by 2032, Research Finds.” Bloomberg, June 1. https://www.bloomberg.com/company/press/generative-ai-to-become-a-1-3-trillion-market-by-2032-research-finds/.

CFI Team. n.d. “First Mover Advantage - Benefits and Drawbacks of Being First.” Corporate Finance Institute. Accessed April 27, 2025. https://corporatefinanceinstitute.com/resources/management/first-mover-advantage/.

Chakraborty, Dipanjana. 2025. “Cross-Licensing Agreements: A Strategic Tool to Minimise Patent Conflicts.” De Penning & De Penning. https://depenning.com/blog/cross-licensing-agreements-a-strategic-tool-to-minimise-patent-conflicts/.

Chapman, Audrey R. 2011. "The right to enjoy the benefits of scientific progress." Health and Human Rights 13, no. 2: 1-8.

Chen, A. 2024. “University of Chicago professor, PhD students create tools to protect art from generative AI.” ABC Chicago, January 24. https://abc7chicago.com/ai-generator-generative-art-university-of-chicago/14357524/.

Cheong, Ben C. 2024. “Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making.” Frontiers.

https://www.frontiersin.org/journals/human-dynamics/articles/10.3389/fhumd.2024.1421273/full.

Chien, Colleen V. 2010. "From Arms Race to Marketplace: The New Complex Patent Ecosystem and Its Implications for the Patent System." Hastings Law Journal 62: 297. https://ssrn.com/abstract=1703557.

Cockburn, Iain, Rebecca Henderson, and Scott Stern. 2018. "The impact of artificial intelligence on innovation." National Bureau of Economic Research. http://dx.doi.org/10.3386/w24449.

DeepSeek's Rise: Chinese AI Startup Challenges US Dominance. 2024. The Outpost, February 10. https://theoutpost.ai/news-story/deep-seek-s-rise-chinese-ai-startup-challenges-us-dominance-11035/.

Drexel University. n.d. "Bayh-Dole Act - Office of Research & Innovation." Drexel University. https://drexel.edu/research/innovation/technology-commercialization/bayh-dole-act/.

Duke University. 2015. "Policy on Inventions, Patents, and Technology Transfer." October draft.

Einfolge. 2023. "Generative AI and Patent Trends: Unlocking Innovation." https://www.einfolge.com/blog/generative-ai-and-patent-trends-unlocking-innovation.

Ezell, Stephen. 2019. "The Bayh-Dole Act's Vital Importance to the U.S. Life-Sciences Innovation System." Information Technology and Innovation Foundation. https://itif.org/sites/default/files/2019-bayh-dole-act.pdf.

Fisher, Forsgren. n.d. "Applying Step Two of the Alice Test (Bitlaw Guidance)." BitLaw. Accessed April 27, 2025. https://www.bitlaw.com/guidance/patent/applying-step-two-of-Alice-Mayo-test.html.

FOSSA. 2025. "Open Source Licenses 101: Apache License 2.0." January 1. https://fossa.com/blog/open-source-licenses-101-apache-license-2-0/.

Fraser, J. 2024. "New Biden admin proposal legitimizes foreign IP theft." Post and Courier, March 6. https://www.postandcourier.com/kingstree/opinion/new-biden-admin-proposal-legitimizes-foreign-ip-theft/article_9106cc72-d711-11ee-b3e4-a34722e21688.html.

Georgia-Pacific Corp. v. United States Plywood Corp., 318 F. Supp. 1116 (S.D.N.Y. 1970).

Georgetown University. 2025. "Research and Innovation Showcase 2025." March 29. https://otc.georgetown.edu/newsevents/research-and-innovation-showcase-2025/.

HackMD. 2022. "TTW: Machine Learning Model Licenses." Last modified September 16. https://hackmd.io/@jending12/HyvMU8sJo.

Heikkilä, Melissa. 2023. “Making an image with generative AI uses as much energy as charging your phone.” MIT Technology Review. https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone/.

Herb Hovenkamp. 2024. Current and Emerging Issues in Law and Technology panel discussion at the University of Pennsylvania Carey Law School, September 24.

Hughes, Rose. 2025. “Insilico Medicine: Lessons in IP strategy from a front-runner in AI-drug discovery.” IPKat. https://ipkitten.blogspot.com/2025/02/insilico-medicine-lessons-in-ip.html.

IAM. 2017. “‘Smallest saleable patent practising unit’ doctrine: developments and challenges.” IAM. https://www.iam-media.com/global-guide/iam-yearbook/2018/article/smallest-saleable-patent-practising-unit-doctrine-developments-and-challenges.

IAM. 2018. “The secondary market in patents: What went right, what went wrong and how to fix it.” IAM. https://www.iam-media.com/global-guide/iam-yearbook/2018/article/the-secondary-market-in-patents-what-went-right-what-went-wrong-and-how-fix-it.

IAM. 2024. “Special Report 2024 Q4: New Dawn for Life Sciences IP Strategy.” IAM. https://www.iam-media.com/report/special-reports/q4-2024.

Insilico Medicine. 2025. “Insilico: linking target discovery and generative chemistry AI platforms for a drug discovery breakthrough.” BioPharmaDealMakers. https://www.nature.com/articles/d43747-021-00039-5.

Inventorship guidance for AI-assisted inventions webinar. 2024. USPTO. https://www.uspto.gov/about-us/events/inventorship-guidance-ai-assisted-inventions-webinar.

Jackson, Jr., John L., and Dawn Bonnell. 2023. “III.E. Patent and Tangible Research Property Policies and Procedures < University of Pennsylvania.” University Catalog. https://catalog.upenn.edu/faculty-handbook/iii/iii-e/.

John Keary, “Prompt Engineering for Technology Transfer Professionals,” Tech Transfer Talk (YouTube video, 32:22), March 25, 2024. https://www.youtube.com/watch?v=_yXt1QIKJwI.

John Swartley, conversation with the author, December 2024.

John Swartley, conversation with the author, December 2024.

Jonathan Frankle. 2024. “How to Train an LLM Model from Scratch.” Talk, February.

Klenow, Peter, and Andréas Rodríguez-Clare. 1997. “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?” NBER Macroeconomics Annual 12: 73–114.

Label Your Data. 2025. "Small Language Models: Balancing Size and Performance." January 7. https://labelyourdata.com/articles/small-language-models.

Larry Jameson, email to the University of Pennsylvania community, March 26, 2025.

Lee, Ho Yeon, and Raymis H. Kim. 2024. “High Court Reaffirms Artificial Intelligence (AI) Cannot Be an Inventor in Korea.” Kim & Chang. https://www.ip.kimchang.com/en/insights/detail.kc?sch_section=4&idx=29908.

Lerman, Craig. 2015. “Patent strategies of technology startups: An empirical study.” SSRN Electronic Journal. https://doi.org/10.2139/ssrn.2610433.

Licenses.ai. 2023. "AI Pubs RAIL Licenses." March 9. https://www.licenses.ai/blog/2023/3/3/ai-pubs-rail-licenses.

Lo, Leo, S. 2023. "The CLEAR path: A Framework for Enhancing Information Literacy through Prompt Engineering." The Journal of Academic Librarianship, 49(4).

Marr, K., and P. Phan. 2020. "The valorization of non-patent intellectual property in academic medical centers." Journal of Technology Transfer 45 (6): 1823–1841. https://doi.org/10.1007/s10961-020-09827-0.

Matsumura, Hajime, and Fumio Takahasi. 2025. “IP High Court Case Regarding Patent Inventorship: Patent Inventor Must Be Human.” Shiga International Patent Office. https://shigapatent.com/en/topics/iphc_dabus/.

McIlwain, Mallorie. 2025. “Q&A: Researchers’ approach aims to make AI-powered systems more efficient.” Penn State University. https://www.psu.edu/news/institute-computational-and-data-sciences/story/qa-researchers-approach-aims-make-ai-powered.

Michigan Technological University. 2018. "Faculty Handbook: Role of Innovation and Commercialization." Revised Senate Proposal 5-18.

MIT FutureTech. 2025. "What Drives Progress in AI? Trends in Compute." January 3. https://futuretech.mit.edu/news/what-drives-progress-in-ai-trends-in-compute.

NESTA. 2009. The Innovation Index: Measuring the UK’s Investment in Innovation and Its Effects. London: NESTA. http://www.nesta.org.uk/library/documents/innovation-index.pdf.

Norman, Jeremy. n.d. “The Statute of Anne: The First Copyright Statute.” History of Information. Accessed April 26, 2025. https://www.historyofinformation.com/detail.php?entryid=3389.

NSF Announces $140 Million Investment In Seven Artificial Intelligence Research Institutes. 2023. University of Maryland, June 20. https://www.umd.edu/news/nsf-announces-140-million-investment-seven-artificial-intelligence-research-institutes.

NYU Langone Health. 2025. "NYUTron: Health System Scale Language Models Are All-purpose Prediction Engines." https://datacatalog.med.nyu.edu/dataset/10633.

Ocado Innovation Ltd. v. AutoStore, No. 20-cv-11060 (D.N.H. 2021).

Ohio State University. 2024. “Research and Innovation Showcase.” October 4. https://research.osu.edu/rishowcase.

Office of the Chief Economist USPTO. 2021. “Artificial Intelligence Patent Dataset.” USPTO. https://www.uspto.gov/ip-policy/economic-research/research-datasets/artificial-intelligence-patent-dataset.

Oermann, Eric. 2025. “NYUTron - AI in Healthcare - Specialty Guides at NYU Health Sciences Library.” NYU Health Sciences Library. https://hslguides.med.nyu.edu/aihsl/nyutron.

OpenAI. 2025. “Introducing NextGenAI: A consortium to advance research and education with AI.” https://openai.com/index/introducing-nextgenai/.

Pamplin College of Business in Association with Virginia Tech Publishing. 2020. "4.4 Intellectual property & isolating mechanisms." August 14.

Pannu v. Iolab Corp., 155 F.3d 1344 (Fed. Cir. 1998).

Penn Center for Innovation. 2024. "Penn's $18 Million AIRFoundry to Revolutionize RNA Research." LinkedIn, September 12. https://www.linkedin.com/posts/penncenterforinnovation_penns-18-million-airfoundry-to-revolutionize-activity-7240016592145678336-tzGx.

Penn Medicine. 2024. “Largest-ever antibiotic discovery effort uses AI to uncover potential cures in microbial dark matter.” Penn Medicine News. https://www.pennmedicine.org/news/news-releases/2024/june/penn-expert-uses-ai-to-uncover-potential-antibiotics-in-microbial-dark-matter.

Penn State to Pay $1.25M to Resolve Claims of Cybersecurity Non-Compliance. 2024. Insurance Journal, October 24. https://www.insurancejournal.com/news/east/2024/10/24/753890.htm.

Phelan, Ryan N. 2022. “EPO Proposes AI Machine Owner as Inventor for AI Inventions.” PatentNext. https://www.patentnext.com/2022/07/european-patent-office-epo-suggests-that-the-owner-of-an-artificial-intelligence-ai-machine-could-be-listed-as-the-inventor-of-an-ai-generated-invention/.

PJ Parker. n.d. “Licensing.” PJ Parker. https://pjparker.com/licensing/#:~:text=What%20is%20%22carrot%22%20licensing%3F,owner%2C%20a%20corresponding%20license%20fee.

Plotkin, Robert. 2025. “AI and the Level of Ordinary Skill: Why Patent Law Must (and Can) Adapt to AI-Augmented Invention.” IPWatchdog. https://ipwatchdog.com/2025/01/07/ai-level-ordinary-skill-patent-law-must-can-adapt-ai-augmented-invention/id=184822/.

Pressman, L., et al. 2015. "The Economic Contribution of University/Nonprofit Inventions in the United States: 1996-2013." BIO and ATUM. https://www.bio.org/sites/default/files/BIO_2015_Update_of_I-O_Eco_Imp.pdf.

President Obama Speaks about Startup America. 2011. National Archives and Records Administration, January 31. https://obamawhitehouse.archives.gov/economy/business/startup-america.

Raza, M. 2024. “LLMs vs. SLMs: The Differences in Large & Small Language Models.” Splunk, March 21. https://www.splunk.com/en_us/blog/learn/language-models-slm-vs-llm.html.

Ratner, Ed. 2023. “AI for drug discovery is booming, but who owns the patents?” Verseon. https://www.verseon.com/news-ai-drug-discovery-who-owns-the-patents.html#:~:text=%E2%80%9CEverybody%20patents%20molecules%2C%E2%80%9D%20says,algorithms%2C%20it's%20a%20choice.%E2%80%9D.

Recentive Analytics, Inc. v. Fox Corp., No. 23-2437 (Fed. Cir. Apr. 18, 2025). https://www.cafc.uscourts.gov/opinions-orders/23-2437.OPINION.4-18-2025_2500790.pdf.

Regents Policy 5105: Policy on Innovation Transfer and Entrepreneurship, University of California, May 13, 2021.

Reilley, Elizabeth. 2025. "Generating College Completion." Complete College America. https://completecollegeamerica.s3.us-east-2.amazonaws.com/wp-content/uploads/2025/01/14134637/CCA_GeneratingCollegeCompletion-1.pdf.

Sarkar, S. 2023. "Generative AI patent data shows Chinese universities' dominance." Managing IP, March 2. https://www.managingip.com/article/2bckchmnueeiqnhwujev4/generative-ai-patent-data-shows-chinese-universities-dominance.

SEC. n.d. "SEC Artificial Intelligence Consortium." https://www.thesecu.com/programs/sec-ai-consortium/.

Segel, Liz H., and Homayoun Hatami. 2023. "On Point." Mckinsey & Company. https://www.mckinsey.com/~/media/mckinsey/email/onpoint/2023/10/2023-10-18b.html.

Singh, E. n.d. "Top Universities Using AI to Accelerate Patenting Activity." Xlscout. https://xlscout.ai/top-universities-using-ai-to-accelerate-patenting-activity.

Singhi, S. 2023. “Penn PRECISE Center research shows how generative AI can be used in clinical practice.” The Daily Pennsylvanian, November 12. https://www.thedp.com/article/2023/11/penn-precise-center-ai-research-medicine.

Smits, J., & Borghuis, T. 2022. "Generative AI and intellectual property rights." In Law and Artificial Intelligence, 323–344. T.M.C. Asser Press. http://dx.doi.org/10.1007/978-94-6265-523-2_17.

Somaya, Deepak. 2002. “Theoretical Perspectives on Patent Strategy.” Robert H. Smith School of Business, University of Maryland, August.

Snyk. 2020. "Open Source Licenses: Types and Comparison." July 16. https://snyk.io/articles/open-source-licenses/.

Sosna, Joerg. 2024. “USPTO Issues 2024 Guidance Update on Patent Eligibility for AI Inventions.” Lexology. https://www.lexology.com/library/detail.aspx?g=309538ce-cc9d-48e0-b91c-c55a93170c1c.

Stallman, Richard M. 2021. “Releasing Free Software If You Work at a University - GNU Project - Free Software Foundation.” October 1. https://www.gnu.org/philosophy/university.en.html.

Steinberg, Bob. 2024. “Are AI Innovations Patent Eligible?” Generative IQ. https://generativeiq.com/are-ai-innovations-patent-eligible/.

Steven Weiner, conversation with the author, April 2025.

Subcommittee of the Association’s Committee A on Academic Freedom and Tenure. 2013. “Defending the Freedom to Innovate: Faculty Intellectual Property Rights after Stanford v. Roche.” AAUP. https://www.aaup.org/report/defending-freedom-innovate-faculty-intellectual-property-rights-after-stanford-v-roche-0.

Susalka, Stephen J. 2019. “AUTM Comments on Patenting Artificial Intelligence Inventions (Docket No. PTO-C-2019-0029).” AUTM. https://autm.net/AUTM/media/About-AUTM/Documents/AUTM-Comments-on-Patenting-Artificial-Intelligence-Inventions_11-8-19.pdf.

Thaler v. Vidal, No. 21-2347 (Fed. Cir. 2022).

The Turing Way Community. 2019. "Licensing Machine Learning Models." The Turing Way, January 1. https://book.the-turing-way.org/reproducible-research/licensing/licensing-ml.

Tompa, R. 2024. “Generative AI develops potential new drugs for antibiotic-resistant bacteria.” Stanford Medicine, March 28. https://med.stanford.edu/news/all-news/2024/03/ai-drug-development.html.

Tran, Bao. 2025. “Legal Considerations for Open Source Software in Deep Tech Patents.” Patent PC. https://patentpc.com/blog/legal-considerations-for-open-source-software-in-deep-tech-patents.

TRP International. 2025. “Global Patent Filing Slowdown? Using Patent Landscapes and R&D Investment Trends to Assess.” https://www.tprinternational.com/global-patent-filing/#:~:text=Researchers%2C%20eager%20to%20share%20their,for%20up%20to%2018%20months.

United States Patent and Trademark Office. 2019. “2019 Revised Patent Subject Matter Eligibility Guidance.” Federal Register. https://www.federalregister.gov/documents/2019/01/07/2018-28282/2019-revised-patent-subject-matter-eligibility-guidance.

United States Patent and Trademark Office. 2024. "Inventorship Guidance for AI-Assisted Inventions." Federal Register 89, no. 30 (February 13): 10043–10052. https://www.federalregister.gov/documents/2024/02/13/2024-02623/inventorship-guidance-for-ai-assisted-inventions.

United States Patent and Trademark Office. 2024. “2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence.” USPTO. https://www.uspto.gov/sites/default/files/documents/ai-sme-update-2024.pdf.

University of Maryland. 2025. "Guidelines for the Use of Generative Artificial Intelligence (GenAI) Tools at UMD." January 15. https://ai.umd.edu/resources/guidelines.

U.S. National Science Foundation. n.d. “IoT4Ag Home Page.” Accessed April 27, 2025. https://iot4ag.us/.

Vanian, J. 2023. “ChatGPT and generative AI are booming, but at a very expensive price.” CNBC, March 13. https://www.cnbc.com/2023/03/13/chatgpt-and-generative-ai-are-booming-but-at-a-very-expensive-price.html.

Varshney, Shobhit. 2025. “DeepSeek's Breakthrough is a Win for Innovation and Accessibility.” IBM. https://www.ibm.com/think/insights/deepseek-breakthrough-is-a-win-for-innovation-and-accessibility.

Volpe, A., & Yüce, O. 2023. “The Creative Future of Generative AI.” Arts at MIT, November 16. https://arts.mit.edu/mit-generative-ai-art/.

Wheelock, Brian. 2018. “Third Time’s The Charm: Federal Circuit Remands for Third Damages Trial Because Patent Owner Did Not Prove Use of Entire Market Value was Appropriate.” Harness IP. https://www.harnessip.com/blog/2018/07/05/third-times-the-charm-federal-circuit-remands-for-third-damages-trial-because-patent-owner-did-not-prove-use-of-entire-market-value-was-appropriate/.

WIPO. 2017. "Innovation and Intellectual Property." https://www.wipo.int/ip-outreach/en/ipday/2017/innovation_and_intellectual_property.html.

WIPO. 2019. "The Story of Artificial Intelligence in Patents." https://www.wipo.int/tech_trends/en/artificial_intelligence/story.html.

Winston & Strawn LLP. 2025. "Recentive Analytics v. Fox-In Precedential Case of First Impression, Federal Circuit Holds Machine Learning Patents That Do Not Improve the Technology Ineligible under Section 101." Winston & Strawn Insights, April 21. https://www.winston.com/en/insights-news/recentive-analytics-v-foxin-precedential-case-of-first-impression-federal-circuit-holds-machine-learning-patents-that-do-not-improve-the-technology-ineligible-under-section-101.

World Intellectual Property Organization. 2019. The Story of Artificial Intelligence in Patents. WIPO. https://www.wipo.int/tech_trends/en/artificial_intelligence/story.html.

World Intellectual Property Organization. 2017. Innovation and Intellectual Property. WIPO. https://www.wipo.int/ip-outreach/en/ipday/2017/innovation_and_intellectual_property.html.