Title: **Generative AI Legal Landscape 2024**
Date: March 2024
Authors:

Introduction

The LegalTech field is at the precipice of major advancements from the introduction of generative AI based capabilities. Legal processes are an opportunity rich space for generative AI, due to the preponderance of unstructured data and natural language in legal work. Recent technical advances have improved the accuracy of generative AI applications leading to an explosion in AI-first LegalTech startups. However, there are serious challenges to adoption of this technology, including some structural hurdles, that startup founders may not recognize initially. A key obstacle to think about deeply is the ownership of legal domain data, where entrenched LegalTech incumbents have a strong advantage. Market segments and sizes indicate some areas may be easier than others for new startup breakthroughs.

Historical background on innovation in LegalTech

When we reflect on technological innovation in the legal domain, we have historically witnessed developments primarily in advancing the efficiency of legal workflow via highly structured approaches. This was true from both the lens of data-driven (e.g. structured data for natural language processing) and rules-based systems (e.g. expert systems). **The particularities of the legal domain have largely cabined technological integrations to individual tasks related to information retrieval.** Such was the case for document and practice management, e-Discovery,

legal research, etc. In one form or another, search has been the focal point in decision-making around innovative choices and considered to be a path of least resistance. Automation in other areas, in particular those related to reasoning, have faced challenges of adoption. Exogenous forces behind the inertia, including the Unauthorized Practice of Law (UPL), lack of broad standardization, lack of quality public datasets available for training, and existing dependency on incumbent tools (e.g. Microsoft Office365), further contributed to the limited types of innovation that were seen as most amenable in the legal industry.

Perhaps most intractable was the fundamental problem that legal knowledge and practices are encased in natural language, specifically a subset (or even dialect) known as legal language. Many of the existing technologies required substantive, laborious, and manual effort to unpack the information contained in such text. These manual processes are data-intensive, and simply lack scalability. From an operational standpoint, the startup costs to build these tools were incredibly high, with barriers of entry that few could overcome. From a development standpoint, the ability to translate legal knowledge was highly complex with challenges not only in managing the fuzziness of the language and concepts, but also in forming consensus on ontological framing.

Renaissance: Perspective on Generative AI and Law

Yet, along came generative AI.

Within the broader umbrella of generative AI, large language models (LLMs) were seen as a particularly fitting technology for the legal sector. In fact, it has been described (and perhaps quite hastily) as an ideal fit for these models. An analysis by the Goldman Sachs' Economic Research team classified individual tasks within the various knowledge professionals' workflows and scored each task with a likelihood of automation by AI. Legal occupations standout in their exposure to AI with over 50% of legal tasks exhibiting potential for substitution or outright automation e.g. reviewing documents and proposed actions for compliance, evaluating complicated insurance claims, or making rulings on complex motions. This is higher exposure than other industries like healthcare, education, arts, and entertainment which are under 30%.

So, why is this the case? As discussed, the legal domain is regarded as encapsulated in natural language. Accordingly, these large language models are not only capable of processing text at high volumes and velocity, but they are also capable of generating text that is grammatically correct and has the semblance of eloquence. At the crux, they reflect processes that, from an outside perspective, are seemingly ones that mirror the legal workflow. For example, legal practitioners are known to search and retrieve, then classify from a large repository of text, and subsequently draft narrative and argumentative text for their clients. This reflects well the incumbent landscape and recent spate of new startups in LegalTech. As issues of scalability became resolvable by the substantive technological improvements to LLMs vs. their neural AI

ancestors (i.e. Natural language processing), the data-rich demands of LegalTech were no longer viewed as an obstacle, given performance of few-shot prompting. Furthermore, the ability to engage with these technologies dialogically, via chat or through other natural language interfaces, encouraged a sense of approachability and accessibility for a domain that is historically technology averse.

Venture funding for LegalTech startups totalled **~$700M** between Jan 2023 - Feb 2024, leaning heavily towards startups leveraging AI. Major startup fund raises included Harvey ($80M, Series B), EvenUp ($50M, Series B), Masttro ($43M, Series A), Darrow ($35M, Series B), Robin AI ($22M, Series B), Spellbook ($27M, Series A), as well as notable raises by Litify, Steno, Trellis, LawVu, PowerLaw, Parrot, Xapien and Paxton AI, demonstrating the substantial influx of new AI capabilities targeting LegalTech. Funding in the LegalTech sector has grown ~30% from $490M in 2021 and shifted significantly towards early stage (pre–Series C) venture investing from private equity. All these trends are heavily influenced by the recent advances in Generative AI.

Although use cases and mainstream adoption have not strayed far from the existing incumbent landscape, this renaissance in interest in LegalTech investment has furthered a number of technical solutions that solve some of the major problems with the use of Generative AI.

Emerging technical solutions

The key technical problems facing Generative AI use cases in Law are: (1) LLMs are not built for accuracy-related tasks; (2) legal data is largely proprietary with a substantive subset bound by client confidentiality; and (3) legal knowledge and practice are not entirely encased in text alone. These problems have played out in a rather public manner through several cases, most prominently *Mata v. Avianca* and *People v. Zachariah C. Crabill*. Nevertheless, these cases have not significantly deterred either existing incumbents or emerging players from investing further and developing these frontier technologies. Instead, there appears to be a rising interest and acceleration of LLM use and application in the legal domain.

Early advances in generative AI focus on the generation of text using LLMs pre-trained on extensive cross-domain data. While these models demonstrate emergent behaviors akin to human intelligence in a variety of tasks, this is not by itself sufficient for enterprise use cases due to the rigor required for legal work. ***Major challenges that plague early versions of this technology include lack of accuracy, limited explainability on generated text, limited privacy protections in using the most performant models, and perhaps, most importantly, difficulty in obtaining and training models on legal domain data, practitioner specific nuances, and integration into existing legal workflows.*** Legal work also may require greater control over the models themselves, where they are hosted, how frequently they are updated and the cybersecurity and uptime of any services they use. Although these challenges continue to limit broad adoption of generative AI in law, there are engineering solutions to most of them that make Generative AI software consumable for some initial applications.

- **Retrieval Augmented Generation** is the main technique for combatting LLM hallucinations / lack of accurate output. RAG is a technique where the model uses reference documents to ground its outputs. RAG can require significant engineering effort to ingest, divide up / chunk, compress, store, index, retrieve, rank and update the relevant data for querying. However, by employing the appropriate data processing, embedding creation, vector database, indexing, search and ranking technologies most use cases become tractable. RAG has emerged as an early pre-requisite to ensure a minimum level of accuracy required for legal work.

- **Reference-ability** further addresses the accuracy challenge. RAG allows for the ability to trace LLM output to reference inputs providing a mechanism to cross check the sources driving the output. This is important for legal work both to ensure accuracy and evaluate model performance.

- **Prompt engineering** employing templates and best practices for prompting LLMs can provide appropriate responses with the desired structure, verbosity and adhere to desired output constraints. Additionally rudimentary reasoning mechanisms such as chain-of-thought and ReAct can be employed to build applications that plan and execute a workflow chaining multiple LLMs, with calls to other special purpose tools as needed for the query.

- **Context length** significant research advancements have enabled state-of-the-art LLMs to have longer context lengths, enabling more text from retrieved documents, prompt templates, and additional tooling calls to be input to the LLM within its context window. This leads to higher quality responses and overall better application performance for legal use cases, where the volume of input can be lengthy.

- **Fine tuning and Alignment** legal professionals often need the LLM to read "legalese", specifically of the kind relevant to a particular subdomain of law. Fine tuning LLMs involves using domain specific data and adjusting model weights and hyperparameters to achieve the desired language, tone and domain expertise. In addition, alignment can be used to further fine tune models in accordance with a constitution or set of values (e.g. Anthropic's Constitutional AI). These may pertain to avoiding bias, toxicity, using a particular tone and type of behavior (e.g., socratic, humorous etc).

- **Specialized training sets** pretraining language models from scratch using specialized data sets can also yield superior performance for certain task specific applications. For example, Pile of Law BERT large 2 was pretrained on Pile of Law, the largest curated corpus of legal and administrative data. This dataset has been helpful across academic legal research, particularly, enabling both greater accuracy as well as relevance to specific legal downstream tasks. Other industry examples include Kelvin Legal DataPack from 273 Ventures, which was most recently used to train their "from-scratch" Kelvin Legal Large Language Model (KL3M).

- **Guardrails** entails defining the structure and nature of both inputs and outputs. This mechanism is used to achieve compliance and reduce the risk of inadvertent data breaches and information sharing outside the client specific boundaries such as bias, toxicity, etc.

- **Privacy and security.** Often LegalTech users have pre-existing privacy and security mechanisms, for example for PII redaction, data loss prevention, and client data isolation. These mechanisms can be used, sometimes with modification, in the context of Generative AI apps. Even with existing mechanisms, legal professionals may prefer to custom build applications in their own private environments using their own data to maintain client confidentiality (e.g., for litigation use cases) and ensure proprietary data advantages. Open source and on device models may be preferred for LegalTech for greater control and privacy.

- **Multi-modal inputs** are on the horizon with foundation models that can combine image, video, voice, text and potentially other types of inputs for pretraining and also produce multi-modal outputs. These additional modalities may enable new use cases (e.g., access to justice).

Additional engineering for model training, deployment, evaluation, monitoring and inference contribute to the cost effectiveness and reliability of the applications and should be selected appropriately for the use case. Challenges in this area include the lack of relevant benchmarks for evaluating LegalTech applications, emerging regulations on data use for model pre-training, hardware shortages and vendor lock-in.

As discussed, there is rapid evolution in the application architecture, and components and solutions are being developed at a cutthroat pace to make LLMs usable for a wide variety of use cases. What is fascinating is that regardless of metrics on actual use and behavioral uncertainty from the legal community, market players are forging ahead with generative AI integrations. While technical evolution in this area is promising, there are industry specific structural factors in the legal field overall that moderate generative AI startup success.

Structural Impediments

There are a number of structural hurdles to the adoption of Generative AI technologies. Specific to the legal industry, several forces continue to work against wholehearted integration and adoption. We classify these in a few categories: (1) architectural retrofitting; (2) vendor dependence; and (3) incumbent control. To start, ***we define architectural retrofitting as the difficulty of reconciling existing structural nuances of the industry with automation. This includes, for example, the infamous billable hour and/or the general lack of standardization and metrics in the provision of legal services.*** Legal professionals take incredible pride in the artisanal components of their work. This is displayed through the immense weight afforded to the individual lawyer over the collective profession. For example, there is a gap between knowledge

management's precedent language and document templates in comparison to those of a specific partner.

As many have discussed previously, the billable hour encourages inefficiency, and in effect, works against technological innovation. However, we are seeing internal change in this regard via new methods of billing. These include the ability to bill time and a half (1.5x) to encourage use and adoption of frontier technologies in existing workflow and processes. Certain, more daring, boutique firms have begun to switch to fixed fee structures. Finally, and perhaps underdiscussed, is the notion that the general workforce, particularly those who are young lawyers entering the profession, may be ill prepared for the types of work that will result from the automation of legal tasks by generative AI. That is, for many junior associates, we may see a slow transition away from drafting and redlining towards more of auditing and quality assurance. This may lead to a fundamental shift in the aforementioned perspective on the role of legal practitioners, unless human-machine collaboration is reflected more seriously. Founders considering building in the legal domain must then be sensitive to how these technologies should simultaneously disrupt yet uphold the structural and substantive integrity of the industry.

In the dramatic fallout of OpenAI mid-November 2023, it became evident that vendor dependence on proprietary models could be problematic. Currently, few incumbents in LegalTech have built and trained models from the ground up. Many are leveraging proprietary models, such as Microsoft Azure/OpenAI's and/or Anthropic's state-of-the-art models. While initially, the fear with using proprietary models was associated with privacy and confidentiality concerns on data storage and prompt injection, an added apprehension became issues of procurement and supply chain dependency. Discussions emerged on leveraging open-source models. Yet for many in the legal industry, confidence remains wavering on their relative performance[1] for enterprise deployment. Founders will need to demonstrate that they are designing their APIs to be source/vendor-agnostic. This would ensure independence from the backend model.

One of the additional complexities of the LegalTech market is the specific added advantage as an incumbent. As the legal community is rather risk averse and largely cautious against rapid change, existing relationships with incumbents would continue to have first mover advantage. That is, incumbents will remain trusted technology providers to legal unless they dramatically fail to innovate. Interestingly, with the rise of generative AI, we have seen incumbents encouraging market consolidation, opting for the "buy" versus "build" option. We will further enumerate how incumbents have reacted to the advent of LLMs in the section below. As a general consideration, ***founders must be sensitive to how incumbents not only have a strong foothold on the existing relationships across the legal industry, but also their relative dominance in data assets and security infrastructure***.

Prior to the advent of LLMs, a clear challenge was the ability to acquire sufficient quantities of legal data to train state-of-the-art models. Currently, access to data remains a key issue, but for

[1] Results in this paper were raw outputs from an LLM. Applications in production leverage retrieval augmented generation, guardrails and other standard mechanisms to overcome several of the well-known challenges with generated content.

different reasons. Rather than quantity, the quality of the data available to new entrants continues to be a hurdle. As mentioned, there are few public and open-source data pipelines that are available to founders who do not have existing relationships with legal institutions (both public and private). These include the Free Law Project, the Caselaw Access Project, and/or the Atticus Project. While they have been incredibly impactful thus far for academic researchers and emerging players, these datasets pale in comparison to those available to incumbents. Yet, specific and deep partnerships with law firms and other established legal service providers have demonstrated that, even against incumbents, there may still be an opportunity for a maverick. That is, private and proprietary data enables training that is better curated and personalized for certain tasks. This enables a much more effective product in the legal space, particularly given the field's advisory nature. As a result, **a legal co-founder with rich, seasoned experience in the industry is not only recommended but is fundamental**.

Despite these structural barriers within the legal space, there appears to be unprecedented opportunity for founders. This is largely owed to (1) the diversity of specialization and practice; and (2) the lack of empirical turn in this space. With the former, the complexity of the field naturally lends itself to plurality. This means that even though the task itself may be classified in a singular, categorical manner (e.g., contract drafting, brief writing), the execution of the same task can differ vastly depending on the specific domain within legal and often, the experience of the professional. With exponential advances in foundation models, it will become clear that beyond accessibility, reasoning, and scalability, personality and hyper-customization will become a possibility that will work congruently with legal.

Moreover, the definitional contours of legal work remain largely undefined. In effect, tasks that may not be perceivably known to the client as end-product work could, nonetheless, still be important functional work for legal professionals. These include operational upkeeping (e.g., time and billing) and, relatedly, internal collaboration. The business landscape is such that individual law firms are increasingly merging together indicating a period of consolidation and suggesting a move towards efficiency and expansion. Fostering efficiency will rapidly enable the industry to transition out of its paper-centric paradigm and towards its empirical turn. Law firms and other legal institutions may be able to establish more concrete metrics around their performance. Importantly, adoption of online legal services has since accelerated due to the pandemic force changes in the industry. Many clients turned to online legal services as convenient and accessible ways to get legal help. In fact, the American Bar Association reports 87% of lawyers work from home and/or have the option to work from home. Accordingly, tools that not only mirror existing workflows and processes, but also provide better context to account for the change of mode in professional collaboration are and will be a key area of opportunity.

As a new founder, it is, therefore, critical to understand the LegalTech market segmentation, sizing, and subdomains where breakthroughs may be easier than others. Below we provide our analysis and insights on where we anticipate future evolution, specially underexplored areas that are ripe for entry.

Market Perspective

In light of the aforementioned structural nuances of the legal industry, there often is the sentiment that LegalTech is a difficult hill to climb. While there have been numerous efforts to map the landscape of innovation, the specific market segmentation often appears rather convoluted. Moreover, these maps will always falter in terms of comprehensiveness and completeness. This, in part, is owed to the rapid dynamisms typical of startup culture, but specific to legal, there is generally a lack of clarity around how emerging players measure against incumbents. Therefore, we consider a different approach. Rather than parsing across the entire ecosystem, we first discuss market segmentation centered around the client. We describe this as the "client-facing market." We then provide market sizing that measures client-driven work against internal and other operational investments the industry is making.

To start, we consider the client-facing market to be divided across three broad lines of work: (1) Research and Analysis; (2) Document Review and Drafting; and (3) Litigation. We intentionally categorized the work involved in preparing for litigation as independent of the first two lines of work, though, in practice, litigation is a hybrid of both. This is because a plethora of specific tools have emerged to support tasks involved in litigation preparation. Perhaps most fascinating is whether this would encourage a trend of segmentation based on specific domains (e.g., immigration) within the industry. We will discuss this in further depth in our section on Future Evolutions.

We currently define the client-facing market across three segments: (1) BigLaw; (2) In-house counsel; and (3) Consumer. We add the caveat that the focus of our analysis is on the former two. This is largely owed to two key reasons: (1) current investments made to generative AI tools have made most traction in these segments; and (2) there remain important open and unresolved questions of regulatory reform associated with the Unauthorized Practice of Law (UPL), pertaining to the Consumer segment. While there are resources available for the consumer end-user, such as alternative legal service providers offering family law services, or for wills and estates, they remain confined to rather strict parameters primarily within the states of Arizona and Utah. Furthermore, experiments of regulatory sandboxing have begun to wane, as new concerns emerged at the advent of LLMs and foundation models.

Sizing

The total LegalTech market in the United States is an estimated ~$13B in 2023 growing at ~3.5% CAGR. It is expected to surpass $15B by 2026. We segment the market into 3 general lines of legal work:[2]

1. Research and Analysis includes technology that enables the process of researching information necessary to support legal work and argument development. Lawyers will spend hours poring over legal documents to analyze the details and historical context

[2] Part of the reason the consumer LegalTech market is smaller ($1B) is due to restrictions in unauthorized practice of law, which is heavily regulated and enforced. Even getting advice on a legal question can be obscured to UPL.

that apply to their cases. Incumbents like LexisNexis, Thomson Reuters, and Bloomberg Law cover the space alongside startups like Fastcase.

2. Document Processing/Contracting is a wide category covering technology that helps lawyers draft complex legal documents, such as term sheets, contracts, briefs, and more. Document-intensive industries continuously engage lawyers in document processing and use collaboration and management tools like Docusign, Icertis, and iManage. This category also includes late-stage startups such as Ironclad and Evisort.

3. Litigation is the largest category with over $4B in annual spend largely by law firms. This category describes software that governs matter management, conflict management, and e-discovery, assistive tools that support lawyers in the construction of arguments used in litigation. Incumbents include Relativity and Logickull alongside startups like Everlaw.

We also size 2 categories of legal operational work which are seeing increased investment:

1. Time and Billing is technology that captures and records all billable activity and expenses often with intuitive UI/UX, credit card capability, and ease of integration with accounting systems.

2. Legal Operations includes legal project management, calendaring/docketing, legal financial management, legal process automation, among other emerging categories that help legal departments manage growing complexity as they scale and modernize.

Applying our segmentation framework to a market sizing survey of law firms and legal departments of various sizes commissioned by Mitratech, an incumbent in the Enterprise Legal Software space, we arrive at an estimated market sizing broken down by segment and work category.

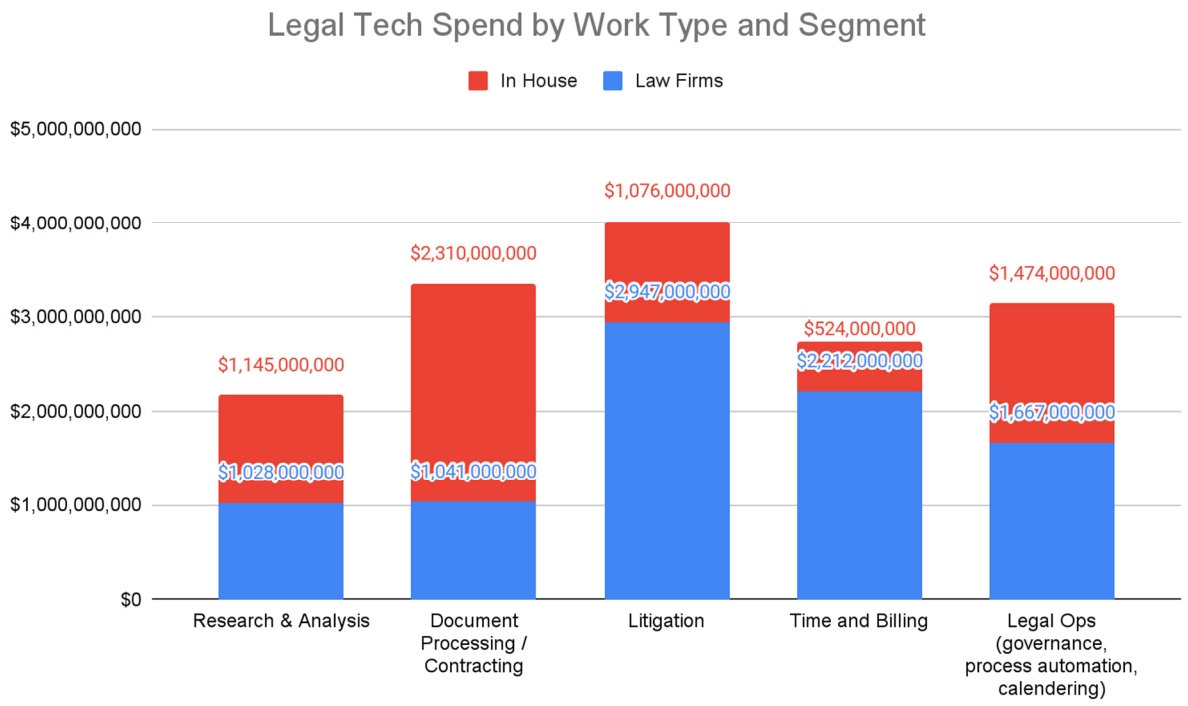

There are three key observations we may draw from our market analysis:

(1) The opportunity for emerging players to break into our aforementioned three lines of work is rather limited. The market in this space is already considerably saturated. While they currently hold the largest share of funding and investment, these areas are dominated by existing players with entrenched, well-negotiated long-term relationships. Furthermore, the moat incumbents continue to possess remains tethered to their access to the highest quality data (both public and private proprietary sources). It appears unlikely for a new entrant to clearly define how they may differ/distinguish themselves from existing solutions offered.

(2) While lawyers will largely place their trust in existing relationships with known vendors for end-product applications, most law firms and legal professionals are keen to increase the efficiency of work that could improve the legal training and education, and/or general operations of their business. These areas provide immense value to how lawyers and partners may be conceptualizing their practice.

(3) An effective startup strategy could begin with compelling horizontal technologies that address internal Legal Ops, or Time and Billing use cases. These horizontal use cases can be less mission critical and may in places require less stringent legal data access. In addition, the in-house legal services market, particularly for document processing and research and analysis, could be reached by tying into general purpose productivity tools.

In effect, ***startups interested in pursuing the legal domain should pay particular attention to how they conceivably position themselves against existing incumbents. If they regard***

themselves as providing distinctive technological capabilities, there may be an opportunity to work together with existing players rather than working as their competitors. Generative AI promises significant new benefits beyond an incumbent's existing stack and could be an entry point either independently or in partnership with an incumbent.

The Role of Incumbents in LegalTech

LegalTech incumbents play a unique role, such that they may be perceivably both a challenge and an opportunity for new entrants. Incumbents hold the trust of most AmLaw 100 firms, and accordingly, they face the critical decision in the wake of generative AI to build in-house solutions or buy existing technologies. This decision significantly impacts operational efficiency, service quality, and market positioning. Building in-house solutions offers customization and control but requires substantial investment from the standpoint of computing infrastructure and acquiring the right talent. In contrast, buying solutions is cost-effective and quick to deploy, but may involve compromises on customization and dependency on external technology partners. Furthermore, strategic decisions will need to be made on how exactly these integrations will interact with their other solutions offerings. Below we enumerate highlights of recent behaviors and decisions taken by incumbent players across our defined market's three lines of work.

Legal Research and Analysis

Thomson Reuters, Bloomberg Law, vLex Group, and LexisNexis are leading the shift towards generative AI technologies in legal research and information services. Thomson Reuters' acquisition of Casetext, known for its AI legal assistant CoCounsel, marks a significant investment in AI in the legal research context. Bloomberg Law is leveraging AI to revolutionize legal research and transactional law practices, employing tools like Brief Analyzer and Contract Solutions to streamline legal brief analysis and contract management. These AI-powered tools enhance efficiency and accuracy in legal work by automating tasks such as document review, docket filing analysis, and contract drafting. The vLex Group, formed by merging vLex and Fastcase, will focus on Vincent, an AI-powered legal research assistant. LexisNexis introduced Lexis+ AI for efficient legal research and launched a program to train 100,000 law students in AI tools.

Document Processing, Generation, and Contracting

With regards to document processing, innovations by iManage and NetDocuments are shaping this domain. iManage AI, integrated into its cloud platform, transforms knowledge search and document management by examining documents to extract and store key points. NetDocuments introduced ndMAX, a series of AI-powered products including PatternBuilder MAX for legal task automation, such as drafting and summarizing documents, extracting information, and facilitating document review. Moreover, NetDocuments acquired Worldox to enhance its document management capabilities.

Companies like Ironclad, DocuSign, Agiloft, Evisort, and Luminance are currently at the forefront of contracting AI. Ironclad launched Contract AI (CAI), which demonstrates the exact steps taken to reach conclusions in complex contract analysis. The company has also introduced Rivet, an open-source visual programming environment for AI product development. DocuSign's AI for eSignature, Contract Lifecycle Management, and Identification tools for businesses enhance the agreement process and minimize risks. Agiloft's recent strategic partnership with Kroll demonstrates its focus on expanding its global services network. Moreover, the company launched AI Trainer, enabling non-technical professionals to use AI tools. Evisort's new AI engine promotes accuracy and responsiveness in its existing solutions, streamlining contract management and analytics. Luminance's AI system, Autopilot, autonomously negotiates contracts, marking a significant advancement in AI application in legal processes.

Litigation

This area is seeing particularly significant change with Everlaw, Reveal, Exterro, DISCO, and Relativity incorporating generative AI technologies and acquiring competitors. Everlaw launched a suite of new tools, including the Everlaw^{AI} Assistant, which enhances legal workflows with features like AI review assistance, AI writing assistance, and unique privacy controls. The company's aim is to improve efficiency in the e-discovery process. Reveal expanded its e-discovery capabilities through acquisitions of Logikcull and IPRO. Exterro acquired Divebell, enhancing data compliance solutions. DISCO announced the release of DISCO Cecilia, an AI chatbot tailored for large-scale e-discovery, which enables lawyers to efficiently interrogate their private e-discovery databases and receive specific citations from private documents. This tool allows direct questioning and provides AI-generated answers with evidence citations. Relativity's new product, aiR for Review, powered by GPT-4, streamlines the e-Discovery process.

Even across these three lines of work, we are consistently seeing new partnerships form, including, most recently, Clearbrief's integration with LexisNexis.[3] These types of partnerships represent a hybrid of both the "build" and "buy" paradigm, whereby incumbents work with mavericks and emerging players to leverage their unique technology with their existing suite.

Ultimately, ***we observe broader trends towards market consolidation from incumbents, whether by purchase or partnership, as opposed to building organically.*** While incumbents may possess the data assets that will set apart the quality of their product, they require equally robust technical capabilities that often come from agile and dynamic movers. These tradeoffs and considerations suggest that for new entrants, the question they may ask is how they'd like to view the incumbent, as their partner or as their competition. Alternatively, and though currently less common, there may be an opportunity for more internal AI developments and integrations via in-house law firm incubators and research labs. We have historically seen examples of these in the UK (e.g. Allen & Overy's Fuse, Mishcon de Reya's MDR Labs), and more recently in the US with Baker McKenzie's BakerML Practice. With continued advancements in generative AI, we may see this as a possible opportunity for innovative leaders and emergent hiring practice.

[3] Clearbrief represents an interesting case, such that, within this past year, it has integrated across multiple incumbent players, including Relativity and vLex/Fastcase.

Future Evolution

Additional Notable Categories

Beyond the aforementioned landscape of generative AI integrations, are there areas in LegalTech that are particularly ripe for entry and/or are currently underexplored? Just as we have seen with dedicated tools for litigation work, symptoms of developments towards specialization appear to be on the horizon. This includes patent and IP, immigration, insurance, and interestingly, regulatory compliance. In effect, this suggests perhaps a renaissance of expert systems and knowledge engineering. Accordingly, there may be great opportunity for the advancement of hybrid tools that leverage both symbolic and neural models. In effect, this enables both technological breakthroughs that extend beyond RAG for accuracy, while also building highly curated products for expert-centered legal domains.

There is also a rising interest in productivity tools for legal purposes. This one may be more challenging as productivity is subjective to the individual lawyer. With limited metrics on performance, optimizing for operational efficiency may be difficult to capture not only due to the incredible variation, but also there remains a lack of understanding of the types of functional, non-client facing work. Still, this area lends itself to innovative promise through Time and Billing.

Finally, there is evidently incredible promise in the use of LLMs and foundation models to enable access to legal services. The issue with uptake and development, however, is not technical and purely structural. First, there needs to be an effort to retrofit and bridge existing self-help services with foundation model technologies. There has been a recurring pattern of replacing existing self-help tools via prompting of LLM-enabled chat interfaces (e.g., ChatGPT, Claude, Bard) for immediate responses to legal questions. This, of course, encourages misuse and undue reliance on these tools. Therefore, and in addition to the above reasons related to UPL, there will need to be partnerships made with legal aid, pro bono efforts, and other court services centers to implement these technologies in a manner that would sufficiently mitigate risk and provide safeguards to the end user. Fortunately, we are seeing an eagerness from these institutions, and in particular both federal and state courts, to leverage generative AI for large-scale access to justice.

Conclusion

The LegalTech market is seeing a renaissance in investment due to the emergence of Generative AI with a 30%+ uptick over the last couple of years, largely in early-stage venture funding. This technology indeed has broad applicability towards solving bottlenecks in a number of areas of legal work with a multi-billion-dollar addressable market. The technical problems with Generative AI, such as hallucinations, explainability and private data access could be tractable if use cases are considered thoroughly. Furthermore, application architectures that appear promising and reliable are emerging. However, there are significant structural hurdles in this industry that limit

the potential for large new startups. Legal workflows are fragmented, and the power of entrenched incumbents is a substantial concern due to their stronghold on access to legal data.

There is potential for startup disruption, but it is not as easy as it may appear at first glance. Founders should pay special attention to the dynamics of how they introduce any new AI driven technology to this market. Careful selection of the market sub-segment and use case is key. Areas that are less data dependent or where incumbents are open to partnerships may be better initial targets. There could be initial footholds in horizontal areas within LegalOps and regulatory compliance from which to build and expand. Another possibility is working directly with law firms as either design partners and/or to contribute to their internal innovation teams. Nevertheless, this may be more of a shift in talent management and how law firms themselves may be strategically considering the skill sets needed to bolster their practices.

All in all, there is undeniably momentum in LegalTech that should not be taken lightly. Ensuring that enthusiasm is met with a measure of caution would unlock incredible opportunities in the domain, offering an unprecedented depth and understanding of legal knowledge and practice.

Appendix

	Big Law	In House	Consumer
Research & Analysis	538M (legal analytics) + 490M (knowledge management) = 1.028B	488M (legal analytics) + 657M (knowledge management) = 1.145B	
Document Processing / Contracting	518 (doc mgmt) + 523 (collab tools) = 1.041B	437 (doc mgmt) + 336 (collab tools) + 769 (contract mgmt) + 768 (ip mgmt) = 2.310B	Creation of wills
Litigation	1.226B (matter mgmt) + 1.207B (conflict mgmt) + 514M (e-discovery) = 2.947B	611M (matter mgmt) + 465 (e-discovery) = 1.076B	
Time and Billing	1.203B (e-billing) + 1.009B (time and billing) = 2.212B	524M (e-billing)	
Legal Ops (governance, process automation, calendering)	514M (legal project management) + 574M (cal/docket) + 556M (legal fin mgmt) + 23M (legal process auto) = 1.667B	825M (GRC) + 649M (legal project mgmt) = 1.474B	
TOTAL	8.414B	5.872B	Entire B2C market is ~$1B