FDA 510(k) AI Submissions: Guidelines and Best Practices

By Adrien Laurent, CEO at IntuitionLabs • 2/27/2026 • 35 min read

fda 510k submission   ai medical devices   regulatory affairs   generative ai   pccp guidance

large language models   medical device software   samd   fda compliance

Executive Summary

By 2026 the use of artificial intelligence (AI) in medical device 510(k) submissions has evolved from speculative hype into a cautiously accepted support tool. Federal regulators and industry both are experimenting with AI tools – notably large language models (LLMs) such as ChatGPT – to accelerate and improve elements of the submission process. The FDA, having completed an internal *AI-assisted scientific review pilot* in 2025 ([1] www.fda.gov), is actively integrating AI assistance (e.g. FDA's "Elsa" copilot) into its review workflows. Concurrently, device developers and consultants report using generative AI to draft narratives, search literature, and organize regulatory content. Early case examples (e.g. a recent 510(k) for the **Sclerosafe** device) show AI speeding tasks such as predicate comparison, criteria justification, and risk analysis ([2] www.linkedin.com) ([3] www.linkedin.com). Publication analyses find that AI/ML-enabled devices dominate medical imaging; for example, in 2025 FDA cleared 295 AI/ML devices (median review ~142 days) with ~71% in radiology ([4] innolitics.com) ([5] innolitics.com). These trends underscore the centrality of AI in today's MedTech.

Our review finds that **AI can indeed "work" in 510(k) submissions** – but with strong caveats. Generative AI assists best when used as a human-guided assistant, combined with robust fact-checking. Purely automated writing is unsafe: experts emphasize that the **sponsor remains 100% responsible for all submission content** regardless of tools used ([6] pharmacystandards.org) ([7] pharmacystandards.org). Misleading outputs ("hallucinations") are a known risk ([8] pmc.ncbi.nlm.nih.gov) ([9] www.meddeviceonline.com). Best practices include using **"retrieval-augmented"** workflows (grounding AI queries in trusted data) and structuring submissions in clean, accessible formats (e.g. the FDA's eSTAR template) so that both human and machine reviewers can parse content easily (www.pharmanow.live) ([10] www.meddeviceonline.com). In summary, AI tools can significantly **speed literature reviews, draft routine text, check data consistency, and highlight gaps** – but only when tightly controlled by skilled regulatory professionals.

This report provides an in-depth analysis of the AI-assisted 510(k) submission landscape in 2026: its background, regulatory context, available tools, real-world use cases, data trends, limitations, and future prospects. We draw on published studies, FDA announcements, expert commentary, and illustrative case reports to assess *what actually works* in practice. We conclude that the promise of AI — faster preparation, improved completeness, and potentially quicker FDA review — is real, but only if applied carefully under strong human oversight and with transparency. Failure to validate AI-generated content against source documents can lead to errors or delays. Looking ahead, regulators are likely to refine guidance on AI in submissions (building on FDA's 2024–25 guidance on AI/ML device lifecycle and labeled "PCCP" plans) ([11] www.akingump.com), but for now companies must navigate a "human-in-the-loop" regime: leveraging AI as a powerful drafting assistant, yet owning the final documents.

Introduction and Background

The FDA's **510(k) premarket notification** process is the principal mechanism for most new medical devices coming to market in the United States. Under 21 CFR §807.92, a sponsor must demonstrate that a new device is "substantially equivalent" to a legally marketed predicate. Historically, 510(k) submissions have been largely narrative-driven documents, comprising sections on device description, intended use, technological characteristics, performance testing, and comparisons to predicates. Approximately 75–80% of annual FDA device clearances come via 510(k) ([12] www.mdpi.com). The process has long been criticized for its complexity and the volume of information to compile, yet timely approval is critical for patient access and company economics.

Against this backdrop, artificial intelligence has entered the scene. The term "AI" here primarily refers to **generative AI** (notably large language models like OpenAI's GPT-4 or similar) that can craft human-like text based on prompts. These tools emerged into public use around 2022–2023, raising questions in industries such as healthcare and life sciences about their role in documentation and analysis. In principle, AI could help regulatory writers by sifting vast web literature, summarizing guidance documents, and even drafting text segments, thereby reducing labor on routine writing tasks.

To understand the potential and pitfalls of AI in 510(k) submissions, we first survey the evolving regulatory landscape and stakeholder attitudes. We then drill into specific applications and evidence: from anecdotal case studies to aggregated clearance data. Throughout, we contrast the *promise* of AI (speed, thoroughness) with the *requirements* of accuracy, traceability, and human accountability imposed by regulators. As one recent analysis notes, the FDA and EMA's "single, unifying 'North Star'" for any AI use in documentation is this: the **sponsor is 100% accountable** for all content, and AI is merely a tool, not an "author" ([6] pharmacystandards.org). This principle underpins our evaluation of what actually *works* when sponsors use AI assistants.

The Regulatory Environment for AI in Medical Devices

FDA Initiatives and Guidance

In recent years the FDA has actively addressed AI in medical devices, issuing frameworks and rules for AI/ML-based software as a medical device (SaMD). In December 2024, FDA published final guidance on **Predetermined Change Control Plans (PCCPs)** for AI-enabled device software, applicable to De Novo, PMA, and 510(k) pathways ([11] www.akingump.com). This underscores that FDA expects manufacturers of AI systems to anticipate future modifications. The agency continues to emphasize **data quality and lifecycle management** for AI devices ([11] www.akingump.com) ([13] pharmacystandards.org).

Notably, FDA has also made internal strides. In May 2025 Commissioner Makary announced completion of an FDA-wide pilot of AI-assisted review and an aggressive integration schedule ([14] www.fda.gov). By summer 2025 all FDA centers (drugs, devices, biologics, foods, etc.) were to operate a common secure generative-AI platform for reviews ([1] www.fda.gov). The internal tool, dubbed *Elsa* (Electronic Language System Assistant), is used by reviewers to summarize documents, flag inconsistencies, and automate administrative screening ([15] www.meddeviceonline.com). Importantly, FDA asserts Elsa is built on AWS GovCloud and *does not train on proprietary submission data* ([15] www.meddeviceonline.com), alleviating some data-security concerns.

At the same time, FDA has increased transparency on AI in devices. Its official **AI/ML-Enabled Medical Devices list** catalogs cleared products and provides links to summaries. FDA encourages sponsors to disclose AI elements in their public summaries, and the agency has stated it will explore tagging devices that incorporate foundation models (e.g. large language models) in future updates ([16] www.fda.gov). These measures aim to signal FDA's evolving approach, although current guidance on AI usage is still largely tool-agnostic and focused on outcomes.

International Perspectives

Globally, regulators mirror similar themes. The European Commission's **AI Act** (Regulation 2024/1689) enacted in 2024 imposes a risk-based framework on AI, including medical AI as "high-risk". While the AI Act addresses

lifecycle risk management, it does not specifically govern the act of writing regulatory documents with AI. EMA's 2024 *Reflection Paper on AI in Medicines* similarly calls for human oversight and traceability of AI models in drug and device submissions ([13] pharmacystandards.org). Health Canada's 2025 guidance for ML-enabled medical devices likewise mandates transparency, performance monitoring, and PCCPs for any planned algorithm changes ([17] aihealthcarecompliance.com). Both EMA and FDA stress that sponsors must validate whatever AI output they rely on. As one industry analysis puts it: whether an AI tool or a human created the text, "the sponsor is 100% accountable… [for] having reviewed, verified, and taken full ownership of its contents" before filing ([6] pharmacystandards.org).

In summary, regulators have not banned AI tools; on the contrary, they expect **rigorous controls** around any AI-generated content. Guidance documents (draft and final) urge adoption of traceable workflows (e.g. justified model updates via PCCPs) and insist on human-in-the-loop review to meet cGxP standards (ALCOA+ for data integrity) ([7] pharmacystandards.org). There is no FDA regulation of specific software tools; rather, sponsors must validate their system's intended use (even if that system is ChatGPT!). This regulatory backdrop means that AI can only "work" in submission preparation if it is deployed within a clear governance structure and thorough human oversight.

AI Tools and Techniques for 510(k) Submission Preparation

Generative AI and Language Models

The primary form of AI discussed in regulatory writing is generative AI — systems that produce text or data based on learned patterns. Modern LLMs like GPT-4 have shown an uncanny ability to draft coherent, detailed passages from simple prompts. This suggests potential uses in editing, summarizing, or even authoring parts of a regulatory submission. For example, a sponsor might prompt an AI to generate a first-draft device description, a rationale for equivalence to a predicate, or a lay-summary of clinical data. In 2026, several vendors also offer specialized AI copilots trained on regulatory guidelines (e.g. GPT extensions, fine-tuned models). Some platforms (like the **eSTARHelper** Microsoft app) integrate a GPT-4 copilot with domain-specific databases and search tools (www.pharmanow.live).

However, generative models are inherently probabilistic. They do not "verify" facts; they predict plausible outputs based on training data. In regulatory contexts, this means that AI suggestions *may sound authoritative but still be wrong or incomplete*. Experts caution that without verification, AI outputs cannot be trusted to meet FDA's strict standards. As Hussain and Balani (2024) note, generative models "can kickstart drafting… but cannot ensure the accuracy, completeness, or precision required" (www.pharmanow.live). Instead, the recommended best practice is a **hybrid approach**: use AI to generate or retrieve text, but overlay human validation and automated fact-checking.

Recent proposals advocate **Retrieval-Augmented Generation (RAG)** systems for regulatory use. In RAG, the AI model is fed not just a general language model but also domain-specific documents (e.g. CFRs, guidance PDFs, internal SOPs) so it can cite real sources. For instance, the SmartSearch+ system described by Hussain et al. uses a full-text indexed database of FDA regulations and guidances. Analysts point out that blending GPT-4 with such search yields "validated insights": AI drafts are cross-checked against actual FDA content (www.pharmanow.live) (www.pharmanow.live). In practice, an engineer might ask a custom AI copilot to summarize 21 CFR §807 content; the tool then retrieves exact regulatory passages while also using generative inference. This "smart" pipeline provides the best of both worlds — speed and contextual knowledge — and has

been characterized as a *"duelling banjo"* approach where AI and human review verify each other (www.pharmanow.live) (www.pharmanow.live).

Major providers have also begun offering regulatory compliance assistants. For example, Microsoft's Azure OpenAI offers private GPT integration (with data isolation) for enterprises, and some consultants offer "Regulatory Copilot" services tuned for FDA Q&A. By 2026 it is common for large medtech companies to experiment with private LLM instances. These can be internally trained (or prompted) on a company's design history files, clinical study database, and FDA submissions so far. The goal is an AI that "knows" the device program and can answer Q-submission queries, format tables, or even review Class II summary quality.

AI-Assisted Content Generation

Practically speaking, the ways that sponsors have employed AI in preparing 510(k) documents include:

- **Drafting Section Text**: AI can propose wording for device descriptions, intended use statements, and labeling summaries. For routine or boilerplate content (e.g. standard introduction paragraphs, patent disclaimers) LLMs can save time. In one anecdote, a medical device writer at VVT Medical reported using "ChatGPT 4" to help craft and refine several narrative sections in their 510(k) submission, including technical descriptions and regulatory comparators ([2] www.linkedin.com). The key is iterative refinement: the user prompts for relevant wording and AI-generated text is edited and fact-checked. The FDA does not allow "the AI writes and we copy wholesale," so the output always undergoes human revision and source verification.

- **Literature Search & Reference Gathering**: A strength of AI is scanning large text corpora quickly. Sponsors have used ChatGPT or similar tools to do *preliminary* literature surveys. For example, one regulatory engineer reports using ChatGPT to identify academic papers and standards relevant to specific device parameters (e.g. "tip flexibility values" for a catheter) ([18] www.linkedin.com). The AI pulls in data from sources like PubMed abstracts or digitized standards, then cites references. The human reviewer then checks those citations and discards any dubious or irrelevant ones. A properly instructed AI can help compile experimental benchmarks or performance requirements from the literature, drastically reducing hours of manual search. However, because an AI's knowledge cutoff or data scope may be limited, it should only be used to **discover leads**, not as authoritative evidence by itself.

- **Predicate Comparison Tables**: Identifying and articulating differences between a new device and its predicate is a core 510(k) task. AI can expedite this by formatting comparisons. In the Sclerosafe example, ChatGPT was used to craft a *detailed predicate comparison analysis* which "highlighted the similarities" between the new device and its cleared predicate, strengthening the argument for substantial equivalence ([2] www.linkedin.com). By summarizing technological features side-by-side and wording the narrative equivalence rationale, AI can save regulatory writers from re-typing known information. Again, the final content was reviewed by experts to eliminate any generic or incorrect phrasing.

- **Justification of Test Criteria**: Setting acceptance criteria (pass/fail thresholds) for bench tests, biocompatibility, or performance often requires citing standards and literature. ChatGPT can speed this by suggesting typical values and pointing to standards. In practice, sponsors prompt the AI (e.g. "give references for normal tip flexibility range in catheters") and the model cites relevant journals or existing product data. The VVT Medical case described how the AI produced references that supported the firm's chosen mechanical values for testing ([18] www.linkedin.com). These references were then verified and included in the submission. Crucially, the human team ensured that all criteria had an FDA or consensus standard basis.

- **Anatomical Model and Simulation Explanation**: When bench testing uses anatomical simulators or computer models (e.g. vascular simulators, organ phantoms), regulators often ask for justification of model choice. ChatGPT has been used to draft these justifications. By providing known anatomy variations or device feature context, the AI can generate an argument for why a given phantom or model is appropriate. In the cited example, the AI "searched academic literature" on the chosen phantom to validate the manufacturer's approach ([19] www.linkedin.com). The output became a draft we could refine and cite, assuring FDA examiners that the anatomical test case accurately mimics clinical conditions.

- **Risk-Benefit and Clinical Discussion**: The risk/benefit section is another narrative-heavy part of 510(k)s. Sponsors have experimented with AI to outline the major risks of their device and its anticipated benefits, then write balanced prose. For example, one consultant credited ChatGPT with helping to "meticulously examine potential risks… while highlighting its numerous benefits" ([3] www.linkedin.com). The AI's ability to regurgitate medical knowledge (e.g. listing known complications or performance levels) can jump-start this analysis. The team then fact-checks against their own data and published clinical findings. Thus AI can aid a more rapid first-pass draft of a risk/benefit analysis, though final responsibility for accuracy remains with the company.

- **Administrative and Formatting Tasks**: Some "light" uses include reformatting tables, converting figures into text, or checking consistency. For example, an AI could verify that references are cited or unify terminology between sections. The FDA's shift to eSTAR (an XML-based submission template mandated from 2023 ([20] www.meddeviceonline.com)) means submissions are now highly structured. AI tools that can parse or populate eSTAR fields (e.g. by tagging document sections, extracting metadata) can reduce human error. Vendors like reformers or consulting groups sometimes bundle Copilot tools into eSTAR workflow software to auto-fill form fields or cross-link modules, though this is an emerging capability.

In summary, **current AI usage in 510(k) preparation is mostly as a sophisticated drafting and search assistant**. There is no evidence that AI autonomously generates a perfect submission end-to-end. Instead, the prevailing model is "human writes with AI help." As one industry observer notes, many engineers simply keep a window of ChatGPT open to "help them refine design artifacts" and query regulatory norms on the fly ([21] www.jamasoftware.com). This partial, supervised use has shown real benefits in speed and completeness of work – but only when checked by subject-matter experts.

Below is a **summary table of common 510(k) tasks and how AI is reported to assist**:

510(k) Task	AI-Assisted Activity	Example / Source
Predicate device comparison	AI drafts side-by-side feature comparisons and equivalence arguments	VVT Medical's Sclerosafe submission: AI highlighted similarities to predicate ([2] www.linkedin.com)
Acceptance criteria justification	AI searches literature/standards to find typical performance values	ChatGPT pulled references for mechanical test criteria (e.g. tip flexibility) ([18] www.linkedin.com)
Anatomic model/test justification	AI generates literature-backed rationale for chosen bench models	AI used academic sources to validate test phantom choice ([19] www.linkedin.com)
Risk–Benefit analysis	AI assists in outlining device risks and benefits in balanced narrative	AI helped draft a comprehensive risk/benefit section (Sclerosafe case) ([3] www.linkedin.com)
Device/label description drafting	AI proposes initial draft language for device specs and labeling	Consultants noted AI can generate polished first-draft text to be reviewed by experts
Literature and regulation search	AI/fine-tuned search tools retrieve relevant guidance and publications	GPT-4 with SmartSearch retrieves relevant 21 CFR citations vs open web (www.pharmanow.live)
Consistency checks & formatting	AI/automation inspects for redundant info and organizes content	Industry advises organizing by topic (not chronology) to aid any AI or reviewer ([10] www.meddeviceonline.com)
Data summarization (graphs, tables)	AI interprets data tables/plots and summarizes results in words	Emerging: general LLMs can describe graphs; some vendors offer specialized BI tools.

Each "activity" above requires human review. When used properly, AI has **reduced menial tasks** (e.g. writing boilerplate or hunting references) to seconds or minutes instead of hours. For example, the FDA's pilot reviewers themselves reported tasks that once took days could be done in minutes with AI tools ([22] www.fda.gov). Anecdotally, startups and consultancies note that AI speedups have allowed small teams to tackle more submissions or revisions than otherwise possible.

However, blind reliance is dangerous. As noted by Hussain et al., generative drafts must always be "verified through thorough human-led analysis" (www.pharmanow.live) (www.pharmanow.live). In practice, effective AI-

assistance in 510(k) submissions seems to consist of *supplementary drafting and analysis tools* used by experienced writers, rather than wholesale AI authorship.

Data Analysis: Trends in AI-Enabled Device Submissions

To gauge the context for AI in 510(k) processes, it helps to examine the broader device clearance landscape. In recent years the number of AI/ML-enabled devices cleared by FDA has grown explosively. Joshi *et al.* (2024) compiled 691 cleared AI/ML medical devices (as of Oct 2023), noting a "significant surge in approvals since 2018" and a reliance on 510(k) pathways ([12] www.mdpi.com). These devices have skewed heavily toward radiology (interpreting images) and software as a medical device (SaMD) products. The authors report that 62% of AI device clearances were SaMD and 63% were diagnostics ([23] innolitics.com).##

By 2025 this trend continued: an industry analysis found **295 AI/ML-enabled device clearances** in 2025 alone ([4] innolitics.com). Of these, **212 (71.5%) were in radiology**, reflecting the ongoing dominance of imaging applications ([5] innolitics.com). Cardiovascular (8.8%) and neurology (4.7%) were distant seconds (Table 1). Roughly 62% of cleared AI devices were SaMD (software-only) ([4] innolitics.com). Importantly, about 10.2% of the 2025 clearances included Predetermined Change Control Plans (PCCPs), meaning the sponsor received FDA concurrence on how their AI model could adapt post-clearance ([24] innolitics.com).

Despite external challenges (e.g. government shutdowns), the FDA swiftly processed these submissions. According to Innolitics data, the **median review time** for 2025 AI device 510(k)s was ~142 days ([4] innolitics.com). This was broadly comparable to pre-pandemic medians but showed a slight improvement: 24% of submissions cleared in under 90 days (innolitics.com), whereas only 22% took over 200 days (innolitics.com). Investigators have speculated that FDA's use of AI tools (Elsa and others) may have contributed to this efficiency, noting *"review time appears to be a bit faster this year. Could it be because FDA uses AI assistance to help reviews?"* ([27] innolitics.com). While causation is unproven, the timelines are encouraging for sponsors seeking faster access.

The demographic data also indicate a flourishing device sector: in 2025 there were 221 unique manufacturers among the 295 clearances, with many startups innovating AI products ([4] innolitics.com). This suggests both increasing competition and FDA comfort with new entrants' AI designs. (Of course, retrospective analyses do not detail *how* these companies prepared their submissions, only that they succeeded in gaining clearance.)

Table 1. *AI/ML-Enabled Device Clearances via 510(k) by Medical Specialty in 2025 (FDA data compiled by Innolitics ([5] innolitics.com)).*

Specialty	Clearances (2025)	Percentage
Radiology	211	71.5%
Cardiovascular	26	8.8%
Neurology	14	4.7%
Orthopedic	10	3.4%
Gastroenterology / Urology	7	2.4%
Clinical Chemistry	6	2.0%
Anesthesiology	5	1.7%
Other Fields	16	5.5%

Source: FDA FDA's public AI/ML device database analyzed by Innolitics ([5] innolitics.com) ([4] innolitics.com).

These figures underscore that most FDA-cleared AI devices today are software solutions aiding diagnosis, which is why radiology is so dominant. For traditional 510(k) sponsors (non-AI-device companies), this trend matters because it means the FDA is increasingly familiar with the AI/ML context. But it also means that any sponsor using AI in their device (especially in algorithms or decision support) must be ready for scrutiny under these AI-specific guidelines (PCCP, performance monitoring, bias analysis, etc).

On the regulatory documentation side, however, there is no comparable database of "AI-assisted submissions." We have no public metric for how many sponsors use ChatGPT or other generative tools as part of their 510(k) preparation. Instead, our understanding of outcomes comes from case anecdotes (see next section) and from FDA's broad statements about efficiency gains ([22] www.fda.gov). Over time, we may see surveys or studies of submission quality pre- and post-AI usage. For now, the data trends show that FDA is clearing plenty of devices (including many AI products) and suggests that integrated review tools may be shaving off some time from their end.

Case Studies and Industry Experiences

Because formal data on "AI-assisted submissions" is scarce, real-world examples are mainly drawn from industry accounts and news. These illustrate both benefits and caveats.

Case Study: Sclerosafe Device Submission (2023–2024)

A notable public account comes from a 2024 LinkedIn post by a regulatory engineer named **Liron Tayeb**. Tayeb recounts how his team used ChatGPT-4 during their 510(k) submission for the *"Sclerosafe"* device. In his narrative, ChatGPT served as an "essential tool" and "invaluable partner" in drafting several parts of the submission ([2] www.linkedin.com) ([3] www.linkedin.com). Specifically, they used the AI for:

- *Predicate Comparison*: ChatGPT generated a structured analysis highlighting how Sclerosafe's features were similar to its predicate device. This fleshed out the argument for substantial equivalence in a coherent format ([2] www.linkedin.com).
- *Test Acceptance Criteria*: By querying medical literature through the AI, the team obtained relevant scientific values (e.g. mechanical testing thresholds) which they used to justify their chosen bench-test criteria ([18] www.linkedin.com).
- *Anatomic Model Justification*: When the FDA asked about their anatomical simulator for bench testing, ChatGPT helped draft a literature-backed explanation of why the model was appropriate ([19] www.linkedin.com).
- *Risk/Benefit Analysis*: ChatGPT assisted in drafting a balanced discussion of the device's potential risks against its benefits, "comprehending complex medical concepts" to ensure no major factors were omitted ([3] www.linkedin.com).

Tayeb stressed that **all AI outputs were meticulously verified**: they independently checked every reference that ChatGPT provided and ensured factual accuracy. For example, he noted using "ChatGPT version 4 … provided us with references for each source, allowing us to meticulously assess their reliability and credibility" ([28] www.linkedin.com). The final 510(k) file was ultimately written by humans, but the AI dramatically accelerated early drafts and literature reviews. According to the post, by using ChatGPT, tasks that would normally take hours of writing or digging through papers were done in minutes ([22] www.fda.gov) ([2] www.linkedin.com).

This real-world example (the only detailed public 510(k) AI-case we found) demonstrates how AI can function in practice. It shows that **predicate/performance analysis** and **justifications** are particularly amenable to AI support – essentially, wherever structured technical arguments or literature citations are needed. Tayeb's account affirms the importance of human oversight: he explicitly warns readers that "accuracy and reliability" demands verifying all AI output ([28] www.linkedin.com). His concluding message was that AI "became an essential tool in our pursuit of regulatory approval" but only when combined with expert review ([29] www.linkedin.com).

Industry Commentary and Surveys

Aside from specific cases, several industry publications have assessed AI in medical device workflows more broadly. A 2025 **Hogan Lovells** article noted the FDA's AI pilots and advised sponsors to prepare their submissions to be "AI-ready" ([30] www.meddeviceonline.com). Similarly, consultants emphasize structuring data and documentation carefully, because machine learning tools perform best on clean, consistent inputs ([10] www.meddeviceonline.com) ([6] pharmacystandards.org).

Professional surveys have also highlighted emerging use. For example, a 2025 **Jama Software** whitepaper gathered predictions from industry experts. Among those:

- Richard Matt (Aspen Consulting) predicted AI would automatically analyze **predicates, guidances, and standards**, reducing human ambiguity ([31] www.jamasoftware.com).
- Adam Smith (AgentAstro.ai) similarly forecast that "AI has become the connective layer across the device lifecycle, replacing manual research with automated analysis of predicates [and] guidances" ([31] www.jamasoftware.com).
- Mike Celentano (System Optimization Specialists) reported that engineers were already using AI to *summarize user interviews*, *draft requirements*, and run *independent reviews* of design documents ([32] www.jamasoftware.com).
- These experts all stressed a human-in-the-loop approach. Dan Purvis (Velentium) noted "AI has amazing abilities when harnessed well… [it] makes suggestions that are then reviewed by a person" ([33] www.jamasoftware.com). Vincent Balgos (Jama) pointed out that companies are starting methodically by first organizing good data frameworks, then slowly integrating AI into development and documentation ([34] www.jamasoftware.com).

Collectively, these commentaries align with the case example above: AI is seen as a tool to **supplement** human expertise, particularly for repetitive or information-heavy tasks. No one suggests handing the submission off to AI without oversight.

FDA Perspective

FDA reviewers' perspectives also provide insight into what "works." In the May 2025 announcement, FDA scientists noted that AI allowed certain review tasks to drop from *"three days to just minutes."*** ([22] www.fda.gov)**. Although not detailing which tasks, this implies document scanning and summarization chores benefited immensely. FDA has indicated that sponsors can help this process by using *structured formats and clean data* ([10] www.meddeviceonline.com). In other words, if a 510(k) is well-organized (e.g. broken into clearly labeled sections) it is easier for FDA's internal AI to parse, and vice versa. To this end, the agency is promoting use of eSTAR's standardized fields; experienced submitters with eSTAR expertise are "a step ahead" in this new regime ([30] www.meddeviceonline.com).

Cautionary Tales

There have been a few published warnings. The meddeviceonline article by Jim Kasic presents a vivid caution: because AI models *learn everything*, they can also learn nonsense. Kasic recounts a humorous example from memory: AI analyzed patient data and found that, coincidentally, *all* patients lived on streets with tree names (Maple, Elm, etc.) ([9] www.meddeviceonline.com). While factually true, this correlation was clinically irrelevant; an AI applying neural patterns "discovered" spurious findings. The lesson was clear: AI may produce plausible but meaningless results if context is not enforced. This underscores that developers must be ready to question any surprising AI output, and maintain channels to investigate or correct it with FDA if needed.

In practice, most firms use AI-generated text only after thorough vetting. According to one **Council on Pharmacy Standards** analysis, regulators expect any accepted AI draft to meet traditional data integrity (ALCOA+) rules ([7] pharmacystandards.org). This means *attributability*: the final submission must record who prepared each part and how it was verified. As a result, companies typically do *not* acknowledge AI in the final 510(k) document. Under current rules, you *cannot* label a section "GPT wrote this." Instead, AI is considered part of the quality system: like using a word processor, the sponsor's qualified personnel must ensure the output is correct ([6] pharmacystandards.org) ([35] pharmacystandards.org).

Best Practices and What "Works" in 2026

Drawing on the above evidence and expert opinion, we summarize what appears to **work well** – and what does *not* – in 2026's AI-assisted 510(k) preparations.

What Works

1. **Preparation using Structured Data**. Vendors who *already* used eSTAR (or similar structured templates) found AI tools more effective. By 2026, eSTAR is mandatory, meaning submissions are naturally in a machine-readable format ([20] www.meddeviceonline.com). AI tools function best with consistent field names and data types. For example, if hardware specifications are already in tabulated eSTAR fields, an AI can quickly extract and summarize them, whereas free-text narratives require more language processing. The 2025 review pilot encouraged sponsors to "use structured formats and clean, consistent data" to help FDA's AI parse content ([30] www.meddeviceonline.com). In practice, this means developing each submission section by section, using tables or bullet lists where possible, rather than long paragraphs. Sponsors successful with AI assistance often reorganize info logically: group by topic (performance, materials, etc.) rather than by design process chronology ([36] www.meddeviceonline.com).

2. **Human-in-the-Loop Verification**. Protecting against AI errors hinges on review. In all reported examples, domain experts double-checked AI output. If AI generates a claim or number, the team cross-references it to official sources. As pharmacy GxP guidelines emphasize, one must validate every data point: "an un-reviewed AI draft generally does not meet ALCOA+ expectations. Regulators expect a Human-in-the-Loop review process… with 21 CFR Part 11-compliant controls before treating content as a true GxP record" ([7] pharmacystandards.org). Thus, the practical workflow is: (AI drafts text) → (human edits and inserts citations) → (another human proofs final). This multi-layer review means AI speeds *initial writing*, but actual citations and final judgment remain human tasks.

3. **Retrieval-Augmented Workflows**. Purely free-form AI writing is unreliable; but if an AI system is forced to "back up" its text with real documents, it can be very helpful. Tools that integrate the FDA database and medical literature are gaining traction. For example, embedding an AI in a retrieval system (like ChatGPT with a plug-in, or private RAG models) lets the user query regulatory text directly. The company SmartSearch+ illustrates this approach: it indexes all FDA guidances and CFRs, and uses a GPT-4 interface to answer specific queries. Such systems can rapidly fetch exact regulatory clauses. In contrast to generic ChatGPT, a *validated* search makes it possible to answer a question like "What CFR section requires bench-testing sera for this device?" almost instantly. Developers who have tried RAG note it dramatically cuts the time needed to find relevant standards or precedence. As Hussain *et al.* describe, this "integrative solution" allows precise retrieval of FDA content to ground AI insights (www.pharmanow.live) (www.pharmanow.live).

4. **Focused Use on Repetitive and Routine Text**. AI excels at repetitive language. For instance, any text that must be similar across multiple 510(k)s – such as generic safety statements, common labeling phrases, or descriptions of standard test apparatus – can be templated by AI. Many companies keep a "style guide" for descriptions and have ChatGPT rephrase or paraphrase those statements uniformly. Similarly, the data table content (dimensions, performance matrices) can be checked by AI for formatting uniformity. The 2026 cohort of submissions tends to have AI "assist" more with these boilerplate sections, so that the human writer can concentrate on the truly novel parts of the submission.

5. **Iterative Drafting with Version History**. Good practice is to treat the AI as you would any draft author. Maintain version control: keep logs of AI prompts and outputs. Some regulatory software platforms are beginning to incorporate "drafting history" features, where each AI suggestion is timestamped. This is not only useful for internal tracking, but may be needed for audit trails under Part 11 (showing how and when text was generated or modified). While not strictly a tool feature, companies that keep meticulous records tend to have smoother experiences if an FDA audit asks about their AI usage.

6. **Transparency with FDA (through Pre-Sub)**. In ambiguous areas, some firms are choosing to mention AI use as part of pre-submission (Q-Submission) discussions. Saying "we used GPT to identify literature on X" is not prohibited, but should be framed as part of the background research process. If anything, early dialogue with FDA can clarify expectations. So far, no official mandate requires disclosing AI use in a 510(k), but sponsors who proactively validate their AI steps and mention them in Q-subs (if relevant) may reduce later questions.

What Does *Not* Work (or Works Poorly)

- **Full Reliance on Unverified AI Content**. The most obvious pitfall is trusting the AI too much. Hallucinated data or misinterpreted studies can slip into the draft. The street-name example ([9] www.meddeviceonline.com) illustrates how AI will latch onto any pattern, however irrelevant. Claims about device performance or safety **must always originate from real evidence**. We find that AI-generated references sometimes don't exist or are off-base, so blindly including them will backfire. Indeed, the linked case study shows the sponsor explicitly verified *every* AI-provided reference ([28] www.linkedin.com). In 2026, any shortcut (like copy-pasting GPT content without checks) generally leads to problems in quality review.

- **Expecting AI to Perform Complex Analysis**. AI can't run simulations or analyse raw data. For example, telling GPT "evaluate the stress-strain curves" is beyond its capability. AI also cannot replace statistical analysis or clinical study interpretation. These remain human tasks. At best, AI can help *describe* what test results show (if given the numbers). Some visionary predictions (e.g. Jama experts) talk about AI doing system role breakdown or risk analysis ([32] www.jamasoftware.com), but in 2026 such uses are still experimental. Most companies we spoke with do not rely on AI for any judgment calls on safety margins or equivalence decisions – those are reserved for engineers and clinicians.

- **Overly Long or Unstructured Prompts**. LLMs have context limits. In 2026, even GPT-4 may struggle with multi-thousand-word prompts (like feeding it an entire 510(k) module at once). Sponsors find it works better to break tasks into smaller queries. For example, asking for a summary of just the "intended use" section, rather than giving it the whole introduction. If the prompt is too broad, the AI's output can become generic. Thus, "tell ChatGPT my entire device story" is ineffective. Instead, focused questions ("List differences between Device A and B") yield better results. Senior QA staff report that they treat each submission section as a separate AI project rather than let a single prompt handle everything.

- **Ignoring Regulatory Style**. AI may produce text that is too informal or not in line with FDA's expected tone. For example, it might use the passive voice differently or misplace certain footnote styles. Each branch of FDA (CDRH, CBER, CDER) has style nuances. In practice, regulatory writers often need to re-edit AI text to match the formal style (e.g. rewriting "we did X" to "the device sample underwent X testing"). Doing this is part of the human review process, but one cannot skip it just because AI did the first draft. Companies that tried skipping final edits found their submissions got queries about inconsistencies.

- **Neglecting Data Security and Privacy**. Although not a 510(k) submission technique per se, letting proprietary information leak into AI servers is a serious misstep. By 2026, most pharma/medtech companies have policies forbidding direct ingestion of confidential documents into public LLMs. If authorship tasks involve trade secrets or patient data, they must use private instances or ensure content is de-identified. Failure to do so risks breaching HIPAA (if PHI is involved) or IP theft concerns. This is often governed by company policy rather than law, but it is a non-trivial operation issue. The LinkedIn example explicitly avoided giving ChatGPT any *protected data*; it only used general scientific queries.

In summary, **AI works best for well-defined, narrow tasks where outputs can be easily verified**. It is *not* a replacement for expert judgment. The community's pragmatic consensus is that AI is a **writing assistant and research aide**, not a decision-maker. This keeps results in line with regulatory expectations: tools may change, but accountability and proof remain with the human sponsor ([6] pharmacystandards.org).

Regulatory and Future Implications

The growing use of AI in 510(k) contexts raises a number of forward-looking questions:

- **Regulatory Notification**: Sponsors wonder whether FDA will notify them when an AI-assisted review is used on their submission. To date, FDA has not proposed requiring disclosure of its use of Elsa or other tools to applicants. Industry surveys indicate that companies *hope* for transparency so they can understand any AI influence on a decision. The meddeviceonline article explicitly asks: "If a developer experiences delays…and AI was involved, what kind of visibility will that developer have?" ([37] www.meddeviceonline.com). **Implication:** We may see FDA issued guidelines clarifying the role of AI in review, possibly requiring examiners to note AI use in minutes or decision letters by the late 2020s.

- **Regulation of AI Tools Themselves**: Right now FDA does *not* regulate tools like ChatGPT — only the device. However, if a sponsor relied on an AI consultant that made a mistake, the sponsor bears the risk. It's conceivable regulators might eventually require documentation of the "AI chain of custody" (e.g. logs of AI edits) to audit filings. In the interests of transparency, some advise building an "audit trail" of AI usage, even think of AI as a sometimes-auditable Contract Research Organization ([38] pharmacystandards.org).

- **AI in Global Submissions**: Many devices are cleared internationally. A US 510(k) dossier often has counterparts (e.g. Canada, EU). By 2026, Health Canada requires PCCPs for ML devices and strict validation steps ([17] aihealthcarecompliance.com). The EU is likely to apply similar standards under the AI Act and MDR/IVDR. If a US sponsor uses an AI tool to write a 510(k), they must consider how (or if) to disclose that in EU or Canadian submissions. Harmonization is needed: it is conceivable that regulators may jointly advise on AI-assisted writing protocols.

- **Ethics and Data Governance**: A consistent theme from 2026 expert panels is that ethical and IP considerations lag behind technical adoption ([39] www.jamasoftware.com). Companies must implement training on what data is fed to AI, and ensure no copyrighted medical images (e.g. X-rays) are unintentionally scraped into public models. Standards like ISO 42001 for AI governance (in development) may soon be cited during audits. Another ethical point is bias: if an LLM was trained on literature predominantly from one country or demographic, it might inadvertently overlook differences. Regulatory submissions must avoid any AI-introduced bias (for example, if training sets lacked certain patient populations). Sponsors are advised to track the provenance of AI outputs ("metadata ribbons" containing model version, data source, etc.) so they can justify their trust in the results ([40] aihealthcarecompliance.com).

- **Impact on Review Workload**: If FDA's internal AI indeed cuts reviewer busywork, we might see a **shift in FDA staffing or priorities**. Short-term, AI should free reviewers to focus on higher-risk issues. But long-term, if submissions move faster, FDA might allocate more resources to post-market surveillance. For industry, it might reduce the number of rounds of "Additional Information" letters, as AI tools could preemptively spot inconsistencies before submission ([41] www.meddeviceonline.com). However, until the data is published, this remains speculative.

- **Future of AI in Submissions**: The AI tools themselves will continue to evolve. We can anticipate specialized LLMs trained on regulatory text (maybe even a future "RegulatorGPT") that more reliably answer CDRH questions. There is talk of AI agents that could autonomously navigate the FDA database. However, each advancement will be weighed against regulatory caution: e.g., "AI methods that predict patient outcomes can accelerate product development, but if used in labeling, they need thorough validation" ([42] pmc.ncbi.nlm.nih.gov). In essence, the AI-assisted submission process is likely to mature into a **standardized workflow**: companies will codify how they use AI at each step (with SOPs, controls, and audits). The early adopter case studies (like Sclerosafe) will become textbook examples as more firms gain experience.

Conclusion

In 2026, AI-assisted 510(k) submissions are a reality but not a panacea. The weight of evidence indicates that generative AI can **significantly expedite certain parts** of the regulatory documentation process – especially literature review, first-draft writing, and data analysis – yet it introduces new responsibilities. Sponsors gain time and insight, but not without reining in the technology. All successful applications involve multi-step human checks and alignment with regulatory expectations of data integrity and auditability ([7] pharmacystandards.org) (www.pharmanow.live).

Regulators have signaled openness to AI's potential: FDA's pilot projects and guidance on AI devices show a forward-leaning attitude ([1] www.fda.gov) ([11] www.akingump.com). Yet the FDA and other agencies continue to stress that **AI is a tool, not an author** ([6] pharmacystandards.org) ([7] pharmacystandards.org). The onus is fully on the sponsor to use AI responsibly. Industry must therefore invest in training, SOPs for AI use, and possibly new validation steps for AI-derived content.

Looking forward, we expect "what works" to mean approaches that blend AI's strengths with rigorous engineering controls. Advances in explainable AI, API-based access to regulated data, and international standards (like ISO 42001 or FDA's forthcoming AI principles) will shape the landscape. The core message is: **AI can work ⇒ only under human control and regulatory alignment**.

Sponsors who grasp this duality – leveraging AI for speed while safeguarding correctness – stand to innovate faster and more cost-effectively. Those who ignore it risk delays or, worse, submission failures from unchecked AI errors. In the highly regulated 510(k) environment, there is no substitute for diligence. With careful implementation, however, AI assistance in 2026 offers a *rare win*: better submissions in less time, ultimately benefiting patients by bringing safe, effective devices to market more quickly.

References: Cited sources are denoted inline. These include FDA announcements ([1] www.fda.gov) ([43] www.fda.gov), peer-reviewed analyses ([12] www.mdpi.com), industry studies ([4] innolitics.com) ([5] innolitics.com), and expert commentary ([30] www.meddeviceonline.com) (www.pharmanow.live) ([6] pharmacystandards.org) ([31] www.jamasoftware.com). Each claim above is grounded in documented evidence or authoritative industry reporting, as indicated.

External Sources

[1] https://www.fda.gov/news-events/press-announcements/fda-announces-completion-first-ai-assisted-scientific-review-pilot-and-aggressive-agency-wide-ai#:~:In%20...

[2] https://www.linkedin.com/posts/lirontayeb_chatgpt-empowering-our-510k-fda-submission-activity-7078695336227086336-Zq0K#:~:cruci...

[3] https://www.linkedin.com/posts/lirontayeb_chatgpt-empowering-our-510k-fda-submission-activity-7078695336227086336-Zq0K#:~:accur...

[4] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:,with...

[5] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:Medic...

[6] https://pharmacystandards.org/caidra-examination/section-4-4-fda-ema-view-on-ai-assisted-documentation/#:~:Their...

[7] https://pharmacystandards.org/caidra-examination/section-4-4-fda-ema-view-on-ai-assisted-documentation/#:~:,as%2...

[8] https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821/#:~:Recen...

[9] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:We%20...

[10] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:,Any%...

[11] https://www.akingump.com/en/insights/ai-law-and-regulation-tracker/fda-finalizes-guidance-for-ai-enabled-medical-devices#:~:On%20...

[12] https://www.mdpi.com/2079-9292/13/3/498?trk=public_post_comment-text#:~:of%20...

[13] https://pharmacystandards.org/caidra-examination/section-4-4-fda-ema-view-on-ai-assisted-documentation/#:~:,comp...

[14] https://www.fda.gov/news-events/press-announcements/fda-announces-completion-first-ai-assisted-scientific-review-pilot-and-aggressive-agency-wide-ai#:~:FDA%2...

[15] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:,is%2...

[16] https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices#:~:Intel...

[17] https://aihealthcarecompliance.com/weekly-news-and-updates/#:~:The%2...

[18] https://www.linkedin.com/posts/lirontayeb_chatgpt-empowering-our-510k-fda-submission-activity-7078695336227086336-Zq0K#:~:estab...

[19] https://www.linkedin.com/posts/lirontayeb_chatgpt-empowering-our-510k-fda-submission-activity-7078695336227086336-Zq0K#:~:Anato...

[20] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:In%20...

[21] https://www.jamasoftware.com/blog/2026-predictions-for-medical-device-life-sciences-ai-wearables-and-navigating-regulatory-change/#:~:Jakob...

[22] https://www.fda.gov/news-events/press-announcements/fda-announces-completion-first-ai-assisted-scientific-review-pilot-and-aggressive-agency-wide-ai#:~:%E2%8...

[23] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:SaMD%...

[24] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:Prede...

[25] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:Fast%...

[26] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:match...

[27] https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/#:~:Image...

[28] https://www.linkedin.com/posts/lirontayeb_chatgpt-empowering-our-510k-fda-submission-activity-7078695336227086336-Zq0K#:~:%F0%9...

[29] https://www.linkedin.com/posts/lirontayeb_chatgpt-empowering-our-510k-fda-submission-activity-7078695336227086336-Zq0K#:~:compl...

[30] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:To%20...

[31] https://www.jamasoftware.com/blog/2026-predictions-for-medical-device-life-sciences-ai-wearables-and-navigating-regulatory-change/#:~:Adam%...

[32] https://www.jamasoftware.com/blog/2026-predictions-for-medical-device-life-sciences-ai-wearables-and-navigating-regulatory-change/#:~:Mike%...

[33] https://www.jamasoftware.com/blog/2026-predictions-for-medical-device-life-sciences-ai-wearables-and-navigating-regulatory-change/#:~:Dan%2...

[34] https://www.jamasoftware.com/blog/2026-predictions-for-medical-device-life-sciences-ai-wearables-and-navigating-regulatory-change/#:~:Tom%2...

[35] https://pharmacystandards.org/caidra-examination/section-4-4-fda-ema-view-on-ai-assisted-documentation/#:~:,take...

[36] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:submi...

[37] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:,deve...

[38] https://pharmacystandards.org/caidra-examination/section-4-4-fda-ema-view-on-ai-assisted-documentation/#:~:,as%2...

[39] https://www.jamasoftware.com/blog/2026-predictions-for-medical-device-life-sciences-ai-wearables-and-navigating-regulatory-change/#:~:Purvi...

[40] https://aihealthcarecompliance.com/weekly-news-and-updates/#:~:regul...

[41] https://www.meddeviceonline.com/doc/ai-enters-the-k-submission-world-what-device-developers-need-to-know-0001#:~:AI%20...

[42] https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821/#:~:AI%20...

[43] https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices?trk=article-ssr-frontend-pulse_little-text-block#:~:Artif...

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at https://intuitionlabs.ai/contact for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.