IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Shadow AI in Biotech: Risks of Unsanctioned Usage

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-03-30 · shadow ai · biotech ai · ai governance · unsanctioned ai · data privacy · chatgpt risks · ai compliance

Original article: https://intuitionlabs.ai/articles/shadow-ai-biotech-risks



In this article

1. Executive Summary
2. Introduction: Background on AI Adoption in Biotech
3. The “Shadow AI” Phenomenon in Biotech
4. Case Studies and Examples of Shadow AI Incidents
5. Analysis of Shadow AI Usage: Data and Surveys
6. Risks and Implications of Shadow AI in Biotech
7. Discussion of Shadow AI in Biotech Workflows
8. Data Analysis: Quantifying the Shadow
9. Case Study: Compliance Implications (Shadow AI in Regulated Workflows)
10. Mitigation Strategies and Future Directions
11. Conclusion

Executive Summary

“Shadow AI” – the unauthorized, informal use of AI tools by employees – is rapidly emerging as a critical issue in the life sciences and biotech industries. A preliminary July 2025 Project NANDA study reported that only **5% of custom enterprise AI tools reached production**, distinguishing these tools from general-purpose chatbots (Project NANDA report). Its deployment estimates are directional and interview-based, with samples varying by tool category and definitions of success differing across organizations. This is not a universal enterprise-pilot failure rate. Employee adoption is a separate measure: IBM’s published analysis concerns **1,000 full-time American office workers** (IBM). Employees may turn to unsanctioned generative AI when approved tools do not meet their needs. Surveys indicate this **shadow AI** phenomenon is widespread: roughly **half of all employees** in large organizations admit to using unapproved AI tools at work (www.csoonline.com), with a large portion of those users occasionally **copying sensitive or proprietary data into free AI chatbots**. For example, one security survey found **49% of workers** use AI tools without employer approval (www.csoonline.com), and LayerX’s 2025 report, based on browsing telemetry from its enterprise customers, found that **77% of users pasted data into GenAI tools**, a general data-pasting measure rather than a sensitive-data-sharing rate (LayerX report). In life sciences specifically, concerns about data privacy and compliance have led many firms to ban ChatGPT usage (e.g. **65% of top pharma companies** had banned ChatGPT by 2024 (www.fiercepharma.com)). Nonetheless, usage persists: in a recent industry survey **over half** of life science professionals reported using ChatGPT at least occasionally (www.fiercepharma.com).

This report examines the phenomenon of shadow AI in biotech in depth. We first define what constitutes shadow AI and why it arises, drawing on industry surveys and case studies. Next, we survey the evidence – both quantitative and qualitative – of how employees in life sciences and related fields are using AI informally, often by pasting proprietary data into consumer-grade tools. We present illustrative examples of recent incidents (drawing from technology, pharma, financial services and healthcare) where unauthorized AI use led to data leaks or compliance breaches. We analyze the specific risks of shadow AI in regulated biotech settings, including loss of intellectual property, violation of HIPAA/GDPR, and erroneous “hallucinations” of AI tools. We

discuss the structural factors (e.g. underinvestment in AI governance, user demand for productivity) that drive informal AI adoption. Finally, we outline mitigation strategies: the need for clear policies, approved AI platforms with enterprise controls, training, and monitoring.

Throughout, this report emphasizes **evidence-based** insights with data, expert commentary, and concrete examples. Key statistics are summarized in tables and discussed in the text. We incorporate perspectives from security researchers, industry surveys (e.g. by BlackFog, LayerX, Salesforce, IBM), news reports (Reuters, tech media) and expert analyses. The narrative presents both the motivations behind shadow AI and its substantial implications for data security and compliance in life sciences. By shedding light on this "hidden revolution" in enterprise technology, we aim to provide biotech leaders and security teams with a comprehensive understanding of the problem and avenues for addressing it.

Introduction: Background on AI Adoption in Biotech

Biotech and life sciences have long been recognized as fertile ground for AI-driven innovation, from accelerating drug discovery to optimizing clinical trials and personalized medicine. Industry analysts predict that generative AI could generate **$60–110 billion per year** in value for the life sciences sector (www.mckinsey.com) (www.salesforce.com). However, unlocking that value is proving elusive: McKinsey reports that while most life science companies have experimented with generative AI, **only about 5% have realized substantial business value** from it (www.mckinsey.com). Similarly, industry surveys show a "long tail" of stalled AI pilots and a lack of enterprise-wide deployment. The preliminary July 2025 Project NANDA report found that **5% of custom enterprise AI tools reached production**, a different measure from the performance of general-purpose chatbots (Project NANDA report). The deployment figures are directional estimates from interviews, with varying category sample sizes and differing definitions of success. Life-sciences decision-makers themselves sense this gap: a 2024 industry poll found that **half of companies had banned ChatGPT outright** due to data security concerns, yet **more than half of professionals still used it (sometimes daily)** (www.fiercepharma.com). In short, **formal AI programs in biotech are often underwhelming**, while employee demand for AI tools soars.

This gap between lofty AI ambitions and real deployments has given rise to "shadow AI". We define shadow AI as the **unsanctioned, ungoverned use of AI tools by employees**, usually consumer or free-tier systems, outside of official IT channels or without explicit management approval. (It is analogous to the earlier concept of "shadow IT" when workers bypassed corporate software restrictions.) In this scenario, employees might log into ChatGPT.com or deploy ChatGPT browser extensions, or use other public LLM-based assistants, to expedite tasks – often pasting work-related content into these tools. The content could be anything from drafting an email to summarizing research, but crucially it may include **sensitive or proprietary data**. Because these tools lie outside corporate control, the data flows they create are invisible to the organization's security and compliance systems.

Several factors drive shadow AI in biotech:

- **Urgent productivity pressures.** Life science workers juggle complex tasks – synthesizing research proposals, analyzing data, writing reports – and often face tight deadlines. Generative AI can dramatically accelerate text-heavy tasks (summarization, drafting, data interpretation). Employees naturally gravitate to the *most convenient* solution: if corporate tools are missing or unwieldy, they turn to the tools they can access immediately (free LLM chatbots, personal AI assistants).
- **Limited approved alternatives.** Many biotech companies have been slow to deploy enterprise-grade AI tools, fearing compliance and accuracy issues. Meanwhile, consumer AI services (ChatGPT, Bard, etc.) have proliferated. IBM's analysis of **1,000 full-time American office workers** found that **80% use AI at work**, while **22% rely exclusively on tools provided by their employers** (IBM). These worker-level measures do not establish that the remaining respondents use only personal tools or use AI without authorization.
- **Lack of guidelines or enforcement.** Few organizations have yet established comprehensive AI policies. One survey noted that **fewer than 60% of life sciences companies have given any guidance** on safe ChatGPT use (www.fiercepharma.com). Without clarity, employees make ad-hoc choices. In some cases, especially C-levels, management may tacitly permit it in exchange for speed gains (www.csoonline.com).
- **Technical blind spots.** Traditional data-loss prevention tools and network monitors often cannot see inside an HTTPS ChatGPT session. Thus businesses are largely unaware of how much data leaks out through shadow AI.

In the regulated environment of biotech – dealing with PHI/HIPAA, clinical trial data, trade secrets, IP – the stakes are high. The use of unapproved, black-box AI tools introduces unique **compliance and IP risks** (discussed in detail below). Nonetheless, as the examples and data in this report will show, shadow AI is not only present but growing. Understanding its contours is essential for any life sciences organization seeking to leverage AI responsibly and protect its most sensitive assets.

The "Shadow AI" Phenomenon in Biotech

Shadow AI in biotech refers to generative AI use by life-science professionals that occurs outside formal IT channels and policies. This often involves **individuals using consumer AI chatbots, free tools, or personal AI accounts with proprietary or sensitive data**. Example scenarios include:

- **Sales or marketing reps** using ChatGPT to draft emails or sales scripts, occasionally pasting product launch timelines or competitor intelligence into prompts.
- **Clinical trial coordinators** feeding portions of patient data or trial protocols into AI for summarization.
- **Research scientists** copying snippets of proprietary research, code, or chemical structures into AI tools to check hypotheses or generate code.
- **Regulatory or quality teams** drafting documents by asking AI tools to parse regulatory text or compile bullet points, sometimes using confidential internal guidelines as input.

In regulated industries like pharma and medical devices, these uses often arise not from malice but from the need to "get work done" more efficiently. A recent profession-focused blog described coordinators **pasting 40-page clinical study proposals into ChatGPT for quick summaries** (www.bplogix.com). Such tasks are

indeed time-consuming if done manually. The blog comments, on behalf of many in medical affairs, note: *"When workflow platforms don't offer AI capabilities, teams don't simply accept slower timelines. They find workarounds."* (www.bplogix.com). This reflects one driver of shadow AI: where official systems lack AI features, creative staff plug the gap with external tools.

Indeed, anecdotal industry evidence suggests shadow AI is pervasive in pharma commercial teams too. As one industry observer notes: *"I didn't meet a single person who hasn't used ChatGPT, Gemini, or some other AI tool"* among pharma field reps (medium.com). They use it for meeting preparation, call summaries, and account research – albeit "quietly" and "not officially" (medium.com). This story is echoed by many in biotech: employees see AI as an "efficiency superpower" and use it implicitly.

Quantitatively, multiple surveys document unsanctioned AI use. BlackFog's research release in 2026 reported that **49% of its 2,000 respondents used AI tools not sanctioned by their employer** (BlackFog). LayerX's 2025 enterprise-customer browsing telemetry found that **77% of users pasted data into GenAI tools**, rather than showing that 77% shared sensitive data (LayerX report). IBM's analysis of 1,000 full-time American office workers reported **80% using AI at work** and **22% relying exclusively on employer-provided tools** (IBM). The latter does not measure the proportion using AI without authorization.

In life sciences, contemporaneous coverage of ZoomRx's survey published on **April 17, 2024** reported that **25% of respondents used ChatGPT more than once a week** (ZoomRx-retained survey coverage). The aggregate policy and usage findings do not establish that individual respondents violated an employer's ban. McKinsey reported that **32% of respondents had taken steps to scale genAI** (McKinsey).

The mechanics of shadow AI make it inherently "invisible". Unlike installing unauthorized software, employees use AI in seconds via web browsers or mobile apps. A risk report explains: *"Shadow AI is nearly invisible: it happens inside a browser tab, in seconds, with zero installation required… no log entry screams 'breach in progress.'"* (www.phishdefense.com). Data exit quietly via encrypted channels. The following discussion uses a historical survey and hypothetical scenarios to frame questions about unapproved AI use.

Case Studies and Examples of Shadow AI Incidents

The historical survey below concerns AI-use policies and reported usage.

Year	Organization / Industry	Survey Findings
2024	Biotech/Pharma (ZoomRx survey)	**25% of respondents used ChatGPT more than once a week**, according to contemporaneous coverage published April 17, 2024 (ZoomRx-retained survey coverage). Aggregate survey results do not establish whether particular respondents violated their employers' policies.

Table 1: *Historical survey findings on AI-use policies and reported usage.*

For training, consider a hypothetical sales representative submitting a client pipeline to an unapproved AI service, or an HR manager submitting salary spreadsheets. These are illustrative scenarios, not documented incidents. A recent enterprise study found **33% of employees admit to sharing research data or code with unapproved AI**, and about **23–27% have fed financial or HR data** into such tools without realizing the security implications (www.csoonline.com). Because these things happen in personal browser tabs or mobile apps, they fall outside corporate DLP systems – Chief Information Security Officers often only discover the problem when regulators or customers lose trust.

Analysis of Shadow AI Usage: Data and Surveys

A wealth of recent surveys and studies documents the rise of unsanctioned AI usage among employees. In aggregate, these data reveal how pervasive shadow AI is and what kinds of data are at risk.

- **Prevalence of Unapproved AI Use:** A BlackFog survey of 2,000 corporate workers found *49% admit to using AI tools at work without employer approval* (www.csoonline.com). Notably, among executives and senior managers the tolerance was even higher (~66–69% "okay" with it) (www.csoonline.com) – indicating that the fad for speed often comes from the top down. Separately, IBM's published analysis covers **1,000 full-time American office workers**: **80% use AI in their roles**, and **22% rely exclusively on employer-provided tools** (IBM). These figures do not establish a three-quarters personal-tool or unauthorized-use rate.
- **Frequency and Volume:** Among those who do use AI, usage is frequent. The BlackFog data note that 86% of employees use AI at least weekly (www.csoonline.com). In LayerX's enterprise-customer browsing telemetry, **ChatGPT reached 43% enterprise-user penetration** (LayerX). This measures the share of users accessing ChatGPT, rather than its share of corporate AI sessions.
- **Types of Data Shared:** Perhaps most alarming, a large share of these AI interactions involve confidential or regulated data. BlackFog's survey found that **33% of employees admitted to putting proprietary research or datasets into unsanctioned AI tools**, 27% copied employee payroll or performance data, and 23% disclosed internal financial details (www.csoonline.com). LayerX's report similarly observed that ~22% of pasted text was sensitive (PII, financial figures, etc.) (www.techrepublic.com). Table 2 below summarizes key statistics from multiple studies on employees' shadow AI habits:

Statistic	Value (Approx.)	Source
Employees using AI at work (any tool)	~80% of office workers (www.ibm.com)	IBM / Censuswide (2025)
Employees admitting to unauthorized AI use	49% (general workers) (www.csoonline.com)	BlackFog (2026)
Employees who think speed outweighs security risk	60% (would accept risk) (www.csoonline.com)	BlackFog (2026)

Statistic	Value (Approx.)	Source
Users pasting data into GenAI tools (any data)	77% in LayerX enterprise-customer browsing telemetry (LayerX report)	LayerX Security (2025); not a sensitive-data-sharing rate
GenAI-facilitated data exfiltration proportion	32% of unauthorized data moves (www.techrepublic.com)	LayerX (2025)
Top pharma firms banning ChatGPT use	65% (top-20 companies) (www.fiercepharma.com)	ZoomRx / Fierce Pharma (2024)
Life-sciences staff using ChatGPT at least monthly	~50% (survey respondents) (www.fiercepharma.com)	ZoomRx / Fierce Pharma (2024)
Custom enterprise AI tools reaching production in Project NANDA's preliminary study	5%; directional interview-based deployment estimate, with varying category samples and success definitions (Project NANDA report)	Project NANDA (July 2025); not a universal pilot-failure or ROI rate

> **Table 2:** *Key statistics on shadow AI and AI adoption in enterprise and life sciences contexts. "Unauthorized AI" refers to use of external AI tools not approved by the company.*

The table underscores that a large fraction of corporate workers use AI tools outside official channels, often sharing regulated or proprietary data. For example, the LayerX data show **71.6% of generative AI access is via non-corporate (personal) accounts** (www.techrepublic.com), entirely bypassing enterprise identity controls. This "shadow SaaS" phenomenon is not limited to AI; it resembles earlier trends where employees used unsanctioned cloud apps for work.

Risks and Implications of Shadow AI in Biotech

Shadow AI poses several acute risks for the biotech and life sciences industries. Unlike benign productivity tools, unsanctioned AI can inadvertently expose a company's most sensitive assets. The main categories of risk include:

- **Intellectual Property (IP) Leakage:** Before submitting proprietary material to an AI service, review the provider, plan and settings, including retention and potential training use. In ChatGPT, users can turn off "Improve the model for everyone"; new conversations then remain in history but are not used to train ChatGPT (OpenAI). Treat external disclosure, retention, training and possible memorization as separate review questions. Consumer privacy controls alone should not be treated as organizational authorization to share confidential information.
- **Compliance and Regulatory Violations:** HIPAA distinguishes PHI from information properly de-identified through Expert Determination or Safe Harbor; the Privacy Rule does not restrict the latter's use or disclosure (HHS de-identification guidance). For a covered entity or business associate using a cloud provider to process ePHI on its behalf, a HIPAA-compliant BAA and compliance with the other HIPAA Rules are required (HHS cloud guidance). An impermissible disclosure requires an incident-

specific assessment: HHS permits documentation of a low probability of compromise or an applicable exception to establish that breach notification is not required (HHS breach guidance). Review separate contractual, research and other applicable requirements before sharing information.

- **Legal and Ethical Liability:** Beyond data protection laws, shadow AI can trigger other legal issues. For example, if AI is used to draft regulatory documents or safety reports without oversight, any factual errors or hallucinations can lead to misinformed decisions. In biotech, imagine using AI to interpret clinical trial inclusion criteria or to prepare regulatory submissions – any mistake could jeopardize compliance.
- **Security Threats:** Exposing internal data to third-party AI models invites malicious exploitation. If proprietary code or credentials are leaked, attackers can use that information to launch tailored attacks (e.g. spearphishing with insider data). Security vendors note that **exposed AI prompts** may be "training data" for adversaries. One vendor warns that leaked data can "profile and target an organization" allowing extortion or network breach (www.csoonline.com). Furthermore, shadow AI expands the surface for social engineering: attackers create fake AI tools to harvest credentials (as noted by PhishDefense (www.phishdefense.com)).
- **Loss of Audit and Quality Control:** Before using AI in a regulated workflow, determine which source documents, prompts, outputs and human review records need to be preserved, and check whether the proposed process preserves them.
- **Competitive Intelligence Leakage:** Shadow AI potentially allows industrial espionage. When an employee inputs strategic plans, M&A documents, or drug pipeline information into an external AI, that data could subsequently be accessed (legally or illegally) by competitors. In highly competitive therapeutic areas, even a few details leaked can be costly. The outsider contractor might be scraped or attacked to retrieve such prompts.

In summary, **shadow AI magnifies the risks already present in biotech**. These industries handle extremely sensitive and regulated data—ranging from clinical trial patient records to novel biomolecular designs. Adding a layer of clandestine sharing with public AI providers threatens far-reaching impacts: from major data breach fines to irreversible IP loss. Yet employees frequently underestimate these dangers. In interviews, many workers expressed the view that "it's no big deal" and expect employers to simply "turn a blind eye" as long as they meet objectives (www.csoonline.com). This attitude increases complacency.

An executive summary from security experts is sobering: "This should raise serious red flags for security teams; there must be greater oversight and visibility into these security blind spots" (www.csoonline.com). For biotech companies, shadow AI is now one of those blind spots. With an average breach cost in pharma around **$5 million** (moduscreate.com), and potential regulatory penalties, even a single incident can be devastating. Moreover, since many pilots and innovation efforts already fail, shadow AI usage often goes undetected; leaders may wrongly assume that limited official adoption means no one is using AI at all.

Discussion of Shadow AI in Biotech Workflows

Shadow AI's specific manifestations in biotech come in various forms. Below we analyze some common use-cases by function or workflow, illustrating both why employees do this and how it can go wrong.

- **Research and Development:** Biotech R&D generates vast textual data (articles, patent text, lab protocols, etc.). Researchers might use ChatGPT to **summarize scientific papers or patents**, speeding literature reviews. For example, an R&D scientist could paste a journal abstract or a PPT on a new protein target into ChatGPT for explanation. While this saves time, it risks embedding any novel hypotheses or unpublished results in the AI model's memory. Similarly, bioinformaticians have tried feeding gene sequences or chemical structures into AI for analysis. If they input proprietary sequence data (e.g. from a non-public genome project), it leaks to the AI vendor. To illustrate, imagine a team researching a new antibiotic compound, feeding its SMILES string to an LLM to predict activity; their formula might then be implicitly shared back out.
- **Clinical Operations:** Shadow AI use is prevalent in trial management and patient data tasks. Intensive administrative paperwork (protocols, regulatory forms, consent documents) might get uploaded to AI by coordinators for auto-completion. Survey data or case reports might be summarized using ChatGPT. The bplogix blog highlighted a scenario where a coordinator uses ChatGPT to parse and summarize 40-page investigator-initiated study (IIS) proposals (www.bplogix.com). For a covered entity, disclosure of PHI must be permitted or required by the Privacy Rule, or authorized in writing by the individual or their personal representative (HHS).
- **Regulatory Affairs / Quality Assurance:** Teams drafting regulatory submissions (eCTDs) or QA reports might ask AI to help structure documents or translate complex regulations. If they paste internal R&D data or proprietary quality metrics, that confidential info escapes. For instance, a team preparing a drug safety report might use ChatGPT to rephrase liability language, inserting key adverse event counts in the prompts. This could expose safety data or trade secrets. In manufacturing QA, even SOP text or non-public stability data could be input, which again risks data egress.
- **Sales, Marketing and Commercial:** Pharma sales reps and marketing analysts are already heavy users of data analytics – shadow AI is a natural extension. Reps might feed HCP (healthcare professional) lists or competitive intelligence into ChatGPT to draft messaging. Market analysts might ask AI to generate campaign emails, inadvertently including proprietary pricing models or strategic plans in prompts. The PhishDefense article notes examples like salespeople pasting **the entire quarterly forecast or pipeline** into AI to get quick summaries for meetings (www.phishdefense.com). In biotech commercial teams, such behavior easily duplicates competitor data exposure risks seen in tech.
- **Finance, HR, and Admin:** While not biotech-specific, companies often see HR or finance staff paste sensitive HR data (salary ranges, performance reviews) or budget spreadsheets into AI for analysis or writing. These were precisely the categories flagged in the BlackFog survey (www.csoonline.com). For

example, an HR manager might ask an AI to draft job descriptions and, in doing so, feed employee performance ratings. This user probably doesn't even know it's malicious – just a productivity hack. But the result is highly confidential data floating outside.

Properly de-identified information is no longer PHI, and the Privacy Rule does not restrict its use or disclosure (HHS).

It's also worth noting that for many users, shadow AI *does succeed at its intent*: it makes tasks faster or better (to an extent). The BlackFog CEO observes that senior leaders prioritize "efficiency gains and personnel cost savings" over privacy concerns (www.csoonline.com). These productivity gains fuel the behavior. If the only focus were on raw efficiency, enforcement would look like fighting gravity. The challenge is balancing these gains with governance.

In CSO's reporting, BlackFog founder and CEO Darren Williams said senior executives were often reluctant to acknowledge their AI use (CSO). Thus shadow AI can be a tacit, if unstable, status quo inside organizations. Contemporaneous coverage of ZoomRx's 2024 survey reported that some 80% of respondents considered AI overrated, while 25% used ChatGPT more than once a week (ZoomRx-retained survey coverage).

Data Analysis: Quantifying the Shadow

Beyond anecdote and individual incidents, the data paints a clear picture of shadow AI's pervasiveness and risk profile in biotech enterprises. We have already cited the broad surveys above. Here we delve deeper:

- **Survey Scope:** BlackFog's research covered employees in the **UK and US**, with 1,000 respondents in each country at companies with more than 500 employees. **49% of the 2,000 respondents reported using AI tools not sanctioned by their employer** (BlackFog research and methodology). These cross-industry findings are not a biotech-specific measurement or an APAC comparison.
- **Data Sensitivity Levels:** We can categorize the leak risk by data classification. At the lowest level, employees have pasted generic corporate information (like product descriptions) into public models. At the next tier, they've entered non-public strategic information (e.g. pipeline priorities). At higher tiers, they have input regulated data categories: patient info, proprietary formulas, or confidential legal documents. From survey data: about **40% of uploaded files contained PII or payment data** (www.techrepublic.com), while **22% of inputs were particularly sensitive (regulatory or patient-related)**. For a covered entity, disclosure of PHI must be permitted or required by the Privacy Rule, or authorized in writing by the individual or their personal representative (HHS).
- **Security Findings:** Tools to detect shadow AI are just emerging. Some security vendors now offer AI-specific DLP: scanning browser queries for keywords or model-access patterns. For instance, LayerX's network monitoring (used for their report) logged **over 50% of pasted AI prompts as containing corporate content** (www.techrepublic.com). This suggests traditional DLP underreports the issue by missing the human copy-paste step. Other threat intelligence notes a surge in attacker interest: e.g. some groups set up fake AI chatbots to phish employees (as PhishDefense points out) (www.phishdefense.com).

- **Shadow AI vs. Official AI Projects:** McKinsey reported that **32% of respondents had taken steps to scale GenAI** (McKinsey). This is an organizational scaling measure, not an estimate of the share of AI activity that is unofficial, decentralized or uncontrolled. It should be considered separately from surveys asking employees about unsanctioned use.

Case Study: Compliance Implications (Shadow AI in Regulated Workflows)

To illustrate the compliance context more concretely, consider the following hypothetical vignette:

> **IIS (Investigator-Initiated Study) Triage.** A pharmaceutical company receives investigator proposals for independent trials. Coordinator Alice opens a 45-page PDF of a new study concept. Each section (endpoints, timeline, budget) must be reviewed. Normally, Alice would manually extract key data into the company's intake system. But today, she copy-pastes whole sections into ChatGPT and asks for a summary of primary endpoints and budget info. The AI returns a neat summary in seconds. Alice then lightly edits it, populates the entry fields, and forwards to reviewers.

In this hypothetical workflow, the company should review whether sharing the proposal is authorized and whether Alice preserves the source material, AI-generated summary and her edits for reviewers. Assess applicable permissions, confidentiality obligations and privacy requirements before drawing a legal conclusion.

Before adopting this workflow, define how reviewers will check the summary against the proposal, document corrections and record any approval decision.

Mitigation Strategies and Future Directions

Given the depth of the shadow AI challenge, biotech organizations must adopt a multi-pronged strategy. This involves not only technology but also policy, education, and a shift in culture. Key approaches include:

- **Establish Clear AI Usage Policies.** Companies must define what AI tools are approved, and what data can be used. For example, categorize data by sensitivity and specify that anything beyond unclassified text requires only enterprise-authorized AI. The policy should clarify that personal or free-tier AI accounts are prohibited for confidential data. This might include rules such as “No PHI/PII allowed in ChatGPT; R&D formulas can only be used in on-prem LLM platforms.” Updating the Acceptable Use Policy to explicitly address generative AI is critical, as recommended by security experts (www.phishdefense.com). As one advisor notes, the first step is to align AI use with existing regulations: *“If human activity without the use of AI is regulated, then the use of AI should similarly be regulated”* (www.paulhastings.com).
- **Deploy Enterprise-Grade AI Solutions.** Rather than have employees resort to consumer tools, provide approved AI platforms that meet security requirements. For example, use ChatGPT Enterprise or specialized biotech AI assistants (with organization-specific fine-tuning and on-prem deployment). Ensure these solutions have disabled model training on corporate data and allow audit logging. The goal

is to offer an official AI experience that is at least as user-friendly as consumer bots. IBM suggests investing in infrastructure to “make internal data cleaner and more accessible to AI models” so that employees prefer the sanctioned AI (www.ibm.com). In practice, this could mean championing federated learning models or secure LLM APIs within the corporate network, or using secure enclaves.

- **Technical Controls and Monitoring.** Integrate DLP and network monitoring tuned for AI usage. Modern CASB/DLP solutions can now flag patterns like ChatGPT access or keywords. For example, monitor for bulk copy-paste actions or uploads to known AI domains. According to expert recommendations, enterprises should log attempted uses of external AI and set alerts (“block suspicious activity patterns”) (centrexit.com). Employing AI-savvy security tools can help: e.g., new products that classify data and flag AI-transmission attempts at the client side.
- **Employee Education.** Many users simply do not realize the risks. Training should go beyond “Don’t paste secrets into ChatGPT.” Workers need role-specific examples (e.g. a drug safety reviewer, a lab scientist) illustrating how certain actions (even well-intended) violate policies. Use hypothetical, role-specific scenarios to make the training concrete. As one expert advises, training must be ongoing (at least quarterly) and hands-on, not a one-time lecture (centrexit.com). The aim is to instill a culture of caution: make employees manually inspect any data before using an AI tool, similar to checking attachments for viruses.
- **Governance Framework Integration.** Life sciences companies should explicitly incorporate AI into their existing quality/compliance frameworks. This means involving AI risk in pharmacovigilance and regulatory reviews (www.bplogix.com) (www.paulhastings.com). A suggested approach is a three-stage governance lifecycle (concept review, design/deploy, continuous monitor) as outlined by Paul Hastings: before introducing an AI use case, evaluate its risk-benefit with legal and regulatory teams; deploy with validation checks; and continuously monitor performance and compliance (www.paulhastings.com). Embedding these steps into project gateways (like requiring an “AI approval” in stage-gate processes) can prevent runaway shadow projects.
- **Leadership Alignment and Change Management.** Perhaps most importantly, the C-suite must acknowledge shadow AI as a real phenomenon (rather than blaming users for a “security breach” mentality). Senior leaders should champion secure AI adoption and communicate both the benefits and risks of generative AI. Analysis shows that **employee enthusiasm for AI is high (97% see productivity upside) but leadership buy-in is often lacking** (www.ibm.com). Closing that gap means showing how official AI platforms will deliver similar speed while protecting IP. It may require dynamic change management, allocating sufficient change budgets (the McKinsey genAI guide suggests roughly 5x the tech spend) to ensure adoption.
- **Risk Assessment and Scenario Planning.** Security teams should conduct threat modeling for AI scenarios. For life sciences, important questions include: How could clinical trial data exit? What would an accidental AI-induced error in toxicology data look like? Companies might run red-team exercises where testers script “An employee has deployed a rogue AI to expedite a task” and see how incident response unfolds. As one analyst notes, tools now exist to simulate phishing via fake AI-landing pages (www.phishdefense.com) – such drills could sensitise staff to the perils of unvetted AI links.

Future directions: The landscape will continue evolving. AI vendors are introducing enterprise offerings (OpenAI's "ChatGPT for Healthcare" in 2026 is one example), which can help mitigate some risks if used properly (www.hipaajournal.com). Regulators are also moving—e.g., FDA's proposed frameworks for AI in submissions (www.fda.gov). Companies that build robust AI governance now may gain competitive advantage (by safely leveraging AI's benefits) while others exposed to shadow AI are more likely to suffer setbacks. The training and systems implemented today will shape whether generative AI remains a hidden liability or becomes a transparent, productive asset in biotech workflows.

Conclusion

Shadow AI represents a critical blind spot for biotech enterprises. In theory, the pharmaceutical and medical industries embrace cutting-edge tech – yet in practice, AI adoption is often piecemeal. IBM reports that nearly 40% of surveyed workers prefer external AI solutions for their better features (IBM). This self-help approach can boost productivity, but it also **sows risk**.

For biotech leaders, this raises urgent questions. How many scientists are inputting unvetted results into ChatGPT right now? Does the legal department know its existence? The chances are high that "shadow AI" use is already part of the fabric of R&D, patient support, compliance and commercial work. Recognizing this, companies must pivot from every-user-will-cheat mentality to one of managed adoption. The pathway forward is clear: robust governance frameworks, enterprise AI solutions, and a culture of security. Only then can organizations capture AI's promise without letting their crown jewels escape in the process.

In summary, shadow AI in biotech is characterized by **informal AI usage, lack of governance, and significant risk exposure**. It demands the same level of attention as any other security threat. Paul Hastings recommends applying existing regulations to AI wherever applicable as the first step in AI governance (Paul Hastings). While the technology landscape is changing fast, the principle remains: data that leaves uncontrolled is data forever lost. By confronting shadow AI through policy, technology, and education (and citing principles from industry research), biotech firms can tame the shadow and turn AI into a disciplined, yet dynamic, asset.

Sources: This report synthesizes findings from industry surveys, security research, regulatory guidance, and case reports. Key references include enterprise AI adoption studies (www.ibm.com) (www.forbes.com), security analyses of unsanctioned AI use (www.csoonline.com) (www.phishdefense.com), and life sciences industry surveys (www.fiercepharma.com) (www.bplogix.com). Read surveys and security-vendor telemetry according to their methods and scope; treat preliminary research as provisional, and distinguish legal commentary and hypothetical examples from documented incidents (www.hipaajournal.com) (centrexit.com). Detailed citations are provided throughout to ensure verifiability.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://files-scs.pstatic.net/2025/08/31/YcL3JEAasR/The%20GenAI%20Divide%20STATE%20OF%20AI%20IN%20BUSINESS%202025.pdf

2. https://www.ibm.com/think/insights/rising-ai-adoption-creating-shadow-risks
3. https://www.csoonline.com/article/4124775/roughly-half-of-employees-are-using-unsanctioned-ai-tools-and-enterprise-leaders-are-major-culprits-2.html
4. https://go.layerxsecurity.com/hubfs/LayerX_Enterprise_AI_and_SaaS_Data_Security_Report.pdf
5. https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey
6. https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry
7. https://www.salesforce.com/news/stories/life-sciences-ai-survey-insights-2025/two-years-to-go_-nearly-100-of-life-science-leaders-say-ai-agents-will-become-essential
8. https://www.bplogix.com/blog/shadow-ai-in-life-sciences-the-compliance-gap-nobody-talks-about
9. https://medium.com/%40dscheine/shadow-ai-is-already-inside-pharma-commercial-teams-15601d345913
10. https://www.blackfog.com/blackfog-research-shadow-ai-threat-grows
11. https://www.zoomrx.com/reports/Endpoints_News_article_17_April_2024.pdf
12. https://www.phishdefense.com/blog/shadow-ai-employees-leaking-company-secrets-to-chatgpt
13. https://layerxsecurity.com/blog/ai-is-now-the-1-data-exfiltration-vector-in-the-enterprise-and-nobodys-watching
14. https://www.techrepublic.com/article/news-employees-share-company-secrets-on-chatgpt
15. https://openai.com/index/how-chatgpt-protects-privacy
16. https://www.hhs.gov/hipaa/for-professionals/special-topics/de-identification/index.html
17. https://www.hhs.gov/hipaa/for-professionals/special-topics/health-information-technology/cloud-computing/index.html
18. https://www.hhs.gov/hipaa/for-professionals/breach-notification/index.html
19. https://moduscreate.com/blog/security-risks-ai-life-sciences
20. https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html
21. https://www.paulhastings.com/insights/client-alerts/building-a-comprehensive-ai-governance-framework-in-life-sciences
22. https://centrexit.com/employee-chatgpt-data-leakage-proprietary-code
23. https://www.hipaajournal.com/is-chatgpt-hipaa-compliant
24. https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions
25. https://www.forbes.com/councils/forbestechcouncil/2026/01/28/95-of-ai-pilots-fail-a-practical-roadmap-for-healthcare-ai

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.