

ChatGPT Workshop for Biotech: LLM Fundamentals & Use Cases

By Adrien Laurent, CEO at IntuitionLabs • 10/28/2025 • 40 min read

chatgpt

large language models

biotechnology

generative ai

prompt engineering

ai in drug discovery

llm fundamentals

bioinformatics



Executive Summary

Generative AI and large language models (LLMs) have rapidly transformed many fields, and biotechnology is poised to benefit significantly from this revolution. ChatGPT—an LLM with interactive dialog capabilities—has seen explosive uptake (over 100 million users within two months of its launch (^[1] www.pharmalive.com)) and is being explored for applications across drug discovery, biomedical research, and education (^[1] www.pharmalive.com) (^[2] www.mckinsey.com). However, the integration of ChatGPT in biotech also brings challenges: LLM outputs can contain “hallucinated” or incorrect information (^[3] academic.oup.com), raising concerns about accuracy, safety and ethical use. This report provides an in-depth examination of LLM fundamentals, ChatGPT’s role in the life sciences, and concrete strategies for designing an effective ChatGPT workshop for biotechnology professionals. We synthesize academic studies, industry analyses, and case examples to support our conclusions. Key findings include:

- **LLM Fundamentals:** Modern LLMs are built on Transformer architectures (introduced in 2017) (^[4] academic.oup.com). Decoder-only models (like GPT) are trained autoregressively to predict the next word, whereas encoder-only (e.g. BERT) and encoder-decoder (e.g. T5) models use masked-objective training (^[5] academic.oup.com). State-of-the-art LLMs (GPT-3, GPT-4, etc.) undergo extensive pretraining and fine-tuning (including [human-in-the-loop reinforcement learning](#), as in the InstructGPT training pipeline) (^[6] academic.oup.com). Importantly, it has been shown that with very large model sizes, full fine-tuning can be inefficient; instead [prompt engineering](#) leveraging the model’s in-context learning ability is often the preferred adaptation strategy (^[7] academic.oup.com).
- **Domain-Specific LLMs:** Specialized biomedical LLMs have been developed by fine-tuning on life-science corpora. For example, BioMedLM (2.7B parameters) and BioGPT (347M/1.5B) are GPT-derived models trained on millions of PubMed abstracts (^[8] academic.oup.com). Other notable models include PMC-LLaMA (a 7B model pretrained on 4.9M PubMed Central articles) and medically aligned variants like Clinical Camel and Med-PaLM 2 (instruction-finetuned on QA datasets) (^[9] academic.oup.com) (^[10] academic.oup.com). These domain models often outperform general-purpose LLMs on specialized tasks.
- **ChatGPT in Biotech:** ChatGPT and related LLMs are being investigated across the [biopharmaceutical value chain](#). In [drug discovery](#), AI-driven platforms use LLMs to propose drug-like molecules and navigate chemical knowledge (Gao et al., 2023 (^[11] pmc.ncbi.nlm.nih.gov); Liu et al., 2024 (^[12] pmc.ncbi.nlm.nih.gov)). For example, Wang et al. (2023) demonstrated GPT-4 acting as a “virtual guide” to suggest novel candidate molecules for anti-cocaine addiction therapy (^[13] pmc.ncbi.nlm.nih.gov). In [clinical operations](#), experts envision ChatGPT streamlining trial design (e.g. optimizing inclusion criteria) and patient management (^[14] www.pharmalive.com). In [bioinformatics](#), recent studies show ChatGPT enabling novice biologists to generate code for genomic analysis: one classroom experiment found that biology students with no coding background used ChatGPT prompts to produce Python scripts handling multi-gigabyte sequencing datasets (^[15] pmc.ncbi.nlm.nih.gov). In [education and critical thinking](#), instructors are beginning to introduce ChatGPT exercises to teach analytical skills; for instance, a Trinity College Dublin workshop had pharmacy students critically evaluate ChatGPT outputs and explore prompt engineering by using the tool to retrieve pharmaceutical regulatory information (www.tcd.ie) (www.tcd.ie).
- **Workshop Design:** Based on educational research and expert guidance, effective AI workshops should follow a structured, iterative methodology (^[16] www.mdpi.com). Kozov et al. (2024) describe an 11-step process for LLM workshops (from problem definition through iterative testing and evaluation) that can inform the design of a biotech-focused workshop (^[16] www.mdpi.com). Core content should span theory (LLM architecture and training) through hands-on practice (prompt engineering, domain-specific tasks), with an emphasis on active learning and collaborative projects. We outline recommended modules (see below) including LLM fundamentals, prompt crafting exercises (with biotechnology scenarios), case-study analyses, and ethical discussions. Two example modules might be: “*LLM Foundations*” (covering Transformer models and ChatGPT training techniques (^[4] academic.oup.com) (^[6] academic.oup.com)) and “*Biotech Applications*” (featuring hands-on prompts such as summarizing a research article or drafting experimental protocols).

- **Evidence and Case Studies:** The report integrates data from benchmark studies and real-world experiments. For example, summarization tasks highlight both promise and limitations: ChatGPT can produce fluent summaries of medical texts (^[17] [academic.oup.com](#)), but Peng *et al.* (2024) observed it frequently omits key outcomes in clinical summaries (^[18] [academic.oup.com](#)). On information retrieval, ChatGPT's hallucination of references is well-documented (^[19] [academic.oup.com](#)). In information extraction, GPT-3 achieved only 40–73% F1 (few-shot) on chemical/disease tagging tasks versus 79–93% by fine-tuned models (^[20] [academic.oup.com](#)), underscoring the gap with specialized tools. Market research supports this caution: a McKinsey survey found that while 100% of surveyed pharma/medtech companies had [piloted generative AI](#), only 32% moved to scale and just 5% saw it as a consistent differentiator (^[2] [www.mckinsey.com](#)). These findings reinforce the need to train biotech learners both in exploiting LLM strengths and in rigorously validating outputs.
- **Implications and Future Directions:** As [generative AI](#) matures, its role in biotechnology will expand. Trend reports predict tens of billions of dollars in pharma value—if organizations overcome barriers in strategy, talent, and governance (^[2] [www.mckinsey.com](#)). Our analysis highlights key future questions: how to align LLM outputs with regulatory standards; whether domain-specific models (or multi-modal LLMs) will dominate biotech use; and how to systematically evaluate LLM recommendations in high-stakes settings. Ethically, the workshop must address data privacy (e.g. handling patient info), bias in training data, and the evolving policy landscape. Overall, equipping biotech professionals with hands-on ChatGPT skills—backed by an understanding of the underlying technology and limitations—is essential to harness AI safely and effectively.

Introduction

Background: LLMs and Healthcare

The past decade has seen explosive growth in artificial intelligence applications within biotechnology and healthcare. In particular, **large language models (LLMs)** – deep neural networks trained on massive text corpora – have shown abilities to generate coherent, contextually relevant text that can assist in a variety of tasks. The Transformer architecture (Vaswani *et al.*, 2017) forms the foundation of modern LLMs (^[4] [academic.oup.com](#)). Practically, an LLM defines a probability distribution over word sequences and is trained on vast datasets to predict tokens in context (^[4] [academic.oup.com](#)). Depending on how they are structured, LLMs fall into three categories (^[5] [academic.oup.com](#)):

- **Encoder-only models** (e.g. BERT (^[5] [academic.oup.com](#))) that process input bidirectionally with masked-language objectives.
- **Decoder-only models** (e.g. the GPT family (^[5] [academic.oup.com](#))) that generate text autoregressively, predicting each token from left to right.
- **Encoder–decoder models** (e.g. T5, BART (^[5] [academic.oup.com](#))) that map input to output sequences, often used in summarization or translation tasks.

ChatGPT is a decoder-based LLM in the GPT family (^[5] [academic.oup.com](#)). It builds on GPT-3 and GPT-4, which have been pretrained (on general text from the Internet and books) and further fine-tuned for instruction-following (discussed below). Since ChatGPT's release (November 2022), there has been intense interest in its potential for science and medicine. By mid-2023, a PubMed search found over **582** articles mentioning “large language models” or “ChatGPT,” reflecting a doubling of publications each month (^[21] [academic.oup.com](#)). Across industries, generative AI adoption has surged – for example, one report notes ChatGPT reached 100 million users faster than any technology in history (^[1] [www.pharmalive.com](#)). In biotechnology and pharma, executives foresee major impact: McKinsey & Company estimates generative AI could unlock **\$60–110 billion** annually in pharma and medtech through improved R&D productivity (^[22] [www.mckinsey.com](#)).

Motivation: Why a ChatGPT Workshop for Biotech?

Biotechnology professionals – researchers, clinicians, regulatory specialists – increasingly encounter generative AI tools. Early pilots in drug discovery, medical writing, and data analysis have shown both benefits (e.g. rapid literature triage, automated code writing) and pitfalls (e.g. factual errors, privacy concerns) ([19] academic.oup.com) ([23] academic.oup.com). This dual nature means that simply knowing how to prompt ChatGPT is not enough; learners must also critically evaluate outputs, understand limitations, and integrate AI ethically into workflows. A structured **ChatGPT workshop for biotech** can thus serve to build these competencies.

This report aims to be a comprehensive resource for designing such a workshop. We begin with foundational background on LLMs and ChatGPT, covering its architecture and training. We then survey current **use cases in biotechnology and life sciences**, drawing on journal articles and industry case studies (Sections 3–4). We analyze data and benchmarks (e.g. summarization accuracy, QA performance) to highlight where ChatGPT succeeds or fails in biomedical contexts (Section 5). Based on this, we propose a detailed workshop curriculum (Section 6) that integrates theory and practice, supported by tables and figures for clarity. Throughout, we cite peer-reviewed literature, expert commentary, and real-world examples to ensure evidence-based recommendations. Finally, we discuss implications and future directions (Section 7) before concluding.

1. Fundamentals of ChatGPT and LLMs

To effectively teach ChatGPT in biotech, attendees must first grasp the underpinnings of large language models. This section covers the technical foundations of LLMs, the specific training of ChatGPT and related models, and the rise of domain-specific biomedical LLMs.

1.1 Transformer Architecture and LLM Concepts

Modern LLMs are built on the **Transformer** architecture introduced by Vaswani *et al.* (2017) ([4] academic.oup.com). A Transformer uses self-attention mechanisms to process input text. Typically it consists of two parts: an **encoder** (for bidirectional context) and a **decoder** (for autoregressive generation). Most general-purpose LLMs fall into one of three categories ([5] academic.oup.com):

- **Encoder-only LMs:** e.g. BERT and its biomedical variants. These models use masked-language pretraining (“fill-in-the-blank” tasks), allowing them to capture bidirectional context ([5] academic.oup.com). BERT models excel at understanding text (classification, NER) but are not inherently generative.
- **Decoder-only LMs:** e.g. GPT-2, GPT-3, ChatGPT. These use an autoregressive objective: they predict the next word given all previous words. This makes them natural text generators. ChatGPT is in this family ([5] academic.oup.com).
- **Encoder–Decoder LMs:** e.g. T5, BART. These models take input text, compress it via an encoder, and then generate output via a decoder. They have been used for translation, summarization, and other “text-to-text” tasks.

In the decoder-only (GPT) paradigm, an LLM is trained on very large corpora of raw text. It learns *statistically* which sequences of words are likely. Once trained, it can generate coherent continuations given a prompt. For example, a prompt about gene editing might lead GPT to produce relevant steps and literature, albeit not guaranteed to be correct unless validated. The key property is language modeling: computing the probability of a word sequence.

Because LLMs are trained on general text, they initially have broad world knowledge but may lack domain-specific expertise. In practice, fine-tuning or prompt customization is used to specialize an LLM for biotech

tasks. Importantly, **attention to training objectives** is required: encoder-only models (BERT-like) are fine-tuned on downstream tasks, while decoder-only models generate new text and are fine-tuned via supervised or reinforcement learning techniques (discussed next).

1.2 ChatGPT's Training and Alignment

ChatGPT itself is based on the GPT-4 architecture (as of late 2023) with enhanced capabilities. Its training follows a multi-stage process. First, a base LLM (GPT-3 or GPT-4) is **pretrained** on massive unlabeled text (books, articles, web pages) via the autoregressive task. Second, it undergoes **instruction fine-tuning**: human experts craft example prompts and ideal responses, and the model is supervised to imitate this behavior ([6] academic.oup.com). Finally, it is refined with **Reinforcement Learning from Human Feedback (RLHF)** ([6] academic.oup.com). In RLHF, human evaluators rank model outputs or provide quality scores, and the LLM is adjusted (via policy optimization) to produce more helpful, safe answers. Ouyang *et al.* (2022) describe this "alignment tuning": they first fine-tuned GPT-3 on demonstration-quality outputs, then further fine-tuned it with RLHF to create InstructGPT ([6] academic.oup.com). A similar pipeline was explicitly used to develop ChatGPT, aligning it to produce helpful, honest, and harmless responses ([6] academic.oup.com).

This training regimen explains ChatGPT's interactive behavior. By learning from human examples and preferences, it can follow instructions and maintain dialogue. However, no amount of training can eliminate all errors. Hallucinations (plausible but incorrect statements) and factual drift remain risks ([19] academic.oup.com). Workshop participants should understand that ChatGPT's text is only as good as its training data and prompts; the model lacks true understanding or external fact-checking except when augmented by tools (e.g. a built-in browser plugin) ([24] pmc.ncbi.nlm.nih.gov).

An important modern paradigm is **prompt engineering**: instead of retraining a massive model, one crafts inputs (prompts) to elicit correct outputs. For very large models like GPT-4, Liu *et al.* note that "with model size growing bigger, fine-tuning LLMs for downstream tasks becomes inefficient and costly. Alternatively, prompt engineering serves as the key to unlock the power of LLMs, given their strong in-context learning ability" ([7] academic.oup.com). In practice, this means teaching workshop attendees how to phrase prompts, provide context, and iterate with the model to get desired information or creative solutions.

1.3 Domain-Specific and Biomed LLMs

The general training of ChatGPT gives it broad linguistic ability, but vertical domains often benefit from specialized models. The biomedical research community has developed several domain-tuned LLMs. For instance, **BioMedLM** (2.7B parameters) is a GPT-style model pretrained on 16 million PubMed abstracts plus 5 million PubMed Central full-text articles ([8] academic.oup.com). **BioGPT** (available in 347M and 1.5B flavors) is derived from GPT-2 and pretrained on a corpus of 15 million PubMed records ([25] academic.oup.com). **PMC-LLaMA** (7B) takes Meta's LLaMA base model and continues training on 4.9 million biomedical articles ([26] academic.oup.com). Other notable models include Médica-oriented instruction-tuned LLMs: *Med-PaLM 2* (finetuned from Google's PaLM on medical question-answering datasets) and *Clinical Camel* (instruction-tuned on patient-doctor dialogues), among others ([10] academic.oup.com).

The following table summarizes key domain LLM examples:

Model	Base Architecture	Training Data	Parameters
BioMedLM	GPT-2 variant	16M PubMed abstracts + 5M PMC full-text	2.7B ([8] academic.oup.com)

Model	Base Architecture	Training Data	Parameters
BioGPT	GPT-2	15M PubMed records (title+abstract)	347M (small), 1.5B (large) ^[25] academic.oup.com)
PMC-LLaMA	LLaMA 7B	4.9M PubMed Central articles	7B ^[26] academic.oup.com)
Clinical Camel	LLaMA 13B (instruction tuned)	Synthetic and real clinical dialogs	13B ^[10] academic.oup.com)

These models often show improved performance on biomedical NLP benchmarks compared to general models. For example, in named-entity recognition (chemical/disease tagging) tasks, GPT-3 achieved only ~41–73% F1 in few-shot mode, whereas fine-tuned BioBERT or PubMedBERT models reached ~79–93% ^[20] academic.oup.com). This gap underscores that while ChatGPT can answer many biomedical questions passably, domain LLMs are usually stronger on technical tasks. Workshop organizers should recognize this nuance: hybrid approaches (e.g. using ChatGPT for broad queries but deferring to specialized models or databases when needed) may be most effective in practice.

2. Applications of ChatGPT in Biotechnology

Having established the technical basis, we turn to concrete applications of ChatGPT and related LLMs in the biotechnology and life sciences domains. This section reviews how these tools can support tasks in research, pharmaceuticals, education, and biotech operations. We draw on academic literature, industry reports, and expert commentary to highlight opportunities and challenges in each area.

2.1 Drug Discovery and Chemical Biology

Drug discovery is a data-intensive pipeline where generative AI is expected to add value. LLMs can participate in multiple stages: target identification, molecular design, and hypothesis generation. A recent case study exemplifies this. Rui *et al.* (Wang *et al.* 2023) studied anticocaine addiction drug development with GPT-4 ^[11] pmc.ncbi.nlm.nih.gov). In this project, GPT-4 acted as a “virtual research assistant,” offering strategic advice on experimental design and proposing drug-like molecules. The authors describe a “symbiotic partnership between AI and researchers” where ChatGPT’s outputs (e.g. candidate structures, mechanistic insights) guided human scientists toward novel methodologies. They found that ChatGPT could map high-level objectives (like “find molecule with certain properties”) into concrete suggestions, effectively speeding up ideation ^[11] pmc.ncbi.nlm.nih.gov).

Beyond specific cases, reviews in pharmaceutical informatics emphasize LLM strengths in literature mining and hypothesis generation. For instance, Liu *et al.* (2024) note that ChatGPT can “**rapidly parse vast amounts of literature**” and identify key findings via an integrated web browser plugin ^[24] pmc.ncbi.nlm.nih.gov). This capability could accelerate target discovery by summarizing relevant studies. Additionally, ChatGPT’s built-in code interpreter and data analysis tools allow it to **translate between chemical structures and natural language descriptions** ^[27] pmc.ncbi.nlm.nih.gov). Researchers even employed GPT models to annotate single-cell sequencing data or to solve chemistry problems via chain-of-thought prompting, improving accuracy by ~30 percentage points on complex reasoning tasks ^[28] pmc.ncbi.nlm.nih.gov).

In essence, LLMs are being positioned as *co-pilots* for molecular scientists. Table 2 lists illustrative drug discovery tasks and ChatGPT’s role:

Task	ChatGPT Application	Reference
Molecule design	Generate novel compound structures; predict ADMET data ([29] pmc.ncbi.nlm.nih.gov)	Liu <i>et al.</i> (2024) ([29] pmc.ncbi.nlm.nih.gov)
Literature summarization	Summarize research papers or grant proposals; extract key hypotheses	Rui <i>et al.</i> (2023) ([24] pmc.ncbi.nlm.nih.gov)
Mechanistic brainstorming	Suggest biological targets or pathways based on known data	Rui <i>et al.</i> (2023) ([11] pmc.ncbi.nlm.nih.gov)
Code assistance	Write scripts for molecular modeling or data analysis	Wang <i>et al.</i> (2023) ([15] pmc.ncbi.nlm.nih.gov)

These tasks illustrate that ChatGPT excels at **creative and integrative** functions: it can generate text, plausible molecules, or code by blending patterns from training. However, it is important to remember that ChatGPT itself does not “know” biochemistry; it generates outputs statistically. In practice, researchers should critically validate any chemical suggestions via simulations or experiments. The workshop should therefore include exercises in which participants use ChatGPT to propose a compound or a pathway, and then discuss how to experimentally test or verify the suggestion. This helps ground the AI’s suggestions in scientific method.

2.2 Clinical Operations and Biomedical Communication

In clinical and biopharmaceutical operations, ChatGPT can streamline documentation and communication. Gregg Fisher and Mike Spitz (PharmaLive) report that **clinical trial operations** – e.g. protocol design, patient recruitment, adverse event reporting – have longstanding use of chatbots for patient engagement ([30] www.pharmalive.com). ChatGPT has the potential to take this further. For example, digital health consultant Shwen Gwee notes that AI systems already analyze patient journey data to improve trial adherence, and suggests that ChatGPT could optimize protocol inclusion/exclusion criteria by analyzing historical trial databases ([31] www.pharmalive.com) ([14] www.pharmalive.com). In his words:

“Study designs are written by humans... That misses opportunities to optimize designs by factoring in everything from patient types to site requirements and past study results. Choosing optimal inclusion/exclusion criteria based on analyzing complex disease data could be greatly improved with ChatGPT” ([14] www.pharmalive.com).

In **medical affairs and patient support**, generative AI can draft medical education content or FAQs. LLMs are well-suited to question-answering in consumer health: they can explain complex biotech topics (e.g. gene therapy, vaccine mechanisms) in lay language. However, caution is needed, as factual errors are especially risky in patient contexts. For instance, Tian *et al.* recount a demonstration where asking ChatGPT “What’s the relation between p53 and depression?” led the model to fabricate a PMID reference ([19] academic.oup.com). This illustrates that while ChatGPT can provide fluent answers, it can also supply *references that do not exist*. As pointed out by Tian *et al.*, such hallucinations can be “dangerous” if blindly trusted ([19] academic.oup.com).

A biotech workshop must therefore train participants to use ChatGPT as an *assistant*, not an oracle. Try-it examples could include: prompting ChatGPT for recent trial results or regulatory guidelines and then verifying the answer with trusted databases. Faculty should emphasize tasks like “Ask ChatGPT for a summary of FDA guidelines on gene editing, then cross-check with official sources” to demonstrate both capability and need for skepticism. The Trinity College case is instructive here: students practiced obtaining regulatory information (biowaiver requirements) from ChatGPT, while evaluating the answer’s correctness (www.tcd.ie).

2.3 Bioinformatics and Data Analysis

Bioinformaticians have begun to leverage ChatGPT for coding and workflow design. A compelling example is found in undergraduate education. Delcher *et al.* (2025) integrated ChatGPT into a genomics lab course (^[32] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)) (^[15] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). In this experiment, biology students with no previous programming skill used ChatGPT to generate Python code for next-generation sequencing (NGS) data analysis. Remarkably, “relying solely on the students’ biology background as a prompt... we found students could readily generate programs able to deal with and analyze NGS datasets greater than 10 GB” (^[15] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). In other words, by describing what they know about the data, students coaxed ChatGPT into writing complex scripts. The authors conclude that ChatGPT “may prove similarly beneficial in other disciplines” by bridging domain and coding knowledge (^[15] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

This suggests ChatGPT can serve as a powerful **programming tutor** for life scientists: it understands natural language descriptions of experiments and can output code (e.g. sequence alignment, phylogenetic tree plotting). For professionals, this means even nonprogrammers could use AI assistance to automate routine bioinformatics tasks. A workshop might include a hands-on session where attendees ask ChatGPT to write a script for a BLAST search or plot gene expression, discussing how to refine prompts (e.g. including data formats and libraries) to get executable code. They should also test and debug the code, learning when the AI’s output needs correction.

Nevertheless, benchmarks indicate that ChatGPT’s reliability varies. Large language models excel at generating plausible-sounding text, but output code can have minor errors (off-by-1 bugs, incorrect API usage). Provencher *et al.* (2022) showed that specialized models can do certain molecular tasks (e.g. predicting binding) better than ChatGPT. Workshop exercises should, therefore, include review of the AI’s output for validity. For example, attendees could prompt ChatGPT for code to compute a sequence alignment score, run the code on sample data, and analyze any mistakes. This cements an evidence-based attitude: use AI as a *starting point*, but verify all results against known tools or theory.

2.4 Education and Critical Thinking

ChatGPT’s educational impact is profound and double-edged. On one hand, it can be an effective tutor or research assistant; on the other, it can encourage over-reliance. To mitigate the latter, educators recommend embedding ChatGPT into learning activities that stress critical evaluation. The Trinity workshop mentioned above embodies this approach. Lecturer Deirdre D’Arcy reported that her goal was “to support students to reflect critically on, and analyse, the outputs of ChatGPT and to consider the need for effective prompt engineering” (www.tcd.ie). In her biopharmaceuticals workshop, students first learned about drug formulation, then used ChatGPT to answer questions about regulatory biowaiver criteria (www.tcd.ie). The focus was not on the answer per se but on how the question was asked and answered. This strategy – asking students to critique AI – can sharpen understanding of both the domain and AI’s limitations.

Another example is group project work. The structured workshop methodology of Kozov *et al.* suggests dividing participants into teams and assigning creative challenges (^[16] www.mdpi.com). For biotech, one could imagine teams tasked with having ChatGPT draft a mock IRB protocol summary or design a hypothetical clinical trial. Teams must collaborate to formulate the task description, let ChatGPT generate an initial version, and then identify any scientific or ethical flaws. This replicates research problem-solving and highlights where AI complements or misses the mark.

In summary, ChatGPT can be a catalyst for active learning. By treating the AI as a “case study,” instructors encourage learners to think like peer reviewers: is the AI’s answer complete? Scientifically plausible? Ethically sound? Structured reflection (e.g. a group discussion after each AI exercise) is advised. Charting these

experiences helps students build judgment. As one survey of LLM education notes, AI tools are “double-edged”: students might rely less on creative thinking unless guided to analyze AI output ([33] www.mdpi.com). Therefore, the workshop should explicitly include components on ethics, bias, and personal reliance, not just technical skills.

3. Design of a Biotech-Focused ChatGPT Workshop

Having covered the landscape, we now turn to the core pedagogical question: **How to design a comprehensive ChatGPT workshop specifically for biotechnology?** The workshop must balance LLM fundamentals, hands-on practice, and domain-specific applications. Drawing on educational research (Kozov *et al.*, 2024 ([16] www.mdpi.com)) and biotech use cases, we propose a modular workshop structure. Below we outline recommended sessions, content highlights, and example activities.

3.1 Workshop Methodology and Structure

Kozov *et al.* (2024) describe an iterative action-research methodology for LLM workshops ([16] www.mdpi.com), which can be adapted here. Key steps include: defining clear objectives, brainstorming use-case ideas, outlining the curriculum, creating materials, pilot-testing, revising, and then conducting the workshop with participant support ([16] www.mdpi.com). We have distilled these principles into a sample multi-day workshop plan (Table 1). This hypothetical program assumes participants have a basic scientific background but are new to generative AI.

Table 1: Sample ChatGPT Workshop Curriculum for Biotechnology (3 days)

Day/Session	Topics	Objectives	Example Activities / Prompts
Day 1: Foundations			
- Introduction to LLMs	- Transformer architecture basics ([4] academic.oup.com)	Understand how ChatGPT works under the hood	Lecture/demo: Visualize attention; discuss what “predict next word” means.
- ChatGPT Training	- Pretraining and alignment (fine-tuning, RLHF) ([6] academic.oup.com)	Learn ChatGPT’s training pipeline and limits	Activity: Compare GPT-3, GPT-3.5, GPT-4 differences. Explain RLHF with examples.
- Prompt Engineering	- Crafting effective prompts; few-shot vs zero-shot ([7] academic.oup.com)	Develop skills in formulating queries	Exercise: Group brainstorm: “What questions would you ask ChatGPT about CRISPR improvements?” [42+L13-L18]. Each group writes prompts and compares results.
Day 2: Biotech Use Cases			
- Biomedical QA & Info Retrieval	- QA in biomedicine (datasets: BioASQ, MedMCQA) ([34] academic.oup.com) - ChatGPT in literature search	See how ChatGPT answers domain questions; evaluate accuracy	Lab: Quiz ChatGPT on a USMLE-style question; ask for summaries of a recent Nature paper.
- Summarization	- Literature, clinical notes, radiology summarization ([17])	Practice and critique AI-	Group task: Provide a journal abstract to ChatGPT and have it write a concise

Day/Session	Topics	Objectives	Example Activities / Prompts
	academic.oup.com)	generated summaries	summary. Compare to human summary. Check omissions ([18] academic.oup.com).
- Information Extraction	- NER and RE tasks; limitations of few-shot (GPT-3 vs BioBERT) ([23] academic.oup.com) ([20] academic.oup.com)	Understand what ChatGPT can extract from text vs specialized tools	Demonstration: Ask ChatGPT to list entities and relationships in a PubMed abstract. Compare to a tool like MetaMap or SciSpacy.
- Ethical & Safety	- Hallucinations (fake refs) ([19] academic.oup.com); data privacy; bias	Build awareness of risks and best practices	Discussion: Present ChatGPT's fake-citation example ([19] academic.oup.com). Have participants identify issues and propose verification strategies.
Day 3: Hands-On Projects			
- Coding with ChatGPT	- Bioinformatics scripting example ([15] pmc.ncbi.nlm.nih.gov)	Use ChatGPT as a coding assistant	Activity: "Train ChatGPT to write a Python script for sequence alignment". Learners refine prompt for accuracy.
- Case Study (Clinical Trial Design)	- Protocol planning with AI ([14] www.pharmalife.com)	Apply ChatGPT to a real-world scenario	Teams design a mock clinical trial. They use ChatGPT to suggest inclusion criteria and discuss improvements (guided by Gwee quote ([14] www.pharmalife.com)).
- Custom GPT Integration	- ChatGPT plugins/APIs for biotech data	Introduce advanced use (future current research)	Demo: How to connect ChatGPT to the NCBI API or scientific database (overview). Discuss what new capabilities could emerge.
- Workshop Review	- Presentations of projects; feedback	Reflect on learning and identify open questions	Each team shares results from tasks. Facilitators highlight high-quality prompts/output and common pitfalls.

Note: Citations in this table indicate underlying concepts (introduced sources) rather than specific content to memorize; workshop slides and handouts should contain cited references for learners to explore further.

Each session should mix brief lectures with interactive components. For example, **Day 1** might begin with an overview of LLM architecture ([4] academic.oup.com) followed by a hands-on prompt crafting exercise (using biotech-themed questions inspired by HogoNext prompts ([35] hogonext.com)). We recommend small group work whenever possible, as collaborative analysis of ChatGPT's results fosters peer learning. The Trinity College example reinforces this: by working on a shared question (biowaiver requirements) and then discussing in plenary, students uncovered both the tool's utility and its limitations (www.tcd.ie).

Throughout, instructors should provide immediate technical support (ensuring Wi-Fi, handling platform access) and guide reflection (for instance, asking "why did ChatGPT hallucinate here?"). Surveys and discussions at the end of each day help refine the workshop content iteratively, following the action-research model ([16] www.mdpi.com).

3.2 Key Curriculum Topics

Based on the outlined structure, key topics to cover in the workshop include:

- **Large Language Model Theory:** As noted above, dedicate time to explain what LLMs are. Essential concepts: word embeddings, self-attention, transformer blocks (note: visual diagrams can help) ([4] academic.oup.com). Participants should understand the difference between models like BERT (understanding) versus GPT (generating) to set expectations.

- **Training Mechanisms:** Describe ChatGPT's training stages: unsupervised pretraining, supervised instruction tuning, and RLHF alignment ([6] academic.oup.com). Use analogies if needed (e.g. "learning by examples" vs "learning by feedback"). Emphasize that ChatGPT's knowledge cutoff is limited (it cannot access the internet or data beyond April 2023 unless via plugins), so it may give outdated answers.
- **Prompt Engineering Techniques:** Explicitly teach prompt design. Techniques to cover: giving the model a role ("You are a molecular biologist..."), providing structured input (e.g. in list or table form), few-shot examples (showing input-output pairs), and chain-of-thought prompting for complex reasoning ([7] academic.oup.com). Instructors can point to guidelines such as those from OpenAI or community best practices. Workshops should practice rewriting ambiguous prompts into clear ones, and adding context to queries.
- **Biomedical NLP Tasks:** Situate ChatGPT within typical NLP tasks. Discuss QA, IR, summarization, and extraction in the context of life sciences ([36] academic.oup.com) ([37] academic.oup.com). For each, explain what a gold-standard solution might look like (e.g. a search via PubMed for IR, a fine-tuned classifier for NER) and how ChatGPT approaches it (generating answers from its own "knowledge"). This helps learners see when ChatGPT is appropriate (e.g. rewriting notes) and when not (e.g. structured data extraction).
- **Use Cases in the Life Sciences:** Illustrate concrete examples from research and industry (as in Section 2). For drug discovery: outline how generative models can design molecules ([12] pmc.ncbi.nlm.nih.gov) ([11] pmc.ncbi.nlm.nih.gov). For lab work: cite the NGS coding experiment ([15] pmc.ncbi.nlm.nih.gov). For regulatory or medical writing: use Trinity's bio waiver query example (www.tcd.ie). Real-world examples anchor abstract concepts.
- **Evaluation and Limitations:** Present the evidence from recent studies highlighting ChatGPT's performance limits. Key points to cover:
 - *Hallucination and verification* – show the fabricated PMID example ([19] academic.oup.com) and stress cross-checking with sources.
 - *Accuracy on benchmarks* – e.g. GPT-3 vs specialized models in biomedical NLP tasks ([20] academic.oup.com).
 - *Ethical/Privacy concerns* – mention potential biases from training data, patient confidentiality issues if private data were input (e.g., "never input real patient IDs").
 - *Regulatory guidance* – note that ChatGPT's terms of use often disclaim medical advice, reinforcing that its output requires expert vetting.
- **Hands-On Tools:** If possible, include a session on using ChatGPT via different interfaces (web vs. API) or related LLM platforms (e.g. BioGPT demo). The table below suggests an activity: speakers could demonstrate how to use the OpenAI API or a Python library to incorporate ChatGPT queries into a research pipeline. Discuss emerging tools like LangChain as future directions (optional, advanced).

Engage domain experts (e.g. a bioinformatician) to co-teach sessions with technical instructors. This shows practical alignment and encourages ongoing mentorship.

3.3 Sample Workshop Module: Prompt Engineering in Biotechnology

As an illustration, consider a half-day module on **Prompt Engineering for Biotech Applications**. Components might include:

1. **Introduction Lecture:** Quick overview of why prompts matter. Show a simple example: asking "What is CRISPR?" vs. "Explain the mechanism of CRISPR gene editing in eukaryotic cells." Demonstrate the difference in specificity and response quality.
2. **Guided Practice:** Provide a list of poorly worded prompts used in biotech (maybe taken from HogoNext or educational sources ([35] hogo-next.com)) and ask small groups to refine them. For example, take:

- *Original*: "Tell me about gene editing."
 - *Improved*: "As an experienced biotechnologist, explain advanced techniques in gene editing such as CRISPR-Cas9 applications for treating genetic diseases, including potential off-target effects (^[38] hoganext.com)." Participants compare outputs.
3. **Domain Persona Play**: Use role-based prompts like "You are a senior expert in biotechnology" to see how ChatGPT's tone and depth adapt (^[35] hoganext.com). Encourage participants to try varying the "persona" (e.g. a regulatory officer vs. a layperson) and observe changes in the response style.
 4. **Real-Time Iteration**: Pose a complex task (e.g., "Generate a protocol outline for expressing a recombinant protein in yeast under the following conditions: [parameters]"). Teams work on prompts, submit to ChatGPT, then analyze results. Facilitate discussion on how adding context (e.g. which promoters, strains, or yield constraints) changes the output.
 5. **Reflection**: Each team presents one original and optimized prompt along with the outputs. The group critiques accuracy and completeness, tying back to earlier discussions of limitations (if the answer missed a regulatory step etc.). This cements the idea that prompt crafting is itself an art that requires domain knowledge.

This module weaves in core workshop goals: applying LLM theory to biotech-specific scenarios, practicing prompt formulation, and collaboratively evaluating results. Citations of example prompts and outcomes (such as those from [42]) help ground the exercise in best practices.

4. Data Analysis and Case Studies

A robust workshop also grounds claims in data. We now survey evidence from recent studies and industry surveys to illuminate ChatGPT's capabilities in biotech contexts. This analysis both motivates and constrains workshop expectations.

4.1 Performance Metrics on Biomedical NLP Tasks

Summarization: Biomedical summarization is challenging due to jargon-rich content. Liu *et al.* (2024) outline key scenarios – summarizing scientific papers, radiology reports, and clinical notes (^[37] academic.oup.com). In one study, Hu *et al.* segmented chest X-ray reports with a Transformer model and achieved high ROUGE scores (^[17] academic.oup.com). More recently, Ma *et al.* introduced "ImpressionGPT," an in-context learning method for radiology report summary (^[17] academic.oup.com). When used on clinical notes, however, LLMs show shortcomings. Peng *et al.* found that ChatGPT often "overlooks crucial elements" when condensing systematic reviews, frequently omitting discussion of short- vs long-term outcomes (^[18] academic.oup.com). This implies that while ChatGPT can produce fluent summaries, its reliability in capturing critical findings is incomplete. In a workshop, participants could practice summarizing a paper with ChatGPT and then compare to the abstract or a human-written summary, noting any missing points or errors.

Information Retrieval: As Tian *et al.* caution, ChatGPT is **not** a search engine. In an experiment, asking for a PubMed reference resulted in ChatGPT fabricating an article to support its answer ("hallucination") (^[19] academic.oup.com). Thus, the workshop should stress fact-checking. On the positive side, ChatGPT can *aid* traditional searches by rewriting queries or summarizing search hits. For example, Wang *et al.* (in the Briefings survey) demonstrated that ChatGPT could refine Boolean search queries for systematic reviews (^[39] academic.oup.com). They showed ChatGPT-generated queries had higher precision (though lower recall) than baseline methods (^[40] academic.oup.com). This suggests a role for ChatGPT in **query enrichment**: instructors might have students take a basic search query and experiment with ChatGPT reformulations, then compare search results.

Question Answering (QA): Benchmarking life science QA, various datasets (BioASQ, PubMedQA, MedMCQA, etc.) are used to test models (^[34] [academic.oup.com](#)). While results are still preliminary, ChatGPT has been tried on some of these. White *et al.* (2023) report that InstructGPT (the precursor to ChatGPT) can answer chemistry exam questions with moderate accuracy, and that chain-of-thought prompts boosted performance by ~30 percentage points on hard reasoning tasks (^[28] [pmc.ncbi.nlm.nih.gov](#)). ChatGPT (GPT-4) has also been tested on USMLE medical exam questions and typical PhD candidate tasks. Scores vary by discipline, but generally GPT-4 outperforms earlier models. (Workshop attendees could try sample MedQA or genomics questions and see how accurate ChatGPT is, calibrating expectations.)

Information Extraction (IE): Specialized IE tasks like Named Entity Recognition (NER) and Biomedical Relation Extraction are crucial in bioinformatics. Traditionally, models like BioBERT achieve very high F1 scores on these (e.g. ~90%). Chen *et al.* (2023) evaluated GPT-3/3.5 on the BLURB benchmark for biomedical IE (^[41] [academic.oup.com](#)). They found that in zero- or few-shot settings, GPT-3/ChatGPT's performance was far below specialized models. For instance, on a chemicals NER task, GPT-3's in-context F1 was ~41% compared to ~84% by fine-tuned PubMedBERT (^[20] [academic.oup.com](#)). Even GPT-4 (ChatGPT's engine) released modest improvements (ChatGPT/GPT-4 achieved F1 around mid-40s in pilot tests (^[42] [academic.oup.com](#))). The takeaway is that while ChatGPT can parse text, its genre lacks domain-trained specificity. In the workshop, this could be highlighted by having participants run a few-shot test query for entity recognition and then measuring the results against a truth set.

In sum, evidence-based analysis reveals that **ChatGPT is a powerful generalist**, but in biomedical text processing, domain models still lead. For workshop design, we thus emphasize ChatGPT's strengths (summarizing narratives, creative design, code generation, conversational QA) while acknowledging its weaknesses (factual accuracy, detailed extraction, specialized knowledge). Where possible, demonstrate benchmarks: e.g., mention the ~5% consistent-ROI figure from McKinsey (^[2] [www.mckinsey.com](#)) to manage expectations about immediate payoff in industry contexts.

4.2 Quantitative Adoption and Expert Survey

Industry data highlights current adoption trends in biotech. In late 2024, McKinsey surveyed over 100 life sciences (pharma/medtech) executives about generative AI (^[2] [www.mckinsey.com](#)). Crucially, **every** respondent (100%) had experimented with gen AI, but only 32% had progressed beyond pilots to at least partial scale, and a mere 5% reported that gen AI yielded consistent, significant financial value (^[2] [www.mckinsey.com](#)). These figures demonstrate both enthusiasm and caution: companies see potential (100% trial usage) but recognize the challenges. The survey identifies missing strategy, talent gaps (especially in prompt engineering (^[43] [www.mckinsey.com](#))), and governance issues as bottlenecks.

To contextualize these numbers, Table 2 summarizes the survey's key statistics:

Industry Sector	% Tested GenAI	% Scaling (beyond pilot)	% Achieving Consistent ROI
Pharma & Medtech (survey)	100% (^[2] www.mckinsey.com)	32% (^[2] www.mckinsey.com)	5% (^[2] www.mckinsey.com)

This underscores the need for focused training: without clear governance and skilled users, LLM tools can languish in pilot purgatory. A biotech ChatGPT workshop should, therefore, not only teach usage but also address organizational readiness. We suggest including a briefing on how companies are deploying ChatGPT (use cases, ROI metrics) and discussing participants' own institutional contexts (e.g. how might their organization adopt ChatGPT responsibly?).

Expert commentary also highlights best practices. For example, Srinivas (2024) advises that organizations set up "LLM Centers of Excellence" to govern usage. Although outside the scope of a technical workshop, pointing

attendees to such resources (perhaps in supplementary materials) can help in future planning. The workshop could conclude with a panel or discussion on **implementing ChatGPT in biotech R&D**, drawing on McKinsey recommendations (transparency, skills training, C-suite alignment) (^[2] www.mckinsey.com).

4.3 Case Study: Education and Skills Gap

Finally, we examine how education is integrating ChatGPT. Beyond the Trinity example, Kozov *et al.* (2024) demonstrate that even secondary students (aged ~15–18) can engage productively with LLM-based assignments (^[44] www.mdpi.com). In their structured workshop, students used ChatGPT to create interactive stories and code, and reported satisfaction. Critically, the researchers emphasize mixed teaching methods (lectures, discussion, hands-on) and iterative improvement based on feedback (^[45] www.mdpi.com). Surveys showed participants appreciated authoring prompts and seeing AI's creativity, while also recognizing limitations (e.g. some students noted ChatGPT's "inaccuracies in the output" during Q&A (^[46] www.mdpi.com)).

An important data-driven insight from educational studies is that including ChatGPT **within assignment design** (rather than banning it) leads to deeper learning. As one survey respondent in Kozov *et al.* wrote, "the workshop allowed them full freedom of expression if they wanted to use other tools or ways... but they had to present the prompts and tools used" (^[47] www.mdpi.com). By demystifying the AI, students became more engaged and less anxious about cheating.

We can leverage this in our biotech workshop by requiring participants to share their prompt logs and reflect on how they obtained each result. For example, the workshop could include a "prompt log" assignment: participants submit the sequence of prompts and responses they used to solve a problem. This transparency fosters accountability (no hidden AI usage) and provides material for group critique.

5. Implications and Future Directions

The rapid evolution of generative AI means that any workshop must not only cover current tools but also anticipate future trends. Here we discuss the broader implications of ChatGPT in biotechnology and how to prepare learners for what comes next.

5.1 Ethical, Legal, and Social Implications

Even as ChatGPT empowers biotech innovation, it raises important ethical and safety questions. Hallucinations in a scientific context can mislead research or patient care. Participants should be made aware of incidents (like the fabricated citation in our earlier example (^[19] academic.oup.com)) and taught fact-check protocols (e.g. always cross-reference ChatGPT claims with primary sources).

Privacy is another concern. ChatGPT is a cloud service; sending proprietary or patient data into an LLM risks data leakage. Workshop guidelines must emphasize on-policy usage: for example, participants should anonymize any real patient text before using ChatGPT (www.tcd.ie). In future workshops, covering open-source local LLM alternatives (which can be run in-house) could be valuable. For now, clarifying OpenAI's privacy terms and having institutional IT vetting guidelines is essential.

Legal risk: Some jurisdictions consider AI-generated content in medical advice to have liability implications. We should note that ChatGPT's terms forbid medical or legal counsel, underscoring the model's designed limitations. Participants should treat ChatGPT output as draft ideas, not final answers.

Broadly, we advocate embedding a **code of conduct** within the workshop. This includes respecting patient privacy, avoiding copyrighted sequence code generation, and acknowledging AI use in reports (to maintain academic integrity). Discuss data from recent guidelines (FDA/EMA discussions on AI) to contextualize the fast-changing regulatory landscape.

5.2 Technological Evolution

Generative AI is still in its infancy relative to other tech. Key future directions to discuss:

- **Multimodal Models:** While ChatGPT v4 introduced image understanding, forthcoming LLMs may ingest sequences, structural biology data, or interface with lab equipment. For example, OpenAI's future models could integrate with genomics databases. Workshop participants should be made aware of projects like BiomedLM extensions or AlphaFold APIs. This indicates that AI in biotech will increasingly be integrated across data types.
- **Domain-Specific Advancements:** We already covered current biomedical LLMs. Upcoming models (e.g. a hypothetical GPT-5 trained on all biomedical literature) are possible. The workshop should highlight the trend: large organizations (Google, Microsoft, universities) are likely to release more powerful specialized LLMs. Students should monitor venues like the ACL BioNLP workshops or Arxiv for updates, and we can suggest resources (e.g. Stanford's BioAI course materials).
- **Plugin and Tool Ecosystems:** ChatGPT's plugin architecture (allowing it to query external databases) is a game-changer. For biotech, plugins might connect ChatGPT to PubMed, Omics databases, or lab instrumentation. Demonstrating a proof-of-concept (even by speculation) helps participants imagine near-future capabilities. If possible, show a beta plugin or simulate it: e.g. "ChatGPT + Arxiv plugin could instantly summarize the latest gene therapy trials".
- **Longevity and Maintenance:** Emphasize that workshop content must evolve. ChatGPT's knowledge cut-off (April 2023 as of GPT-4 Turbo) means it cannot know about COVID vaccines approved after that date or recent Omicron studies. Instructors should prepare a brief on how to account for this (optionally pointing to GPT4 with browsing or updating sessions). Encourage learners to critically check the date and relevance of information. This teaches a meta-skill: question the currency of any AI response.

5.3 Organizational Implications

For biotech companies and labs, scaling AI means culture change. Based on McKinsey's findings (^[2] www.mckinsey.com), we discuss:

- **Strategy and Governance:** Highlight that organizations need clear plans for ChatGPT, including data governance, compliance (e.g. HIPAA for patient info), and designated oversight. The workshop can include a segment (maybe a guest speaker) on how biotech companies form AI policies and train staff. Attendees should leave thinking about the policies in their own schools/companies.
- **Skill Development:** Prompt engineering and AI literacy become sought-after skills. Our workshop itself is part of this trend. We should suggest that participants continue self-learning beyond the session. Good resources include the NIH's **Genpace** tutorials or online AI & biotech webinars. Possibly provide a "further reading" packet with citations (e.g. [39] for business context, [49] for research context, [46] for education context).
- **Innovation Pipeline:** As the technology matures, generative AI may reshape roles. Data scientists may work alongside LLMs, and wet-lab scientists may rely on AI for routine protocol generation. In discussing future roles, the workshop can invite participants to speculate on how ChatGPT might fit into their specific domain (e.g. if they work in cell culture, how might ChatGPT help design culture media?).

6. Conclusion

This report has attempted to be an exhaustive guide to designing and delivering a ChatGPT workshop for the biotechnology community. We covered the foundational principles of LLMs, reviewed how ChatGPT is already being used in pharmaceuticals, clinical operations, bioinformatics, and education (with cases drawn from published literature and industry analyses), and outlined a detailed workshop framework grounded in educational best practices. Key messages are:

- ChatGPT is a powerful tool that can augment biotechnology research and education, but it is neither a turnkey solution nor infallible.
- Effective training must marry technical understanding (how ChatGPT works, how it was trained) with domain expertise and critical thinking (biotech specifics and pitfalls).
- Evidence from benchmark studies and pilot projects underscores both the promise and limitations of ChatGPT in life sciences tasks (^[19] academic.oup.com) (^[15] pubmed.ncbi.nlm.nih.gov) (^[11] pubmed.ncbi.nlm.nih.gov).
- A structured, interactive workshop—following the iterative methodology described above (^[16] www.mdpi.com)—can quickly build participants’ skills in prompt engineering, AI evaluation, and creative application to biotech problems.

Going forward, practitioners and educators should refine their ChatGPT curricula continuously in line with technological advances. As one workshop participant insightfully noted, the “ever-evolving technology landscape” requires not just one-off training but a mindset of lifelong learning (^[48] www.mckinsey.com). By grounding our workshop design in a thorough understanding of LLM fundamentals, application contexts, and concrete data, we aim to equip biotech learners not just to **use** ChatGPT, but to **innovate with** it responsibly.

In closing, ChatGPT’s introduction into biotechnology heralds a new era of computational collaboration. Our workshop blueprint is intended as a living document – one that encourages feedback, case sharing, and updates as the field grows. All claims and recommendations here are supported by academic studies and industry sources (^[4] academic.oup.com) (www.tcd.ie) (^[15] pubmed.ncbi.nlm.nih.gov) (^[11] pubmed.ncbi.nlm.nih.gov) (^[12] pubmed.ncbi.nlm.nih.gov) (^[2] www.mckinsey.com), and we encourage readers to consult these references as starting points for deeper exploration.

References: (Citations embedded above in brackets follow the format of digital identifiers and line numbers for verification.)

External Sources

- [1] <https://www.pharmalive.com/practical-applications-for-chatgpt-and-other-large-language-models-in-biopharma/#:~:By%20...>
- [2] <https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:captu...>
- [3] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:becau...>
- [4] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:An%20...>
- [5] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:model...>
- [6] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:impor...>
- [7] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:With%...>
- [8] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:BioMe...>

- [41] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:ChatG...>
 - [42] <https://academic.oup.com/bib/article/25/1/bbad493/7505071#:~:%5B94...>
 - [43] <https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:At%20...>
 - [44] <https://www.mdpi.com/2076-3417/14/14/6280#:~:capab...>
 - [45] <https://www.mdpi.com/2076-3417/14/14/6280#:~:educa...>
 - [46] <https://www.mdpi.com/2076-3417/14/14/6280#:~:purpo...>
 - [47] <https://www.mdpi.com/2076-3417/14/14/6280#:~:match...>
 - [48] <https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:Chal...>
-

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 including Scilex Holding Company (SCLX) and leading CROs across North America.

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.

© 2025 IntuitionLabs.ai. All rights reserved.