

Understanding ChatGPT: Architecture, Function, and Generative AI

By IntuitionLabs • 6/10/2025 • 70 min read

chatgpt

large language models

generative ai

gpt architecture

transformer model

neural networks

unsupervised learning

artificial intelligence



ChatGPT in Life Sciences: A Comprehensive Guide for Industry Professionals

Understanding ChatGPT and How It Works

ChatGPT is a conversational AI system built on [large language model \(LLM\)](#) technology – essentially a **** Generative AI **** that produces human-like text responses from prompts. It was developed by OpenAI and is powered by the **GPT (Generative Pre-trained Transformer)** architecture [flexos.work](#). As a transformer-based neural network, it has been trained on massive amounts of text (internet articles, books, scientific papers, code, etc.), enabling it to understand context and generate coherent answers in natural language [flexos.work](#). The “Transformer” backbone – first introduced in the 2017 paper *Attention Is All You Need* – uses self-attention mechanisms to model the relationships between words in a sequence [ibm.com](#). This architecture allows GPT models to capture complex linguistic patterns and long-range dependencies more effectively than prior RNN-based approaches. In fact, autoregressive Transformer LLMs like OpenAI’s GPT-3/GPT-4 (which ChatGPT is based on) have **catalyzed the modern era of generative AI** [ibm.com](#), demonstrating unprecedented capabilities in language understanding and generation.

In practice, ChatGPT works by predicting the next word in a sentence based on the context of all prior words – a capability learned from its vast pre-training. It has on the order of hundreds of billions of parameters that were adjusted (via *unsupervised learning* on text corpora and subsequent fine-tuning with human feedback) to produce useful answers. The result is a model that can perform a wide range of language tasks: from answering questions and explaining concepts to drafting documents and writing code [sumble.com](#). It’s important to note that ChatGPT was further refined using **Reinforcement Learning from Human Feedback (RLHF)**, where human reviewers taught it to follow instructions and maintain conversational etiquette. This training process makes ChatGPT adept at interactive dialogue, allowing it to respond with contextual awareness and even admit uncertainty or rephrase when asked. However, because it generates responses probabilistically based on patterns in training data, it **does not “reason” or “know” facts in a human sense** – it can sometimes produce incorrect or nonsensical answers if prompted with ambiguous or complex queries. Understanding these fundamentals will help life science professionals use ChatGPT more effectively and cautiously.

In summary, **** ChatGPT is a large language model chatbot **** (based on GPT-3.5 and GPT-4) **that can generate human-like text** in response to user prompts [sumble.com](#). It falls under the

umbrella of *generative AI*, meaning it creates new content (sentences, paragraphs, answers) rather than simply retrieving exact answers from a database. Its power lies in the generalization learned from enormous datasets – it has “read” millions of biomedical papers, clinical trial reports, textbooks, and websites during training, giving it a broad (if sometimes superficial) knowledge across domains. With the right guidance, ChatGPT can **summarize complex scientific information, translate technical jargon, draft well-structured text, and even brainstorm hypotheses**. The following sections delve into how this technology can be applied in various subfields of the life sciences industry, and how [life science professionals](#) – from biotech R&D scientists to clinical trial managers and public health officials – can get started with ChatGPT in their workflows.

Applications and Use Cases Across Life Sciences

The life sciences industry spans diverse domains including biotechnology, pharmaceutical R&D, clinical research, genomics, public health, and more. **ChatGPT’s versatility means it can assist in almost all of these areas**, though the specific use cases and benefits vary by role. Below, we explore how generative AI is being leveraged (or piloted) in key subfields and job functions, with examples to illustrate real-world applications.

Drug Discovery and Preclinical Research (Biotechnology)

One of the most impactful areas for ChatGPT (and LLMs in general) is in **drug discovery and early-stage research**. R&D scientists in biotech and pharma are using ChatGPT as an intelligent research assistant to help sift through the ever-growing volume of scientific literature. For example, a researcher can ask ChatGPT to *summarize recent findings on a protein target or disease pathway*, saving hours that would otherwise be spent reading dozens of papers. Pfizer and AstraZeneca have employed generative AI to scan vast libraries of publications and data to identify new drug targets or molecular designs [intuitionlabs.ai](#). In surveys of biopharma companies, **drug discovery was cited as the #1 application of AI** to date [intuitionlabs.ai](#), reflecting the high value placed on tools that can accelerate early research.

ChatGPT can quickly **summarize the state of research on a given gene or protein** in minutes, highlight potential knowledge gaps, and suggest next experiment ideas – effectively augmenting a scientist’s ability to generate and prioritize hypotheses [intuitionlabs.ai](#). Importantly, it can draw connections between disparate pieces of information (e.g. linking a pathway involved in one disease to a novel use in another) by synthesizing knowledge from different sources.

Researchers at smaller biotechs like Recursion Pharmaceuticals have even **integrated GPT-4 models into their discovery platforms** [intuitionlabs.ai](#) to aid in hypothesis generation and data interpretation. Another emerging use is in *medicinal chemistry*: generative models can propose novel chemical structures for drug candidates. While still in early stages, **** LLMs **** coupled with cheminformatics can suggest new compounds to test by learning from databases of known

chemicals intuitionlabs.ai. This reduces some of the trial-and-error in molecule design by algorithmically brainstorming structures that fit certain target profiles.

In practical terms, a bench scientist might use ChatGPT to:

- **Literature review & ideation:** *"What are the known mechanisms of Drug X's toxicity?"* or *"Summarize the latest approaches to inhibit Enzyme Y."* The model's response can surface key points from myriad papers ptglab.com, giving the scientist a starting point for deeper investigation. It might even point out patterns or hypotheses (with appropriate prompts) that spark a new experiment idea.
- **Protocol refinement:** ChatGPT can suggest improvements to experimental protocols or troubleshooting tips. For instance, *"How can I improve the yield of my PCR reaction for gene Z?"* might yield advice (drawn from its training on molecular biology texts and forums) about optimizing primer design, Mg²⁺ concentrations, or annealing temperatures.
- **Data interpretation (with caution):** By providing summary statistics or observations to ChatGPT, researchers can get a narrative interpretation. E.g., *"I observed that compound A inhibits cell growth by 45% in high glucose conditions but not in low glucose – what could be going on?"* ChatGPT might propose possible biological explanations (perhaps invoking metabolic pathways or gene regulatory effects), which the scientist can then evaluate. While **the AI's suggestions must be validated**, they can inspire new angles to test.

It's worth noting that specialized biomedical LLMs (like **BioGPT** or fine-tuned versions of GPT on chemical or genomic data) are also emerging to complement ChatGPT. But even the general ChatGPT has ingested a huge amount of biomedical knowledge. Properly guided, it can be a powerful amplifier of a scientist's productivity – functioning like a knowledgeable but error-prone lab colleague who is very well-read. As McKinsey analysts noted in early 2024, generative AI has the potential to **transform nearly all aspects of pharmaceutical R&D**, potentially saving billions by accelerating compound identification, preclinical research, and lead optimization mckinsey.com. In summary, for biotech R&D teams, ChatGPT offers a way to **mine the scientific literature, brainstorm ideas, and reduce rote work**, allowing researchers to focus on designing and executing critical experiments.

Clinical Research and Trial Operations

Moving from the lab to the clinic, **ChatGPT is being applied in clinical research and development to streamline trial design, documentation, and analysis**. Clinical trials generate vast amounts of text – protocols, investigator brochures, patient eligibility criteria, regulatory reports, etc. LLMs excel at handling such text, making them well-suited as assistants in trial planning and execution.

A key use case is **clinical trial protocol design and optimization**. Companies like AstraZeneca have found that AI helps in crafting smarter trial protocols, for example by suggesting more precise inclusion/exclusion criteria or identifying potential study endpoints from prior literature intuitionlabs.ai. In one case, AstraZeneca reported using AI to help design eligibility criteria that

improve patient selection for trials intuitionlabs.ai, potentially leading to more efficient enrollment and clearer results. Researchers can prompt ChatGPT with a draft protocol (minus any confidential details) and ask for improvements or risk identification: *“Review this trial design for a Phase II study in oncology; suggest any potential logistical challenges or biases.”* The AI might point out, say, that a specific patient monitoring schedule is very frequent (which could hurt compliance), or that a certain subgroup might be underrepresented given the criteria – insights that the human team can then consider.

Another area is **patient recruitment and engagement**. Recruiting patients who meet complex eligibility criteria can be slow; generative AI tools are being used to match patient records or Real World Data against trial criteria in a more flexible, natural language way. For example, Sanofi partnered to develop an AI tool (nicknamed “Muse”) that uses OpenAI models to *find eligible patients faster* by scanning health records and disease registries intuitionlabs.ai. ChatGPT can also help draft patient-facing materials: e.g., simplifying a trial description into layperson language for consent forms or creating outreach emails that are compassionate and clear. While these drafts require review by ethics and compliance teams, they **speed up the process of generating trial communications**.

Adverse event reporting and pharmacovigilance in trials is another domain seeing AI assistance. During a study, investigators submit narrative reports of any adverse events. ChatGPT can standardize and summarize these narratives. Indegene, a pharma services firm, noted that ChatGPT can intake adverse event narratives and *draft pharmacovigilance case summaries more efficiently* than the manual process intuitionlabs.ai. By ensuring each report is structured and clear, it helps pharmacovigilance specialists review safety signals faster. Similarly, AI can assist medical monitors by quickly comparing an event to known side effect profiles: *“Does the combination of symptoms X, Y, Z in a trial patient resemble any known syndrome or drug reaction?”* – ChatGPT might instantly flag a known condition or drug-drug interaction from literature, which the medical monitor can verify.

Generative AI is also being leveraged to **analyze clinical trial data and results**. With the introduction of tools like *ChatGPT’s data analysis mode (formerly Code Interpreter)*, researchers can feed data (in a secure environment) and have the AI produce visualizations or statistical summaries intuitionlabs.ai. Moderna, for instance, used GPT’s data analysis capabilities in a tool called “DoseID” to analyze clinical dosing data intuitionlabs.ai. In a Nature article, scientists highlighted using AI to *write first drafts of trial results and even to analyze data in real-time* during studies intuitionlabs.ai. By simulating “digital twin” patients or outcomes, generative models might help predict trial outcomes under different scenarios, potentially reducing the need for some control patients intuitionlabs.ai (an idea still experimental but intriguing).

It’s important to emphasize that **human experts remain firmly in the loop**. ChatGPT can draft an interim clinical study report, but biostatisticians and medical writers then refine it. It can suggest that a trial arm be dropped for futility based on patterns, but leadership will make the decision after thorough analysis. The AI serves as a *co-pilot*: accelerating tedious tasks and illuminating patterns, but not flying solo on decisions that impact patient safety or regulatory

compliance intuitionlabs.ai. Early results from industry pilots are promising – companies report cutting down the time to generate trial documentation and even finding novel insights in data intuitionlabs.ai. By automating mundane documentation and providing analytical augmentation, ChatGPT is **making clinical operations more efficient**, helping new therapies reach patients faster.

Regulatory Affairs and Medical Writing

Professionals in regulatory affairs, medical writing, and related functions deal with **massive documentation workloads** – from regulatory submissions (INDs, NDAs, CTDs) and clinical study reports to literature reviews and standard operating procedures. Here, ChatGPT’s talent for generating and structuring text can offer tremendous productivity boosts, *if used carefully*.

Drafting and reviewing regulatory documents: Pharma and biotech companies have already begun using generative AI to produce first drafts of various sections of regulatory filings. For example, **Eli Lilly and Merck reportedly used ChatGPT to draft portions of clinical study reports and submission documents** intuitionlabs.ai, which were then edited by experts. Sanofi’s CEO has publicly stated they expect AI to write first drafts of FDA submission documents for upcoming drug filings, calling LLMs an “insane opportunity” to streamline R&D documentation intuitionlabs.ai. By letting the AI assemble a baseline text (e.g., summarizing efficacy results for an NDA module or compiling a drug’s risk/benefit discussion), medical writers can focus on polishing language and verifying accuracy rather than starting from scratch. This **cuts down tedious writing time and ensures consistency** in style and formatting across documents intuitionlabs.ai. One must be cautious, however – any content generated must be rigorously checked against source data to meet regulatory standards. Companies mitigate risk by using *enterprise-secure versions* of ChatGPT (either OpenAI’s ChatGPT Enterprise or self-hosted models) so that sensitive data isn’t exposed intuitionlabs.ai. For instance, Merck’s internal “GPTeal” platform provides a gated, secure environment where employees can use ChatGPT and other LLMs for confidential work without information leaving the firewall intuitionlabs.ai.

Medical writing and literature review: Medical writers are finding ChatGPT useful for creating **initial drafts of manuscripts, white papers, or slide decks**. Given a set of key results or an outline, ChatGPT can expand it into prose with appropriate scientific tone. It might draft an *abstract* summarizing a journal article or generate a coherent introduction for a literature review. Writers then edit for accuracy, add references, and adjust style. The model’s ability to maintain a formal tone and logical flow is a big advantage – it can take bullet points and turn them into well-structured paragraphs in seconds. As a trivial example, a prompt like: *“Draft an introduction for a review article on the current advances in CAR-T cell therapy for lymphoma”* can yield a surprisingly solid starting paragraph, covering the background of CAR-T, recent successes, and open challenges. The writer can then fact-check each statement and insert specific citations. **Speeding up the grunt work of writing** allows medical writers to spend more time on analysis and critical thinking. In fact, ChatGPT’s use in generating plain-language summaries is already

being explored – e.g., companies using it to create lay summaries of clinical results for trial participants or regulatory public disclosure.

Another common task is **literature search and summarization**. Regulatory teams compiling **Literature References sections** or doing pharmacovigilance literature monitoring have to summarize dozens of papers about a drug's safety/efficacy. ChatGPT can assist by summarizing each paper or even synthesizing trends across them: *"Summarize any reports of liver toxicity associated with Drug X in the literature."* It might output a synopsis: *"Drug X was associated with elevated liver enzymes in two case reports intuitionlabs.ai; one 2019 study (Smith et al.) reported reversible hepatitis in a patient. No large-scale studies have flagged significant hepatotoxicity."* – providing a quick overview that the specialist can verify and cite properly. **Pharmacovigilance writing** is similarly being turbocharged: transforming raw safety narratives into structured reports. As noted earlier, generative AI can draft *adverse event case narratives* by extracting key info (patient, event, outcome) and organizing it logically intuitionlabs.ai. This ensures consistency and frees up pharmacovigilance officers to focus on analysis rather than wording.

One must remain vigilant about **accuracy and compliance**. All content for regulatory purposes must be correct and not overstate or omit critical information. ChatGPT, if not explicitly guided, may "hallucinate" – e.g., fabricate a non-existent study or exaggerate a finding. Thus, best practice is to use ChatGPT *only with robust human review*. Many firms have internal SOPs: AI-generated text must be labeled as draft and cannot be used verbatim in filings without verification. Despite these precautions, the time savings are significant. Early adopter companies report that **writing time for certain regulatory documents dropped by well over 50%** when writers used AI to generate initial versions intuitionlabs.ai. Moreover, by training staff on AI use and establishing governance (as Johnson & Johnson did, training 56,000+ employees on AI and creating usage guidelines), organizations can **harness ChatGPT's productivity while mitigating risks** intuitionlabs.ai intuitionlabs.ai.

In summary, **regulatory affairs specialists, medical writers, and safety experts can leverage ChatGPT as a documentation assistant**, accelerating the preparation of high-quality drafts for the myriad reports and filings required in life sciences. The key is to pair the AI's speed with human expertise in fact-checking and regulatory knowledge. Done right, this can shorten submission timelines and reduce the drudgery of compliance paperwork – a significant win in an industry where every month saved means patients get access to therapies sooner.

Genomics and Personalized Medicine

The genomics revolution has led to an avalanche of data – sequencing results, variant databases, genome-wide association studies, etc. Making sense of this data and tailoring insights to individual patients (the essence of personalized medicine) is a daunting task. **ChatGPT and similar LLMs offer new ways to interpret and communicate genomic information.**

Analyzing genetic data and literature: While ChatGPT cannot analyze raw DNA sequences (that requires specialized bioinformatics tools), it can assist in *interpreting results* and connecting them to biological knowledge. For example, a genomic scientist might ask: *"What is the significance of the BRCA1 variant c.5266dupC (5382insC) in breast cancer?"* A well-prompted ChatGPT could respond with an explanation that this specific mutation is a known pathogenic variant in BRCA1 associated with a high risk of breast and ovarian cancer, perhaps mentioning it's one of the founder mutations in certain populations. It can summarize what is known about that variant from literature (again, we caution: the scientist must verify details from primary sources, as the model might not have the latest studies or could err in specifics). Researchers are also using LLMs to **generate hypotheses from genomic data** – for instance, suggesting which genes in a given list from an omics experiment might be *master regulators* or which pathways are enriched and worth investigating ptglab.com. By asking ChatGPT to contextualize a set of genes (e.g., "These 5 genes came up in our RNA-seq study of diabetes; what common pathways are they involved in?"), one might get a quick narrative: *"These genes are all related to inflammatory signaling – three of them (GeneA, GeneB, GeneC) are part of the TNF-alpha pathway, which is known to be involved in insulin resistance..."*, providing a starting point for deeper analysis.

Clinical genomics and variant reporting: In clinical settings, when a patient gets their genome or exome sequenced, physicians receive a report describing any clinically significant variants. ChatGPT could help **draft these genomic reports in plain language**. For example, given a variant and some annotation, it could produce a paragraph for a report: *"The patient carries a mutation in the DPYD gene (c.1905+1G>A), which is known to reduce the body's ability to break down the chemotherapy drug 5-FU. This means they may be at higher risk of toxicity from standard doses of 5-FU; a dose reduction or alternative therapy should be considered."* A genomic specialist would feed the relevant details, and the AI can structure a clear explanation, which the specialist then verifies and includes in the official report. This can save significant time in personalizing reports for each patient. Similarly, genetic counselors or medical geneticists might use ChatGPT to *explain genetic concepts to patients*. For instance, "Explain in simple terms what it means to have a 'pathogenic variant in the LDLR gene'" – the model can draft a patient-friendly explanation about familial hypercholesterolemia that the counselor can refine.

Drug discovery and pharmacogenomics: On the research side, genomics intersects with drug discovery in identifying new targets and understanding disease mechanisms. Generative AI can keep scientists up-to-date by summarizing *the latest genomics studies*. If you prompt, *"Summarize any recent discoveries in genomics related to Alzheimer's disease risk"*, ChatGPT might retrieve (from its training memory) that variants in genes like TREM2, APOE, and others have been implicated, describing their impact on amyloid or tau pathology (note: as of GPT-4's knowledge cutoff and slight updates, it may know up to 2023 findings, but anything very recent might be missing unless connected to a live data plugin). Nonetheless, it provides a broad review which a genomics scientist can then augment with a targeted literature search.

Furthermore, **personalized medicine** often involves combining genomic data with clinical data to tailor treatments. ChatGPT can support this by analyzing *complex patient profiles*. For example, an oncologist might use a secured version of ChatGPT (with proper de-identification) to ask: *"Given a 55-year-old male patient with metastatic lung cancer who has an EGFR L858R mutation and a new TP53 loss-of-function mutation, what therapies or trials should be considered?"* The AI could summarize: *"EGFR L858R is sensitizing to first-line EGFR inhibitors (like erlotinib); however, acquired TP53 mutations often indicate more aggressive disease. Standard care would be an EGFR inhibitor, but close monitoring for resistance is needed; consider combination trials that address TP53 loss (though no approved targeted therapy for TP53). You may also consider newer generation EGFR inhibitors if resistance emerges."* Such an answer (which must be vetted against up-to-date clinical guidelines) can help the physician ensure they haven't missed an angle, or prompt them to look up a trial that the AI alluded to. In practice, tools like **ChatGPT are being integrated with clinical decision support systems** – for example, some EHR software providers are experimenting with GPT-based assistants that doctors can query during patient care. In personalized medicine, where data is complex and patient-specific, an AI assistant can synthesize information on-the-fly in a useful way.

Of course, in **genomics and personalized medicine, accuracy is paramount**. Any gene or variant interpretation error can have serious consequences. Thus, ChatGPT's use here is primarily *assistant and drafter*; final interpretations rely on human experts and databases (like ClinVar, gnomAD, etc.). Still, by automating parts of the interpretation and explanation process, ChatGPT can make genomic medicine more scalable – a crucial benefit as we enter an era where sequencing might become routine for many patients.

Public Health and Epidemiology

Public health professionals and epidemiologists can also harness ChatGPT to improve population health outcomes and communications. This domain often involves **analyzing large datasets, surveillance reports, and policy documents**, as well as **disseminating information to the public or policymakers** – tasks well suited to a language model's capabilities.

Data analysis and outbreak response: Epidemiologists deal with surveillance data (e.g., disease incidence rates, mortality statistics, survey data) that needs to be quickly analyzed and reported. While statistical software is used for heavy analysis, ChatGPT can assist in *interpreting and communicating results*. For example, after crunching numbers on an increase in influenza cases, a public health analyst might ask ChatGPT: *"Interpret the following trend: influenza cases rose 20% in the past month in region X, mostly among age group Y. What factors could be contributing to this rise?"* The AI might respond with a logical explanation such as: increased travel or social gatherings, lower vaccination rates in that age group, or a new strain – essentially creating a draft situational report that the analyst can refine with actual data references. In a time-sensitive outbreak scenario (think of the early days of COVID-19 or a localized measles outbreak), **ChatGPT could help draft situation updates, press releases, or internal briefs** summarizing the latest numbers and recommended actions. By rapidly generating human-

readable summaries from raw data, it allows public health officials to respond more quickly and consistently.

There is also exploration of AI for **disease surveillance** on non-traditional data sources. For instance, monitoring social media or news for early signs of outbreaks – while this is more the realm of specialized AI, an LLM like ChatGPT could be used to *summarize unstructured reports*. A tool might feed in dozens of news snippets about a “mystery pneumonia” and ask ChatGPT to summarize the key details and assess if they sound like a new outbreak. The model could highlight common symptoms reported, locations, etc., giving epidemiologists a faster synthesis than reading all sources individually. Indeed, AI algorithms have been used for outbreak alerts (e.g., BlueDot for COVID-19), and we can envision LLMs providing readable analyses of those alerts.

Health education and communication: Public health relies heavily on effectively communicating guidelines and advice to the public. ChatGPT can be a **tremendous aid in drafting clear, accessible health communications**. Whether it’s creating pamphlet text, public service announcement scripts, or social media Q&A, the model can tailor language to different reading levels and cultural contexts. For example, *“Write a 5-point public advisory on how to prevent dengue fever, aimed at a general audience.”* ChatGPT might output bulleted advice about removing standing water, using mosquito repellent, etc., in simple language. Public health officials could then tweak the tone or add locale-specific information. The speed at which such messages can be generated (and re-generated in different languages) is valuable during emergencies when consistent messaging is needed quickly across multiple channels.

Moreover, ChatGPT can assist in **training and educational content** for health professionals. It could generate quiz questions for epidemiology students, or draft case study scenarios for public health workshops. For instance: *“Create a realistic scenario exercise for training health workers in contact tracing for a tuberculosis outbreak.”* The model might produce a narrative of a TB case, a list of contacts, some twist like one contact traveling, etc., which instructors can then use or refine for their training session. This saves educators time in content creation.

Policy development and grant writing: At a higher level, public health experts often need to write policy briefs or grant proposals. ChatGPT can help by summarizing evidence for policy options or drafting sections of a grant application. For example, *“Draft a background section for a grant proposal on improving vaccination coverage in rural communities, including problem statement and literature evidence.”* It will compile a generic but structured background that the writer can then fill with specific data and citations. This can be a huge time-saver when working on tight deadlines for funding opportunities, ensuring no section is blank due to writer’s block.

When using ChatGPT in public health, **ethical considerations** are front and center. Any advice given to the public must be accurate and culturally sensitive. Miscommunication can be dangerous. Therefore, while ChatGPT can draft messages or reports, **public health officials must review and approve all content**. Additionally, data privacy is crucial – any use of actual health data (even if de-identified) with ChatGPT needs to comply with regulations like HIPAA.

Generally, non-confidential aggregate data or published statistics are safe to use with such tools, but anything involving personal health information should be handled in secure, approved systems only.

In summary, for public health professionals, ChatGPT can act as a **speedy analyst and writer**: turning epidemiological data into insights, and turning insights into clear messages. It helps bridge the gap between complex data and actionable knowledge for communities. By leveraging it, public health agencies can potentially respond faster and communicate more effectively – whether it’s for everyday health promotion or crisis management during a global pandemic.

Knowledge Management, Collaboration, and Other Use Cases

Beyond the specific domains above, **ChatGPT is making inroads in the everyday knowledge work and collaboration tasks within life science organizations**. Large companies, especially, are deploying it as a general-purpose assistant across departments:

- Internal knowledge bases and support:** Companies like AstraZeneca and Novartis have built internal chatbot assistants (e.g., “AZ ChatGPT”, “NovaGPT”) that are **integrated with proprietary databases and documents** intuitionlabs.ai intuitionlabs.ai. Scientists and employees can query these systems in natural language to find information that previously might be buried in SharePoint sites or intranets. For instance, “Where can I find the protocol for the ABC123 mouse study we did last year?” could return a link or summary from the internal ELN/LIMS. Or a researcher could ask, “What were the key outcomes of Project Gemini’s Phase I trial?”, and the assistant (having been fed the internal report) can present the results. This turns siloed corporate knowledge into a conversational resource. Johnson & Johnson leadership even described a vision of a “bilingual” employee – fluent in both their domain *and* in using AI tools – meaning staff who routinely consult ChatGPT-based assistants to augment their work intuitionlabs.ai. By democratizing access to information, these tools can reduce time wasted searching through documents or relying on knowing the right person to ask.
- Collaborative writing and email drafting:** Many life science companies are using ChatGPT to help draft internal communications – from meeting summaries to policy updates. Novartis’s NovaGPT, for example, has been used by HR to draft policy documents and job descriptions, saving considerable time intuitionlabs.ai. Likewise, Merck’s GPTeal provided employees with a way to generate first drafts of emails and memos intuitionlabs.ai. In highly regulated environments, even internal emails about regulated products must be written carefully; ChatGPT can ensure a baseline of clarity and professionalism in such correspondence. Imagine a regulatory affairs manager using it to outline a project update email: “Compose a formal email to the clinical team summarizing the key feedback from the FDA on our recent submission, and next steps.” The draft would include the main points and a polite tone, which the manager can adjust. This reduces the friction of writing and lets employees focus on content rather than wording. **Meeting minutes and documentation** are another example – a ChatGPT integration in a conferencing tool might transcribe a project meeting and produce a summary of action items, which the team lead can verify and send out.

- **Sales and marketing content creation:** While the focus of this guide is on scientists and technical roles, it's worth noting that **commercial teams in pharma also utilize generative AI**. Pfizer's "Charlie" platform is a notable case – a custom ChatGPT solution to generate and fact-check marketing content intuitionlabs.ai. It helps create draft copy for brochures, slide decks for product presentations, and even **tailored messages for healthcare providers**. For example, a medical liaison could ask, *"Create a one-page summary of Drug X's mechanism of action and clinical benefits to share with cardiologists."* The AI will draft a concise, compliant summary (assuming it has been trained on approved content). Of course, all promotional content goes through medical legal review, but if the first draft is 80% there, that's significant efficiency. ChatGPT can also personalize communications: sales reps could input notes from a doctor meeting and get a suggested follow-up email that smoothly incorporates those details (e.g., *"Thank you Dr. Smith for discussing your experiences with diabetic patients on Drug Y..."*), again staying within approved claims intuitionlabs.ai. These use cases show that from R&D to sales, generative AI is permeating the industry.

Finally, **cross-cutting use cases** include coding and data analysis tasks. Many scientists do programming (in R, Python, etc., for bioinformatics or analysis). ChatGPT can assist by generating code snippets, debugging errors, or even helping design an analysis pipeline (via its Code Interpreter capabilities). For instance, a bioinformatician could prompt: *"Write a Python script to read a CSV of gene expression data and perform a t-test between two groups for each gene, correcting for multiple testing."* ChatGPT will produce a code outline (often correct or close to correct), which can be refined. This is like having a junior programming assistant and can be a boon for scientists who are not expert coders but need to manipulate data.

To summarize this section, ChatGPT's applications in life sciences are **extensive and growing**. Virtually any role that involves reading, writing, or analyzing natural language data can find some use for this AI assistant. From speeding up laboratory documentation to improving cross-team communication and aiding decision-making, generative AI is becoming a ubiquitous co-worker in the life sciences industry. McKinsey estimates that, taken together, such use cases could unlock **\$60–110 billion per year in value across the pharma/medtech value chain** by boosting productivity and accelerating innovation mckinsey.com. The next sections will discuss how to get started with ChatGPT in a responsible way, covering the crucial aspects of ethics, validation, integration with existing tools, and tips for crafting effective prompts to ensure scientific accuracy.

Getting Started: Ethical Considerations, Validation, and Data Security

Adopting ChatGPT in life sciences requires not just technical access to the tool, but also a thorough understanding of the **ethical and practical considerations** involved. Given the sensitive and high-stakes nature of biomedical work, careful steps must be taken to use generative AI responsibly. This section outlines key considerations – from data privacy to model accuracy – and offers guidance on validation and risk mitigation.

1. Data Privacy and Security: Perhaps the most immediate concern for companies is protecting confidential information. **Never input sensitive patient data, proprietary research results, or trade secrets into the public ChatGPT interface** – it's a cloud service and your inputs could be seen by the model developers or used for further training. Within3, a pharma tech company, bluntly stated that the publicly available ChatGPT "isn't the best fit for pharma" precisely because sending proprietary data into it is *unwise*, and there are risks of that data being exposed within3.com. The solution is to use *enterprise-grade offerings* or on-premise models. OpenAI now offers **ChatGPT Enterprise** which promises that it does not train on your data and provides encryption – many firms are opting for this. Others, like Moderna and Merck, have built internal platforms (e.g. mChat, GPTeal) using OpenAI's API but within their secure IT environment intuitionlabs.ai. Such instances ensure that any queries and responses remain within the company's firewall. It's critical to work with your IT and compliance teams to choose an approved method of ChatGPT access. In regulated sectors, also consider if any personal health information (PHI) might accidentally be included in prompts – if yes, you must follow HIPAA and other regulations to anonymize or avoid those prompts altogether. In summary: treat ChatGPT as you would any external software service – *don't share what you wouldn't email to a stranger*. And whenever possible, use a **secured, enterprise implementation** that aligns with your company's data protection policies within3.com.

2. Model Accuracy and Hallucinations: ChatGPT can **sometimes generate incorrect or fabricated information**, known as "hallucinations." This can range from minor factual errors to completely made-up references or data. In the life sciences context, such errors can be dangerous – imagine an AI that *invents* a fake clinical trial result or misquotes a dosage. As a systematic review in *NPJ Digital Medicine* noted, a distinctive concern with LLMs in healthcare is their tendency to produce content that is **convincingly written but inaccurate or even harmful** nature.com. Therefore, validation is crucial. **Never trust ChatGPT's output blindly**, especially for any decision-making. Always double-check critical facts against primary sources. A best practice is to **ask ChatGPT for sources or context**: for example, if it provides a claim ("Drug A improves 5-year survival by 20% in condition X"), prompt it with *"Can you provide the source or study for that statement?"* – sometimes it may cite a real paper. Other times, it might fumble and produce a fake citation, which itself is a red flag that the information needs verification. For important workflows, establish a *human review step*: e.g., if ChatGPT drafts a section of a regulatory document, a qualified professional must review every line. This is analogous to having junior staff write a draft – you always review their work; in this case the "junior staff" is an AI with no real-world judgment.

One effective mitigation is to **constrain the model's behavior through prompt instructions**. You can and should instruct ChatGPT *not* to make guesses. For instance, prefacing your prompt with *"You are a scientific assistant. If you don't know an answer or it's uncertain, do not fabricate information; instead say you are unsure."* can reduce the incidence of hallucinations. As an example, Certara's AI experts suggest phrases like *"Do not make things up if you don't know. Say 'I don't know' instead."* within your prompt certara.com. While ChatGPT won't literally say "I don't know" in normal chat (it tends to always attempt an answer), this instruction **does help it**

avoid overly confident falsehoods. Also, be specific in your queries to avoid ambiguity that might lead to nonsense. If you ask a broad question like “What are the side effects of drug X?” you might get a generic list. But if you specify “according to clinical trials” or “based on FDA-approved labeling” in the prompt, the AI is more likely to pull actual listed side effects rather than guessing.

3. Bias and Fairness: AI models can reflect or even amplify biases present in their training data. In healthcare and science, this could mean biased recommendations or perspectives (for example, overlooking data from under-represented populations). The *NPJ Digital Medicine* review highlighted concerns about **fairness and bias** [nature.com](https://www.nature.com/articles/s41746-023-01000-0). An example might be if an AI is used to draft patient education materials, it might unknowingly use language that is not culturally sensitive for certain groups, or presume a reading level that is too high for a target audience. Always review outputs with an eye for such biases. It may be wise to include in your prompt any context that helps mitigate bias: e.g., “*Draft this patient leaflet in a culturally inclusive manner and at a 6th-grade reading level.*” Transparency is also important – be open with colleagues or the public when AI is used in creating content. Internally, maintain logs of AI-generated contributions so that if a mistake is found later, you can trace back how it was introduced and correct it.

4. Compliance and Ethical Use: For regulated areas (like drug marketing or clinical decision support), additional rules apply. **AI should not be making decisions that require licensed medical judgment.** For instance, using ChatGPT to advise treatment for a patient directly is not appropriate – it’s not a medical device cleared for that use, and it has no accountability. Always keep a human in the loop for any clinical recommendations. If ChatGPT is used to draft promotional content, ensure it doesn’t introduce unapproved claims or language that compliance wouldn’t allow. Actually, many pharma companies have *banned certain uses of ChatGPT* or at least placed guardrails, especially after some early missteps. (E.g., some top companies briefly banned use altogether until internal policies and training were in place.) A good approach is to develop **Standard Operating Procedures (SOPs)** for generative AI use: define what types of tasks it can be used for (e.g., literature summaries, coding assistance, administrative writing) and what it cannot (e.g., new scientific analysis without validation, personal data processing, final content for external release without review). Also, consider the ethical dimension of automation – if ChatGPT is used to draft papers or reports, give proper attribution. There have been debates in scientific publishing about whether AI can be listed as a co-author (most journals say no, but you should disclose use in methods). Adhering to honesty and transparency about AI’s role is part of ethical use.

5. Model Limitations and Monitoring: Understand that **ChatGPT’s knowledge has cut-off dates** (for instance, GPT-4’s training data is mostly through 2021, with limited knowledge of 2022–2023 events unless updated). It won’t know the latest drug approvals or guidelines unless manually fed. So, don’t rely on it for up-to-the-minute facts. If you need current information (say the result of a conference last month), you’ll have to provide that context yourself or use tools like the browsing plugin. Additionally, monitor how the model performs on your specific tasks

over time. It can be useful to do a pilot phase – e.g., have multiple team members test ChatGPT on a known task (like summarizing the same article) and evaluate the outputs. This can reveal any common errors or gaps that you need to be mindful of. Continuously update your strategies and prompts based on these observations. In many organizations, Centers of Excellence or AI steering committees have been formed to share lessons learned from initial AI projects, so that best practices spread.

In essence, **getting started with ChatGPT in life sciences means starting carefully and thoughtfully**. Treat the AI as a powerful tool that, like any lab instrument, must be calibrated, supervised, and used by a trained operator. Start with low-risk tasks first (internal-facing uses, things that have strong oversight). Ensure everyone on the team understands the dos and don'ts – for example, conduct training sessions on how to craft safe prompts and how to detect when the AI's output is likely wrong. J&J's approach is instructive: they rolled out training to tens of thousands of employees and instituted governance programs to *"safely enable tools like ChatGPT"* intuitionlabs.ai. This kind of program is highly recommended. When done right, the **risks can be managed** (no AI hallucinations in final documents, no data leaks, no compliance issues), and the rewards – in efficiency and insight – can be substantial intuitionlabs.ai.

Integrating ChatGPT with Existing Tools and Workflows

To maximize ChatGPT's usefulness in the life sciences, organizations are looking to **integrate it directly into their existing software, lab systems, and workflows**. Instead of using ChatGPT as a standalone webpage, integration means it becomes embedded in the platforms scientists and professionals already use – from electronic lab notebooks to clinical data systems and CRMs. Such integrations can streamline processes by bringing AI assistance to the user's fingertips in context. Here we explore some integration opportunities and examples in labs, research, and business operations.

Electronic Lab Notebooks (ELNs) and Laboratory Information Management Systems

(LIMS): Modern R&D labs often use ELN/LIMS platforms (like Benchling, Labguru, LabVantage, etc.) to record experiments, manage samples, and track data. Imagine having ChatGPT *inside* the ELN – so while documenting an experiment, a scientist can get AI help with a click. This is already happening. Several ELN providers have announced AI features; for example, Labii and Genemod (providers of ELN/LIMS solutions) have integrated GPT-based assistants into their platforms genemod.net genemod.net. **LabiiGPT** is one such AI assistant that helps researchers with note-taking, protocol generation, and more labii.com. When a scientist is writing an experiment entry, they can simply describe the experiment in natural language, and the assistant will **generate a detailed protocol or notes automatically** labii.com. This can include materials, steps, and even safety precautions, all drafted from a simple description. Labii demonstrated generating step-by-step experimental protocols by just providing the name of the protocol – the AI fills in the standard steps labii.com. Similarly, **Genemod's AI-powered ELN** touts an "AI Chat" sidebar that researchers can ask questions or request summaries from, without leaving the ELN

interface genemod.net. For instance, if a scientist has a result chart in the ELN, they might ask, “Summarize the key findings from this result graph,” and the AI panel could output: “The new compound showed a 50% higher efficacy than control at day 7, with no observed toxicity...”. By integrating GPT, these lab systems **streamline data access and research support** – Genemod’s platform notes that users can obtain “immediate, accurate scientific insights” from its AI which is connected to an extensive database genemod.net. In essence, **common lab tasks (writing methods, analyzing results, documenting experiments) can be partially automated**, saving researchers time and ensuring more complete record-keeping. An example screenshot from Genemod’s ELN (shown below) illustrates how an AI chat panel might sit alongside experiment notes, ready to answer queries in real-time during research sessions.

Screenshot: An example of an AI assistant integrated into an Electronic Lab Notebook (Genemod’s ELN). The researcher can chat with the AI (right panel) while documenting experiments (left panel). This integration allows scientists to ask for protocol steps, data explanations, or literature summaries without switching context, thus streamlining lab workflows.

Integrating ChatGPT with LIMS can also enable **natural language queries of laboratory data**. Researchers or lab managers could ask, “How many samples do we have remaining from Project ABC and where are they stored?”, and the system could translate that into a database query and answer, “Project ABC has 12 samples left, stored in freezer 4 (rack B3).” This is far more user-friendly than manually running queries or combing through interfaces. Some vendors are enabling exactly that: Sapio Sciences, for example, has an AI assistant named **ELaiN** which lets scientists “chat” with their ELN/LIMS to set up experiments or retrieve data drugdiscoverynews.com. This kind of **conversational interface for lab management** can reduce training burdens (new staff can just ask in plain English instead of learning complex software menus) and improve efficiency.

Integration with Data Analysis and Bioinformatics Tools: Many R&D organizations use computational notebooks (like Jupyter), data visualization tools, or statistical software. Integrating ChatGPT via APIs into these environments means scientists can get coding help or data interpretation on the fly. Microsoft’s Azure OpenAI service, for example, allows companies to embed GPT models in custom applications. A bioinformatics team might integrate GPT into their data pipeline such that after running an analysis, the AI automatically generates a report of the results. Or within a Jupyter notebook, a researcher could highlight a block of results and trigger a “ExplainResults” AI function. This would be akin to having a virtual data analyst looking over your shoulder. Furthermore, combining GPT with domain-specific libraries – say having it call a protein database – could let users ask, “Find any known PDB structures for protein XYZ and summarize their active site characteristics.” The integration layer would fetch relevant data and let GPT compose the summary. These are custom solutions, but powerful: essentially **AI becomes part of the scientist’s toolset, embedded in their computational workflow**.

Customer Relationship Management (CRM) and Commercial Tools: On the commercial side (e.g., sales, medical liaison activities), integrating ChatGPT with CRM systems (like Veeva or

Salesforce used in pharma) can greatly enhance productivity. A sales rep using the CRM could have an AI assistant that generates a call summary after each doctor visit: *"Draft a call note based on my inputs: Doctor is mainly interested in new data about Drug X's side effects in elderly patients."* The AI could output a concise note to log in the CRM. Or when planning a meeting, the rep could query, *"What were the last 3 messages we sent this physician about Drug X?"*, and the assistant (integrated with CRM data) can summarize, ensuring continuity in communication. Pfizer's "Charlie" platform essentially integrates generative AI into their content supply chain for marketing intuitionlabs.ai, indicating such solutions can be custom-built. Even **chatbots for external stakeholders** (like healthcare professionals or patients) can use a ChatGPT backbone integrated with a company's specific knowledge base. For example, a pharma company might have a chatbot on its website for healthcare providers that, behind the scenes, uses GPT to answer questions about a drug – but it's constrained to only use approved prescribing information and published data. Integration means the chatbot can handle a wide range of phrasing in questions (thanks to GPT's language understanding) yet only provide vetted answers (by fetching from a controlled database). This offers a far better user experience than rigid FAQ systems, while maintaining accuracy.

Enterprise Collaboration Tools: Many life science companies rely on SharePoint, Microsoft Teams, Confluence, and other collaboration platforms. Microsoft's introduction of **Copilot for Office/Teams** is effectively bringing GPT into these apps. Thus, one integration scenario is using ChatGPT in project management or reporting. For instance, within a project management tool, a team could use an AI assistant to generate a project update summary from a collection of task notes, or to answer questions like "Which milestones are at risk this month?" by analyzing the project data. These enterprise integrations often use Azure OpenAI (which is the same GPT technology but offered for custom app integration). In fact, OpenAI has a specific focus on life sciences as indicated by their hiring (e.g., a GTM lead for life sciences platforms openai.com), meaning we can expect even more out-of-the-box integrations and templates tailored to our industry's needs.

Connecting to internal databases and knowledge is a recurring theme. The best ROI from ChatGPT comes when it's not just drawing on generic public knowledge (which it has up to 2021 or so), but on *your organization's specific knowledge*. Integration can involve using retrieval augmentation: e.g., linking GPT with a vector database of your documents so it can pull up relevant snippets to form its answers (this prevents hallucination and ensures it uses real data). AstraZeneca's AZ ChatGPT did this by leveraging internal data – it could answer complex R&D questions by having access to proprietary knowledge bases intuitionlabs.ai. So when integrating, consider pipeline where the query first fetches relevant internal info, and then GPT composes an answer. Many vendors and open source tools (LangChain, etc.) support building this architecture.

From the above, it's clear that **integration amplifies ChatGPT's utility**. Instead of being a separate tool that someone has to consciously go use, it becomes a seamless part of everyday software – a bit like having an AI "mode" or assistant in every application. The benefits include:

users stay in their flow (no context switching), the AI can use contextual data from the app (e.g., the document you're editing, the data you're viewing), and IT can enforce usage policies more easily at the integration level (for example, stripping out any sensitive fields before sending a query to the AI API). However, integration projects do require technical investment – APIs, possibly middleware to mediate between the AI and existing systems, and testing to ensure reliability. A phased approach is wise: perhaps start by integrating ChatGPT into a non-critical system or as a pilot with one team's tools, gather feedback, and then scale up.

Leading organizations that have done these integrations report improved efficiency. Moderna's mChat (their internal ChatGPT) achieved **over 80% employee adoption rapidly**, with employees building 750+ custom mini-AI assistants for tasks like trial dose selection (DoseID) intuitionlabs.ai. This illustrates how, once the AI is readily available in their workflow, users themselves will innovate and find new uses for it. It becomes part of the digital fabric of the company. As one Genentech scientist commented (anecdotally), "Having GPT in my ELN is like having a genius lab partner who's always available." That captures the promise of these integrations – the AI becomes a ubiquitous collaborator across software tools, enhancing every step of work from planning to execution to reporting.

Best Practices for Prompt Engineering in Scientific Work

Using ChatGPT effectively is an art and science of its own – often referred to as **prompt engineering**. Crafting your prompts well can be the difference between a useless answer and a hugely valuable one. Especially in scientific and technical domains, where precision matters, following best practices in prompt design will help ensure accuracy and usefulness. Here are some **prompt engineering best practices tailored for life science professionals**:

- Be Specific and Provide Context:** A vague question will yield a generic answer. Always try to **frame your prompt with enough detail** so the model understands exactly what you need. Instead of asking *"Explain CRISPR"*, you'd get a far more targeted answer with *"Explain how CRISPR-Cas9 gene editing works to a molecular biology graduate student, focusing on the role of guide RNA."* The latter sets context (molecular biology grad student), specifies the focus (guide RNA's role), and the format (explanation). Similarly, if you want a summary of a paper, include the text or key points of that paper if possible, and say *"Summarize the above findings in 3-4 bullet points."* The more you guide the AI, the better it performs.
- State the Role or Style if Needed:** You can ask ChatGPT to adopt a persona or style that fits your task. For scientific accuracy, you might start with *"You are an expert pharmacologist..."* or *"Act as a clinical data analyst..."*. This often leads to answers with the appropriate tone and depth. For example, prompting *"You are a medical writer tasked with summarizing a clinical trial result for an FDA submission. Explain the efficacy outcome clearly and formally."* will produce a more regulatory-style summary than a casual tone answer. Defining the role can also help with **terminology** – an "expert pharmacologist" will likely use the correct jargon, whereas a "patient educator" will intentionally simplify language. Choose the role that matches your audience.

- Include Instructions to Reduce Hallucinations:** As mentioned in the ethical section, you can embed instructions like *"If you are unsure of an exact fact, do not fabricate it."* or *"Cite any numbers you mention with their source from the input."* While ChatGPT doesn't have a true database of sources to cite (unless you provided one in the prompt), telling it not to make things up can somewhat reduce false outputs certara.com. Another trick: ask for answers in a specific format that discourages guessing. For instance, *"List three possible explanations for X, and for each, state if it's supported by data (yes/no). If no data, say it's hypothetical."* Forcing that structure may make the model more cautious before claiming something is supported by data.
- Use Step-by-Step or Chain-of-Thought for Complex Tasks:** If the question is complex, you can explicitly instruct the model to reason it out step by step. For example: *"First list the known pathways involved in inflammation in diabetes. Then identify which of those pathways Drug A affects, and finally summarize how that could explain the trial results."* By breaking the task into steps, you guide the model's thinking process. This is akin to chain-of-thought prompting, which often yields better reasoning. You can also use the strategy of asking the model if it needs additional info before answering: *"Do you have all the necessary data to answer or would you like me to provide X?"* – sometimes it will ask for clarification, which you can then give, leading to a more accurate response.
- Iterate and Refine: Prompting is an iterative process.** You might not get the ideal answer on the first try. Don't be afraid to refine your prompt and ask again. For instance, if an answer comes back too shallow, you can prompt further: *"Thanks, can you elaborate on the second point, providing more details from the literature?"* Or if it was too verbose, *"Please repeat that more concisely."* Each iteration can bring the answer closer to what you need. Certara's experts note you need not craft the perfect prompt at first shot – start small and then build it up certara.com. If you'll reuse a prompt often (like a template for summarizing papers), invest time in refining it through trial and error. If it's a one-off question, a quick two-pass approach may suffice (initial answer, then a follow-up for clarification or expansion).
- Constrain Length or Format as Needed:** When you expect a long report or a list, tell the model explicitly what format and length you want. For example: *"Provide a one-paragraph summary (~100 words) followed by 3 bullet points of key data."* This avoids situations where the model might write pages if not told otherwise. It also ensures the output fits into your document or slide. When generating tabular data or code, you can say *"Output as a markdown table"* or *"provide code only, no explanation"*, respectively, to get the format you want. Being precise might also prevent the model from meandering off-topic – it knows exactly what to deliver.
- Provide Examples (One-shot/Few-shot Prompting):** One of the most powerful techniques is to **show the model an example of the output you expect**. For instance, if you want it to rewrite text in a certain way, give it a sample input and the desired rewritten output as a guide (that's one-shot). Or provide a couple of QA pairs so it learns the style. For example: *"Q: What is the capital of France? A: Paris. Q: What is the largest organ in the human body? A: The skin. Q: \${your question}\$".* In a scientific context, if you want an analysis in a specific structure, you can write a short dummy example: *"E.g., Input data: | [some dummy data]. Analysis: | [demonstrate the style of reasoning]"*. The model will infer that pattern. Certara's blog demonstrated how giving an example significantly improved the accuracy of the model's response to a tricky prompt certara.com certara.com. Use this to your advantage whenever you have a clear idea of the format/logic you want.

- Double-Check Numerical or Factual Answers:** If ChatGPT provides numerical results (say, from calculations or data interpretation), it's wise to verify them. You can actually ask ChatGPT to double-check its math explicitly by prompting something like *"Double-check the above calculation step by step."* It sometimes finds its own arithmetic mistakes. For critical calculations, though, rely on a proper statistical tool or calculator. Think of ChatGPT's math ability as that of an average person – it might slip up on more than trivial arithmetic or when many steps are involved (although GPT-4 is much better at math than earlier versions). For factual claims, a neat trick is to ask the model to provide references (if you have the browsing plugin or if you feed it a reference text yourself). If it confidently states a fact without reference, consider doing a quick manual literature search or asking it *"What is the source for that information?"*. If it can't provide one, you should be skeptical of the fact.
- Be Mindful of Prompt Length and Information Overload:** There's a limit to how much you can and should stuff into a prompt. If you provide huge texts, the model might focus on the wrong parts or get confused. When giving context like an article to summarize, ensure it's relevant and maybe break it into chunks if very long. Also, remember the model has a token limit (for GPT-4, it can handle several thousand words of combined input+output). If you approach that limit, the model might truncate or forget earlier parts of the input. In such cases, summarizing or splitting the task into parts is better.
- Avoid Ambiguity and Open-Endedness (Unless Intentionally Brainstorming):** If you ask something like *"Tell me about cancer research"*, you'll get a rambling general answer. Narrow it down: *"Compare the mechanisms of action of CAR-T cell therapy vs immune checkpoint inhibitors in cancer treatment."* This yields a structured comparison. Only leave the prompt open-ended if your goal is indeed to brainstorm or explore broadly. When brainstorming, you can explicitly say *"List as many ideas as possible about X"* or even encourage creativity: *"Give 5 creative hypotheses why our assay might be showing inconsistent results."* By specifying number or framing as a brainstorming, you signal the model to produce divergent answers rather than a single focused one.
- Use Temperature and Other Settings if via API:** If you happen to be using ChatGPT via the OpenAI API or a tool that allows parameter tweaking, remember that *temperature* controls randomness (0 for very deterministic, 1 for more creative). For scientific accuracy, a lower temperature (0–0.3) is often better because it will stick to more likely responses and not get too wild. However, for creative brainstorming, a higher temperature might produce more novel ideas. Also, the *max_tokens* setting can limit length – ensure it's set high enough for your needs or you might see the answer cut off.
- Keep a Prompt Log and Learn:** Finally, treat prompt engineering as a learning process. Keep note of what kinds of prompts worked well for your purposes. Often, you'll develop reusable templates. For instance, a med writer might have a go-to prompt structure for summarizing a study: *"You are a medical writer. Summarize the study | [TITLE] by | [AUTHORS] in 3 paragraphs: Background, Methods/Results, Conclusion, highlighting efficacy and safety outcomes. Use a formal tone."* This could be saved and each time just fill in the specifics. Over time, you'll refine these. Share effective prompts with colleagues (some companies even build internal "prompt libraries" for common tasks).

By following these best practices, you can significantly **improve the quality and reliability of ChatGPT's outputs**. Prompt engineering is a key skill for leveraging AI in any professional domain – akin to knowing how to query databases or search engines, but in a more conversational and creative way. With some experience, you'll find that you can get ChatGPT to

do exactly what you need most of the time, whether that's writing a crisp summary, analyzing an experimental result, or providing insightful explanations. The next section will illustrate many of these principles in action, by providing real-world example prompts and explaining why they are constructed the way they are.

Real-World Prompt Examples for Life Science Professionals

To solidify the concepts, here are **20 example prompts** spanning various roles and tasks in the life sciences industry. Each prompt is accompanied by a brief explanation of its context and why it's useful. You can try these (with appropriate modifications to fit your specific case) to jumpstart your use of ChatGPT in your work:

1. **Prompt:** *"Summarize the key findings of the latest research on CRISPR-based therapies for sickle cell disease in 3-4 sentences."*
Explanation: This prompt asks for a concise summary of a specific topic. A **research scientist** or R&D manager could use it to get a quick overview of recent developments (for example, the success of CRISPR-edited cell therapies in clinical trials) without reading dozens of papers. Specifying *3-4 sentences* forces ChatGPT to be brief and focus on the most important points.
2. **Prompt:** *"You are a medicinal chemist. Propose three novel chemical scaffolds that could potentially inhibit the ABC kinase (involved in cancer growth), and explain your reasoning for each."*
Explanation: This is a brainstorming prompt for **drug discovery**. By assigning the role "you are a medicinal chemist," the user nudges the model to use appropriate technical language. Asking for three novel scaffolds with reasoning encourages creative hypothesis generation (divergent thinking) while also requiring an explanation (so the suggestions aren't random). A chemist could use the output as inspiration for real compound designs to consider (with the caveat that these are hypothetical ideas).
3. **Prompt:** *"Explain the mechanism of action of pembrolizumab in treating cancer to a non-specialist audience. Use an analogy if possible."*
Explanation: This prompt is useful for a **medical affairs or medical science liaison** professional who needs to communicate complex science to healthcare providers or patients in simpler terms. By requesting an analogy (e.g., "it takes the brakes off the immune system's T cells, a bit like releasing a parking brake so the immune system can move forward to attack cancer cells"), the response can be more relatable. It's instructive to specify "non-specialist audience" to control the complexity of language.

4. **Prompt:** *"Generate an outline for an FDA briefing document section: 'Clinical Efficacy Results' for our Phase III trial of Drug X in rheumatoid arthritis. Include subheaders for study design, endpoints, outcomes, and subgroup analyses."*

Explanation: A **regulatory affairs specialist or medical writer** could use this to overcome the blank-page syndrome when writing a big document. The prompt asks for an outline with specific subheaders, which helps ensure the structure covers all necessary elements. ChatGPT will produce a structured outline (e.g., *Study Design, Primary Efficacy Endpoint Results, Secondary Endpoint Results, Subgroup Outcomes, etc.*) that the writer can then fill in with actual data from the study.

5. **Prompt:** *"Given the following data on adverse events from our trial (provide data in a short table or list), draft a paragraph describing the safety profile of Drug Y versus placebo."*

Explanation: This combines providing data with asking for narrative. A **clinical researcher or pharmacovigilance specialist** can use this after assembling the key safety data. ChatGPT will take the numbers and turn them into sentences like "Drug Y was generally well-tolerated with a similar overall adverse event rate to placebo. The most common side effects were headaches (Drug Y 10% vs placebo 8%) and nausea (7% vs 5%). Importantly, serious adverse events were rare and occurred in 2% of Drug Y patients compared to 1.5% on placebo intuitionlabs.ai." This saves time and ensures the description flows well, which the specialist can then tweak.

6. **Prompt:** *"List five potential reasons why our ELISA assay might be showing inconsistent results, and suggest a way to address each reason."*

Explanation: This is an example of using ChatGPT for **troubleshooting and problem-solving** in a lab setting. A **lab scientist or research associate** facing an experimental issue (variable ELISA results, in this case) can ask the AI to brainstorm causes and solutions. The output might include reasons like reagent degradation, plate reader calibration issues, operator technique variability, etc., each with a suggested fix (e.g., "ensure all reagents are at room temp and mixed thoroughly" or "run a calibration curve on the plate reader"). It's like having an experienced colleague to bounce ideas off of.

7. **Prompt:** *"Translate the following technical paragraph into layperson language (8th-grade reading level): [insert a dense paragraph about a clinical study outcome]."*

Explanation: Public health officials or medical writers often need to **communicate science to the general public**. This prompt explicitly asks for a translation to 8th-grade level, which will strip out jargon and simplify sentences. It's extremely useful for writing patient education materials or press releases about scientific findings. ChatGPT excels at adjusting tone and complexity when instructed.

8. **Prompt:** *"Create a step-by-step checklist for conducting a GLP-compliant stability study for a pharmaceutical product."*

Explanation: For **quality assurance or lab management**, having SOPs and checklists is essential. This prompt asks for a procedural checklist, which ChatGPT can generate by drawing on guidelines for Good Laboratory Practice (GLP). The result would be something like: "1. Prepare stability indicating assay, 2. Calibrate all instruments, 3. Label samples with batch and timepoint, ..." etc. It provides a starting checklist that can be reviewed and modified to the specific SOPs of the organization.

9. **Prompt:** *"As a clinical data manager, you need to explain a protocol deviation in simple terms during a meeting. How would you describe a situation where a patient missed a visit window, and what steps were taken?"*

Explanation: This is tailored to a **clinical data manager or clinical operations** context. It asks the model to formulate an explanation of a technical issue ("protocol deviation – patient visit out of window") in simple terms suitable for a broad team meeting. The answer will likely yield a concise explanation like: "We had a protocol deviation: one patient's follow-up visit happened outside the allowed timeframe. In other words, their scheduled visit was a week late. To address this, we documented the reason (the patient was traveling) and allowed them to come at the earliest possible date. The site also re-trained staff on visit scheduling to prevent this in future." This helps the user articulate the issue clearly to stakeholders.

10. **Prompt:** *"Summarize this research article on genome-wide association studies (GWAS) into 5 bullet points of key takeaways."*

Explanation: A **genomics researcher or any scientist** could use this to quickly distill a paper. By feeding the article text (or just the abstract) into the prompt, and asking for 5 bullet point takeaways, the model will extract major points: e.g., population studied, number of loci found, biggest implicated gene and its effect, etc. This is very handy for literature reviews – one can process many papers this way to build a summary table of findings. Do remember to cross-check the bullets with the source for accuracy, but it dramatically speeds up the first pass of comprehension.

11. **Prompt:** *"Draft an email to the project team updating them on our IND submission status. Mention that Module 4 is complete and we are awaiting feedback from toxicology reviewers. Keep the tone professional and positive."*

Explanation: Professionals in **regulatory or project management** often need to send status updates. This prompt essentially outsources the first draft of a project update email. ChatGPT will produce a nicely worded email, e.g., "Dear Team, I wanted to provide an update on our IND submission for Project Alpha. We have successfully completed Module 4 (Nonclinical/Toxicology) and submitted it to the agency last week. We are now awaiting feedback from the toxicology reviewers, which we expect in the next few weeks. This is a major milestone, and I want to thank everyone for their contributions... etc." The user saves time and only needs to tweak details or style as needed.

12. **Prompt:** *"Outline a protocol for an animal study to test the efficacy of a new antibiotic. Include objectives, animal model, dosing strategy, and outcome measures."*

Explanation: This is useful for a **preclinical scientist or research manager** who is designing a study. ChatGPT will generate a structured protocol outline with sections like Objectives (e.g., "to evaluate efficacy of Compound Z in a mouse sepsis model"), Animal Model (which species/strain, infection method), Dosing (doses of antibiotic, frequency, controls), and Outcomes (survival, bacterial load, etc.). While it won't be tailored to your exact compound without more details, it gives a solid framework that can be filled in. It ensures you don't forget key elements when drafting a new protocol.

13. **Prompt:** *"Regulatory Q&A: Provide a clear answer to the question – 'Does the product contain any materials of animal origin?' – assuming our vaccine uses a recombinant protein expressed in yeast. Answer as if responding in a regulatory document."*

Explanation: In regulatory submissions, there are often question-answer sections (like in Module 3 for quality, or authorities asking clarifications). This prompt sets up such a scenario. ChatGPT will respond with a formally worded answer: "The product does not contain materials of animal origin. The antigen is a recombinant protein expressed in *Saccharomyces cerevisiae* (yeast), and no animal-derived raw materials are used in the fermentation, purification, or formulation processes. All excipients are of synthetic or plant origin." A **regulatory affairs professional** can use that as a draft and ensure it fits the actual manufacturing info. It's particularly useful because the AI's formal tone matches what authorities expect.

14. **Prompt:** *"Brainstorm potential market applications for our AI-driven bioinformatics platform beyond oncology research. List at least four distinct use cases (e.g., in agriculture, environmental science, etc.)."*

Explanation: This is a more business-oriented prompt, perhaps for an **innovation team or product manager** at a biotech informatics company. It asks ChatGPT to think of tangential or expanded use cases. The model might come up with: 1) crop disease resistance gene analysis in agriculture, 2) analyzing microbial genomes for environmental bioremediation, 3) personalized nutrition genomics, 4) epidemiological genomics for tracking pathogen evolution, etc. Such brainstorming can expand the team's perspective and perhaps reveal an application they hadn't considered. The key here is the phrase "beyond oncology research," directing the AI away from its assumption if the platform is currently used in oncology.

15. **Prompt:** *"Create a brief SOP (Standard Operating Procedure) for cleaning and calibrating a pH meter in the laboratory. Use numbered steps."*

Explanation: Lab managers or **quality control personnel** can use ChatGPT to draft SOPs or work instructions. The prompt requests numbered steps, which yields a clear stepwise procedure: "1. Rinse the electrode with distilled water; 2. Prepare calibration buffers (pH 4, 7, 10)... etc." The result can serve as a first draft SOP that then gets reviewed for compliance with internal standards. It ensures completeness (cleaning, calibration, storage of electrode) because the model has likely seen many SOP texts.

16. **Prompt:** *"We have a 70% dropout rate in our clinical trial's control arm. Ask ChatGPT: What are the possible reasons for such a high dropout, and how can we mitigate this in future trials?"*

Explanation: This is an analytical prompt for **clinical operations or trial design teams**. A high dropout rate is a big problem. ChatGPT will likely list reasons: e.g., control arm might be on placebo and not perceiving benefit so they quit, or side effects if the control is standard of care that's poorly tolerated, or logistic issues like too many visits, etc. It will also suggest mitigations: improving patient engagement, better communication, perhaps using some incentive or making trial participation easier (telemedicine visits, etc.), or trial design changes like cross-over arms. This kind of analysis can help the team ensure they've thought of all angles when addressing the issue.

17. **Prompt:** *"Explain the difference between positive predictive value and sensitivity in diagnostics, and give an example of why each matters in a COVID-19 test."*

Explanation: This prompt is education-focused. A **public health professional or clinical educator** might use it to clarify concepts to colleagues or students. ChatGPT will articulate that sensitivity is the true positive rate (detecting disease when present) whereas positive predictive value (PPV) is the probability that a positive test is a true positive, and then tie it to COVID-19: e.g., sensitivity matters to catch as many cases as possible, PPV matters especially when prevalence is low (to avoid too many false alarms). The user gets a nice explanation with an applied example that can be used in training materials or discussions.

18. **Prompt:** *"Draft a short introduction for a peer-reviewed journal article where we used ChatGPT in the methods. Acknowledge the use of an AI assistant and the need for human validation of its outputs."*

Explanation: There is growing interest in **how to document AI usage in research**. If a team used ChatGPT to help write or analyze, journals might require an acknowledgement or description. This prompt has ChatGPT write that statement itself. It will likely produce something like: "In this study, we employed OpenAI's ChatGPT, a generative language model, as an aid in drafting portions of the manuscript (specifically, the initial outline and summary of related work). All content generated by the AI was subsequently reviewed and validated by the authors for accuracy and completeness before inclusion in the paper. We acknowledge the use of this tool in accordance with journal guidelines on AI assistance." This saves the researcher time formulating the wording and ensures transparency.

19. **Prompt:** *"Our biotech startup is preparing a pitch. Summarize our technology (CAR-T therapy with a novel co-stimulatory domain) in one compelling sentence for a lay investor audience."*

Explanation: This prompt is about **communication and marketing**. Crafting a punchy one-liner for a pitch is hard – ChatGPT can try dozens quickly. It will compress the concept into something like: "We reprogram a patient's own immune cells to become 'super immune cells' that hunt down and destroy cancer, using a new built-in signal that makes them stronger and longer-lasting than current therapies." – which a layperson investor can grasp. The user can then refine the favorite attempt. This demonstrates how AI can assist in distilling complex biotech concepts into elevator pitches.

20. **Prompt:** *"Generate 5 quiz questions (multiple-choice) to test understanding of good clinical practice (GCP) guidelines, with answers."*

Explanation: This is for **training and development** purposes. Perhaps a clinical trial manager wants to quiz the team after a GCP training. ChatGPT will produce Q&A like: "Q1. What is the primary purpose of GCP? A) Protect trial integrity and participant rights (correct), B) Speed up drug approvals, ..." etc., complete with answers. It dramatically reduces the time needed to create training materials. The user should still double-check that the answers are correct and the questions align with their training focus, but it's a quick first draft of a quiz.

Each of the above examples demonstrates how to structure prompts to get useful output from ChatGPT. They cover a range of scenarios – summarization, explanation, creative brainstorming, drafting communications, troubleshooting, and even generating educational content. By studying these examples, you can adapt the patterns to your own needs. Notice how they often specify the audience or depth, request a certain format (bullets, list, email, etc.), and sometimes

feed context or data to the model. These techniques help steer ChatGPT to provide high-quality, relevant answers, making it a practical partner in daily life science workflows.

Conclusion

ChatGPT and similar generative AI models represent a significant new toolset for life sciences professionals. As we have explored, **ChatGPT can accelerate literature reviews, aid in experimental design, streamline clinical trial documentation, assist with regulatory writing, enhance genomic interpretations, and improve public health communications – among many other applications.** In an industry where knowledge is vast and the stakes are high, having an AI assistant that can digest information and generate coherent, context-aware content is akin to adding a highly skilled team member who never sleeps. Early case studies from leading organizations (Moderna, Pfizer, J&J, AstraZeneca, and others) show that when deployed thoughtfully, ChatGPT can save substantial time and even unlock new insights intuitionlabs.ai intuitionlabs.ai.

However, it's equally clear that **success with ChatGPT in life sciences depends on responsible usage.** Ethical considerations – ensuring patient privacy, avoiding AI hallucinations, and maintaining compliance – are not just checkboxes but fundamental requirements. The guide above emphasized steps for validation and risk mitigation at every turn: always keep a human in the loop, double-check facts, and start with secure implementations within3.com nature.com. As regulatory frameworks catch up with AI and as internal company policies mature, these precautions will become standard operating procedure. Professionals who stay informed about best practices and contribute to shaping AI governance in their organizations will help ensure that AI deployment is both impactful and safe.

From a technical standpoint, **integrating ChatGPT into existing tools and workflows** will likely drive the next wave of productivity gains. As seen, embedding AI in ELNs, LIMS, CRMs, and knowledge portals brings the technology to users' fingertips genemod.net intuitionlabs.ai. Life science companies should evaluate their software ecosystems and identify high-friction points where AI assistance could make a difference – whether it's writing up an experiment, querying a database, or preparing a slide deck. Even without deep technical integration, teams can start by creating template prompts for common tasks (some of which we exemplified) and sharing them internally. This can quickly raise the baseline productivity and ensure consistency across work products.

Best practices in prompt engineering will continue to evolve, but the core idea is that we can talk to our computers in plain language now – and get meaningful work done through that dialogue. This is a profound change from traditional software constrained by menus and forms. For scientists and experts, it means learning a new skill: how to ask effectively. The better we articulate our needs to ChatGPT, the better it can assist us certara.com certara.com. Fortunately,

domain experts already excel at structured thinking, so with some experimentation, they often become adept prompters in short order.

In conclusion, getting started with ChatGPT in the life sciences is both an exciting opportunity and a journey of adaptation. By deeply understanding what ChatGPT is (and isn't), recognizing where it can add value in our diverse industry, and by implementing it with care and oversight, professionals can harness generative AI to **enhance innovation and efficiency in biotechnology, pharmaceuticals, clinical research, genomics, and public health**. The promise – as evidenced by early adopters – is a future where tedious groundwork is reduced, human creativity and decision-making are amplified, and perhaps even scientific breakthroughs are accelerated by our new AI collaborators mckinsey.com. Embracing this tool responsibly will be a key competitive advantage for life science organizations in the years ahead, and a valuable skill set for the scientists and experts who drive these fields forward.

Sources:

- IBM – *What is a transformer model?* (Think Blog) ibm.com ibm.com
- Sumble Tech Insights – ChatGPT Overview sumble.com
- FlexOS Future of Work – *Everything You Need to Know About ChatGPT* flexos.work flexos.work
- McKinsey Report – *Generative AI in Pharma: Hype to Reality* mckinsey.com
- Intuition Labs – *ChatGPT Adoption in Life Sciences* (Industry case studies) intuitionlabs.ai intuitionlabs.ai intuitionlabs.ai intuitionlabs.ai intuitionlabs.ai
- Within3 – *Generative AI and ChatGPT in Pharma* (blog) within3.com
- *npj Digital Medicine* (Nature) – Systematic Review on Ethics of ChatGPT in Healthcare nature.com
- Certara – *Best Practices for AI Prompt Engineering in Life Sciences* certara.com certara.com
- Labii & Genemod – AI integration in ELN/LIMS labii.com labii.com genemod.net genemod.net
- Intuition Labs – Company AI adoption details (Moderna mChat, Merck GPTeal, etc.) intuitionlabs.ai intuitionlabs.ai
- Proteintech Blog – *ChatGPT and BioGPT in Life Science Research* ptglab.com ptglab.com
- Additional references embedded in text for specific facts and examples intuitionlabs.ai intuitionlabs.ai, etc.

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 an AI software development company specializing in helping life-science companies implement and leverage artificial intelligence solutions. Founded in 2023 by [Adrien Laurent](#) and based in San Jose, California.

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.