

OpenAI Codex for Biotech and Bioinformatics Research

By Adrien Laurent, CEO at IntuitionLabs • 4/11/2026 • 30 min read

openai codex

bioinformatics


biotech research

ai coding agents

computational biology

data analysis

python programming



OpenAI Codex for Biotech and Bioinformatics Research

intuitionlabs.ai

Executive Summary

OpenAI's Codex is an advanced AI-driven coding agent that dramatically accelerates software development tasks, including those crucial to biotechnology research. It integrates with [ChatGPT](#) and programming environments to write, test, and debug code based on [natural-language prompts](#). This report provides a comprehensive analysis of how Codex can be leveraged for scientific research in biotech. We review the historical context of AI in biology, the technical capabilities of Codex, and practical applications in genomics, proteomics, and [drug discovery](#). We include data from academic studies and industry reports showing how Codex and similar models can automate common bioinformatics tasks (e.g. sequence analysis, pipeline development) and yield functional code even for non-programmers (^[1] [pmc.ncbi.nlm.nih.gov](#)) (^[2] [pmc.ncbi.nlm.nih.gov](#)). For example, in an educational setting two dozen biology students successfully used ChatGPT/Codex to generate Python scripts for DNA sequence manipulation (reverse-complementing and fragmentation) with 100% success (^[1] [pmc.ncbi.nlm.nih.gov](#)) (^[3] [pmc.ncbi.nlm.nih.gov](#)) and to count reads in multi-GB FASTA files (^[2] [pmc.ncbi.nlm.nih.gov](#)). We analyze these and other case studies, compare Codex to earlier tools like GitHub Copilot, and discuss Codex's strengths (multi-file memory, code execution, debugging) and limitations ([non-deterministic output](#), error rates around 30–50% per attempt (^[4] [pubs.rsc.org](#))). We also cover future implications: Codex promises to transform how biotech researchers code data analyses (freeing scientists to focus on biology above syntax (^[5] [pubs.rsc.org](#))), but safety and reproducibility must be ensured. In summary, Codex is poised to become a powerful collaborator in biotech R&D, dramatically speeding up coding tasks while requiring careful oversight (^[6] [openai.com](#)) (^[4] [pubs.rsc.org](#)).

Introduction and Background

The intersection of advances in artificial intelligence and biotechnology has grown increasingly significant. Modern biotech research—spanning genomics, proteomics, drug discovery, and synthetic biology—relies heavily on computational analysis. Historically, biologists have depended on specialized software and programming by bioinformaticians to process vast biological datasets. However, programming expertise can be a barrier for domain scientists. The advent of [large language models \(LLMs\)](#) and AI coding assistants promises to bridge this gap. Notably, OpenAI's Codex (building on GPT architectures) can generate code from natural language descriptions, enabling researchers to automate routine tasks without writing every line themselves (^[7] [pubs.rsc.org](#)) (^[1] [pmc.ncbi.nlm.nih.gov](#)).

Historical Context. Over the past decade, AI tools have gradually entered scientific workflows. Early efforts applied machine learning to problems like protein folding (e.g., AlphaFold) and genomics (predicting regulatory elements), while NLP models began assisting literature review. In programming, OpenAI's GPT-3 and later Codex (2021–2022) marked breakthroughs by generating functional code from text prompts (^[7] [pubs.rsc.org](#)) (^[4] [pubs.rsc.org](#)). In parallel, code-assistants like GitHub Copilot (2021) helped engineers auto-complete code. These tools, however, lacked the ability to run or test code comprehensively. In 2025, OpenAI introduced Codex as a cloud-based coding agent integrated into ChatGPT (the “research preview” launched May 2025 (^[8] [openai.com](#))), capable of executing tasks and managing multi-file projects (^[9] [bridgeinformatics.com](#)) (^[10] [help.openai.com](#)). This evolution arrived at a time when biotech data volumes – from next-generation sequencing, high-throughput assays, and multi-omics – were exploding, creating an urgent need for more scalable coding solutions.

What is OpenAI Codex? Codex is an [AI agent](#) specialized for software engineering tasks. It is powered by advanced models (e.g. GPT-5.3-Codex) trained on gigantic corpora of code and text. According to OpenAI, Codex-1 (the core model) was trained by reinforcement learning on real code tasks to “mimic human style” and iteratively test code until success (^[11] [openai.com](#)). Unlike static code-completion tools, Codex can autonomously navigate a code repository, make edits, run commands, and execute tests in a secure sandbox environment (^[10] [help.openai.com](#)) (^[12] [help.openai.com](#)). OpenAI emphasizes that Codex “can write functions, answer questions about your code, debug issues, and propose pull

requests” for review ⁽⁸⁾ [openai.com](#)). It runs on cloud servers (accessible via the ChatGPT interface, or through a command-line interface (CLI) and IDE plugins) and can operate in parallel on different tasks ⁽¹³⁾ [help.openai.com](#)).

Codex (and its evolution into GPT-5.3-Codex ⁽¹⁴⁾ [openai.com](#)) thus represents a new paradigm: a conversational coding assistant with memory across files and automated execution capabilities. Bridge Informatics notes that in pharmaceutical and biotech contexts, Codex is “more than a code completion tool, [it] is an autonomous, cloud-based coding assistant that can execute scripts, troubleshoot errors, validate outputs, and even propose changes across multiple files and repositories” ⁽¹⁵⁾ [bridgeinformatics.com](#)). Compared to earlier AI aids like ChatGPT’s basic code interpreter or GitHub Copilot, Codex enables end-to-end software workflows: one can ask it to refactor pipelines, test edge cases, and output a validated script for human review, all in one session ⁽¹⁶⁾ [bridgeinformatics.com](#) ⁽¹⁰⁾ [help.openai.com](#)).

This report explores how such capabilities can support scientific research in biotechnology. We cover technical details of Codex’s design, illustrate specific use cases (from sequence analysis to lab automation), analyze data and studies on its performance, and discuss implications for researchers. Throughout, we cite peer-reviewed research, industry case studies, and official documentation to ground our discussion.

Technical Overview of OpenAI Codex

Model Architecture and Performance. Codex’s underlying models (e.g. GPT-5.3-Codex) blend cutting-edge language understanding with specialized code generation prowess. As OpenAI describes, the latest version “advances both the frontier coding performance of GPT-5.2-Codex and the reasoning and professional knowledge capabilities of GPT-5.2” ⁽¹⁴⁾ [openai.com](#)). It achieves state-of-the-art results on industry benchmark suites (SWE-Bench Pro and Terminal-Bench) measuring real-world coding and CLI tasks ⁽¹⁷⁾ [openai.com](#)). GPT-5.3-Codex also exhibits “agentic” capabilities: it can handle extended interactions with follow-up prompts (e.g. iterative bug fixes) without losing context ⁽¹⁸⁾ [openai.com](#) ⁽¹⁹⁾ [openai.com](#)). Importantly, Codex can use heuristic understanding to default to sensible design choices: for example, generating more complete website templates compared to prior versions ⁽¹⁹⁾ [openai.com](#)).

The Codex model is proprietary to OpenAI and not open-sourced. Access is via ChatGPT’s Codex app or CLI, not via a public API endpoint (as of 2026). The Help Center notes Codex is included with all ChatGPT Plus/Pro/Enterprise plans ⁽²⁰⁾ [help.openai.com](#)). Data on exact model parameters is scarce, but its capabilities imply many billions of parameters fine-tuned on code.

Features and Tools. Codex’s primary interface is through ChatGPT: users can select the “Code Interpreter” or “Codex” mode to input prompts describing coding tasks. OpenAI also provides a dedicated Codex app (for Windows/MacOS) which can spawn multiple Codex agents working in parallel on separate projects ⁽²¹⁾ [help.openai.com](#)). Additionally, Codex offers a CLI and IDE extensions (e.g. for VSCode, [Cursor.ai](#), [Windsurf](#)) allowing in-situ interaction: “Starting from a prompt or spec, Codex navigates your repo to edit files, run commands, and execute tests” ⁽¹⁰⁾ [help.openai.com](#)). It can both pair with a developer locally (via prompt-and-response) or run tasks unattended in the cloud (each task isolated in a sandbox) ⁽¹⁰⁾ [help.openai.com](#) ⁽¹²⁾ [help.openai.com](#)). Notably, Codex has a feature to automate code reviews: it can be set to review pull requests on GitHub repositories, providing feedback similar to senior engineers ⁽²²⁾ [help.openai.com](#)).

Codex leverages internet access for research tasks: users can enable an “internet browsing” mode enabling Codex to fetch documentation or APIs beyond its training data. Internally, Codex has memory of multiple files and can revise code across them in one session, a leap beyond simple code completion. It is trained to follow instructions closely and can “iteratively run tests until success” ⁽¹¹⁾ [openai.com](#)).

Coding Languages and Libraries. Codex supports virtually all mainstream programming languages encountered in biotech: Python, R, MATLAB, C/C++, etc. In practice, Python is most common in scientific contexts, and Codex is proficient with libraries like NumPy, pandas, scikit-learn, BioPython, and others in its training corpus. It even understands domain-specific pseudo-code or query languages; for example, code queries to NCBI APIs (as in GeneGPT ⁽²³⁾

pmc.ncbi.nlm.nih.gov)) or BLAST command-line arguments. Codex's multi-language training allows it to suggest code for diverse tasks, from data parsing to simulation scripting.

Performance Characteristics. Codex's outputs are probabilistic. On standard coding problems, research suggests a single attempt yields a correct, runnable solution roughly 30–50% of the time ([4] pubs.rsc.org). Specifically, Hocky and White (2022) report that on real programming problems, Codex generates *exactly* correct code about 30% of the time with one prompt, improving to above 50% if multiple prompts/resamplings are attempted ([4] pubs.rsc.org). It can generate syntactically valid code not always following best scientific practices ([24] pubs.rsc.org), so user oversight is crucial. Codex can catch some of its own mistakes by writing code and unit tests, but direct testing by the user remains advised. The effectiveness of Codex depends heavily on prompt phrasing; precise, detailed instructions yield much better outcomes.

Table 1 compares Codex's high-level capabilities with prior AI tools:

Tool / Model	Code Generation	Code Execution / Testing	Multi-File/Memory	Context/Knowledge	Notes
GitHub Copilot (Oct 2021)	Autocomplete/in-line code suggestions in IDE ([25] bridgeinformatics.com)	X (no execution)	Limited to current file	Code-trained only	Limited to one-file context; no run. ([25] bridgeinformatics.com)
ChatGPT (typical)	Textual code snippets (with GPT-4, many languages)	X (unless Code Interpreter used)	Short context (~8k tokens)	General knowledge + fine-tuned	Good for explanation, not running.
ChatGPT Code Interpreter (Jul 2023)	Generates Python code; uses full language model	✓ (runs Python in sandbox)	Limited (one session, file upload)	Bases on GPT-4; has Python libs	Cannot navigate user codebase; only input files.
OpenAI Codex (May 2025+)	Generative code from detailed prompts ([1] pmc.ncbi.nlm.nih.gov)	✓ (multiple languages, sandbox)	✓ (remembers multiple files)	Code-specific (Codex-1); high reasoning	Cloud agent that can run tests and propose PRs ([9] bridgeinformatics.com).

Table 1. Comparison of AI coding tools. Data from OpenAI and Bridge Informatics ([26] bridgeinformatics.com) ([10] help.openai.com). ✓ indicates capability present; X indicates missing.

Using Codex in Bioinformatics and Biotech Workflows

Biotechnology research often involves coding tasks in data analysis and automation. Codex can support many such tasks:

- **Data Parsing and Transformation.** Converting and cleaning large biological data files (e.g. FASTQ, FASTA, VCF, PDB). For example, Codex can generate Python scripts to strip unwanted characters from sequences, reverse-complement DNA strings, or reformat data into standard formats ([1] pmc.ncbi.nlm.nih.gov) ([7] pubs.rsc.org).
- **Pipeline Development.** Building end-to-end workflows for sequencing analysis (quality control, alignment, variant calling, etc.). Codex can scaffold pipeline code by stitching together tools like samtools, BWA, and summarizing outputs, even generating shell scripts or Nextflow/Snakemake configurations. Bridge Informatics anticipates tasks like “refactor a pharmacogenomics pipeline ... output a clean, validated script” being handed to Codex ([16] bridgeinformatics.com).
- **Statistical Analysis and Visualization.** Writing analysis scripts in R or Python (pandas, SciPy) for experiments: computing differential expression, plotting results, or running machine learning models. Codex can use natural language to generate code that produces plots or statistical tests (e.g. ANOVA, t-tests) of experimental data ([7] pubs.rsc.org).
- **Database Access and Bio-APIs.** Querying scientific databases via APIs. As GeneGPT demonstrates, Codex can be prompted to formulate API calls to NCBI or elsewhere, parse results, and integrate them into answers ([23] pmc.ncbi.nlm.nih.gov). A researcher could ask, “fetch and count all human genes on chromosome 17 from Ensembl”

and Codex could generate code using Ensembl's REST API. GeneGPT achieved a GeneTuring score of 0.83 by using Codex to call NCBI tools (^[23] pmc.ncbi.nlm.nih.gov), exemplifying this capability.

- **Machine Learning and Modeling.** While advanced modeling might exceed simple prompts, Codex can still scaffold ML code. For instance, generating Python code to train a classification on omics data or to define deep learning models in TensorFlow/PyTorch, given appropriate prompts about layer architectures and loss functions.
- **Lab Automation and Robotics.** In principle, Codex can write control scripts for lab instruments (e.g. gating microfluidic devices, controlling pipetting robots, data collection loops). For example, a user might prompt Codex: "Write Python code to control a Biomek liquid handler to pipette sample columns A1 through A12 into a 96-well plate, using XYZ library," and Codex could produce a valid script. This is supported by its generality and execution capability. (Although direct case studies are sparse, similar usage is envisioned for coding tasks across disciplines (^[27] bridgeinformatics.com) (^[13] help.openai.com).)
- **Documentation and Reporting.** Beyond code, Codex can draft README files, documentation, or reports summarizing code, saving researcher time. In drug development, it can even help format regulatory reporting code (e.g. data submissions), as Bridge Informatics suggests, handling "clinical data wrangling" and submissions (^[28] bridgeinformatics.com).

Critical for these uses is the iterative nature of Codex: researchers can refine prompts and ask the agent to debug or optimize code. Codex's ability to maintain context means it can adjust a pipeline codebase when requirements change.

Workflow Integration

Researchers can access Codex through:

- **ChatGPT (AI Coding Partner).** The simplest way for biologists is via ChatGPT's Codex interface on chat.openai.com (Plus/Pro subscription). Here, users input prompts in conversational style: e.g. "Generate Python code to remove all non-ATCG characters from a FASTA file and produce its reverse complement." The response is executable code which can be run locally or pasted into an environment. ChatGPT also supports file uploads and the new Codex app mode, allowing multiple file contexts (^[10] help.openai.com).
- **Codex CLI and IDE Extensions.** For advanced use, the Codex CLI/extension can be used in local development environments. For instance, in VSCode, a researcher could open an existing project, then use a prompt like "Refactor the variant_analysis.py script to improve performance and add unit tests," and Codex will edit files and run tests. This mode is analogous to having an AI collaborator integrated into the codebase (^[10] help.openai.com).
- **Notebook Integration.** While not official, one can copy Codex-generated code into Jupyter notebooks or similar to execute and visualize results immediately, iterating in that interactive environment.

Best Practices. OpenAI's guidance emphasizes writing clear, specific prompts to get reliable code (^[1] pmc.ncbi.nlm.nih.gov) (^[10] help.openai.com). Researchers should review and test all generated code. Suggestions include asking Codex to produce unit tests or verifying outputs on known data (^[4] pubs.rsc.org). Because Codex's underlying model may not know lab-specific constraints or newer libraries, user oversight is vital.

Data Analysis Evidence

Multiple studies illustrate Codex's utility in bioinformatics tasks. One controlled experiment had undergraduate/graduate biology students use ChatGPT to perform sequence analysis coding. In this pilot course (University of South Alabama, 2025), students with little to no coding background used ChatGPT (essentially Codex through the interface) to generate Python scripts for genomic tasks (^[1] pmc.ncbi.nlm.nih.gov) (^[2] pmc.ncbi.nlm.nih.gov):

- DNA Sequence Manipulation:** All students achieved the goal of writing code to “remove all non-nucleotides and reverse complement a DNA sequence” ([1] pmc.ncbi.nlm.nih.gov). Each student’s ChatGPT-generated code was functional. This task is fundamental (cleaning raw reads), and Codex handled it with 100% success in this class. One student even combined two steps into one script, showing creative code structuring.
- Sequence Fragmentation:** The majority of students succeeded in using ChatGPT to generate a script splitting a long DNA sequence into 100-nt fragments in FASTA format ([3] pmc.ncbi.nlm.nih.gov). FASTA formatting (with headers) was executed correctly in most cases, enabling downstream analysis. Some students submitted partial attempts and later obtained working code from peers, but overall success was high.
- Counting Reads in Large FASTA:** In analyzing real sequencing data, students asked ChatGPT to write a program counting the number of reads in a huge FASTA file. All students produced working code for this, capable of handling 10–20 GB files in seconds ([2] pmc.ncbi.nlm.nih.gov). This code forms the basis for computing reads-per-million normalization in RNA-seq projects. (Note: three students’ solutions were notably more efficient, highlighting variation in Codex prompts.)
- Complex BLAST Parsing:** Students attempted to extract information from BLAST output (.xlsx) and retrieve matching sequences. Only 3 of 12 students succeeded in generating fully working code for this task ([29] pmc.ncbi.nlm.nih.gov). The difficulty underscores that more complex multi-step bioinformatics tasks can exceed a novice user’s ability to prompt effectively, or may require hand-off between peers. Yet those who did generate code were able to transform millions of alignments into usable FASTA subsets (see Figure 7, ([29] pmc.ncbi.nlm.nih.gov)). Working solutions were then shared class-wide, enabling completion.

Overall, this case demonstrates that even non-programmers can, via ChatGPT/Codex prompts, automate many routine bioinformatics steps. The study authors reported that integrating ChatGPT in teaching “bridge [s] a critical gap” between biologists and data analysis ([30] pmc.ncbi.nlm.nih.gov). Notably, students’ ChatGPT prompts were akin to actual research queries (e.g. “Split given DNA into 100 bp fragments.”). Figures from the study (3–7) show input prompts, generated code, and execution outputs ([1] pmc.ncbi.nlm.nih.gov) ([2] pmc.ncbi.nlm.nih.gov). This suggests Codex can serve as an on-the-fly coding tutor and workhorse.

Table 2: Example Tasks from Educational Case Study and Codex Outcomes

Task	ChatGPT/Codex Outcome	Citation
Remove non-nucleotides & reverse-complement DNA	100% success (12/12) – all students generated correct code ([1] pmc.ncbi.nlm.nih.gov)	
Split a DNA sequence into 100 bp fragments (FASTA)	≥83% success (10/12+) – majority succeeded ([3] pmc.ncbi.nlm.nih.gov)	
Count reads in a large FASTA file	100% success – all students generated working code ([2] pmc.ncbi.nlm.nih.gov)	
Extract unique read names from BLAST output & sequences from FASTA	25% success (3/12) – only a few students generated fully functional code ([29] pmc.ncbi.nlm.nih.gov)	

Table 2. Results from Delcher et al. (2025) pilot course using ChatGPT for genomics tasks. Codex successfully automated basic sequence manipulations, though more complex BLAST parsing proved challenging ([1] pmc.ncbi.nlm.nih.gov) ([2] pmc.ncbi.nlm.nih.gov).

Beyond education, emerging research shows Codex-like models aiding scientific inquiry more broadly. For example, OpenAI reports that GPT-5 (the next-generation model) significantly improved a molecular cloning protocol – the model introduced a novel recipe that increased cloning efficiency by **79x** compared to baseline ([6] openai.com). While this specific experiment focused on wet-lab AI reasoning (not code-writing), it illustrates the power of AI-suggested improvements in biotech workflows. Combined with Codex’s coding accelerations, future research could see integrated pipelines where AI suggests experimental protocols and generates the analysis code to interpret them.

Case Studies and Real-World Examples

GeneGPT: Domain-Specific Codex for Genomics

Researchers at the NIH introduced **GeneGPT**, an implementation of Codex fine-tuned with domain knowledge and tools for genomics question-answering. GeneGPT uses Codex (code-davinci-002) with carefully engineered prompts that include NCBI API documentation and examples (^[23] pmc.ncbi.nlm.nih.gov) (^[31] pmc.ncbi.nlm.nih.gov). By teaching Codex how to use real web APIs (e.g. Entrez, BLAST) through in-context learning, GeneGPT can answer complex genomics queries. On the GeneTuring benchmark (question-answer tasks about genes), GeneGPT achieved an average score of **0.83**, far outperforming ChatGPT's 0.44 and other baseline models (^[23] pmc.ncbi.nlm.nih.gov). This demonstrates that Codex can be guided to act like an expert bioinformatics agent, chaining API calls and processing results to deliver precise answers. For instance, GeneGPT may interpret a prompt like "Use NCBI to find mutations associated with BRCA1" by generating code that calls the appropriate services. According to Jin et al., GeneGPT's methodology increased success rates by 88% over previous best systems (^[32] pmc.ncbi.nlm.nih.gov). While not all biotech research is Q&A, this case shows Codex's potential when combined with domain resources.

AI-Assisted Research (Dog Vaccine Example)

High-profile anecdotes illustrate AI's emerging role in biotech. OpenAI CEO Sam Altman recounted a story of a tech founder (Paul Conyngham) who, faced with his Labrador Rosie's aggressive cancer, used ChatGPT to help devise a plan (^[33] tech.yahoo.com). Guided by AI, he sequenced the dog's tumor DNA, pored through mutation data, and even consulted structural models (AlphaFold was involved) to design a **custom mRNA immunotherapy** candidate (^[34] tech.yahoo.com). While Codex itself didn't generate the vaccine sequence, ChatGPT (a close relative) acted as a "smart research assistant" suggesting hypotheses and literature (^[35] tech.yahoo.com). This anecdote highlights the broader trend: biologists (or bio-entrepreneurs) are using LLMs and AI coding agents as research tools. Such tools can analyze genomic data, propose experiments, and even write scripts – tasks that previously needed expert teams. It underscores the message that Codex and siblings can accelerate discovery in life sciences.

Bridge Informatics Pharma Case (Pipeline Refactoring)

In the pharmaceutical industry, Bridge Informatics published a primer forecasting Codex's impact on drug development workflows. They note that Codex can "accelerate everything from genomic analysis pipelines to clinical data reporting systems" (^[36] bridgeinformatics.com). For example, a researcher could instruct Codex: "*Refactor our R-based RNA-seq analysis pipeline to handle new multi-omics inputs*", and Codex would propose code edits across the repository, run tests, and produce a merged pull request (^[16] bridgeinformatics.com). Bridge expects use cases like automated testing of bioinformatics scripts, iterative prototyping with scientists in the loop, and context-aware code edits for FDA/NDA document generation (^[27] bridgeinformatics.com). They emphasize Codex's advantage over previous tools: it *executes* and iterates on code, rather than merely suggesting lines (^[26] bridgeinformatics.com). While still anecdotal, Bridge's analysis indicates industry excitement and identifies concrete scenarios (e.g. pharmacogenomics pipelines, NGS data wrangling) where Codex can cut development time.

Detailed Use Cases in Biotechnology

In this section, we delve deeper into specific domains of biotech and how Codex can be applied, with examples and citations.

Genomics and Sequence Analysis

Genomics research is arguably the most code-intensive area of biotech, making it ripe for Codex. Common genomics tasks Codex can assist with include:

- **DNA/RNA Sequence Processing.** Cleaning raw sequencing outputs (FASTQ/FASTA) is routine. Researchers might prompt Codex: “Write Python code to trim all bases below quality 20 from FASTQ reads using Biopython, and write trimmed reads to a new file.” Codex can leverage libraries like Biopython to implement trimming logic. As shown in the educational study, ChatGPT-produced code handled trimming (removing non-nucleotide symbols) and reverse-complementing easily (^[1] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).
- **Alignment and Assembly Pipelines.** Codex can script workflows invoking BWA or Bowtie2 for alignment, breaking pipelines into steps. For example, “Create a shell script to run a BWA alignment of sample.fastq to human_g1k_v37, then use samtools to sort and index the SAM file.” Codex will produce a multi-command script.
- **Variant Calling and Annotation.** After alignment, variant callers like GATK or FreeBayes are used. Codex can generate code to run GATK, then annotate VCFs. For instance, “Generate a Python script that loads sample.vcf, filters variants with depth < 10 and quality < 30, and outputs a summary table of remaining SNPs.” Even if Codex makes minor mistakes, it can often provide a solid template for filtering and summarization.
- **Motif and Pattern Search.** Searching for DNA/RNA motifs (e.g. promoter regions) can be coded by Codex given a description. Example prompt: “Find all occurrences of the promoter motif TATA [A/T]A [A/T] [A/G] in the given sequence file and output their positions.” Codex can use regex or scanning algorithms. Students in the case study had ChatGPT generate code to count occurrences of k-mers across files .
- **BLAST and Homology Analyses.** Codex can automate calling NCBI BLAST (via command line or API) and parsing results. In the student project, ChatGPT-generated code extracted unique read names from BLAST output and retrieved corresponding sequences (^[29] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). A researcher might similarly ask Codex to: “Run BLASTN of all FASTA reads in sample.fasta against the P22 phage genome, then count how many hits each read has.” Codex would output a script using NCBI BLAST+ command-line tools and parse its TSV or JSON output.

Case Study: Students Analyzing RNA-Seq Data

A concrete example comes from the RNA-seq class project. (^[2] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)) Students used Codex to assist in their research on Salmonella sequencing data. One task was simply counting the number of reads in large FASTA files (several gigabytes). Prompting ChatGPT to “write code to count reads” yielded scripts that completed this task instantly (^[2] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). This enabled quick calculation of reads-per-million (RPM), a key normalization metric, without manual effort.

More ambitiously, students had a BLAST dataset (millions of matches). Two specific subtasks were to (a) extract all unique read names from the BLAST table, and (b) extract those reads from the raw FASTA for further analysis (^[29] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). For (a), Codex-generated Python code (using pandas/openpyxl) pulled column A (read IDs) from an .xlsx BLAST output and wrote them to a file. For (b), another Codex script read the large FASTA file and wrote out only sequences whose headers matched the extracted names. Three students succeeded end-to-end with working code; the others (expecting school credit) turned in partial results or peer-shared solutions. This resulted in a full pipeline from raw RNA-seq to a focused dataset, much of it prototyped by Codex. It demonstrates Codex can string together I/O and filtering operations on large files tractably.

Tools and Libraries

Codex knows popular bioinformatics libraries. For example, `Biopython` (for sequence I/O), `pysam` (for SAM/BAM), `Bioconductor` and `data.table` in R, or `scikit-bio`. A prompt like “Using Biopython, get GC content of each sequence in a FASTA file” yields Python code with `seqIO` and string methods. It can also use general tools: e.g., producing bash loops with `grep` and `awk` for simple patterns. CODex’s ability to incorporate best practices (like streaming files to avoid memory overload) depends on prompt detail. Users can refine by follow-ups, e.g. “Make the code memory-efficient.”

Proteomics and Structural Biology

Beyond nucleic acids, Codex can also support protein and small-molecule analysis:

- **Protein Sequence and Structure.** Calculating protein properties (isoelectric point, molecular weight) from FASTA inputs is straightforward in Python. Codex can import Bio.SeqUtils and generate such scripts. For structural data, prompts like *"Write code to read a PDB file, calculate distances between all cysteine pairs, and highlight disulfides"* can be answered using libraries (e.g. MDAnalysis or BioPython PDB module). Even elegantly invoking AlphaFold or RoseTTAFold from Python is possible in principle (e.g. *"write code to call AlphaFold on sequence X"*), though running those requires GPU backend, so Codex might sketch a subprocess call.
- **Mass Spectrometry Data.** Codex can help parse spectral data if given format hints. For example: *"Using pyOpenMS, create a Python script that reads a mzML file, extracts all peaks above intensity 1000, and plots their m/z distribution."* It may produce code with pyOpenMS or similar libraries if prompted correctly.
- **Molecular Docking and Screening.** A user might ask for code to automate a docking workflow: e.g., *"generate an AutoDock Vina docking script to test ligand.sdf against target.pdb."* Codex can write the shell or Python wrapper calling Vina, exhaustively sampling.

While literature examples in proteomics are fewer, these tasks follow the same pattern: give Codex a domain-specific problem, and it will often produce working analysis code.

Synthetic Biology and Lab Automation

In synthetic biology, experiments often require both simulation and physical execution steps. Codex may assist in designing genetic circuits (producing pseudocode or Python) or controlling lab equipment:

- **Gene Circuit Design.** A prompt could be *"Write a COBRAbios Python script to simulate flux balance in an E. coli metabolic network with a new gene knockout."* Codex would use COBRAPy library to load a genome-scale model, knock out a gene, and run FBA. Results (growth rate, pathway fluxes) would be printed. This sort of metabolic engineering task can be largely automated by Codex as long as the language description is clear.
- **CRISPR Guide Design.** For instance, *"Generate Python code to design gRNA sequences targeting gene X with no off-targets in genome Y"*. Codex might use Biopython to fetch gene sequences and a CRISPR library (though likely incomplete knowledge). Alternatively, hooking to a CRISPR API through Codex's code-writing could be done (similar to GeneGPT).
- **Automating Lab Instruments.** One could imagine prompts such as *"Using Opentrons Python API, write a protocol to transfer 100 μ L from wells A1–A8 of a deep-well plate to wells B1–B8 of a 96-well plate"*. Codex would script the Opentrons deck and pipette instructions. Indeed, Opentrons has an official Python API, and Codex's broad training should include examples. This is a concrete example of Codex bridging in silico plans to hands-on experiments. Such integration is forward-looking but technically feasible given Codex's capability to output instrument code.

Drug Discovery and Computational Biology

Codex can contribute to early-phase drug research tasks:

- **Virtual Screening Pipelines.** Generating workflows to screen compound libraries (e.g. calling OpenEye OMEGA or RDKit). Prompt: *"Write a bash script to use RDKit to filter SDF for molecules with MW < 500 and logP < 5, then output smiles."* Codex can chain RDKit commands in Python or shell.
- **Chemoinformatics Data Analysis.** Statistical analysis of assay results (linear regression for a dose-response curve, logistic regression for activity classification). Codex can create scripts in R or Python (using scikit-learn) to model this data.
- **"Wet lab" Planning.** Not coding per se, but by generating code to simulate experiment design. e.g. simulating population growth under inhibitors.

- **Documentation for Regulatory Submission.** Codex can format data analysis outputs or create compliance reports if prompted with relevant standards (like converting data frames to CSV with specific columns).

These tasks leverage Codex's general coding power; specifics (e.g. accurate pharmacology models) still need expert knowledge.

Performance, Limitations and Best Practices

Accuracy and Reliability. As noted, Codex's success is not guaranteed. According to Hocky & White, generated code is correct (meets specification exactly) about 30% on a single try, improving when multiple outputs are allowed (^[4] [pubs.rsc.org](#)). In practice, one may need to re-prompt Codex if the first answer is flawed. For research purposes, every piece of AI-generated code must be validated on test cases. Peer scientists advise that Codex is best used as a rapid prototyping tool: it can draft substantial portions of code quickly, but humans should verify logic and results.

Common error modes include off-by-one indexing bugs, misunderstanding ambiguous prompts, or relying on deprecated library functions. The RSC perspective warns that Codex can produce code that "does not follow best scientific practice" for complex tasks (^[24] [pubs.rsc.org](#)). However, RSC also found that even when mistakes occur, using Codex encourages scientists to articulate task requirements clearly (writing good prompts) and review results critically (^[5] [pubs.rsc.org](#)). Overall, they concluded Codex makes domain scientists *better programmers* by removing boilerplate work (^[5] [pubs.rsc.org](#)).

Security and Bias. Relying on large models and code-generation poses some risks. Codex has no awareness of data privacy; code execution happens in the cloud, and any proprietary data or code could be sent to OpenAI's servers. Users must be cautious about confidential datasets. Additionally, Codex's training data includes a mix of licensed, public, and user-submitted code, raising potential IP and licensing concerns for generated code (^[4] [pubs.rsc.org](#)). OpenAI's policies aim to mitigate this by disallowing certain uses, but organizations should still review legal aspects.

Scientific Best Practice. Any AI-assistance should augment, not replace, scientific judgment. Researchers should use Codex to automate routine coding and focus on interpreting results. It is recommended to maintain version control (e.g. Git) of auto-generated code and to perform code reviews. Also, documenting prompts and outputs is good practice for reproducibility. The NIH GeneGPT team's success was partly due to *in-context* documentation: including API docs in the prompt (^[23] [pmc.ncbi.nlm.nih.gov](#)). Similarly, a research team using Codex should embed comments and references in code to ensure interpretability later.

Discussion and Future Directions

OpenAI's Codex is at the frontier of AI-assisted programming, and its implications for biotech research are profound.

Accelerating Innovation. Codex can drastically shorten development cycles. Tasks that could take weeks of coding (especially for non-experts) can be done in minutes. This speedup can translate into faster discoveries: e.g. generating analysis code immediately after data collection, rather than waiting for a specialist. As OpenAI noted in a wet-lab experiment, GPT-5 accelerated molecular cloning experimentation by 79× (^[6] [openai.com](#)). In software terms, Codex could similarly amplify throughput. Biotech companies investing in AI are already seeing such gains: Bridge Informatics reports use of AI-coding can "move [projects] faster through the scientific literature and ... propose plausible mechanisms" (^[37] [openai.com](#)).

Accessibility and Education. Tools like Codex democratize data science. The case study showed novices could accomplish scripting tasks previously out of reach (^[1] [pmc.ncbi.nlm.nih.gov](#)) (^[2] [pmc.ncbi.nlm.nih.gov](#)). In the future, biologists with minimal coding background may routinely rely on AI to do heavy lifting, shifting training priorities toward

prompt engineering and result interpretation. This can broaden participation in computational biology. However, it also means curricula may need to adapt, teaching students how to critically use AI tools along with traditional methods.

Integration with Lab Automation. Looking ahead, we expect deeper integration between coding assistants and laboratory rigs. A “digital lab assistant” could use Codex to translate experimental designs into instrument protocols. Combined with robotics and automated experimentation platforms, this could create closed loops: design (via AI), execute (automated lab), analyze (AI), iterate faster than human cycles allow.

Multi-Modal AI Agents. Codex itself already spans code and text. Future AI agents (like GPT-5.3 and beyond) will further blur the line between code, data, and reasoning. The ongoing trend is toward unified agents: per OpenAI’s announcement, GPT-5.3-Codex now supports “all work in the software lifecycle” including writing PRDs, doing user research, and analyzing data in spreadsheets (^[38] [openai.com](#)). In biotech, this could mean an AI that not only writes analysis code but also understands lab notebooks, generates slides for reports, and even converses with humans about results.

Challenges and Risks. However, there are open challenges. **Reliability:** Ensuring scientific correctness is paramount; an AI’s plausible-sounding answer might be subtly wrong. **Reproducibility:** Black-box AI models complicate reproducibility of results. Journals and regulators may require clear documentation of AI involvement in research. **Security:** The Dog-vaccine example highlights biosecurity concerns: powerful AI could conceivably suggest how to engineer pathogens if misused. OpenAI itself conducted GPT-5 tests on benign tasks with strict controls (^[39] [openai.com](#)), but the community must develop guidelines for safe AI use in biology.

Compare with Other Pathway Developments. It’s worth noting that OpenAI is not alone: Google’s DeepMind recently launched “AlphaGenome” (June 2025), a new DNA-sequence model for variant effect prediction, and others are pursuing tool-enabled LLMs. The ecosystem will likely feature multiple AI tools for biotech. However, Codex’s niche is coding: no other major model currently offers the same level of automated code writing and execution. (Google’s “AlphaEvolve” project hints at AI-assisted algorithm discovery, but its biotech focus is not documented publicly as of 2025).

Conclusion

OpenAI Codex is a transformative technology for biotechnology research. As an AI coding agent, it bridges the gap between natural-language scientific questions and executable code. It allows researchers to offload rote programming tasks—such as data cleaning, pipeline scripting, and basic analysis—unlocking their time for higher-level insight. Empirical examples from education and industry confirm that Codex can indeed produce working code for genomics and bioinformatics tasks (^[1] [pmc.ncbi.nlm.nih.gov](#)) (^[2] [pmc.ncbi.nlm.nih.gov](#)), accelerating workflows and enabling non-programmers to contribute to computational projects.

The combination of Codex’s **multi-file context**, **built-in execution/runtime**, and **iteration capabilities** sets it apart from earlier tools (^[26] [bridgeinformatics.com](#)) (^[10] [help.openai.com](#)). It is essentially like having a 24/7 junior developer who understands biology-specific needs. However, Codex is not infallible. Studies warn of error rates (roughly 30–50% first-try correctness) (^[4] [pubs.rsc.org](#)) and document that users must critically review AI-generated code. Ethically and practically, researchers must ensure AI assistance boosts reliability, not undermines it, by thoroughly validating results and assessing biases.

Looking ahead, the trajectory is clear: Codex will become more capable (already GPT-5.3 is faster and smarter), and its adoption in biotech will grow. We will see more case studies of rapid prototyping in genomics, automated experimental design, and integrated lab-AI cycles. As OpenAI’s research suggests, AI can dramatically speed up biological research even in the lab (^[6] [openai.com](#)). Combined with Codex’s coding prowess, the future research lab may look very different: a collaboration between human biologists and AI agents. The implications reach from fundamental science to healthcare innovation, making tools like Codex vital to the next generation of biotech breakthroughs.

IntuitionLabs - Industry Leadership & Services

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

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

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

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

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

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

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

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

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

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

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

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

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

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

DISCLAIMER

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

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

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

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

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

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

© 2025 IntuitionLabs.ai. All rights reserved.