

ScienceIO

by Founder & CEO








Machine Learning

Large Language Models

Details

Review Date	02/07/2024
Purchase Date	Q3'23
Implementation Time	1 week
Product Still in Use	Yes
Purchase Amount	pay per API call
Intent to Renew	100%
Review Source	Vendor

Product Rating

Product Overall		5.0
Use Case Fit		5.0
Ease of Use		5.0
API		5.0
Integrations		N/A
Support		5.0
Value		5.0

About the Reviewer

Purchasing Team

Implementation Team

Reviewer Organization

N/A

Reviewer Tech Stack

GPT-4

Other Products Considered

N/A

Summary

- **Product Usage:** The product is used for enhancing clinical documentation, specifically for SOAP notes, leveraging both ScienceIO's and OpenAI's APIs.
- **Strengths:** The product provides reliable service, high-quality output, effectively extracts structured data from transcripts even with transcription errors.
- **Weaknesses:** While not significant, there has been mention of potential improvements in automating the fine-tuning of the output to make it more contextual and selectively filtered.
- **Overall Judgment:** The product has been beneficial in significantly enhancing the quality of clinical documentation, is reliable, and there hasn't been seen any competitors providing the same services.

Review

So today we're chatting about SciencelO and how it's used at your company. Before we jump into that, could you give a brief overview of the company and your role there?

We're a developer tools and technology company. Approximately 25% of our revenue is derived from the healthcare sector. We help our customers create innovative solutions using our developer tools. I'm a cofounder and engineer, and I supervise multiple product teams within the company.

What was the core business problem you were looking to solve with a product like SciencelO?

We needed to enhance clinical documentation, specifically for SOAP notes, due to increasing customer demand. Our customers required assistance in setting up clinical documentation workflows. We began prototyping SOAP notes generation and workflows using our developer tools. Although we achieved good results using large language models like GPT-4, we aimed for excellence. We explored various solutions, including combining SciencelO's structured medical data language models with GPT-4 via an API, to improve the quality of clinical documentation.

Were there any specific requirements you were looking for when you were initially evaluating SciencelO?

Our priorities were ease of integration, HIPAA compliance, data privacy, and in-depth healthcare expertise. We needed a vendor we could confidently recommend to our customers. And support was essential; we wanted to be able to inquire about the inner workings, especially due to the regulatory and compliance requirements in healthcare and the novelty of clinical documentation AI workflows. For us to be able to iterate successfully in production, it was important to be able to get engineering insights from SciencelO. Although their public documentation was comprehensive, the ability to engage with their engineering team provided an additional layer of confidence for us, even though it turned out to be unnecessary due to the reliability of their tools.

Were there any alternative solutions or other directly competitive vendors that you were looking into when evaluating SciencelO?

I was unable to find any direct competitors to SciencelO's healthcare-specific large language models. We had discussions with the hyperscalers who provide our infrastructure, but they don't offer anything similar to SciencelO. We've also networked with startups and health tech investors but only found companies with narrowly focused healthcare models, not the broad developer tooling that SciencelO provides. Testing of Google's public medical models showed they're inadequate for our needs. Interestingly, GPT-4 outperforms these models in medical contexts due to its advanced general capabilities. SciencelO is valuable for us because it generates highly effective prompts for GPT-4 so we can leverage GPT-4's strengths.

Since we launched and announced our SOAP notes developer tooling, I've engaged with many people and have yet to identify any direct competitors. An interesting development is the increased availability of datasets, which is beneficial for companies like ours that can fine-tune models. However, when it comes to an API-based solution designed to extract structured information from raw medical transcripts, I haven't come across any other offerings in the market.

Can you describe SciencelO's pricing structure?

The pricing structure is pretty standard. It's a self-serve SaaS with a pay-per-API-call system. Their team was helpful in clarifying how costs would scale with increased usage, which is vital when considering future growth and adoption by a larger customer base. We wanted to confirm that, as usage grows, the cost per API call decreases. We're satisfied with the agreement that, as our product's usage scales, the pricing becomes more economical.

How did you find the sales process?

It was fairly standard. We engaged with sales and sales engineering, gained API access, conducted testing, inquired about details, analyzed pricing, and ultimately signed a flexible contract that provided transparency on pricing and scale for both parties.

How was the onboarding and implementation process?

We had no issues. We successfully utilized the APIs as they were documented, and we efficiently executed the Business Associate Agreement (BAA) and processing agreements.

Can you describe the workflows or use cases in which you're leveraging the SciencelO models?

We have transcripts from telehealth sessions available in our system. We send them to SciencelO's data API and receive a JSON response. After trimming the JSON information to suit our needs and reduce latency and costs, we enhance it with metadata from the telehealth session and tailored prompts for various medical contexts. This enriched data is then input into GPT-4, which generates a draft clinical document. We subsequently provide this to our customers for integration into their user workflows. Essentially, we're an intermediary, leveraging both SciencelO's and OpenAI's APIs along with our enhancements before delivering the final product to our clients.

We're also giving our customers the option of utilizing SciencelO's Redact API—an API focused on anonymizing patient information. This is considered next-generation technology in some of our client discussions and could be extremely beneficial for them. Although we're not using it yet in production, its uniqueness stands out, as I haven't come across an exact match for this set of APIs, especially not in the medical context. We prefer to use medically specific models in healthcare workflows. Looking ahead, I anticipate that we will integrate these APIs into the developer tools we provide our healthcare customers in 2024.

How do you measure how much SciencelO is helping your quality of output?

We have a test data set that we process through different workflow versions. This output is then rated by clinicians or target users on a one-to-five-star scale. We've conducted this evaluation with and without SciencelO using OpenAI's flagship models, other open-source models, and some specifically fine-tuned by us. The results consistently show that incorporating SciencelO significantly enhances the quality of the results. It makes sense to use SciencelO to provide the highest quality output possible to our customers.

Interestingly, we found that SciencelO's efficiency could allow us to utilize less resource-intensive models than GPT-4 while still maintaining high-quality output, as indicated by healthcare provider ratings. However, we currently opt for the best possible output without compromise. As the landscape evolves, with considerations like cost, the need for low latency, and desire for quicker turnaround times, SciencelO could be instrumental in finding the right balance between performance and resource utilization.

What would you consider as the overall strengths and weaknesses of SciencelO?

Strengths include the reliability of the service and the quality of output from its API. The system can effectively extract structured data from transcripts, despite transcription errors such as incorrectly transcribed drug names. As we've started to expect from AI in 2023 and 2024, it's able to clean up the transcript and pull out the important information.

As far as weaknesses, there has been some discussion about automating the fine-tuning of the output, making the JSON output more contextual and selectively filtered. However, this is not a significant issue, since the data can be manually trimmed to meet specific needs. I expect that more flexibility like this will be developed over time as breadth of usage increases.

Are there any other aspects of the product worth calling out?

It's uniqueness. SciencelO transforms unstructured information into structured data, with exceptional proficiency in medical contexts. It's a general-purpose model with specialized competence in healthcare terminology and data structuring. Its ability to accurately identify and contextualize drug names, even from poorly transcribed inputs, is particularly impressive. During tests with our dataset, I often had to rely on multiple Google searches to decipher drug names from raw transcripts, whereas SciencelO's technology managed this accurately in one attempt.

How was the developer experience and documentation?

All the documentation is accessible on the public website, and the APIs do what they say they're going to do. Integration and development were straightforward. However, we were surprised by the volume of information that's returned. You probably don't need everything that is returned in the JSON, so you can select only the necessary data points for your specific workflow.

Can you share any learnings from how your transformation phase of the workflow has evolved over time?

We started simply but evolved as usage increased. We focused on optimizing token count, which is a significant consideration in AI workflows. JSON's verbosity uses extra tokens, so we restructured our data into a more token-efficient format for better workflow efficiency.

How have you found the account management and support teams?

We quickly got up and running, and the process felt almost self-serve, though we did sign the documentation required for healthcare standards. As we progressed toward production, we engaged the account management team to secure long-term pricing visibility and integrate the necessary documentation into our privacy policy for new customers. The process was streamlined, and their operation meets the stringent requirements of the healthcare sector.

Do you think you made the right decision in moving forward with SciencelO?

Yes, definitely.

Where do you see interesting growth areas for the product?

The company likely has an extensive roadmap for enhancing their platform, which has a solid base in converting unstructured to structured data. Potential improvements could include customizable outputs and more efficient data processing, but they're not essential at the moment. I assume SciencelO is experiencing significant growth

across various healthcare sectors, since their technology for handling unstructured data is universally applicable and beneficial.

Do you have any general advice for folks who may be evaluating this sort of product today?

I haven't seen any competitors for what ScienceIO does. If you need to convert transcripts into a programmatically operable format, it does it really well.