HARNESSING GENERATIVE AI: LEVERAGING OPENAI'S APIS TO ENHANCE ACOUSTIC FIELD NOTE-TAKING

S Su   iFieldnotes, Edinburgh, UK

1   INTRODUCTION

In less than two years, Generative AI (GenAI) has rapidly evolved from a fascinating novelty to a widely used tool among acoustic professionals[1]. However, most practitioners engage with GenAI through ready-made user interfaces like ChatGPT and Claude on an individual basis. While these user-friendly interfaces are easily accessible, this approach falls short of realising the technology's full potential. These interfaces, while convenient, often struggle with complex, multi-step tasks and lack the scalability required for widespread implementation in professional workflows.

This paper, written from the perspective of an acoustic consultant, presents a case study of iFieldnotes, a mobile app powered by OpenAI APIs designed for field note-taking. Through this example, it illustrates how integrating GenAI APIs as a development tool can unlock GenAI's substantial potential. It demonstrates that GenAI APIs enable the execution of more complex tasks on a system-wide scale, offering a balance between customisability and scalability of AI solutions.

The primary objective of this study is to highlight the advantages of leveraging GenAI's power through its APIs, aiming to broaden the perspective of acoustic professionals beyond the limitations of current user interfaces. The exploration of the capabilities of API integration is to inspire innovative applications of GenAI that can revolutionise existing workflows and methodologies in the field of acoustics.

2   BACKGROUND

2.1   Methods of applying AI in acoustics

From the AI in Acoustics conference[1] held in May, the application of AI in acoustics has currently been approached through two primary methods: building custom machine learning models from scratch, and utilising Generative Pre-trained Transformers (GPTs), hereafter referred to as GenAI.

Building custom machine learning models involves developing and training AI algorithms specifically for acoustic applications, such as noise pattern analysis and source identification. This approach can be highly effective for a specific problem but also presents significant challenges. It typically requires vast amounts of high-quality, domain-specific data, which can be costly and time-consuming to collect and process. Additionally, developing effective models necessitates substantial expertise in AI and data science. As a result, this method has been predominantly feasible for specialised AI firms and large engineering companies with extensive resources. Furthermore, the energy consumption associated with processing such large datasets further raises sustainability concerns.

In contrast, the application of GenAI leverages existing AI models that have been trained on massive and diverse datasets. This approach offers several distinct advantages. It allows users to achieve meaningful results without the need for extensive domain-specific datasets. The availability of pre-trained models makes advanced AI capabilities more accessible to a broader range of professionals, including smaller firms and individual consultants. These models can be used for a variety of tasks, from generating reports to assisting with problem-solving in acoustics. Although these general-

purpose models may lack specialised acoustic knowledge, they provide a versatile starting point that can be fine-tuned or customised for specific applications. This flexibility enables quicker implementation of AI solutions in acoustic projects, and reduce computing energy compared to building machine learning models from scratch.

Overall, using pre-trained GenAI models is more accessible for acoustic professionals in industry and offers significant potential for innovation, presenting exciting opportunities for the future of the field.

2.2 Practical pathways of leveraging GenAI for acoustic tasks

Pre-trained GenAI models possess digital power comparable to the energy generated by electric power stations. Figure 1 illustrates an analogy between the transmission of digital power to acoustic applications and the transmission of electrical power to appliances.

Electric power transmission from electric power stations to appliances

Digital power transmission from pre-trained GenAI models to acoustic applications
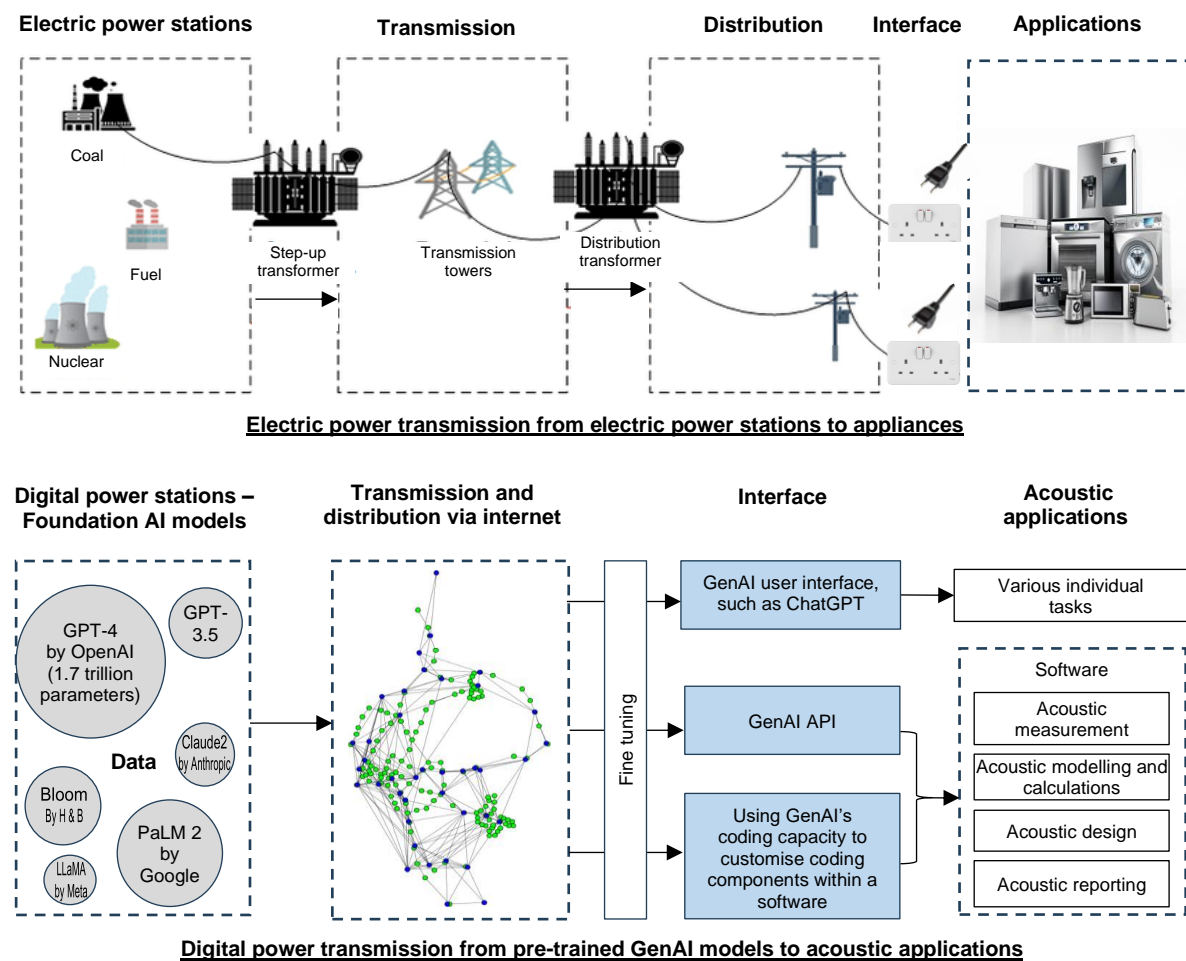
Figure 1. An analogue between the transmission of digital power to acoustic applications and the transmission of electric power to appliances (The sizes of the pretrained GenAI models are proportional to their parameter sizes, as referenced from Wikipedia[2] and Khawaja[3])

In this analogy, pre-trained GenAI models function as digital power stations, where data serves as the fuel. The relative sizes of the circles in Figure 1 reflect the parameter sizes of these models, indicating their digital power capacity. Internet infrastructure acts as the transmission grid, while the fine-tuning process is akin to electric transformers, adjusting the power for specific needs. Acoustic applications, in turn, are comparable to electric appliances, tailored to utilise this digital power.

But what connects this "digital power grid" to acoustic applications? The author has identified three main pathways: using a user interface like ChatGPT, utilising GenAI's API calls such as OpenAI APIs, and employing customisable coding components within software to harness GenAI's coding capabilities.

At last year's annual Institute of Acoustics conference, the author presented a paper exploring the third method. The study involved using ChatGPT to generate Grasshopper Python code for optimising the angle of a reflector in a theatre to maximise sound projection to the audience. While this method shows promise in environments where software packages offer customisable coding components or development kits, it addresses only specific, isolated problems rather than offering scalable, system-wide solutions. Consequently, its broader impact remains limited.

Building on this, the following section compares the first two pathways: using a GenAI user interface and leveraging GenAI's API calls. To illustrate these two pathways, ChatGPT and OpenAI APIs serve as representative examples.

2.3   ChatGPT vs using OpenAI API calls

Both ChatGPT and OpenAI APIs are powered by the same family of pre-trained AI models developed by OpenAI, sharing core capabilities such as natural language understanding, text generation, and task completion across various domains. They both rely on "prompt engineering", where text inputs guide the AI models to effectively complete tasks.

The key differences are as follows:

- ChatGPT offers a straightforward chat interface, whereas OpenAI APIs are programming interfaces.
- ChatGPT is user-friendly and accessible to the general public, while OpenAI APIs require programming expertise to utilise.
- ChatGPT offers a fixed interface with predefined functionalities, whereas OpenAI APIs provide greater flexibility, allowing for customisation and integration into specific applications or workflows.
- Users of ChatGPT have limited control over the model's behaviour. In contrast, API users can fine-tune parameters, manage the context, and even fine-tune the model on domain-specific data.
- ChatGPT is designed for individual user interactions, whereas APIs can handle large-scale, automated processing of data or queries.
- ChatGPT functions as a standalone tool, while APIs can be directly integrated into existing software systems, including specialised tools for acoustic analysis.
- ChatGPT provides a general-purpose interface, while APIs allow developers to create custom interfaces tailored to specific needs.

Figure 2 illustrates the interfaces of ChatGPT and an example of OpenAI API calls. The process for using OpenAI APIs involves importing the OpenAI library, setting up an API key, and using the Chat Completion method to obtain results. Users can further customise parameters of the GenAI models by selecting engines, defining the maximum allowable tokens, and adjusting the randomness of the generation.

If the individual functions of GenAI are analogous to the pieces of a "magic seven-piece" Tangram, then ChatGPT's simple user interface is like a rigid square frame that constrains its potential. Using OpenAI APIs liberates these pieces, allowing them to be customised into numerous combinations, much like a Tangram, to suit different workflows. These combinations can then be duplicated and scaled, leading to the creation of more powerful applications.

Currently, much of the focus of acoustic professionals has been on GenAI's user interface for individual tasks, while the potential of GenAI as a developer tool through API calls remains underappreciated.
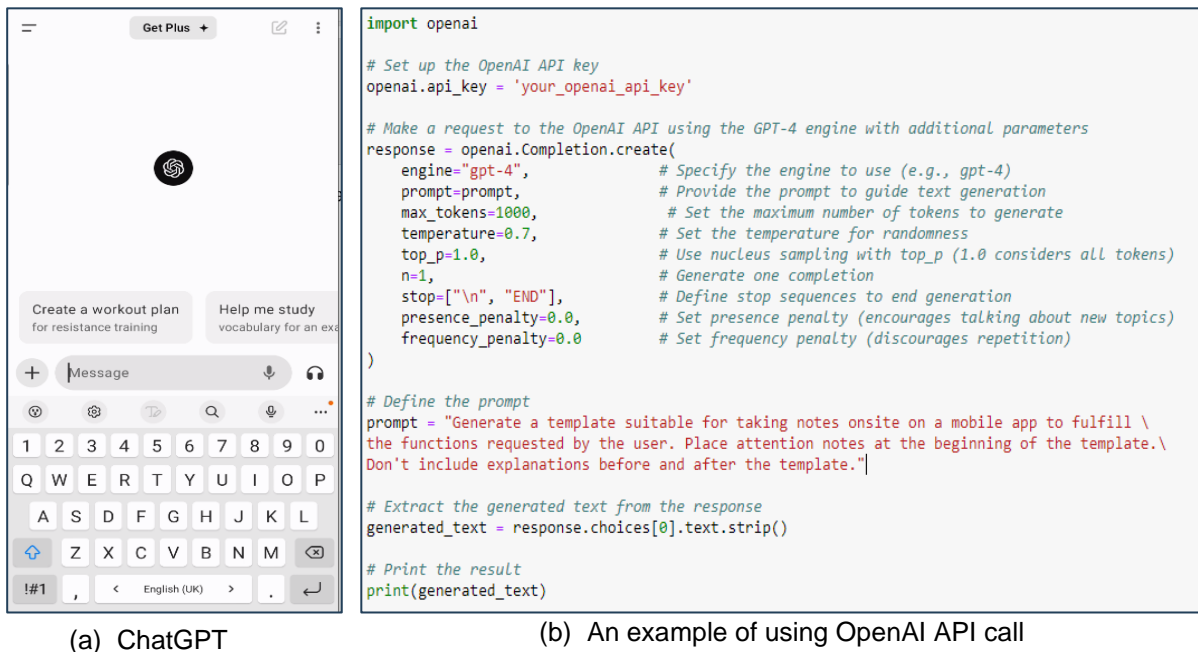
(a) ChatGPT

(b) An example of using OpenAI API call

Figure 2. Interfaces of ChatGPT and OpenAI API

3 NECESSARY OF ENHANCING ACOUSTIC FIELD NOTE-TAKING

The application of GenAI should always begin with an analysis of existing workflows to identify practical needs, rather than adopting the technology solely because it is trendy.

As an acoustic consultant, the author has constantly sought opportunities to improve work efficiency and quality throughout the entire circle of acoustic consultancy, while experimenting with the capabilities of GenAI and making connections between the two. The circle of acoustic consultancy spans from the preparation of fee proposals, collection of relevant information both offsite and onsite, and data analysis, to the preparation of reports, communications, and project management throughout the entire process.

The author identifies the tasks with the following characteristics have great potential to be optimised or reimagined due to the rise of GenAI:
- Repetitive: Tasks that are repeated frequently, often with little variation.
- Standardised: Processes that follow a set pattern or format.
- Monotonous: Tasks that are dull and uninspiring.
- Non-digitalised: Information or data that is not in a digital format, which is essential for training AI models. Notably, it is projected that publicly available data could be exhausted between 2026 and 2032.[4] Therefore, high-quality, proprietary, domain-specific data will become increasingly vital for future AI development.[5,6]

The current field note-taking practices exhibit the above characteristics. The process is generally standardised for each type of survey, and it is sometimes repetitive and monotonous, and lacks digitalisation.

A pilot survey conducted by the author shows that half of acoustic consultants use pen and paper while half use a generic mobile app for field notes, complemented by mobile phones for capturing

photos, videos, and audio recordings. The participants of the survey recognised the following challenges of the current onsite note-taking practice:

- Integrating photos/sketches/multi-media files with notes. Scattered information makes it challenging to consolidate and communicate information effectively among team members
- The quality of site notes is often compromised due to time constraints
- Difficulty in organising notes
- Sharing notes with team members
- Environmental factors, such as cold weather and dark nights
- Inability to take notes due to having full hands
- Handwriting is often messy and difficult for others to read.

Field notes, which document observations and contextual details of a survey site, form the cornerstone of any acoustic analysis. Improving their quality is particularly valuable in light of current AI advancements. GenAI's versatile capabilities offer a significant opportunity to enhance the note-taking process by intelligently processing shorthand notes and multimedia files. However, fixed user interfaces, such as those found in applications like ChatGPT, are not well-suited for integration into specialised workflows. In contrast, GenAI's APIs, such as those offered by OpenAI, provide the flexibility to customise AI functions alongside non-AI features, tailoring solutions to specific needs.

iFieldnotes, an innovative mobile app for field note-taking, powered by OpenAI's APIs, has been developed with this inspiration.

4 OVERVIEW OF CAPABILITIES OF OPENAI APIS

To effectively leverage OpenAI API calls for enhancing acoustic field note-taking, it is crucial to understand the full range of their capabilities. OpenAI's APIs provide developers with access to a suite of pre-trained large language models (LLMs) and audio AI models. Additionally, OpenAI offers APIs that integrate external tools, such as code interpreters, function calling, and file search, further enhancing the functionality of their AI models.

Below is an overview of the key capabilities of OpenAI API calls relevant to note-taking, including processing text, images, and audio recordings, as well as integrating tools with AI models.

4.1 Text

One of the core capabilities of the OpenAI API is processing text. This feature is highly versatile and includes a wide range of applications such as:
- Creating text from scratch or based on minimal input to improve writing and generate content.
- Providing detailed answers to questions about specific text.
- Converting text from one language to another.
- Categorising content into different topics, which is useful for content management systems or research databases.

4.2 Images

The models like GPT-4, GPT-4 mini, and GPT-4 Turbo have vision capabilities, meaning they can interpret and analyse images. These models can understand information in an image with a resolution similar to what an average human can perceive. If the information is not visible to the average human eye, the models will not be able to interpret it either.

4.3 Audio

OpenAI's Audio API provides access to Whisper, a robust audio AI model that can be used for:

- Converting audio into text in the same language as the audio.
- Translating audio from its original language into English while transcribing it.

The Audio API currently supports file uploads in formats such as mp3, mp4, mpeg, mpga, m4a, wav, and webm. File sizes are limited to 25 MB, which corresponds to approximately a 26-minute recording at 128 Kbps in mp3 format for a mono track, or a 30-minute recording at 96 Kbps in M4A format.

4.4   Multi-model functionality

The multi-modal functionality of OpenAI's API allows for the use of different types of inputs and outputs, including text, images, audio, and even code, within the same framework. This flexibility makes the API highly versatile and adaptable for various applications across multiple domains, enhancing its usefulness for complex tasks that require processing diverse data types.

4.5   Code interpreter

OpenAI's APIs also include features specifically designed to enhance and streamline workflows involving Code Interpretation, function calling, and file search. The code interpreter function enables the execution and interpretation of code in various programming languages, which is particularly useful for tasks such as data analysis and performing complex calculations.

The code interpreter can process up to 20 files simultaneously, each up to 512 MB in size and containing no more than 5,000,000 tokens. By default, the total size of all files uploaded in a project cannot exceed 100 GB, though this limit can be increased by contacting the OpenAI support team.

4.6   Function calling

The function calling feature allows the API to interact with user-defined functions or external systems. This capability is particularly valuable for extending the API's functionality beyond text generation, enabling it to dynamically interact with other software and services, such as
- Fetching data from an internal system before it can generate the response to the user
- Enabling actions based on user preferences and calendar availability.
- Building rich workflows
- Modifying an applications' user interface

4.7   File search

The file search function enables searching and retrieving information from various file types, including PDFs, Word documents, Excel spreadsheets, JSON, and more. This function is useful for:
- Extracting text from uploaded files.
- Searching for specific keywords, phrases, or patterns within files to quickly locate relevant information.
- Providing summaries of content from large documents, making it easier to understand key points quickly.
- Extracting structured data from documents, such as tables in Excel or form fields in PDFs, to facilitate data analysis and manipulation.

The file search function can process up to 10,000 files simultaneously, each up to 512 MB in size and containing no more than 5,000,000 tokens. By default, the total size of all files uploaded in a project cannot exceed 100 GB, though this limit can be increased by contacting the OpenAI support team.

By understanding and leveraging the aforementioned capabilities, iFieldnotes has been developing solutions using OpenAI API calls to create more efficient solutions that enhance workflows and improve the quality of notes. The sections below showcase some example AI functions.

5   APPLICATIONS OF OPENAI APIS IN IFIELDNOTES

iFieldnotes uses OpenAI APIs to provide two types of AI functions. The first, called "Custom GPTs", allows users to generate new content, such as performing acoustic calculations. The second, known as "Ask AI menus", enables users to query the AI about the content within a note, including text, images, audio, one or more notes. A user can customise both GPTs and Ask AI menus to suit individual needs and diverse projects.

5.1   Custom GPTs

By integrating OpenAI's chat completion API into an app, developers can redesign the user interface to optimise workflow with OpenAI's GPT. In iFieldnotes, a dedicated section for "Custom GPTs" allows users to tailor GPTs by defining names, system prompts, and instructions for specific requests. These customised GPTs can then be used when writing notes or creating templates in iFieldnotes, with the generated content seamlessly integrated into the documents.

5.1.1   Usage example - Automate note-taking templates

The cornerstone of acoustic surveys is to capture observations of the survey site and contextual details of measurement data. It is beneficial to have templates appropriate for different types of surveys prepared in advance. This ensures that all necessary details are systematically captured, reducing the risk of missing critical information and improving the organisation of the data. It also helps ensure that surveys are conducted in compliance with relevant standards and guidelines.

The user can customise a GPT for automating templates. An example is shown in Figure 3 (a), named Template Generator. When a user wants to automate a specific type of template in the Template Pad, they only need to select "Template Generator" as shown in Figure 3 (b) and (c), and input the specific type of template they want to generate, such as "Road traffic noise survey to CRTN," as shown in Figure 3 (d). If necessary, the user can edit the GPT-generated template in the Template Pad of iFieldnotes. A user onsite can quickly start a note by inserting the required type of template with one click, as shown in Figure 3 (f).
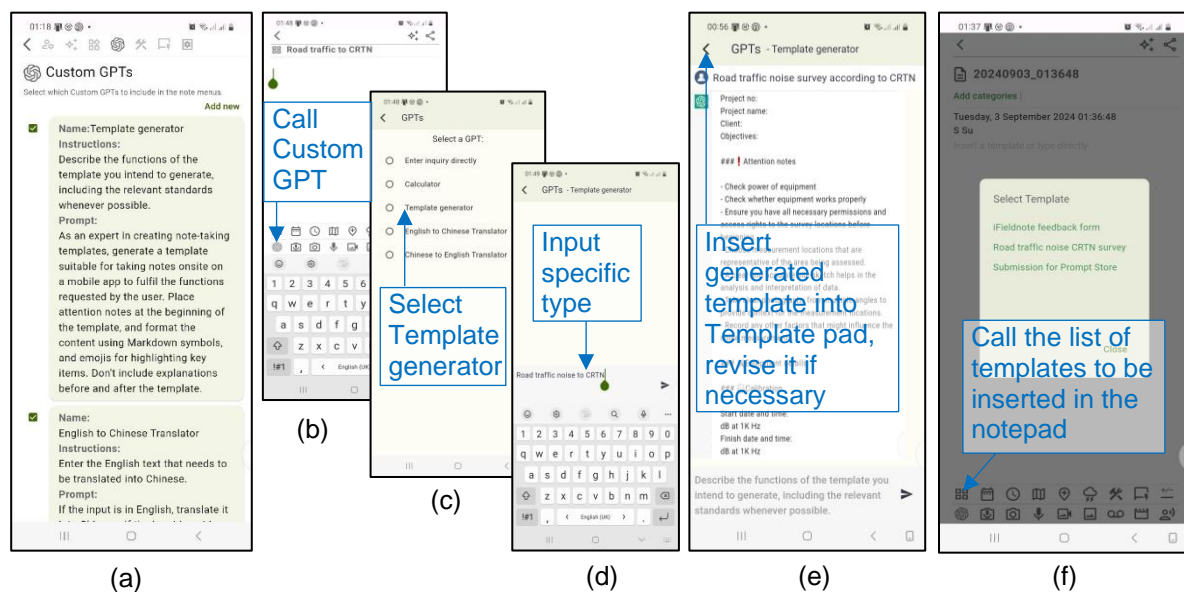

Figure 3. The customised workflow and user interface for automating a template and inserting it into a note, with assistance from the OpenAI API.

5.1.2 Usage example - Carry out quick acoustic calculations onsite

Custom GPTs in iFieldnotes can be configured to create a powerful calculator using GPT-4o API combined with the Code Interpreter tool via the Assistants API. This setup allows for real-time acoustic calculations in the field, enabling surveyors to quickly assess sound levels and other important parameters without relying on external tools or waiting for post-survey analysis. This capability facilitates immediate adjustments and responses on site.

For example, after measuring noise levels from various sources, a surveyor can use the Calculator to instantly determine overall noise levels. This enables immediate assessments, such as checking if cumulative noise levels exceed acceptable limits or evaluating the effectiveness of noise mitigation measures, and deciding if further measurements are needed.

While ChatGPT can perform similar functions, Custom GPTs in iFieldnotes streamline the process by eliminating the need for experimenting with prompts to find the most suitable one onsite. It also save the time of typing lengthy prompts and allowing seamless integration of results into notes without switching apps.

Note that GPT-4o has limited calculation capabilities, particularly for dB calculations. The Code Interpreter tool via the Assistants API converts requests into Python code, ensuring accurate and reliable results.

5.2   Ask AI menu feature

By incorporating OpenAI APIs for text, vision, and audio models into iFieldnotes, users can ask the AI about the text, images, or audio at the cursor's location. The AI can be prompted to analyse the entire note and multiple notes. Users can customise the "Ask AI" menu by defining custom names and prompts for each option. Figure 4 illustrates examples of "Ask AI" menu configurations and their use within a note.
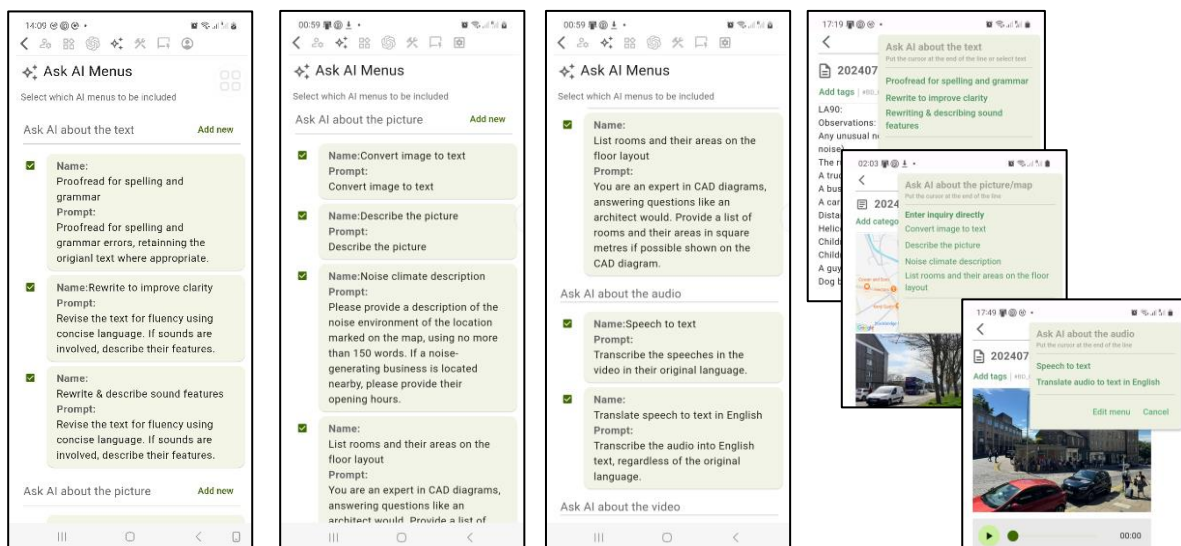
(a) Example of Ask AI menu configurations   (b) Use of Ask AI menus in note

Figure 4. Example settings of Ask AI menus and their use within a note

5.2.1 Ask AI about Image

Visual data, such as photos and maps, is essential in acoustic surveys for contextualising the sound environment. Traditionally, surveyors write down notes on paper or use a generic app, and store

maps and photos separately. With OpenAI API for GPT-4o's vision model, surveyors can generate detailed descriptions from a map or a photo and revise it if necessary. This improves note quality and saves time by seamlessly combining visual and textual data, offering a more comprehensive assessment of the acoustic environment.

Using OpenAI APIs, text from printed material or handwriting in a photo can be easily converted into digital notes. For example, during a plant noise survey, you can take a photo of a plant item's label and use the "Ask AI" menu to extract and convert the text. Similarly, handwritten notes can be incorporated by photographing them and using the same menu to turn them into text.

The "Ask AI" menu, with prompts to OpenAI's vision API, allows iFieldnotes to interpret inserted images, such as floor plans. It can list all the rooms and calculate floor areas based on dimensions or known room sizes.

5.2.2 Ask AI about Audio

In many acoustic surveys, taking detailed written notes sometimes may not always be practical. Instead, surveyors can use iFieldnotes' audio recording function to capture their observations. The "Ask AI About Audio" menu can then automatically transcribe these voice notes into text, allowing surveyors to document their findings quickly and accurately without interrupting their workflow. This is particularly useful in challenging conditions, such as cold weather or nighttime, where writing is difficult. Non-native English speakers can also record in their native language, with the speech automatically transcribed into English.

For accurate transcription, the speech-to-noise ratio (SNR) should be above zero; otherwise, GPT may hallucinate content.

5.2.3 Ask AI about Text

It is often the case that acoustic professionals can only quickly jot down observations in shorthand due to time constraints. Using the Ask AI menu, powered by OpenAI's GPT, these shorthand notes can be transformed into well-structured, grammatically correct text. This is particularly useful for improving the readability and professionalism of survey notes. It also saves time in the preparation of the final report.

5.2.4 Ask AI about a note

As OpenAI GPT-4o includes multi-model functionality, the Ask AI menu can be prompted to rewrite, analyse, and summarise a note covering not only the text and but also the images.

5.2.5 Ask AI about selected notes

Customising a search engine for private files, like an internal Google, is usually expensive. OpenAI's assistant API offers a low-cost alternative with its file search function, which can also provide overviews and analyses of searched notes.

The file search function is limited to 10,000 files at a time, with each file up to 512 MB and no more than 5,000,000 tokens. While this may be impractical for searching all proprietary documents within an organisation, it is ideal for note-taking, as it's unlikely that individual notes will exceed these limits. To further optimise cost and accuracy, iFieldnotes introduces a category filter, allowing users to select specific notes for AI search and analysis (this feature is in development and is not included in the current beta testing version).

5.3   Scalability of iFieldnotes' AI solutions

To maximise the value of AI solutions, scaling them is essential. One way to achieve this is by packaging customised AI prompts into user-friendly applications like iFieldnotes. These apps can then be distributed to potential users through widely accessible platforms such as Google Play and the Apple Store.

Developing highly effective AI prompts often requires time, trial and error, and multiple iterations. Recognising this, iFieldnotes has designed features that not only utilise AI but also allow users to scale and share their own AI solutions. iFieldnotes enables users to export their custom settings, including AI configurations (Custom GPTs and Ask AI menus) as well as non-AI settings like Insert Widgets, Equipment List, and Quick Text options. This exportability allows users to easily share their optimised setups with team members, enhancing collaboration and efficiency.

Furthermore, iFieldnotes provides a marketplace on its website where users can sell their custom configurations. Potential buyers can import these pre-configured settings with a single click, saving time and benefiting from others' expertise. A rating system will be introduced to help find the most useful prompts. This ecosystem fosters innovation and knowledge sharing while providing potential monetisation opportunities for users who develop particularly effective AI solutions.

6   DISCUSSION

As detailed above, the application of OpenAI APIs in iFieldnotes streamlines the note-taking process by reducing the time required for manual input while simultaneously improving the overall quality of the notes. Additionally, the APIs enable automatic analysis, offering insights that would traditionally require manual effort. Despite these advantages, it remains essential for a human acoustic expert to oversee the process to ensure that AI-generated outputs are accurate and align with professional standards. In this context, AI acts as a supportive tool or "co-worker" rather than a replacement for human expertise in acoustic consultancy.

When using GenAI APIs, note data must be submitted to GenAI cloud services. So ensuring data privacy and security is crucial when choosing reliable providers. OpenAI assures users that data processed via their APIs will not be used for model training. Additionally, some companies are developing on-device AI systems for mobile phones, providing an extra layer of protection.

The growth of GenAI brings significant innovations, but it also has costs. One of these is increased electricity demand, which conflicts with the goal of achieving net-zero greenhouse gas emissions[7]. Therefore, it is essential to remain aware of GenAI's environmental impact. As GenAI users, the following steps can help reduce its carbon footprint:

- Choose vendors that source electricity from renewable energy as much as possible.
- When feasible, use existing generative models and fine-tune them to meet specific needs, rather than creating entirely new ones. Reusing models and resources can also help reduce the environmental cost.
- When using AI file search functions, only process the data required to meet the needs of a specific use case.

The development of iFieldnotes stemmed from the author's experience in acoustic consulting. Recognising the potential of AI, the author has explored how AI capabilities can address specific challenges within the field. However, as the software has not yet been widely tested, pilot studies are necessary to demonstrate measurable improvements in acoustic field note-taking.

Although it is too early to fully assess the impact of iFieldnotes, it is evident that domain knowledge is crucial for effectively leveraging AI. As Mollick[8], a prominent author and New York Times bestseller, argues, innovation in AI applications within organisations often starts at the periphery—among individual professionals—rather than in centralised digital departments. He advocates for a decentralised approach to AI adoption, encouraging professionals from various roles to experiment

with AI and discover its most valuable applications through hands-on experience. As AI continues to advance, many acoustic engineers are likely to evolve into digital acoustic innovators, identifying novel ways to integrate AI into their work.

The development of iFieldnotes demonstrates how OpenAI APIs can be used to create a customisable and scalable AI-based software application. To achieve this functionality, acoustic professionals must have some coding knowledge and, at least for now, collaboration with skilled programmers. However, in the AI era, this collaboration has evolved. Unlike traditional software development, which often requires specialised expertise, the AI-related aspects of coding have become more accessible. Acoustic engineers with basic coding knowledge can now take a more active role in the development process, leading to more efficient and collaborative efforts to build AI-driven tools like iFieldnotes. As AI continues to advance, the reliance on professional programmers is expected to diminish.

7   CONCLUSION

Leveraging the power of OpenAI APIs, iFieldnotes—an innovative mobile app for field note-taking—can perform a wide range of tasks. These include converting shorthand notes into complete sentences, generating summaries, interpreting images and floor plans, performing OCR, transcribing voice notes, and translating voice notes into English from multiple languages. These features address the limitations of traditional note-taking methods, which rely on pen and paper or generic mobile apps. While iFieldnotes is still in the testing phase and its full impact has yet to be measured, its development demonstrates the potential of GenAI APIs in achieving both customisation and scalability.

The development of iFieldnotes is rooted in the author's own experience, highlighting how acoustic professionals can take a more active role in software development thanks to the rise of AI. Although collaboration with professional programmers is still necessary, advancements in AI are expected to gradually reduce this dependence. In the future, many acoustic engineers are likely to evolve to digital acoustic innovators, discovering new ways to incorporate AI into their work.

8   REFERENCES

1.   S. Dance and V. Wills, AI in acoustics – An IOA London Branch event, Acoustics Bulletin, Vol. 50 No 4 (July/August 2024).
2.   Wikipedia, GPT-4. https://en.wikipedia.org/wiki/GPT-4
3.   R. Khawaja, Best large language models (LLMs) in 2023, https://datasciencedojo.com/blog/best-large-language-models/# (26 July 2023)
4.   P. Villalobos, A. Ho, J. Sevilla, T. Besiroglu, L. Heim, M. Hobbhahn, https://epochai.org/blog/will-we-run-out-of-data-limits-of-llm-scaling-based-on-human-generated-data (06 June 2024)
5.   R. Rohm, The Future of AI: Capitalising the value of proprietary data, https://www.eu-startups.com/2023/12/the-future-of-ai-capitalising-the-value-of-proprietary-data/ (22 December 2023)
6.   L. Cascio, Why It's The Year Of Proprietary Data, Not AI, https://www.forbes.com/councils/forbesfinancecouncil/2024/04/25/why-its-the-year-of-proprietary-data-not-ai/ (24 April 2024)
7.   N. Bashir, P. Donti, J. Cuff, S. Sroka, M. Ilic, V. Sze, C. Delimitrou, and E. Olivetti, The Climate and Sustainability Implications of Generative AI, An MIT Exploration of Generative AI, https://doi.org/10.21428/e4baedd9.9070dfe7 (27 March 2024)
8.   E. Mollick, Latent Expertise: Everyone is in R&D (20 June 2024)