

Prompt Engineering for Voice AI agents: A Practical Guide to Engaging Conversations

Author: Marc Johnson; Date: 2025-02-14

Introduction

Voice AI agents are increasingly used in customer support and interactive experiences. Crafting effective prompts for these voice agents is crucial to ensure they provide an engaging, informative, and human-like experience. A well-designed prompt guides the AI's behavior – defining its **persona, style, and conversation flow** – so that the voice agent can delight users with a seamless interaction. This report outlines best practices for prompt engineering with Voice AI (VAPI) agents, focusing on AI voice agent/assistant scenarios. Key areas include **structuring prompts for immersion, iterative testing and refinement, handling unpredictable user inputs, setting the right tone, and insights from community and real-world implementations**. Each section provides guidelines backed by VAPI's official documentation and expert discussions.

1. Prompt Structuring

Designing the prompt with a clear structure helps the AI voice agent/assistant deliver an immersive and balanced experience. Experts recommend **breaking the system prompt into distinct sections**, each defining a specific aspect of the AI's behavior¹

. For example, you might include:

- **Identity** – Define the AI's persona and role. For an assistant, this could be something like *"You are an enthusiastic and knowledgeable assistant."* This anchors the AI in a specific character and context.
- **Style** – Set the tone and style guidelines for responses. For instance, instruct the guide to *"be friendly and descriptive, with a conversational tone, but also concise due to the voice interaction"*. This ensures the agent is informative and polite yet avoids long monologues (since users cannot read back voice output).
- **Response Guidelines** – Include formatting or length preferences to balance information and engagement. You might limit each description to a few sentences and encourage the agent to ask a question or check if the user wants more details, rather than delivering a one-sided lecture.² Voice **users appreciate quick, digestible answers**, so **breaking complex information into smaller chunks** keeps the conversation engaging. If the user is curious for more, the agent can offer to dive deeper – for example, *"I can tell you more about this painting's history if you're interested."* (retellai.com) This approach prevents overloading the listener while maintaining an interactive feel.

¹ VAPI.ai. "Voice AI Prompting Guide | Vapi." Accessed March 14, 2025. <https://docs.vapi.ai/prompting-guide>.

² Retell.ai. "How to Write Voice Bot Scripts That Engage Users | Retell AI," 2025.

<https://www.retellai.com/blog/how-to-write-voice-bot-prompt>. **1. Overloading Users with Long or Complex Responses: Users interact with voice AI bots expecting quick, straightforward answers. Long-winded or overly detailed responses can overwhelm them, leading to frustration and disengagement.**

- **Tasks/Flow** – Outline the flow step-by-step. A prompt can **explicitly enumerate the sequence of actions** the guide should take (greeting, introducing the location, highlighting key facts, inviting questions, moving to the next point, etc.)³ docs.vapi.ai
Structuring the **conversation flow as an ordered list** helps the AI stay organized. It can also **include conditional branches** (e.g., **if the user asks about something specific, address it before returning** to the main goal). By having a clear script or plan, the AI provides a cohesive guidance while remaining flexible to user interactions.

Engaging and Immersive Experience: To create an immersive UX, **prompts should encourage the AI to use vivid, sensory language and personalize the content to the user**. In a study of a GPT-3.5-based tour guide system, developers **carefully crafted prompts to elicit the user's interests and tailor the tour accordingly**, enabling the AI to provide detailed, contextually relevant information about each point of interest.⁴ These prompts “foster an engaging and informative dialogue” by making the AI ask about the visitor's preferences and then weaving those into the tour narrative. For example, if a user shows interest in architecture during a visit, the AI might give extra details on a building's design or construction story, creating a more personalized experience. This **combination of factual information with a narrative style keeps users immersed**.

Clarity through Markdown Formatting: Interestingly, even though the final output is spoken, **formatting the prompt using Markdown or structured lists can improve clarity for the model**.⁵ The VAPI documentation suggests using **Markdown headings, bullet points, and indentation** in your prompt definition to **clearly separate sections and instructions**. A structured prompt is easier for both humans and the AI model to understand. For instance, formatting the system prompt with labeled sections like `[Role]`, `[Context]`, `[Instructions]` on new lines (as shown in VAPI's examples) helps ensure the AI “knows” which part is which. Using **bullet points or numbered lists** for complex instructions can also make them more digestible for the AI.⁶ This way, the AI voice agent/assistant's behavior is well-organized: it knows its role, the style to adopt, how to format responses, and the sequential steps of the experience. Overall, a well-structured prompt – with an **explicit persona, conversational style rules, and a logical flow** – lays the foundation for an engaging guided experience.

2. Iterative Development & Testing

Creating an effective AI voice agent/assistant prompt is an iterative process. You will rarely get the perfect prompt on the first try; instead, developers **refine prompts over multiple rounds** based on testing and user feedback.⁷ VAPI's official guide notes that as your use-case grows in complexity, you must “experiment and iterate on your prompt to improve your success rate” – where *success*

³ VAPI.ai. “Voice AI Prompting Guide | Vapi.” Accessed March 14, 2025. <https://docs.vapi.ai/prompting-guide>.

⁴ Ababneh, Alaa. “GPT-3.5 as a Personalized Tour Guide in Biosphere Reserves.” *Advances in Machine Learning & Artificial Intelligence* 6, no. 1 (February 21, 2025): 1–9. <https://doi.org/10.33140/AMLAI.06.01.06>. “These prompts are carefully crafted to **elicit relevant information about the user's interests, preferences, and background**, enabling the system to personalize the tour experience.”

⁵ VAPI.ai. “Voice AI Prompting Guide | Vapi.” Accessed March 14, 2025. <https://docs.vapi.ai/prompting-guide>. **Using [Markdown](#) formatting in prompts is beneficial because it helps structure your content, making it clearer and more engaging for readers or AI models to understand.**

⁶ Moulton, Luke. “How to Write Prompts for AI Voice Agents - Best Practices,” November 19, 2024. <https://leadagents.ai/prompt-engineering/prompt-ai-voice-agents/>.

⁷ VAPI.ai. “Voice AI Prompting Guide | Vapi.” Accessed March 14, 2025. <https://docs.vapi.ai/prompting-guide>.

rate means the percentage of user requests the AI can handle fully without human intervention. In practice, this means after deploying the assistant, you should monitor how conversations go: *Does the AI handle all user questions well? Are there points where it gets confused or users get disengaged? Use these insights to tweak the prompt.*

Design → Test → Refine: A structured development cycle is key to prompt engineering. Start by *designing* a clear initial prompt that encodes the role and behavior. Next, *test* it with real or simulated user interactions – for example, have the AI conduct a sample and observe its responses. Then *refine* the prompt instructions based on what you learned, and repeat. This could involve adding new instructions for edge cases, rephrasing parts the AI misunderstood, or adjusting the level of detail. Following this loop of testing and refinement leads to continuous improvement. “Regular testing and refinement of prompts are crucial for enhancing AI performance,” as one best-practices guide puts it.⁸ **By reviewing real user conversations, you can identify where the assistant might be talking too much, not enough, or misinterpreting queries.** Each iteration of prompt tuning should bring the AI’s behavior closer to the desired experience.

Learning from User Feedback – Case Study: A real-world example of iterative improvement comes from a voice-guided tour research project using GPT-3.5. The system was designed to adapt and improve over time by analyzing user interactions. Specifically, it would leverage user feedback and past conversation logs to **generate improved prompts** for subsequent calls. In other words, the AI was not static – after encountering misunderstandings or new user questions, the developers updated the prompt (and underlying logic) so that next time the AI would handle those cases better. Thanks to this feedback loop, the GPT-3.5 tour guide could correct its mistakes without requiring complete retraining. This case illustrates the power of **prompt iteration**: by tuning instructions in response to actual user behavior, the AI’s tours became more accurate and satisfying over time.

A/B Testing Prompt Styles: One advantage of AI voice agents is that they are programmatic, allowing for rapid A/B testing of different prompt approaches.⁹ Rather than guessing which style of prompt will engage users best, you can experimentally compare options. For instance, you might **A/B test a conversational, flexible prompt vs. a very strict, scripted prompt** to see which yields better user engagement. One recent article suggests not to “set and forget” your voice agent, but to continually optimize it by testing such variations. In a tour guide context, you could prepare **two versions** of the AI’s prompt: one where the agent has some *personality and freedom to improvise descriptions*, and *another where the wording of each tour segment is fixed*. By split-testing these with different user groups, you can measure which version leads to more user questions, higher satisfaction ratings, or better retention of information. Industry experts note that *locking the AI to a rigid script can lead to robotic-sounding experience*, whereas *giving the AI some flexibility* (e.g. *example phrases* rather than exact scripts) *often produces more natural interactions*. **Other elements to A/B test** include the agent’s voice tone or speed – for example, does an upbeat enthusiastic voice keep users more engaged than a calm neutral voice? Does speaking a bit slower improve comprehension? Experimenting methodically with these factors and gathering data will help

⁸ Moulton, Luke. “How to Write Prompts for AI Voice Agents - Best Practices,” November 19, 2024. <https://leadagents.ai/prompt-engineering/prompt-ai-voice-agents/>.

⁹ “5 Must Run A/B Tests for Your AI Voice Agent.” Accessed March 14, 2025. <https://www.regal.ai/blog/a-b-tests-for-ai-voice-agent>. “By A/B testing different elements of your AI agent, you can enhance engagement, improve user experience, and fine-tune responses. If you’re just “setting and forgetting” your AI Voice Agent, you’re not taking advantage of this transformational new technology.” 1. Voice Gender, 2. Voice Emotion (neutral vs expressive), 3. Conversational Style: Exact Language vs. AI-Driven Flexibility, 4. Personality vs. Neutral Professionalism, 5. Voice speed.

pinpoint the most engaging style for your AI voice agent/assistant. The key is to be **data-driven and unafraid to experiment**: continuously measure user responses, and iterate on the prompt and voice settings to optimize the experience.

3. Handling Dynamic User Inputs

Real users won't always follow the script – they may ask unexpected questions, go on tangents, or respond in unanticipated ways. A robust AI voice agent/assistant needs to handle dynamic user inputs gracefully to keep the experience on track. **Best practices for managing unpredictable queries** start with anticipating them in your prompt design. **Always include fallback and error-handling instructions** in the system prompt so the AI knows what to do when confronted with an input it doesn't understand or that falls outside its knowledge. The VAPI documentation explicitly advises adding a dedicated **[Error Handling]** section (or similar) in your prompt with guidance for such situations. For example, you might instruct: *“If the user’s response or question is unclear or off-topic, politely ask for clarification or gently steer back to the main goal/theme.”* A sample error-handling prompt might read: *“If the customer’s response is unclear, ask a clarifying question. If you encounter any issue, apologize and politely ask them to repeat.”*¹⁰ By scripting these behaviors, your AI will be less likely to get stuck or respond inappropriately when surprised by user input.

Preparing for Edge Cases: Community experts emphasize planning for edge cases to avoid dead-end conversations.¹¹ Not all callers will respond as expected – one might give a very vague answer to a guide’s question, another might bring up a completely unrelated topic (“By the way, what’s the weather there?”). If such cases aren’t planned for, the bot could fall silent or give a generic error, breaking the immersion. To prevent this, **anticipate less-common inputs and questions** and decide how the AI should handle them. For instance, if the user asks something unrelated to the content, the AI could be prompted to **briefly address it and then transition back**. An AI assistant might say, *“Hmm, I’m not sure about the weather, but I do know a lot about this museum. Was there something about the exhibit you wanted to explore?”* – **politely refocusing the conversation**. In a voice bot example for a receptionist, the script explicitly included a rule: *“If the caller asks for information unrelated to the clinic, politely redirect or provide limited assistance”*. Applying this to an AI assistant: **if a user asks a question outside the tour’s scope, the AI can offer a quick helpful snippet if possible, or gently steer back to the tour**. The goal is to **never leave the user feeling ignored or at a loss** – **fallback responses** should **always offer a path forward, whether that’s answering the question, clarifying what was said, or returning to the last known topic**.

Maintaining Context: Handling dynamic inputs also means the AI should remember context and resume the visit smoothly after addressing a user’s query. A common challenge is when an AI loses track of earlier details – for example, forgetting a user’s name or that they already visited a particular exhibit – which can happen in multi-turn conversations. Best practices to maintain context include **using the platform’s memory features or variables to store important information** (like the current stop or the user’s preferences) and **explicitly instructing the AI to acknowledge previous info**. For example, if a user asked a question about a painting’s artist and the AI answered, the next part

¹⁰ VAPI.ai. “Voice AI Prompting Guide | Vapi.” Accessed March 14, 2025. <https://docs.vapi.ai/prompting-guide>.

¹¹ Retell.ai. “How to Write Voice Bot Scripts That Engage Users | Retell AI,” 2025. <https://www.retellai.com/blog/how-to-write-voice-bot-prompt>. 1. Break down complex information into smaller, digestible chunks. 2. Provide **concise answers and offer follow-up questions for additional details** if necessary.

of the prompt could remind the AI to tie back into the narrative: *“After addressing the question, resume the tour from where it left off (e.g., ‘Now, continuing with our tour...’).”* This prevents the guide from either looping back unnecessarily or skipping ahead. VAPI’s system allows maintaining state across the call, and your prompt can leverage that to ensure continuity. As a developer, plan and test these transitions: does the AI properly handle a user interrupting the AI voice with a question, and then continue correctly? If not, refine the prompt or use VAPI’s tools (like functions or state variables) to fix it.

Error Recovery: Despite preparation, the AI might still get tripped up by something unforeseen. In these cases, how it recovers is crucial. The **prompt should encourage the agent to fail gracefully** – meaning **if it doesn’t understand, it should say something polite and helpful rather than just going silent or repeating nonsense**. A good fallback utterance might be, *“I’m sorry, I didn’t catch that. Could you please rephrase your question?”* or *“Let me try explaining that differently.”* Including a few of these polite recovery phrases in the prompt gives the AI options to handle errors without derailing. In more severe cases (e.g., the user’s requests are entirely off-topic or the system encounters an internal error), the prompt can direct the AI to either offer an alternative (like, *“Perhaps I can answer that at the end. For now, shall we continue?”*) or even escalate to a human guide if available. Some implementations route to a human operator if the AI cannot handle the request – if that’s an option, the prompt can instruct the AI how to hand off gracefully. The overarching principle is **robustness**: the AI voice agent/assistant should be prepared for many contingencies so that even when the unexpected happens, the experience remains as seamless and helpful as possible.

4. Tone & Voice Modulation

Setting the right conversational tone is especially important, as it directly impacts user engagement. The tone should match the context and the expectations of the audience. For example, a historical museum tour might use a warm, storytelling tone to captivate the listener, whereas a support tour of a software product might stick to a friendly but professional voice. Developers should **clearly define the bot’s personality and voice in the prompt to ensure consistency**. This can involve training writers and developers on the desired tone or brand guidelines, as one best-practice article suggests. In the prompt itself, you might include a line like *“Tone: Enthusiastic, friendly, and informative, as if you truly love the subject you’re describing.”* By giving the AI this kind of guidance, you align its style with the immersive experience you want to create.

Persona and Character: Introducing a distinct personality can make the AI voice agent/assistant more relatable and engaging. For instance, you might design the agent to have the persona of a knowledgeable individual with a bit of humor, or an expert curator who’s passionate about the exhibits. Many experts advocate for **adding such personality to voice agents** – it humanizes the interaction and can delight users. A comparative insight from industry: A personality-driven voice AI can create a more engaging experience than a purely neutral, robotic one. For example, when asked about an artifact, a neutral AI might simply state facts, whereas an AI assistant **with personality** might add *“Isn’t that amazing?”* or a fun anecdote, sparking more interest. However, it’s important to match the tone to the context; an overly cheerful or joke-cracking bot might not be suitable for a solemn art gallery or a brand that requires formality. The best approach is to define the desired tone in the prompt (under a Style or Persona section) and ensure it aligns with the brand or environment. **Consistency builds trust** – users will feel more at ease if the AI’s voice and phrasing **stay true to its character** throughout.

Controlling Verbosity: In real time, the AI voice agent/assistant should modulate how much it says to keep users engaged without overwhelming them. This is partly handled by prompt instructions (as discussed, e.g., *limiting response length in the prompt*), but it's worth highlighting as a tone/voice consideration too. The **agent should be conversational, not lecturing**. One guideline is to **keep sentences short and clear, and pause to allow the user to interact**. The prompt can explicitly tell the AI to be concise: for example, *"Keep your explanations to 2-3 sentences at a time, unless the user asks for more detail."* In VAPI's travel assistant example, one style rule was *"Be concise, as you are currently operating as a Voice Conversation."* This reminds us that unlike text chat, where a user can scroll up to read long answers, in voice the information must be absorbed on the fly. If the AI dumps a huge paragraph of text verbally, the user may miss parts or lose interest. In fact, voice AI users expect the bot to get to the point quickly.¹² To manage this in real-time, an AI assistant might **break a description into segments, and use a natural stopping point to check in** (e.g., *"Shall I continue?"* or it could listen for a brief moment for any interruption before proceeding). Additionally, the AI should **avoid overly complex vocabulary or jargon** unless appropriate for the audience – part of tone management is **making sure the language level fits the user**. A **conversational tone with simple, engaging language generally works best** for broad audiences, whereas a specialist (say a technical factory tour for engineers) might allow more technical terms. **These choices should be encoded in the prompt's style guidelines.**

Voice Modulation and Expressiveness: Beyond words, how the AI's voice sounds (if using TTS – text-to-speech) can be adjusted to improve user experience. Many voice AI platforms (including VAPI) let you **choose voices** and even **control parameters like speaking speed and emotional tone**. It's wise to **experiment with voice settings** to find what feels most natural for your use case. For example, you can **A/B test a slightly slower speaking rate** versus a normal rate to see which one users find clearer and more pleasant. An AI voice agent might slow down when explaining something complex or speed up slightly during transitions – some systems can do this dynamically. Another element is *expressiveness*: a monotone delivery will bore users, so consider using a more expressive TTS voice that has variation in pitch and emphasis. One industry piece notes that an *"expressive, energetic voice can make AI interactions feel more human-like and solicit more conversational engagement from customers,"* whereas a too-flat voice might lead users to disengage or respond with one-word answers. In the context of a tour, an **expressive voice conveys enthusiasm** for the content, which can be infectious. However, keep in mind the downside: highly expressive synthesized voices sometimes mispronounce words or inflect strangely, so **test the voice thoroughly with your script**. Striking the right balance in voice temperature (neutral vs lively) and speed can significantly enhance the overall ambiance.

Emotional Intelligence and Adaptability: Tone modulation isn't just about sounding happy or sad; it's also about responding appropriately to the user's mood and the moment. If a user seems confused or asks for clarification, the AI should **adopt a reassuring tone**, not a jovial one. In customer service scenarios, **showing empathy** is key. For an AI assistant, empathy might mean **acknowledging the user's reactions** ("I hear your excitement!" or "No worries if that was a lot to take in, I can repeat anything if needed."). The prompt can equip the AI with some of these adaptive phrases. In VAPI's tips, there's a concept of **"Emotional Prompting"** – using language that sets a certain mood or emotion for the AI's response. For example, you could prompt the AI with, *"You*

¹² Retell.ai. "How to Write Voice Bot Scripts That Engage Users | Retell AI," 2025.

<https://www.retellai.com/blog/how-to-write-voice-bot-prompt>. **Users interact with voice AI bots expecting quick, straightforward answers.** Long-winded or overly detailed responses can overwhelm them, leading to frustration and disengagement.

“speak with enthusiasm and warmth, as if you’re truly delighted to share these stories.” Such instructions shape the AI’s tone across the interaction. Additionally, to sound more natural and human-like, some developers even **introduce natural speech patterns** intentionally. This might include the occasional **interjection** like “well,” **slight pauses**, or even **mild stutters** if it fits the persona (though for a confident AI assistant, we might use fewer stammers). These elements, when used sparingly, can prevent the AI’s speech from sounding too robotic. For instance, adding **an ellipsis** “...” in a prompt response can cue a brief pause in speech, making it sound like the guide is thinking – a touch of realism. Overall, controlling tone and voice is about **making the AI sound engaging but also authentic**. By tuning the persona, verbosity, and vocal qualities of the AI, developers ensure the experience holds users’ attention and delivers information in a pleasant, conversational manner.

5. Community Insights & Real-World Applications

The rapidly evolving community around voice AI and VAPI has been actively sharing insights, success stories, and challenges. Across forums, blogs, and professional networks, there’s a consensus on many best practices – and a focus on overcoming common hurdles in building voice agents.

Echoes from the VAPI Community: VAPI’s own documentation and community discussions reinforce the importance of the principles discussed above. In summary, creating a successful voice AI agent requires **well-structured prompts**, a good dose of **personality and natural speech patterns**, **graceful error handling**, and ongoing quality improvements. Developers on platforms like Reddit and LinkedIn often discuss how **prompt engineering is the key** to voice agent performance. Many report that getting the AI to behave correctly is initially tricky, but by applying the structured approach (defining role, style, etc.) and iterating, they can achieve reliable and human-like interactions. The availability of **knowledge bases and retrieval tools** (which VAPI supports) is another topic in the community: to handle a wide range of user questions, developers integrate FAQs or documents so that the AI can pull factual answers when needed, rather than winging it. This helps especially in specific applications, where users might ask detailed questions – connecting the voice agent to a curated knowledge source can improve accuracy. VAPI’s Knowledge Base feature has its own best practices (like keeping answers concise and relevant) which align with what we’ve covered.

Common Challenges and Solutions: Despite best efforts, developers frequently face certain challenges when deploying AI voice agent/assistants. One is the issue of AI **hallucinations** or misinformation – the AI might occasionally make up an answer that sounds plausible but is incorrect. This is a known problem with generative models. Community experts advise placing **guardrails** in the prompt to mitigate this.¹³ For example, include a guideline like *“Do not invent information. If you don’t know the answer, say you’re not sure or redirect the question.”* Reinforcing this in the prompt and **testing the agent on factual queries** can reduce the chance of it going off-script. Another challenge is the AI’s unpredictability; because the model may generate varied responses, ensuring consistent behavior is hard. A solution here is to **make prompts as specific as possible** (to limit randomness) and **use a lot of example dialogues** during development to see how

¹³ Moulton, Luke. “How to Write Prompts for AI Voice Agents - Best Practices,” November 19, 2024. <https://leadagents.ai/prompt-engineering/prompt-ai-voice-agents/>.

the AI responds in different situations. Some advanced users even **adjust model parameters** (like **temperature settings**) to **find a balance between creativity and reliability** in responses.¹⁴

Maintaining context in long interactions is another issue that comes up in real-world use. As mentioned, forgetting context leads to the bot asking redundant questions or contradicting itself. Developers share techniques like **using session memory** or repeating key context in the prompt for each response cycle. One practical tip is to have the AI occasionally paraphrase or summarize the current state affairs (e.g., *“So far, we’ve seen the main hall and the ancient artifacts exhibit...”*) which both helps the user keep track and reinforces the AI’s memory of where it is in the flow.

A/B Testing in Practice: In community discussions, developers have found surprising results through A/B tests. For instance, some shared that a more *“enthusiastic”* voice actually improved user feedback scores for a travel tour guide bot, while for a corporate office tour, a slightly more *“professional and calm”* voice tested better. This shows the value of **tailoring the tone to your audience**. Others have **experimented with script styles and discovered that giving the AI a bit of freedom** (instead of reading a fixed script) made the conversation feel more natural and increased user engagement – aligning with what Regal’s study suggested about not over-scripting. On the other hand, **completely open-ended prompts tended to let the AI drift off-topic**, so the sweet spot was a guided yet flexible prompt. These real-world experiments underscore that there’s no one-size-fits-all prompt; you have to fine-tune it for your specific context.

Expert Opinions: Industry experts often share their insights on platforms like LinkedIn and AI forums. A common piece of advice is to **treat your AI agent almost like a new team member** you’re training. You need to **coach it, set expectations** (through the prompt), **monitor its performance, and give it feedback** (by refining the prompt or providing corrected examples). As one expert summarized: effective prompting involves **understanding the AI’s limitations, giving clear instructions, and continuously testing and refining** the prompts. Embracing this iterative, hands-on approach will eventually “unlock the full potential” of your AI voice agent. In the domain of customer support, this means a smoother, more interactive caller experience that can scale to many users. Companies pioneering voice AI (like those using VAPI) have reported significant improvements in customer satisfaction and efficiency by deploying these best practices in their voice agents.

In conclusion, the community’s insights and real deployments highlight that building a great AI voice agent/assistant is an ongoing journey. By **structuring prompts clearly, refining them through real-world feedback, handling user surprises gracefully, and fine-tuning tone and personality**, developers can create AI voices that are both informative and a delight to converse with. The combination of official guidance and grassroots knowledge-sharing in this space provides a rich toolkit for anyone looking to leverage voice AI in customer support or interactive tour applications. By following these best practices and learning from others’ experiences, you’ll be well on your way to delivering engaging, immersive AI Voice experiences that captivate your users.

¹⁴ “5 Must Run A/B Tests for Your AI Voice Agent.” Accessed March 14, 2025. <https://www.regal.ai/blog/a-b-tests-for-ai-voice-agent>. **Tuning your AI’s emotional expression – controlled via the temperature setting of your TTS voice – may significantly affect how users engage.**

6. An email response by Dr. Alaa Ababnaas providing extra examples, tips and context for prompting effective AI Voice agent.

I suggest to you some draft processes I do in my bot to minimize dilution in your AI model; you should develop interactive AI. When I was developing my bot, I focused on this prompt response. I suggest many scenarios; AI encounters **a question it can't answer**, it can respond with, "I'm not sure about these specific details, but you can visit the reserve's website for more information." Similarly, **if the query is too vague**, the AI can suggest related topics by saying: "I couldn't find this information. Would you like to learn more about the animals in the biosphere reserve?" To ensure clarity, the AI should also **prompt users to rephrase** their questions, such as: "Can you identify which aspect of the biosphere reserve interests you? Is it the wildlife, the trails, or the conservation efforts?" After submitting a response, it's recommended to **ask users how helpful they found the information with a question** like "Was this information helpful?" (Yes/No), as this helps guide future responses. Additionally, **allowing users to leave comments about specific answers** helps gather more information about the issues, aiding improvement. Keeping records of queries that result in **unsatisfactory responses** allows for identifying common issues and facilitating improved AI training data. Regularly **analyzing user feedback** can reveal patterns in AI performance, while **redirecting users to other sources**, such as visitor centers, local guides, or trusted websites, **ensures they still receive the information they need if the AI fails to provide it accurately**. Regular model updates based on accumulated user data, error logs, and feedback will improve response accuracy and relevance, while **training the AI on diverse inputs and scenarios to better handle a wide range of queries**.

To effectively guide interactions, a range of personalized prompts can be used. AI can **greet** users by saying, "Welcome to the biosphere reserve! What interests you most? Wildlife, hiking trails, or conservation efforts?" **or asking**, "Do you have any specific questions about the reserve or its activities that I can help with?" Information-gathering questions aim to tailor recommendations to the user's interests, such as, "To give you the best suggestions, can you tell me about your interests or what you'd like to experience here?" or "Are you looking for family-friendly activities, adventure hikes, or educational tours?" **Pique curiosity is crucial**, such as, "Did you know that this particular animal/plant is only found in this biosphere reserve? Would you like to learn more about it?" and "What season are you planning to visit? I can provide insights into the best experiences during that season." **Follow-up questions can enhance user engagement**, with recommendations such as, "If you're interested in hiking, would you like me to recommend the best trails based on difficulty and scenery?" or "Would you like to learn more about the reserve's conservation efforts?" Finally, **fun facts and trivia add a fun touch to interactions**, such as: "Here's an interesting fact: The biosphere reserve is large and home to several unique species. Would you like to learn more about them?" or "Many visitors love the activities. Would you like to try them?"

All of these suggestions must be interactive and prompt when developing your bot.

Regards

Alaa