

Flow AI Documentation

As of 9/7/2023 release (This document was updated with various clarifications on 9/30/2023).

We're thrilled to give you access to our new AI features. This is just the beginning, and your insights are invaluable to us, so as you explore, we encourage you to share your thoughts on how we can elevate your experience.


Here's a list of initial features:

- Get context from any selected dot
- Engage in a visual conversation with AI
- Extract context from PDF documents
- Command the swarm with your voice

We'll delve deeper into each of these shortly. We're committed to continuously enhancing our app, so stay tuned for more exciting updates in the coming weeks and months.


Enabling AI and the AI buttons

First, enable AI for any particular Flow in the Settings, accessed through the menu in the upper right of the screen:



Within the Settings, find the “Enable AI features” checkbox.


Image not found or type unknown



You may also want to check the “Show AI conversation” in order to see the history of the conversation.


To actually start using AI, you create a new button here, using the “+” icon. You can also create a new button by clicking the "Create/New Object" button and then selecting "Button".

Image not found or type unknown



A new button editor pane will appear:

Image not found or type unknown




Click the checkbox to “Enable AI.” This will show the options to control the AI and put the button on the “heads-up display” at the bottom of the screen. Other button configuration options are not available for AI buttons.

There are three options for the Action for response:

- Ask AI
- Ask PDF
- Edit swarm


Image not found or type unknown



Please note that additional options will be added over time.

With perhaps additional context for your prompt (see the "GPT prompt" section below), your editor screen may look like this, with a new button in the center of the 3D canvas.


Image not found or type unknown



Also, in this screenshot, you can see the partial AI response above the "Ask AI" button.

Here is a view on the Desktop, after you have set up a couple of buttons with labels to "Ask AI" and "Change", here is what it might look like when Presenting the Flow:

Image not found or type unknown



Here is what the button area of the screen looks like in a VR or AR headset:

[image.png](#) image not found or type unknown

GPT prompt:

Using LLM prompts is a whole sub-specialty these days.

Prompt variables are used to include user input into the AI prompt. The variables are:

\$spokenInput - The user's verbal input

\$visiblePopupData - the text within the currently visible swarm popup. The popup must be a 3D popup. Information in Overlay popups will not be included in the GPT prompt.

You don't need to get carried away with prompting, but you might find some inspiration from these examples:

Generic example:

\$spokenInput| Tell me more about \$visiblePopupData| Do not respond with apologies as an AI large language model. Keep the answers to three short concise paragraphs.

Example from a historical biographical Flow:

\$visiblePopupData, \$spokenInput.


Respond in the voice of the person referenced, including linguistic characteristics of this person's writings.

Don't apologize about being an AI model, but if there are no extant writings by this person, end your response with a short phrase such as "I've done my best to speak in my voice with limited material."

AI response window

On the desktop, the AI response is shown in a popup window that is scrollable and movable.


Image not found or type unknown



Each time a response is provided, the window will open. It can also be opened by clicking on the short response directly above the AI buttons.

In an XR headset, the window appears a bit to the left of the user:

Image not found or type unknown



Cancelling an on-going spoken command

Whenever a `$spokenInput` is present on a visible button, then an accompanying “x” button will appear beside the button to cancel or reset listening. This is great when you misspeak and want to start the query over, or when AI makes a mistake and does not correctly recognize your speech.

Edit Swarm

Edit Swarm is used to manipulate the swarm, as shown below where the user request to filter the data.

The following current prompts are supported, but remember that simple variations on these phrases work as well because AI is not expecting an exact verbal match:

- Filter to a category: “Filter to Vietnam”
- Filter to a time frame or a numeric range: “Filter from 2016 to 2023”
- Change swarm color: “Make the color a gradient from green to red”
- Set swarm column on a specific axis for the scatterplot: “Set the depth axis to population”
- Toggle axis on and off: “Hide the depth axis”
- Change dot size: “Make the dots bigger”
- “Remove filter”

The mechanism currently works with only one filter, so each filter request will replace the prior AI-generated filter with a new filter. In the current early stage of AI integration we do have the ability to generate complex filters (no AND or OR filtering). In the future, we’ll need a more complex UI visible to the user to help the user to indicate when they want to replace or add a filter.

As of this release, error reporting is not extensive, so the AI may fail silently if it becomes confused. It is also possible for AI to create a filter that shows no dots. In this case, ask AI to “remove filter.”

Ask PDF

Flow has the ability to upload a PDF document (just one for now) and use that as the source material for ChatGPT to query against.

Upload the PDF in the “Manage Data” popup.

Suggested prompt for PDF documents:

Only respond with information from this document. If the request is outside the scope of this document, then respond with several suggested questions that are within the scope of the document. \$spokenInput

As we work with PDF documents, we expect to learn how best to structure PDF documents for optimal access by ChatGPT. Early results indicate that written text documents are well-understood by ChatGPT, and that pure database or spreadsheet extracts of columnar information are not as well understood. As we gain additional experience in passing PDF content along to ChatGPT, we will provide additional guidance on this topic.

LLM model

Currently, this release supports OpenAI’s ChatGPT 3.5 and 4.0 models. If using 4.0, you also have the ability to add system context. We can support additional models upon request for enterprise customers.

Conversational AI experimentation

We have been experimenting with Conversational AI.

There is now a checkbox in the Settings in the UI that enables the “conversation” with the AI to be captured and displayed. The conversation is displayed as a swarm, offset from the origin by a significant distance. Currently, this swarm is not very configurable, but we look forward to receiving feedback, and we will be extending this functionality soon.

Revision #10

Created 6 September 2023 17:36:09 by Jason Marsh

Updated 6 October 2023 04:19:24 by Bill