

AI Assistants for a Nonprofit Community Center,
Grounded in the Organization's Own Program Material

Nine assistants, and a separate decision about where each one may look

The assistants that answer families are bounded to what the center has published and show their sources. The one that helps staff draft is deliberately not bounded at all.

9

Assistants created, five of which carried a question

5 of 6

Live assistants answering only from the center's own material

6 of 6

Live assistants with anti-hallucination enabled

799

Questions asked, raw count, across 14 months

THE CHALLENGE

Three jobs that have nothing in common, and one small team holding all of them

A parent the night before the first day of camp, working out what to pack and what time drop-off actually starts. A camp director with next season's program to write and a stack of last season's material to write it from. A development officer with a grant application due and the same organizational background to restate for the fourth time this year.

Two of those people need an answer that is already true. The third needs something that does not exist yet. A nonprofit this size has one team covering all three, and one thing that cannot be allowed to happen: an assistant telling a parent something about a children's program that the center never said. The question was never whether to ground the assistants. It was whether every assistant had to be grounded the same way.

WHAT GOT BUILT

A camp assistant for families.

Bounded to the center's own material, with citations shown on the answer, and a written persona rather than a default one.

A grant drafting assistant.

Bounded to the center's own material, so the organizational background it reuses is the center's own.

A curriculum assistant for program staff.

The one assistant deliberately allowed to draw on general knowledge, because drafting is the job.

A successor camp assistant.

Built for the following season rather than by editing the one that had already done its year.

THE REQUIREMENTS

Five things a nonprofit answering to families would not compromise on

A community center holds a particular kind of trust. The people asking are the people whose children are in the building, and an answer that sounds official is treated as official.

- 1 Nothing said in the center's name that the center did not publish.
- 2 Silence is better than a guess.
- 3 Show the source when a family is asking.
- 4 A different rule for making things.
- 5 Quick enough to build one for a single season.

SWORN TO SECRECY

A nonprofit community center in the western United States with roughly two hundred staff, running year-round programs and a summer camp. Shared anonymously. Every figure comes directly from the center's own deployment.

AT A GLANCE

SECTOR

Nonprofit community programs. A community center in the western United States running year-round programs and a summer camp.

DEPLOYMENT

Nine assistants created over fourteen months. Six live, five of which have carried a question.

ACCESS

Shareable assistants for families and for staff. No agentic actions enabled anywhere.

HOW IT WAS BUILT

Built one at a time, each for a job that had come up

Start from what already exists.

Program documents were uploaded as they were. Two assistants were pointed at the center's own website instead.

Set the answer source per assistant.

The single decision this study is about, made separately each time rather than once for the whole deployment.

Write the persona by hand.

Every assistant that carried a question has a persona the center wrote, and every one of those was revised after it went live.

Let a finished one rest.

Nothing is retired and nothing is rebuilt. An assistant goes quiet between seasons and is there when the season comes back.

THE ANSWER SOURCE, THE CAPABILITY THAT DECIDED IT

One setting, decided separately for every assistant

Each assistant is set either to answer only from what the center has uploaded and published, or to draw on the model's general knowledge. Five of the six live assistants are held to the center's own material. The sixth, which drafts program curriculum for staff, is the one exception, because a tool that can only repeat last season cannot help write the next one. Where that choice is made once for a whole deployment, a nonprofit has to pick between a safe assistant and a useful one. Here it is a per-assistant decision, so a parent and a program director can be served by the same platform under different rules.

THE RECORD

Five assistants, and what each one was allowed to look at

WHAT THE NUMBERS SAY

This is a seasonal organization and the usage admits it. Questions crowd into the weeks around a camp season and thin out to nothing in between. The figure to read is the curriculum assistant's: 316 questions across its first season, then quiet, then 82 more when the next season came around. Nobody rebuilt it, retrained it or migrated it. It sat there for the better part of a year and was picked up again by the people writing the following summer.

WHY CUSTOMGPT.AI CAME OUT ON TOP

The grounding decision is per assistant.

One deployment can hold a strictly bounded assistant and an open one, without a second deployment or a second tool.

Refusing came as a setting.

Anti-hallucination is on for every live assistant here, which is a choice made once rather than prompt-engineered each time.

Building one is not a project.

Four assistants were created and never used, and that is a sign of a low bar rather than a failure.

WHAT THIS MEANS FOR YOU

Decide what each assistant may look at, one assistant at a time

Separate answering from drafting. They want opposite things from a model, and a platform that makes you choose once for everything will have you compromising on both.

Count the ones nobody used. If building an assistant is a project, the abandoned ones are waste. If it is an afternoon, they are how you found the two that worked.

LOOKING AHEAD

The parent asking the night before gets an answer drawn from what the center actually published, with the source shown. The camp director writing next season has a drafting tool that is allowed to suggest something new. The development officer stops restating the same organizational background from memory. Three jobs, one platform, three different sets of rules.

ABOUT CUSTOMGPT.AI

CustomGPT.ai is a no-code platform for building AI agents that answer only from an organization's own content: documents, websites and knowledge bases. Every answer is grounded in verified sources with citations, and the platform is SOC 2 Type 2 and GDPR compliant. Organizations use it for customer support, internal knowledge search, website chat and more, across help centers, apps and the API.