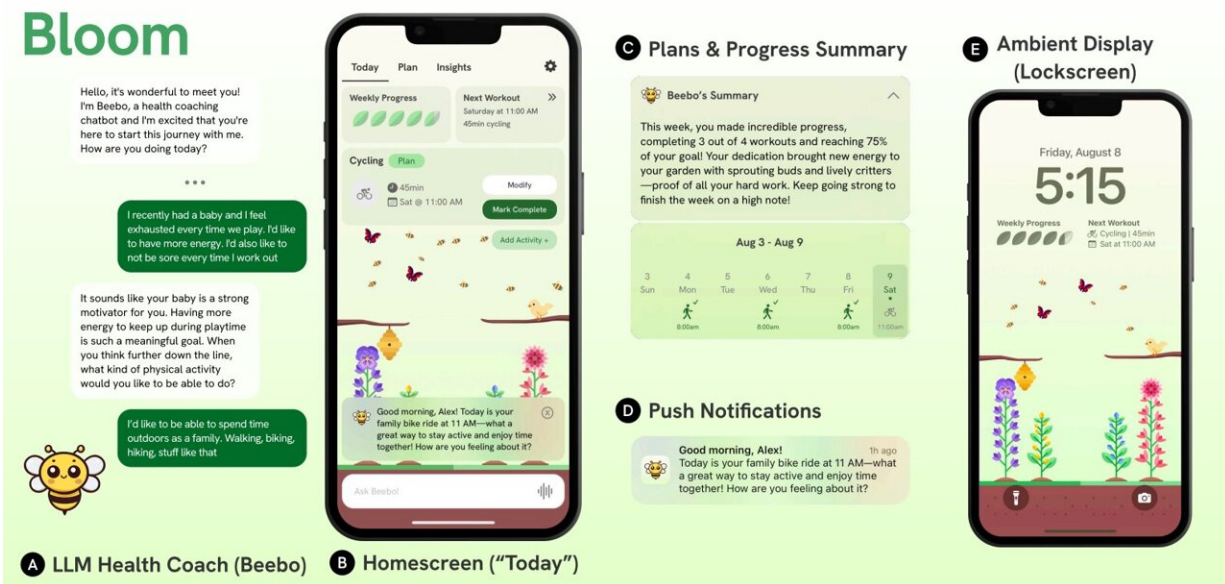


This AI coach is quietly rewriting how people think about exercise

April 28 2026, by Katharine Miller



Bloom integrates an LLM-based physical activity coaching chatbot (Beebo) with established behavior change interactions. Credit: *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (2026). DOI: 10.1145/3772318.3790506

With so many health apps and wellness coaching gadgets to choose from, it might seem as if the market is already saturated. But commercial health apps often treat coaching as an optimization problem, says Matthew Jörke, a graduate student in computer science at Stanford University. Given enough data collected from an app, the chatbot coach

will prescribe a workout plan and expect the user to follow it, he says. "That's not how human behavior works."

According to health coaches Jörke has interviewed, "being facilitative, empowering, and nonprescriptive works best for people, especially those who are just starting out with physical activity."

To create an LLM-based health coaching app that adheres to that expert advice, Jörke and his colleagues at Stanford designed Bloom, which combines an aesthetically pleasing, state-of-the-art fitness app with a large language model (LLM) coaching agent called Beebo, symbolized by a little bee icon in the app.

The agent asks questions about a user's goals and what they've already tried, and offers praise and empathy. With users' permission, Beebo also proposes and calendars weekly workout plans that users can easily modify through a chat interface.

"The app privileges human autonomy and behavior change in a way that supports users to identify their own goals and achieve them," says Emma Brunskill, associate professor of computer science at Stanford and senior co-author on a [paper](#) about Bloom, which won a Best Paper award at the 2026 ACM CHI Conference on Human Factors in Computing Systems ([CHI 2026](#)), held in Barcelona, April 13–17. The open source code is [posted](#) on GitHub.

In that study, the researchers found that users who had access to the LLM coaching feature increased their physical activity by about the same amount as users given a version of the app that lacked all LLM features, but their mindsets had noticeably shifted. "People start recognizing that physical activity is something they can do, that it's good for them, and it makes them feel better," Jörke says.



Credit: Unsplash/CC0 Public Domain

As more and more people turn to LLMs for advice in other arenas besides health, Jörke and his colleagues' work suggests a different way forward. "The design strategies we learned through building this health coach could transfer to building any other type of AI advisor that helps you figure out what matters to you."

Creating the LLM agent's coaching style

Researchers with a human-computer interaction (HCI) background, such as Jörke and his collaborators, including senior author James Landay, the

Denning Director of the Stanford Institute for Human-Centered AI, start from a place of wanting to deeply understand what works or doesn't work for real people.

So, the team's first step in building an LLM coach was to interview 12 health experts and 10 potential health app users about what would make for a helpful LLM coach. Key insights that became design principles for Beebo: Take a facilitative approach so that clients take ownership of their health; tailor advice to the user's unique personal situation; and use a nonjudgmental, supportive tone.

The team also designed Beebo to align with Stanford Active Choices, a counseling program developed by Abby C. King and colleagues at Stanford Medical School that has proven effective at helping people of all ages and fitness levels become more active. A key first step in that program is the initial interview, in which the counselor gathers relevant information about the client's long-term goals, past experiences, barriers, and resources.

Bloom implements the onboarding conversation using an approach called motivational interviewing—a style of interviewing that helps people tap into their own motivations—to gain an understanding of the client's goals and limitations.

But transferring that approach to an AI coach wasn't simply a matter of prompting a vanilla LLM to administer the initial interview. "That didn't work," Jörke says. "It couldn't stay on topic, was very prone to giving unsolicited advice, and made assumptions about who you are and what you want."

Rather than gathering a large dataset of onboarding conversations to fine-tune the LLM, the team tried prompting the LLM in various ways to see if they could get the model to cover the desired topics using the

motivational interviewing style.

Their preferred approach relies on two prompt chains. The first, a dialog state prompt chain, ensures that the LLM thoroughly covers each Stanford Active Choices topic before moving to the next. And the second enforces motivational interviewing principles through a two-step process: One AI agent selects an appropriate conversational strategy (such as asking an open-ended question or making a simple reflection) from 11 options, while a second agent crafts the actual response using that strategy.

"The LLM needs that extra scaffolding to use motivational interviewing the way we want it to," Jörke says. Beebo also includes safety filters to protect against harmful content such as negative feedback or inappropriate body image comments.

Changing mindsets

In a four-week field study, roughly half of 54 subjects were asked to use a version of the Bloom app that removed all LLM features, while the treatment group used the app with Beebo.

The control condition's version of the app is as good as most health apps out there, Jörke says. It allows people to set goals, plan their activity for the week, receive reminders, and see progress toward their goals as flowers grow on a lovely home screen ambient display.

The treatment group's version of Bloom has all of that but more. First, there's the onboarding interview that gives Beebo a deep understanding of the user. After that, Beebo offers to design a workout plan together with the user, adds it to the user's calendar, and sets up reminders and alerts. Beebo also checks in by chat to see how things are going, provide encouragement, and create new plans as appropriate. And those plans are

flexible: In Beebo's chat function, users can ask for modifications if something comes up, like their schedule changes, it's raining, or they aren't feeling well.

"We specifically wanted to study what a nonprescriptive LLM would add on top of a high-quality health app," Jörke says. "And remarkably, it wasn't a whole lot more exercise. It was a change in mindset."

In surveys, Beebo users expressed increased satisfaction with their activity levels and an increased belief that their activity is sufficient and beneficial, which typically boosts people's desire to keep moving. Some said that Beebo helped them realize that even gardening or moving around the kitchen can count as exercise steps. Others reported a sense that Beebo's support was personal, treating daily messages as invitations for a morning chat.

"Slowly, over time, Beebo helped them see that they are capable, that activity is beneficial and not that hard," Jörke says.

According to Landay, some features in the baseline application, such as an ambient display, have been shown to be highly impactful in his team's previous research studies but are still rarely found in most commercial apps. "By including them in the baseline, we set the comparison up to be quite hard for the LLM version of Bloom to beat. So, the positive results we found in the treatment condition are quite promising," he says.

Lessons for LLM coaching

Jörke sees potential in adapting Bloom to new contexts such as sleep, nutrition, or life coaching, but admits that would require researchers to develop expert-informed dialog state prompt chains to keep the onboarding conversation on topic.

"We need to develop a simpler way to prompt or fine-tune a model so that it does a better job of balancing the need to strategically gather relevant information with the urge to deliver advice and recommendations," he says.

Jörke also wonders if Beebo's adorableness played a role in the app's appeal. Researchers are concerned about how much to [anthropomorphize](#) these sorts of agents because of risks that people will form inappropriate attachments to them. But users didn't go to Beebo for medical advice or general conversation, or describe it as their new best friend, Jörke notes.

Perhaps that is because Beebo was a little critter rather than a humanoid and had a distinct role as a coach. "The metaphors that we use for these chatbots and how we communicate their capabilities to people might actually make their interactions better overall," he says.

While acknowledging that Bloom is still a research prototype, the team plans to make it available to the public in the near future. Based on comments by people who learn about Bloom, the demand seems to be there.

"I think this project is tapping into a deep need," Brunskill says.

"Everyone knows that exercise is important, and people want to be more active, but they are also balancing an enormous number of demands on their time. Bloom tries to address this in a way that centers human agency."

More information: Matthew Jörke et al, Bloom: Designing for LLM-Augmented Behavior Change Interactions, *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (2026). [DOI: 10.1145/3772318.3790506](https://doi.org/10.1145/3772318.3790506).

Provided by Stanford University

Citation: This AI coach is quietly rewriting how people think about exercise (2026, April 28)
retrieved 28 April 2026 from
<https://sciencex.com/news/2026-04-ai-quietly-rewriting-people.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.