Mental Models

AI-powered systems can adapt over time. Prepare users for change—and help them understand how to train the system. This chapter covers:

Which aspects of AI should we explain to our users?

How should we introduce AI to the user initially—and thereafter?

What are the pros and cons of introducing our AI as human-like?

Want to drive discussions, speed iteration, and avoid pitfalls? Use the worksheet.

What's new when working with AI

A mental model is a person's understanding of how something works and how their actions affect it. People form mental models for everything they interact with, including products, places, and people. Mental models help set expectations for what a product can and can't do and what kind of value people can expect to get from it. Mental models can also serve as bridges between experiences. For example, if you know how to steer a bicycle, you know something about how to steer a motorcycle.

However, users' mental models may not always match what a product can actually do. Mismatched mental models can lead to unmet expectations, frustration, misuse, and product abandonment. Often, product creators unintentionally set incorrect mental models for users by not considering the early user experience with a product or not fully explaining how the product works. Key considerations:

① **Set expectations for adaptation**. AI allows for more systems to adapt, optimize, and personalize for users, and probability-based user experiences have become more common over time. Building on the familiarity of existing mental models can help users feel comfortable.

② **Onboard in stages**. When introducing users to an AI-powered product, explain what it can do, what it can't do, how it may change, and how to improve it.

③ **Plan for co-learning**. People will give feedback to AI products, which will adjust the models and change how people interact with them — which will change the machine learning models further. Users' mental models will similarly change over time.

④ **Account for user expectations of human-like interaction**. People are more likely to have unachievable expectations for products that they assume have human-like capabilities. It's important to communicate the algorithmic nature and limits of these products to set realistic user expectations and avoid unintended deception.

① Set expectations for adaptation

Most products are highly static. You can bet that the hammer you buy today will be the same hammer tomorrow. There have also been responsive products for quite a while — products that can adapt how they respond based on user input over time. These systems keep track of whether or not an output was useful, and update how they respond going forward. Using a digital example, many streaming media services will adjust their recommendations based on your interaction with previous ones. This can in turn create an expectation that other products will adapt based on user interactions.

With AI becoming more prevalent in products and experiences, we can expect to see more experiences that change in response to users. One of the biggest opportunities for creating effective mental models of AI products is to build on existing models, while teaching users the dynamic relationship between their input and product output.

Identify existing mental models

Start by thinking about how people currently solve the problem that your product will use AI to address. That existing solution will very likely inform their initial mental model for your product. For example, if people currently label their email messages manually the assumption could be that an AI-powered email product — somehow — follows the same process that the user does. In this case, that would be: read the email, think about the meaning, context, and importance, and attach a label accordingly. Therefore, it may surprise users that the ML model could use other signals—time of day the message was sent or length of the email—to determine the label.

Key concept

To understand the context for the user's relationship to your AI product, work through some of the questions below:

- What is the user trying to do?

- What mental models might they carry over to your product?

- What is the step-by-step process that novice users currently use to accomplish the task?

- How uniform is this process between different users?

Apply the concepts from this section in Exercise 1 <u>in the worksheet</u>

② Onboard in stages

Onboarding is the process of helping a new user or customer get to know a product or service. The onboarding experience begins before users purchase or download your product or even visit your website, and continues indefinitely. As with any product, it's important to consider the different stages of introducing your AI and how mental models form and change along the way.

Introduce and set expectations for AI

After identifying your users' existing mental models, imagine how information your user received before their first interaction with the product — including marketing messages, ads, or manuals — has shaped their expectations. Collaborate closely with your marketing team to develop appropriate and consistent messaging.

Many products set users up for disappointment by promising that "AI magic" will help them accomplish their tasks. This kind of messaging can establish mental models that overestimate what the product can actually do. Though product developers may intend to shield users from a product's complexity, hiding how it works can set users up for confusion and broken trust. It's a tricky balance to strike between explaining specific product capabilities, which can become overly technical, intimidating, and boring, and providing a high-level mental model of your AI-powered product.

RUN

4.5 ★★★★⯪ (1,348,231)

RUN is a running app that adapts to your fitness levels and designs personalized workouts to help you improve your running.

| | RUN | Download |

RUN

2.1 ★★☆☆☆ (4,651)

RUN is the only intelligent running app that uses sophisticated deep neural net machine learning to make your run smarter because we believe in ML driven workouts.

| | RUN | Download |

Aim for

Emphasize how the app will benefit users. Learn more

Avoid

Don't emphasize the underlying technology.

Here are some messaging guidelines for setting the right expectations for your product:

- Be up front about what your product can and can't do the first time the user interacts with it, ideally in your marketing messages.

- Offer examples of how it works that clarify the value of the product.

- Let people know up front that it may need their feedback to improve over time.

- Communicate why people should continue to provide feedback, focusing on the value to them.
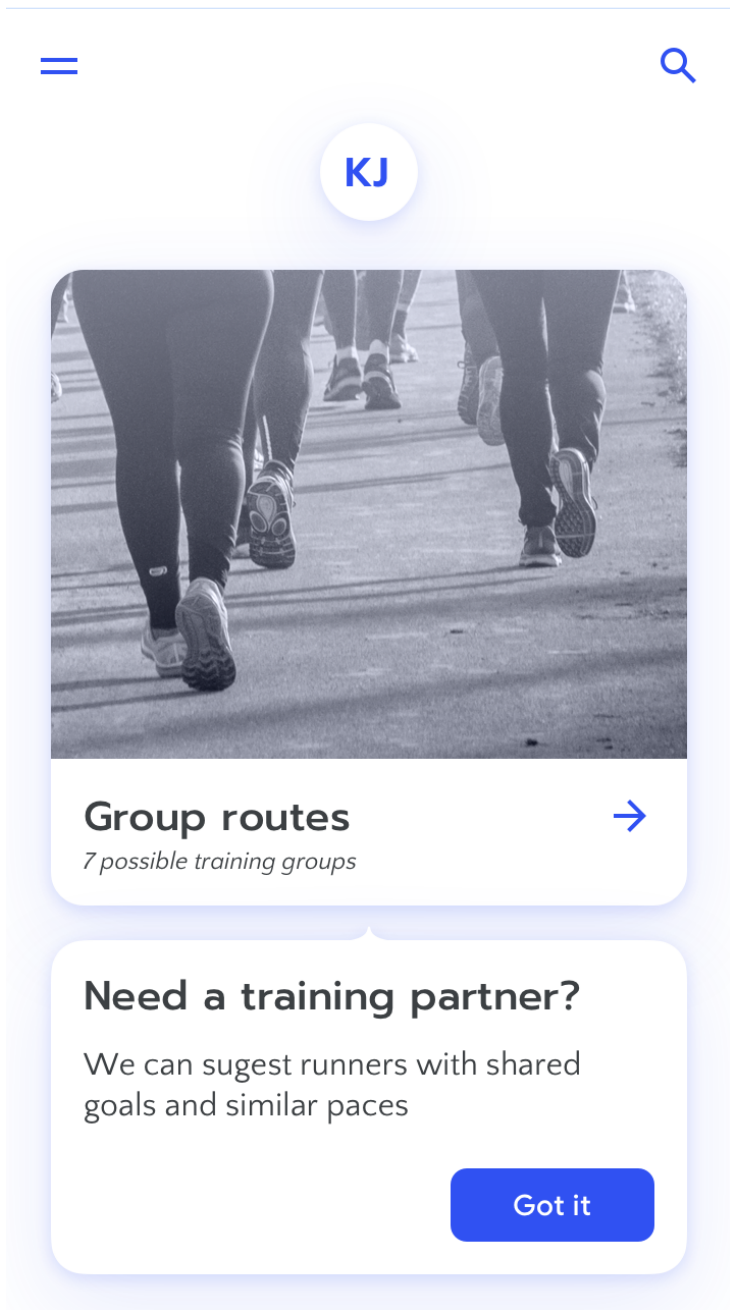
Explain the benefit, not the technology

Often as product creators, we're fascinated by the underlying technologies that make products and experiences possible. This is especially true if we've cracked a hard technical problem for the first time. However, make sure to evaluate which details users need to build a good mental model. If they're interested in understanding the underlying technology of your product, you can always provide more detail with tooltips and progressive disclosure . If you do talk about the AI, focus on how it specifically makes part of the experience better or delivers new value.

See more about explaining AI at the right level of detail in the <u>Explainability + Trust</u> chapter.
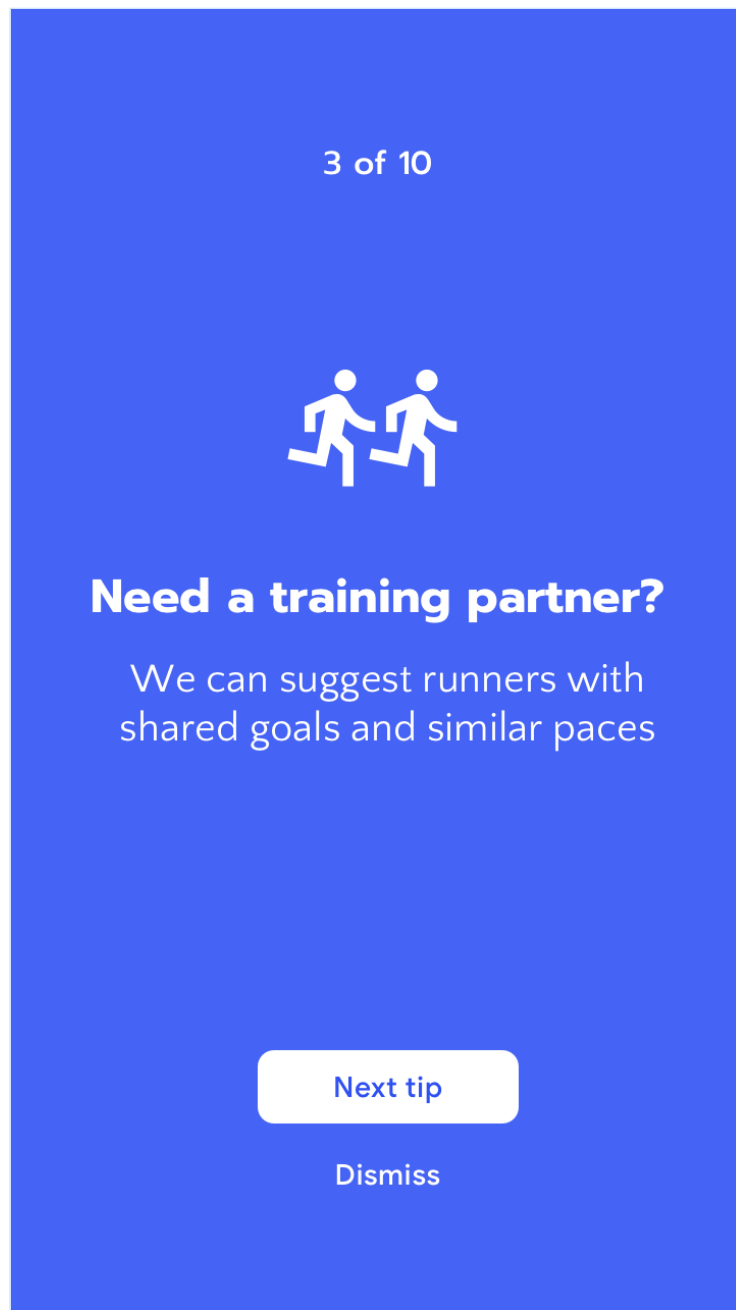
Only introduce new features when needed

As users explore the product, use relevant and actionable "inboarding" messages to help them along. Try to avoid introducing new features when users are busy doing something unrelated. This is especially important if you're updating an existing product with new AI features that change the function or user experience. People learn better when short, explicit information appears right when they need it.

KJ

Group routes

→

7 possible training groups

Need a training partner?

We can sugest runners with shared goals and similar paces

Got it

Need a training partner?

We can suggest runners with shared goals and similar paces

Next tip

Dismiss

Aim for

Introduce an AI-driven feature at the moment it is relevant to the user. Learn more

Avoid

Don't introduce AI-driven features as part of a long introductory list of product features.

Design for experimentation

Many people learn best by tinkering with a new experience. People sometimes skip onboarding steps because they're eager to start using the system, and reading even a few screens feels like it's in the way. By keeping onboarding short, you'll let them get right to it. Suggest a low-risk or reversible action they can try right away — users are often curious about how an AI-powered feature or product will behave, so encourage that with a small, contained initial experimentation experience. For example, applying photo filters is easy to test out and undo with a tap.

One caveat is that a user's willingness or ability to spend time experimenting depends on their goal in using your product. For example, an average consumer who purchases a new smart speaker might enjoy spending time experimenting with different commands and questions. In contrast, a busy enterprise user might regard testing commands and functions as just one more chore in a busy day.

Regardless of their goals, point users towards where they can quickly understand and benefit from your product. Otherwise, they may find the boundaries of the system by experimenting in ways it isn't prepared to respond to. This can lead to errors, failure states, and potentially erosion of trust in your product.

Learn more about helping users get back on track in the <u>Errors + Graceful Failure</u> chapter.

Your first run

We'll average your first few runs together before we start making recommendations

Learn more   **Let's go**

Welcome

To get started, lace up and head out on your first intelligent run

Learn more   **Let's go**
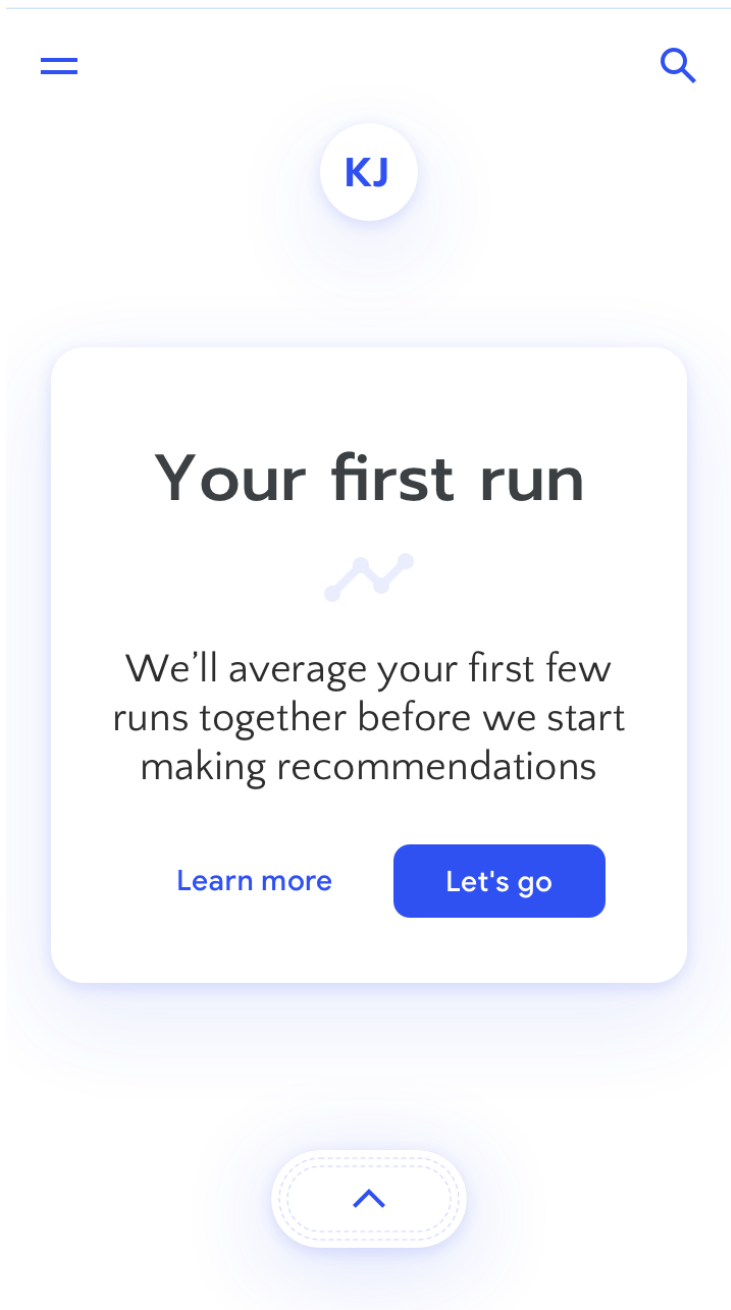
Aim for

Encourage experimentation and reassure users that experimenting won't dictate their future experiences. Learn more
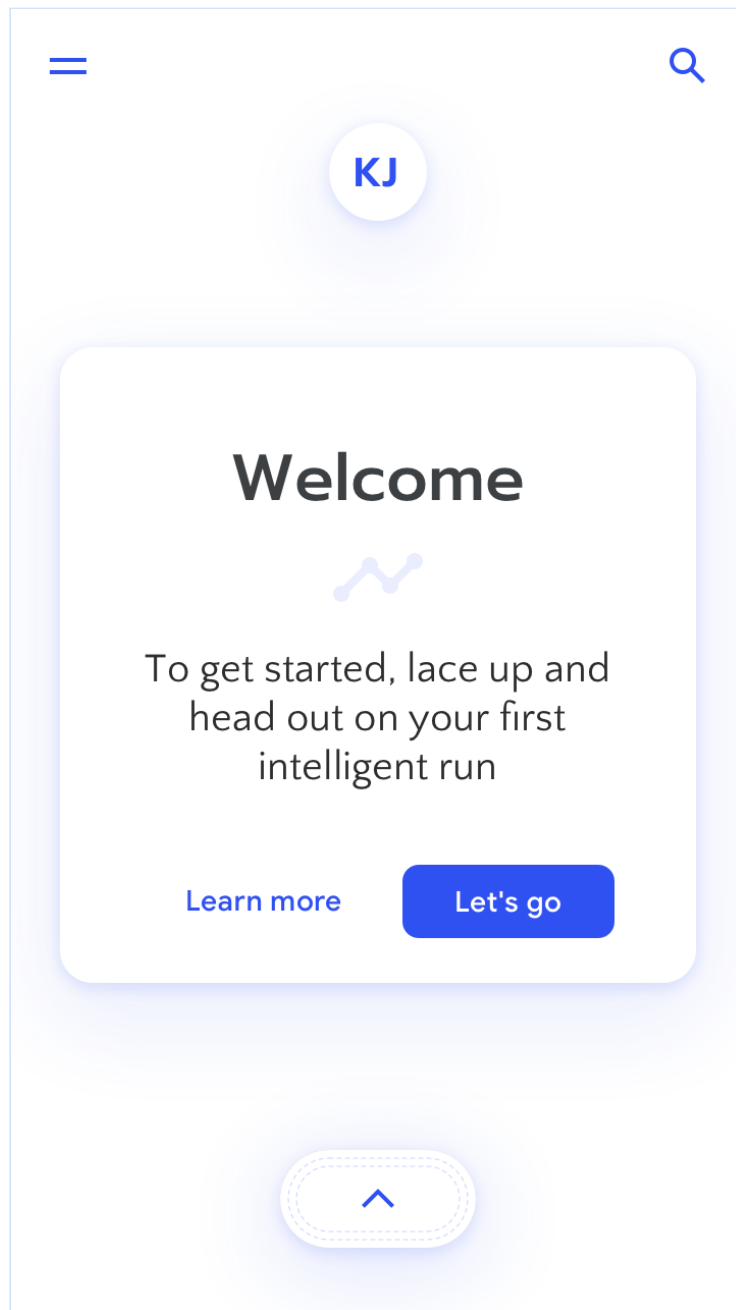
Avoid

Don't assume users want the AI to start learning from the first use.

Key concept

Onboarding is all about setting up the interaction relationship between the user and your product. Here's a simple messaging framework to get you started:

This is *{ your product or feature }*,

and it'll help you by *{ core benefits }*.

Right now, it's not able to *{ primary limitations of AI }*.

Over time, it'll change to become more relevant to you.

You can help it get better by *{ user actions to teach the system }*.

Apply the concepts from this section in Exercise 2 <u>in the worksheet</u>

③ Plan for co-learning

Because AI-powered products can adapt and get better over time, the user experience can change. Users need to be prepared for that, and adjust their mental model as necessary.

Connect feedback with personalization

In onboarding, let users know how the feedback they provide helps the AI personalize their experience. You can tie this to the user benefit with phrasing like "you can improve your experience by giving feedback on the suggestions you receive", and letting them know where and how to do so.

There are two ways to collect feedback:

Implicit feedback is when people's actions while using the product help improve the AI over time. There should be a place in your product where users can see which signals are being used to what end, and this should be disclosed in your terms of service.

For example, choosing to listen to the next suggested song in a music app confirms the model's prediction that the song is relevant to you. That fact should be part of the definition, user benefit, and terms of how the app works.

Explicit feedback is when people intentionally give feedback to improve an AI model, like picking categories of music they're interested in. This kind of feedback can help the user feel more in control of the product. If you can, explain precisely what impact the feedback will have on your AI, and when it will take effect.
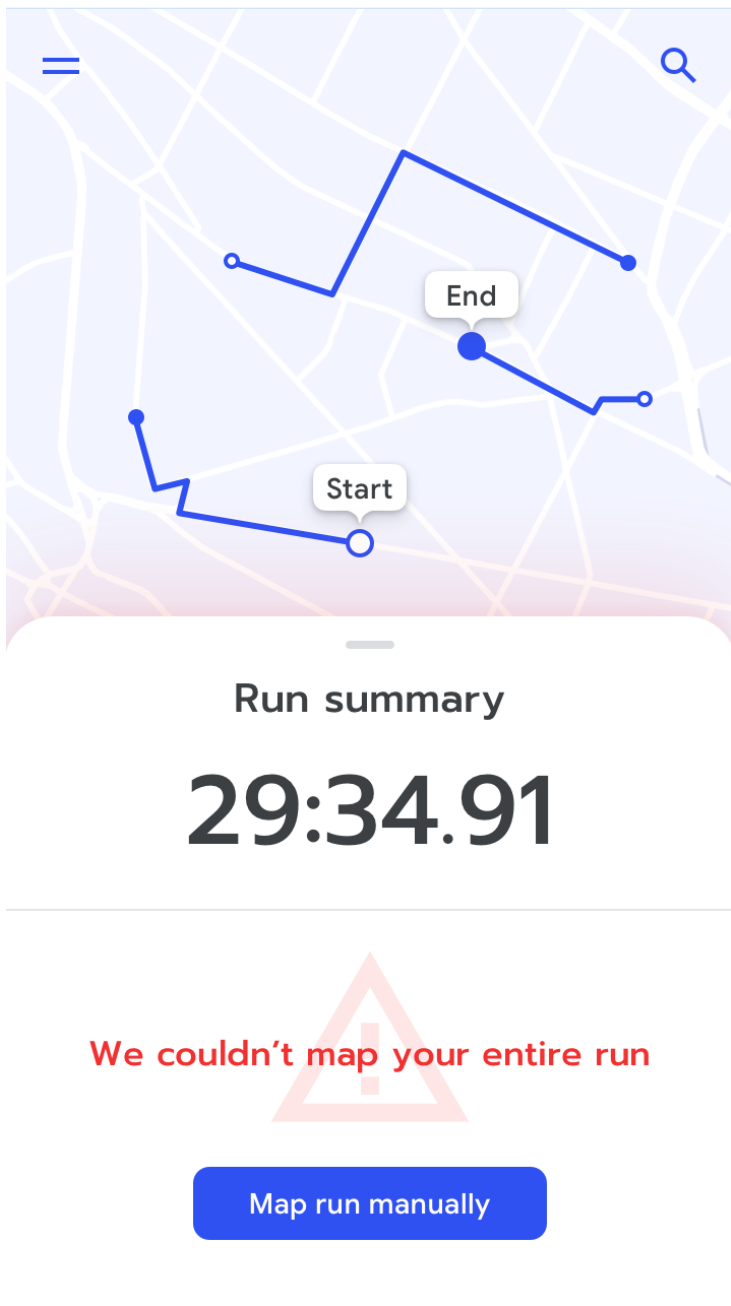
When the system collects feedback, explain how continually teaching the system benefits the user. Be clear about what information will help your AI learn and how it will improve the product output. See examples and much more detail in the Feedback + Control chapter.

Fail gracefully

The first time the system fails to meet expectations, the user will likely be disappointed. However, if the mental model includes the idea that the system learns over time, and learns better with the right input, then failure, especially the first failure the user encounters, becomes an opportunity to establish the feedback relationship. Once this relationship is set, users will see each failure not only as more forgivable, but also something that they can help fix. This buy-in can help cement the mental model of co-learning.

When your system isn't certain, or can't complete a request, make sure there's a default user experience that doesn't rely on AI. That way, the burden of educating your AI doesn't stop users from getting things done. When your product fails gracefully, it doesn't get in the user's way, and they see feedback as a way to make their objectives even easier over time, while still being able to use your product right now.
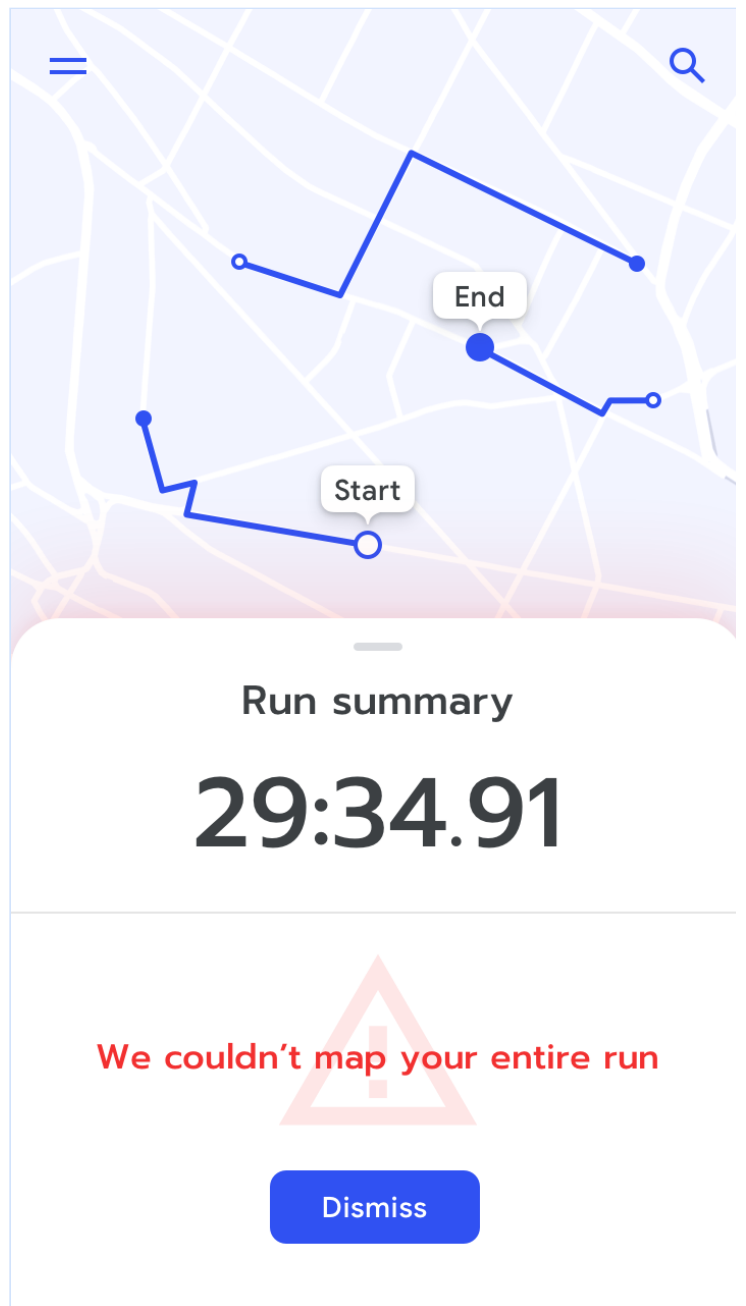
See examples and much more detail in the Errors + Graceful Failure chapter.

Run summary

29:34.91

We couldn't map your entire run

Map run manually

Run summary

29:34.91

We couldn't map your entire run

Dismiss

Aim for

Let users know an error occurred and give them a way to complete the task manually. Learn more

Avoid

Don't create dead-ends when an AI feature fails. These miss the opportunity to shape the user's mental model about the system's limitations.

Remind, reinforce, and adjust

Sometimes products become part of a user's everyday routine, so their mental models get formed and reinforced by ongoing use. But some products are only meant to be used occasionally. For these products, mental models might erode over time, so it's helpful to consider ways to reinforce them, or to remind users of the basics.

You can also help strengthen mental models across AI products by maintaining consistent messaging about the user benefits of improving AI with feedback. Over time, users may adopt a common mental model that recognizes AI solutions and their strengths and weaknesses, becoming more comfortable with what they'll get and how they can shape their experience. There are a few things you can do to increase the odds of that happening:

Keep track of user needs

Monitor how the product is being used. Reviewing your product logs can show you behavior or use trends that point to user confusion or frustration. This can help you determine when you might need to help users rebuild or adjust their mental models. If your product is only meant to be used for a short time or to achieve a specific goal, that will determine how frequently mental models should be reinforced or updated.

See more about metrics in the User Needs + Defining Success chapter.

Adapt to the evolving user journey

If how a feature works changes or improves significantly — enough that a user would notice — consider whether your users need "re-boarding" to the new experience. Re-boarding is also useful when adding new features, or if your system starts using new or different data for an existing feature.

If the system is simple and the mental model is clear and memorable, it's possible that only a little reinforcement is required. A quick user study with people who have used the product previously, but not in the last month or so, could reveal what kind of nudge, if any, might be most helpful.

④ Account for user expectations of human-like interaction

A number of products have launched in recent years that are designed to be anthropomorphic or human-like, such as Cortana, Alexa, Google Assistant, or Siri. This choice has advantages and disadvantages that should be weighed carefully. It's true that people tend to reflexively infer human characteristics from voice interfaces, and some interactions, such as conversational interfaces, are inherently human-like. However, if the algorithmic nature and limits of these products are not explicitly communicated, they can set expectations that are unrealistic and eventually lead to user disappointment, or even unintended deception.

When users confuse an AI with a human being, they can sometimes disclose more information than they would otherwise, or rely on the system more than they should, among other issues. Therefore, disclosing the algorithm-powered nature of these kinds of interfaces is a critical onboarding step. Specifically, your messages should make it extremely clear that the product is not a human, in a way that's accessible to all users regardless of age, technical literacy, education level or physical ability.

This topic is the subject of ongoing research, and these considerations are just a first step. Expect more on this topic in future editions of the Guidebook.
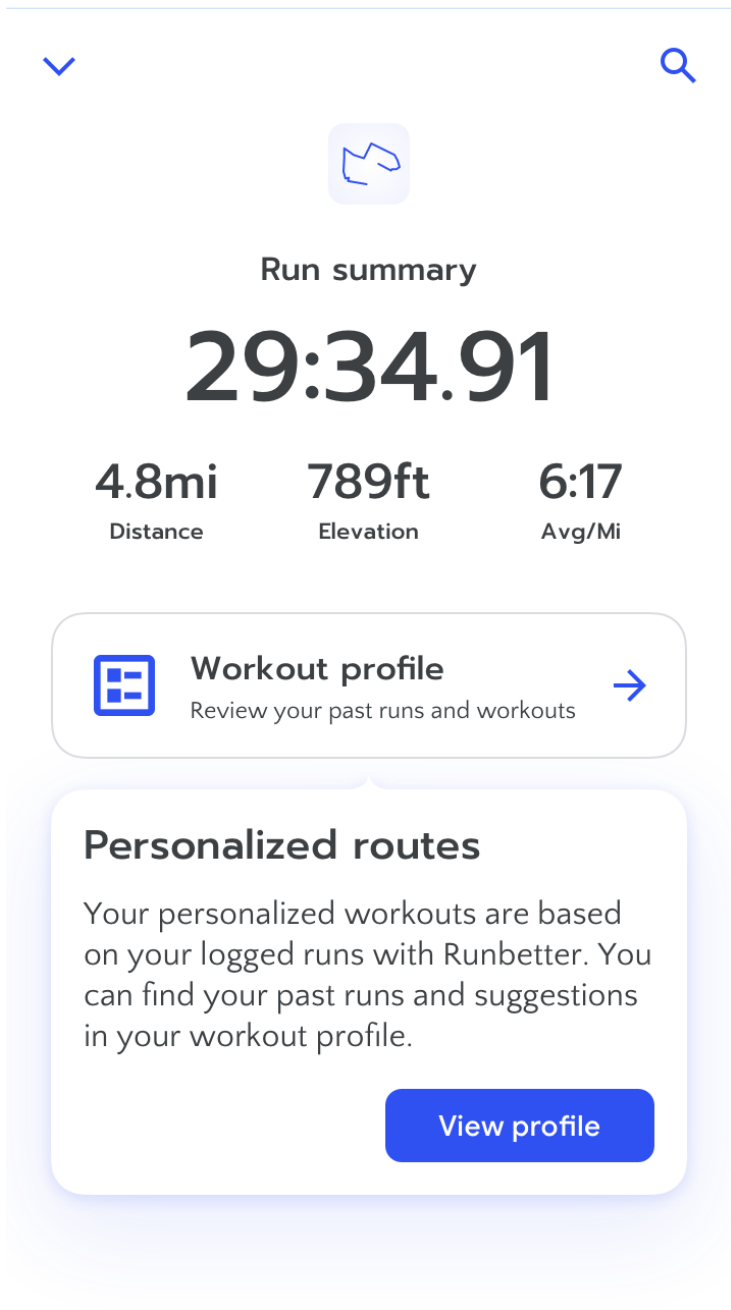
Clearly communicate AI limits and capabilities

People may struggle to form an accurate or useful mental model of an anthropomorphized AI-powered product because the way these systems execute tasks is inherently different than the way a person would. On the surface, AI-powered functionality might seem similar to the manual method but the mental model might not map precisely.

For example, "automatic photo tagging" sounds like the tool is tagging photos the same way a person would, just "automatically". If someone hasn't used this tool before, their mental model of this process is by default, a human one. The app may be able to find and tag all the pictures of a particular friend, just as a person would, but miss some that show her from the back. This is an unexpected break: of course most people can recognize their friends from multiple angles, but this tool can't.

A disconnect like this doesn't necessarily mean that the product itself or the mental model is broken. Not all humans have the same abilities — sight and hearing are not a given — and the product is still doing an incredible task that humans can't do: scanning thousands of photos, identifying the subjects, and labeling them. The key is to communicate the system's limits and capabilities in a way that doesn't create or support expectations of super-human abilities. AI doesn't do anything the same way people do, so though this model is convenient, it's pretty fragile.

Often the idea of a generalized "helper AI" is easier to grasp and more inviting for users, but the risk of mistrust is high when the system's limits aren't clear. When users can't accurately map the system's abilities, they may over-trust the system at the wrong times, or miss out on the greatest value-add of all: better ways to do a task they take for granted. Choose the level of humanization based on how well your AI's capabilities match the user's perceptions of what a human can do.

Run summary

29:34.91

4.8mi	789ft	6:17
Distance	Elevation	Avg/Mi

Workout profile
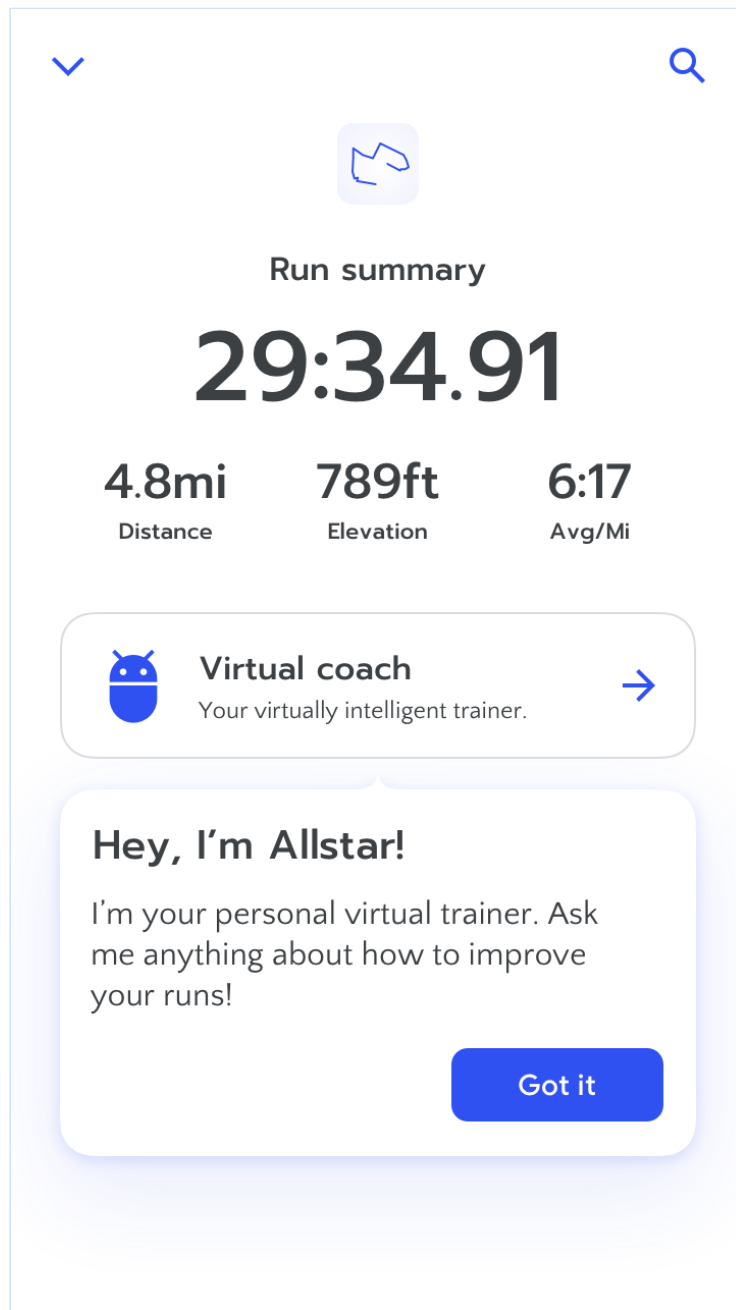Review your past runs and workouts →

Personalized routes

Your personalized workouts are based on your logged runs with Runbetter. You can find your past runs and suggestions in your workout profile.

View profile

Run summary

29:34.91

4.8mi	789ft	6:17
Distance	Elevation	Avg/Mi

Virtual coach
Your virtually intelligent trainer. →

Hey, I'm Allstar!

I'm your personal virtual trainer. Ask me anything about how to improve your runs!

Got it

Aim for

Describe AI features in terms of helping people improve while setting the right expectations for what the AI can do. Learn more
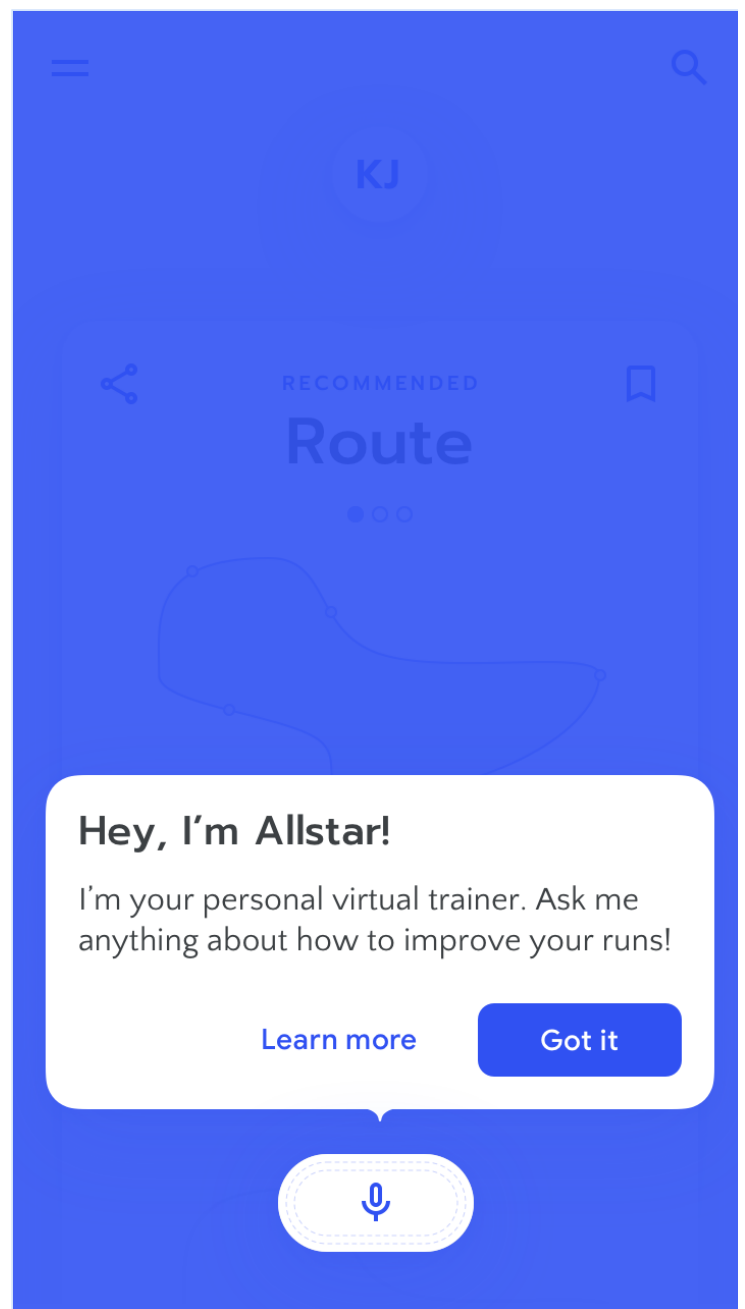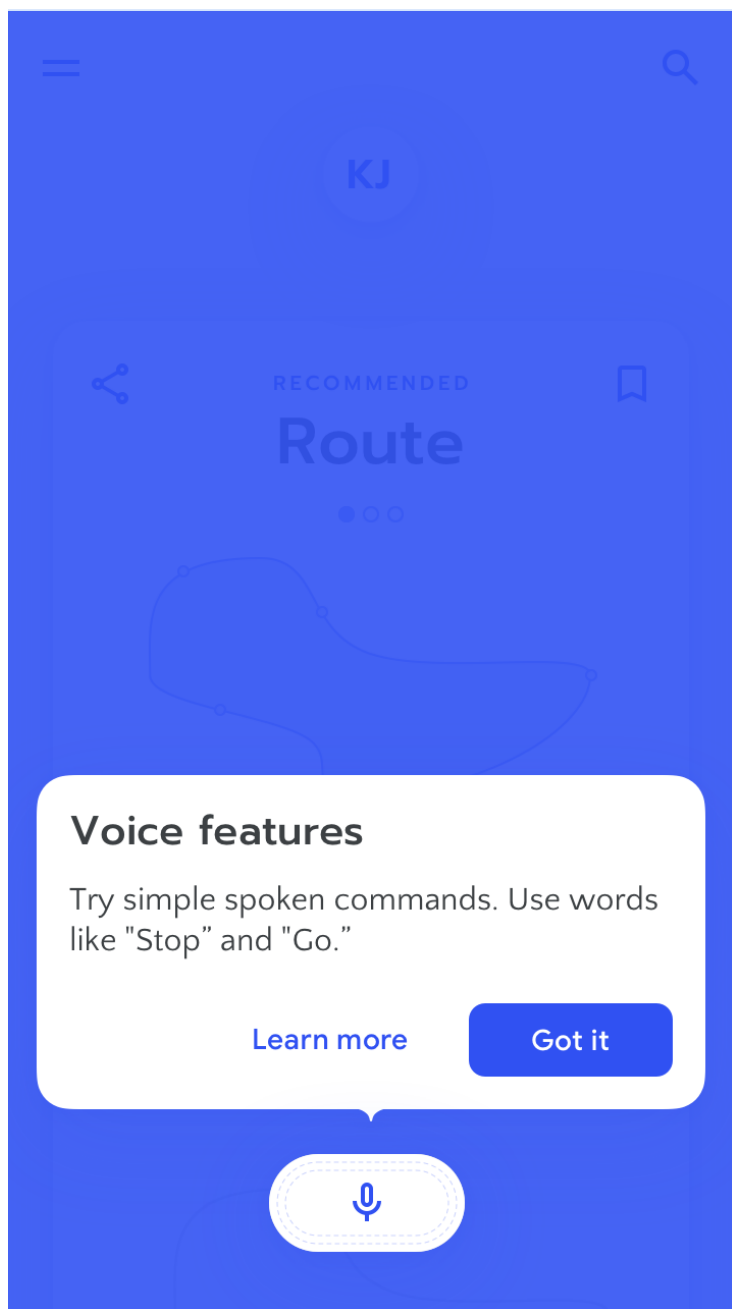
Avoid

Don't create unrealistic expectations by presenting the AI as human-like when it can actually do far less than a person.

Cue the correct interactions

Leveraging human characteristics to build mental models is particularly useful if your product interactions rely on distinctly human behaviors, such as conversation. Using the first person in chatbots and voice interactions can help people intuitively understand how to use your system. It's much easier for users to understand that they can talk to something that invites them to in a conversational way.

However, this approach has its risks. Specifically, if your conversational AI refers to itself as "I", the corresponding user mental model includes near-perfect natural language processing, which your AI may not be able to pull off yet. It's a delicate balance between cueing the right type of interaction while trying to limit the level of mismatched expectations or failures.

Aim for	Avoid
Set expectations for the kind of commands the AI can understand to reinforce the right mental models. [Learn more](#)	Don't set unrealistic expectations about what the AI can do, especially compared to humans.

Summary

Mental models for AI-driven products are influenced by multiple factors including: existing mental models for similar features or products, marketing messages from your team, onboarding and expectations setting, and the feedback relationships in your product. When you set out to help users construct the right mental models for your AI, consider the following:

① **Set expectations for adaptation**. Help people get the most out new AI uses by identifying and building on existing mental models. Ask yourself questions like "What is the user trying to do?", "What mental models might already be in place?," and "Does this product break any intuitive patterns of cause and effect?"

② **Onboard in stages**. Set realistic expectations early. Describe user benefits, not technology. Describe the core value initially, but introduce new features as they are used. Make it easy for users to experiment with the AI in your product.

③ **Plan for co-learning**. Connect feedback to personalization and adaptation to establish the relationship between user actions and the AI output. Fail gracefully to non-AI options when needed.

④ **Account for user expectations of human-like interaction**. Clearly communicate the algorithmic nature and limits of these products to set realistic user expectations and avoid unintended deception.

Want to drive discussions, speed iteration, and avoid pitfalls? Use the worksheet