

APPENDIX

In this appendix, we illustrate each of our 18 human-AI interaction design guidelines with three example applications and three example violations provided by our user study participants when testing the principles against popular AI-infused products. Each application was rated as either a “Clear Application” or an “Application” by participants, while each violation was rated as either a “Clear Violation” or a “Violation”. For each example we indicate the product category the participant was testing, but obscure the specific product names.

Guideline 1.

Make clear what the system can do.

Help the user understand what the AI system is capable of doing.

Example applications:

- [Navigation, Product #1] “Specific examples of things I could search for are in the search bar”
- [Activity Trackers, Product #1] “Displays all the metrics that it tracks and explains how. Metrics include movement metrics such as steps, distance traveled, length of time exercised, and all-day calorie burn, for a day.”
- [Social Networks, Product #2] “[The product] communicates to users that it will evaluate and provide to you potential people to follow based on your interests and the community of people you follow.”

Example violations:

- [Social Networks, Product #1] “This guideline is violated because I cannot even tell what this news feed can/will show.”
- [Photo Organizers, Product #1] “We know the AI is able to detect and associate an image with a category, but the user does not know all the categories available.”
- [Voice Assistants, Product #1] “When [the assistant] was invoked verbally...I was not given any indication of what commands I could request.”

Guideline 2.

Make clear how well the system can do what it can do.

Help the user understand how often the AI system may make mistakes.

Example applications:

- [Music Recommenders, Product #1] “A little bit of hedging language: ‘we think you’ll like’.”
- [Email, Product #2] “While the interface does not give any indication about the level of performance possible, the help page is an application of this heuristic. It sets the expectation that it will start working right away, but will get better with use, making it clear that mistakes will

happen and you can ‘teach’ [the product] to better perform, and failing that, you can set overrides.”

- [Email, Product #1] “There are two things this help text tells me. First, by calling it ‘magic’ - to me that means they don’t know exactly how to explain what’s happening (which means it could do crazy things I didn’t expect). Second, since I can ‘teach’ it that it’s not important, it emphasizes that it might be wrong about things (which is ok cause I can help fix it).”

Example violations:

- [Navigation, Product #1] “There is no indication of accuracy of the time estimates or how conditions may be changing. There is no measure given of how well the AI predictions matched the result once you arrive at the destination.”
- [Voice Assistants, Product #1] “Aside from the ‘Hi, how can I help?’, [the product] does not promise anything more. No expectation of quality is set.”
- [Social Networks, Product #1] “For some of the ads, there is a ‘suggested post’ on the top to indicate that this is just a suggestion. But for the rest of the posts, there is no clue for me to tell the quality.”

Guideline 3.

Time services based on context.

Time when to act or interrupt based on the user’s current task and environment.

Example applications:

- [Navigation, Product #1] “In my experience using the app, it seems to provide timely route guidance. Because the map updates regularly with your actual location, the guidance is timely.”
- [Autocomplete, Product #1] “Suggestions are always present when you might need them (whenever the keyboard is up)”
- [Social Networks, Product #1] “If the user has not accessed [the product] in a while, the application will let the user know that there is something new to be explored - a story, video, etc.”

Example violations:

- [Activity Trackers, Product #2] “Context is very basic, it notifies when I approach my goal; hit my goal; or exceed my goal. The timing of it is not clear, however. The timestamps are varied, too... It feels pretty arbitrary; my interpretation of the reasoning behind the notification can’t be described by my activity or proximity to the goal.”
- [Email, Product #2] “Sending notifications for unimportant messages likely something most people will not want as an interruption.”
- [Voice Assistants, Product #2] “There is no indication of when [the assistant] will actually remind you of the set reminder. She just confirms that the reminder has been set

(mind you this I regardless if you have this Reminders app downloaded or not).”

Guideline 4.

Show contextually relevant information.

Display information relevant to the user’s current task and environment.

Example applications:

- [E-commerce, Product #2] “The feature assumes I’m about to buy a gaming console and shows accessories and games that would go with it, and it features those items prominently (above the product information) on the webpage.”
- [Web Search, Product #2] “Searching a movie title returns show times near my location for today’s date”
- [Navigation, Product #1] “When I use [the product] for driving directions, it remembers where I parked my car. Next time when I open the app, it suggests routing me back to my car.”

Example violations:

- [Activity Trackers, Product #2] “They chose to have a uniform view regardless of context. When I’m moving or static, the view is the same. The only change is that the step counter advances (or not).”
- [E-commerce, Product #1] “If I start looking at a new item (like paper towels), when I scroll down to the end of the page, I get recommendations related to the recent items I viewed (tennis balls)... it doesn’t take into account that I am currently looking for paper towels and I have already purchased the tennis balls.”
- [Email, Product #2] “What goes into the [tabs] is the same all the time. It does not change based on the context, for example to...emails related to the meeting I’m attending or the message I’m reading.”

Guideline 5.

Match relevant social norms.

Ensure the experience is delivered in a way that users would expect, given their social and cultural context.

Example applications:

- [Photo Organizers, Product #1] “[The product’s] album suggestions feature is able to recognize people’s pets and uses the verbiage “important cats & dogs”, understanding that people’s pets are important to users and are like family even.”
- [Voice Assistants, Product #1] “[The assistant] uses a semi-formal voice to talk to you - spells out “okay” and asks further questions.”
- [Navigation, Product #1] “If you select walking, the AI avoid(s) busy roads and searches for trails.”

Example violations:

- [Activity Trackers, Product #1] “Provides a reminder to stand up without understanding my social context (e.g., in a meeting, having lunch with a friend etc)... Does not consider the social context prior to sending notifications for activity and does not use tone appropriately - just says “time to stand!” no matter what”
- [Email, Product #2] “the system does not follow the social norms of a workplace. For example, one norm is to pay attention to your manager. However, even with access to company hierarchies, it isn’t clear that the system will put messages from one’s direct manager in the [appropriate tab] automatically.”
- [Voice Assistants, Product #1] “[The assistant] does not match expected conversation norms. When asked “[wake command], set a reminder for next week” the AI does not recognize that the remind me command has been invoked and responds “Sorry, I can’t set reminders yet.” Only when the specific command syntax “Remind me to [reminder] on [date] at [time]” is used does [the assistant] understand to set a reminder...It seems [the assistant] is unable to interpret conversational language and instead requires a very specific command syntax.”

Guideline 6.

Mitigate social biases.

Ensure the AI system’s language and behaviors do not reinforce undesirable and unfair stereotypes and biases.

Example applications:

- [E-commerce, Product #2] “The feature does not unfairly assume gender biases in some search results that could potentially introduce them. For example, a search for tools or diapers could accidentally serve related products that are gender biased. The system seems to provide highly specific recommendations of the very same product type.”
- [Web Search, Product #2] “a search for CEO or Doctor shows somewhat diverse people in the resulting images...The images are pretty diverse in terms of gender and ethnicity, although still lack in some respects such as disability”
- [Autocomplete, Product #2] “The autocomplete feature clearly suggests both genders [him, her] without any bias while suggesting the text to complete.”

Example violations:

- [Voice Assistants, Product #1] “When asked “...can you change your voice” [the assistant] responds in a male voice saying “Here is an example of my other voice. Would you like me to use this one?”... While it’s nice that a male voice is given as an option, the default...voice is female, which reinforces stereotypical gender roles that presume a secretary or receptionist is female.”

- [Autocomplete, Product #2] “autocomplete suggests names correctly for Western sounding names but falls short for ethnic-sounding names.”
- [Navigation, Product #1] “Regards the “Walking” transport there’s no way to set an average walking speed. [The product] assumes users to be healthy.”

Guideline 7.

Support efficient invocation.

Make it easy to invoke or request the AI system’s services when needed.

Example applications:

- [Voice Assistants, Product #1] “I can say [wake command] to initiate.”
- [E-commerce, Product #1] “In addition to the system giving you recommendations as you browse, you can go to your “Browsing history > Manage history > More like this” to get recommendations specific to a particular product.”
- [Web Search, Product #1] “user can highlight a specific part of an image to search for that specific piece. This shows up on every image, so user can use it anytime they like.”

Example violations:

- [Navigation, Product #1] “[The product] remembers where you parked your car. However if it fails to remember, or I want it to remember something else (e.g., where I chained up my bike), it is not possible (or at least not easily discoverable) to invoke the capability when I need it.”
- [Navigation, Product #2] “Guideline is violated because user cannot ask the system for alternative routes if they are not detected initially.”
- [E-commerce, Product #2] “Many of the products I searched for did not show the “Customers also considered” AI feature. There is no way to invoke this feature manually.”

Guideline 8.

Support efficient dismissal.

Make it easy to dismiss or ignore undesired AI system services. Example applications:

- [E-commerce, Product #2] “Feature is unobtrusive, below the fold, and easy to scroll past...Easy to ignore.”
- [Social Networks, Product #2] “[The product] allows the user to easily hide or report ads that have been suggested by the AI by tapping the ellipses at the top right of the ad.”
- [Voice Assistants, Product #1] “I can say “nevermind” to dismiss it once I have said [wake command]. I can also just not say anything and it stops listening.”

Example violations:

- [Autocomplete, Product #2] “I didn’t see a dismiss button. I can dismiss it by dismissing the whole keyboard”

- [Activity Trackers, Product #1] “The guideline is violated because it is not clear within the Steps page, how to turn off the background/ambient step tracking functionality.”
- [Navigation, Product #1] “Suggested locations based on calendar entries can’t be removed from the suggestions.”

Guideline 9.

Support efficient correction.

Make it easy to edit, refine, or recover when the AI system is wrong.

Example applications:

- [Navigation, Product #1] “If [the product] is wrong about where I parked my car, it provides an easy way to edit the location by dragging on the map.”
- [Web Search, Product #2] “automatically ‘corrects’ spelling errors, etc. but gives option at top to return to query as originally typed...Notes that the query had been corrected and is one click to revert back to original”
- [Voice Assistants, Product #2] “Once my request for a reminder was processed I saw the ability to edit my reminder in the UI that was displayed. Small text underneath stated “Tap to Edit” with a chevron indicating something would happen if I selected this text.”

Example violations:

- [E-commerce, Product #1] “I already recently bought the items which are in my recommendation list & there is no message to discontinue nor option for users to deselect.”
- [Activity Trackers, Product #1] “As far as I can tell, there is no way for the user to edit the number of steps collected. The user can delete the data point altogether...There is just no way to manually input or change the data.
- [Web Search, Product #1] “Searches can be easily corrected with a new query (which are sometimes suggested by [the product] itself). However, editing a seemingly AI system override to interpret “Sea of” to “SEA to” is not possible.”

Guideline 10.

Scope services when in doubt.

Engage in disambiguation or gracefully degrade the AI system’s services when uncertain about a user’s goals.

Example applications:

- [Navigation, Product #1] “If more than one line takes the same route the user can choose between the preferred line.”
- [Autocomplete, Product #1] “It usually provides 3-4 suggestions instead of directly auto completing it for you”
- [Voice Assistants, Product #2] “If I didn’t respond or if I spoke quietly, [the assistant] let me know they had trouble hearing me”

Example violations:

- [Navigation, Product #1] “When searching for a specific restaurant, if the right place can’t be found, it fails to recognize that you are searching for food and does not suggest restaurants in your area...The AI doesn’t recognize the broader user goal and suggest possible alternatives when it can’t find the exact result.”
- [Email, Product #2] “This heuristic is violated. There is no indication if the system is unsure. The system does not ask for assistance (e.g. request feedback on an uncertain classification). It is not clear how it decides what to do with unclear messages, or what the current quality of the model is, beyond the information from the help screen which says the system will improve over time.”
- [Music Recommenders, Product #2] “The system doesn’t ask for user input about its recommendations, nor does it provide a way for users to provide input.”

Guideline 11.

Make clear why the system did what it did.

Enable the user to access an explanation of why the AI system behaved as it did.

Example applications:

- [E-commerce, Product #1] “Clicking “Why recommended” explains why they have recommended that particular item to you.”
- [Music Recommenders, Product #2] “I think this applies because each of recommendation has some information as to which songs are displayed on it - similar to the song, from the same artist, from the same album etc.”
- [Navigation, Product #2] “The route chosen by the app was made based on the Fastest Route, which is shown in the subtext.”

Example violations:

- [Music Recommenders, Product #2] “The system provides no information about why the recommended songs/artists are chosen.”
- [Email, Product #2] “There is no indication of why a message is classified as it is, and I cannot find a way to find out. A message could be classified as [unimportant] because of the content, the sender, or any number of other reasons and no details are available, either per message or as a summary of the [unimportant category] rules which have been learned.”
- [Activity Trackers, Product #2] “There are no explanations in the app about how walking detection and step count are measured. The system is pretty deterministic in how it makes measurements about my walking behaviour and then how it matches them with the initial goal I set.”

Guideline 12.

Remember recent interactions.

Maintain short term memory and allow the user to make efficient references to that memory.

Example applications:

- [Navigation, Product #1] “Opening the app shows a list of recent destinations, as well as allows you to access “favorite” locations.”
- [Web Search, Product #1] “[The search engine] remembers the context of certain queries, with certain phrasing, so that it can continue the thread of the search (e.g., “who is he married to” after a search that surfaces Benjamin Bratt)”
- [Voice Assistants, Product #1] “[The assistant] seems to remember conversation context at least one command back. When asked “[wake command], what’s the reminder?” she announces the last unheard remind me.”

Example violations:

- [Social Networks, Product #1] “This guideline is violated since there is no indication of “what you have read”
- [Web Search, Product #2] “Although many food/recipe searches were made right before, a search for “stir fry” has a music video as a top result”
- [Voice Assistants, Product #2] “Nope set the same reminder twice to drink water at 4:15 and it had no idea”

Guideline 13.

Learn from user behavior.

Personalize the user’s experience by learning from their actions over time.

Example applications:

- [Music Recommenders, Product #2] “I think this is applied because every action to add a song to the list triggers new recommendations.”
- [Social Networks, Product #1] “I am pretty sure they do this as I have observed that I get more ads of a certain type if I accidentally or deliberately hovered over the ad for a while. I also have a suspicion that friends are also preferred that are more likely to generate a like or comment from me. It is less clear how negative signals are taken into account.”
- [Email, Product #1] “(My guess is) the system learns from what previous emails have attracted more attention from me (i.e. longer/more frequent reply, reading time taken, longer email threads, etc.) and infer email importance.”

Example violations:

- [Web Search, Product #1] “The search results do not take into account the user’s previously visited sources. In this example from real life: I search for recipes on [the search engine] all the time, and I only ever click on recipes from

4 sources: Serious Eats, Pioneer Woman, 101 Cookbooks, and Smitten Kitchen. I often search for these sites as navigational queries on [the search engine]. However, [the search engine] rarely surfaces recipes from these sites for me.”

- [Activity Trackers, Product #1] “The app notifies/nudges users to push them to achieve goals, but I haven’t seen a change in behavior even though my patterns have changed.”
- [Navigation, Product #1] “If you don’t take the suggested route because you don’t want stop and go traffic on the highway, it never learns. You can have it suggest routes to avoid tolls and highways, but you have to manually select those preferences.”

Guideline 14.

Update and adapt cautiously.

Limit disruptive changes when updating and adapting the AI system’s behaviors.

Example applications:

- [Music Recommenders, Product #2] “Once we select a song they update the immediate song list below but keeps the above one constant.”
- [Social Networks, Product #1] “Think this is good. When I unfollow someone it shows there stuff for a little bit? But after a day or so its gone. Or once I reload. But not RIGHT away.”
- [Web Search, Product #2] “After clicking and returning from a search result, the order of the search results didn’t change. The updated “people also searched for” that was relevant to the clicked result was highlighted and contained”

Example violations:

- [E-commerce, Product #1] “Maybe this is just my experience, but when I accidentally browse an item suddenly my entire recommendation list changes to things relevant to that new item. The change is really jarring, especially when the browsing is a result of a curious moment or an accident. For example, I just clicked on one camera after viewing a lot of tennis balls and all my recommendations changed from tennis related items to electronic devices.”
- [Social Networks, Product #1] “The refresh (pulling down action) can sometimes be triggered unintentionally and thus causes some posts disappeared.”
- [Social Networks, Product #2] “If I’m in the middle of watching a video post, every now and then they will scroll me up to the first post.”

Guideline 15.

Encourage granular feedback.

Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.

Example applications:

- [Social Networks, Product #2] “[The product] allows the user to “Hide an Ad,” and then when doing so, solicits feedback to improve relevancy of future ads. ”
- [Music Recommenders, Product #1] “Love/dislike buttons are prominent and easily accessible.”
- [Email, Product #1] “The user can directly mark something as important, when the AI hadn’t marked it as that previously.”

Example violations:

- [Voice Assistants, Product #2] “Once [the assistant] performed the task I had asked of it, there was no additional ability to customize the experience or give feedback on my satisfaction; even when I chose to remove the reminder right after I verbally requested it.”
- [Photo Organizers, Product #2] “There is no feedback channel for the user to validate the album suggestions.”
- [Navigation, Product #2] “The system does not do this. For example, if a user were to consistently pick a route home from work that was not the main route, then there might be patterns in traffic or circumstances that might be considered. If the user (instead of spouse) picked up a child on the way home from work because traffic was bad on the normal route, then the system should learn this habit, or allow the user to program it in... Doesn’t let the user tweak the routes based on prior behavior.”

Guideline 16.

Convey the consequences of user actions.

Immediately update or convey how user actions will impact future behaviors of the AI system.

Example applications:

- [Music Recommenders, Product #1] “Tapping the like/dislike button results in immediate popups informing that the user will receive more/fewer recommendations like it.”
- [Social Networks, Product #2] “[The product] communicates that hiding an Ad will adjust the relevancy of future ads.”
- [Web Search, Product #1] “With different filters, the search results are auto updated.”

Example violations:

- [Social Networks, Product #1] “You can unfollow or like or interact but how that affects you isn’t clear. It just sorta happens.”
- [Email, Product #1] “It’s clear there is some consequence for the user actions, but it’s not well-defined. So while it exists, the lack of clarity on the exact outcomes for behavior are unclear.”

- [Photo Organizers, Product #2] “There isn’t any messages to confirm that the system will learn from my dismiss or save action.”

Guideline 17.

Provide global controls.

Allow the user to globally customize what the AI system monitors and how it behaves.

Example applications:

- [Web Search, Product #2] “It has settings such as...private results that help users get results that are more relevant to them.”
- [Photo Organizers, Product #1] “[The product] allows users to turn on your location history so the AI can group photos by where you have been.”
- [Navigation, Product #2] “A few options to adjust my preferences for how I want to get directions. Uses common patterns for multi-select and single-choice.”

Example violations:

- [Email, Product #2] “The only option is to turn the system on or off. It otherwise applies to all messages at all times. It is not clear how the system works, or what data it monitors, but presumably it applies to entire email contents.”
- [Music Recommenders, Product #2] “[The product] does not provide a mechanism to turn off tracking of listening data (at least not in the process of building a playlist) or to impact what it learns.”
- [Navigation, Product #1] “There is no apparent way to customize what information the system has access to or what it learns. Marked locations can be removed, but otherwise, I see no customizability with regard to data access or behavior.”

Guideline 18.

Notify users about changes.

Inform the user when the AI system adds or updates its capabilities.

Example applications:

- [Email, Product #2] “The help tab for the interface features a “What’s new” section which could be used to inform the user about AI system additions or capability updates.”
- [Social Networks, Product #2] “Updates of privacy and regulations...They do post their updated privacy or legal regulations and make the UI inaccessible until the user agrees.”
- [Navigation, Product #2] “I don’t have a way to show this, but it does provide small in-app teaching callouts for important new features. New features that require my explicit attention are pop-ups.”

Example violations:

- [Music Recommenders, Product #1] “The algorithm feels like it constantly updates, with a slightly different feel to my recommendations every week. However there is no explanation of what has changed.”
- [Social Networks, Product #1] “It’s always a mystery when [the product] updating the ranking algorithm.”
- [Web Search, Product #2] “there are no notifications or mentions when the search algorithm changes or new capabilities are added, such as disambiguation, the presentation of recipes directly, etc...Updates or changes are only noticeable with use (like new AI feature applications) while others are completely unknown (like changing the search algorithm to return more diverse CEO images)”