The Big House Dataset: Desired Applications and Interactions

Meghan Clark
University of California, Berkeley
mclarkk@berkeley.edu

Mark W. Newman
University of Michigan
mwnewman@umich.edu

Prabal Dutta
University of California, Berkeley
prabal@berkeley.edu

ABSTRACT

What do normal, everyday people want to do with consumer IoT systems in their homes? How do different IoT interfaces affect what users think the system can do? We deployed four questionnaires to collect information on the interactions and applications that typical home occupants desire from smart home IoT technologies. We received over 1,500 responses, about 600 of which are users' descriptions of IoT applications they would like in their home, and about 900 of which are users' interactions with a smart home AI. This dataset was released publicly along with a paper describing key findings on the priming effects of common IoT system interfaces. However, the data is a rich source of additional information related to what people want to do and how they want to do it. Researchers in both academia and industry can benefit from the insights this dataset has to offer about consumer IoT applications, user-centric system design, and trade-offs between interfaces.

1 INTRODUCTION

We collected this dataset to better understand what users want to do in smart homes and how the interface design impacts those desires. By analyzing the data, we found that users are heavily primed by the primitives and metaphors provided by the interface, which impacts the system's workload and the amount of value that users ultimately extract from their IoT systems [1]. Agents and datastreams emerged as more promising conceptual models than devices for helping users take advantage of available IoT capabilities. We also found that the responses contained many nouns but only a few core verbs, suggesting that actions may be better suited than devices to be the fundamental building blocks around which to standardize IoT interfaces and APIs.

We released the dataset publicly as supplemental material associated with the paper. However, the ACM page obscures the location and nature of the supplemental material, Better awareness of this dataset could help the community discover that there is an open-access dataset available, free of charge, that can provide:

- User interest in different categories of IoT applications.
- User interest in different categories of IoT devices.
- Preferred conversational agent gender and name.
- Training data for smart home natural language interfaces.
- The entities, capabilities, attributes, and logical control flows that best match end-user mental models.
- All of the above filtered by various demographic criteria.

2 THE DATA

This data was collected with IRB approval from native English-speaking Amazon Mechanical Turk workers using four different questionnaires. Each of these four questionnaires described the same hypothetical smart home using a unique representation, and then asked the respondents what application they would like in their smart home (Figure 1). The four representations were:

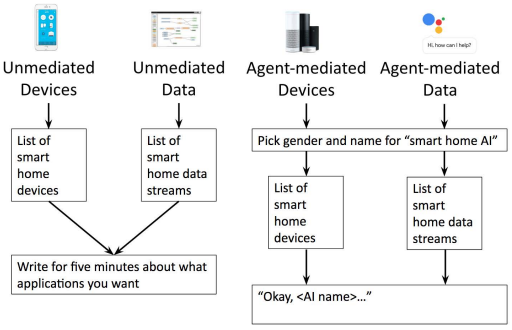

Figure 1: The structure of the four questionnaires. The questionnaires described the same hypothetical smart home in four different but functionally equivalent ways (representing different classes of IoT interfaces), then asked the user to either write about what applications they wanted or to write what they wanted to say to their smart home AI. Each box corresponds to a separate page displayed to the user.

Unmediated Devices (313 responses): The smart home was described by providing a list of smart devices available in the home. This is the most common conceptual model presented to IoT users, epitomized by the many apps that display icons of each device that users directly manipulate.

Unmediated Data (302 responses): The smart home was described by presenting a list of data streams available in the home. This conceptual model is epitomized by dataflow-based programming interfaces like IBM Node-RED.

Agent-mediated Devices (442 responses): Respondents were asked to assign a gender (male/female/androgynous) to an intelligent agent and pick a name for it, then were presented with a list of devices that the smart home agent had access to (the same list as used in Unmediated Devices).

Agent-mediated Data (478 responses): Respondents were asked to assign a gender (male/female/androgynous) to an intelligent agent and pick a name for it, then were presented with a list of data streams that the smart home agent had access to (the same list as used in Unmediated Data).

After receiving one of the four descriptions above, the respondent was presented with a text box. Those given an unmediated description received the prompt, "Write in the box below for five minutes, describing different applications that you would want in your smart home." Those given an agent-mediated description were given the prompt, "Write in the box below what you want <AI name> to make your smart home do for you. 'OK, <AI name>, ...' "

DATA'18, November 4, 2018, Shenzhen, China
2018. ACM ISBN 978-1-4503-6049-4/18/11.
https://doi.org/10.1145/3277868.3277878

PUBLIC DOMAIN To the extent possible under law, the authors have waived all copyright and related or neighboring rights to this work. This work is published from the United States.

Label	Includes	Excludes
Immediate action	Actions that happen immediately, like "turn on the lights."	Commands to obtain information, like "tell me my weight." (See indirect question.) Actions that may not happen right away. (See conditional action.)
Conditional action	Actions that are conditioned on something like an event or time, e.g. "turn on the light when I get home." Conditions can also be signaled by words like "if," "while," "before," "unless," "at [time]," and "every." Would this action always happen right away? If no, then it is likely a conditional action.	Notifications, reminders, alerts, and other requests for information. Examples: "let me know when my husband gets home," "alert me if the temperature gets above 70." (See notification.)
Direct question	Any phrase that ends (or should end) with a question mark.	Any phrase without a question mark.
Indirect question	Immediate requests for info that are phrased as commands rather than questions. E.g., "tell me how much I weigh," "let me know what the weather is outside."	Phrases that start with phrases like "tell me" or "let me know" but include conditions that may trigger a notification later. (See notification.)
Notification	Requests for information that may not be answered immediately. In other words, alerts, notifications, and reminders. Things like "notify me if [event]" "remind me at [time]" "tell me when [event]" "let me know when [event]." The conditions are signaled by words like those listed in the conditional action description.	Requests for information that are fulfilled immediately. Sometimes the differences between indirect questions and notifications can be subtle. For example, "tell me when my husband gets home" is likely a notification, whereas "tell me when my husband will get home" is an indirect question.
Wants device	Phrases that express a desire to have or use a particular device or devices. E.g., "smart thermostat," "I would like smart lights," "I would like a smart watch in order to control my lights," "I would like lights that turn on and off automatically."	Sometime it can be hard to tell whether the category should be "Wants device" or "Wants app." Imagine that you are a genie. After hearing a person say this phrase, would you give them a new device ("Wants device") or would you enchant their devices to behave in a particular way ("Wants app")?
Wants automation	Phrases that express a desire for a particular application or system behavior that requires automation. Basically, if the system takes action based on some rule or trigger other than the user's manual control. E.g., "I would like an app that turns my lights on and off when I'm not home" or "automatic temperature control." or "I would like my lights to turn on and off automatically."	If a phrase includes both devices and behavior, it can be hard to tell whether the category should be "Wants device" or "Wants automation." Imagine that you are a genie. After hearing a person say this phrase, would you give them a new device ("Wants device") or would you enchant their devices to behave in a particular way ("Wants automation")?
Wants remote control	Phrases that express a desire for a remote control ability. When a person wants the ability to manually initiate an action. Example: "I want to be able to turn my lights from my car."	Phrases like "turn on the lights" are immediate actions, where want remote control covers phrases like "I would like to be able to turn on the lights."
Wants to know	Phrases that just express a desire to know something, receive alerts, or monitor. For example: "I would like to know..." or "it would be great to know..." or "alerts for traffic" or "weather conditions" or "I would like to monitor my home."	Actual questions or requests for information. (See direction question, indirect question, and notification.)
None of the above	Anything that isn't covered by one or more of the above categories.	Things that could be multiple categories. Pick the closest option that best captures the spirit of the phrase.

Figure 2: Inclusion/exclusion criteria for utterance labels.

For researchers interested in what devices or applications people say they want in a smart home, the responses to the two unmediated questionnaires will be the most helpful. For researchers interested in how users would speak directly to an intelligent agent in the home, the responses to the two agent-mediated questionnaires will be the most helpful. The questionnaires also collected demographic and computing expertise information, so researchers can filter responses based on population characteristics like CS expertise or age.

2.1 Table Descriptions

The dataset includes the database, a Markdown file that explains the database schema, and an example Python script for grabbing database records by subpopulation and preparing them for processing. The database has three tables: *responses*, *sentences*, and *annotations*.

2.1.1 Responses. Each record in the *responses* table contains each response, information about the associated prompt, and information about the demographics of the respondent. The demographic information can be used to select subpopulations from the data, and the response ID can be joined with the other tables to fetch the corresponding sentences, parses, and annotations.

2.1.2 Sentences. Each record in the *sentences* table contains information about a particular sentence. It contains the response ID for the response that the sentence came from, the sentence itself in lowercase, and the parse tree with parts of speech (POS) tags for that sentence as determined by the Stanford Parser (version 3.5.1). This table is useful for analyzing parts of speech like nouns or verbs, as well as syntactic information like subordinate clauses, which often signal triggers in trigger-action style automation rules.

2.1.3 Annotations. We took the first sentence of each response and asked three trained experts to select labels for each of the sentence's independent clauses from a set of labels generated through an inductive coding process. The labels and their inclusion/exclusion criteria are shown in Figure 2. This table is useful for visualizing the distribution of interaction types under each prompt, or analyzing particular types of responses.

3 ACKNOWLEDGEMENTS

This work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, as well as the TerraSwarm Research Center. The work was also supported in part by the NSF Graduate Research Fellowship and NSF award CNS-1824277. Special thanks to Catherine Finegan-Dollak and Laura Wendlandt for their help with annotations.

REFERENCES

[1] Meghan Clark, Mark W Newman, and Prabal Dutta. 2017. Devices and Data and Agents, Oh My: How Smart Home Abstractions Prime End-User Mental Models. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 44.