Mental Elf: An Alexa Skill for Mental Health Advocacy

Dhanush B K, Andrew Hoang, Ben McCready, Sara Ng, Naomi Tachikawa Shapiro, Yifan Zhuang

University of Washington

OUR VISION

- Promote accurate information on mental health conditions
- Engage users in naturalistic conversation
- Break stigma surrounding the discussion of mental health

Target Users:

We intend our target audience to grow as Mental Elf evolves. In its current state, Mental Elf targets users who are curious about various mental health conditions and their symptoms, as this is our skill's starting knowledge base. However, as the system evolves, we would also like to attract users who are also interested in Mental Health *awareness* – i.e., discussions around how we can better approach and view mental health care. This would accompany Mental Elf's ability to *advocate* for those living with mental health issues.

SYSTEM SKILLS

QA Capabilities:

- Condition overviews
- Symptoms
- Treatment options

Conversational Capabilities:

- Elicit user opinions about provided information
- Recite real opinions and stories from users of a public online mental-health forum
- Control flow of conversation with user to navigate between QA and other conversational modules

SAMPLE DIALOGUES

DATA

The following conditions can be described via overview, treatment, and symptoms:

- Autism
- Alcohol use disorder
- Kleptomania
- (Postpartum) depression
- PTSD
- Social anxiety disorder
- ADHD
- Atypical depression
- Schizophrenia
- Binge-eating disorder
- Bulimia
- BPD
- Dissociative disorders
- Panic attacks
- Panic disorders
- Rett syndrome
- SAD
- Bipolar disorder
- Tourette syndrome

These conditions have associated anecdotal data:

- Seasonal Affective Disorder (SAD)
- Depression
- PTSD
- Panic attacks
- Bulimia
- Post-partum depression

While we have initialized a small database, we are also able to draw upon outside information based on user input. For example, if a condition is queried that is not covered in the database, our backend will attempt to crawl the data sources and find the requisite information for overview, symptoms, and treatment.

INTENTS

GiveOverview(CONDITION): Gives basic information about specific mental health conditions.

GiveSymptoms(CONDITION): Returns a list-like phrase of possible symptoms of the condition.

GiveTreatment(CONDITION): Returns a list-like phrase of possible treatments (if any).

GiveForum(CONDITION): Provides an anecdote from someone discussing the condition on an online public-facing mental health forum.

AffirmPassive: Detects affirmative responses (e.g., *Yes* in response to the question *Would you like to hear about treatment for PTSD?*) and positive opinion (e.g., *I think that's interesting*).

Negative: Detects negative responses (e.g., *No thank you* in responses to *Would you like me to elaborate?*).

Other various Easter-Egg type intents – Try it out yourself and see! ☺

SYSTEM ARCHITECTURE

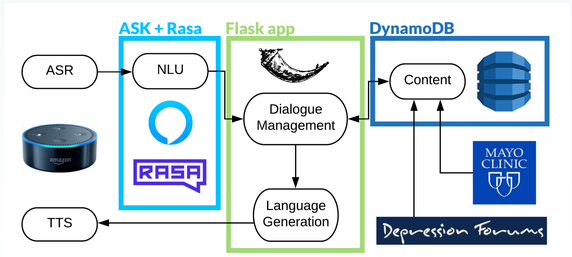

User inputs undergo intent classification by the Alexa Skills Kit, with unknowns aided by Rasa NLU. On the backend, a Flask-ASK web app processes the intent and user input, reading content from DynamoDB to help generate an appropriate output.

DIALOGUE MANAGEMENT

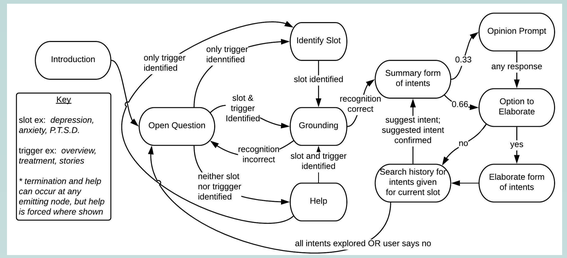

In this system, dialogue moves within a single condition, exploring information from the database, as well as eliciting user responses.

We have chosen to centralize dialogue management within the skill, rather than allowing the user to control most of the dialogue flow. Admittedly, this might seem frustrating for some users. However, we feel that this better provides the user with information they may not otherwise think to ask for and allows them space to voice opinions.

EVALUATION METRICS

Team member Andrew interacts with the skill via an Amazon Echo Dot during the testing phase. We have found that curated conversations from team members work well.

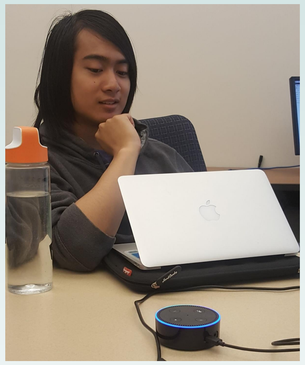

- **TTS Performance**: Handled for us by the Alexa system
- **ASR Performance**: The Alexa platform also provides most of the ASR implementation; we needed to give samples of technical terms to help the ASR.
- **Task Ease**: This can be tracked by how many times users repeat the same/similar utterances, implying that the question they asked was not answered the first time.
- **Interaction Pace**: As our intention is not to restrict the amount of information that we or the user provides, there is no set constraint on the pace of the conversation. However, we want to limit the amount of erroneous corrections that will slow the pace of the conversation.
- **User Expertise**: Tracking how well users know what utterances our bot can handle can be traced by the number of user utterances that go un-handled by our intent handler.
- **System Response**: This metric will largely align with the speed at which our system can access our DynamoDB database.
- **Expected Behaviour**: During evaluation, matching expected intent with the triggered intent will be key in assessing how well our system's behaviour aligns with the users' expectations.

PRELIMINARY PERFORMANCE

Merits:

- Correctly follows prescribed dialogue flow
- The ASR is correct with slot sample conditions

Issues:

- TTS mispronounces unknown words from the database
- Failure to allow users to break dialogue chain
- Unable to correctly identify problem on conditions that are not classified as conditions by Alexa's slot filling functionality

WHAT WE LEARNED

- How to create a basic interaction model using the Alexa Skills Kit
- How to store and query information using DynamoDB
- How SSML can be used to make output utterances more natural
- How to bootstrap an understanding model using Rasa NLU
- Various approaches to dialogue flow
- Best practices in data curation
- Best practices in pair programming
- Skills in working with a web service backend

CHALLENGES

- Some system dependencies with ngrok were difficult to overcome (read: WSL problems)
- Overlap of many tasks made task allocation difficult
- Development requires a very stable internet connection
- It was difficult to find *clean* mental health, so we had to hand-curate a lot of data
- Integrating moving parts
- Many ethical considerations in which data to scrape and what kind of dialogues we are willing to have

NEXT STEPS

- Allow for more user-oriented control of dialogue flow
- Perform sentiment analysis on user opinion input to inform both dialogue flow and nature of output provided
- Include broader range of topics (e.g., prevalence, stigma, and advocacy)
- Point users to relevant online resources or local healthcare advocates
- Consult in-field experts to better understand successful dialogues for advocation and engagement
- Better tailor experiences to individual users

ACKNOWLEDGMENTS

Thank you to our teacher Hao Fang for guidance and to Amazon for providing us with Amazon Echo Dots.

CONTACTS AND RESOURCES

Source files for our skill can be found at github.com/tsnaomi/mental-elf.

Dhanush – dhanush@uw.edu
Ben – bmccr@uw.edu
Naomi – tsnaomi@uw.edu
Andrew – andrewh3@uw.edu
Sara – sbng@uw.edu
Yifan – zhuang93@uw.edu