

1 • Introduction

Motivation

Finding a good fit for a new job requires clear knowledge of personal preferences as well as position availability. One of the most effective ways people find and learn about jobs is by talking about their search with others. We aim to emulate this dialogue flow with a chatbot that gleans relevant information through conversation and then suggests real, open job positions for the user to consider.

Target User

Broadly, we target job seekers who are unsure on how to find a good job fit but have at least some preferences. Our target users may include:

- New job seekers (i.e., students)
- People seeking to change careers
- Job seekers who are new to the country

Chatbot Overview

Our aim is to narrow down user preferences using a **slot filling approach** wherein the slots are decided by available fields in the jobs database. HireMe learns user preferences one by one while suggesting specific job opportunities from the database whenever the user is unable to provide a preference that is present in the database. The bot can provide suggestions for most related job title as per the user's skillset.

Tech Stack: Python, AWS Lambda, Alexa Skills Kit, DynamoDB, boto3

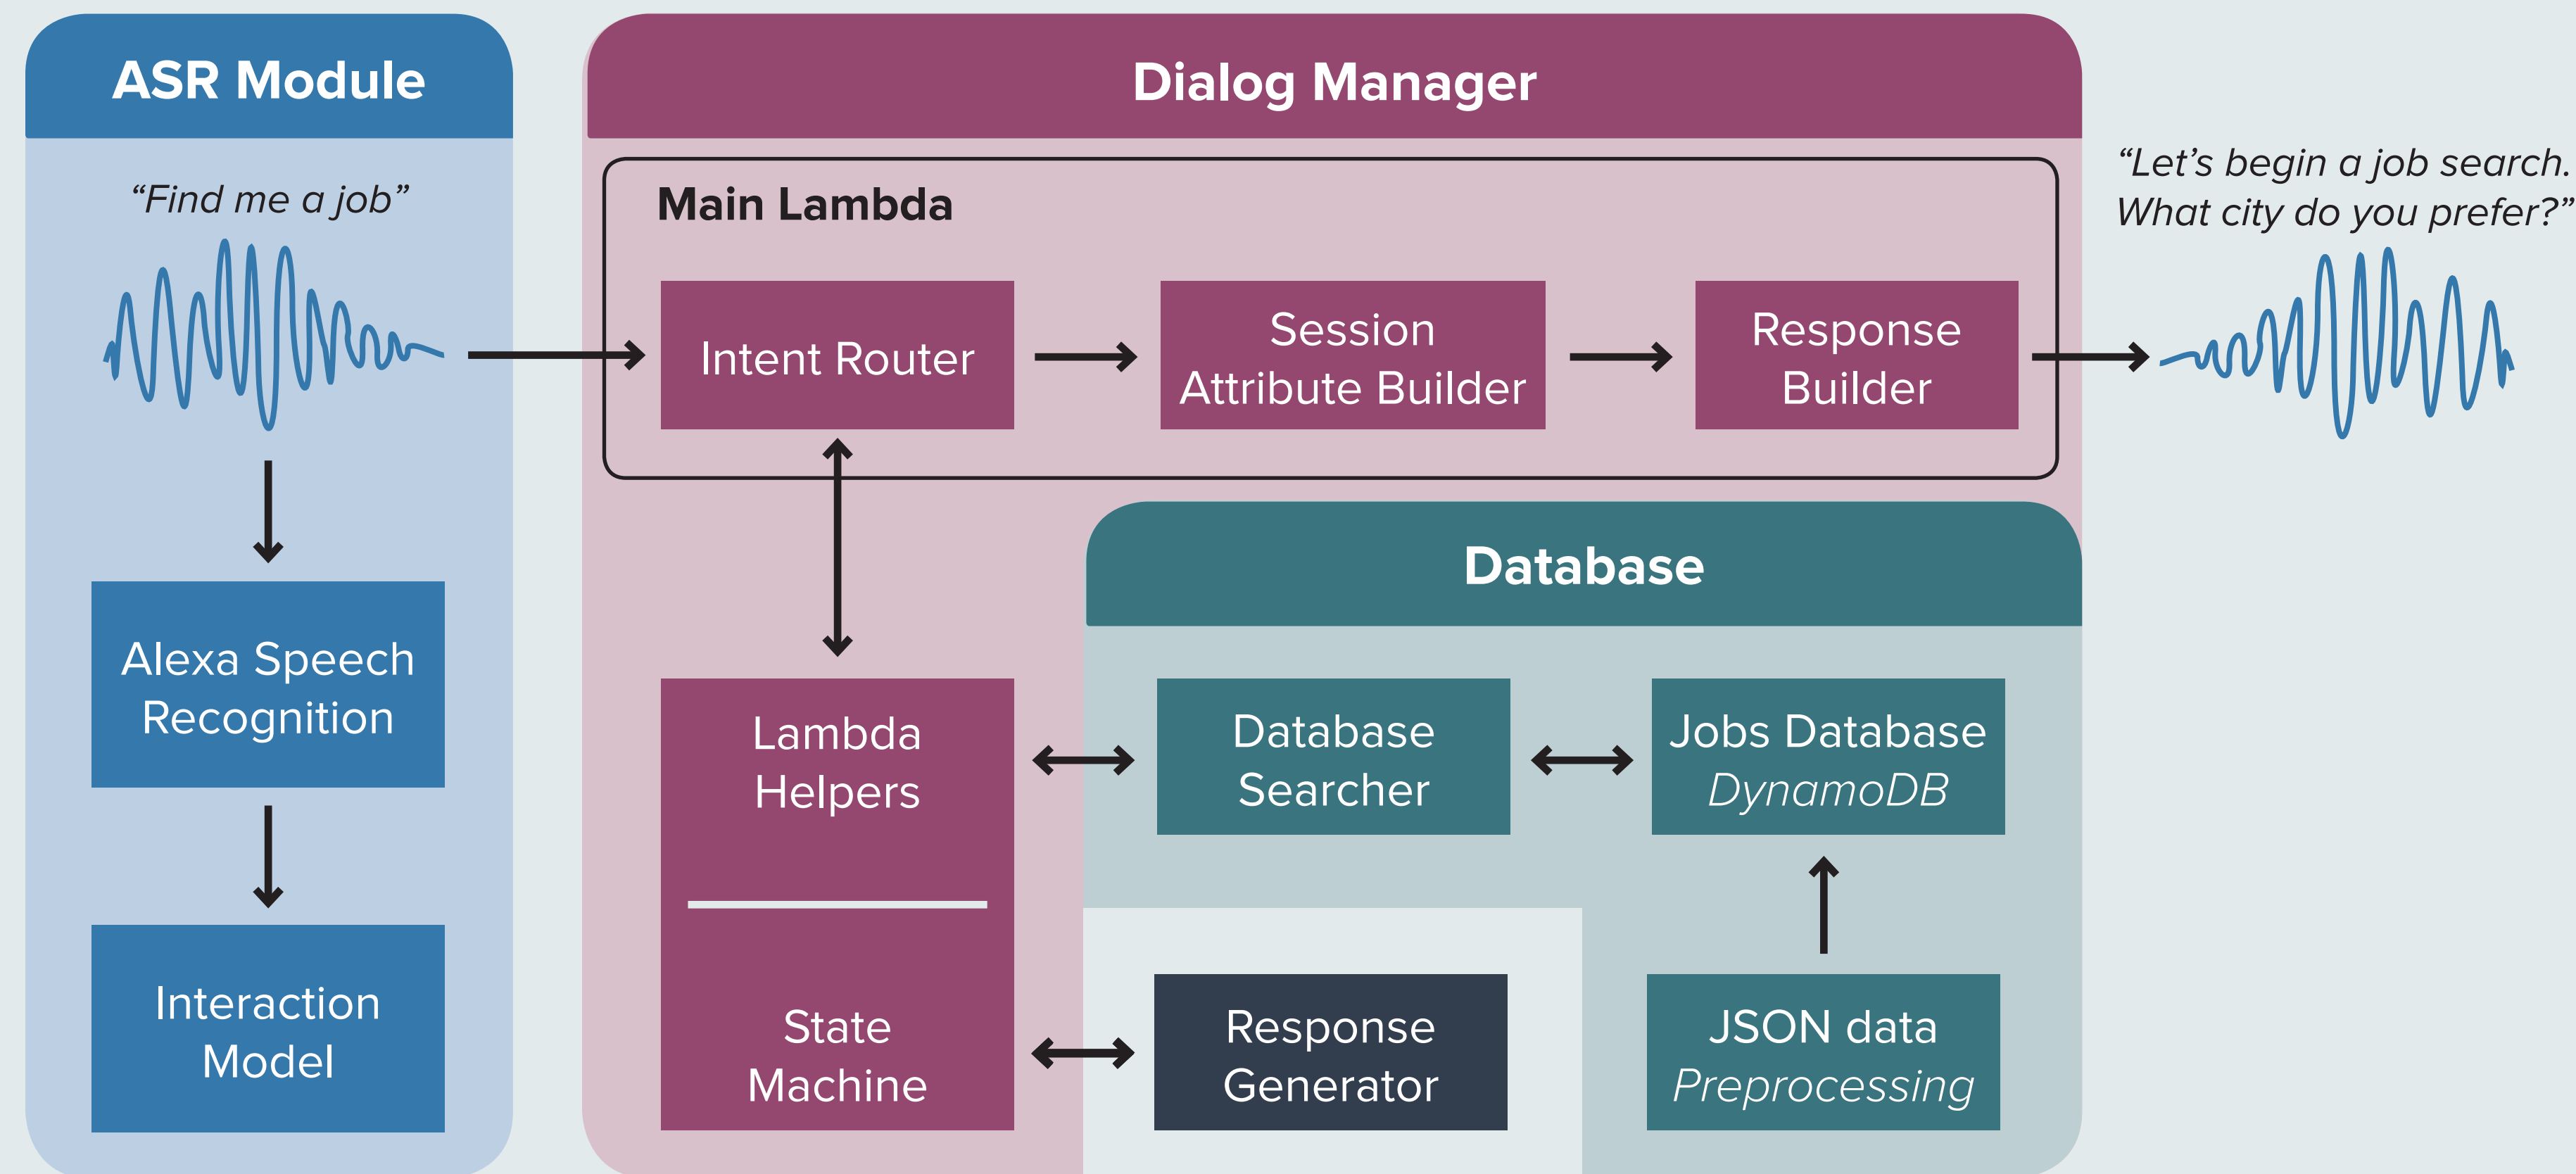
2 • Sample Dialog



3 • System Architecture

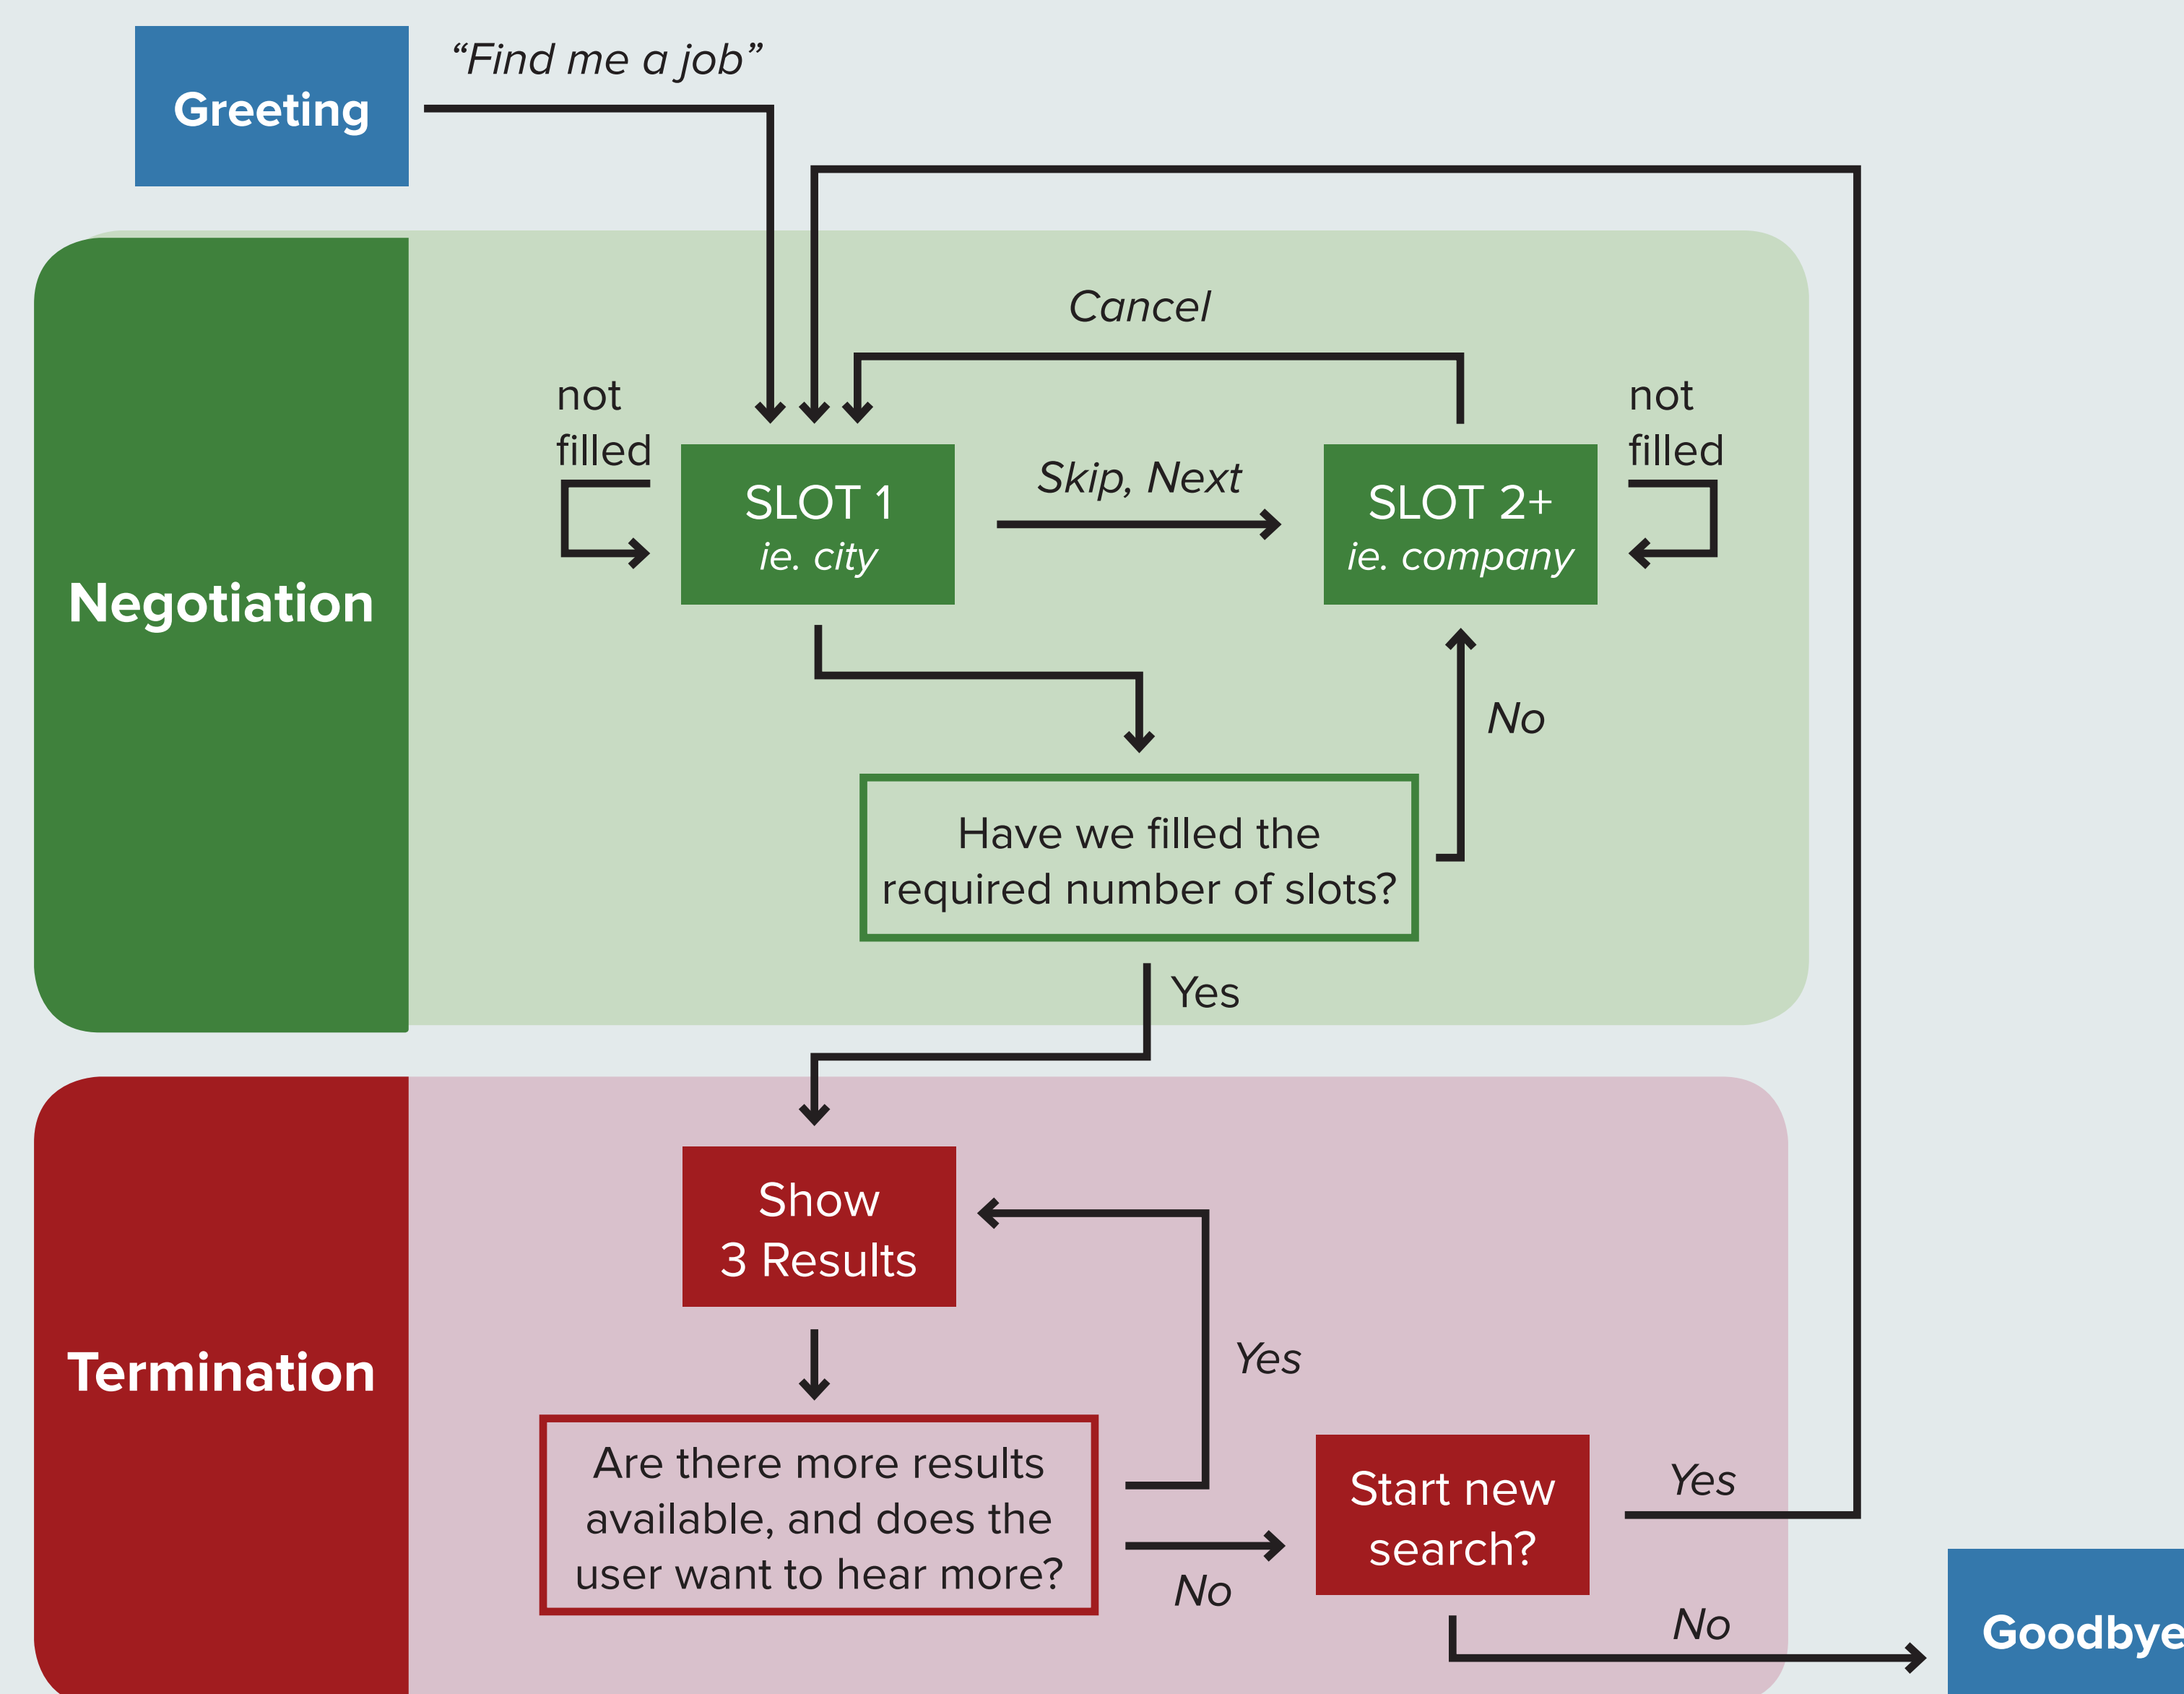
The core of the **Dialog Manager (DM)** is context management across turns. The Intent Router directs user utterances to either one of the Conversation Intent, Yes, No, Skip or Cancel intents. All intents attempt to fill context slots through context reprompts, elaboration or displaying options in accordance with the State Machine detail (below) and illicit the corresponding response using a rule-based Response Generator. The DM uses three primary database capabilities:

- 1) Validation of user utterances by context categories before actually adding the user choice to the context,
- 2) Retrieving exemplars that the Response Generator can leverage when crafting suggestions for a context slot, and
- 3) Generating job postings for user preferences for context slots



State Machine

The **State Machine (SM)** guides the goals of the Dialog Management as shown above. Greeting transitions into Negotiation state, which collaborates with the user to map specific preferences to available context slots like City, Company Name, Skill and so on. The Conversation Intent will attempt to fill a context slot. Other intents, like "Skip" or "Cancel" will shift the focus (working slot) as indicated. The current SM will sequentially transition through the available slots. Other transition strategies are supported by this design.



4 • Observations

- Alexa intent recognition algorithm is heavily dependent upon:
 - Number of words in the utterances
 - The utterance vocabulary
 - Number of samples in each slot
- Alexa ASR does not recognize sentence boundaries
- Alexa ASR does not perform well with named entity words that are commonly referred to with abbreviations, i.e., AT&T
- Alexa ASR does not support utterances over approximately 45 words
- Built-in intents handle disfluencies well when supplemented with sufficient sample utterance data, but does not handle backup and corrections, i.e., "I want to work in Portland uh I mean Seattle"
- A slot-filling approach is a naive but effective approach to building an information retrieval chatbot
- 'Amazon' cannot be recognized as a standalone utterance by Alexa
- Wake words such as 'Alexa' cannot be recognized in any utterance

5 • Future Work

We propose the following features & improvements towards building a better system:

Dialog Policy

- Transform system to mixed-initiative system
- Incorporate Markov processes to maintain information state
- Use reinforcement learning to evaluate dialog flow and guide it towards taking fewer turns; i.e., negative reward if number of turns is less than twice the number of context slots, and positive reward otherwise

NLU Robustness

- Allow user to fill multiple context slots in a single utterance
- Map user preferences using context slot clustering in database
- Use RNNs to identify user utterances and map them to intents

NLG Improvement

- Change current rule-based system to a neural sequence model

Entity Resolution

- Disambiguate between named entities

Additional Features

- Allow user to change existing context slot values at any point
- Send e-mail to user containing job application links

Suggested Evaluation Metrics

- Allow user to change existing slot values at any point
- % of search restarts
- % of turns per slot
- % of session "quits" without entering Termination state

6 • Acknowledgements

We'd like to thank **Hao Fang** for running the course, and our team members for their undeterred effort.

