

Data Curation, Crowdsourcing, Integrations

2/12/2020

Objectives

- Create effective crowdsourcing jobs to gather data.
- Learn about basic data curation needs
- Understand how to integrate with Cline

Crowdsourcing

Gathering Classification Data

Scenario

- Set up a scene
- Ask worker what would they say in this scenario
- Gathers a wider data set

Scenario-driven

Suppose you have an intelligent assistant that can handle _____. Given the following scenario, supply 5 creative ways of asking the intelligent assistant to help with your situation.

“ _____ ”

Paraphrase

- Provide examples
- Ask worker to provide different expression with similar meanings
- Gathers targeted data

Paraphrasing

Given the following sentence, supply 5 creative ways for rephrasing it and expressing the same meaning.

“ _____ ”

Classification Job Examples

Scenario

Scenario-driven

Suppose you have an intelligent assistant that can handle _____. Given the following scenario, supply 5 creative ways of asking the intelligent assistant to help with your situation.

“ _____
JOB PROMPT _____ ”

Generic

“You’re looking for a specific set of transactions in your bank account”

“You want to learn about your spending history”

Paraphrase

Paraphrasing

Given the following sentence, supply 5 creative ways for rephrasing it and expressing the same meaning.

“ _____
JOB PROMPT _____ ”

Specific

“Tell me about my transactions during the month of November”.

“What was my spending on food and drink?.”

Crowdsourcing Classification

The screenshot shows a web interface for adding a classification job. The title is "Add Classification Job". There are two tabs: "Default Settings" (selected) and "Advanced Settings".

Job Name: A text input field containing "flight_booking".

Worker Job Title: A text input field containing "Provide examples of booking a flight/travel plans".

Job Description: A text area containing the text: "Imagine you are talking to a travel agent about your upcoming vacation. You want to find any available flights that will fit your busy schedule. What would you say?".

Classification: A dropdown menu with "Select Intent" selected.

Examples: A list of three example utterances in rounded rectangular boxes:

- What flights are going to Toronto on the 5th? I don't care about price
- Can you show me all the Delta trips to Berlin that will arrive before Chr
- Get me the cheapest and quickest flight to San Diego right now please

Below the examples is a button labeled "+ Add Another Example".

Number of Utterances: A text input field containing "50".

Reward per Worker (Dollars): A text input field containing "0.20".

Tips

- Gather small data sets first
- Provide multiple examples
- Use diverse, creative examples
- Use different slots
- Track your jobs and iterate

Classification Data Curation

- **Example:** *“Pretend you need to transfer money from checking to savings, what would you say”*
 - Intent: transfer
 - Slots: to_account, from_account, amount
- **Common crowdsourcing problems:**
 - **Noisy data:** data that is not relevant
 - *“qwerty”*
 - *“Online bank transfers. Log in to your online account and select the option for making a payment.”*
 - **Incorrect data:** wrong intent, or not in-scope
 - *“Can I speak to someone in person?”*
 - *“What is my balance?”*
 - **Homogeneous data:** data that only contains relatively few unique words
 - *“Transfer from checking to savings”*
 - *“Transfer money from checking to savings”*
 - *“Transfer money from savings to checking”*

Fixing Homogeneous Data

- **Homogeneous data:** data that only contains relatively few unique words
 - *“Transfer from checking to savings”*
 - *“Transfer money from checking to savings”*
 - *“Transfer money from savings to checking”*
- Classifier will be “overfit” to certain words (bad)
 - Overfit: only queries with certain words will work
- This can be alleviated by:
 - **Using varying crowdsourcing prompts**
 - *“Pretend you need to transfer money from checking to savings, what would you say?”*
 - *“Pretend you want to move money between two of your bank accounts, what would you say?”*
 - **Using slot replacement**
 - checking → [checking, savings, money market, college savings, Premier account, ...]
 - **Using synonym replacement**
 - transfer → [transfer, move, ship, relocate, shift, siphon, ...]

Curating Classification Data

<input checked="" type="checkbox"/>	When is the earliest first class flight to London on the 24th of the month, please?
<input type="checkbox"/>	When is the earliest first class flight to London on the 24th of the month?
<input checked="" type="checkbox"/>	When is the earliest first class flight to London this Saturday?
<input type="checkbox"/>	When is the earliest first class flight to London next Friday?
<input checked="" type="checkbox"/>	When is the earliest first class flight to London on New Year's Eve?
<input checked="" type="checkbox"/>	Are there any flights to Paris available next week?
<input checked="" type="checkbox"/>	I need to get to Berlin tomorrow. Can you find me any flights?
<input checked="" type="checkbox"/>	It's an emergency to get to Tampa tonight- I don't care about the seat or price. What can you find for me?
<input checked="" type="checkbox"/>	I need to go to San Francisco on Friday, are there any seats at all available?
<input checked="" type="checkbox"/>	Can you find me any flights to Detroit for tomorrow night?
<input checked="" type="checkbox"/>	I want a cheap and quick flight to San Diego
<input checked="" type="checkbox"/>	I want the cheapest flight to San Digo
<input checked="" type="checkbox"/>	I would love a very cheap flight to San Diego
<input type="checkbox"/>	I wanna get a very cheap flight to San Diego
<input checked="" type="checkbox"/>	I would like one cheap flight to San Diego
<input checked="" type="checkbox"/>	When is the earliest flight to Chicago on the 15th of the month, please?

General Concepts

- Ensure that all utterances are within competency scope
- Do **not** delete typos, slang, incorrect grammar, etc.
- Delete copy cat utterances
 - “Is this teaching the AI something new?”
- Consider the AI’s response --- does the training data cover all avenues to AI’s response?

Classification

Uniqueness Sorting Tool

Uniqueness Sorting Tool helps you identify irrelevant utterances in classification data.

For most unique utterances:

- Are there irrelevant data?
- Are there underrepresented data (i.e. specific keywords) we need more of?

For least unique:

- Are there phrases being overfitted?

The screenshot shows the 'account_transfer_start' interface. At the top right are 'Import', 'Export', and 'Save' buttons. Below that are 'Select' and 'Crowdsource' buttons, and a search bar. A dropdown menu is set to 'account_transfer_...' and there is an 'Add' button. The main area is titled 'UTTERANCES' and has a 'SORT BY' dropdown set to 'Uniqueness ↓'. A table lists 12 utterances with their uniqueness scores and delete buttons. The first row, 'Is this FDIC insured' with a score of 1.88, is highlighted with a red border.

utterance	score
account_transfer_star ↓ Is this FDIC insured	1.88
account_transfer_star ↓ Where is your ATM?	1.20
account_transfer_star ↓ Move 40 to MMA	1.18
account_transfer_star ↓ When was my last transfer?	1.12
account_transfer_star ↓ Put 10 in college fund	0.96
account_transfer_star ↓ Take 59.75 from checking and pay my cell phone bill	0.87
account_transfer_star ↓ I want to transfer 10k	0.84
account_transfer_star ↓ Transfer \$789.98 to auto loan account	0.82
account_transfer_star ↓ Transfer 10000 from roth to ira	0.81
account_transfer_star ↓ Move \$2,798 to mortgage account	0.80

Uniqueness

- The uniqueness sort tool works by first mapping each query to a high dimensional semantic embedding space. A mean embedding representation is then computed, and finally the tool sorts each query by distance from the mean embedding.
- A query's uniqueness score is the query's distance to the mean embedding.
- Uniqueness sort draws attention to the most unique utterances so that you can identify weak spots in data or potentially mislabeled or out of scope data

Uniqueness Sort Tool

- Uniqueness Sort (or anomaly detection) is a data insight tool that Clinc provides to analyze **classifier** data to identify errors or underrepresented training samples
 - Errors are training samples that are either mislabeled or unnecessary noise
 - Underrepresented training samples are good but unique compared to the rest of the samples

The screenshot displays the Clinc interface for the 'uniqueness_sort_demo' project. The main panel shows a list of utterances for the 'my_balance_start' state, sorted by 'Uniqueness'. The utterances and their uniqueness scores are:

Utterance	Uniqueness Score
how do asdf	0.21
How do I reset my PIN number	0.14
speak back checking account cash amount	0.13
balance please	0.11
I'd like to know my savings balance	0.11
balance I have in bank is what	0.00

Annotations on the screenshot:

- Red lines point to the utterances 'how do asdf' and 'How do I reset my PIN number', which are labeled as 'Errors'.
- A blue line points to the utterance 'balance I have in bank is what', which is labeled as 'Underrepresented'.

Accessing Uniqueness Tool

- Navigate to the intent classifier page
- Select Sort By -> Uniqueness

The screenshot displays the Uniqueness Tool interface for the 'my_balance_start' intent. The left sidebar shows the project structure, including 'uniqueness_sort_demo' and 'my_balance'. The main area shows the intent name 'my_balance_start' and a list of utterances. A 'SORT BY' dropdown menu is open, showing 'Recently Added' and 'Uniqueness' options. The 'Uniqueness' option is highlighted, indicating it has been selected. The right sidebar shows the 'my_balance' intent details, including 'STATES', 'UTTERANCES', 'SLOTS', and 'RESPONSES'.

uniqueness_sort_demo
AI Version

my_balance_start

Train Import Export Save

Select Internal Crowdsourcing Crowdsourcing

my_balance_start Add utterances of what a user might say... Add

UTTERANCES

my_balance_start	my bank balance is how much	×
my_balance_start	balance i have in bank is what	×
my_balance_start	balance of checking account	×
my_balance_start	balance please	×
my_balance_start	can you tell me my balance	×
my_balance_start	how much cash do i have in my checking	×

SORT BY Recently Added

my_balance

STATES

UTTERANCES

root

my_balance

TRANSITIONS INCOMING

root

my_balance_start 21

SLOTS

UTTERANCES

You don't have any slots yet.

Template Slots

RESPONSES

DEFAULT Edit


Recap: Classification Data Curation

- Data curation is an iterative process!
 - Gather data
 - Curate
 - Test
- Use fine-grained intent labels
- Test your classification first! Similar to incremental development, you want to test what you build in as granular chunks as possible. You want to test to make sure that your classifications work as expected before doing too much more additional work to minimize potential re-work.

Front End Integrations

What is a Cline Integration?

Enabling users to experience Cline's rich conversational experiences and sophisticated NLU across various channels and devices



How do Apps Integrate with Clinc?

Configuration

Configuration data for the app itself and Clinc's APIs

Authentication

Methods and settings for linking external accounts and authenticating requests to Clinc's APIs

Query Endpoint Connection

Sending and receiving HTTP requests to Clinc's APIs that use the AI version and Business logic

App Configuration

Clinic Configuration

- UUID of the AI Version being used
- Credentials to request auth tokens
 - Username
 - Password

Integration Specific Configuration

- Application set up
 - Example: Creating an Alexa Skill or Web app
 - Any configuration needed by the application itself.

App Authentication

Clinic Authentication

- Clinic's APIs require Bearer Tokens be set in the request header
- To obtain a token, HTTP POST to `/v1/oauth` with credentials
- Credentials set in body

External Authentication

- External user accounts can be linked in the integration by passing auth tokens in the body to the query endpoint

HTTP POST `/v1/oauth`

Header

```
Content-Type: application/x-www-form-urlencoded
```

Body

```
username : bob  
password : $3kr3tp@$w0rd  
grant_type : password
```



Response

```
{  
  "access_token": "acdefghijklmnopqrstuvwxyz12345",  
  "expires_in": 36000,  
  "token_type": "Bearer",  
  "scope": "user_read user_write query",  
  "refresh_token": "12345678901234567890123456789"  
}
```

Query Endpoint Requests

Query Endpoint Requests

- To use the Cline API, HTTP POST to `/v1/query`
- Set Bearer token in the header
- **Key items in request body:**
 - `query`: User's text query for processing by the AI and Business Logic (BL)
 - `ai_version`: UUID of the AI version being used
 - `dialog`: Optional token used to maintain conversation state
 - `token`: Optional token used to link external accounts - can be used in BL

HTTP POST `/v1/query`

Header

```
Content-Type: application/json
Authorization: Bearer abcd1234abcd1234abcd1234abcd1234
```

Body

```
query: "Can I transfer $100 from savings to checking?"
ai_version: "123abc123abc123abc123abc123abc123abc"
lat: 42.2730207
lon: -83.7517747
time_offset: 300
dialog: "abcdefg123456789"
device: "alexa"
```

Query Endpoint Response

slots

- Object of keys and values collected by the conversation
- Useful when application needs neat parseable data that has been gathered from AI Version and Business Logic.

dialog

- Token to maintain conversation state
- Send **dialog** token in the body of request to the Query Endpoint to preserve data in the conversation so far.

```
"slots":{
  "_ACCOUNT_TRANSFER_": {
    "type": "string",
    "values": [
      {
        "amount": "$50",
        "src_account": "savings",
        "src_account": "checking",
        "confirmed": 1
      }
    ]
  },
  "_USER_": {
    "type": "string",
    "values": [
      {
        "first_name": "Peter",
        "last_name": "Gibbons",
        "member_id": "123456789"
      }
    ]
  }
}
```

Query Endpoint Response Continued

`classifier_state`

- A state is a node on the state graph (seen in the platform) which can be classified into given an intent
- Useful when application needs information on what the user is doing
- Example:
 - User opens the application and says “Hello”
 - `classifier_state` is `greeting`
 - The application could display a “Welcome” page or animation

Query Endpoint Response Continued

`speakableResponse`

- Text response suited for a voice experience
- Value is passed to Text to Speech (TTS) services

`formattedResponse`:

- Text response for displaying and reading in the application
- Useful when the response is too long and bulky to be heard aloud or contains sensitive data that might not want to be broadcasted

Other Considerations & Example Project

Automatic Speech Recognition (ASR) and Text to Speech (TTS)

- Some integrations, like web app integrations require external libraries and APIs for TTS and ASR for a voice first experience.
- WebKit Speech, Google Cloud Speech, Responsive Voice

Example Project

- Alice wants to build a Clinic Integrated Web Application
 - ASR
 - TTS
 - Authentication
 - Query Endpoint

Alice's Web Integration

