Tip of Your Tongue: Methods for an Effective Reverse Dictionary Model

Arman Aydin, Cem Gokmen

Department of Computer Science, Stanford University

Stanford Computer Science

Problem

- A **reverse dictionary** is a tool that finds a suitable target word given a rough description or definition, with the goal of helping the **"tip of your tongue" problem** where one needs to use a word in their writing but cannot remember it.
- It is known that pre-trained Transformer models like **BART** and **T5** can be fine-tuned to perform well on **sequence-to-sequence tasks** like reverse dictionary.
- However, no "reverse dictionary" dataset exists, and while such datasets can be generated by simply switching the input/output pairs in a real dictionary, this causes **subpar performance** due to the **large divergence** between the distribution of such a dataset (formal definitions) and the kinds of prompts users might give a reverse dictionary.
- The goal of this project is to investigate data augmentation and sequence generation methods that can improve the accuracy of a reverse dictionary model.

Background

- LSTM models, which can learn contextual meaning in the input phrases through an encoder-decoder architecture, have been used as a method for **reverse dictionary tasks**.
- **Wiktionary** is a multilingual, web-based project to create a free dictionary of terms and their definitions.
- **WordNet** is a lexical database that links words into semantic relations.

Methods

- **Model Architecture**: For our reverse dictionary model, we propose Transformer-based models as our foundation. These models have several advantages:
  - they are known to work well even for much more complicated sequence-to-sequence tasks (e.g. summarization)
  - pretrained weights are readily available
- **Data Augmentation using Metadata:** Beyond simply using the word:definition pairs from the dictionary, we generate additional queries using synonym, antonym, part-of-speech metadata.
- **Better sampling using Hamming Diversity:** Being able to generate diverse predictions rather than similar forms of the same senses is critical. To this end, we use the Hamming Diversity Penalty generation method where beams are separated into groups during search and each group is penalized if its beams contain similar tokens to other groups.

Experiment Setup

- Training was performed on **6 Nvidia RTX 2080 Ti GPUs** on pre-trained HuggingFace transformer model **"facebook/bart-base"** after comparing to other models such as different T5 versions and picking this one as a good size/performance tradeoff.
- Test dataset contains **human-generated queries** for words **randomly picked from a dictionary**.
- Top-K accuracy metrics (top-10 unless otherwise noted) are used for evaluation.

Experiment Results

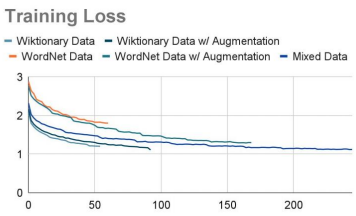

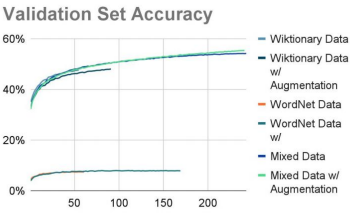

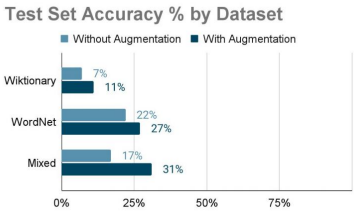

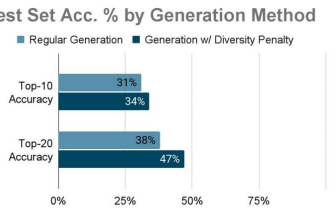

Key Findings

- **Without the use of data augmentation and diverse generation methods**, model top-10 accuracy remains extremely low at **7%**.
- Adding data from the **WordNet dataset** increases the accuracy to a still-low **17%**.
- Applying additional **data augmentation in the form of synonym, antonym, and part-of-speech queries** increases accuracy to **31%**.
- Applying **generation with diversity penalty** increases top-10 accuracy to a further **34%**, and shows an even more drastic improvement in **top-20 accuracy to 47%**.
- Altogether, combining these methods allows us to train a reverse dictionary model that can be used effectively on in-the-wild user queries.

References

- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, & Luke Zettlemoyer. (2019). BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, & Alexander M. Rush. (2020). HuggingFace's Transformers: State-of-the-art Natural Language Processing.
- McCrae, J., Rademaker, A., Rudnicka, E., & Bond, F. (2020). English WordNet 2020: Improving and Extending a WordNet for English using an Open-Source Methodology. In Proceedings of the LREC 2020 Workshop on Multimodal Wordnets (MMW2020) (pp. 14–19). The European Language Resources Association (ELRA).