

Towards a General-Purpose Linguistic Annotation Backend

Graham Neubig[†], Patrick Littell[‡], Chian-Yu Chen[†], Jean Lee[†],
Zirui Li[†], Yu-Hsiang Lin[†], Yuyan Zhang[†]

[†]Language Technologies Institute, Carnegie Mellon University

[‡]National Research Council Canada

gneubig@cs.cmu.edu

1 Introduction

Language documentation is inherently a time-intensive process; transcription, glossing, and corpus management consume a significant portion of documentary linguists' work. Advances in natural language processing can help to accelerate this work, using the linguists' past decisions as training material, but questions remain about how to prioritize human involvement.

In this extended abstract, we describe the beginnings of a new project that will attempt to ease this language documentation process through the use of natural language processing (NLP) technology. It is based on (1) methods to adapt NLP tools to new languages, based on recent advances in massively multilingual neural networks, and (2) backend APIs and interfaces that allow linguists to upload their data (§2). We then describe our current progress on two fronts: automatic phoneme transcription, and glossing (§3). Finally, we briefly describe our future directions (§4).

2 Overall Framework

The final goal of our project is to create a linguistic annotation backend (LAB), that will take in raw or partially annotated linguistic data, and provide annotation candidates for a linguist (or other interested party) to peruse. Candidates for the types of services to provide are automatic phoneme transcription (Adams et al., 2018; Michaud et al., 2018), speech-to-text alignment (Johnson et al., 2018), word segmentation (Peng et al., 2004; Goldwater et al., 2009), morphological analysis (Yarowsky and Wicentowski, 2000), syntactic analysis (Nivre, 2005), automatic glossing (Riding, 2008), or linguistic typology prediction (Daume III and Campbell, 2007). The LAB will be hosted on a server and exposed as an API that can be linked to popular annotation software

such as ELAN¹ or FLEX.²

The obvious difficulty in creating such an interface is data scarcity in the languages in question. In order to overcome these barriers, we plan to take advantage of recent advances in NLP that allow for multilingual modeling (Täckström et al., 2012; Johnson et al., 2016) and multi-task learning (Caruana, 1997), which allow models to be trained with very little, or even no data in the target language (Neubig and Hu, 2018). We also plan to utilize active learning (Settles, 2009), which specifically asks the linguists to focus on particular examples to maximize the effect of linguists' limited time when working with field data. While there is still no alternative to significant human engagement when processing data, many of the decisions a linguist is faced with when transcribing, glossing, organizing, or searching a corpus are relatively rote – decisions that could be deducible from past decisions or from similar languages.

3 Current Progress: A Backend/Interface for Automatic Phoneme Transcription and Glossing

As first steps towards realizing our final goal, we have currently developed a backend for two tasks (automatic phoneme transcription and glossing), which is integrated with a simple example interface.

3.1 Backend Overview

The current LAB is based on a simple three-step process:

Data Upload The linguist uploads any existing annotated data to the interface.

¹<https://tla.mpi.nl/tools/tla-tools/elan/>

²<https://software.sil.org/fieldworks/>

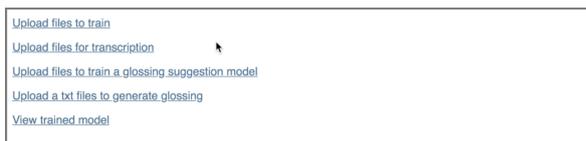


Figure 1: The prototype of the automated interface supporting transcription and glossing.



Figure 2: Uploading the training data to the automated interface for transcription model training.

Model Training A model is trained to process this data. This training could potentially utilize other data sources for multilingual and multi-task training.

Data Annotation The linguist uploads unannotated data, and the trained model proposes annotations for the linguist to accept or edit.

An example of an overall interface exposing this functionality for the currently implemented tasks of transcription and glossing is shown in Figure 1.

3.2 Phoneme Recognition

The automatic phoneme transcription component provides an interactive online interface for users to manage speech recognition models and transcribe speech recordings. The speech recognition model can be any one of the user’s choosing as long as it supports the API. In our current system, we use Persephone (Adams et al., 2018) as our transcription backend, which is designed for low-resource language transcription. Through the API, the users can upload a batch of speech recordings along with the corresponding transcriptions as the training data to train a transcription model tailored to the language and speakers of their interest. The system is equipped with some default model and training configurations so that the users are not required to have expert knowledge of the transcription model and training. The model obtained from each training session will then be stored for later use. Figures 2 to 4 show the work flow of training a transcription model.

The users can upload speech recordings they want to transcribe to the interface, and perform the automatic transcription using previously trained models (Figures 5 and 6). The interface shows the transcription results to the users, and the users can



Figure 3: Training a transcription model using the training data uploaded by the user.



Figure 4: Training a transcription model using the training data uploaded by the user.

optionally edit the transcription results to fix errors or make model improvements (Figure 7). The refined transcriptions can then be downloaded by the users. If the user’s data privacy preferences allow, the system can also collect them along with the original speech recordings as extra training data to further fine tune the model.

3.3 Automatic Glossing

The interface also supports making glossing suggestions. Glosses are generated word-by-word with Moses (Koehn et al., 2007), a statistical machine translation system. The system takes parallel data as input, which could be either the language and translations, or the language and glosses. Using this parallel data, we learn a word alignment with a statistical model, specifically the IBM alignment models (Brown et al., 1993) as implemented in GIZA++ (Och and Ney, 2003). Then we perform phrase extraction (Koehn, 2010), which gives us a translation probability distribution for each word or phrase in the combined corpus. We then display translations with high probability as glossing suggestions. An example of how the automatic glossing suggestion works on the interface can be seen in Fig. 8.

4 Future Plans

Working with field data is highly rewarding, but on a moment-to-moment basis the work is not usually particularly *engaging*; most of the individual decision events that a linguist makes during field corpus creation do not fully engage their reasoning capacity. Our goal is to maximize the effects of human engagement with data by maximizing the time the linguist spends on interesting and relevant



Figure 5: Uploading the speech recordings to transcribe.



Figure 6: Transcribing the speech recordings using the model previously trained.

decisions. We intend to explore this question with respect to both low-level decisions (“What word was said here?”) and high-level decisions (“These utterances exemplify ergativity in this language; are there other examples in this corpus?”). Our future work towards this goal will take a three-pronged approach: developing a general-purpose linguistic annotation API and integrating it with popular annotation frameworks, developing new methods to perform multi-lingual and multi-task learning to train effective models even in a paucity of training data, and working with linguists to help refine and prioritize our work in these areas. In particular, for the third goal we are actively seeking collaborators who would be interested in testing and giving advice about the utility of the proposed approach.

Acknowledgements

We thank Alexis Michaud for his useful comments and help in preparation of data, Oliver Adams for his help with Persephone, and Antonis Anastasopoulos for helping us access and prepare the Griko data. This material is based upon work supported by the National Science Foundation under Grant No. 1761548.

References

Oliver Adams, Trevor Cohn, Graham Neubig, Steven Bird, and Alexis Michaud. 2018. Evaluating phonemic transcription of low-resource tonal languages for language documentation. In *Language Resources and Evaluation Conference (LREC)*, Miyazaki, Japan.

Antonios Anastasopoulos, Marika Lekakou, Josep Quer, Eleni Zimianiti, Justin DeBenedetto, and David Chiang. 2018. Part-of-speech tagging on an endangered language: a parallel Griko-Italian resource. In *Proc. COLING*.

Peter F. Brown, Vincent J. Della Pietra, Stephen A. Della Pietra, and Robert L. Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19:263–312.

Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.

Hal Daume III and Lyle Campbell. 2007. A Bayesian model for discovering typological implications. In *Proc. ACL*, pages 65–72. Association for Computational Linguistics.

Sharon Goldwater, Thomas L. Griffiths, and Mark Johnson. 2009. A Bayesian framework for word segmentation: Exploring the effects of context. *Cognition*, 112(1).

Lisa M Johnson, Marianna Di Paolo, and Adrian Bell. 2018. Forced alignment for understudied language varieties: Testing prosodylab-aligner with tongan data. *Language Documentation and Conservation*.

Melvin Johnson et al. 2016. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *arXiv preprint arXiv:1611.04558*.

Philipp Koehn. 2010. *Statistical Machine Translation*. Cambridge Press.

Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. ACL*, pages 177–180.

Alexis Michaud, Oliver Adams, Trevor Anthony Cohn, Graham Neubig, and Sverine Guillaume. 2018. Integrating automatic transcription into the language documentation workflow: Experiments with na data and the persephone toolkit. *Language Documentation and Conservation*.

Graham Neubig and Junjie Hu. 2018. Rapid adaptation of neural machine translation to new languages. In *Proc. EMNLP*, Brussels, Belgium.

Joakim Nivre. 2005. Dependency grammar and dependency parsing. *MSI report*, 5133(1959):1–32.

Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.

Fuchun Peng, Fangfang Feng, and Andrew McCallum. 2004. Chinese segmentation and new word detection using conditional random fields. In *Proc. COLING*.

Jon D Riding. 2008. Statistical glossing, language independent analysis in bible translation. *Translating and the Computer*, 30.

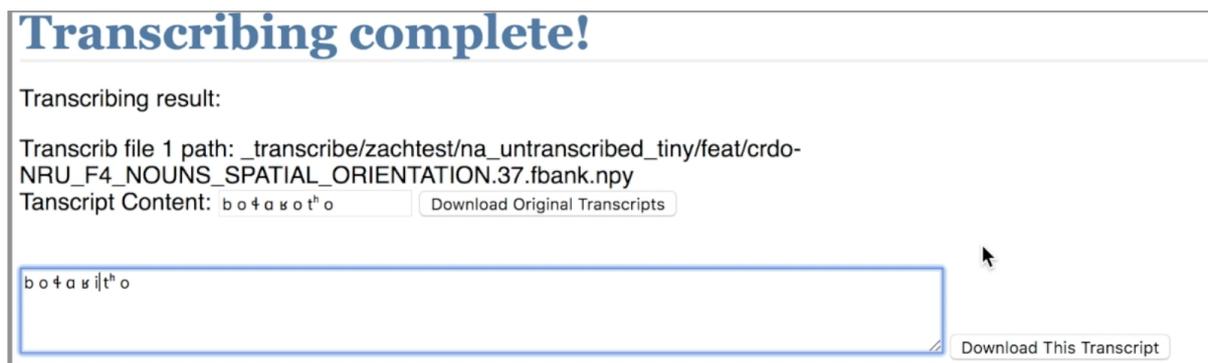


Figure 7: The users can examine the transcription results and optionally edit the results to correct errors.

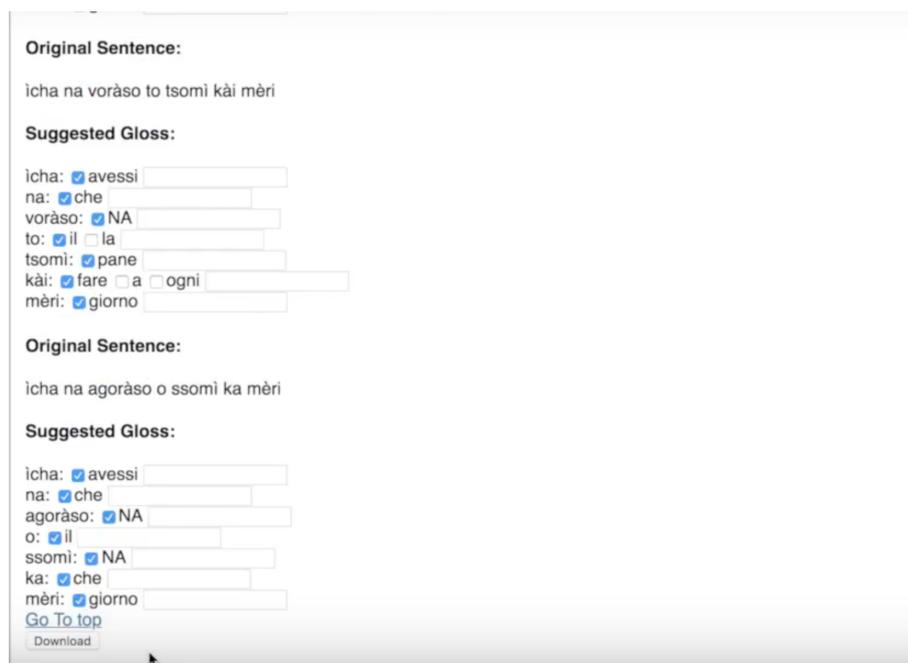


Figure 8: An example of generated glosses, for Griko language data from Anastasopoulos et al. (2018).

Burr Settles. 2009. Active learning literature survey. Computer Sciences 1648, University of Wisconsin–Madison.

Oscar Täckström, Ryan McDonald, and Jakob Uszkoreit. 2012. Cross-lingual word clusters for direct transfer of linguistic structure. In *Proc. NAACL*, pages 477–487, Montréal, Canada. Association for Computational Linguistics.

David Yarowsky and Richard Wicentowski. 2000. Minimally supervised morphological analysis by multimodal alignment. In *Proc. ACL*.