We are at an especially opportune time in the history of the study of human language. Brain imaging technology allows us to directly observe and model brain activity associated with human language. Techniques from statistics and machine learning allow us to construct quantitative computational models that describe these brain activities. Furthermore, they allow us to construct mental state decoders that can accurately predict certain aspects of thought from measured brain activity (Mitchell et al., 2008; Chang et al., 2010). In addition to the scientific impact of better understanding the representation and processing of human language, this research will lead to many applications and broad impact. For example, a brain-computer interface (BCI) device that could decode internal speech could enable locked-in patients to communicate.

I am particularly interested in computational neurolinguistics, an emerging research area which integrates recent advances in computational linguistics and cognitive neuroscience, with the objective of developing cognitively plausible models of language and gaining a better understanding of the human language system. I have been instrumental in the development of computational neurolinguistics, co-organizing the first two workshops in the area (NAACL-HLT 2010 workshop on computational neurolinguistics; NIPS 2011 workshop on interpretable decoding of higher cognitive states from neural data), and am interested to continue my efforts. In particular, the field requires techniques that are capable of taking advantage of temporally coordinated activity in the brain. To take language comprehension as an example, speech is received at 3-4 words per second; acoustic, semantic and syntactic processing can occur in parallel. Two non-invasive brain imaging methods that are sensitive to the fine-grained temporal patterns of multiple processes are electroencephalography (EEG) and functional near-infrared spectroscopy (fNIRS). EEG uses electrodes on the surface of the scalp to measure the voltage signal arising from large areas of coordinated neural activity, whereas fNIRS uses light to measure changes in cerebral blood oxygenation associated with brain activity. For future goals, I am interested to use portable brain imaging devices to help improve intelligent tutoring systems and decode human speech.

The ultimate automated tutor could peer directly into students' minds to identify their mental states (e.g. engagement, competencies, and intentions) and decide accordingly what and how to teach at each moment. Recent advances in brain imaging technologies have brought upon several portable EEG headsets that are commercially-available and simple enough to use in schools (NeuroSky, 2011; Emotiv, 2010; BCInet, 2010). Using EEG signals recorded from adults and children reading text and isolated words, both aloud and silently, we train and test classifiers to tell if students are reading easy or hard sentences, and to distinguish among easy words, hard words, pseudo-words, and unpronounceable strings (Mostow, Chang & Nelson, 2011). Better-than-chance performance shows promise for tutors to use EEG at school. This development makes it feasible to record longitudinal EEG data in authentic school settings. I will co-author and co-PI a three-year, NSF Cyberlearning: Tranforming Education grant that the program officer has informally notified us she intends to fund (Design and Implementation Projects: Exploiting Longitudinal EEG Input in a Reading Tutor). I plan to extend my previous work on student modeling (Chang et al., 2006) and use the longitudinal EEG signals as input to dynamic Bayes nets to trace student's knowledge across different skills. An EEG-based student model allows direct assessment of student's competencies that can be administered unobtrusively and in a timely fashion.

One ambitious project that I want to work toward is to develop neural accounts of sentence processing. In my thesis work, I have elucidated how the distributed patterns of brain activity encode the meaning of nouns, adjective-noun and noun-noun phrases. The next step is to collect brain activity for broader range of words with different syntactic categories (e.g. subjects, objects) and part-of-speech tags (e.g. verbs). This will enable thought-to-text systems that can identify the basic subject-verb sentence structure that are akin to automatic speech recognizer (ASR). Jou and Schultz (2009) showed that it is possible to use surface electromyographic (EMG) signals that are generated by the human articulatory muscles to build a phone-based ASR in a 100 word vocabulary recognition task. Consequently, I am motivated to see if portable brain imaging devices such as EEG and fNIRS can also be used to help decode human speech. One direction is to derive "neural features" that are akin to the acoustic features derived from speech input. Alternatively, we can use the brain activity to help improve language models. More specifically, ASR benefits from an accurate language model. One challenge faced by ASR is that human dialogues change in topics - people may talk about politics in one moment and sports in the next. A brain imaging device that can distinguish topics of conversation can help the ASR to pick the right language model and produce better recognition results.

Finally, I would like to take an active role in engaging and assisting other LTI colleagues to utilize brain imaging methods. The cognitive plausibility of linguistic models has primarily been evaluated against collections of subjective intuitions (e.g. semantic feature norms, grammaticality judgments). Recent improvement in function brain imaging has made it possible, for the first time, to ground linguistic theories by the patterns of brain activity – the implementational level of Marr's three levels of computational modeling. For example, Chang et al. (2009) provided a neural account of how people use adjectives to modify the meaning of the noun. I look forward to further brain imaging studies shedding new light on the nature of human language.

References

BCInet, 2010. Retrieved from http://www.bcinet.com/.

Chang, K. M., Beck, J. E., Mostow, J., Corbett, A., 2006. A Bayes Net Toolkit for Student Modeling in Intelligent Tutoring Systems. *Proceedings of the 8th International Conference on Intelligent Tutoring Systems*, Jhongli, Taiwan, 104-113.

Chang, K. M., Cherkassky, V. L., Mitchell, T. M., Just M. A., 2009. Quantitative modeling of the neural representation of phrases: How vector-based models of semantic composition can account for fMRI activation. *Proceedings of the 47th Annual Meeting of the ACL and the 4th IJCNLP of the AFNLP*, Suntec, Singapore, 638-646.

Chang, K. M., Mitchell, T. M., Just M. A., 2010. Quantitative modeling of the neural representations of objects: How semantic feature norms can account for fMRI activation. *NeuroImage* 56, 716-727.

Emotiv, 2010. Retrieved from http://www.emotiv.com/.

Jou, S. C., Schultz, T., 2009. Automatic Speech Recognition Based on Electromyographic Biosignals. *Biomedical Engineering Systems and Technologies: Communications in Computer and Information Science* 25 (3), 305-320.

Mitchell, T., Shinkareva, S.V., Carlson, A., Chang, K.M., Malave, V.L., Mason, R.A., Just, M.A., 2008. Predicting human brain activity associated with the meanings of nouns. *Science* 320, 1191-1195.

Mostow, J., Chang, K. M., Nelson, J., 2011. Toward Exploiting EEG Input in a Reading Tutor. *Proceedings of the 15th International Conference on Artificial Intelligence in Education*, Auckland, New Zealand, 230-237. [Nominated for Best Paper].

NeuroSky, 2011. Retrieved from http://www.neurosky.com/.