TongAI: Helping Neuroradiologists Do Better Things

Jose Miguel Giron
jmgiron@stanford.edu

Nikita Demir
ndemir@stanford.edu

Motivation

- Neuroradiologists spend a large portion of their time recommending an imaging protocol based on a doctor's patient description.
- In conjunction with the Stanford Hospital, we compile a novel dataset and implement state of the art Deep Learning NLP techniques to automatically make the protocol recommendation.
- We find that a combination of FastText embeddings and a custom FastText model variant provide great overall results.
- Our research indicates that this is the first time someone has attempted to automate this task.

Problem Definition

Problem: Text Classification

- Input: Patient description, age, gender
- Output: Brain Protocol Classification (11 classes)

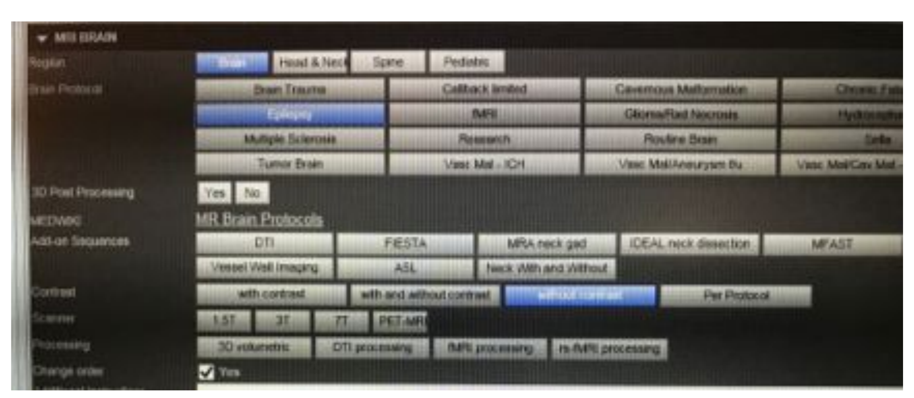

Related Works

- [1] Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. Transactions of the Association for Computational Linguistics, 5:135–146.
- [2] Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T. (2016). Bag of tricks for efficient text classification. arXiv preprint arXiv:1607.01759.
- [3] Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In Adv. NIPS
- [4]Wang Y, Liu S, Afzal N, Rastegar-Mojarad M, Wang L, Shen F, et al.(2018) A Comparison of Word Embeddings for the Biomedical Natural Language Processing.
- [5]Dubois S, Romano N, Kale, D, Shah N, Jung K. (2018) Efficient Representations of Clinical Text
- [6]Yoon Kim. (2014) Convolutional Neural Networks for Sentence Classification

Dataset

Age	Gender	Reason for imaging	Protocol
86	f	Progressive dementia	routine brain
49	f	49 yo female with h/o pituitary mass	sella
69	f	Surgical planning. Please include t1 and t2 fiducials	stereo/cyberknife
69	m	Lung mass evaluate for metastatic disease	mets
47	m	Follow up for ms	Multiple sclerosis

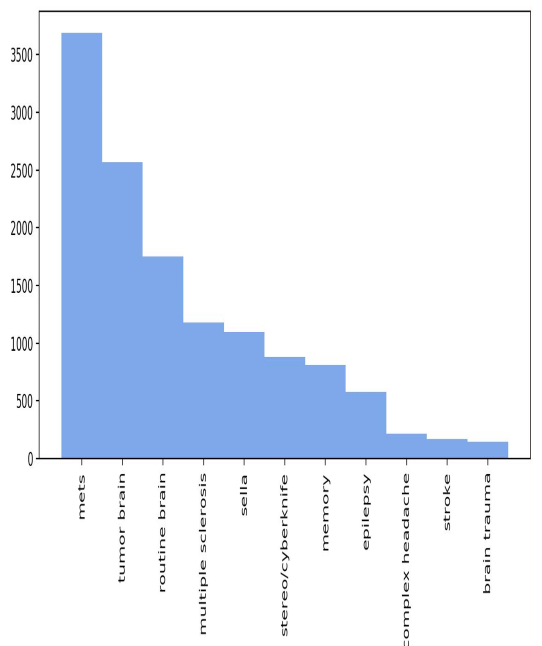

Class Distributions

- We see differences in the distribution of protoc certain classes.

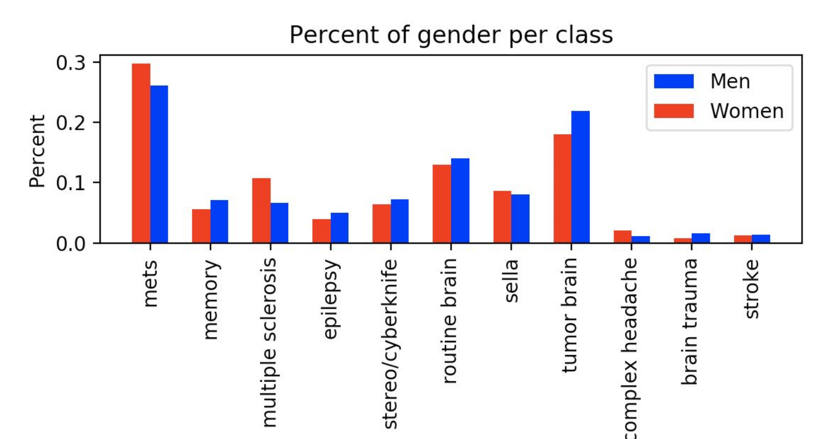

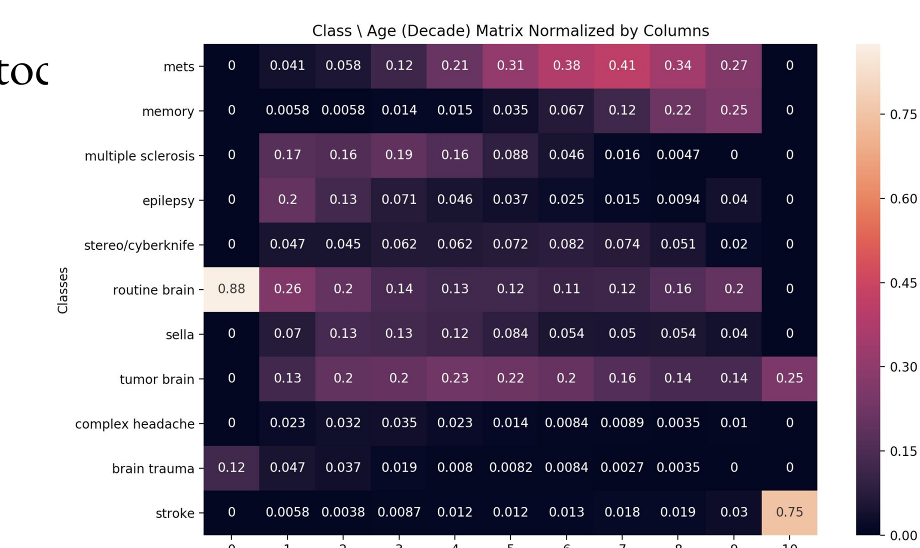

Model

Naive Bayes

- Powerful text classification baseline
- Bayes rule + naive conditional independence assumption

FastText

- 300-dim FastText Embeddings (pretrained)
- 3x min-mean-max document representation
- Concatenated with age and gender one-hot vectors
- Single Fully-Connected layer

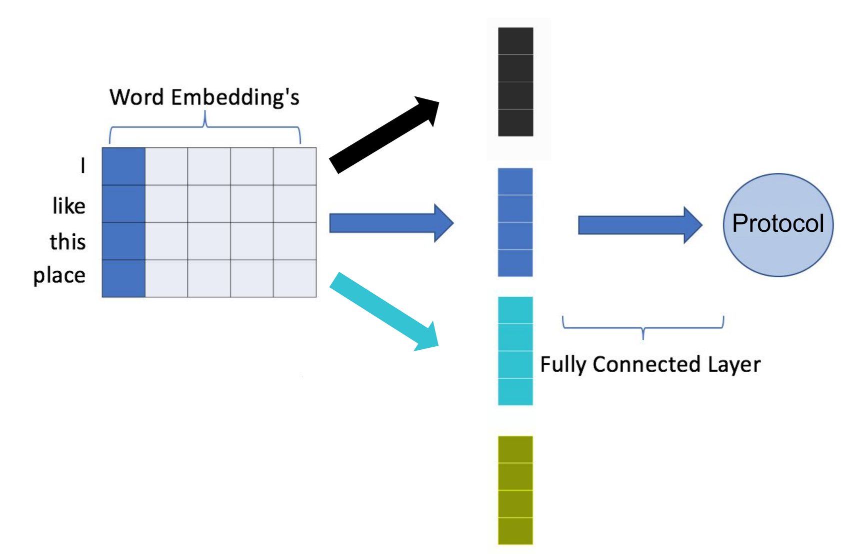

Results and Analysis

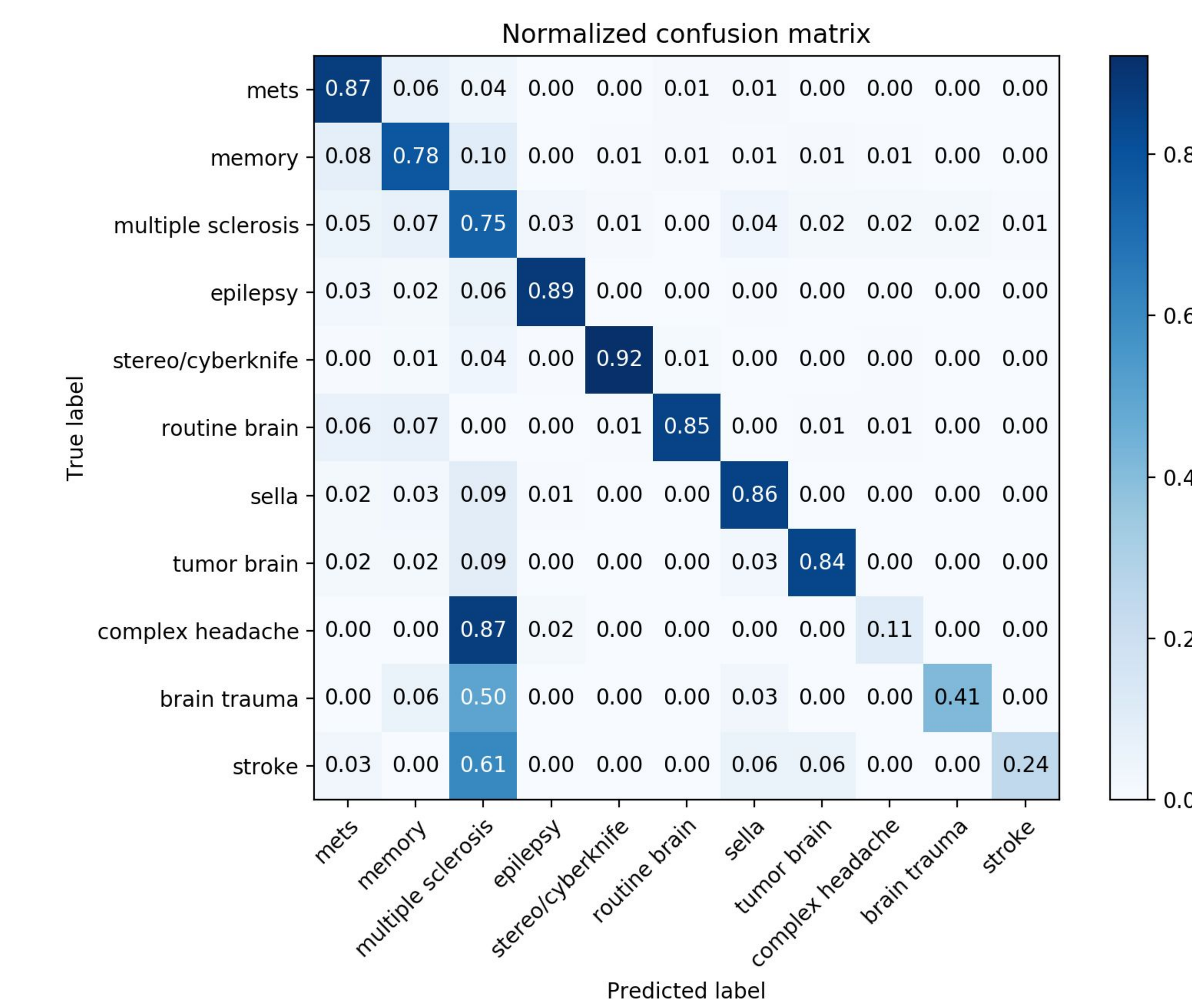

Model	Accuracy	Precision	Recall	F-Score
Random	15.45%	0.170	0.155	0.155
Age	31.79%	0.195	0.318	0.229
Gender	28.35%	0.089	0.284	0.134
Naive Bayes	76.73%	0.777	0.767	0.760
CNN	64.20%	0.642	0.639	0.633
Vanilla FastText	78.60%	0.790	0.786	0.778
TongAI	**82.06%**	**0.829**	**0.821**	**0.816**

Future Work

- Clean up dataset to find new examples for brain
- Create datasets for other regions
- Train a the model on more regions
- Provide contrast recommendations