ReGAL: Rule Generative Active Learning for Model-in-the-Loop Weak Supervision

David Kartchner, Wendi Ren, Davi Nakajima An, Chao Zhang & Cassie Mitchell
Georgia Institute of Technology, Atlanta, GA

Introduction

- Labeling is a main cost in extending ML to new tasks, especially low resource domains (health, legal, etc.)
- *Weak supervision* uses labeling rules to create scalably create noisy labels.
- *Active learning* creates a diverse labeled set by labeling the most relevant data points
- **ReGAL** performs active learning over model-generated labeling rules

Contributions

- Method for creating model-generated labeling functions
- Framework for active learning on Rules

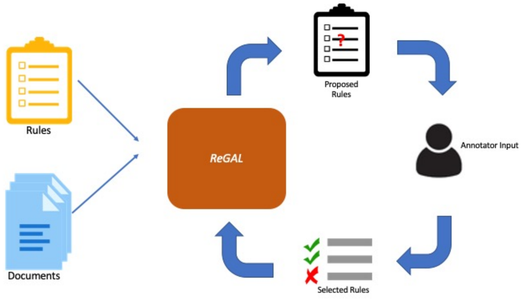

Methodology

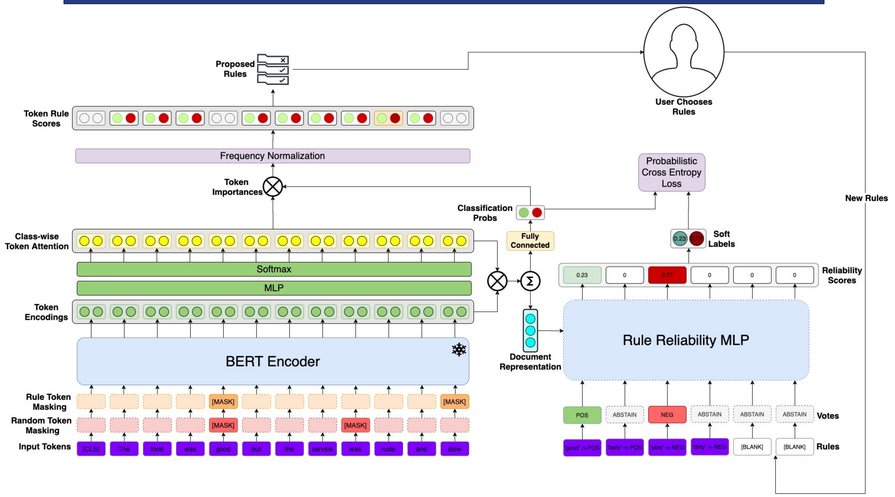

Rule Generation

- Aggregate and normalize class-wise attention over tokens
- Mask rule tokens to force discovery of new labeling signals

Rule Denoising

- Rule reliability conditioned on instance-level semantic content
- Weighted voting gives probabilistic class labels

Rule Generation

Round	Politics	Sports	Business	Tech
Seed	'war' , 'prime minister', 'president', 'commander', 'military', 'militant'	'baseball','soccer', 'basketball', 'football', 'world cup', 'olympics'	'sales', 'stock', 'market', 'shareholder', 'money', 'business'	'tech', 'engineering', 'scientist', 'processor', 'cpu', 'compute'
1	'election', 'bush', ~~'warned'~~	'uefa', 'league', 'olympic'	'shares', 'prices', 'oil'	'apple', 'intel', 'internet'
2	'iraq', 'iraqi', ~~'warns'~~	'coach', 'sports', 'season'	'crude', 'investors', 'company'	'space', 'microsoft', ~~'security'~~
3	~~'warn'~~, ~~'warning'~~, ~~'warnings'~~	'champions', 'manager', 'game'	'price', 'financial', 'companies'	'nasa', 'cyber', 'nuclear'
4	'putin', 'vladimir', ~~'leader'~~	'games', 'players', 'championship'	'profit', 'sale', 'economic'	'atomic', 'uranium', 'science'
5	~~'warming'~~, ~~'warm'~~, ~~'climate'~~	'hockey', 'team', ~~'night'~~	'futures', 'economy', 'trading'	'shuttle', ~~'enrichment'~~, ~~'spy'~~

Quality of Generated Rules

- ReGAL produces class-specific keyword rules
- Rule quality degrades over time due to overfitting and spurious patterns

Error Analysis

- Disproportionate rule bias towards tokens containing a current rule as a substring (e.g. “warn”, “warm”, “warning” from *“war”*)
- Rule sets are highly correlated, leading to errors from concept drift, e.g. ”scientists” (seed) > ”nuclear” (round 3) > “atomic” (round 4) > ”iran” (round 6)

Rule Denoising

Model	Seed			Seed + ReGAL		
	Yelp	IMDB	AGnews	Yelp	IMDB	AGnews
Majority voting	50.23	55.12	40.26	62.05	62.84	59.79
WeSTClass	76.90	71.87	79.42	**76.90**	71.87	79.42
Epoxy	73.03	74.00	61.55	75.55	74.68	64.23
BERT	79.08	72.76	79.72	75.84	73.28	80.99
MSWS	80.79	76.48	80.11	76.16	74.67	80.69
ReGAL	**85.21**	**78.40**	**80.15**	65.66	**75.24**	**83.85**
Fully Supervised BERT	91.1	90.7	87.2	91.1	90.7	87.2

ReGAL beats previous weakly supervised models, but its performance degrades due to its inductive biases and overfitting

Proposed Improvements

Rule Generation

- Generate complex rules from contextualized phrases and syntactic patterns
- Slice-specific rule generation (e.g. unmatched samples, overlapping classes, etc.)

Rule Denoising and Active Learning

- Simultaneously curate instance labels with rules to learn instance-level trustworthiness

Acknowledgements

Funding
This work was supported by NSF CAREER grant 1944247, NIH grant R21-CA232249, Whitaker Biomedical Engineering Seed Grant for COVID-19, Alzheimer’s Association Research Grant Award, Childrens Hospital of Atlanta Aflac Pilot Grant Award, National Science Foundation award III-2008334, an Amazon Faculty Award, and a Google Faculty Award.

References

1. Alexander Ratner, Stephen H. Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. “Snorkel.” In *Proceedings of the VLDB Endowment*, 11(3):269–282, 11 2017.
2. Wendi Ren, Yinghao Li, Hanting Su, David Kartchner, Cassie Mitchell, and Chao Zhang. “Denoising Multi-Source Weak Supervision for Neural Text Classification.” In *Findings of EMNLP*, 2020.
3. Yu Meng, Jiaming Shen, Chao Zhang, and Jiawei Han. "Weakly-supervised neural text classification." In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pp. 983-992. 2018.
4. Mayee Chen, Daniel Y. Fu, Frederic Sala, Sen Wu, Ravi Teja Mullapudi, Fait Poms, Kayvon Fatahalian, and Christopher Ré. "Train and You'll Miss It: Interactive Model Iteration with Weak Supervision and Pre-Trained Embeddings." *arXiv preprint arXiv:2006.15168* (2020).