


TL;DR: Collecting training sets in a dynamic adversarial fashion leads to higher quality training data and more robust models

Problems With Crowdsourcing




Lack of dataset quality!

- easy to solve examples
- spurious correlations
- lack of example diversity

Big question: How can we encourage crowd workers to write better training examples?

Dynamic Adversarial Data Collection



Possible improvements to data

- gamifies data collection
- harder examples
- increasingly diverse and difficult

Our paper: What happens if we collect many rounds of adversarial data?

Our Dataset



- Use NLI as test bed
- Small set of premises
- Many data rounds
- w/ and w/o adversarial

	No Model	Static Model	Dynamic Model
# Rounds	20	20	20
# Hypo.	11,000	11,000	11,000
# Verified Hypo.	7,684	7,102	6,911

	No Model	Static Model	Dynamic Model
<i>Diversity</i>			
Unique Unigrams	4.0k	4.2k	4.3k
Inter-example Sim.	41.2	41.9	39.5
<i>Complexity</i>			
Syntax	2.0	2.1	2.3
Reading Level	4.9	5.4	5.9
<i>Artifacts</i>			
Hypo-only Acc %	75.4	69.3	69.7

Adversarial data is more diverse, complex, and less prone to artifacts

Key Results



Over few rounds, adversarial data provides limited gains

Over many rounds, adversarial data provides big gains