

Current Research in NLP

Mausam

Plan (first 25%)

- Classical papers/problems in IE
 - NELL, Open IE, event extraction
- Important techniques for IE
 - distant supervision, joint inference

Plan (from then on: in flux)

- Machine learning concepts
 - Pre-trained language models
 - Transfer learning over pre-trained language models
 - Graph neural nets for NLP
 - Reinforcement Learning for NLP
 - Adversarial attacks in NLP
 - Generative Adversarial Nets for NLP

Plan (from then on: in flux)

- NLP topics
 - Debiasing NLP models
 - Distant supervision for Information Extraction
 - Tensor Factorization and Random-Walks for Knowledge-base Completion
 - Retrieval+ranking for Question Answering
 - Knowledge-base based Question Answering
 - Goal-oriented Dialog Systems
 - Other Knowledge-based Dialog Systems
 - Automated Summarization
 - Other Text Generation models

Signup on IE

- NELL
- Open IE
- Joint event extraction
- Distant supervision using PGMs
- Distant supervision using CNNs
- Distant supervision using RL

Workload

- Read ~1.5 papers every week (about 21-25 total)
 - Review them before midnight
- Make 1 presentation
 - You don't have to make original slides – credit the original authors!
- Midterm (max work)
 - Write a survey
- 2 assignments/1 project
 - Which one?
- 1 final

Grading

- 35% project/assignments
- 20% final exam
- 20% regular reviews
- 10% one presentation
- 15% midterm survey
- Extra credit: participation

Other Course Policies

- No audit!
- Request you to make a presentation even if you drop!

Structure of a review

- about 250 words is recommended
- (1) very short summary
- (2) pros
- (3) cons
- (4) how to improve the work
- (5) how to extend the work
- (6) dialogue: respond to previous reviews

Presenter: no review; but incorporate all reviews in your presentation

Projects

- Monthly milestones
 - 1st milestone: decide on a project and have the dataset/baseline: end of Jan.