

CSE 291:

Trends in Recommender Systems and Human Behavioral Modeling

Assignment

Assignment

- Very similar in structure to CSE258 assignment:
 - Identify one (or more) academic papers to implement and extend
 - Implement the idea described in the paper/s and reproduce their results
 - Describe a possible extension or modification of the paper in question
 - Implement and evaluate your proposed extensions/modifications
 - Discuss your results. Did your method succeed and why?
- Difference is mainly in focus on *state-of-the-art* methods and techniques from *academic literature*

Assignment

- Identify one (or more) academic papers to implement and extend
 - Choose a technical paper whose specific formulation could be modified
 - Could adapt a task to a new domain
 - Could take a non recsys paper/idea and apply it to recsys tasks (or vice versa, but ask me if you're proposing something like this!)
 - Add different forms of content/context to an existing model (e.g. how to use social/temporal/textual/visual/geographical information for a method that currently doesn't?)
 - Suggest a new type of evaluation for an existing method
- Write a (short) literature review of papers closely related to the one you chose
- Can use any paper, including those not covered in class, but do ask me!

Assignment

- Implement the idea described in the paper/s and reproduce their results
 - Think about what's feasible to implement!
 - Is there **code** available?
 - If not, is the technique simple enough that you can re-implement it based on what's described? Are parts of the description vague or obscure?
 - Is there **data** available?
 - If not, is there reasonably similar data you could use to validate the method?

Assignment

- Implement and evaluate your proposed extensions/modifications
- Discuss your results. Did your method succeed and why?
 - Were you able to reproduce what was described? If not, how did the results differ? What modifications did you have to make to the described method to get it to work?
 - Did your proposed modifications improve upon the previous solution? Are the results statistically significant? What types of data does your model work well or poorly for?

Examples

- Shoe recommendation++
 - **Paper:** Recommending Product Sizes to Customers
 - **Extension:** The paper only made use of 2-d (observed) size features. Could a higher-dimensional, but latent model, work better?
 - **Challenges:** Code and data is proprietary, evaluate on a similar dataset instead?

Examples

- Temporally & Sequentially-aware recommendation
 - **Paper:** Factorizing Personalized Markov Chains for Next-Basket Recommendation
 - **Extension:** The method makes use only of sequential data, but not timestamps. How can both be used together
 - **Challenges:** Data is proprietary. What are some similar datasets, where you expect both of these attributes to help?

Examples

- Content-based metric recommendation
 - **Paper:** Translation-based Recommendation and/or PRME
 - **Extension:** The method makes use only of latent attributes. Might it work better given features such as text?

Examples

- Subjective Q/A w/ attention mechanisms
 - **Paper:** Addressing Complex and Subjective Product-Related Queries with Customer Reviews
 - **Extension:** The method makes use of bag-of-words representations, could it be improved via the use of RNNs, or attention mechanisms?