

Your Name:

Your Andrew ID:

Homework 3

0 Introduction

0.1 Collaboration and Originality

1. Did you receive help of any kind from anyone (other than the instructor or TAs) in developing your software for this assignment (Yes or No)? If you answered Yes, provide the name(s) of anyone who provided help, and describe the type of help that you received.
2. Did you give help of any kind to anyone in developing their software for this assignment (Yes or No)? If you answered Yes, provide the name(s) of anyone that you helped, and describe the type of help that you provided.
3. Did you examine anyone else's software for this assignment (Yes or No)? Do not describe software provided by the instructor.
4. Are you (or the course instructor) the author of every line of source code submitted for this assignment (Yes or No)? If you answered No:
 - a. identify the software that you did not write,
 - b. explain where it came from, and
 - c. explain why you used it.
5. Are you the author of every word of your report (Yes or No)? If you answered No:
 - a. identify the text that you did not write,
 - b. explain where it came from, and
 - c. explain why you used it.

0.2 Instructions

Some experiments require you to set parameters or weights. You must explain why you chose particular values, and how your choices relate to how the technique works. We look for good experimental design – parameters that explore interesting issues or hypotheses (even if the hypothesis turns out to be wrong).

You must analyze the experimental results. Don't just summarize the numbers contained in the table – we can read the table ourselves. Instead, explain what conclusions you can reach based on the experiment. We are primarily interested in your observations about general trends, not quirky queries. Usually a good analysis addresses several issues. Show that you understand what the results mean, based upon what we have discussed in class.

Leave the page breaks between sections, as shown in this file. This will reduce the number of duplicate pages in your graded report.

Instructions are shown in a red italic bold font. Do not include instructions in your report. For example, delete this page. There is a 2 point deduction for leaving instructions in your report (because it creates extra work for the TAs).

1 Learning to Rank Baselines

Use your learning-to-rank software to train four models that use different groups of features.

1.1 Experimental Results for Learning to Rank Baselines

Your .zip / .tgz file must include files named HW3-Exp-1.1a.qry, HW3-Exp-1.1a.param, etc., in the QryEval directory. The experimental results shown above must be reproducible by these files and the parameter values shown in the table.

SVM _{rank}					
	BM25 (Exp-1.1a)	Indri SDM (Exp-1.1b)	IR Fusion (Exp-1.1c)	Content- Based (Exp-1.1d)	LTR Base (Exp-1.1e)
P@10	0.0000	0.0000	0.0000	0.0000	0.0000
P@20	0.0000	0.0000	0.0000	0.0000	0.0000
P@30	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@10	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@20	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@30	0.0000	0.0000	0.0000	0.0000	0.0000
MAP	0.0000	0.0000	0.0000	0.0000	0.0000

Coordinate Ascent					
	BM25	Indri SDM	IR Fusion (Exp-1.2c)	Content- Based (Exp-1.2d)	LTR Base (Exp-1.2e)
P@10	0.0000	0.0000	0.0000	0.0000	0.0000
P@20	0.0000	0.0000	0.0000	0.0000	0.0000
P@30	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@10	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@20	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@30	0.0000	0.0000	0.0000	0.0000	0.0000
MAP	0.0000	0.0000	0.0000	0.0000	0.0000

ListNet					
	BM25	Indri SDM	IR Fusion (Exp-1.3c)	Content- Based (Exp-1.3d)	LTR Base (Exp-1.3e)
P@10	0.0000	0.0000	0.0000	0.0000	0.0000
P@20	0.0000	0.0000	0.0000	0.0000	0.0000
P@30	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@10	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@20	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@30	0.0000	0.0000	0.0000	0.0000	0.0000
MAP	0.0000	0.0000	0.0000	0.0000	0.0000

1.2 Parameters for Learning to Rank Baselines

Briefly describe the parameter settings that were used to obtain these results.

1.3 Discussion for Learning to Rank Baselines

Discuss the trends that you observe; whether the learned retrieval models behaved as you expected; how the learned retrieval models compare to the baseline methods; and any other observations that you may have.

2 Custom Features

2.1 Descriptions of Custom Features

Describe each of your custom features, including what information it uses and its computational complexity. Explain the intuitions behind your choices. This does not need to be a lengthy discussion, but you need to convince us that your features are reasonable hypotheses about what improves search accuracy, and not too computationally expensive to be practical.

2.3 Parameters for Custom Features Experiments

Briefly describe the parameter settings that were used to obtain these results.

2.4 Discussion of Custom Features Experimental Results

Discuss the trends that you observe; whether the learned retrieval models behaved as you expected; how the learned retrieval models compare to the baseline methods; and any other observations that you may have.

3 Experiment 3: Feature Combinations

3.1 Experiment Table

Your .zip / .tgz file must include files named HW3-Exp-3.1a.qry, HW3-Exp-3.1a.param, etc., in the QryEval directory. The experimental results shown above must be reproducible by these files and the parameter values shown in the table.

	All (Baseline) (Exp-3.1a)	Comb ₁ (Exp-3.1b)	Comb ₂ (Exp-3.1c)	Comb ₃ (Exp-3.1d)	Comb ₄ (Exp-3.1e)
P@10	0.0000	0.0000	0.0000	0.0000	0.0000
P@20	0.0000	0.0000	0.0000	0.0000	0.0000
P@30	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@10	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@20	0.0000	0.0000	0.0000	0.0000	0.0000
NDCG@30	0.0000	0.0000	0.0000	0.0000	0.0000
MAP	0.0000	0.0000	0.0000	0.0000	0.0000

3.2 Parameters

Briefly describe the parameter settings that were used to obtain these results. Be sure to specify which learning algorithm you used.

3.3 Discussion

Describe each of your feature combinations, including its computational complexity. Explain the intuitions behind your choices, including the learning algorithm. This does not need to be a lengthy discussion, but you need to convince us that your combinations are investigating interesting hypotheses about what delivers good search accuracy. Were you able to get good effectiveness from a smaller set of features, or is the best result obtained by using all of the features? Why?

Be sure to discuss the effectiveness of your custom features. This should be more insightful than “They improved P@10 by 5%”. Discuss the effect on your retrieval experiments, and if there is variation in the metrics that are affected (e.g., P@k, MAP), how those variations compared to your expectations.