Social Recommendation and External Resources for Book Search

Romain Deveaud[*], Eric SanJuan[*] and Patrice Bellot[**]

[*] LIA – University of Avignon

[**] LSIS – Aix-Marseille University

UNIVERSITÉ D'AVIGNON

Introduction

- new book search task this year
  - searching books with social information and opinion instead of book content

- whole new collection
  - amazon.com and LibraryThing documents
  - topics extracted from LT forums

Introduction (2)

- 3 different approaches experimented this year

- weighted query expansion with Wikipedia
  - performed well last year
  - new thematic graph approach
  - #fail

Introduction (3)

- recommendation systems
  - using user reviews and ratings
  - good performances

- baselines
  - modeling unigrams and bigrams occurrences
  - using user tags and Amazon/LT browse nodes
  - best results (at least with crowdsourcing judgements)

Outline

- Introduction
- Baselines
- Query Expansion with Wikipedia
- Social Recommendation
- Conclusions and future work

Baselines

- language modeling approach to retrieval
- modeling unigrams and bigrams of the query
- Sequential Dependance Model
  - special case of the Markov Random Field [Metzler2005]
  - Dirichlet smoothing with default value ($\mu = 2500$)

Baselines (2)

- weighting query terms
  - unigram matches $f_T(q, D) = log\ P(q_i|D)$
  - bigram exact matches $f_O(q_i, q_{i+k}, D) = log\ P(\#1(q_i, ..., q_{i+k})|D)$
  - bigram matches within an unordered window of 8 tokens $f_U(q_i, q_{i+k}, D) = log\ P(\#uw8(q_i, ..., q_{i+k})|D)$

$$score_{SDM}(Q, D) = \lambda_T \sum_{q \in Q} f_T(q, D) + \lambda_O \sum_{i=1}^{|Q|-1} f_O(q_i, q_{i+1}, D) + \lambda_U \sum_{i=1}^{|Q|-1} f_U(q_i, q_{i+1}, D)$$

- $\lambda_T = 0.85$, $\lambda_O = 0.10$, $\lambda_U = 0.05$

- best run overall using Amazon MT judgements

Baselines (3)

- using classification information?
  - user tags, Amazon/LT category labels

```xml
<book><isbn>0673993280</isbn><title>Microeconomics (6th Edition)</title>
...
<tags>
    <tag count="1">Paperback</tag>
    <tag count="1">textbook</tag>
    <tag count="1">economics</tag>
    ...
</tags>
...
<browseNodes>
    <browseNode id="3">Business &amp; Investing</browseNode>
    <browseNode id="53">Nonfiction</browseNode>
    <browseNode id="1000">Subjects</browseNode>
    <browseNode id="2581">Economics</browseNode>
    ...
</browseNodes>
...
</book>
```

Evaluation using amazon mechanical turk™ Artificial Artificial Intelligence judgements

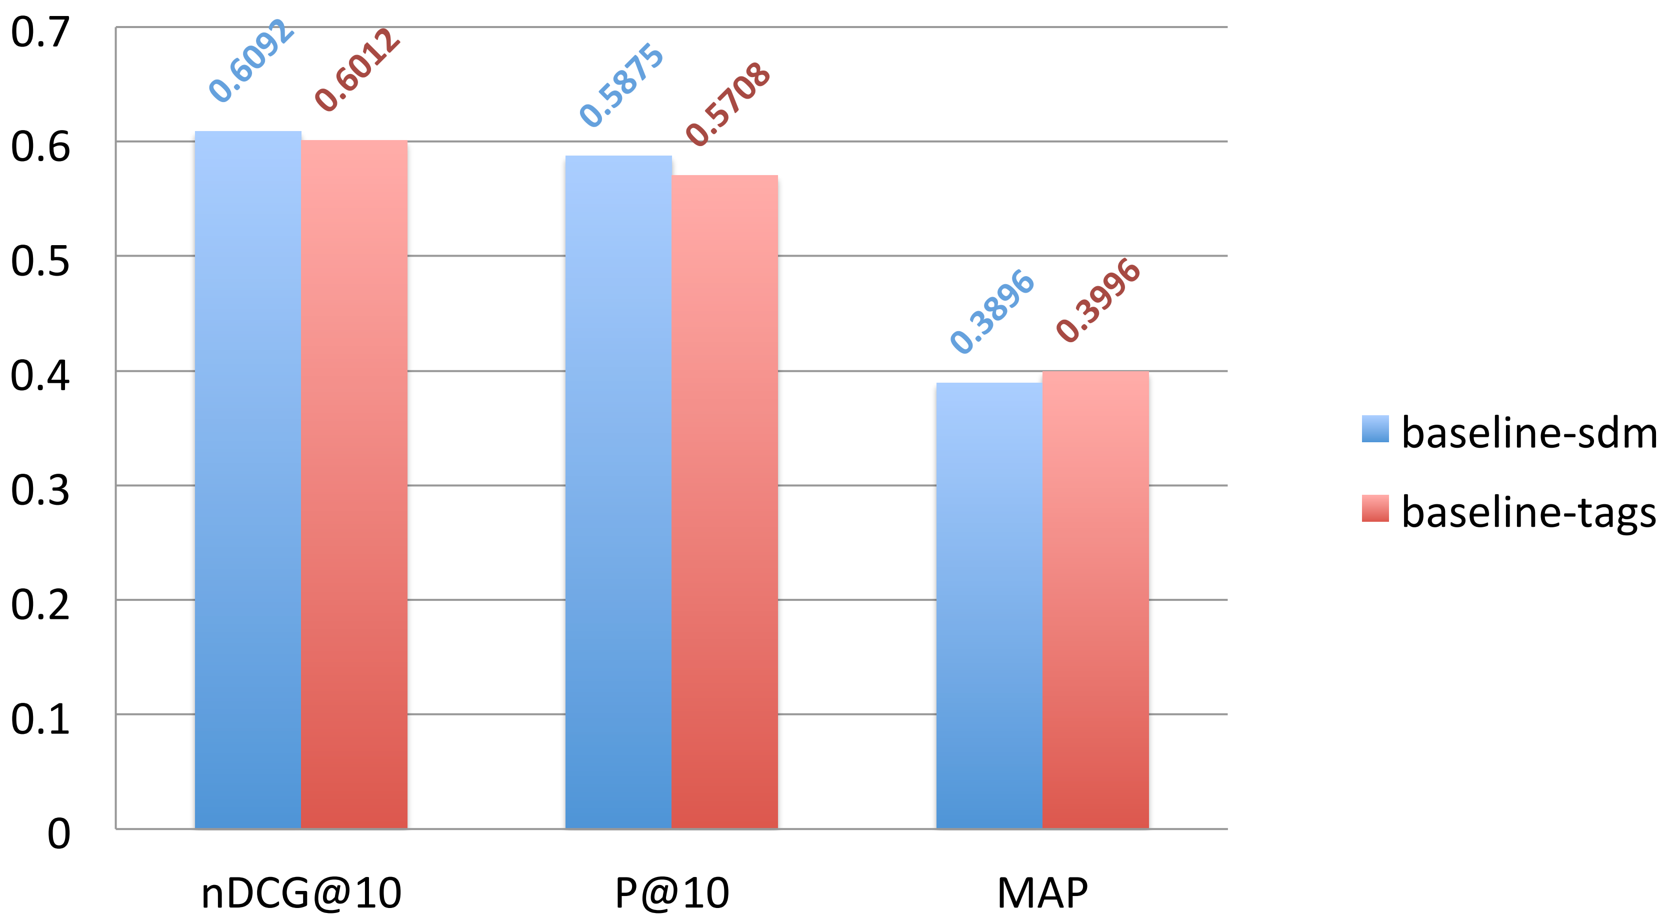

Outline

- Introduction
- Baselines
- Query Expansion with Wikipedia
- Social Recommendation
- Conclusions and future work

Query Expansion with Wikipedia

- few « real » text in the documents
  - mainly reviews or blurbs
  - very different from last year!

- expanding queries with terms from Wikipedia
  - expect to retrieve well-known books relevant to the query and extract useful terms

Query Expansion with Wikipedia (2)

– terms must be weighted in order to reflect their relative importance inside the expansion

– using *entropy* as a measure of informativeness

$$H_{\mathcal{W}}(w) = -\sum_{w \in \mathcal{W}} p_{\mathcal{W}}(w) \cdot \log p_{\mathcal{W}}(w)$$

– combining expansion terms with SDM

$$score(Q, D) = score_{SDM}(Q, D) + \lambda_{\mathcal{W}} \sum_{w \in \mathcal{W}} H_{\mathcal{W}}(w) \cdot f_T(w, D)$$

Query Expansion with Wikipedia (3)

- terms are extracted from the best ranked Wikipedia article for a given query

- using Wikipedia API for retrieving articles
  - stopwords removal with a 429 terms list, including HTML-related elements (nbsp, amp, http...)

- selecting the 20 words with best *entropy* for the expansion
  - performed well last year on Book and Ad Hoc tracks

Query Expansion with Wikipedia (4)

- problem: not enough « good » informative terms are selected

- possible solutions:
  - selecting more terms from pages: can introduce even more noise
  - using more pages: can cause some « topic drift »
  - using more pages thematically related to the best ranked page

Query Expansion with Wikipedia (5)

- Wikipedia hyperlinks and their anchor texts are defined by (expert) users
  - an hyperlink between two Wikipedia articles often models a thematic link
  - e.g. Heineken article have a link to the Beer article

- if an anchor text contains expansion terms, the linked article may be a good expansion document
  - beer is an expansion term from the Heineken article
  - more expansion terms can be extracted from the Beer article

Query Expansion with Wikipedia (6)

- can be viewed as an oriented graph
  - represents thematic hyperlinks from an article to another
  - query-driven tree

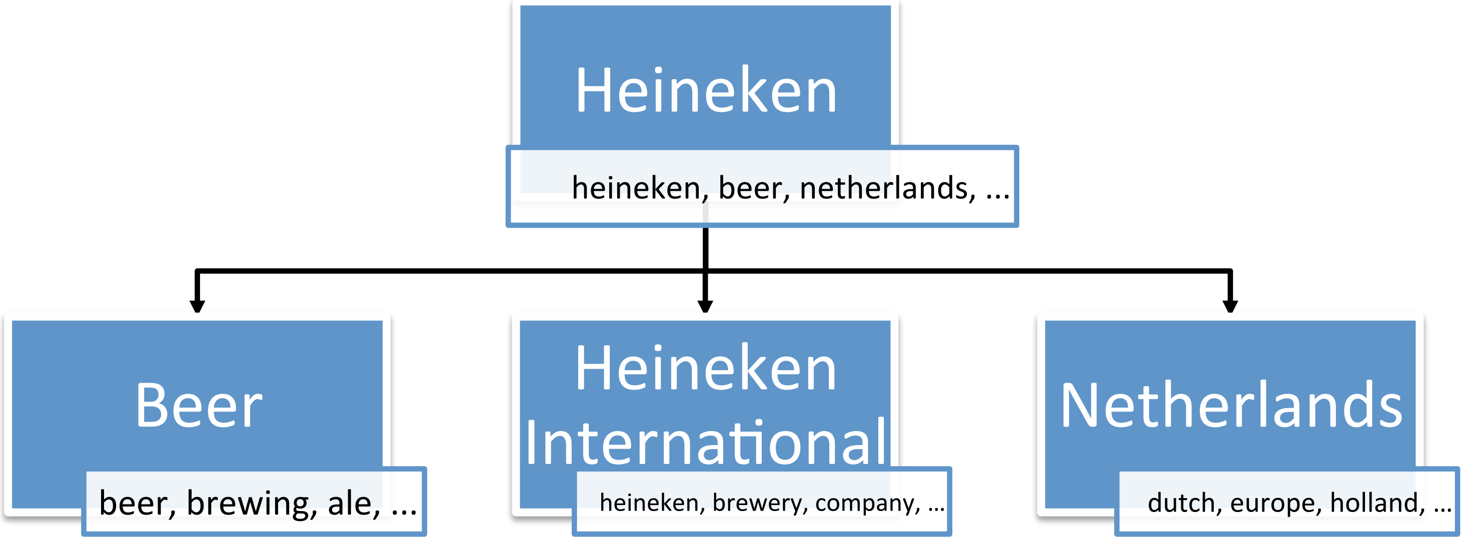

Query Expansion with Wikipedia (7)

- relies on the good selection of the first article
  - P@1 must be 1.0

- can cause topic drifts
  - down-weighting sub-articles ($w = 0.5$)
  - limiting the number of sub-articles ($N_{sub} = 5$)

Evaluation using amazon mechanical turk™ Artificial Artificial Intelligence judgements

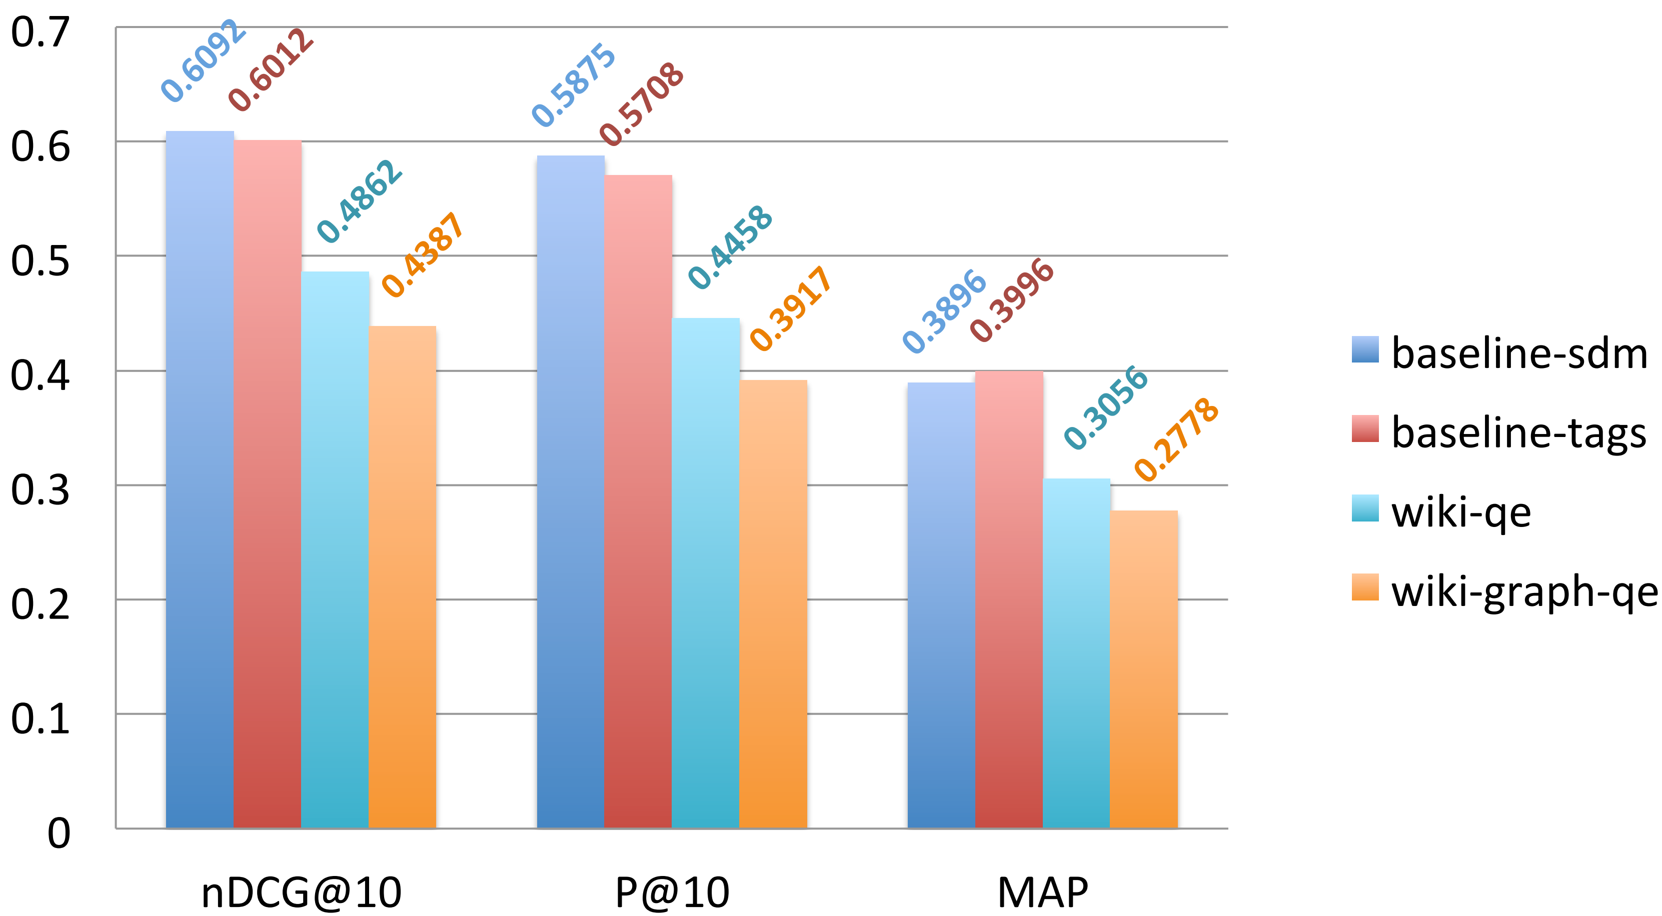

Outline

- Introduction
- Baselines
- Query Expansion with Wikipedia
- Social Recommendation
- Conclusions and future work

Social Recommendation

- documents are Amazon pages
  - users comment and review and rate products
- intuition #1: a high reviewed product must be relevant
  - or at least popular...
  - PageRank-like [Bao2007]
- intuition #2: a high rated product must be relevant

Social Recommendation (2)

source: xkcd.com

Social Recommendation (3)

- often a small amount of ratings for a book
  - how significant is the contribution of each rating?
  - X_R a random set of « bad » ratings in [1,3]
  - X_U the set of user ratings for a given book

- evaluate significant differences between X_R and $X_R + X_U$ (X_{R+U})
  - small $|X_U|$ but very good ratings
  - average ratings but large $|X_U|$

Social Recommendation (4)

- using Welch's *t*-test
  - test whether the population means are different
  - i.e. whether user ratings are useful

- popularity and quality in a single estimate
  - probability that users do not displease a book

- combining it with SDM (baseline) estimate
  - probability that a book is relevant to a query

Social Recommendation (5)

$$s_{recomm}(Q,D) = \lambda \ score_{SDM}(Q,D) + (1-\lambda) \ \frac{\bar{X}_{R+U} - \bar{X}_U}{s_{\bar{X}_{R+U} - \bar{X}_U}}$$

- based on observations over test topics:

$$\lambda = 1 - \frac{SDM_{MAX} - SDM_{100}}{N_{Results}}$$

- experiments with field retrieval
  - only `<title>` and `<content>` instead of all the text
  - field retricted smoothing
    - 6 for the title, 157 for the content of the reviews

Evaluation using amazon mechanical turk™ judgements

Artificial Artificial Intelligence

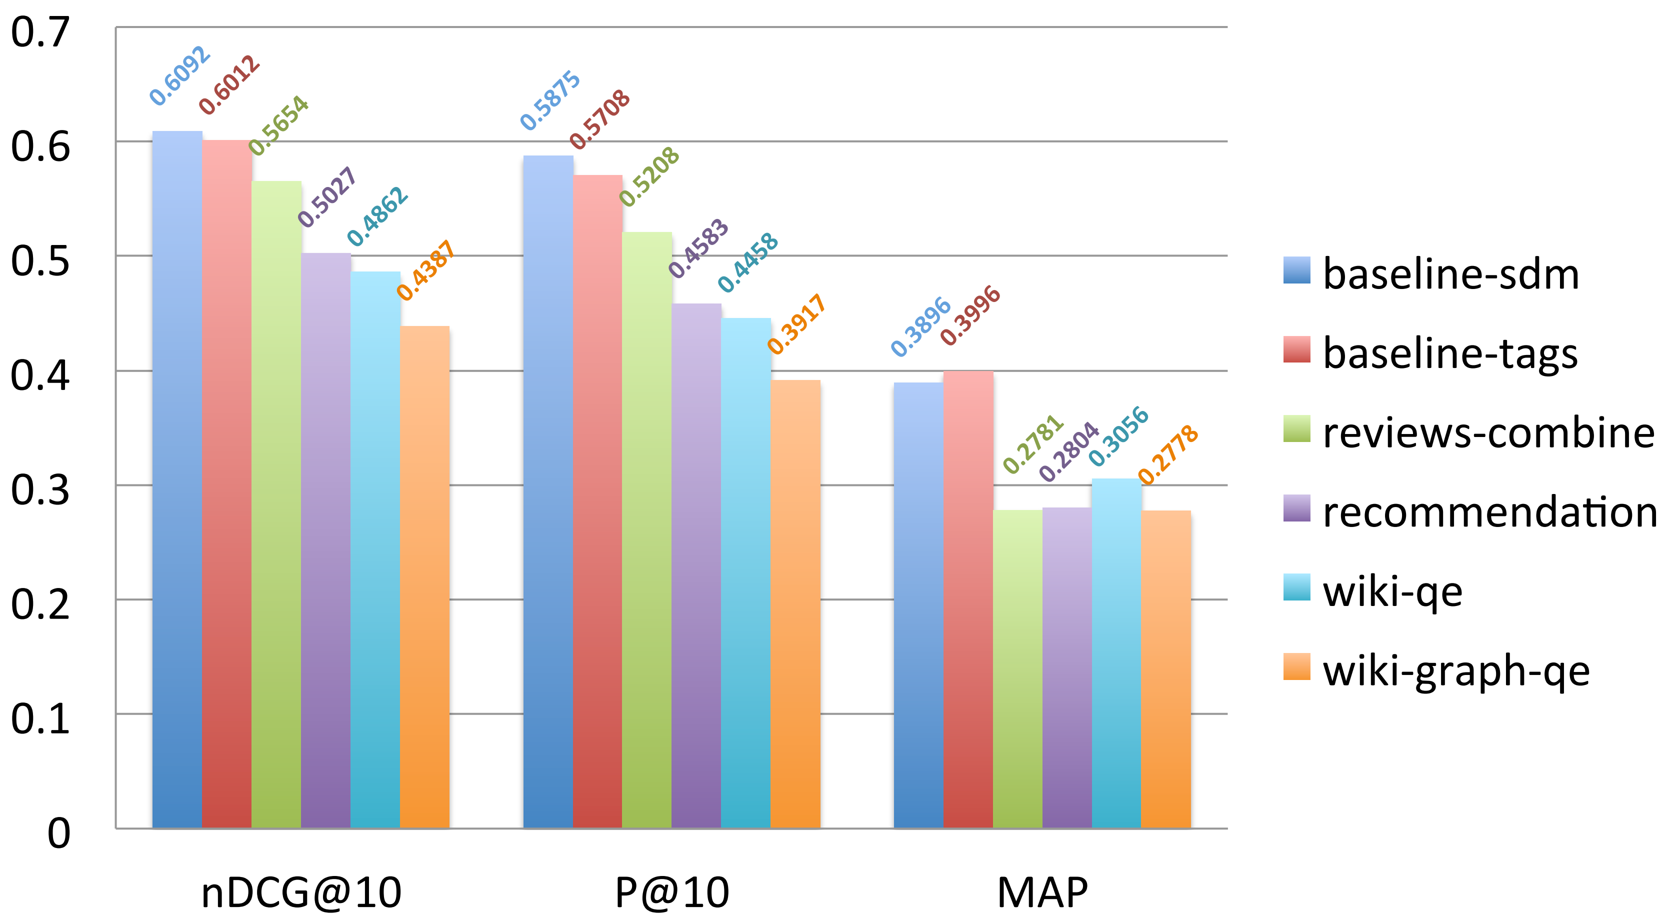

LibraryThing vs AMT judgements ?

- LibraryThing
  - derived from LT forums
  - 2377 judgements for 211 topics
  - 67 topics have less than 5 judgements
  - 74 topics have 10 or more judgements

amazon mechanical turk™
Artificial Artificial Intelligence

  - 1426 judgements for 24 (mixed) topics
  - 601 documents judged relevant

Evaluation using LibraryThing recommendations

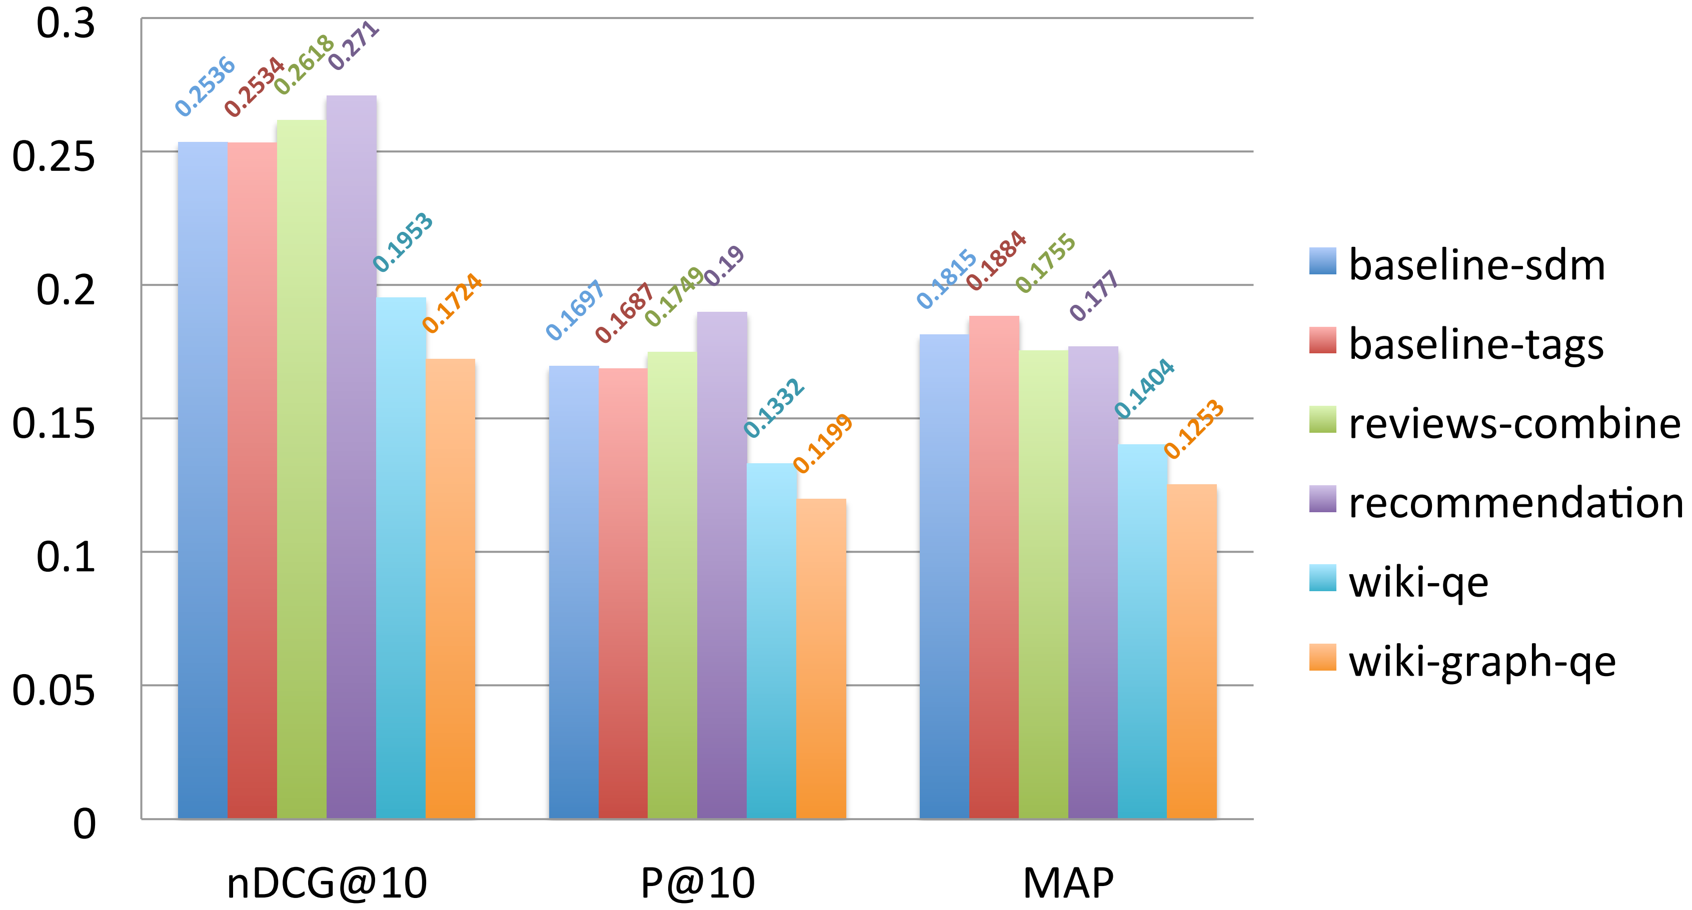

Outline

- Introduction
- Baselines
- Query Expansion with Wikipedia
- Social Recommendation
- Conclusions and future work

Conclusions and future work

- query expansion with weighted words extracted from Wikipedia does not work

- recommendation systems based on social ratings performs well
  - depending on the relevance judgements set

- LM approach with a SDM on the entire text of the Amazon pages achieves best results

Conclusions and future work (2)

- extending the relevance judgements?
  - concatenation?
  - more crowdsourcing?

- using more resources
  - social and opinion-oriented
  - combining several sources of user reviews

- apply to a general product search system thank you for your attention