

Intent Triage

Quantifying the Severity of Poor Performance on Intent
Classes

Fernando Diaz

Yahoo! Labs

July 23, 2010

query intent: a detailed, unambiguous representation of the user's information need.

query intent class: a group of related query intents (e.g. shopping, travel research, checking email).

Outline

Prior Work

Retrieval Performance on Different Intent Classes

Summary

State of the Art

Query Intent Class Detection

- I³R¹
- query classification²
- task intent classification³
- vertical intent classification⁴
- ...

¹ Croft and Thompson, "The use of adaptive mechanisms for selection of search strategies in document retrieval systems".

² Jansen et al., "Real Life Information Retrieval: A Study of User Queries on the Web"; Silverstein et al., *Analysis of a Very Large AltaVista Query Log*; Broder, "A taxonomy of web search"; Kang and Kim, "Query type classification for web document retrieval"; Beitzel et al., "Automatic web query classification using labeled and unlabeled training data".

³ Dai et al., "Detecting online commercial intention (OCI)"; Li, Wang, and Acero, "Learning query intent from regularized click graphs"; Jones and Klinkner, "Beyond the session timeout: automatic hierarchical segmentation of search topics in query logs"; Murdock and Croft, "Task orientation in question answering".

⁴ Arguello et al., "Sources of Evidence for Vertical Selection".

State of the Art

Responding to Query Intent Class

- specialized ranking models⁵
- vertical presentation⁶
- query suggestion⁷
- personalization⁸
- ...

⁵Dong et al., "Towards recency ranking in web search"; Bian et al., "Ranking with Query-Dependent Loss for Web Search", "Ranking Specialization for Web Search: A Divide-and-Conquer Approach by Using Topical RankSVM".

⁶Diaz, "Integration of News Content Into Web Results".

⁷Cao et al., "Context-aware query suggestion by mining click-through and session data".

⁸Teevan, Dumais, and Liebling, "To personalize or not to personalize: modeling queries with variation in user intent".

Managing a Taxonomy of Intent Classes

- Sources of evidence
 - corpus: defines the possible means of satisfying intents
 - query, browsing logs: defines the types of intents users have
 - ...
- Defining intent classes
 - data mining: analyze the available evidence
 - feedback driven: explicit/implicit user feedback (e.g. 'report a problem' link)
 - ...

$$\text{DCG}(q, \theta) = f(\text{Intent-Class}(q))$$

The performance of a retrieval system may be influenced by the query's intent class.⁹

⁹Bian et al., "Ranking with Query-Dependent Loss for Web Search".

How do we compare performance across different intent classes?

Approach 1: Uniform Impression Weighting

$$\begin{aligned}\frac{1}{|Q'|} \sum_{q \in Q'} \text{DCG}(q, \theta) &= \frac{1}{|Q'|} \sum_{q \in Q'} f(\text{Intent-Class}(q)) \\ &= \sum_{i \in \mathcal{I}} \frac{|Q'_i|}{|Q'|} f(i)\end{aligned}$$

- **Assumption:** Each impression is equally important.
- **Implication:** Each intent class is weighted by its volume in traffic.

Approach 2: Uniform Query Weighting

$$\begin{aligned}\frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \text{DCG}(q, \theta) &= \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} f(\text{Intent-Class}(q)) \\ &= \sum_{i \in \mathcal{I}} \frac{|\mathcal{Q}_i|}{|\mathcal{Q}|} f(i)\end{aligned}$$

- **Assumption:** Each query is equally important.
- **Implication:** Each intent class is weighted by its representation in the query set.

Approach 3: Uniform Intent Weighting

$$\begin{aligned}\sum_{i \in \mathcal{I}} \frac{1}{|\mathcal{I}|} f(i) &= \sum_{q \in \mathcal{Q}} \frac{1}{|\mathcal{I}|} \left(\frac{1}{|\mathcal{Q}_{\text{Intent-Class}(q)}|} f(\text{Intent-Class}(q)) \right) \\ &= \sum_{q \in \mathcal{Q}} \frac{1}{|\mathcal{I}| |\mathcal{Q}_{\text{Intent-Class}(q)}|} \text{DCG}(q, \theta)\end{aligned}$$

- **Assumption:** Each intent class is equally important.
- **Implication:** Each query is weighted by the number of other queries with its same intent class.

Approach 4: Revenue-Based Intent Class Weighting

$$\sum_{i \in \mathcal{I}} \frac{\$i}{\sum_{i' \in \mathcal{I}} \$i} f(i)$$

- **Assumption:** Each intent class should be weighted according to its impact on revenue.
- **Implication:** Must carefully select horizon for $\$i$.

Approach 5: User-Based Intent Class Weighting

$$\sum_{i \in \mathcal{I}} u_i f(i)$$

- **Assumption:** Each intent class should be weighted according to its impact on the user.
- **Implication:** Must carefully define u_i .

Measuring User Impact

$u_i \propto$

- delight: user delight from system success for i
- frustration:¹⁰ user frustration from system failure for i
- recovery effort:¹¹ user effort to recover from system failure for i
- quality of life impact: quality of life impact to the user as a result of system failure for i
- ...

¹⁰Feild, Allan, and Jones, "Predicting searcher frustration"; White and Dumais, "Characterizing and Predicting Search Engine Switching Behavior".

¹¹Feild, Allan, and Jones, "Predicting searcher frustration".

Opportunity

- The most catastrophic impact is for users unable to get information during an emergency crisis (e.g. earthquake, hurricane).
- Information is the most disorganized at this point; difficult to recover with reformulation, deep inspection, etc.
- Information retrieval techniques are designed to support this kind of task.

Summary

- Can we be more formal about defining, detecting, organizing intent classes?
- Should we reason about intent class importance when computing system performance?
- What are the most important queries for users which systems are not currently supporting?

...

Arguello, Jaime et al. "Sources of Evidence for Vertical Selection". In: *SIGIR 2009*. 2009.

Beitzel, Steven M. et al. "Automatic web query classification using labeled and unlabeled training data". In: *SIGIR '05: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*. Salvador, Brazil: ACM, 2005, pp. 581–582. ISBN: 1-59593-034-5. DOI: <http://doi.acm.org/10.1145/1076034.1076138>.

Bian, Jiang et al. "Ranking Specialization for Web Search: A Divide-and-Conquer Approach by Using Topical RankSVM". In: *WWW '10: Proceedings of the 19th international conference on World wide web*. 2010.

Bian, Jiang et al. "Ranking with Query-Dependent Loss for Web Search". In: *Proceedings of the Third ACM International Conference on Web Search and Data Mining*. 2010.

Broder, Andrei. "A taxonomy of web search". In: *SIGIR Forum* 36.2 (2002), pp. 3–10. ISSN: 0163-5840. DOI: <http://doi.acm.org/10.1145/792550.792552>.

Cao, Huanhuan et al. "Context-aware query suggestion by mining click-through and session data". In: *KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. Las Vegas, Nevada, USA: ACM, 2008, pp. 875–883. ISBN: 978-1-60558-193-4. DOI: <http://doi.acm.org/10.1145/1401890.1401995>.

Croft, W. Bruce and Roger H. Thompson. "The use of adaptive mechanisms for selection of search strategies in document retrieval systems". In: *Proceedings of the 7th annual international ACM SIGIR conference on Research and development in information retrieval*. 1984.

Dai, Honghua (Kathy) et al. "Detecting online commercial intention (OCI)". In: *WWW '06: Proceedings of the 15th international conference on World Wide Web*. Edinburgh,

Scotland: ACM, 2006, pp. 829–837. ISBN: 1-59593-323-9. DOI: <http://doi.acm.org/10.1145/1135777.1135902>.

Diaz, Fernando. “Integration of News Content Into Web Results”. In: *Proceedings of the Second ACM International Conference on Web Search and Data Mining*. 2009.

Dong, Anlei et al. “Towards recency ranking in web search”. In: *WSDM '10: Proceedings of the third ACM international conference on Web search and data mining*. New York, New York, USA: ACM, 2010, pp. 11–20. ISBN: 978-1-60558-889-6. DOI: <http://doi.acm.org/10.1145/1718487.1718490>.

Feild, Henry A., James Allan, and Rosie Jones. “Predicting searcher frustration”. In: *SIGIR '10: Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. Geneva, Switzerland: ACM, 2010, pp. 34–41. ISBN: 978-1-4503-0153-4. DOI: <http://doi.acm.org/10.1145/1835449.1835458>.

Jansen, Bernard J. et al. "Real Life Information Retrieval: A Study of User Queries on the Web". In: *SIGIR Forum* 32.1 (1998), pp. 5–17.

Jones, Rosie and Kristina Lisa Klinkner. "Beyond the session timeout: automatic hierarchical segmentation of search topics in query logs". In: *CIKM '08: Proceeding of the 17th ACM conference on Information and knowledge management*. Napa Valley, California, USA: ACM, 2008, pp. 699–708. ISBN: 978-1-59593-991-3. DOI: <http://doi.acm.org/10.1145/1458082.1458176>.

Kang, In-Ho and GilChang Kim. "Query type classification for web document retrieval". In: *SIGIR '03: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*. Toronto, Canada: ACM, 2003, pp. 64–71. ISBN: 1-58113-646-3. DOI: <http://doi.acm.org/10.1145/860435.860449>.

Li, Xiao, Ye-Yi Wang, and Alex Acero. "Learning query intent from regularized click graphs". In: *SIGIR '08*:

Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval. Singapore, Singapore: ACM, 2008, pp. 339–346. ISBN: 978-1-60558-164-4. DOI:

<http://doi.acm.org/10.1145/1390334.1390393>.

Murdock, Vanessa and W. Bruce Croft. “Task orientation in question answering”. In: *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. Tampere, Finland: ACM, 2002, pp. 355–356. ISBN: 1-58113-561-0. DOI: <http://doi.acm.org/10.1145/564376.564441>.

Silverstein, Craig et al. *Analysis of a Very Large AltaVista Query Log*. Tech. rep. SRC-TN-1998-014. HP Labs Technical Report, 1998.

Teevan, Jaime, Susan T. Dumais, and Daniel J. Liebling. “To personalize or not to personalize: modeling queries with variation in user intent”. In: *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on*

Research and development in information retrieval.
Singapore, Singapore: ACM, 2008, pp. 163–170. ISBN:
978-1-60558-164-4. DOI:

<http://doi.acm.org/10.1145/1390334.1390364>.

White, Ryen W. and Susan T. Dumais. “Characterizing and Predicting Search Engine Switching Behavior”. In: *CIKM '09: Proceeding of the 18th ACM conference on Information and knowledge management*. New York, NY, USA: ACM, 2009, pp. 87–96.