

Controlling the search for expanded query representations by constrained optimization in latent variable space

Kevyn Collins-Thompson
Microsoft Research
1 Microsoft Way
Redmond, WA 98052 U.S.A.
kevynct@microsoft.com

Joshua V. Dillon
College of Computing
Georgia Institute of Technology
jvdillon@gatech.edu

ABSTRACT

We investigate a new family of algorithms for finding reliable query representations based on pseudo-relevance feedback, by replacing the traditional E-step in EM methods with a more general convex optimization step that allows constraints and objectives to bias the search by direct influence in the space of latent variables. This bias amounts to a penalized form of maximum likelihood objective, and we investigate adding a translation kernel and a word diversity constraint. Combined with simple generative models for IR, constraining the posterior distribution of the latent variables is also a natural way to incorporate evidence of user intent and feedback, such as observations on words, documents, or specific word-document pairs. More generally, we believe such an approach provides an interesting starting point for unified probabilistic approaches in which multiple sources of context are used to guide parameter estimation in models of query intent under task- or domain-specific constraints.

Categories and Subject Descriptors: H.3.3 [Information Retrieval]: Retrieval Models

General Terms: Algorithms, Experimentation

Keywords: Query expansion, convex optimization

1 Introduction

It is well known that the effectiveness of pseudo-relevance feedback can be highly variable across queries, and sensitive to a number of parameters, such as the number of terms, or number of top-ranked documents chosen. Such instability of query expansion and other query reformulation algorithms motivates *robust algorithms* that try reduce instability by finding reliable values for these parameters automatically for each query, reducing the risk inherent in committing to a single operational setting for all queries.

In general, when developing algorithms for estimating query models, it is common to work with objectives and constraints in the space of parameters of the model. For example, in a language modeling approach, smoothing or other adjustments are typically applied to the parameters, which are

usually estimated using maximum likelihood. Motivated by the need for more reliable algorithms, in this paper we investigate guiding model search directly in the *latent variable space* of simple generative models, instead of modifying model parameters directly. As an example, we study an extended form of EM in which the basic E-step is replaced by a more general convex optimization step that amounts to a penalized version of the basic maximum likelihood objective. The idea is that task-specific knowledge be used to express expected relationships over the latent variables, which often have an intuitive interpretation and make more sense to work with for obtaining detailed control over the parameter estimation process. Feedback observations on either words or documents, or even specific words in specific documents, are also easily and directly incorporated as constraints in the latent variable space.

For example, suppose we assume a generative model that uses binary latent variables Z , with one Z_{dw} for each word w , document d pair, where $P(Z_{dw} = 1)$ indicates how likely it is that word w was generated from a ‘relevant topic’ model in the context of document d . A simple constraint might be that $z_{w1} > z_{wj}$ if we have feedback that word w document 1 is preferred by the user to its use in other documents d_j , say because w is used in a non-relevant sense in those documents. In this paper, we consider a more general *diversity constraint* over the z_{dw} for a given document. In standard EM formulations common to IR, there is no term dependency model that influences the parameter updates. Thus, we also add the possibility of a word-to-word *translation model* in the latent variable estimation. We do a basic evaluation of the effects of these generalizations on the risk-reward tradeoff of the query model estimation algorithm.

2 Optimization framework

We derive our optimization model in three steps. First, we describe a generative model introduced by Tao and Zhai (TZ) that conditions queries and document observations on latent variables. Second, we give some background on the TZ algorithm for estimating the parameters of the model using EM and a regularized maximum likelihood objective. Third, we discuss specific generalizations for the E-step that modify the maximum-likelihood objective to bias the model search. The overall result is an extremely flexible optimization framework for constructing effective searches for high-quality query representations that can account for meaningful problem structure or task knowledge in the latent variable space.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

```

GENERATE-RANDOM-DOCUMENT( $L, \alpha; \theta_T, \theta_B$ )
1  $c \leftarrow \text{ZEROS}(1, K)$  // observed rv
2  $z \leftarrow \text{ZEROS}(1, K)$  // latent rv
3 for  $j \leftarrow 1$  to  $L$ 
4 do  $f \leftarrow \text{FLIP-HEADS-BIASED-COIN}(\alpha)$ 
5   if  $f$  is HEADS
6     then  $w \leftarrow \text{ROLL-BIASED-DIE}(\theta_T)$ 
7          $z[w] \leftarrow z[w] + 1$ 
8     else  $w \leftarrow \text{ROLL-BIASED-DIE}(\theta_B)$ 
9          $c[w] \leftarrow c[w] + 1$ 
10 return  $c$ 

```

Figure 1: The Tao-Zhai (TZ) generative procedure for a single document in the feedback set. This model posits documents are a two-part mixture of multinomials (on a per-word basis). L denotes the length of the document, i.e., number of multinomial samplings, and α its relevant/nonrelevant mixture. Documents are assumed independent, given the parameters.




Figure 2: Graphical model depiction of (above) generative procedure (for n documents). Nodes are random variables, with shaded nodes being observed (user’s query θ_q , fixed corpus background model θ_B , word counts c_{dw}). Plates indicate conditionally independent replications.

2.1 A basic generative model of queries and documents

Tao & Zhai [11] introduced a simple generative model and optimization algorithm for pseudo-relevance feedback¹. One distinctive feature of their model is that it jointly optimizes both query and document weights simultaneously. The generative model is shown in Figures 1 and 2. Top- n feedback documents² $F_n(q)$ are modeled as being generated from a mixture of two multinomial language models: a background model θ_B and a ‘relevant topic’ model θ_T using vocabulary \mathcal{V} of dimension K . The topic model θ_T is assumed to have a fixed Dirichlet prior, $\text{Dir}(1 + \mu\theta_q)$. The feedback documents are also endowed with ‘relevance’ weights α_d which are to be learned jointly with the word probabilities of θ_T , by introducing hidden variables Z_{dw} which are meant to capture whether term w in document d has been generated from θ_T instead of θ_B . We denote our main observations, the counts of term w in document d , by c_{dw} and also denote $P(Z_{dw} = 1)$

¹Our notation mostly follows that used in Tao & Zhai, with some extensions. Omitting the subscript indicates the matrix/vector, as opposed to the specific element, e.g. α is a vector of α_i and $[\theta_B]_w = \theta_{Bw}$. Also, the value of a variable, e.g. α , after k iterations is denoted by $\alpha^{(k)}$.

²The ranking function used to obtain $F_n(q)$ is not critical here, but for purposes of discussion we assume it is the Kullback-Leibler (KL) divergence model [8].

by z_{dw} . Note that we view z_{dw} as a matrix Z of variables having K rows and n columns.

To estimate α_d and θ_T , Tao and Zhai give a standard EM algorithm with closed-form E and M steps, yielding the following two-step iterative procedure:

E-step:

$$z_{dw}^{(k)} \leftarrow \frac{\alpha_d^{(k)} \theta_{Tw}^{(k)}}{\theta_{Tw}^{(k)} \alpha_d^{(k)} + \theta_{Bw}^{(k)} (1 - \alpha_d^{(k)})} \quad (1)$$

M-step:

$$\alpha_d^{(k+1)} \leftarrow \frac{\sum_{j=1}^V c_{dj} z_{dj}^{(k)}}{\sum_{j=1}^V c_{dj}} \quad (2)$$

$$\theta_{Tw}^{(k+1)} \leftarrow \frac{1}{\xi} \left(\mu \theta_{qw} + \sum_{i=1}^n c_{iw} z_{iw}^{(k)} \right) \quad (3)$$

with Lagrangian ξ the appropriate normalization of $\theta_{Tw}^{(k+1)}$. Typically the E- and M-steps are repeated until convergence: the TZ algorithm changes this slightly to impose a schedule of decay on parameter μ , i.e., $\mu \leftarrow \mu \delta^k$ where $\delta \in (0, 1)$.

2.2 Constraining the E-step

The above EM algorithm seeks a parametrization that maximizes the posterior likelihood under the query driven Dirichlet prior. Here, we seek an alternative procedure that still retains maximum likelihood as part of a more general objective, but adds additional flexibility in the E-step. The key idea is that at the k -th step, we seek the closest latent variable matrix $X^{(k)}$ to the ‘default’ E-step update matrix $Z^{(k)}$, but subject to additional conditions. By ‘closest’, for this paper we use the Frobenius norm, a standard distance measure for matrices. We apply two extensions: a *diversity constraint* over the latent variables z_{dw} for all words w in a given document d , and a *term dependency model* which can be used to influence the least-distance objective. The complete algorithm is shown in Figure 3. To use it in the EM algorithm we simply replace the normal E-step with our convex program, and use the optimal solution matrix \hat{X} instead of the default matrix Z .

2.2.1 Diversity constraints

In the search for reliable solutions, we may consider that relying too heavily on only a small number of uncertain latent variables is too risky. Instead, we could implement a *diversity constraint* over the words in each document, so that no more than η_w -percent of the total probability mass can be allocated to the top r_w terms. It turns out that this can be expressed as a linear constraint ([2], p.279) via auxiliary vectors u_j of size $|V|$ and a scalar variable t_j for each document d_j , along the columns of $X^{(k)}$. This diversity constraint appears superficially similar to standard smoothing methods, in that it acts to redistribute probability mass from higher-probability events to lower-probability ones. However, unlike standard smoothing, relative changes in latent variable mass can change significantly from iteration to iteration, due to the nature of the top- k criterion and hard upper-bound on the mass³.

³We note that a similar form of diversity constraint could be applied to *documents* by constraining the rows of $X^{(k)}$, which hold the latent variables for occurrences of a single word w across all documents. We might prefer states in which there must be stronger evidence across multiple documents instead of relying on a single source. Variance-based Markowitz-type diversity [4] is another possibility. These are

$$\begin{aligned}
& \text{minimize } \|X - (I + \lambda \Sigma_T)Z\|_F && \text{E-step dist.} && (6) \\
& \text{subject to } \Sigma_i X_{ij} = \Sigma_i Z_{ij} && \text{Doc mass invariant} && (7) \\
& r_W \cdot t_j + \mathbf{1}^T u_j \leq \eta_W && \text{Diversity constr I} && (8) \\
& t_j + u_j \geq x_j / c_j && \text{Diversity constr II} && (9) \\
& c_j = \Sigma_{i=1}^V Z_{ij} && && (10) \\
& u_j \geq 0 && && (11) \\
& 0 \geq X \geq 1 && \text{Label consistency} && (12)
\end{aligned}$$

Figure 3: The basic constrained E-step for finding the closest matrix X to the default E-step matrix Z , while respecting diversity constraints over a document’s latent variables for words, and using a translation kernel Σ_T in the objective. Here, $j = 1 \dots F$ over the set of F feedback documents.

2.2.2 Adding term dependency information via translation kernels

The fact that the basic TZ algorithm does not incorporate a term dependency model suggests biasing the model search by adding term dependency information into the objective.

One effective method of estimating semantic term dependency is to define a statistical translation process between terms using *translation kernels* [6]. We create a translation kernel by first computing a similarity graph between all pairs of terms. For vertices u and v , the edge weight $e(u, v)$ is defined as a function of $f_u(w)$, the co-occurrence frequency of term u with term w in the top-ranked documents, giving a matrix E with entries

$$e_{uv} = \exp\left(-\frac{1}{\sigma^2} \arccos^2 \sum_w \sqrt{f_u(w)f_v(w)}\right) \quad (4)$$

where the sum is taken over all words w in the vocabulary \mathcal{V} . The graph heat kernel is computed via the matrix exponential of the normalized graph Laplacian

$$\mathcal{L} = D^{-1/2}(D - E)D^{-1/2} \quad (5)$$

where D is a diagonal matrix with $D_{ii} = \sum_j e_{ij}$. The matrix exponential $\Sigma_T = \exp(-t\mathcal{L})$ models the flow of heat across the graph as a function of time parameter t , which controls the amount of translation. For small t , $\Sigma_T \approx I$ and for large t , Σ_T is approximately uniform. We use a translation strength parameter λ to combine Σ_T with Z using $(I + \lambda \Sigma_T)$. We set the dilation factor σ^2 and time parameter t of the translation kernel to 0.75 and 5 respectively.

2.2.3 Implementation

In practice, because the initial z_{dw} are very close to zero, for numerical reasons we first run a small number L of iterations (in these experiments $L = 3$) of the standard EM algorithm before switching to the Constrained EM version. Second, because the diversity constraint is a hard constraint, the convex program is occasionally infeasible at some iteration⁴. When this happens, we simply terminate the search and use the last known good solution.

3 Related work

Optimization frameworks have been applied before to query representations. Collins-Thompson introduced the use of topics for future work.

⁴Problems such as infeasibility and failure to converge are detected automatically by the solver and reported with a status code.

portfolio theory, applying it to optimize a convex mean-variance criterion over the set of words used for query expansion [4] [3]. Their approach operates as a post-process in the space of expansion terms and assumes little about the underlying retrieval model. It may return the *empty set* of expansion terms if expansion is deemed too risky for a given query. The TZ algorithm, on the other hand, jointly solves a non-convex objective for both term and document weights, but does not model term dependencies or have the ability to constrain its solution space. Subsequently, Xu and Akella [13] replaced the TZ two-mixture generative model with a Dirichlet Compound Multinomial, using a different latent variable model and closed-form E-step based on simulated annealing. It would be interesting to explore the implications of the latter work in our optimization framework.

We have also drawn inspiration from a key reference work by Graca et al. [7], who proposed modifying EM using a constrained E-step in order to model posterior constraints. They also gave theoretical results that give a penalized maximum likelihood interpretations to their framework, and gave examples of several natural-language applications, including statistical translation. The application of heat-transfer translation kernels to query expansion is another contribution of this paper. The closest previous work on random walk models for query expansion [5] also used a term dependency graph in which word co-occurrence was one of several dependency types. We also note that Mei & Zhai proposed a framework [10] for smoothing language models on graphs. Finally, in a recent study, Udapa et al. [12] presented evidence for the importance of set-level properties in finding high-quality expansion sets. This is relevant to our work because desirable set-level properties can often be expressed or approximated in terms of convex constraints or objectives and thus easily incorporated in our framework.

4 Evaluation


We provide a brief analysis of the effect that changes in the E-step have on query model quality: the addition of a translation kernel to the objective and linear diversity constraints.

For this summary evaluation, we present results for wt10g (TREC topics 451-550). Indexing and retrieval were performed using the Indri 2.8 system in the Lemur toolkit [9]. We used the title fields of the TREC topics and phrases were not used. We also did not use stopping or stemming. The initial queries consisted of the title words wrapped with the `combine` operator, with the top 1000 documents retrieved using Dirichlet query smoothing with $\mu = 2000$. We limited the maximum number of iterations for both TZ and Constrained EM to 50, since both algorithms have generally converged by then. We used the 50 top-ranked documents and 20 expansion terms.


We use risk-reward curves [3] to show an algorithm’s achievable tradeoff between average precision gain and the variance of expansion failures⁵. Such curves are able to summarize regimes of performance for two expansion methods that might appear identical by their equal MAP performance but where one is more prone to catastrophic failures than the other. One important difference is that our risk-reward curves are generated as a function of the *number of iterations* instead of, say, an interpolation parameter.

Figure 4 shows two effects. In a), increasing the amount of translation from $\lambda = 0$ to $\lambda = 5$ gives a dramatic improvement in the risk-reward tradeoff: with the translation kernel on this collection, the algorithm never hurts query

⁵An expansion failure is a query for which the expansion algorithm gives worse performance than the original query.



(a) Effect of translation parameter λ , with constant diversity $\eta = 0.9$.



(b) Effect of diversity parameter η , with no translation ($\lambda = 0$)

Figure 4: Risk-reward tradeoff curves show the effect of increasing the word translation factor λ with fixed diversity constraint $\eta_W = 0.90$, compared to Tao-Zhai baseline (TZ). Curves are a function of iteration, with each iteration as a dot and every 10th iteration numbered. Because the initial query is the starting point and the y -axis shows relative MAP gain, curves will start at the origin and trace out a risk-reward tradeoff with each iteration. Tradeoff curves that are *higher and to the left* are better.

performance on average, giving its maximum MAP at convergence. The baseline TZ algorithm, on the other hand, deteriorates quickly after about 20 iterations. In b), making the diversity constraint more strict by decreasing the allowable η percentage available to the top-ranked words acts to flatten the risk-reward curve slightly. Future analysis work includes sensitivity of the results to parameter changes, and measuring interaction effects between diversity and translation or other components. All algorithms have more or less

converged after 50 iterations, and the expansions found by the TZ and Constrained EM are quite different, so improved performance is not due to simply slower convergence toward a similar solution.

5 Discussion and Future Directions

Although we believe that most computational effort for modeling query intent is appropriately spent training large-scale models offline, Web search engines must operate an increasingly complex decision environment, and so we also foresee that an on-line learning component that solves per-query optimization problems in real time could be a powerful complement to offline training. Such online components could perform context-sensitive noise reduction or feature weighting on predictor inputs, or ‘course correction’ via posterior constraints on their output, incorporating the kind of user interaction feedback, or new intent evidence that is only visible at query time.

This paper is a first step in exploring useful constraints and objectives for such query-focused optimization algorithms. Because of its simplicity and flexibility, we believe that a Constrained EM framework is a good starting point to explore estimation algorithms for unified probabilistic models that utilize data from multiple queries and documents [1], or multiple sources of evidence for inferring user intent. Beyond words, our feature set could be extended to richer query representations that include proximity or specialized operators. We believe that a general trend in software systems is that simpler and more powerful algorithms are eventually preferred over methods designed for efficiency in special cases, so that the advantages of flexible optimization frameworks in information retrieval systems will soon outweigh their moderate computational costs.

6 References

- [1] D. Bodoff and S. E. Robertson. A new unified probabilistic model. *JASIST*, 55(6):471–487, 2004.
- [2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [3] K. Collins-Thompson. Reducing the risk of query expansion via robust constrained optimization. In *Proceedings of CIKM 2009*, pages 837–846.
- [4] K. Collins-Thompson. Estimating robust query models using convex optimization. In *Advances in Neural Information Processing Systems 21 (NIPS)*, 2008.
- [5] K. Collins-Thompson and J. Callan. Query expansion using random walk models. In *CIKM 2005*, pages 704–711.
- [6] J. Dillon, Y. Mao, G. Lebanon, and J. Zhang. Statistical translation, heat kernels, and expected distances. In *Proc. of UAI 2007*, 2007.
- [7] J. Graça, K. Ganchev, and B. Taskar. Expectation maximization and posterior constraints. In *NIPS*, 2007.
- [8] J. D. Lafferty and C. Zhai. Document language models, query models, and risk minimization for information retrieval. In *SIGIR '01: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 111–119, 2001.
- [9] Lemur. Lemur toolkit for language modeling & retrieval. 2002. <http://www.lemurproject.org>.
- [10] Q. Mei, D. Zhang, and C. Zhai. A general optimization framework for smoothing language models on graph structures. In *Proceedings of SIGIR 2008*, pages 611–618, New York, NY, USA, 2008. ACM.
- [11] T. Tao and C. Zhai. Regularized estimation of mixture models for robust pseudo-relevance feedback. In *SIGIR 2006*, pages 162–169.
- [12] R. Udupa, A. Bhole, and P. Bhattacharya. A term is known by the company it keeps: On selecting a good expansion set in pseudo-relevance feedback. In *Proceedings of ICTIR 2009*, Advances in Information Retrieval Theory. Springer, 2009.
- [13] Z. Xu and R. Akella. A new probabilistic retrieval model based on the dirichlet compound multinomial distribution. In *SIGIR '08*, pages 427–434, New York, NY, USA, 2008. ACM.