

A Method to Incorporate Temporal Seasonality into Search Ranking for Improved Relevance and User Engagement*

Shreya Mahapatra¹, Anandita Chopra¹, Liping Zhang¹, Shabhareesh Komirishetty¹ and Tracy Holloway King¹

¹Adobe Inc. (USA and India)

Abstract

Search behavior is often influenced by seasonal trends such as Christmas, Halloween, and other recurring cultural or temporal events. During these periods, users tend to search for specific topics or content aligned with the seasonal moment. However, traditional ranking systems typically overlook these signals. In this paper, we devise a framework for identifying and incorporating seasonality signals into search ranking models. Seasonality signals capture recurring patterns in user behavior over time derived from historical interaction data. Our approach detects these recurring patterns using a forecasting model and integrates them into the retrieval and ranking stages of the search pipeline. Using seasonality, the system can predict and surface content that aligns with emerging user interests based on the current temporal and regional context. Furthermore, our method accounts for cultural variations across regions, allowing more localized and relevant seasonal experiences.

Keywords

multi-modal search, seasonality and temporality, search ranking

1. Introduction

In search systems for images, templates, and other multi-modal content, ranking approaches primarily rely on embedding-based retrieval and user behavioral signals. These methods focus on measuring the semantic similarity between the user query and candidate content such as image. In addition, multi-modal search also incorporates engagement metrics such as clicks, downloads or exports to promote popular or trending results. Although these methods work well for capturing contextual matching and content relevance, they are essentially reactive. These methods model what users have interacted with in the past, rather than anticipating what users are likely to look for in the near future. As a result, such systems often fail to capture temporal intent or predict future user interests, especially those driven by recurring seasonal trends.

The predictive approach is the main emphasis of this paper, where we attempt to forecast future user behavior. We propose a framework that forecasts user behavior by leveraging seasonality signals. Seasonality signals are crucial for understanding and predicting seasonal fluctuations in various contexts, such as user behavior or content relevance. By analyzing seasonal trends and patterns, we can tailor search results and recommendations to align with current and upcoming events, such as holidays or special occasions. This ensures that users receive suggestions that match their immediate needs. Our work is in the context of Adobe Express template search which provides users with multi-modal templates that they can edit and remix for use on social media and printing (see Figure 1 for an example); all examples in this paper are in that context. For instance, during the winter holidays, users searching for “holiday invitations” or “party flyers” will receive tailored recommendations featuring Christmas, New Years or winter. Similarly, as summer approaches, users looking for “event banners” or “vacation photo templates” will be presented with vibrant, summer-themed options. For events these predictable seasonal events, we can dynamically adjust search results and recommendations to feature relevant templates, ensuring that users find the most appropriate and timely content for their needs. Furthermore, our work considers the user’s cultural context (or their country). For example in mid-June, US results

ECOM’25: SIGIR Workshop on eCommerce, Jul 17, 2025, Padua, Italy

*shmahapa@adobe.com, anan@adobe.com, lipingz@adobe.com, skomirishett@adobe.com, tking@adobe.com



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

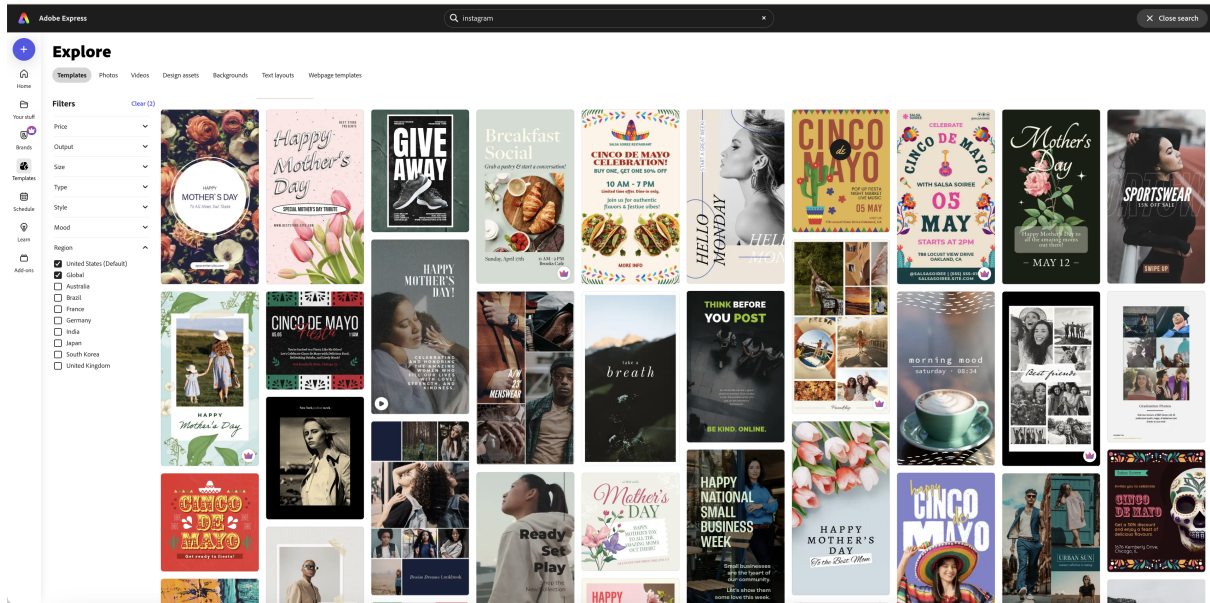


Figure 1: Adobe Express results for the query *instagram* from late April in the US. In addition to generic Instagram results, seasonal results for Mother’s Day (May 11th) and Cinco de Mayo (May 5th) are highly ranked. (<https://new.express.adobe.com/>).

show content related to US Independence Day (July 4th), while India results how content linked to Yoga Day (June 21st). Our work is divided into parts:

- Predicting the seasonality signal
- Incorporating the seasonality signal into the search retrieval and ranking

This work was conducted to improve search ranking of templates in Adobe Express. Adobe Express is a content creation tool developed by Adobe, where templates are rich objects which contain many visual layers and text boxes [1]. The platform follows a freemium model, i.e. users with free accounts have access to free content plus premium that has to be paid for individually, while paid users have unlimited access to both free and premium content. This tiered structure affects user engagement patterns, highlighting the need to analyze and incorporate seasonal trends across different user segments. The proposed algorithm is generic and can be applied to other search ranking systems.

2. Prior Art

There has been recent research on enhancing the search relevance by incorporating seasonality signal. Chen et al. [2] addresses the challenge of incorporating seasonality into online product search by utilizing a season enhanced BERT model to capture seasonal semantic relevance between queries and product. However, the paper only uses a six-season classification system to define seasons for relevance modeling. It does not incorporate seasonal moments such as holidays or festivals into its model. Verma et al. [3] uses a feed forward neural network model to predict seasonality scores based on query text and month. It then uses a L2 ranker to integrate the predicted seasonality scores and re-rank auto complete suggestions. This model works at a monthly granular level.

In contrast, our work uses a time-series model like TBATS [4] to model seasonal trends (section 3.4). Moreover, our work operates at a daily granular level, thereby capturing finer seasonal patterns. Additionally, our intent taxonomy includes a wider diversity of seasonal intents such as major festivals and holidays. Unlike prior work that focuses on global seasonal pattern, our model effectively captures seasonality trends at the country level.

3. Model and Data

3.1. Template Data

Adobe Express has a rich corpus of templates. Templates are editorially curated, customizable layouts present within Adobe Express. A template can be thought of as a multi-modal asset that includes shapes, icons, text, and images. In addition to the structured visual content, each template is associated with a curated title. Insights regarding the use of templates are also provided via user engagement metrics like clicks, exports, edits, and so forth.

To predict the user behavior, we first devise an algorithm to classify the templates into relevant seasonal intents. Our taxonomy [5] includes ~300 intents, out of which 135 are categorized as seasonal, covering both global events (e.g. New Years Eve, summer (which takes into account northern and southern hemispheres)) and culturally specific local events (e.g. Diwali, Chinese New Year). While our current taxonomy is manually curated, our framework is designed to be easily extensible to new intents.

3.2. Template Intent Classification

We classify the templates into intents using the intent taxonomy described above. An unsupervised learning technique is employed for this classification, leveraging textual information and metadata such as title. Sentence transformers generate embeddings for both the template and the predefined intent taxonomy, where each intent is treated as a textual description. To determine the most relevant intent for each template, we compute the cosine similarity between the template embedding and each intent embedding. The template is assigned the intent corresponding to the closest match. Finally, the assigned intent is encoded as a 512-dimension one hot vector, where a non-zero entry corresponds to the confidence of the most likely intent, with additional capacity included for future seasonality.

3.3. Few Shot Learning with Sentence Transformer

Each intent described in section 3.2 is associated with a set of templates. We select ~24000 templates across different intents. The templates with high confidence score are chosen as the base dataset for the training. Each template is associated with the textual description and an intent which acts as its label. This dataset comprises templates from varying geographical regions and languages.

SetFit [6] is an efficient few shot tuning algorithm for sentence transformers. It achieves high accuracy with relatively little labeled data and so is appropriate for our use case which requires a model that can accurately classify the user behavior into different intents with little labeled data. We trained a SetFit model with the Adobe-licensed templates' textual data utilizing the dataset described in section 3.1. The dataset was divided into two parts based on the locale of the content used — English and non-English (multi-lingual) training data. We utilized the 'sentence-transformers/all-MiniLM-L6-v2' model as our base model and trained it on Adobe-licensed templates' English textual data. Similarly, 'sentence-transformers/paraphrase-multilingual-MiniLM-L12-v2' was utilized for non-English textual data. The CoSENTLoss (Contrastive Sentence Embedding Loss) [7] loss function was employed for training.

The model achieved an accuracy of 0.85 on the multi-lingual test set and an accuracy of 0.95 on the English test set. This fine-tuned model understands the Adobe Express template data and can accurately predict the user behavior since users typically remix these templates to create their projects.

3.4. Seasonal Forecast Prediction

Seasonal forecasting is performed using the TBATS algorithm [4], which is particularly effective for time series data with complex seasonal patterns. TBATS takes as input a time series of intents from the past four years.¹ We extract the sequence of project-level intents created by users aggregated over time and converted it to a time series data. For each intent, we construct a time series representing the

¹Four years was selected based on the availability of accurate user behavioral data in Express template search.

count of the occurrence of each intent over the four year window. The TBATS model learns patterns from historical trends and predicts the seasonal intent for a time period. This modeling is performed for each geographical region (e.g. larger regions like the United States, Europe and Asia-Pacific and/or specific countries). Thus, our model predicts the relevant seasonal moment and intent for a given date and region. The implementation steps are outlined below:

1. For each pre-defined region:
2. Select the last four years of data for prediction.
3. Predict the intent of the user behavior data using the model defined in section 3.3.
4. Filter out intents with a very low confidence value to reduce noise. Only user behavior data with high confidence value are chosen.
5. Aggregate the predicted intent data on a weekly basis, using Monday as the first day of each week as depicted in table 1.

Table 1
Aggregation of Intents

Week	Intent	Count
Week 1	Intent 1	4000
	Intent 2	8000
	Intent 3	2000
	Intent 7	20
Week 2	Intent 1	4000
	Intent 2	8000
	Intent 4	10000

6. Compute z-score for each intent’s occurrences via formula (1).

$$\frac{\text{count} - \text{mean}}{\text{standard deviation}} \quad (1)$$

7. Use the TBATS model with a seasonal period of 52.18 weeks. The TBATS model predicts the forecasts for the next 104 weeks or 2 years. Forecast scores are generated for each intent across weekly intervals.
8. Threshold filtering is applied to the predicted forecast score to focus on statistically significant seasonal patterns. Specifically, only the intents with scores exceeding the 90th percentile are chosen.
9. Select the top most likely forecasted intent for each week.

The prediction also covers seasons (i.e. spring, summer, fall (autumn), and winter) across the world. That is, it understands seasons based on geo-location and date. We derive the value of the current ongoing season for all top-tier countries based on the hemisphere and date. For instance, if the date is March 2nd and the country is Japan, we would categorize the season as ‘spring,’ since Japan is in the northern hemisphere. In contrast, a user from a country in the southern hemisphere on the same date would experience ‘fall’, reflecting the opposite seasonal pattern. These seasons are added to the list of forecasted intents from the model.

The algorithm takes into account the cultural aspects of a country. There are certain festivals which have greater significance in a particular country. For instance, both International Women’s Day and Holi occur in March. International Women’s Day is a global event, whereas Holi has more cultural significance in India. Therefore, Holi should be ranked higher than International Women’s Day in India. The system ranks culturally significant intents higher than global events. This is achieved by manually tagging certain intents from the taxonomy as country specific and which countries they apply to. Moreover, we leverage Large Language Models (LLMs) to predict the cultural significance of the intent for a given country. As a result, the cultural significance of an intent is derived by a combination of manual curation and LLM’s knowledge.

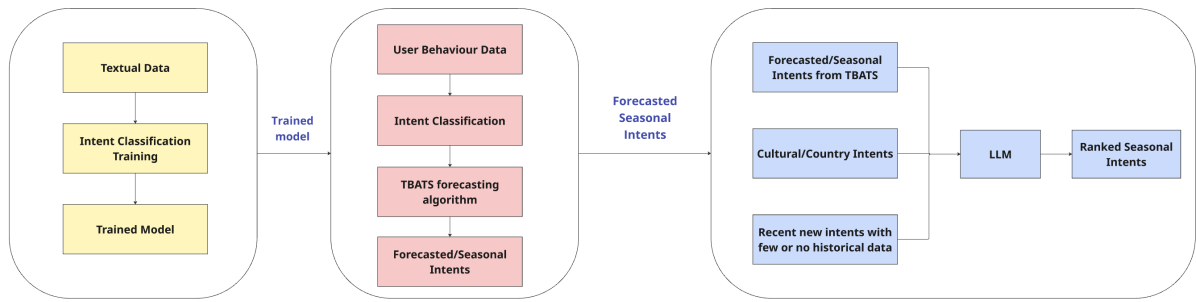


Figure 2: High level overview of Seasonal Forecast Prediction




Figure 3: Seasonal intent distribution for *Independence Day* and *New Year's Day* in the United States

Some events exhibit limited historical behavior but demonstrate strong relevance because of recent user behavior signal (i.e. trending signals). The system incorporates these signals by manual addition of new intents based on region and date. For intents not captured by the TBATS algorithm due to limited historical data, we leverage LLMs to estimate their relevance scores. The LLM is provided with an input of existing intents along with their relevance score for the specified region and date, i.e. the ones forecasted by the TBATS algorithm for the given date and region. Based on the provided context, the LLMs generate a relevance score for the new intent. If a new trending signal, which has limited historical data, emerges, we simply need to include the new intent 3.1, re-classify the templates and apply the same LLM-based technique to forecast its seasonality score.

The overall seasonal forecasting system described above is shown in Figure 2

The figure 3 shows the distribution of Independence Day and New Year in the US over the past two years. As expected, the intents exhibit a cyclical pattern, becoming more prevalent as their actual date approaches.

4. Seasonality in Search

Adobe Express template search leverages a RankNet-based learning-to-rank model [8] to determine the relevance of templates. This model is trained on user engagement metrics such as clicks and exports, the content embedding and the manually curated tags and title (see [1] on how the embedding-based features are combined with user behavior and metadata features for Express matching and ranking). These features help the model to learn user preferences and rank templates accordingly. However, the template ranking is based on the past interaction of the users and cannot anticipate the future usage of the templates, including seasonally predictable preferences (see Figure 1 where a broad query like *instagram* should include seasonally relevant templates, not just historically popular ones).

To further enhance the search experience, this work incorporates a seasonality signal into the current ranking framework. This signal, as derived from the forecasting algorithm described in section 3.4, forecasts the relevant seasonal moment at a particular time of the year. Integrating the seasonal signal into the search ranking not only reflects historical usage preferences but also anticipates the seasonal usage, leading to a more context-aware search experience.

4.1. Dataset and Training

Each row of the training dataset consists of a unique combination of query, country and date. The columns consist of features such as user engagement metrics, template embeddings and metadata-related fields. The dataset is enhanced by including the seasonality signal. The seasonal signal is incorporated through the following steps:

1. For a given date and country, predict the seasonal intents using the forecasting model (section 3.4). This results in a multi-hot vector, where each non-zero value corresponds to a specific intent and represents the model's confidence score for that intent.

Example: Given a user searching for a template on December 24th in the US, the forecast model may identify high-relevance seasonal intents such as *Christmas* (0.8), *Winter* (0.5), and *New Year* (0.7). These are encoded as a seasonal intent vector with several intents with confidence scores but many zeros representing extremely unlikely intents for that date:

$$[0, 0.8, 0, 0, 0.5, 0, \dots, 0, 0.7, 0, \dots, 0]$$

2. Each template is associated with a seasonal embedding, represented as a one-hot vector (section 3.2). The non-zero value in this vector represents the confidence of the intent.

Table 2

Template Seasonal Embedding

Template ID	Seasonal Embedding
template_id_1	[0.6, 0, 0, 0, ...]
template_id_2	[0, 0.8, 0, 0, ...]
template_id_3	[0, 0, 0, 0.5, ...]

3. To compute the seasonal relevance of a template for a given query, we take the dot product between the template's seasonal embedding vector and the forecasted intent vector generated for the query's date and country. The dot product serves as the scalar score, signifying how well a template aligns with the query's forecasted intent. A higher score signifies higher alignment with the forecasted intent.
4. Seasonal events are more relevant for broad, high level queries such as *flyer*, *invitation*, *sale* or *poster*, where users are looking to create event-driven content. To effectively capture this seasonal relevance, the existing training dataset is enhanced by including more broad, high level queries. This is achieved by filtering the behavioral dataset to include more broad queries, thereby increasing the model's exposure to seasonally influenced search behavior.

5. This dataset is further enriched by combining it with other user engagement signals. This data, combined with seasonality, provides a good balance between showing templates that are seasonal, relevant to the user’s context and recently created. The user engagement signals include:
- Recency: Captures how recently the template was published.
 - Clicked Template Id and BM25 Scores: Captures the textual relevance between the user’s query and the clicked templates’ metadata (e.g. title, content) by leveraging BM25 scores.
 - Template Embeddings: Encoding each template into a vector space to capture the semantic relationship between the user’s search query and template [1].
 - Template Metadata: This includes key attributes which provide contextual information about the template, such as: User Segments (Enterprise, Education K-12), Task type (instagram-story, poster), region (US, India, France etc). These attributes are associated with a template and signify its intended audience and purpose.

We trained a linear SVM model using the RankNet loss function in PyTorch. The clicked template Ids correspond to the positive samples, while negative samples are drawn from the recall set of the search query that were not clicked. The resulting model produces a weight score corresponding to each feature, including seasonality, based on their contribution to ranking relevance.

4.2. Applying Seasonal Boost at Query Time

At query time, for a given query’s date and country, the forecast model returns the forecasted intent. To incorporate this seasonal information, we compute a weighted boost as follows:

$$\text{Seasonal_Boost} = \text{Forecast_Confidence} \times \text{SVM-learned_Weight} \quad (2)$$

The weighted score reflects both the seasonal alignment as well as the learned impact on user behavior. During template re-ranking, this boost is applied to rescore the template’s score. A dot product is computed using the weighted forecast intent vector and template’s seasonal embedding to derive a seasonal relevance score. A higher dot product indicates a greater match with the current date forecasted intent, and hence boosts the template’s ranking in search results.

The seasonality boost is applied along with other factors, which strike a balance between seasonal relevance, contextual relevance, and recency. The adjustments help provide a diverse user experience by promoting newly created templates along with the seasonal and contextually relevant templates.

5. Evaluation

5.1. Offline Large-scale Evaluation

Our offline large-scale evaluation uses the historical query logs to determine the effect of the new ranking on the user experience. We computed Mean Reciprocal Rank (MRR) from 15000 randomly sampled search queries. Table 3 shows the impact of adding the seasonality signal on MRR in the US region. For all query frequencies (head, torso, tail) the MRR increases, suggesting that adding seasonality would reduce the amount of scrolling that users would have to do to find relevant results. Both with and without seasonality, MRR is higher for more frequent queries. This at first seems counter-intuitive since broad head queries have many relevant results that users have to inspect to select from. However, the relatively limited selection of templates for tail queries combined with the fact that templates are editable and hence can be easily modified by users means that users have to scroll to find the ideal template for niche intents, while for head queries they can select from the top results.

5.2. Offline Manual Evaluation

We conducted an internal side-by-side evaluation, where users, including Adobe product managers, were asked to label which search experience they preferred based on the top 20 templates for a randomly

Table 3
Impact of Seasonality on MRR in the US

Query Type	MRR without Seasonality	MRR with Seasonality
Head	0.3871	0.391
Torso	0.25	0.267
Tail	0.172	0.186

selected set of queries. Table 4 shows the statistics of the evaluation. Incorporating seasonality into search improves the search experience, with 56% of search queries indicating that the seasonality enhanced search experience was much better.

Table 4
User Evaluation Results for Seasonality-Enhanced Search

Better Without Seasonality	Better With Seasonality	Both Relevant	Both Irrelevant
9%	56%	22%	13%

6. Conclusion and Future work

By integrating seasonality signals into the search ranking pipeline, we aim to enhance the search experience by providing more seasonal content, thereby diversifying the content in the search. This method not only enriches the search experience, but also provide a scalable and robust method to extract seasonality signal and incorporating them into various systems. Many Adobe Express users leverage templates to create social media content for their e-commerce business. In the context of e-commerce, seasonally relevant search results are particularly helpful because seasonal moments such as back to school, holidays, and regional festivals greatly influence user purchasing behavior and intent. By enhancing search results with seasonality signals, we can anticipate user needs and surface relevant results based on the seasonal pattern.

Our next steps are to run an AB test for Express template search to determine whether user behavior on Adobe Express reflects the offline predictions (section 5). In addition, we are integrating the seasonality signal into template collection browse. These template collections are similar to broad queries (e.g. collections for Instagram Square Posts and for Posters) and should benefit from having seasonally-relevant templates as opposed to ranking based on popularity and recency.

As future work, this model can be extended to other Adobe products such as Adobe Stock, which is an image and video marketplace. Incorporating seasonality into Adobe Stock could help to surface more contextually relevant seasonal images and videos. Moreover, the current model only leverages text for intent classification. We can leverage multi-modal models that can incorporate visual features to classify templates and user behavior prediction. Furthermore, the seasonality signal can be incorporated into many other workflows such as auto complete, and email campaigns.

Declaration on Generative AI

During the preparation of this work, the author(s) have used ChatGpt for Grammar and spelling check. After using these tool(s)/service(s), we have reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] C. Aroraa, T. H. King, J. Kumar, Y. Lu, S. Sharma, A. Srikantan, D. Uvalle, J. Valls-Vargas, H. Vardhan, Smart multi-modal search: Contextual sparse and dense embedding integration in adobe express, in: Proceedings of The 1st Workshop on Multimodal Search and Recommendations (CIKM MMSR '24), CEUR, 2024. URL: <https://ceur-ws.org/Vol-3859/paper3.pdf>.
- [2] H. Chen, Y. Chen, J. Meng, Y. Jiao, Y. Ni, Y. Gao, M. Momma, Y. Sun, Improving product search with season-aware query-product semantic similarity, in: Companion Proceedings of the ACM Web Conference 2023, 2023, pp. 864–868.
- [3] P. Verma, S. Zhong, X. Liu, A. Rajan, Seasonality based reranking of e-commerce autocomplete using natural language queries, arXiv preprint arXiv:2308.02055 (2023).
- [4] A. M. De Livera, R. J. Hyndman, R. D. Snyder, Forecasting time series with complex seasonal patterns using exponential smoothing, Journal of the American statistical association 106 (2011) 1513–1527.
- [5] S. Sharma, T. H. King, J. Kumar, M. Poddar, K. Blank, Augmenting knowledge graph hierarchies using neural transformers, in: Proceedings of ECIR24 Industry Track, 2024.
- [6] L. Tunstall, N. Reimers, U. E. S. Jo, L. Bates, D. Korat, M. Wasserblat, O. Pereg, Efficient few-shot learning without prompts, arXiv preprint arXiv:2209.11055 (2022).
- [7] S. Jianlin, Cosent: A more efficient sentence vector scheme than sentence-bert, 2022. URL: <https://kexue.fm/archives/8847>.
- [8] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, G. Hullender, Learning to rank using gradient descent, in: Proceedings of the 22nd International Conference on Machine Learning (ICML), ACM, New York, NY, USA, 2005, pp. 89–96. URL: <https://doi.org/10.1145/1102351.1102363>. doi:10.1145/1102351.1102363.