

Retaining Data from Streams of Social Platforms with Minimal Regret

Nguyen Thanh Tam¹, Matthias Weidlich², Duong Chi Thang¹,
Hongzhi Yin³, Nguyen Quoc Viet Hung^{4*}

¹ École Polytechnique Fédérale de Lausanne, ² Humboldt-Universität zu Berlin,
³ The University of Queensland, ⁴ Griffith University
{tam.nguyenthanh, thang.duong}@epfl.ch, matthias.weidlich@hu-berlin.de,
h.yin1@uq.edu.au, quocviethung1@gmail.com

Abstract

Today’s social platforms, such as Twitter and Facebook, continuously generate massive volumes of data. The resulting data streams exceed any reasonable limit for permanent storage, especially since data is often redundant, overlapping, sparse, and generally of low value. This calls for means to retain solely a small fraction of the data in an online manner. In this paper, we propose techniques to effectively decide which data to retain, such that the induced loss of information, the regret of neglecting certain data, is minimized. These techniques enable not only efficient processing of massive streaming data, but are also adaptive and address the dynamic nature of social media. Experiments on large-scale real-world datasets illustrate the feasibility of our approach in terms of both, runtime and information quality.

1 Introduction

Current social platforms such as Twitter, Facebook, and Yelp, produce data streams with an unprecedented rate. For example, about half a billion tweets are generated every day [Mathioudakis and Koudas, 2010]. To make sense of these streams, one can typically only retain a small fraction of the data for analysis, due to storage limits and the cognitive load induced by the sheer data volume.

Against this background, traditional methods for the analysis of social platforms perform data summarization, e.g., based on relevance detection or measures of information diversity [Chang *et al.*, 2016]. However, these approaches are inherently limited to a static setting: The data is crawled and stored, before the top- k most important data items are selected as a data summary. Even if feasible, such an approach incurs high storage cost and does not avoid the problem of retaining only a fraction of the data once its volume exceeds a storage limit.

In this paper, we consider the natural setting of social platforms, where data is dynamic and available as a stream. Then, retaining of data becomes more challenging compared to one-off summarization, as data selection has to be repeated every time new data arrives. Instead of considering the whole historical data, summarization now works on the retained data

(i.e., a previous summary) and the new data. This further degrades the informativeness of the original data since the summary of the retained data already induces some loss of information. Specifically, the lack of historical data leads to a biased assessment of data importance: Data that was considered unimportant and thus discarded in the past may retrospectively turn out to be important.

To minimize the regret of discarding important data, two requirements have to be met. First, a compact data sketch needs to be maintained, in addition to the actual data summary, to capture the long-term history of data and enable a precise assessment of data utility over time. Yet, for data stemming from social platforms, this sketch needs to be adaptive to changes in the data stream. Second, a protocol needs to be specified to decide which data items to retain and to discard, such that the total regret in data utility is minimal. To cope with the data stream volume and velocity of social platforms, this protocol needs to be very efficient.

In this paper, we tackle these requirements and propose a novel statistical model, which does not only capture the traditional context of social data (importance of topics, user influence, information diffusion), but also embeds the dynamics of this context over time. For example, topics are not considered to be static. They may emerge or disappear over time and relate to recurring events. Striving for online processing of streaming data, we develop a scalable learning mechanism to quickly update the model with new data. We further show how the statistical model is used to define a utility function to assess the representativeness of a data summary. Minimizing the regret of discarding data then becomes the problem of minimizing the difference between the utility of the retained data and the utility of whole historical data. Finally, we present a progressive algorithm to select which data to retain, with guarantees on the induced regret factor. This algorithm scales linearly in time and space, solely in the size of the data summary (not the whole data stream).

Our contributions and the paper structure are summarized as follows. After outlining the retaining problem with minimal regret (§2), we present (i) a statistical model to sketch data properties; (ii) a utility function to assess data representativeness (§3); and (iii) a progressive algorithm to solve the retaining problem (§4). Using diverse datasets derived from Twitter, we demonstrate improvements of five orders of magnitude in efficiency and up to 42% in information quality of

*Corresponding author

our approach over state-of-the-art baselines (§5). We then review related work (§6) and conclude the paper (§7).

2 Problem Statement

We model the stream of data stemming from a social platform by an infinite set of textual data items $E = \{e_1, e_2, \dots\}$. The items are totally ordered based on their occurrence time, denoted by the subscript. By $E_t = \{e_1, \dots, e_t\}$, we denote the set of items until time t . Acknowledging that not all items from E can be stored permanently, a representative subset of E of size k shall be retained. Here, the parameter k depends on the application context and typically reflects the storage limit. We further postulate a non-negative function $f : 2^E \rightarrow \mathbb{R}_{\geq 0}$ to quantify the utility of a set of items $S \subseteq E$, capturing how well S represents E according to some objective. Given that E is continuously extended with new data items, the *retaining problem with minimal regret* is to select k items, such that the regret ratio—the normalized difference between the utility of the retained items and the utility of the whole data stream—is minimal.

Problem 1 (Retaining Problem with Minimal Regret). *Given a data stream E_t until time t , a current set of retained items $S_t \subseteq E_t$, a window of new data items $W = \{e_{t+1}, \dots, e_{t+|W|}\}$, the problem is to construct a new set of retained items S_{t+w} , such that:*

$$S_{t+w} = \arg \min_{S \subseteq (S_t \cup W), |S|=k} \frac{f(E_t \cup W) - f(S)}{f(E_t \cup W)}. \quad (1)$$

The problem setting is illustrated in Figure 1. Upon the arrival of a window of new items, the set of retained items is updated. This is done by selecting items from the old set of retained items and from the window.

The figure further illustrates why the retaining problem with minimal regret cannot be addressed by applying traditional data summarization each time a window of new data items arrives. Traditional summarization would consider solely the current set of retained items and the items of the new window. Yet, the data stream history in terms of items discarded in the past would be neglected. Consequently, when constructing a new set of retained items, the utility of possible candidate sets cannot be assessed accurately.

As illustrated in Figure 1, therefore, a concise sketch of historical data needs to be maintained. It captures essential properties of data items that have been discarded in the past, thereby enabling an accurate assessment of the regret ratio of a potential set of retained items. To realise such a sketch, §3 presents a statistical model and also shows how to assess




Figure 1: Illustration of the retaining problem ($k = 5$, $|W| = 3$).

the representativeness of an item set using an utility function. In §4, we then present a progressive algorithm to solve the retaining problem with minimal regret under this model.

3 Model and Approach

Below, we first propose a statistical model to sketch the historical data of a stream, before turning to the question of how to assess the utility of a set of retained data items. Finally, we discuss a simple strategy to solve the retaining problem.

3.1 A Statistical Model for Social Data

For textual data items that originate from social platforms, topics are a fundamental concept to understand the co-occurrence relations of words [Cha *et al.*, 2013]. We thus capture information on topics as a statistical means to select representative data items from important clusters of words. However, the dynamic nature of streaming data from social platforms prevents us from knowing a specific distribution of topics in some future state. Rather, we face the following phenomena:

- *Emergent topics*: In a streaming setting, new topics may emerge over time. Hence, topic modelling techniques such as LDA [Blei *et al.*, 2003] or pLSI [Hofmann, 1999] that fix a pre-defined number of topics are not applicable. A small number of topics may lead to information loss, as different words might end up in the same cluster, whereas a large number of topics may imply sparse clusters, destroying data regularities.
- *Emergent vocabulary*: An evolving collection of topics implies that the vocabulary of words changes over time. Words may only be invented at a specific time [Zhai and Boyd-Graber, 2013] (e.g., ‘brexit’ during the events in the UK in June 2016). A fixed vocabulary as in traditional topic models does not capture these dynamics and tends to be inefficient due to unused and redundant words.
- *Recurring topics*: Traditional data summarization typically spans a short period of time [Chakrabarti and Punera, 2011], due to data storage limits. In contrast, when processing data streams of social platforms, a long history of data items is considered, so that topics will recur over time [Herrera-Damas, 2014] (e.g., the topic of ‘football’ shows seasonal patterns). Such effects influence the decision of which data items to retain.

Against this background, we propose a non-parametric [Hollander *et al.*, 2013] probabilistic model to sketch the properties of past data items. It features a potentially unbounded number of topics and words that are learned from textual data items over time. It also considers recurrent topics by means of temporal cluster variables, for which the time granularity can be customised.

Formally, a textual data item e is modelled as a multiset of words $\{v_1, \dots, v_{|e|}\}$, where v_i is a word from a dynamic vocabulary V . A multiset of words further defines a semantic topic z . We model the dynamics of words per topic by means of a Dirichlet process to generate the word distribution ϕ_z and the vocabulary ρ_z of a topic z . Further, at time t , a topic distribution θ is used to generate the topics of a new data item e_t . To ensure that the temporal aspect of evolving topics is reflected in the generation, we use a hierarchical

$$f(S) = \sum_{e \in S} \sum_{v \in V} n(v, e) p(v, e) \log \frac{1}{p(v, e)} g(e) h(e) \quad (2)$$

Using this formulation, an algorithm to select data items will prefer sets with high entropy, i.e., sets with items that cover diverse topics and preserve the evolving topic distribution. As proven in our technical report [Nguyen et al., 2017], the above measure is monotonic (selecting more items increases utility) and submodular (marginal gains by selecting more items start to diminish due to saturation of the utility objective).

Proposition 1. $f(\cdot)$ is a monotonic function.

Proposition 2. $f(\cdot)$ is a submodular function.

3.3 A Simple Retaining Algorithm

A straight-forward approach to solve Problem 1 under the above model applies traditional data summarization [Zhuang et al., 2016] on the retained items and the content of a new window. Then, the new set of retained items is selected as:

$$\max_{S \subseteq S_t \cup W, |S|=k} f(S) \quad (3)$$

which is equivalent to Eq. 1 since the value of $f(E_{t+w})$ is constant in terms of selecting any S . However, this problem is known to be NP-complete [Hung et al., 2015]

Due to the computational complexity, greedy approximation algorithms (inspired by the knapsack problem) are commonly employed. They start with the empty set $S^{(0)} = \emptyset$, and at each iteration i over the current data, choose an item $e \in S_t \cup W$ maximizing the utility, i.e.,

$$S^{(i)} = S^{(i-1)} \cup \arg \max_{e \in S_t \cup W} f(S^{(i-1)} \cup e) - f(S^{(i-1)}) \quad (4)$$

For a monotonic and submodular function (as our utility function defined above), this greedy algorithm yields a $(1 - 1/e) \approx 0.63$ approximation [Nemhauser and Wolsey, 1981]. However, this algorithm has an update time complexity of $\mathcal{O}(k(k + |W|))$, which is undesirable for streaming applications. Also, the update process needs to be repeated every time new data arrives.

4 A Progressive Retaining Algorithm

Given the above results on hardness of the retaining problem and the time complexity of a simple greedy algorithm, we now present a progressive algorithm. It is tailored to the stream processing setting and shows linear time and space complexity. The algorithm comprises of (i) an incremental inference mechanism to update the sketch of historical data; and (ii) a mechanism to select a new set of retained items.

4.1 Updating the Data Sketch

To update the parameters Θ of our model, we realise an online learning mechanism. When data arrives, we compute the observed variables and propagate back the information to the model parameters. Once the parameters have been updated, the conditional and marginal probabilities for the utility function are computed following the generative process.

Many online learning techniques have been developed based on Markov chain Monte Carlo sampling (e.g., incremental Gibbs sampling [Canini et al., 2009]). However,

these techniques either reduce model complexity (losing the guarantee to converge for the complete model) or have a space complexity that is linear in the number of items to analyse [Hoffman et al., 2010]. To overcome these issues, we rely on stochastic variational inference [Hoffman et al., 2013]. Here, the idea is to minimize the evidence lower bound (ELBO) of the expected difference between the observed distribution and the latent distribution, defined as:

$$L(Z) = \mathbb{E}_{q(Z)} \left[\ln \frac{p(Z, E)}{q(Z)} \right]$$

where Z is the set of all latent variables in the model (except model parameters and observed variables); and $q(Z)$ is an approximate distribution of $p(Z|E)$, which can be factorised over the distributions of model parameters in the generative process. Then, we apply stochastic optimisation to the ELBO function over a data stream, which basically updates the new parameter values from the previous ones following the direction of the ELBO gradient of new data with a fixed step size. The more data is received, the more the model parameters will converge to minimize the ELBO function.

Instead of single-item update, our approach considers multiple observations per update to reduce noise, which also aligns with window-based processing to avoid order distortion in data streaming settings. Receiving data as a series of windows W_b , $b = 1, 2, \dots$, of items, we proceed as follows:

1. We update the local parameters of the variational distribution of the word-topic variable z and the topic-cluster variable c . This requires us to maintain additional M_z local parameters for possible values of z , and M_c local parameters for possible values of c . Here, M_c, M_z are ‘prior beliefs’ on the maximum number of topics and recurring topic clusters. Yet, their effects are marginal, as the updates are dominated by the observed information, so that they can be safely set to large constant values (e.g., 1000).
2. We compute the natural gradients using previous parameter values $\Theta^{(t-1)}$ and the above local parameters of the ELBO function decomposed over each item $e \in E_b$ [Hoffman et al., 2010]. Formally, we obtain $\nabla_{\Theta^{(t-1)}} L_e$ as a vector of gradients for each parameter in $\Theta^{(t-1)}$.
3. The new values for model parameters are computed from their previous value:

$$\Theta^{(t)} = \Theta^{(t-1)} + w_b \frac{1}{|W_b|} \sum_{e \in W_b} \nabla_{\Theta^{(t-1)}} L_e \quad (5)$$

where w_b is the learning rate to control the learning quality and convergence of the inference. w_b is often modelled as a power function of b with a forgetting rate r [Hoffman et al., 2013]. Setting $r \in (0.5, 1]$ guarantees convergence, while larger values often lead to higher learning quality and faster convergence (but not monotonically).

Note that windows of large size reduce the number of updates, but may lead to a poor estimation of model parameters. To further improve scalability, updating of model parameters can be parallelised by exploiting the conditional independence property. When the global variables (i.e., the most outer parameters in the model) are given, the updates to local variables (i.e., inner parameters) become independent and can

thus be computed concurrently. Also, the computation of semantic information, decay in interestingness, and social features per data item, see §3.2, is independent once the model parameters are updated and thus can be parallelised.

4.2 Retaining Data Items

We now turn to the selection of data items to retain. In essence, at each step, a new set of items is selected as the one with the highest utility among all candidate sets. The candidate sets are created by swapping at most one new item with a retained item, if the utility increases. However, the arrival of a new item may change the model parameters, influences the decay in interestingness of old items, and may introduce new social features. Hence, the utility of the retained set of items has to be updated, before the swapping procedure is started.

Retaining algorithm. Our progressive retaining algorithm is formalised in Alg. 1. We illustrate the algorithm with a window size of $|W| = 10$ (line 6). The algorithm starts with the empty set $S_0 = \emptyset$. As long as no more than k elements e_1, \dots, e_t have arrived, all of them are kept, i.e., $S_t = S_{t-1} \cup \{e_t\}$, for $t \leq k$. For each new item e_t , where $t > k$, we update the utility value of $f(S_{t-1})$ and compute the semantic importance $p(e_t)$, the decay in interestingness $g(e_t)$ and the social features $h(e_t)$. Then, we check whether swapping this item and an item in S_{t-1} will increase the utility value. If so, the one that maximizes the utility is selected for swapping.

Theorem 1. *Alg. 1 does a single pass over data stream, uses $\mathcal{O}(k)$ memory, and has $\mathcal{O}(k)$ update time per item.*

The proof of Theorem 1 can be found in our technical report [Nguyen et al., 2017]. Correctness of Alg. 1 is established as follows: First, the decay in interestingness is a monotonic decreasing function. Thus, an optimal selection remains optimal after the utility has been updated. Also, even if a new item does not increase utility in terms of entropy and social features, the algorithm still swaps it with an old item to preserve the freshness of the retained set of items. Second, adding a new item increases the entropy of the topic distribution. Thus, while the algorithm favours new items, it still preserves topics by ensuring an even distribution of the selected items across all topics (old and new). Moreover, an optimal selection remains optimal as the entropy increases for all old items.

Incremental utility computation. The above complexity result assumes that the computation of utility (line 14), when swapping data items, is done in constant time. This indeed holds true, since utility can be computed incrementally, i.e., $f(S_i)$ is derived from $f(S_{i-1})$ in constant time as follows:

$$f(S_{i-1} \setminus \{e\} \cup \{e'\}) = f(S_{i-1}) - \sum_{v \in V} n(v, e) p(v, e) \log \frac{1}{p(v, e)}$$

$$g(e) h(e) + \sum_{v \in V} n(v, e') p(v, e') \log \frac{1}{p(v, e')} g(e') h(e').$$

Here, values $p(\cdot, \cdot)$, $n(\cdot, \cdot)$, $g(\cdot)$, $h(\cdot)$ have been computed already in the previous steps of the algorithm (lines 9 to 13).

5 Empirical Evaluation

Below, we first elaborate on the experimental setup, before we analyse our method’s efficiency and effectiveness.

Algorithm 1: A Progressive Retaining Algorithm

```

input : An infinite sequence  $E$  of data items
output: A selected set  $S_t$  of size  $k$  of data items at any time  $t$ 
1  $S_0 = \emptyset$ ;
2 for  $t = 1$  to  $k$  do  $S_t = S_{t-1} \cup \{e_t\}$ ;
3  $W = \emptyset$ ; ▷ Window
4 for  $t = k + 1$  to  $|E|$  do
5    $W = W \cup \{e_t\}$ ;
6   if  $|W| < 10$  and  $t < n$  then
7      $S_t = S_{t-1}$ ;
8   else
9     // Incremental learning of model parameters
10    Compute  $\Theta^{(t)}$  from  $\Theta^{(t-1)}$  and  $W$ ;
11    Update  $p(v, e)$  and  $n(v, e)$  by new parameter  $\Theta^{(t)}$ ,  $\forall v \in V, e \in S_{t-1}$ ;
12    Update  $g(e) \forall e \in S_{t-1}$ ;
13    Update  $f(S_{t-1})$ ;
14    // Prepare computation of utility of new items
15    Compute  $p(v, e')$ ,  $n(v, e')$ ,  $g(e')$ ,  $h(e')$  for all  $e' \in W$  and  $v \in V$ ;
16    // Find swapping pair
17     $e^*, e_b = \arg \max_{e \in S_{t-1}, e' \in W} f(S_{t-1} \setminus \{e\} \cup \{e'\})$ ;
18    if  $f(S_{t-1} \setminus \{e^*\} \cup \{e_b\}) \geq f(S_{t-1})$  then
19       $S_t = S_{t-1} \setminus \{e^*\} \cup \{e_b\}$ ;
20    else  $S_t = S_{t-1}$ ;
21     $W = \emptyset$ ; ▷ Reset for new window

```

5.1 Experimental Setup

Datasets. We extracted datasets using the Twitter Streaming API. Over a year, we considered five different domains (climate change, vaccination, processed food, genetically modified organism, general public) and randomly selected 1 million English tweets per domain. Furthermore, a total of five important social features (e.g., user influence, retweet score, and affective language) had been extracted for each data item using existing frameworks [Zhuang et al., 2016].

Baselines. We compare our approach with several baselines:

- *Traditional summarization*: a state-of-the-art summarization technique [Zhuang et al., 2016] for social data.
- *Greedy*: the simple greedy algorithm (see §3.3) to select the items to retain.
- *Offline learning*: an iterative algorithm to compute model parameters using deterministic variational inference [Jordan et al., 1999; Blei et al., 2006], which requires a full pass of the data in each iteration.
- *Static*: our retaining algorithm tailored to the traditional, static setting of data summarization: *offline learning* to compute the sketch of historical data and the *greedy* algorithm to select the items to retain.

Environment. All results have been obtained on an Intel i7 3.8GHz system (4 cores, 16GB RAM). Following [Hoffman et al., 2013], we vary the forget rate in $(0.5, 1]$, choose a stable window size = 10 and report average values.

5.2 Efficiency

We evaluate the update time of our approach, when new data arrives. To assess the average time per window needed to update the model parameters, we compare our online learning algorithm with its offline version. The latter considers the whole data received so far when computing the parameters upon the arrival of a new window. Figure 3 illustrates the results averaged over all datasets, reporting the average

update time until 100K, 500K, and 1M data items are received. Here, the update time of our progressive approach remains constant and small (< 0.01 s), whereas the baseline yields a high and increasing runtime (up to 10^3 s).

Focusing on how to select the items to retain, we compare the efficiency of our progressive algorithm (*fast*) with the *greedy* algorithm. We realise the online setting as above and vary the size of the set of retained items ($k = 0.2\%$ to 1%). Figure 4 depicts the update time of the algorithms, averaged over all datasets. Our progressive algorithm outperforms the *greedy* one. It also scales better to large data summaries.




Figure 3: Model update





Figure 4: Retaining data

5.3 Effectiveness

We compare the quality of model parameters, in terms of utility, obtained with our online learning algorithm and its offline version. To mitigate the randomness of the Twitter streaming API, we select 100K items \tilde{E} from the original datasets and construct a set of retained items S as the $k = 1\%$ oldest items in \tilde{E} . We stream \tilde{E} and learn model parameters online ($\Theta_{\tilde{E}}$). Offline learning considers all data received so far, the result being $\Theta'_{\tilde{E}}$. We then assess the relative difference in utility of S , computed with either method, i.e., $\frac{|f_{\Theta}(S) - f_{\Theta'}(S)|}{f_{\Theta'}(S)}$.



(a) Updating model parameters (b) Retaining data items

Figure 5: Effectiveness relative to amount of processed data

The difference in utility relative to the amount of processed data is shown in Figure 5a (averaged over 100 runs of different \tilde{E} and the five datasets). Due to the stochastic property of online learning (data is needed to converge in the model parameters), the utility increases initially. Also, the difference between the online and offline learning results is less than 15% in general, underlining the usefulness of our approach.

We further compare the utility of retained items selected by the greedy and our progressive algorithm. Similar to the above setting, we stream all data and use online learning to update the model parameters. Both algorithms start from the set of $k = 1\%$ old items and update it upon the arrival of a new window. We then measure the relative difference of the obtained sets of retained items at each step. The results in Figure 5b (averaged over 100 runs and the five datasets) show that the utility difference increases with the arrival of

data. This is because the progressive algorithm accumulates some loss of information. However, the difference is small ($\leq 15\%$) and converges with more data. As such, data quality is not compromised too much.

Finally, we compare the overall utility ratio (utility of output over utility of whole data) of our retaining algorithm (*dynamic*) and two baselines: the *static* version of our algorithm and traditional summarization (*sum*). The utility ratios obtained after processing all data of the five datasets are shown in Table 1, for different sizes of the set of retained items ($k = 0.1\%$ and $k = 1\%$). While the *static* approach outperforms traditional summarization (*sum*), we need to acknowledge that both, *static* and *sum*, are inapplicable for streaming data. However, the results of the *dynamic* technique are relatively close to those of the *static* approach, highlighting its usefulness for online processing. Also, its decrease in utility for a smaller number of retained items (k) is less drastic compared to traditional summarization.

Table 1: Overall utility ratio

k		climate	vacc	food	gmo	public
1%	static	0.84	0.83	0.84	0.81	0.86
	dynamic	0.70	0.74	0.73	0.74	0.76
	sum	0.72	0.70	0.69	0.71	0.71
.1%	static	0.74	0.76	0.73	0.72	0.75
	dynamic	0.69	0.67	0.68	0.65	0.68
	sum	0.51	0.53	0.48	0.51	0.52

6 Related Work

Social data analysis is often based on topic modelling, feature extraction [Nguyen *et al.*, 2017], and temporal-aware information processing [Chang *et al.*, 2016]. Methods for topic modelling, e.g., Latent Dirichlet Allocation, hierarchical Dirichlet processes [Griffiths and Tenenbaum, 2004], or word modelling [Zhai and Boyd-Graber, 2013], are not applicable for a streaming setting, since they require multiple passes over the data. Streaming versions of these techniques [Canini *et al.*, 2009], in turn, ignore the dynamics of social data. Our model follows a non-parametric approach, where the number of topics and vocabulary words is learned from the data rather than specified in advance. Moreover, our model incorporates social features [Ernst *et al.*, 2014] when assessing data utility. Also, methods to query a streams of social data [Kedzie *et al.*, 2016] are not applicable for data summarization, since the query is not known in advance.

Traditional data summarization works on offline data [Nguyen *et al.*, 2015; Hung *et al.*, 2013] and, even if temporal aspects are considered [Alsaedi *et al.*, 2016], on the whole data. Existing streaming algorithms for data summarization [Badanidiyuru *et al.*, 2014] also rely on access to the complete data, as they sample the data for an estimation of the utility. A relaxed version of the retaining problem has been addressed in [Epasto *et al.*, 2016], which finds subsets of items with maximal utility using a sliding window. Yet, different from our problem formulation, this method targets solely the recent data bounded by a fixed window size, discarding all old items. Whereas, our approach retains old items as long as they are valuable. Also, unlike [Epasto *et al.*, 2016], our approach summarizes the entire history of a data stream. Finally, their algorithm assumes a priori knowledge

of an upper bound of utility. While this assumption may be reasonable for some types of data streams, it is unrealistic for dynamic data produced by social platforms.

7 Conclusions

This paper proposed a technique to retain a representative set of items from a stream of social data. That is, we acknowledge the online nature of data produced by social platforms, which prevents us from storing the complete data stream. This led the retaining problem with minimal regret, where a protocol decides which data to retain, such that the loss of utility is minimized. To address this problem, we proposed a light-weight, adaptive sketch of historical data and a progressive algorithms for the selection of data items. Experiments on large-scale real-world data showed that our approach is efficient (five orders of magnitude faster than the baseline) and effective (less than 15% reduction in the utility ratio).

References

- [Alsaedi *et al.*, 2016] Nasser Alsaedi, Peter Burnap, and Omer Farooq Rana. Automatic summarization of real world events using twitter. In *ICWSM*, pages 511–514, 2016.
- [Badanidiyuru *et al.*, 2014] Ashwinkumar Badanidiyuru, Baharan Mirzasoleiman, Amin Karbasi, and Andreas Krause. Streaming submodular maximization: Massive data summarization on the fly. In *KDD*, pages 671–680, 2014.
- [Blei *et al.*, 2003] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *JMLR*, pages 993–1022, 2003.
- [Blei *et al.*, 2006] David M Blei, Michael I Jordan, et al. Variational inference for dirichlet process mixtures. *Bayesian analysis*, pages 121–143, 2006.
- [Blei *et al.*, 2010] David M Blei, Thomas L Griffiths, and Michael I Jordan. The nested chinese restaurant process and bayesian non-parametric inference of topic hierarchies. *JACM*, pages 7:1–7:30, 2010.
- [Canini *et al.*, 2009] Kevin Robert Canini, Lei Shi, and Thomas L Griffiths. Online inference of topics with latent dirichlet allocation. In *AISStats*, pages 65–72, 2009.
- [Cha *et al.*, 2013] Youngchul Cha, Bin Bi, Chu-Cheng Hsieh, and Junghoo Cho. Incorporating popularity in topic models for social network analysis. In *SIGIR*, pages 223–232, 2013.
- [Chakrabarti and Punera, 2011] Deepayan Chakrabarti and Kunal Punera. Event summarization using tweets. *ICWSM*, pages 66–73, 2011.
- [Chang *et al.*, 2016] Yi Chang, Jiliang Tang, Dawei Yin, Makoto Yamada, and Yan Liu. Timeline summarization from social media with life cycle models. In *IJCAI*, pages 3698–3704, 2016.
- [Epasto *et al.*, 2016] Alessandro Epasto, Silvio Lattanzi, Sergei Vassilvitskii, and Morteza Zadimoghaddam. Submodular optimization over sliding windows. *arXiv preprint arXiv:1610.09984*, 2016.
- [Ernst *et al.*, 2014] Patrick Ernst, Cynthia Meng, Amy Siu, and Gerhard Weikum. Knowlife: a knowledge graph for health and life sciences. In *ICDE*, pages 1254–1257, 2014.
- [Griffiths and Tenenbaum, 2004] DMBTL Griffiths and MIJJB Tenenbaum. Hierarchical topic models and the nested chinese restaurant process. *NIPS*, pages 17–24, 2004.
- [Herrera-Damas, 2014] Susana Herrera-Damas. Recurring topics in the social media policies of mainstream media. *AJMS*, pages 155–173, 2014.
- [Hoffman *et al.*, 2010] Matthew Hoffman, Francis R Bach, and David M Blei. Online learning for latent dirichlet allocation. In *NIPS*, pages 856–864, 2010.
- [Hoffman *et al.*, 2013] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *JMLR*, pages 1303–1347, 2013.
- [Hofmann, 1999] Thomas Hofmann. Probabilistic latent semantic indexing. In *SIGIR*, pages 50–57. ACM, 1999.
- [Hollander *et al.*, 2013] Myles Hollander, Douglas A Wolfe, and Eric Chicken. *Nonparametric statistical methods*. John Wiley & Sons, 2013.
- [Hung *et al.*, 2013] Nguyen Quoc Viet Hung, Hoyoung Jeung, and Karl Aberer. An evaluation of model-based approaches to sensor data compression. *TKDE*, pages 2434–2447, 2013.
- [Hung *et al.*, 2015] Nguyen Quoc Viet Hung, Duong Chi Thang, Matthias Weidlich, and Karl Aberer. Minimizing efforts in validating crowd answers. In *SIGMOD*, pages 999–1014, 2015.
- [Jordan *et al.*, 1999] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, pages 183–233, 1999.
- [Kedzie *et al.*, 2016] Chris Kedzie, Fernando Diaz, and Kathleen McKeown. Real-time web scale event summarization using sequential decision making. In *IJCAI*, pages 3754–3760, 2016.
- [Mathioudakis and Koudas, 2010] Michael Mathioudakis and Nick Koudas. Twittermonitor: trend detection over the twitter stream. In *SIGMOD*, pages 1155–1158, 2010.
- [Morstatter *et al.*, 2013] Fred Morstatter, Jurgen Pfeffer, Huan Liu, and Kathleen M Carley. Is the sample good enough? comparing data from twitter’s streaming api with twitter’s firehose. In *ICWSM*, pages 400–408, 2013.
- [Mukherjee *et al.*, 2014] Subhabrata Mukherjee, Gerhard Weikum, and Cristian Danescu-Niculescu-Mizil. People on drugs: credibility of user statements in health communities. In *KDD*, pages 65–74, 2014.
- [Nemhauser and Wolsey, 1981] George L Nemhauser and Laurence A Wolsey. Maximizing submodular set functions: formulations and analysis of algorithms. *NHMS*, pages 279–301, 1981.
- [Nguyen *et al.*, 2015] Thanh Tam Nguyen, Quoc Viet Hung Nguyen, Matthias Weidlich, and Karl Aberer. Result selection and summarization for web table search. In *ICDE*, pages 231–242, 2015.
- [Nguyen *et al.*, 2017] Quoc Viet Hung Nguyen, Chi Thang Duong, Thanh Tam Nguyen, Matthias Weidlich, Karl Aberer, Hongzhi Yin, and Xiaofang Zhou. Argument discovery via crowdsourcing. *JVLDB*, pages 1–25, 2017.
- [Nguyen *et al.*, 2017] Thanh Tam Nguyen et al. Retaining data from streams of social platforms with minimal regret. *EPFL-CONF-227978*, 2017.
- [Zhai and Boyd-Graber, 2013] Ke Zhai and Jordan L Boyd-Graber. Online latent dirichlet allocation with infinite vocabulary. *ICML*, pages 561–569, 2013.
- [Zhuang *et al.*, 2016] Hao Zhuang, Rameez Rahman, Xia Hu, Tian Guo, Pan Hui, and Karl Aberer. Data summarization with social contexts. In *CIKM*, pages 397–406, 2016.