

Gibbs Sampling Strategies for Semantic Perception of Streaming Video Data

Yogesh Girdhar¹ and Gregory Dudek²

Abstract—Topic modeling of streaming sensor data can be used for high level perception of the environment by a mobile robot. In this paper we compare various Gibbs sampling strategies for topic modeling of streaming spatiotemporal data, such as video captured by a mobile robot. Compared to previous work on online topic modeling, such as o-LDA and incremental LDA, we show that the proposed technique results in lower online and final perplexity, given the realtime constraints.

I. INTRODUCTION

Making decisions based on the environmental context of a robot's locations requires that we first model the context of the robot observations, which in turn might correspond to various semantic or conceptually higher level entities that compose the world. If we are given an observation model of these entities that compose the world then it is easy to describe a given scene in terms of these entities using this model; likewise, if we are given a labeling of the world in terms of these entities, then it is easy to compute the observation model for each individual entity. The challenge comes from doing these two tasks together, unsupervised, and with no prior information. ROST [1], a realtime online spatiotemporal topic modeling framework attempt to solve this problem of assigning high level labels to low level streaming observations.

Topic modeling techniques were originally developed for unsupervised semantic modeling of text documents [2] [3]. These algorithms automatically discover the main themes (topics) that underly these documents, which can then be used to compare these documents based on their semantic content.

Topic modeling of observation data captured by a mobile robot faces additional challenges compared to topic modeling of a collection of text documents, or images that are mutually independent.

- Robot observations are generally dependent on its location in space and time, and hence the corresponding semantic descriptor must take into account the location of the observed visual words during the refinement, and use it to compute topic priors that are sensitive to changes in time and the location of the robot.
- The topic model must be updated online and in realtime, since the observations are generally made continuously at regular intervals. When computing topic labels for a




Fig. 1. Spatiotemporal Topics: As a robot observes the world, we would like its observations to be expressed as a mixture of topics with perceptual meaning. We model the topic distribution of all possible overlapping spatiotemporal regions or neighborhoods in the environment, and place a Dirichlet prior on their topic distribution. The topic distribution of the current observation can then be inferred given the topic labels for the neighborhoods in the view. Modeling neighborhoods allows us to use the context in which the current observation is being made to learn its topic labels. To guarantee realtime performance, we only refine a constant number of neighborhoods in each time step, giving higher priority to recently observed neighborhoods.

new observation, we must also update topic labels for previous observations in the light on new incoming data.

ROST[1] extends previous work on text and image topic modeling to make it suitable for processing streaming sensor data such as video and audio observed by a robot, and presents approximations for posterior inferencing that work in realtime. Topics in this case model the latent causes that produce these observations. ROST has been used for building semantic maps [4] and for modeling curiosity in a mobile robot, for the purpose of information theoretic exploration [5]. ROST uses Gibbs sampling to continuously refine the topic labels for the observed data. In this paper we present various variants of Gibbs sampling that can be used to keep the topic labels converged under realtime constraints.

II. PREVIOUS WORK

A. Topic Modeling of Spatiotemporal Data

Given images of scenes with multiple objects, topic modeling has been used to discover objects in these images in an unsupervised manner. Bosch et al. [6] used PLSA and a SIFT based [7] visual vocabulary to model the content of

¹Yogesh Girdhar is at Deep Submergence Laboratory, Woods Hole Oceanographic Institution, Woods Hole, MA 02542, USA yogi@whoi.edu

²Gregory Dudek is at Center for Intelligent Machines, McGill University, Montreal, QC H3A0E9, Canada dudek@cim.mcgill.ca

images, and used a nearest neighbor classifier to classify the images.

Fei-Fei et al. [8] have demonstrated the use of LDA to provide an intermediate representation of images, which was then used to learn an image classifier over multiple categories.

Instead of modeling the entire image as a document, Spatial LDA (SLDA) [9] models a subset of words, close to each other in an image as a document, resulting in a better encoding of the spatial structure. The assignment of words to documents is not done *a priori*, but is instead modeled as an additional hidden variable in the generative process.

Geometric LDA (gLDA) [10] models the LDA topics using words that are augmented with spatial position. Each topic in gLDA can be visualized as a pin-board where the visual words are pinned at their relatively correct positions. A document is assumed to be generated by first sampling a distribution over topics, and then for each word, sampling a topic label from this distribution, along with the transformation from the latent spatial model to the document (image). These transformations are all assumed to be affine, to model the change in viewpoints.

LDA has been extended to learn a hierarchical representation of image content. Sivic et al. [11] used hierarchical LDA (hLDA) [12] for automatic generation of meaningful object hierarchies. Like LDA, hLDA also models documents as a mixture of topics; however, instead of the flat topics used in LDA, topics in hLDA correspond to a path in a tree. These topics become more specialized as they travel farther down from the root of the tree.

III. SPATIOTEMPORAL TOPIC MODEL

An observation word is a discrete observation made by a robot. Given the observation words and their location, we would like to compute the posterior distribution of topics at this location. Let w be the observed word at location x . We assume the following probabilistic model for the observation words:

- 1) word distribution for each topic k :

$$\phi_k \sim \text{Dirichlet}(\beta),$$

- 2) topic distribution for words at location x :

$$\theta_x \sim \text{Dirichlet}(\alpha + H(x)),$$

- 3) topic label for w :

$$z \sim \text{Discrete}(\theta_x),$$

- 4) word label:

$$w \sim \text{Discrete}(\phi_z),$$

where $y \sim Y$ implies that random variable y is sampled from distribution Y , z is the topic label for the word observation w , and $H(x)$ is the distribution of topics in the neighborhood of location x . Each topic is modeled by distribution ϕ_k over V possible word in the observation vocabulary.

$$\phi_k(v) = \mathbf{P}(w = v | z = k) = \alpha n_k^v + \beta, \quad (1)$$




Fig. 2. Each cell shown corresponds to a spatiotemporal bucket containing all the observation from that region. We refine the topic label for a word w_i in an observation by taking into account the spatiotemporal context G_i of the observation.

where n_k^v is the number of times we have observed word v taking topic label k , and β is the Dirichlet prior hyperparameter. Topic model $\Phi = \{\phi_k\}$ is a $K \times V$ matrix that encodes the global topic description information shared by all locations.

The main difference between this generative process and the generative process of words in a text document as proposed by LDA [2], [3] is in step 2. The context of words in LDA is modeled by the topic distribution of the document, which is independent of other documents in the corpora. We relax this assumption and instead propose the context of an observation word to be defined by the topic distribution of its spatiotemporal neighborhood. This is achieved via the use of a kernel. The posterior topic distribution at location x is thus defined as:

$$\theta_x(k) = \mathbf{P}(z = k | x) \propto \left(\sum_y K(x - y) n_y^k \right) + \alpha, \quad (2)$$

where $K(\cdot)$ is the kernel, α is the Dirichlet prior hyperparameter and, n_y^k is the number of times we observed topic k at location y .

IV. APPROXIMATING NEIGHBORHOODS USING CELLS

The generative process defined above models the clustering behavior of observations from a natural scene well, but is difficult to implement because it requires keeping track of the topic distribution at every location in the world. This is computationally infeasible for any large dataset. For the special case when the kernel is a uniform distribution over a finite region, we can assume a cell decomposition of the world, and approximate the topic distribution around a location by summing over topic distribution of cells in and around the location.

Let the world be decomposed into C cells, in which each cell $c \in C$ is connected to its neighboring cells $G(c) \subseteq C$. Let $c(x)$ be the cell that contains points x . In this paper we only experiment with a grid decomposition of the world in which each cell is connected to its six nearest neighbors, 4 spatial and 2 temporal. However, the general ideas presented here are applicable to any other topological decomposition of spacetime.

```

Initialize  $\forall i, z_i \sim \text{Uniform}(\{1, \dots, K\})$ 
while true do
  foreach cell  $c \in C$  do
    foreach word  $w_i \in c$  do
       $z_i \sim \mathbf{P}(z_i = k | w_i = v, x_i)$ 
      Update  $\Theta, \Phi$  given the new  $z_i$  by updating
       $n_k^v$  and  $n_G^k$ 
    end
  end
end

```

Algorithm 1: Batch Gibbs sampling

The topic distribution around x can then be approximated using cells as:

$$\theta_x(k) \propto \left(\sum_{c' \in G(c(x))} n_{c'}^k \right) + \alpha \quad (3)$$

Due to this approximation, the following properties emerge:

- 1) $\theta_x = \theta_y$ if $c(x) = c(y)$, i.e., all the points in a cell share the same neighborhood topic distribution.
- 2) The topic distribution of the neighborhood is computed by summing over the topic distribution of the neighboring cells rather than individual points.

We take advantage of these properties while doing inference in realtime.

V. REALTIME INFERENCE USING GIBBS SAMPLING

Given a word observation w_i , its location x_i , and its neighborhood $G_i = G(c(x_i))$, we use a Gibbs sampler to assign a new topic label to the word, by sampling from the posterior topic distribution:

$$\mathbf{P}(z_i = k | w_i = v, x_i) \propto \frac{n_{k,-i}^v + \beta}{\sum_{v=1}^V (n_{k,-i}^v + \beta)} \cdot \frac{n_{G_i,-i}^k + \alpha}{\sum_{k=1}^K (n_{G_i,-i}^k + \alpha)}, \quad (4)$$

where $n_{k,-i}^w$ counts the number of words of type w in topic k , excluding the current word w_i , $n_{G_i,-i}^k$ is the number of words with topic label k in neighborhood G_i , excluding the current word w_i , and α, β are the Dirichlet hyper-parameters. Note that for a neighborhood size of 0, the above Gibbs sampler is equivalent to the LDA Gibbs sampler proposed by Griffiths et al.[3], where each cell corresponds to a document. Algorithm 1 shows a simple iterative technique to compute the topic labels for the observed words in batch mode.

In the context of robotics we are interested in the online refinement of observation data. After each new observation, we only have a constant amount of time to do topic label refinement. Hence, any online refinement algorithm that has computational complexity which increases with new data, is not useful. Moreover, if we are to use the topic labels of an incoming observation for making realtime decisions, then

```

while true do
  Add new observed words to their corresponding
  cells.
   $T \leftarrow 0$  (current time)
  Initialize  $\forall i \in M_T, z_i \sim \text{Uniform}(\{1, \dots, K\})$ 
  while no new observation do
     $t \sim \mathbf{P}(t|T)$ 
    foreach cell  $c \in M_t$  do
      foreach word  $w_i \in c$  do
         $z_i \sim \mathbf{P}(z_i = k | w_i = v, x_i)$ 
        Update  $\Theta, \Phi$  given the new  $z_i$  by
        updating  $n_k^v$  and  $n_G^k$ 
      end
    end
  end
   $T \leftarrow T + 1$ 
end

```

Algorithm 2: Realtime Gibbs sampler

it is essential that the topic labels for the last observation converge before the next observation arrives.

Since the total amount of data collected grows linearly with time, we must use a refinement strategy that efficiently handles global (previously observed) data and local (recently observed) data.

Our general strategy is described by Algorithm 2. At each time step we add the new observations to the model, and then randomly pick observation times $t \sim \mathbf{P}(t|T)$, where T is the current time, for which we resample the topic labels and update the topic model.

We discuss the choice of $\mathbf{P}(t|T)$ in the following sections.

A. Now Gibbs Sampling

The simplest way of processing streaming observation data to ensure that the topic labels from the last observation have converged is to only refine topics from the last observation till the next observation has arrived.

$$\mathbf{P}(t|T) = \begin{cases} 1, & \text{if } t = T \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

We call this the Now Gibbs sampler. This is analogous to o-LDA approach by Banerjee and Basu [13].

If R is our computation budget, defined as the expected number of observation time-steps our system can refine between the arrival times of two consecutive observations, and $r(t)$ be the number of times observations in M_t have been refined after time T , then this approach gives each observation R amount of resources.

$$\mathbb{E}\{r(t)\} = R \quad (6)$$

Although this sounds fair, the problem is that no information from the future is used to improve the understanding of the past data.

B. Uniform Gibbs Sampling

A conceptually opposite strategy is to uniform randomly pick an observation from all the observations thus far, and refine the topic labels for all the words in this observation.

$$\mathbf{P}(t|T) = 1/T \quad (7)$$

This is analogous to the incremental Gibbs sampler for LDA proposed by Canini et al.[14].

Let M_t be the set of cell containing observations at time t , R be the number of observations our system can refine between two observations, and $r(t)$ be the number of times observations in M_t have been refined after time T . The expected value of $r(t)$ is then:

$$\mathbb{E}\{r(t)\} = R \left(\frac{1}{t} + \frac{1}{t+1} + \dots + \frac{1}{T} \right) \quad (8)$$

$$\approx R(\log T - \log t). \quad (9)$$

We see that older observations are sampled disproportionately higher than newer observations, and topic labels of new observations might take a long time to converge. In fact, if τR is the expected number of iterations it takes for topic labels of an observation to converge, where $\tau < 1$ is a constant, then all observations after time $t' = 1/\tau$ would never be able to converge in the time before the next observation arrives. This is a big problem for a real-time system, where we need the topic labels of the last observations to actuate the robot.

C. Age Proportional Gibbs Sampling

A seemingly good in-between approach might be to bias the random sampling of observations to be refined in favor of picking recent observations, with probability proportional to its timestamp.

$$\mathbf{P}(t|T) = \frac{t}{\sum_{i=1}^T i} \quad (10)$$

Then, the expected number of times this observation is refined is given by:

$$\mathbb{E}\{r(t)\} = R \left(\frac{t}{\sum_{i=1}^t i} + \frac{t}{\sum_{i=1}^{t+1} i} + \dots + \frac{t}{\sum_{i=1}^T i} \right) \quad (11)$$

$$\approx 2R \frac{(T-t)}{T}. \quad (12)$$

When a new observation is made, the expected number of refinements it will gets before the next observation arrives is $Rt/\sum t \approx 2R/t$, which implies that if t' is the time after which it will not have sufficient number of refinements, then:

$$\frac{2R}{t'} = \tau R \quad (13)$$

$$\implies t' = \frac{2}{\tau} \quad (14)$$

Hence, we see that this strategy, although better than uniform random sampling (for which we computed $t' = 1/\tau$), is still not useful for long term operating of the robot.

D. Exponential Gibbs Sampling

Using a geometric distribution we can define the probability of refinement of timestep t , at current time T

$$\mathbf{P}(t|T) = q(1-q)^{T-t}, \quad (15)$$

where $0 < q < 1$ is a parameter. Using this distribution for picking refinement samples ensures that on average qR number of refinements are spent on refining the most recent observations, and the remaining $(q-1)R$ refinement iterations are spent on refining other recent observations. In the limit $T \rightarrow \infty$, observations in each time-step are refined $\mathbb{E}\{r(t)\} = R$ number of times, similar to Now Gibbs Sampler. This approach, however, allows new information to influence some of the recent past observations, resulting in lower global perplexity of the learned model.

E. Mixed Gibbs Sampling

We expect both Now and Exponential Gibbs samplers to be good at ensuring the topic labels for the last observation converges quickly (to a locally optimal solution), before the next observation arrives, whereas Uniform and Age-proportional Gibbs samplers are better at finding globally optimal results.

One way to balance both these performance goals is to combine these global and a local strategies. We consider four such approaches in this paper:

Uniform+Now:

$$\mathbf{P}(t|T) = \begin{cases} \eta, & \text{if } t = T \\ (1-\eta)/(T-1), & \text{otherwise} \end{cases} \quad (16)$$

AgeProportional+Now:

$$\mathbf{P}(t|T) = \begin{cases} \eta, & \text{if } t = T \\ (1-\eta) \frac{t}{\sum_{i=1}^{T-1} i}, & \text{otherwise} \end{cases} \quad (17)$$

Uniform+Exp:

$$\mathbf{P}(t|T) = \eta q(1-q)^{T-t} + (1-\eta)/T \quad (18)$$

AgeProportional+Exp:

$$\mathbf{P}(t|T) = \eta q(1-q)^{T-t} + (1-\eta) \frac{t}{\sum_{i=1}^T i} \quad (19)$$

Here $0 \leq \eta \leq 1$ is the mixing proportion between the local and the global strategies.

VI. EXPERIMENTS

1) *Dataset*: We evaluated the performance on ROST in analyzing videos using three different datasets with millions of visual words. We used a mixed vocabulary to describe each frame, with 5000 ORB words, 256 intensity words (pixel intensity), and 180 hue words (pixel hue), for a total vocabulary size of 5436. Although it is difficult to substantiate the optimality of the vocabulary, our experiments have suggested that once the vocabulary size is sufficiently large, there is limited sensitivity to its precise value [15].

Some key statistics for these datasets is shown in Table I.

Name	size	T	N(words)	$\frac{N(\text{words})}{T}$	V
2objects	720x480	1158	1741135	1503	5436
aerial	640x480	3600	8190231	2275	5436
underwater	1024x638	2569	4809869	1872	5436

TABLE I
VIDEO DATASETS FOR EVALUATING ROST

The *2objects* dataset show a simple scenario in which two different objects appear on a textured (wood) background randomly, first individually and finally together.

The *aerial* dataset was collected using Unicorn UAV over a coastal region. The UAV performs a zig-zag coverage pattern over buildings, forested areas and ocean.

The *underwater* dataset was collected using Aqua as it swims over a coral reef. The dataset contains a variety of complex underwater terrain such as different coral species, rocks, sand, and divers.

The video files corresponding to these datasets, and some examples of ROST in action are available at ¹.

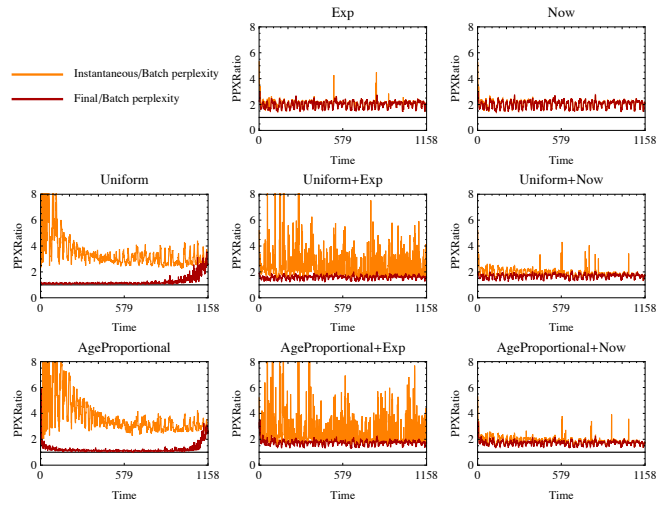
To focus on analyzing the effects of spatiotemporal neighborhoods, and various Gibbs samplers, we fixed all other parameters of the system. We used cells of size 64x64 pixels with temporal width of 1 time step, Dirichlet parameters $\alpha = 0.1, \beta = 0.5$, number of topics $K = 16$.

A. Realtime Gibbs Samplers


To evaluate the proposed realtime Gibbs samplers on real data, we performed the following experiment. For each video dataset, and for each Gibbs sampler, we computed the topic labels and perplexity online, with 10 random restarts. We then compared the mean perplexity of words, one time step after their arrival (instantaneous), and after all observations have been made (final), with the perplexity of topic labels computed in batch. For a fair comparison, we used the same refinement time per time step (T_R) for both batch and online cases. The resulting perplexity plots are shown in Figures 3, 4, and 5. The mean perplexity scores for the entire datasets are shown in Tables III (instantaneous perplexity), and II (final perplexity). Note that instantaneous perplexity is computed on a new image, given the model learnt online from all previous data. Hence this perplexity score serves the same purpose as computing perplexity on held out data when evaluating topic modeling on batch data.

From our experiments we find that although Uniform and Age Proportional Gibbs samplers perform well when it comes to final perplexity of the dataset, they however perform poorly when measuring instantaneous perplexity. Low instantaneous perplexity, which is measured one time step after an observation is made, is essential for use of topic modeling in robotic applications. We would like to make decisions based on current observations, and hence low instantaneous perplexity is crucial. We find that the mixed Gibbs samplers such as Uniform+Now perform consistently well. Note that all experiments with the mixed Gibbs samplers were performed with a fixed mixing ratio $\eta = 0.5$,

¹<http://cim.mcgill.ca/mrl/girdhar/>



(a) Refinement time $T_R = 40$ ms.



(b) Refinement time $T_R = 160$ ms.

Fig. 3. 2Objects dataset – ratio of instantaneous and final perplexity to batch perplexity, for each time step

giving equal weight to local and global refinement. We are confident that better tuning of this variable will result in even better performance of ROST.

VII. CONCLUSION

Topic modeling techniques such as ROST, model the latent context of the streaming spatiotemporal observation, such as image and other sensor data collected by a robot. In this paper we compared the performance of several Gibbs samplers for realtime spatiotemporal topic modeling, including those proposed by o-LDA and incremental LDA.

We measured how well the topic labels converge, globally for the entire data, and for individually for an observation, one time step after its observation time. The latter measurement criterion is useful in evaluating the performance of the proposed technique in the context of robotics, where we need to make instantaneous decisions. We showed that the proposed mixed Gibbs samplers such as Uniform+Now

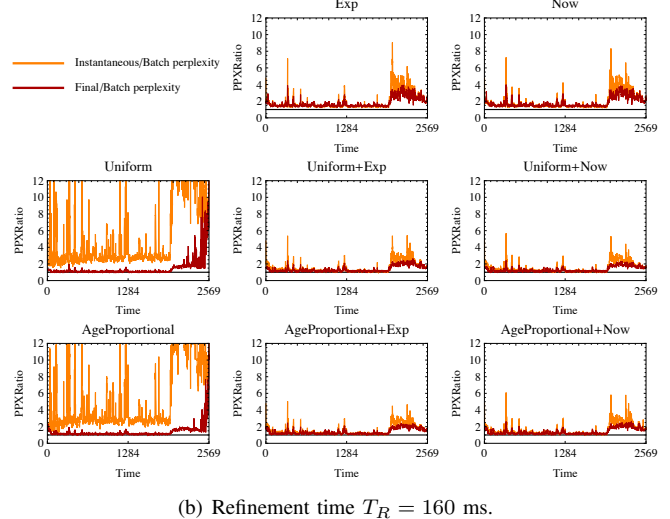
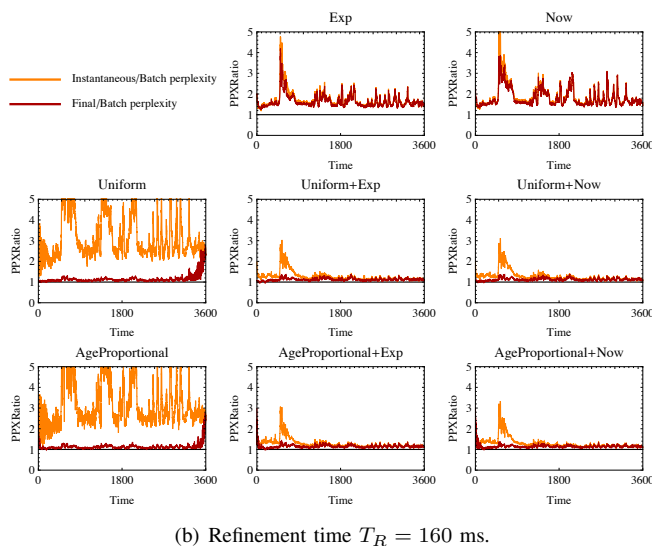
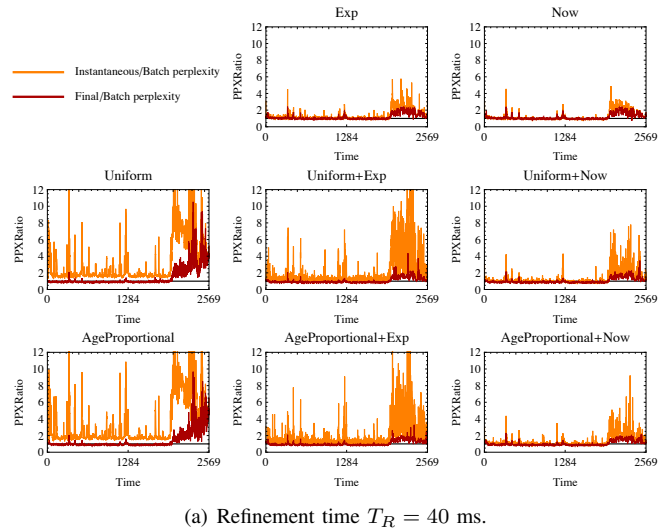
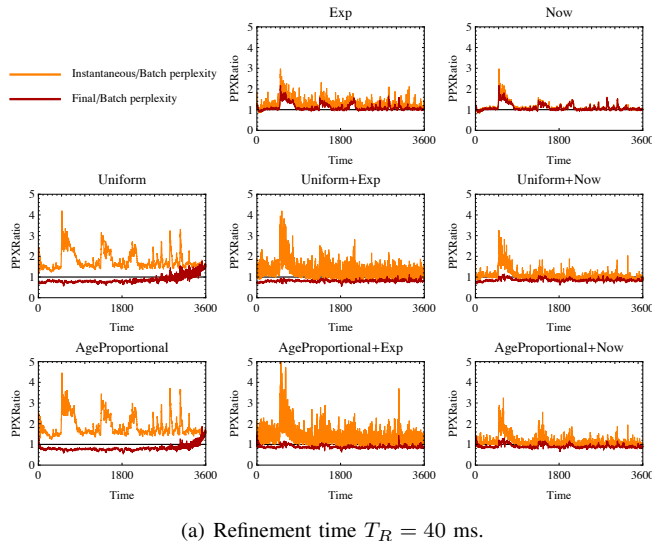


Fig. 4. Aerial dataset – ratio of instantaneous and final perplexity to batch perplexity, for each time step

perform consistently better than other samplers, which just focus on recent observation, or which refine all observation with equal probability.

ACKNOWLEDGMENT

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) through the NSERC Canadian Field Robotics Network (NCFRN). Yogesh Girdhar is currently supported by the Postdoctoral Scholar Program at the Woods Hole Oceanographic Institution, with funding provided by the Devonshire Foundation and the J. Seward Johnson Fund. Authors would like to thank Julian Straub at MIT for helpful discussion.

REFERENCES

- [1] Y. Girdhar, P. Giguere, and G. Dudek, “Autonomous adaptive exploration using realtime online spatiotemporal topic modeling,” *The International Journal of Robotics Research*, vol. 33, no. 4, pp. 645–657, Nov. 2013. [Online]. Available: <http://ijr.sagepub.com/cgi/doi/10.1177/0278364913507325>
- [2] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent dirichlet allocation,” *The Journal of Machine Learning Research*, vol. 3, pp. 993–1022, 2003.

Fig. 5. Underwater dataset – ratio of instantaneous and final perplexity to batch perplexity, for each time step

- [3] T. L. Griffiths and M. Steyvers, “Finding scientific topics,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, no. Suppl 1, pp. 5228–5235, 2004.
- [4] Y. Girdhar, D. Whitney, and G. Dudek, “Curiosity Based Exploration for Learning Terrain Models,” in *IEEE International Conference on Robotics and Automation*, Oct. 2014, p. 7. [Online]. Available: <http://arxiv.org/abs/1310.6767>
- [5] Y. Girdhar and G. Dudek, “Exploring Underwater Environments with Curiosity,” in *2014 Canadian Conference on Computer and Robot Vision*. Montreal: IEEE, May 2014, pp. 104–110. [Online]. Available: <http://www.cim.mcgill.ca/~yogesh/publications/crv2014.pdf> <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6816831>
- [6] A. Bosch, A. Zisserman, and X. Muñoz, “Scene Classification Via pLSA,” in *Computer Vision – ECCV 2006*, ser. Lecture Notes in Computer Science, A. Leonardis, H. Bischof, and A. Pinz, Eds. Springer Berlin / Heidelberg, 2006, vol. 3954, pp. 517–530.
- [7] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International Journal of Computer Vision*, vol. 60, no. 2, 2004.
- [8] L. Fei-Fei and P. Perona, “A Bayesian Hierarchical Model for Learning Natural Scene Categories,” in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, vol. 2. IEEE, Jun. 2005, pp. 524–531. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1467486>
- [9] X. Wang and E. Grimson, “Spatial Latent Dirichlet Allocation,” in *Advances in Neural Information Processing Systems*, vol. 20,

Dataset \rightarrow	2objects		aerial		underwater	
Alg. \downarrow , $T_R \rightarrow$	40	160	40	160	40	160
Uniform	1.215	1.155	0.881	1.141	0.993	1.133
AgeP.	1.207	1.269	0.832	1.120	0.969	1.147
Exp.	1.799	2.272	1.065	1.548	1.005	1.488
Now	1.791	2.270	1.056	1.662	0.998	1.491
Uni+Now	1.555	1.420	0.842	1.109	0.885	1.152
AgeP+Now	1.599	1.575	0.893	1.134	0.894	1.193
Uni+Exp	1.480	1.398	0.808	1.1082	0.887	1.163
AgeP+Exp	1.609	1.576	0.874	1.137	0.906	1.185

TABLE II
MEAN FINAL PERPLEXITY

Dataset \rightarrow	2objects		aerial		underwater	
Alg. \downarrow , $T_R \rightarrow$	40	160	40	160	40	160
Uniform	3.500	4.826	1.665	2.798	1.907	3.140
AgeP.	3.714	5.091	1.715	2.855	1.997	3.135
Exp.	1.859	2.310	1.170	1.598	1.073	1.490
Now	1.829	2.312	1.089	1.711	1.006	1.502
Uni+Now	1.756	1.770	0.994	1.226	0.988	1.242
AgeP+Now	1.739	1.883	1.028	1.247	0.996	1.270
Uni+Exp	2.267	1.784	1.279	1.217	1.357	1.249
AgeP+Exp	2.320	1.873	1.339	1.238	1.375	1.256

TABLE III
MEAN INSTANTANEOUS PERPLEXITY PERPLEXITY

- 2007, pp. 1577–1584. [Online]. Available: <http://people.csail.mit.edu/xgwang/papers/STLDA.pdf>
- [10] J. Philbin, J. Sivic, and A. Zisserman, “Geometric LDA: A Generative Model for Particular Object Discovery,” in *In Proceedings of the British Machine Vision Conference*, 2008.
- [11] J. Sivic, B. C. Russell, A. Zisserman, W. T. Freeman, and A. A. Efros, “Unsupervised discovery of visual object class hierarchies,” in *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Jun. 2008, pp. 1–8. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4587622>
- [12] D. M. Blei, T. L. Griffiths, M. I. Jordan, and J. B. Tenenbaum, “Hierarchical topic models and the nested Chinese restaurant process,” *Advances in neural information processing systems (NIPS)*, vol. 16, pp. 106–114, 2004.
- [13] A. Banerjee and S. Basu, “Topic Models over Text Streams: A Study of Batch and Online Unsupervised Learning,” in *SIAM International Conference on Data Mining*, 2007, p. 6. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/1.9781611972771.40>
- [14] K. R. Canini, L. Shi, and T. L. Griffiths, “Online Inference of Topics with Latent Dirichlet Allocation,” *Proceedings of the International Conference on Artificial Intelligence and Statistics*, vol. 5, no. 1999, pp. 65–72, 2009.
- [15] Y. Girdhar and G. Dudek, “Online Visual Vocabularies,” in *CRV '11: Proceedings of the 2011 Canadian Conference on Computer and Robot Vision*. IEEE Computer Society, May 2011, pp. 191–196. [Online]. Available: <http://cim.mcgill.ca/~yogesh/publications/crv2011.pdf>