Approximate Distributed Spatiotemporal Topic Models for Multi-Robot Terrain Characterization

Kevin Doherty^{1,2}, Genevieve Flaspohler^{1,2}, Nicholas Roy¹, and Yogesh Girdhar²

Abstract—Unsupervised learning techniques, such as Bayesian topic models, are capable of discovering latent structure directly from raw data. These unsupervised models can endow robots with the ability to learn from their observations without human supervision, and then use the learned models for tasks such as autonomous exploration, adaptive sampling, or surveillance. This paper extends single-robot topic models to the domain of multiple robots. The main difficulty of this extension lies in achieving and maintaining global consensus among the unsupervised models learned locally by each robot. This is especially challenging for multi-robot teams operating in communication-constrained environments, such as marine robots.

We present a novel approach for multi-robot distributed learning in which each robot maintains a local topic model to categorize its observations and model parameters are shared to achieve global consensus. We apply a combinatorial optimization procedure that combines local robot topic distributions into a globally consistent model based on topic similarity, which we find mitigates topic drift when compared to a baseline approach that matches topics naïvely. We evaluate our methods experimentally by demonstrating multi-robot underwater terrain characterization using simulated missions on real seabed imagery. Our proposed method achieves similar model quality under bandwidth-constraints to that achieved by models that continuously communicate, despite requiring less than one percent of the data transmission needed for continuous communication.

I. Introduction

Unsupervised machine learning techniques can enable adaptive robotic systems that are robust to unexpected changes in their environment. Furthermore, the operation of robots in environments like the deep benthic sea, where they may encounter species and terrains not previously documented, demands the flexibility to novel observations afforded by unsupervised learning. A team of underwater robots capable of unsupervised scene categorization would enable more efficient ocean surveying and exploration. However, because unsupervised learning methods do not have a priori defined labels, combining the local models discovered by individual robots into a globally cohesive model is a hard combinatorial optimization problem. Achieving a globally cohesive scene model is critical for both postmission analysis and so that multi-robot teams can make better global planning and exploration decisions in real-time. This paper presents a novel approach to multi-robot learning

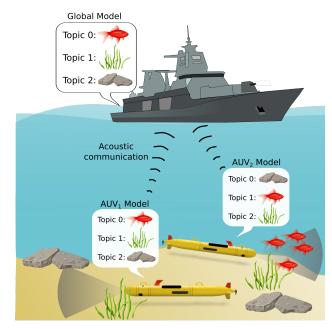


Fig. 1: Illustration of the multi-robot terrain characterization problem. Each individual AUV builds its own local topic model. Model parameters from all vehicles are transmitted top-side via acoustic modem, where the models must be combined into a consistent global model. However, there is no guarantee that the individual robots will agree on which visual category corresponds to each topic ID. This correspondence issue must be addressed to achieve global consensus.

that obtains a globally consistent scene categorization under communication constraints, endowing a team of exploratory robots with the ability to reach consensus in their semantic description of the world.

In this work, we consider the problem of underwater terrain characterization. Formally, given a sequence of images, we would like to predict for each image the distribution over a set of latent categories that generated the data. Previous approaches to this problem for single robots [1] make use of spatiotemporal variants of topic models like latent Dirichlet allocation (LDA) [2] and the nonparametric hierarchical Dirichlet process (HDP) [3]. The extension of these methods to the multi-robot setting presents several challenges. We aim to have every robot build a model that discovers thematic structure within an image stream that coincides well with human semantics, but also for each robot's model to be *consistent* with the other robots in a multi-robot exploration

¹Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology, 77 Massachusetts Ave, MA 02139, USA {kjd, geflaspo, nickroy}@csail.mit.edu

²Department of Applied Ocean Physics and Engineering, Woods Hole Oceanographic Institution, 86 Water St, Woods Hole, MA 02543, USA yogi@whoi.edu

team. We define consistent topic models to be those with similar topic-word distributions.

Semantically meaningful, consistent topic models can be achieved by allowing the robots to share raw sensor observations directly with one another continuously; however, in many practical scenarios this approach is infeasible due to communication constraints. Additionally, it does not allow the individual robots to build their own models for local control. We opt instead to build a local topic model on each robot and provide an approach allowing the correspondences between topic models to be identified and fused without sharing the raw sensor observations. An illustration of our problem scenario is depicted in Figure 1, where we show seafloor exploration with multiple autonomous underwater vehicles (AUVs) communicating acoustically with a central top-side modem. Topic models learned by each robot are transmitted top-side, where the models are combined to build a consistent global model that is then sent back to all vehicles.

The primary contributions of this work are as follows:

- We present a novel algorithm for multi-robot real-time online spatiotemporal topic modeling.
- We experimentally validate our approach by simulating a multi-robot underwater mission where we seek to categorize terrains and fauna observed on the seafloor without human supervision or a priori defined labels.
- We examine the performance of the proposed approach under different communication constraints, and show that our approach produces consistent topic models even under substantial communication delays, as well as when one robot observes topics that go completely unobserved by another.
- We additionally show that the proposed model achieves similar model quality under bandwidth-constraints to that achieved by models that continuously communicate, despite requiring <1% of the data transmission needed for continuous communication.

To our knowledge, this work is the first application of distributed topic models to streaming data from a multi-robot mission. Consequently, this work provides new avenues for intelligent, coordinated multi-robot exploration.

The remainder of the paper is organized as follows: in Section II we discuss related work in machine learning and robotics. In Section III we briefly describe the latent Dirichlet allocation topic model and real-time online spatiotemporal topic modeling (ROST) [1], and then present our approach, approximate distributed ROST (AD-ROST). In Section IV we present both qualitative and quantitative results from topic modeling experiments on a simulated multi-robot mission using data collected at the Hannibal Bank Seamount off the coast of Panama. Finally, in Section V we discuss several compelling directions for future work in this area.

II. RELATED WORK

Unsupervised machine learning models have been of interest in the robotics community for several years. In the area

of visual topic models, both parametric and nonparametric models have been used for analysis of seabed imagery [4], [5], underwater exploration [6], [1], and navigation [7]. Our approach is based on the *realtime online spatiotemporal topic modeling* (ROST) framework developed by Girdhar et al. in [1]. More recently, the seabed imagery analysis and anomaly detection work in [5] was adapted to use features learned by a deep convolutional auto-encoder [8]. The parametric visual topic models developed in the aforementioned works could all be directly extended to the multi-robot setting using our proposed methods.

Within the machine learning community, there are several approaches for distributing topic models. Newman et al. [9] provide methods for distributing LDA and HDP on multiple processors in an offline setting. Our baseline approach is an adaptation of their approximate distributed LDA algorithm to the online setting, while inheriting the spatial and temporal modeling capabilities of ROST. The bipartite matching approach we take toward addressing the component correspondence issue is grounded in methods for multi-processor learning, and has been applied in distributed learning literature to LDA and HDP in centralized, sampling-based [9] and variational inference [10], [11], as well as in the decentralized case [12].

Our problem differs from the traditional distributed learning setting fundamentally in that our data is partitioned naturally by virtue of being collected on independent robots. In contrast, in a traditional learning setting, we would have access to the full dataset and partition the data into minibatches across multiple processors, for example by sampling randomly from the dataset. Our data also contains significant spatial and temporal correlations, since they consist of consecutive observations from a sensor stream taken as a robot navigates through the environment. Beyond that, we are subject to the practical constraints imposed by acoustic communication on our ability to send updates to a global model. The consequence of this is not only that topic drift is more likely, but also that some robots may make observations of visual categories that have not been observed by others. Our work focuses on addressing these topic correspondence issues as they apply to learning on multi-robot systems.

III. METHODS

In the following sections, we review the latent Dirichlet allocation (LDA) topic model and its extension to spatiotemporal domains, highlighting the permutation symmetry problem as it pertains to these two models. We then provide two methods for multi-robot topic modeling. The first is a naïve adaptation of parallel LDA for multiple processors [9] to the multi-robot setting, in which we merge topics directly by their ID, without considering topic consistency. In the second approach, we explicitly account for the permutation symmetry problem during inference in order to improve global consistency between models.

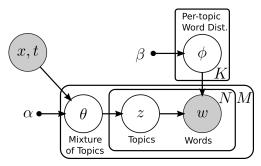


Fig. 2: Graphical model representation of the ROST framework. ROST incorporates both spatial and temporal relationships into the traditional LDA formulation.

A. Latent Dirichlet Allocation

The latent Dirichlet allocation (LDA) topic model [2] is used to extract latent thematic structure from a corpus of documents. In the LDA model, a document d consists of a set of discrete words $\{w_1, w_2, \ldots, w_N\}$ from a vocabulary of size V. In LDA, we infer the distribution over a set of latent "topics" for each document and the likelihood of all words given these topics by computing a topic assignment z_i for each word w_i in a document. With this factorization into the distributions $p(z_i|d)$ and $p(w_i|z_i)$, the likelihood of an observed word w_i given its document d is computed by marginalizing over the topics as follows:

$$p(w_i \mid d) = \sum_{k=1}^{K} p(w_i \mid z_i = k) p(z_i = k \mid d).$$
 (1)

We will often abbreviate the topic mixing proportions for a document p(z|d) as θ_d and the per-topic word distribution p(w|z=k) as ϕ_{kw} . Dirichlet priors are placed over both distributions, i.e. $\theta_d \sim \mathrm{Dir}(\alpha)$ and $\phi_{kw} \sim \mathrm{Dir}(\beta)$ where α and β are hyperparameters.

B. Real-time Online Spatiotemporal Topic Models

The traditional formulation of LDA does not consider correlations between observed words in a document. Spatial LDA [13] incorporates spatial correlations between visual words within an image. The real-time online spatiotemporal topic model (ROST) [1] further extends this method to image streams by additionally incorporating correlations between visual word observations within the same temporal neighborhood. In the ROST framework, depicted graphically in Figure 2, we consider a document, here denoted G_i , as the set of all observed words in a spatiotemporal neighborhood. Neighborhoods are determined by cells of fixed-size in space and time. Similar to LDA, ROST places Dirichlet priors over ϕ_{kw} and the topic mixing proportions, which we denote θ_{G_i} to make explicit the spatiotemporal context. The likelihood of a word w_i can then be written

$$p(w_i \mid x, t) = \sum_{k=1}^{K} p(w_i \mid z_i = k) p(z_i = k \mid x, t)$$
 (2)

Algorithm 1 AD-ROST-ID

```
\label{eq:continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous_continuous
```

where x and t denote the location of w_i in an image and the time of its observation, respectively.

Exact inference in the ROST model and LDA more generally is intractable. Instead, we perform approximate inference using Gibbs sampling [14]. Gibbs sampling is performed by sampling a topic z_i for every word $w_i = v$ conditioned on the set of all words \mathbf{w} and the current set of topics assigned to all *other* words, denoted \mathbf{z}_{-i} . This sampling distribution is computed as follows:

$$p(z_i = k \mid \mathbf{z}_{-i}, \mathbf{w}) \propto \left[\frac{n_{k,-i}^{(v)} + \beta}{n_{k,-i}^{(\cdot)} + V\beta} \right] \left[\frac{n_k^{(G_i)} + \alpha}{n_{-i}^{(G_i)} + K\alpha} \right], \quad (3)$$

where $n_{k,-i}^{(v)}$ is the count of assignments of topic k to every other observation of word v, $n_{k,-i}^{(\cdot)}$ is the count of all current assignments of topic k, $n_k^{(G_i)}$ is the count of all assignments of topic k in spatiotemporal neighborhood G_i , and $n_{-i}^{(G_i)}$ is the corresponding total count of topic assignments to all other words in G_i . Gibbs sampling is performed continuously as new images are observed by the robot.

The maximum-likelihood estimates for the per-topic word distribution and topic mixing proportions can be calculated using the counts recorded during Gibbs sampling as follows:

$$\hat{\phi}_{kw} = \frac{n_k^{(w)} + \beta}{n_k^{(\cdot)} + V\beta} \tag{4}$$

$$\hat{\theta}_{G_ik} = \frac{n_k^{(G_i)} + \alpha}{n^{(G_i)} + K\alpha}.\tag{5}$$

C. Approximate Distributed ROST with Topic IDs

The marginalization over topics in the generative models of LDA and ROST (Equations 1 and 2 respectively) causes the likelihood of a word to be invariant to the permutation of topics in these models. For a single topic model, this permutation symmetry does not warrant consideration. When we seek to distribute the model, however, this permutation symmetry becomes vital: two models can be equivalent in terms of the likelihood, but be inconsistent in terms of topic

correspondence. This problem of resolving correspondence between latent variables in unsupervised models is a general problem in distributed unsupervised learning.

The baseline approach that we consider is an adaptation of the parallel LDA formulation of Newman et al. [9] to the ROST model. We make the naïve assumption that the topic IDs assigned by each robot directly correspond, i.e topic 1 on robot 1, topic 1 on robot 2, through topic 1 on robot R are topics corresponding to the same visual category, where R is the number of robots. Under this assumption, it is straightforward to derive a global update rule based on the per-topic word counts for each agent. Letting N_{kw} be the global $K \times V$ matrix of per-topic word counts obtained from Gibbs sampling across all agents, and letting N_{kwr} be the per-topic word count matrix on robot r, we combine the local counts by summation, taking care to avoid duplicates of the global count matrix:

$$N_{kw} \leftarrow N_{kw} + \sum_{r=1}^{R} (N_{kwr} - N_{kw}),$$
 (6)

where R is the number of robots. The complete procedure for distributing ROST with ID-based topic matching, AD-ROST-ID, is described in Algorithm 1. Each robot r receives the global count matrix N_{kw} , then begins processing the next Timages, where T is a fixed number of observations between global synchronizations of all robot models. The procedure ExtractWords produces a set of visual words consisting of extracted feature descriptors. RefineTopics then performs Gibbs sampling over the data given the new observations. After we have added T new observations, we synchronize models by communicating local counts to a central node, which combines the counts according to the update in Equation 6. On the next cycle, each robot retrieves the new global count matrix, replacing its local model with the global model. We repeat this process until every agent ceases to record new observations. If one agent finishes before the others, its final updates are made and the remaining agents proceed as usual, relying on the central node to handle synchronization appropriately.

It may seem counter-intuitive to assume that models combined by ID in this manner would converge to a globally consistent set of topics describing the same visual categories. Consistency between models is achieved by the repeated mixing of model parameters into an averaged global model and the resetting of the local models. When mixing happens frequently enough, ID-based matching performs Gibbs sampling from an approximation of the true posterior for the combined dataset [9]. As a consequence, successfully arriving at the appropriate topic correspondences in the case of ID-based topic matching rests on the frequency of global model updates. In a traditional distributed learning setting, it is not unreasonable to expect that we can synchronize processors after only a few iterations of Gibbs sampling. However, on a team of robots, communication delays may be substantial, and models may converge to different sets of topics locally between synchronizations.

Algorithm 2 AD-ROST-SIM

```
repeat
     // Local model updates
    for each robot r in parallel do
          // Receive global counts N_{kw}
          N_{kwr} \leftarrow N_{kw}
          for t from t_{curr} to t_{curr} + T do
               \mathbf{w}, \mathbf{x}, \mathbf{t} \leftarrow \text{ExtractWords}(I_t)
              \mathbf{z}, N_{kwr} \leftarrow \text{RefineTopics}(\mathbf{z}, \mathbf{w}, \mathbf{x}, \mathbf{t})
          end for
     end for
     // Global model updates
     Synchronize // Retrieve each N_{kwr}
     for each robot r do
          C_{r1} \leftarrow \text{ComputeCost}(\hat{\phi}_{kwr}, \hat{\phi}_{kw1})
          \pi_r^* \leftarrow \operatorname{Hungarian}(C_{r1})
    N_{kw} \leftarrow N_{kw} + \sum_{r} \pi_r^* (N_{kwr} - N_{kw})
until no new observations
```

D. Approximate Distributed ROST with Topic Similarity

In order to mitigate the potential topic drift issues caused by naïve ID-based matching in a direct adaptation of parallel LDA [9], we now propose an alternative approach. We would like to obtain local topic models on each robot that are *consistent* in the sense that their topics are appropriately aligned and the same topic ID on multiple robots corresponds to the same semantic visual category. Accounting for this consistency should result in a visual topic model that better corresponds with human semantics. If the probability distribution over the vocabulary is similar for two topics, we say that the topics are similar. We adopt the heuristic in [11], measuring the similarity between two topic models in terms of the sum of squared distances between these distributions¹

$$D(\hat{\phi}_{kw}, \hat{\phi}'_{kw}) = \sum_{k=1}^{K} \|\hat{\phi}_{kw} - \hat{\phi}'_{kw}\|^2.$$
 (7)

Given this measure of model similarity, we pose model consistency as an explicit optimization objective during the inference procedure. Each time the global model is updated, we permute each robot's topic-word distribution to optimize model consistency. There are K! possible permutations of topics each time we perform model synchronization. If we synchronize S times during a mission, there are $K!^S$ total possible sequences of permutations over the entire mission. Consequently, it is intractable to maintain a full distribution over all possible permutations of topics for every model synchronization. Instead, we optimize the permutations greedily at every model update.

Let $\hat{\phi}_{kwr}$ denote the empirical topic-word distribution on robot r and let $\pi(\hat{\phi}_{kwr})$ denote a permutation of topics in

¹Although we use the sum of squared Euclidean distance heuristic, other measures of topic similarity for related models have been presented, such as symmetric Kullback-Leibler divergence [9].

 $\hat{\phi}_{kwr}$. We aim to solve the following optimization:

$$\pi_r^* = \operatorname*{arg\,min}_{\pi(\hat{\phi}_{kwr})} f(\pi(\hat{\phi}_{kwr})) \tag{8}$$

$$\pi_r^* = \arg\min_{\pi(\hat{\phi}_{kwr})} f(\pi(\hat{\phi}_{kwr}))$$
(8)
$$f(\pi(\hat{\phi}_{kwr})) = \sum_{k=1}^K \|\pi(\hat{\phi}_{kwr}) - \hat{\phi}_{kw1}\|^2,$$
(9)

where π_r^* is the optimal permutation of the topic-word distribution on robot r. In our case, the optimal permutation minimizes the sum of squared Euclidean distances between the topic-word distributions on robot r and those on the first robot. Since we optimize permutations with respect to the first robot, we take π_1^* to be the identity.

We minimize this objective for R robots using the Hungarian algorithm [15], which provides the optimal assignment in $\mathcal{O}(K^3)$, giving an overall complexity of $\mathcal{O}(RK^3)$ for the optimization procedure at each model synchronization step.

After finding the optimal permutation for each robot, π_r^* , we update the global topic-word counts according to the following rule:

$$N_{kw} \leftarrow N_{kw} + \sum_{r=1}^{R} \pi_r^* (N_{kwr} - N_{kw}),$$
 (10)

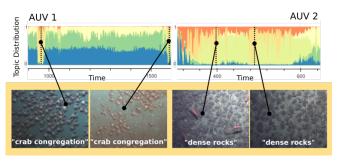
where, as before, N_{kw} and N_{kwr} are the global and local counts of topic k for word w, respectively. After subtracting the set of previous global topic-word counts, we permute the model updates from each robot using its respective optimal permutation from Equation 8. The new update is exactly the update from Equation 6 when π_r^* is identity for all r, which is the only scenario where ID-based matching is optimal with respect to the objective posed in Equation 9.

This approach, which we call AD-ROST-SIM, is summarized in pseudocode in Algorithm 2. The individual vehicles perform the same operations as in the case of AD-ROST-ID. The primary distinction is that the central node, after receiving the model parameters of the robots, then computes for each robot the pairwise distances between word distributions for each topic on that robot and that of the first robot (i.e. the procedure ComputeCost). Optimization of costs produces the optimal permutation of topics on robot r to match the topics on the first robot. Differences in counts are permuted, then incorporated into the global model by summation.

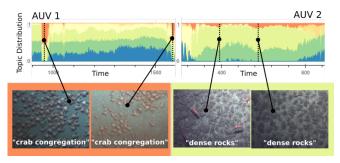
IV. EXPERIMENTAL RESULTS

A. Experimental Setup

We compare the proposed distributed learning approaches using real video data collected by a single AUV equipped with a downward-facing camera at the Hannibal Bank Seamount, off the coast of Panama [16]. To simulate a multirobot mission, we treat data from multiple deployments of a single robot as though it were retrieved simultaneously by multiple vehicles. Each simulated "robot" communicates with a central node to update a globally synchronized model in order to produce a consistent scene categorization model shared between all agents in the system. To aid in evaluation, each mission has been annotated using a set of thirteen



(a) Topic distributions predicted by AUV 1 and AUV 2 using the AD-ROST-ID method. AD-ROST-ID results in different visual categories being assigned to the same topic.



(b) Topic distributions predicted by AUV 1 and AUV 2 using the AD-ROST-SIM method. Matching by topic similarity helps to ensure that distinct visual categories are appropriately assigned to different topics.

Fig. 3: Comparison of topic distributions inferred by AD-ROST-ID (a) and AD-ROST-SIM (b) with global updates every 100 observations. Neglecting to consider permutation symmetry in this limited communication setting results in different visual categories being placed in the same topic. Accounting for permutation symmetry with AD-ROST-SIM causes allocation of separate topics for these categories.

possible ground-truth terrain labels, including "crab congregation," "water column," and "sparse rocks."

In the first simulated mission, the first deployment, "AUV 1," contains 2,296 image frames, each taken four seconds apart over approximately 2.5 hours. In this mission, the robot descends through the water column until it arrives near the seafloor, which primarily consists of smooth, sandy terrains. Notably, however, these visually nondescript sandy terrains are punctuated by several dense crab swarms [16]. Finally, the vehicle ascends through the water column. The second deployment, "AUV 2," contains 1,117 image frames taken four seconds apart, with a total duration of approximately 1.25 hours. In the second mission, the robot observes several distinct seafloor terrains. It initially descends through the water column, then observes a rocky terrain, followed by porous, sandy seafloor, before ascending back through the water column. Since the second AUV mission ends before the first AUV mission, the vehicle stops sending updates to the global model and the first vehicle continues its mission alone until it receives no new observations.

In the second simulated mission, the first deployment,

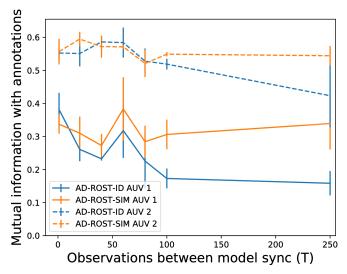


Fig. 4: AD-ROST-SIM achieves higher mutual information than AD-ROST-ID as we increase the period between model synchronizations. Despite requiring a fraction of the data transmission, the mutual information of AD-ROST-SIM with T=250 is comparable to that of either method with T=1.

"AUV 1," contains 1,244 frames over about 1.25 hours in which the robot descends through the water column, travels over sandy and rocky seafloor terrains, and then ascends back through the water column. The second deployment, "AUV 2," contains 2,733 image frames with a total duration of approximately 3 hours. In this deployment, after the vehicle's descent, it simply observes a variety of nondescript sandy terrains over the course of the mission, before ascending back through the water column. In this mission, the first AUV deployment finishes before the second, so the first vehicle stops sending updates to the global model while the second vehicle continues.

Since our focus is on the relative performance of each method under a variety of communication constraints, we fixed the hyperparameters $K=7,\,\alpha=0.1,\,$ and $\beta=10$ in all experiments. In practice, hyperparameters may be optimized in cross-validation by comparing mutual information between maximum-likelihood topics and a set of ground-truth labels, if annotations are available [8]. We evaluated models with $T=1,\,20,\,40,\,60,\,80,\,100,\,$ and 250 observations between global model updates. Because the Gibbs sampling process may produce different estimates when run multiple times on the same data, we trained each model three times and show average performance and standard deviation for all quantitative results.

The topic models used in our experiment make use of visual words consisting of ORB features [17]. Our approach is also compatible with other methods of feature extraction, including features learned by a deep neural network [8]. The visual vocabulary we use has size V=6500 and was built using k-means clustering of ORB features extracted from unrelated video. At run time, ORB features extracted from

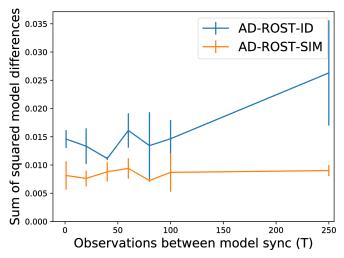


Fig. 5: Sum of squared differences in topic-word distributions on each model evaluated on the *first* mission. As we increase the number of observations between model synchronization steps, the topic distributions begin to diverge in the case of AD-ROST-ID. AD-ROST-SIM is comparatively robust to decreases in communication rate.

images count as instances of their nearest visual vocabulary word, quantified by Euclidean distance.

B. Evaluation Metrics

We tested the correspondence between the discovered topics and human annotation by measuring the mutual information between the most likely topic labels for each mission and ground-truth annotations. Mutual information is the amount of entropy reduced in a random variable Y after the observation of a random variable X, and is formally defined as

$$I(X;Y) = \sum_{x} \sum_{y} p(x,y) \log \frac{p(x,y)}{p(x)p(y)}.$$
 (11)

Mutual information here measures the degree to which the maximum-likelihood topic labels correlate with annotations. Specifically, normalized mutual information equal to 1 signifies that maximum-likelihood topics are in one-to-one correspondence with annotated labels. A normalized mutual information score of zero indicates the opposite: that the maximum-likelihood topics and annotated labels are statistically independent.

We quantify topic divergence by examining the similarity between the topic-word distributions for the two vehicles at the end of the mission. Specifically, we measure the sum of squared distances between topic-word distributions for each topic, consistent with the heuristic distance $D(\hat{\phi}_{kw}, \hat{\phi}'_{kw})$ posed in Equation 7. A large value of $D(\hat{\phi}_{kw}, \hat{\phi}'_{kw})$ indicates that the pair of models has drifted from one another, resulting in dissimilar topic-word distributions.

C. Results and Discussion

In Figure 3 we show, for both vehicles in the first mission, the topic distributions produced by AD-ROST-ID and AD-

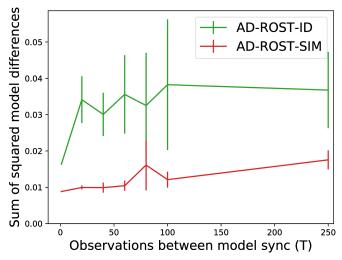
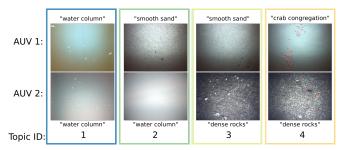


Fig. 6: Sum of squared differences in topic-word distributions on each model evaluated for the *second* mission. As communication becomes less frequent, AD-ROST-ID quickly begins to show topic model divergence that persists to 250 iterations between synchronizations. AD-ROST-SIM performance degrades slightly with 250 iteration gaps in communication, but consistently performs much better than AD-ROST-ID.

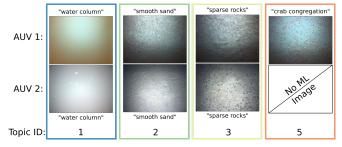
ROST-SIM with synchronizations every 100 observations. The portion of the mission highlighted in Figure 3 demonstrates that when updates to the global model are infrequent (e.g. T=100), merging topics naïvely based on ID (ADROST-ID) causes topic distributions to drift. This problem is most apparent when comparing the maximum-likelihood topics near t=950 on the first AUV and t=400 on the second AUV. The second AUV observes a rocky terrain at t=400, but at t=950, the same maximum-likelihood topic corresponds to a crab colony on AUV 1. On the other hand, merging topics based on similarity (AD-ROST-SIM) appropriately places these distinct visual categories into different topics, despite also synchronizing the models only every 100 observations.

In Figure 4, we show the mutual information between the maximum-likelihood topic predictions and a set of human annotations of the data from the first mission. We observe that as we increase the period between synchronizations, the performance of AD-ROST-ID degrades. With AD-ROST-SIM, we achieve nearly the same model quality when we synchronize every 250 observations as when we synchronize after every observation, despite the fact that we transmit *less than one percent* of the data required to synchronize after every observation.

In Figure 5, we show the sum of squared differences between the per-topic word distributions on each robot at the end of the first mission as we increase the number of observations between global updates. We observe that the sum of squared model differences begins to grow in the case of AD-ROST-ID. This quantifies the phenomenon that we observed in Figure 3, when the same topic was assigned to



(a) Topic exemplars, grouped by topic ID, from the model inferred using AD-ROST-ID with T=100 show several instances where the robots assign the same topic to visually different categories.



(b) Exemplars from the model inferred using AD-ROST-SIM with $T=100\,$ show that the model successfully assigns different topics to different visual categories. Additionally, when a particular category is observed by one AUV but not the other, such as in the case of the "crab congregation," there is no image from AUV 2 where the corresponding topic occurs with maximum likelihood.

Fig. 7: Annotated exemplar images for each topic. Topics that did not occur with maximum-likelihood on any image are denoted "No ML Image." This occurs when a visual category is only observed by one of the two AUVs. Mismatches where the individual robot topic models assign the same topic label to images of visually distinct categories arise with AD-ROST-ID, but not when using AD-ROST-SIM.

both images of a crab colony and the rocky seafloor. Consistent with our previous observations, AD-ROST-SIM is robust to infrequent model communication up to intervals of 250 observations between synchronizations, which corresponds to synchronization roughly every 15 minutes.

For the second mission, we show the sum of squared differences in per-topic word distributions on each of the two vehicles in Figure 6. We similarly observe topic divergence in the case of AD-ROST-ID, which happens much more quickly than in the first mission and persists to 250 iterations between global model synchronizations. Distance in the topic models grows in the case of AD-ROST-SIM as well around 250 iterations, but the resulting model similarity is comparable to AD-ROST-ID with synchronization every iteration.

Another way to examine the topics learned by these models is through "exemplar images" of each topic. To further compare the effects of model drift on the topic correspondences between the two vehicles, we provide high-likelihood exemplars from four of the discovered topics in the first mission in Figure 7. Topics that never occur with maximum-likelihood are not displayed. We show that AD-

ROST-ID incorrectly matches several topics between the two vehicles. Most notably, we observe again the phenomenon where AUV 1's topic representing a congregation of crabs is matched to a rocky terrain image on AUV 2. Using AD-ROST-SIM, we see that the topic exemplified by an image of crabs has no maximum-likelihood match on AUV 2. Since we know that the crab congregation never appears throughout the second AUV's mission, this is appropriate. Furthermore, the remaining topic exemplars from the similarity-based matching procedure are in direct correspondence with one another, in contrast to AD-ROST-ID.

V. FUTURE WORK

The multi-robot topic modeling framework presented here enables many interesting directions for future research. We provide a multi-agent topic modeling framework, enabling distributed variants of unsupervised exploration algorithms [1], [5], or perplexity-based navigation methods, as in [7]. In general, leveraging our globally consistent, unsupervised models for multi-agent robotic exploration is a compelling area for future research.

We have not considered the situation where robots can communicate directly with one another absent a global, centralized node. Application of decentralized variants of topic models, either using variational inference [12] or Gibbs sampling [18], to the multi-robot domain would greatly improve the robustness of these methods to unpredictable communication delays or node failures.

Another avenue for future work is the extension of these methods to the case where the number of topics is not chosen *a priori* and consistently across robots, such as in the case of the hierarchical Dirichlet process (HDP) model [3]. A topic matching procedure similar to Algorithm 2 could be used with HDP models in communication-constrained settings, for example by allocating new global topics using a threshold on distance between topic-word distributions as in [9]. The notion of category discovery is vital for the sort of lifelong learning required of systems operating in novel environments, such as in underwater exploration. Extension of our methods to nonparametric unsupervised models is a priority for further work on this problem.

VI. CONCLUSION

We have presented a novel algorithm for multi-robot realtime online spatiotemporal topic modeling. We applied our algorithm in the context of underwater terrain characterization. By explicitly incorporating an optimization procedure to match topics during the distributed inference process, we achieve more globally consistent models despite the robots' having never shared their raw sensor observations due to communication constraints. Beyond resolving issues due to topic drift, we also found that our method is robust to situations where certain visual categories were observed by one robot but not by the other. Finally, enforcing topic consistency can lead to improved correspondence between maximum-likelihood topic labels and human annotations, as measured by mutual information. The proposed method, AD-ROST-SIM, infers semantically meaningful, consistent topics comparable to those of models requiring constant communication, while transmitting substantially less data.

ACKNOWLEDGMENTS

We thank James Kinsey, Brian Claus, and Trevor Campbell for helpful discussions of ideas presented in this paper. This project was supported in part by National Science Foundation (NSF) NRI Award #1734400, and NOAA Award #NA17OAR0110207.

REFERENCES

- 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, 2014
- [2] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," Journal of Machine Learning Research, vol. 3, pp. 993–1022, 2003.
- [3] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Sharing clusters among related groups: Hierarchical dirichlet processes," in *Advances* in neural information processing systems, 2005, pp. 1385–1392.
- [4] D. M. Steinberg, O. Pizarro, and S. B. Williams, "Hierarchical bayesian models for unsupervised scene understanding," *Computer Vision and Image Understanding*, vol. 131, pp. 128–144, 2015.
- [5] Y. Girdhar, W. Cho, M. Campbell, J. Pineda, E. Clarke, and H. Singh, "Anomaly detection in unstructured environments using bayesian nonparametric scene modeling," in *Robotics and Automation (ICRA)*, 2016 IEEE International Conference on. IEEE, 2016, pp. 2651–2656.
- [6] Y. Girdhar and G. Dudek, "Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring," *Autonomous Robots*, vol. 40, no. 7, pp. 1267–1278, oct 2016. [Online]. Available: http://link.springer.com/10.1007/s10514-015-9500-x http://cim.mcgill.ca/mrl/girdhar/rost/aqua_curiosity.mp4
- [7] R. Paul and P. Newman, "Self help: Seeking out perplexing images for ever improving navigation," in *Robotics and Automation (ICRA)*, 2011 IEEE International Conference on. IEEE, 2011, pp. 445–451.
- [8] G. Flaspohler, N. Roy, and Y. Girdhar, "Feature discovery and visualization of robot mission data using convolutional autoencoders and bayesian nonparametric topic models," in *Intelligent Robots and Systems (IROS)*, 2017 IEEE/RSJ International Conference on. IEEE, 2017, pp. 1–8.
- [9] D. Newman, A. Asuncion, P. Smyth, and M. Welling, "Distributed algorithms for topic models," *Journal of Machine Learning Research*, vol. 10, no. Aug, pp. 1801–1828, 2009.
- [10] T. Campbell, J. Straub, J. W. Fisher III, and J. P. How, "Streaming, distributed variational inference for bayesian nonparametrics," in Advances in Neural Information Processing Systems, 2015, pp. 280–288.
- [11] M. Rabinovich, E. Angelino, and M. I. Jordan, "Variational consensus monte carlo," in *Advances in Neural Information Processing Systems*, 2015, pp. 1207–1215.
- [12] T. Campbell and J. P. How, "Approximate decentralized bayesian inference," in *Proceedings of the Thirtieth Conference on Uncertainty* in *Artificial Intelligence*. AUAI Press, 2014, pp. 102–111.
- [13] X. Wang and E. Grimson, "Spatial latent dirichlet allocation," in Advances in Neural Information Processing Systems, 2008, pp. 1577– 1584.
- [14] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proceedings of the National academy of Sciences*, vol. 101, no. suppl 1, pp. 5228–5235, 2004.
- [15] H. W. Kuhn, "The hungarian method for the assignment problem," Naval Research Logistics (NRL), vol. 2, no. 1-2, pp. 83–97, 1955.
- [16] J. Pineda, W. Cho, V. Starczak, A. F. Govindarajan, H. M. Guzman, Y. Girdhar, R. C. Holleman, J. Churchill, H. Singh, and D. K. Ralston, "A crab swarm at an ecological hotspot: patchiness and population density from auv observations at a coastal, tropical seamount," *PeerJ*, vol. 4, p. e1770, 2016.
- [17] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *Computer Vision (ICCV)*, 2011 IEEE international conference on. IEEE, 2011, pp. 2564–2571.
- [18] P. Smyth, M. Welling, and A. U. Asuncion, "Asynchronous distributed learning of topic models," in *Advances in Neural Information Process*ing Systems, 2009, pp. 81–88.