

Autonomous Adaptive Underwater Exploration using Online Topic Modeling

Yogesh Girdhar, Philippe Giguère, and Gregory Dudek

1 Introduction

Exploration of underwater environments, such as coral reefs and ship wrecks, is a difficult and potentially dangerous tasks for humans, which naturally makes the use of an autonomous robotic system very appealing. This paper presents such an autonomous system, and shows its use in a series of experiments to collect image data in an underwater marine environment. We presents novel contributions on three fronts. First, we present an online topic-modeling based technique to describe what is being observed using a low dimensional semantic descriptor. This descriptor attempts to be invariant to observations of different corals belonging to the same species, or observations of similar types rocks observed from different viewpoints. Second, we use the topic descriptor to compute the *surprise* score of the current observation. This is done by maintaining an online summary of observations thus far, and then computing the surprise score as the distance of the current observation to the summary, in the topic space. Finally, we present a novel control strategy for an underwater robot thats allows for intelligent traversal; hovering over surprising observations, and swimming quickly over previously seen corals and rocks.

Exploration, in the context of robotics, has been studied before. Work has been done on autonomous mapping of challenging environments [19], [12], frontier expansion [22], minimizing uncertainty [21], and utility based exploration[8]. Our focus is instead on traversing an environment similar to how a tourist might do so in a new city; stopping and recording any surprising sights, while moving fast when

Yogesh Girdhar and Gregory Dudek
Centre for Intelligent Machines, McGill University
308-3480 University Street, Montreal, QC, Canada, H3A 0E9.
e-mail: {yogesh,dudek}@cim.mcgill.ca

Philippe Giguère
Département d'informatique et génie logiciel, Université Laval
1065, Avenue de la Médecine, Québec, QC, Canada, G1V 0A6.
e-mail: philippe.giguere@ift.ulaval.ca

nothing new is in sight. This is similar to the vacation snapshot problem described in [2].

We used an untethered amphibious robot (Aqua[15]) with an in-house designed autopilot to carry the exploration task. Images were taken with a downward-looking camera, with all computations performed onboard. Its propulsion is based on six flippers that can provide motion in five degrees of freedom. By using a novel combination of gaits, the robot was able to move at various speeds while maintaining its orientation, despite external disturbances. This was necessary in order to complete this exploration task.

2 Approach

The abstract task of controlling the motion of a robot, based on surprise, can be broken down into three sub-problems: computing surprise of a new observation, having a meaningful descriptor for observations over which surprise could be computed, and finally, controlling the robot given the surprise score of the current observation. We describe these sub-problems in the following sections.

2.1 Summaries and Surprises

Summarizing observations made by a robot has recently gained popularity in robotics[13, 6]. Our goal, however, is to compute a summary which assists in evaluating the novelty of a new observation. We do this by maintaining a summary that is representative of all of observations made thus far, and then compute the surprise score as the distance to this summary.

Let $\mathbf{M}^t = \{M_1, \dots, M_t\}$ be the set of all observations till time t . We maintain a subset of k observations as the summary $\mathbf{S} = \{S_1, \dots, S_k\}$, $\mathbf{S} \subseteq \mathbf{M}^t$, such that the maximum distance of an observation to its closest summary sample is minimized. The cost function is thus defined as:

$$\text{Cost}(\mathbf{S}|\mathbf{M}^t) = \max_i \min_j d(M_i, S_j), \quad (1)$$

where d is the distance function, which measures distance as the symmetric KL divergence between the corresponding topic distributions, which we describe in Sec 2.2. Such a summary is sometimes called as an *Extremum Summary* [6], because minimizing the above cost function is essentially minimizing the distance of the worst outlier to the summary. This is different from a more typical a k -medoids clustering based summary, which tries to minimize the mean distance of an observation to the closest summary.

The novelty or surprise of a new observation $\xi(Z_t|\mathbf{S})$ can then be defined as its Hausdorff distance to the summary [5].

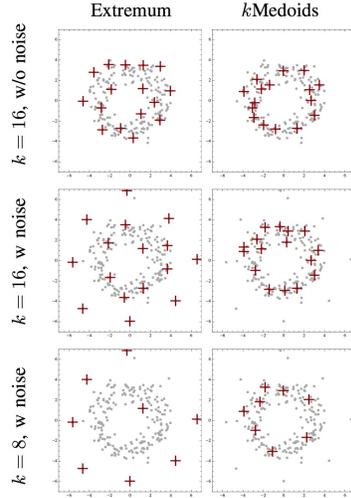


Fig. 1 Extremum vs k -Medoids Summaries. The dataset consists of 200 points generated randomly around a circle in \mathbb{R}^2 . The summaries generated by the two algorithms are shown in the first row. Since there are no outliers in the dataset, the summaries seem similar. In the second row, we add 8 extra samples from a different distribution, which are all outliers in the context of the other points. Adding these outliers highlights the differences between the two strategies. We see that extremum summary favors picking the outliers, whereas the k -medoids summary ignores these outliers completely. In the last row, we reduce the summary size and see the differences exaggerated even more. The extremum summary is almost entirely made up of the outliers, whereas the k -medoids summary is only representative of the mean.

$$\xi(M_t|\mathbf{S}) = \min_j d(M_t, S_j). \quad (2)$$

If the distance function obeys the triangle inequality, which is true in our case, then not only is this problem NP-hard, but Huse and Nemhauser [11] showed that any α -approximation of this problem is also NP-hard for $\alpha < 2$. Gonzalez [7] proved that the simple greedy solution of recursively picking the farthest samples, has an approximation ratio of 2, which is likely the best we can do unless P=NP.

In the online case, Charikar et al. [4] have proposed a simple strategy where after each pick, the picking threshold is doubled. This leads to a summary which is guaranteed to have a cost less than $8 \times$ 'optimal'. However, since the topic assignment of samples in the summary are continuously being refined, we instead set the threshold dynamically to $2 \times$ 'minimum inter-sample distance in \mathbf{S} ', as illustrated in Fig. 2.

To control the summary size, we simply use the greedy offline summarization algorithm on the summary to keep the summary of desired size. In our prior work[6], we have studied the rate of growth of the summary, when threshold is set to the mean distance of a summary sample, to the remaining summary. This is useful in the case when we want the summary size to grow with the data.

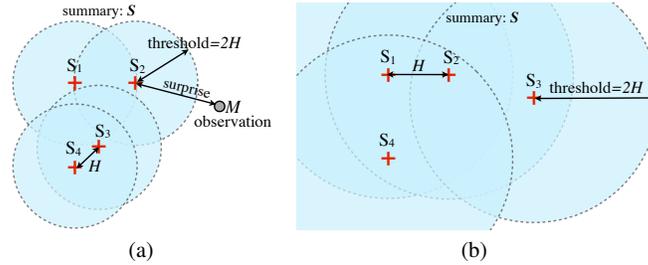


Fig. 2 (a) Given a summary, represented here by the ‘+’ sign, we define the threshold score for updating the summary as twice the smallest inter-sample distance H . When a new observation M arrives, we compute its surprise: the distance to the closest summary sample. If the surprise exceeds the threshold $2H$, then the summary is updated to include the new observation. The updated summary and the threshold are shown in (b).

Figure 1 highlights the characteristic difference in summaries generated by the extremum summary algorithm, and the k -Medoids algorithm. The summaries generated by the two algorithms are shown in the first row. Since there are no outliers in the dataset, the summaries seem similar. In the second row of Figure 1, we add 8 extra samples from a different distribution, which are all outliers in the context of the other points. Adding these outliers highlights the differences between the two strategies. We see that extremum summary favors picking the outliers, whereas the k -medoids summary ignores these outliers completely. In the last row of Figure 1, we reduce the summary size and see the differences exaggerated even more. The extremum summary is almost entirely made up of the outliers, whereas the k -medoids summary is still only representative of the mean.

Although a k -medoids summary might be useful when we want to model the mean properties of an environment, if however we are interested in identifying the range of what was observed, then an extremum summary is more useful, since its objective function ensures that each observations is close to at least one of the summary samples.

2.2 Online Spatiotemporal Topics

To have meaningful summaries, and thus a meaningful surprise score, we must use an image descriptor that is sensitive to thematic changes in the scene, while being immune to low level image changes. We do this via the use of a topic modeling framework, using which we describe an incoming observation with a low dimensional distribution over topics.

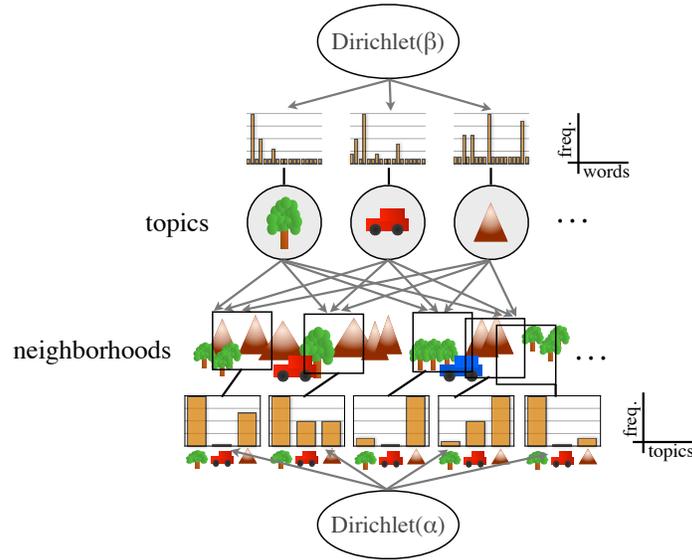


Fig. 3 Spatiotemporal Topics: As a robot observes the world, we would like its observations to be expressed as a mixture of topics with perceptual meaning. We model the topic distribution of all possible overlapping spatiotemporal regions or neighborhoods in the environment, and place a Dirichlet prior 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.

2.2.1 Topics

Topic modeling methods were originally developed for text analysis. Hoffman [10] introduced the idea of probabilistic Latent Semantic Analysis (pLSA) for text documents, which modeled the probability of observing a word w_i in a given document m as:

$$\mathbb{P}(w_i|m) = \sum_{k=1}^K \mathbb{P}(w_i|z_i = k) \mathbb{P}(z_i = k|m). \quad (3)$$

The central idea being the introduction of a latent variable z , which models the underlying topic, or the context responsible for generating the word. The topic variable z takes a value from $1 \dots K$, where K is a small number compared to the size of the vocabulary. The distribution of these topics in a document gives us a low dimensional semantic description of the document. The words in our context corresponds to visual words [17], and documents correspond to observed images. Latent Dirichlet Allocations [1] improve upon pLSA by setting Dirichlet priors on the distribution of words in a topic, and the distribution of topics in a document.

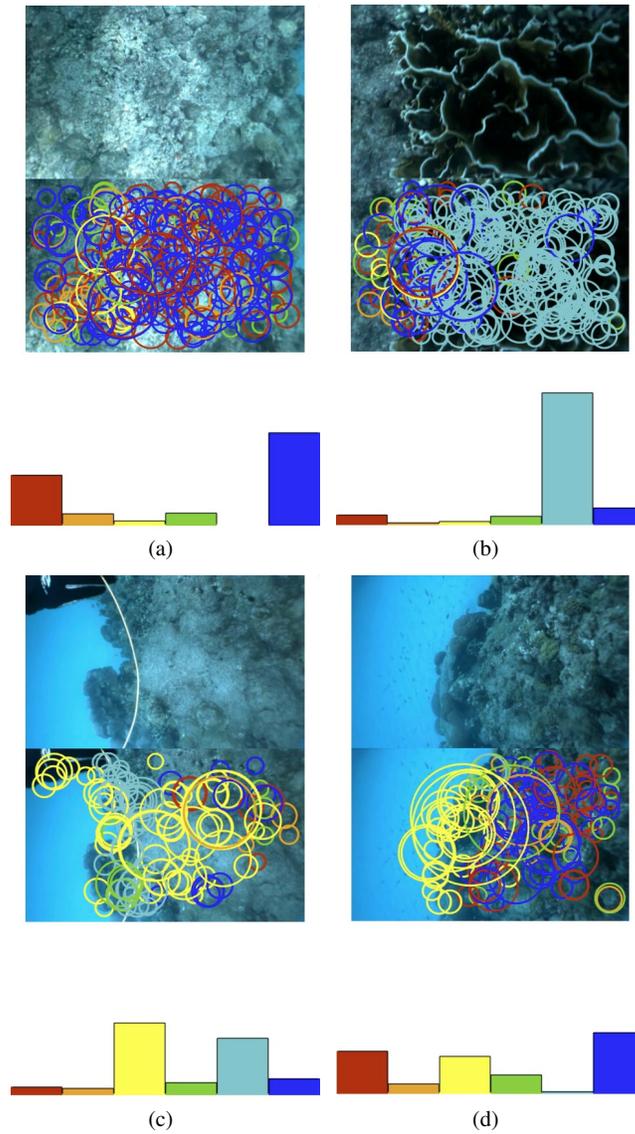


Fig. 4 Example of topics learned on images of the ocean floor taken by the Aqua robot, for a single trajectory. Each visual word is marked by a circle, the size of which corresponds to the size of the visual feature. Histograms depicting the content of each color-coded topic are shown below.

2.2.2 Spatiotemporal Topics

In this work, we use a novel online spatiotemporal LDA for computing topic labels, which not only is able to efficiently converge the topic labels for a new observation,

but also update topic labels for previous observations in the light on new incoming data. Moreover, we take into account the location of the observed visual words during the refinement, using them to compute topic priors which are sensitive to changes in time and location of the robot. Our algorithm’s iteration runs in constant time, and hence is suitable for realtime use on autonomous vehicles with limited computing capabilities. Figure 4 shows examples of topics which were learnt by running the above topic model on an underwater image sequence containing 2000 images.

Each observation is a tuple (w_i, \mathbf{x}_i, t_i) consisting of an observed visual word $w_i = 1 \dots V$ from a fixed vocabulary of size V , associated coordinates \mathbf{x}_i , and timestamp t_i . In the presence of location information, \mathbf{x}_i is the location of the word w_i in world coordinates. In absence of location information, we can simply use the pixel coordinates of the visual words. The neighborhood of an observation at (\mathbf{x}, t) , denoted by $G(\mathbf{x}, t)$, is the set of observations in its spatiotemporal neighborhood. This could either be defined using k nearest neighbors, or using a radius search. Instead of computing topic distributions over documents in a traditional LDA [1], or image windows in Spatial-LDA [20], we compute topic distributions over these spatiotemporal neighborhoods (Figure 3). Modeling topic distribution over neighborhoods allows us to use spatiotemporal context in which an observation is being made, which in turn results in much faster convergence as is shown later in our results.

Given a location and time (\mathbf{x}_i, t_i) , we use the following generative model for the corresponding observed word w_i :

1. word distribution for each topic k : $\phi_k \sim \text{Dirichlet}(\beta)$,
2. neighborhood for an observation at (\mathbf{x}_i, t_i) : $G(\mathbf{x}, t) \sim$ uniformly from all neighborhoods which contain (\mathbf{x}_i, t_i) ,
3. topic distribution the neighborhood $G(\mathbf{x}, t)$: $\theta_{G(\mathbf{x}, t)} \sim \text{Dirichlet}(\alpha)$,
4. topic label for location (\mathbf{x}_i, t_i) : $z_i \sim \text{Discrete}(\theta_{G(\mathbf{x}, t)})$,
5. word observed at location (\mathbf{x}_i, t_i) : $w_i \sim \text{Discrete}(\phi_{z_i})$,

where $x \sim Y$ implies that random variable x is sampled from distribution Y .

2.2.3 Gibbs Sampling

Similar to the Gibbs sampler proposed by Griffiths et al.[9], we can define the posterior topic distribution of a neighborhood $G(\mathbf{x}, t)$:

$$\mathbb{P}(z_i = k | \mathbf{z}_{-i}, w_i = w, \mathbf{w}_{-i}, G(\mathbf{x}, t)) \propto \frac{n_{k,-i}^w + \beta}{\sum_{w=1}^V (n_{k,-i}^w + \beta)} \cdot \frac{n_{G(\mathbf{x}, t), -i}^k + \alpha}{\sum_{k=1}^K (n_{G(\mathbf{x}, t), -i}^k + \alpha)}, \quad (4)$$

where $n_{k,-i}^w$ counts the number of words of type w in topic k , excluding the current word w_i , and $n_{G(\mathbf{x}, t), -i}^k$ is the number of words with topic label k in neighborhood $G(\mathbf{x}, t)$, excluding the current word w_i .

Several different strategies exist in the literature to do online refinement of the topic assignment in a given streaming dataset [18, 3]. The general idea is to initialize the topic label of the current observation with random labels, and then do a batch refinement of the entire dataset. This allows for previous topic assignments to be updated in the light of new observed data. Convergence is guaranteed because in the limit of time going to infinity, the algorithm behaves like a batch Gibbs sampler.

Since the number of neighborhoods is extremely large, we cannot refine word labels individually. In offline mode, sampling a neighborhood uniformly from all neighborhoods, computing its topic distribution, and refining all the word labels (using Eq 4) in it will approximate the proposed generative model. Online, we can do the same, however, in a robotics context, the number of refinements between two observation needs to be constant. We propose to sample the neighborhoods from a $\text{Beta}(a, 1)$ distribution, with $a > 1$, giving higher picking probability to recently observed regions. This ensures that new observations quickly converge, while older observations are less likely to change their labels. In this work, we set $a = 2$ for all experiments, however, increasing the value of a with time might lead to better results for long experiments. Algorithm 1 shows an iteration of the topic refinement strategy.

```

while no new observation do
   $W \leftarrow |\mathbf{w}|$  (* total number of observed words *)
  Randomly sample  $r \sim \text{Beta}(a, 1)$ 
   $i \leftarrow \lfloor W * r \rfloor$  (* index of the observation in the center of the neighborhood *)
  foreach  $j$  in  $G(\mathbf{x}_i, t_i)$  do
    (*update the topic label for each observation in the neighborhood *)
     $z_j \sim \mathbb{P}(z_j = k | \mathbf{z}_{-j}, w_j = w, \mathbf{w}_{-j}, G(\mathbf{x}_i, t_i))$ 
  end
end

```

Algorithm 1: REFINETOPICS ($\mathbf{z}, \mathbf{X}, \mathbf{w}, \mathbf{t}$). Refine topic labels, given the current assignment of topics (\mathbf{z}) for the set of all observed words (\mathbf{w}), their locations (\mathbf{X}), and observation times (\mathbf{t}).

2.3 Robot Control

Let $q(t) = \xi(M_t | \mathbf{S}) / 2H$ be the normalized surprise score of an incoming observation at time t . We then set the speed (v) of the robot by mapping the surprise score through a sigmoid function:

$$v(t) = \frac{1}{1 + e^{-\gamma(q(t) - 0.5)}}, \quad (5)$$

where γ controls the *responsiveness* of the robot. A higher γ made the scheduling of the forward velocity v more aggressive. We calculated γ empirically, and found $\gamma = 10$ to perform well during our sea trials.

The attitude (pitch, roll and heading) of the vehicle were kept stable during the experiments via simple PD controllers. In order to maintain depth, two different strategies were employed. Previously [14], we employed pitch corrections to maintain constant depths. This strategy was viable at high forward constant velocities. For the experiments in this paper however, the robot had to be able to control its depth, even when the speed v was 0. Early trials confirmed that at lower speeds ($v < 0.2$), controlling depth by changing the pitch angle was not possible. This is due to the fact that depth change via pitching is only possible for significant forward velocity v . Otherwise, the robot simply pitches up. Consequently, a heaving motion from a previously developed hovering gait [16] was used to maintain depth at lower speeds, or when the robot was standing nearly still over 'interesting' areas.

3 Results

We tested the proposed underwater system above a coral reef, in open waters. This demonstrated the applicability of our system to a real-life scenario that included unpredictable water currents, image noise due to floating particles (sediment and plankton) as well as illumination changes due to sunlight variations.

3.1 Spatiotemporal Topics based Surprise

We set both summary size and topic size to 6 for our experiments. The hyperparameters for the LDA were determined empirically. A video demonstration of the robot as it traverses a path and reacts to surprising observations is available online¹. There we can see the robot stopping over different, previously unobserved visual features, and then moving on at higher speeds when there is nothing of surprise. Some examples of the learned topic labels are shown in Fig 4. We see that the topics are representative of underlying physical phenomenon, and do well in describing scenes where a mixture of these exist. Red and blue topics are being used to represent rocks in the dataset, yellow for the sand-rock boundary, and cyan for the fire coral and the white rope.

Fig. 5(a) shows an example of the final summary generated by our online topics based summarizer from a sample trajectory. The corresponding histograms show the distribution of topics in the image. Fig. 5(b) shows uniformly sampled images over the same trajectory, presented here for comparison. We see that the proposed algorithm is able to recognize different species of corals (images 2 and 3), and the

¹ <http://cim.mcgill.ca/~yogesh/iser2012>

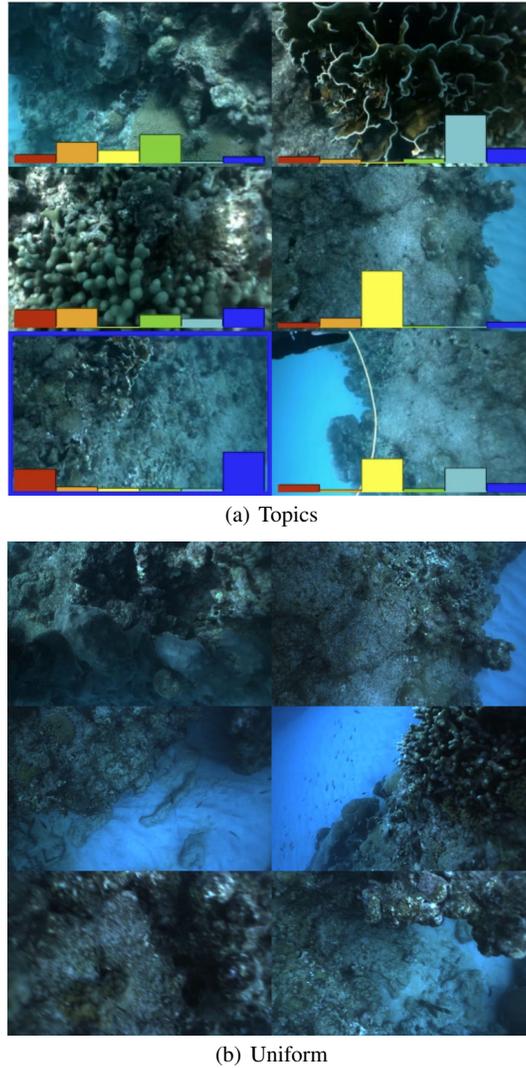


Fig. 5 (a) A summary of six images generated online by the system. The histogram shows the distribution of visual topics in the image, each color corresponding to a different topic. (b) For comparison we show images sampled uniformly over the robot trajectory.

accidental inclusion of a diver’s hand with a rope (image 6). When these images were observed, the robot evaluated them as surprising and as a result slowed down to a halt. Once these images are added to the summary, the surprise score falls instantly, and the robot continues forward in search for new surprises.

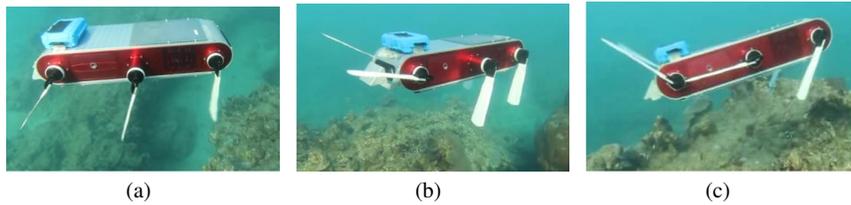


Fig. 6 Pictures showing the flippers' angle due to the action of the autopilot system, during one of the sea trials. (a) the robot is performing a heave-up maneuver to maintain depth and attitude at zero forward speed. (b) the robot is executing a combined heave up, pitch up and slow forward speed maneuver. (c) the robot is performing a pitch-up maneuver at high speed.

3.2 *Hovering Autopilot*

Fig. 6 shows the flipper placement configurations employed in the hovering gait and used to maintain attitude and depth, at different forward speeds v . As mentioned earlier, we employed two different strategies to maintain depth. At higher forward velocities ($v > 0.2$), depth was maintained via pitch angle changes, as depicted in Fig. 6 (c). By pitching up or down, the forward travel of the robot induced a change in depth, thus allowing depth control. When the robot had no forward velocity ($v = 0$), maintaining depth required the use of a heaving motion. This motion is accomplished by having the 6 flippers pointing upward or downward, as illustrated in Fig. 6 (a). This way, the net thrust produced by the oscillating flippers does not induce forward motion. Attitude stabilization is still possible with this leg configuration, by means of a forward/aft differential thrust for pitch corrections and left/right differential thrust for roll corrections. For low velocities ($v < 0.2$), the robot flippers were placed so as to generate both heaving and forward motion (Fig. 6 (b)). All of these pictures were taken from a single trial, demonstrating the need to adapt the locomotion strategy in order to satisfy motion requirements.

During the trials, we noted a decreased performance of depth and attitude control at lower velocities. This is a common phenomenon, as controllability of a vehicle operating in a fluid decreases with its velocity, due to the reduced lift forces on the control surfaces (flippers). This reduction in controllability can be seen for depth in Fig. 7 d). At commanded speed $v > 0.5$ (corresponding roughly to forward velocities above 20 cm/s), the depth error is reduced. For example, at $t = 63\text{s}$ where the speed command v is above 0.5, the depth static error is less than 0.1 m . At low velocities v , this static error in depth increased and was generally positive (i.e., deeper than commanded), due to the fact that the robot was negatively buoyant, i.e., had a tendency to sink when standing still.

One issue plaguing the autopilot was unwanted roll oscillations of $\pm 10^\circ$ at low velocities v . These oscillations had never been observed with the autopilot system before, when the robot was operated at medium to high velocities (above 20 cm/s). We suspect that these were caused by an unforeseen change in the hydrodynamic

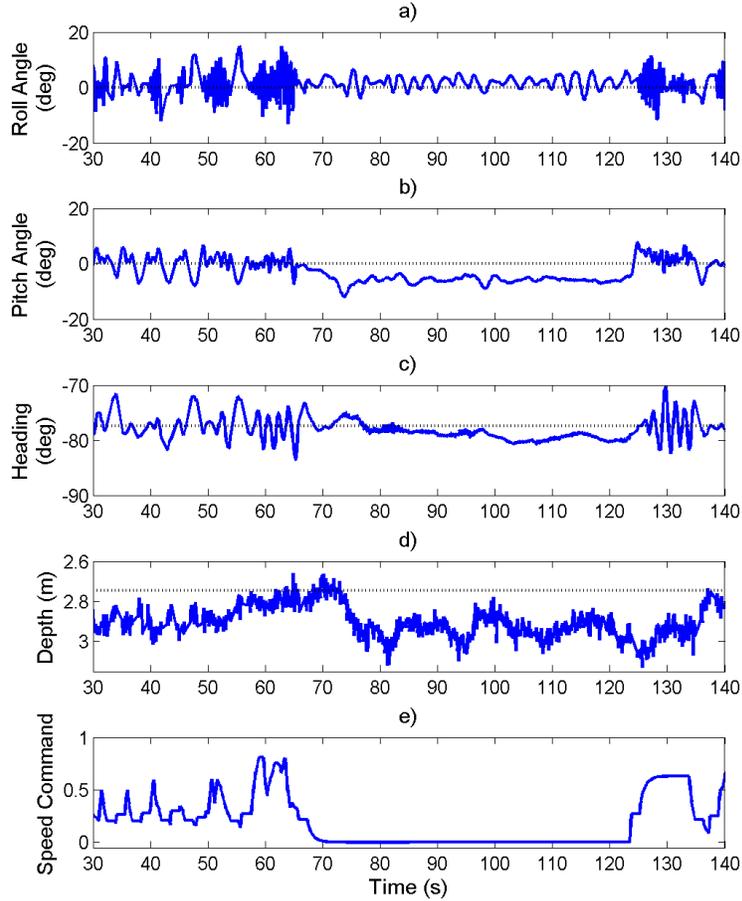


Fig. 7 Performance of the autopilot system in hovering mode, with robot attitude a)-c) and depth d) over time, during one of the sea trials. The dotted lines in a)-d) correspond to the target values sent to the autopilot system. The open-loop speed command v in e) provided by the surprise module is unit-less: a value of 1.0 corresponded to maximum forward velocity of the vehicle.

behavior of the vehicle and its flippers at low velocity. As a consequence, controller improvements at low velocities are under investigation.

4 Conclusion

We have demonstrated a novel autonomous robotic system that can be used to assist in exploring dangerous underwater environments. Our use of an online spatiotem-

poral topic modeling was an attempt to model semantic surprise, which is sensitive to presence of high-level patterns in the scene, such as different coral species, rocks, and sand. Given a fixed trajectory, the robot traversed it with a non-uniform speed, stopping at locations containing surprising observations, and moving at high speeds over seemingly boring or previously observed regions. From a locomotion point of view, these experiments helped validate the use of heaving motion for depth control at low velocities, as well as highlighted certain deficiencies of our roll control system in that regime. The resulting summaries produced by our system is able to capture the visual diversity of the underwater environment. Our ongoing future work is focused on developing better realtime online topic modeling techniques, such as the use of nonparametric hierarchical Dirichlet processes, and their use in control of different robotic platforms for exploration tasks.

References

1. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *The Journal of Machine Learning Research* **3**, 993–1022 (2003)
2. Bourque, E., Dudek, G.: Automated Image-Based Mapping. In: *IEEE Computer Vision and Pattern Recognition (CVPR)–Workshop on Perception of Mobile Agents*, pp. 61–70 (1998)
3. Canini, K.R., Shi, L., Griffiths, T.L.: Online Inference of Topics with Latent Dirichlet Allocation. *Proceedings of the International Conference on Artificial Intelligence and Statistics* **5**(1999), 65–72 (2009)
4. Charikar, M., Chekuri, C., Feder, T., Motwani, R.: Incremental clustering and dynamic information retrieval. In: *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing - STOC '97*, pp. 626–635. ACM Press, New York, New York, USA (1997). DOI 10.1145/258533.258657. URL <http://dl.acm.org/citation.cfm?id=258533.258657>
5. Girdhar, Y., Dudek, G.: ONSUM: A System for Generating Online Navigation Summaries. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS)*, pp. 746–751 (2010). DOI 10.1109/IROS.2010.5650315. URL <http://cim.mcgill.ca/yogesh/publications/iros2010.pdf>
6. Girdhar, Y., Dudek, G.: Efficient on-line data summarization using extremum summaries. In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3490–3496 (2012). DOI 10.1109/ICRA.2012.6224657. URL <http://cim.mcgill.ca/yogesh/publications/icra2012.pdf>
7. Gonzalez, T.F.: Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science* **38**, 293–306 (1985). DOI 10.1016/0304-3975(85)90224-5
8. Grabowski, R., Khosla, P., Choset, H.: Autonomous exploration via regions of interest. In: *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, vol. 2, pp. 1691–1696. IEEE (2003). DOI 10.1109/IROS.2003.1248887. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1248887>
9. Griffiths, T.L., Steyvers, M.: Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America* **101**(Suppl 1), 5228–5235 (2004). DOI 10.1073/pnas.0307752101. URL <http://www.pnas.org/content/101/suppl.1/5228.abstract>
10. Hofmann, T.: Unsupervised Learning by Probabilistic Latent Semantic Analysis. *Machine Learning* **42**(1), 177–196 (2001). DOI 10.1023/A:1007617005950
11. Hsu, W.L., Nemhauser, G.L.: Easy and hard bottleneck location problems. *Discrete Applied Mathematics* **1**(3), 209–215 (1979). DOI DOI: 10.1016/0166-218X(79)90044-1. URL <http://www.sciencedirect.com/science/article/B6TYW-46HW0N4-5/2/1a9ef82be9ca95988c9d75e76c216024>

12. Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J.J., Dellaert, F.: iSAM2: Incremental smoothing and mapping using the Bayes tree. *The International Journal of Robotics Research* **31**(2), 216–235 (2011). DOI 10.1177/0278364911430419. URL <http://ijr.sagepub.com/cgi/content/abstract/31/2/216>
13. Paul, R., Rus, D., Newman, P.: How was your day? Online Visual Workspace Summaries using Incremental Clustering in Topic Space. In: *IEEE International Conference on Robotics and Automation (ICRA)* (2012)
14. Plamondon, N., Nahon, M.: A trajectory tracking controller for an underwater hexapod vehicle. *Bioinspiration and Biomimetics* **4**(3), 036,005 (2009)
15. Sattar, J., Dudek, G., Chiu, O., Rekleitis, I., Giguère, P., Mills, A., Plamondon, N., Prahacs, C., Girdhar, Y., Nahon, M., Lobos, J.P.: Enabling Autonomous Capabilities in Underwater Robotics. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS)*, pp. 3628–3634. Nice, France (2008). DOI 10.1109/IROS.2008.4651158. URL <http://cim.mcgill.ca/yogesh/publications/iros2008.pdf>
16. Sattar, J., Giguere, P., Dudek, G.: Sensor-based behavior control for an autonomous underwater vehicle. *I. J. Robotic Res.* **28**(6), 701–713 (2009)
17. Sivic, J., Zisserman, A.: Video Google: Efficient Visual Search of Videos. In: J. Ponce, M. Hebert, C. Schmid, A. Zisserman (eds.) *Toward Category-Level Object Recognition, Lecture Notes in Computer Science*, vol. 4170, pp. 127–144. Springer Berlin / Heidelberg (2006). DOI 10.1007/11957959_7. URL http://dx.doi.org/10.1007/11957959_7
18. Song, X., Lin, C.Y., Tseng, B.L., Sun, M.T.: Modeling and predicting personal information dissemination behavior. In: *Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining - KDD '05*, p. 479. ACM Press, New York, New York, USA (2005). DOI 10.1145/1081870.1081925. URL <http://dl.acm.org/citation.cfm?id=1081870.1081925>
19. Thrun, S., Thayer, S., Whittaker, W., Baker, C., Burgard, W., Ferguson, D., Hanel, D., Montemerlo, M., Morris, A., Omohundro, Z., Reverte, C.: Autonomous exploration and mapping of abandoned mines. *IEEE Robotics & Automation Magazine* **11**(4), 79–91 (2004). DOI 10.1109/MRA.2004.1371614. URL <http://ieeexplore.ieee.org/xpl/freeabs.all.jsp?arnumber=1371614>
20. Wang, X., Grimson, E.: Spatial Latent Dirichlet Allocation. In: *Advances in Neural Information Processing Systems*, vol. 20, pp. 1577–1584 (2007). URL <http://people.csail.mit.edu/xgwang/papers/STLDA.pdf>
21. Whaite, P., Ferrie, F.: Autonomous exploration: driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(3), 193–205 (1997). DOI 10.1109/34.584097. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=584097>
22. Yamauchi, B.: A frontier-based approach for autonomous exploration. In: *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*, pp. 146–151. IEEE Comput. Soc. Press (1997). DOI 10.1109/CIRA.1997.613851. URL <http://ieeexplore.ieee.org/xpl/freeabs.all.jsp?arnumber=613851>