

Exploring Underwater Environments with Curiosity

Yogesh Girdhar and Gregory Dudek
Centre for Intelligent Machines
McGill University, Montreal, Canada
{yogesh,dudek}@cim.mcgill.ca

Abstract—This paper presents a novel approach to modeling curiosity in a mobile robot, which is useful for monitoring and adaptive data collection tasks. We use ROST, a realtime topic modeling framework to build a semantic perception model of the environment, using which, we plan a path through the locations in the world with high semantic information content. We demonstrate the approach using the Aqua robot in a variety of different scenarios, and find the robot be able to do tasks such as coral reef inspection, diver following, and sea floor exploration, without any prior training or preparation.

I. INTRODUCTION

Gaining knowledge about our environment is a never-ending quest for humanity. Direct exploration by humans although tempting, puts strong limitations on what can be explored. Fortunately through the use of robotics, we can continue this tradition of exploration, without putting human lives at risk. Use of autonomous robots is especially required for space and ocean exploration due to the inherent communication bottlenecks which do not allow remote control of the vehicles [1].

A simple approach for autonomous collection of environment data is to use space filling paths through the environment. This, however is not ideal, because the amount of information collected, associated with the different spatial phenomena, is proportional to the spatial area covered by them. Underwater this might mean that most of the data collected, only contains observations of sand or rocks, and very occasionally we might have a few samples corresponding to something interesting such as thermal vents, marine life, or archeological sites. A better strategy for collecting data is to have the robot behave like an explorer, or a vacationing tourist; moving swiftly over regions with familiar sights, while paying much more attention, i.e., collecting more data when something novel or interesting is in view. In this paper we describe such a techniques, and demonstrate its functioning on an underwater robot.

In our previous work we have presented ROST [12] a realtime online topic modeling framework that can be used by a robot for high level perception of its environment. In [13], we showed that using ROST we can identify exploration paths of high information gain that result in learning better terrain models. In a continuation of our previous work, in this paper we demonstrate the exploration system described in [13] working on Aqua [10] amphibious robot.

Environment sensing is performed primarily through the use of a wide angle camera mounted in front of the robot.

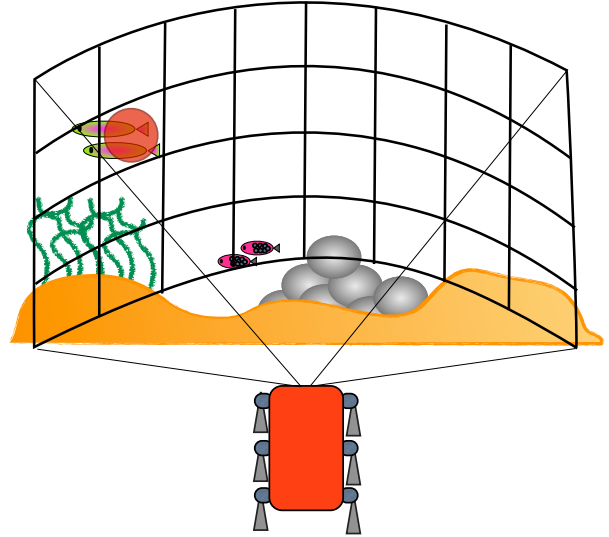


Fig. 1. The robot explores the environment by finding novel observations, while building a topic model of its experience.

The image stream captured by the robot is processed by ROST, which gives different parts of the image topic labels that are representative of the high level scene constructs such as corals, plants, diver and rocks. Low level visual patterns that commonly occur together in space and time are more likely to be given the same topic label. ROST uses Dirichlet distributions to model the topic distribution in different spatiotemporal regions of the video stream, and uses a real-time Gibbs sampler to keep the topic model in a converged state. Given the description of the current scene in topic space, the robot then identifies the part of the scene with most novelty. We use a winner-take-all strategy to determine this area of attention with maximum information gain, and then navigate the robot laterally in the direction of this point. To mitigate the chances of the robot getting stuck in a local maxima, we keep the robot speed constant, while only controlling the yaw.

II. RELATED WORK

In the following sections we briefly look at some common variants of the exploration problem.

A. Coverage of Known Environments

If we have prior knowledge about the world then perhaps the simplest form of exploration is coverage, where the goal is to make the robot pass through every point in the given

spatial region of interest. If the space is free of obstacles, then we can simply use a zig-zag path, sometimes known as a boustrophedon path through to cover the world. In case of known obstacles, Choset et al. [7] proposed boustrophedon cell decomposition of the world such that each cell can be covered by a simple boustrophedon path, and then given this decomposition, plan a path through all the cells. This would result in complete coverage.

Mannadiar and Rekleitis [19] later proposed splitting some boustrophedon cells so that robot does not need to move over previously covered cells, resulting in paths guaranteeing optimal coverage. These paths have been extended for use with the general class of non-holonomic robots such as aerial vehicles [30].

B. Exploration for Improving Navigation

Navigating a robot through free space is a fundamental problem in robotics. Yamauchi [31] defined exploration as the “act of moving through an unknown environment while building a map that can be used for subsequent navigation”. Yamauchi’s proposed solution involved moving the robot towards the frontier regions in the map, which were described as the boundary between known free space and the uncharted territories.

If we have an inverse sensor model of the range sensor, it is possible to compute locations in the world which would maximize the utility of the sensor reading in resolving obstacle position and shape. Grabowski [14] proposed such an exploration strategy where the goal is to maximize the understanding of obstacles rather than the exposure to free space. In this approach, the robot identifies the location with next best view in space where a sonar sensor reading would have the greatest utility in improving the quality of representation of an obstacle.

If there is no external localizer available to the robot, then it is desirable that robot explores, maps and localizes in the environment at the same time. Sim, Dudek and Roy [24] take the approach of finding trajectories at each step that explore new regions while minimizing the localization uncertainty of the robot as it re-enters a previously mapped region.

Bourgault [5] and Stachniss [28] have proposed an exploration strategy which uses gradient ascent to move the robot towards areas of high entropy which would maximize map information gain, while still keeping the robot localized.

Kollar and Roy [18], [17] formulated the exploration problem as a constrained optimization problem, where the goal is to find a path that maximizes map accuracy with the constraint of complete map coverage. To do this the algorithm first identifies the locations in the map that are essential for coverage, and then uses these locations to constraint the trajectory that maximizes map accuracy.

C. Exploration for Monitoring Spatiotemporal Phenomenon

In underwater and aerial environments, obstacle avoidance and map building tasks are typically not of primary concern.

Binney [3] has described an exploration technique to optimize the monitoring spatiotemporal phenomena by taking advantage of the submodularity of the objective function. Bender [2] has proposed a Gaussian process based exploration technique for benthic environments, which uses an experiment specific utility function. Das et al. [8] have presented techniques to autonomously observe oceanographic features in the open ocean. Hollinger et al. [16] have studied the problem of autonomously studying underwater ship hulls by maximizing the accuracy of sonar data stream. Smith et al. [26] have looked at computing robot trajectories which maximize information gained, while minimizing the deviation from the planned path.

III. REALTIME TOPIC MODELING OF SPATIOTEMPORAL DATA

Given the quantized observations (words) and their location, we would like to compute the posterior distribution of the corresponding topic labels at this location that are representative of the underlying cause responsible for emitting the words. Let w be the observed word at location x . We assume the following generative process for these observation words:

- 1) Word distribution for each topic k :

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

- 2) Topic distribution for words at location x :

$$\theta_x \sim \text{Dirichlet}(\alpha)$$

- 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 , and z is the topic label for the word observation w . Each topic is modeled by distribution ϕ_k over V possible word in the observation vocabulary.

$$\phi_k(v) \propto n_k^v + \beta, \tag{1}$$

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 Latent Dirichlet Allocation (LDA)[4] is in step 2. The context of words in LDA are 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. Topic distribution at location x is thus defined as:




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

$$\theta_x(k) \propto \left(\sum_y K(x-y)n_y^k \right) + \alpha, \quad (2)$$

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

A. 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 us to keep track of the neighborhood topic distribution for every observation word, which is computationally infeasible for any large dataset. For the special case when our kernel is a uniform distribution over a finite region, we can assume a cells 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 grid decomposition of the world where each cell is connected to its four nearest neighbors. However, the general idea presented here are applicable to any other topological decomposition of spacetime.

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) topic distribution of the neighborhood is computed by summing over topic distribution of the neighboring cells rather than individual points.

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

B. Realtime Inference using Gibbs Sampling

Given a word observation w_i , and 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:

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

where $n_{k,-i}^w$ counts the number of words of type w in topic k , excluding the current word w_i , and $n_{G_i,-i}^k$ 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, and the above Gibbs sampler is equivalent to the LDA Gibbs sampler proposed by Griffiths et al.[15], where each cell corresponds to a document.

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 [27], [6], is not applicable. Hence we must use a refinement strategy that only partially updates the topic labels after each time step while still ensuring the labels from the last observation have converged. Given the cellular decomposition of the data, we do this by randomly picking cells by sampling from a distribution which favors recent observations, and then refining the topic labels for the words within. To ensure that the topic labels for the last observation have converged, at each time step, for each of the R iterations, with constant probability τ we refine the last observation, and with probability $(1 - \tau)$ refine another observation picked randomly with probability proportional to its age.

This strategy ensure that each observation is refined for at least τR number of times, and as a result all observations have topic labels which have converged. This strategy is hence suitable for use in long-running real-time systems. Empirically, we found $\tau = 0.5$ to work well in most cases, however on faster machines, τ could be set to a lower value, which would encourage better globally optimal topic labels. Algorithm 1 summarizes the proposed realtime topic refinement strategy.

IV. CURIOSITY BASED EXPLORATION

Assume a cellular decomposition of the world, with C cells, the goal then is to plan a continuous path $P \subseteq C$, which allows us to learn the topic model $\Phi = \{\phi_k\}$ that best describes the world by labeling each observation at each location with a representative topic label.

At time t , let the robot be in cell $p_t = c$, and let $G(c) = \{g_i\}$ be the set of cells in its neighborhood. We would like to compute a weight value for each g_i , such that the probability of the robot taking a step in this direction is proportional to this weight.

$$\mathbb{P}(p_{t+1} = g_i) \propto \text{weight}(g_i). \quad (5)$$

```

while true do
  while no new observation do
     $a \sim \text{Bernoulli}(\tau)$ 
    if  $a == 0$  then
      (*pick the cell from the last timestep*)
       $t \leftarrow T$ 
    else
      (*pick a cell with probability proportional to its
      timestamp*)
       $t \sim \mathbb{P}(t = j) \propto j, 1 \leq j \leq T$ 
    end
    foreach word  $(w_i, x_i)$  in  $c_t$  do
      (*update the topic label for word in the
      observation *)
       $z_i \sim \mathbb{P}(z_i = k | w_i = v, x_i)$ 
    end
  end
   $T \leftarrow T + 1$ 
  Add new observed words to their corresponding cells.
end

```

Algorithm 1: Keep topic labels up-to-date as new observations arrive.

When there is no location information present we propose a weight function that biases the next step towards the cell which has high topic perplexity.

$$\begin{aligned} \text{weight}(g_i) &= \text{TopicPerplexity}(g_i) & (6) \\ &= \exp\left(-\frac{\sum_i^W \log \mathbb{P}(z_i = k | P)}{W}\right). & (7) \end{aligned}$$

To compute topic perplexity of the words observed in g_i , we first compute topic labels z_i for these observed words by sampling them from the distribution in Eq. 4, without adding these words to the topic model. These temporary topic labels are then used to compute the perplexity of g_i in topic space.

When there is location of the robot is known, we add a repulsion term to the weight function, which prevents the robot from being stuck in a local maxima.

$$\text{weight}(g_i) = \frac{\text{TopicPerplexity}(g_i)}{\sum_j n_j / d^2(p_i, g_j)}. \quad (8)$$

Here n_j is the number of times we have visited cell g_j , and $d(p_i, g_j)$ is the Euclidean distance between these two cells.

V. IMPLEMENTATION

A. Robot Platform

Aqua [10] is an amphibious six legged robot capable of autonomous operations. Although vehicles like Dorado [23] and Hugin [20] are extremely capable for long distance and deep water exploration, Aqua’s propulsion is based on six flippers that can provide motion in five degrees of freedom is most suitable for tasks requiring high maneuverability, such as coral reef exploration. By using a novel combination of

gaits, the Aqua can move at various speeds while maintaining its orientation, despite external disturbances [11], [21].

Aqua is equipped with four cameras: a fish-eye camera in the front for environment awareness, a pair of stereo cameras in the front for depth perception, and a back camera. It is possible to take downward-looking images via the use of a mirror mount.

Aqua is also capable of walking on land, and through the use of a recently designed flippers [9] can perform both swimming and walking using the same set of flippers.

B. Visual Words

For modeling visual data observed by the robot, instead of text words, we use three different kinds of visual words: Oriented BRIEF (ORB) [22] based visual words [25] that describing local visual features, texton words [29] in Lab color space to describe texture properties of a region, and pixel hue words.

To generate an ORB word vocabulary, we first extract features from an unrelated dataset¹, with high visual diversity. These features are then clustered using the k-means algorithm, with $V = 5000$ clusters corresponding to the desired vocabulary size. The cluster centers of these V clusters represent the visual words in the vocabulary. We use ORB instead of SURF or SIFT because these descriptors are binary, and the distance between two feature vectors can hence be computed by taking the Hamming distance between the bit strings, which can be done very efficiently using XOR operations.

We computed texton words by computing Gabor filter response for 4 different scales, 4 different orientations, and 2 different color channels, for a total of 32 dimensions. Then, similar to ORB words, we clustered these texton features into 1000 clusters to generate a vocabulary.

Apart from ORB and texton words, we also use “pixel color words” to represent the average color around a spatial region. We use 1000 pixel color words per image, distributed uniformly over the image.

VI. RESULTS

A. Simulated Exploration on a 2D Map

To validate our claim that the proposed approach would result in collection of more information, we experimented with simulated exploration on a 2D map of a coral reef, shown in Fig. 3(a). The 2560x2560 pixel map was decomposed into square cells of width 32 pixels. We computed exploration paths of varying length, with 20 different random restart locations for different exploration algorithms. Each time step was fixed at 200 milliseconds to allow the topic model to converge. While observing a cell we only use randomly chosen 25% of the word observations in that cell to better simulate repeated observations at the same location.

We compared the proposed curiosity driven exploration with two other exploration algorithms: Brownian motion, which is

¹We used the documentary movie Baraka(1992) for extracting visual feature, because of its rich visuals from many different contexts

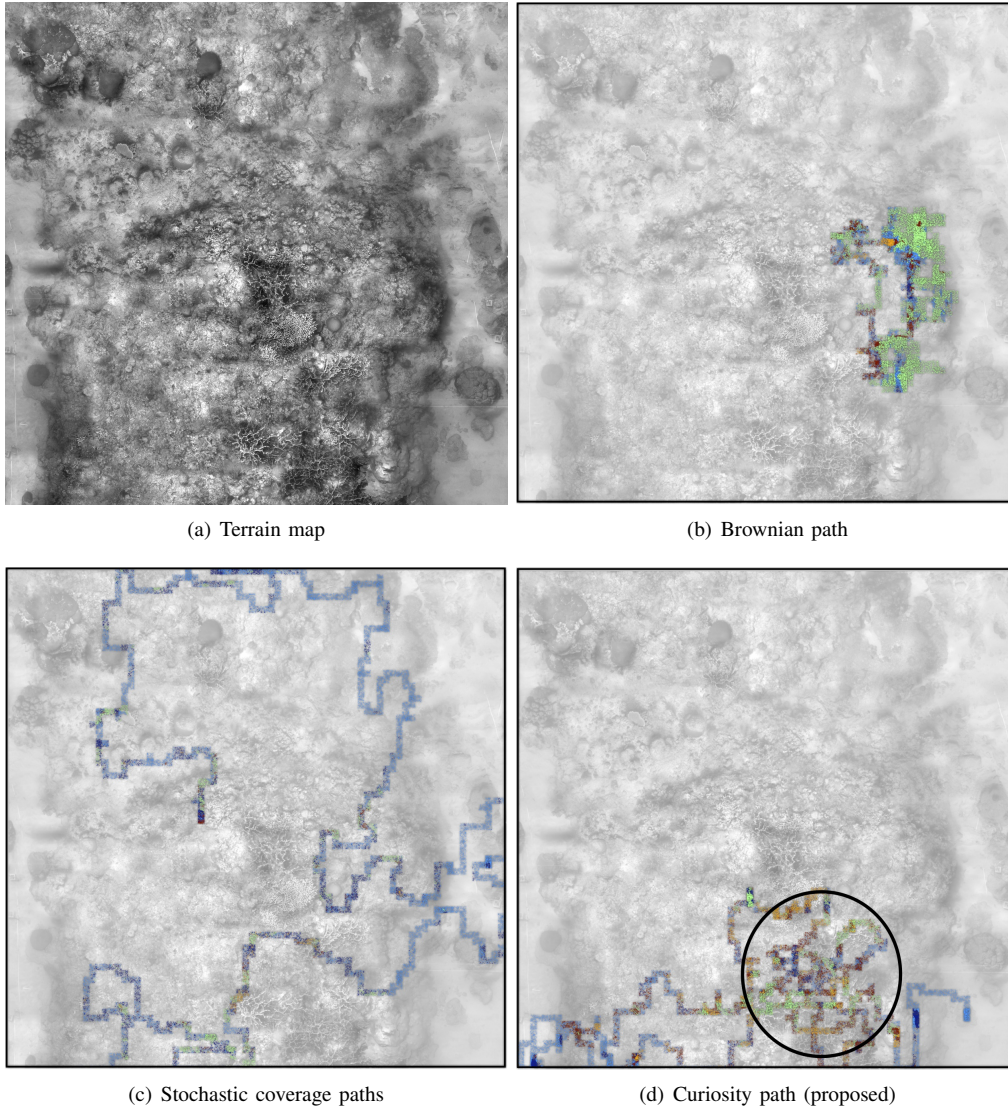


Fig. 3. Examples of different exploration paths of length 640 steps. Points on the paths are colored according to the topic labels of the words observed at that location. (a) Map used for the simulated exploration experiment. (b) A typical brownian motion exploration path with minimal coverage. (c) A typical stochastic coverage exploration path. We see that these paths are likely to encounter much more diversity of terrain. (d) A typical curiosity driven exploration path. We see that unlike stochastic coverage the proposed algorithm collects more data at locations where there is more information, such as the circled region containing corals.




Fig. 4. Information gained in topic space, as a function of path length.

equivalent to setting the weight function in Eq. 8 to a constant value, implying no bias towards any neighbor; and stochastic coverage, which is equivalent to setting topic perplexity in Eq. 8 to 1, implying attraction solely towards previously unvisited location while ignoring their content.

Fig. 3 shows typical paths from these three exploration algorithms. We see that the proposed curiosity based algorithm favors locations with high information content, which in this case correspond to locations with corals. Given these exploration paths, we computed the amount of information gained, measured as Shannon entropy of the topic labels observed along the exploration path. In 4 we show this information gain for the three approaches and varying path lengths. We see that the proposed technique performs significantly better than other

techniques for all path lengths.

B. Underwater Exploration using Aqua

We implemented the proposed curiosity modeling system on Aqua amphibious robot, and tested it in many different underwater scenarios, three of which are shown in the video attachment². In this video we see the robot exploring its environment from two different points of view. We split the observed image into 8x4 cells. We color the cells in robot's view with blue, and change the opacity based on the perplexity score. A cell marked with more opaque blue circle has higher topic perplexity score, and the cell with the highest score is marked with a red color. Figure 5 shows some examples of these high perplexity regions in observed images by the robot. For all our experiments, we fixed the number of topics to $K = 64$, and set Dirichlet hyper-parameters $\alpha = 0.1$, $\beta = 0.1$, refinement bias $\tau = 0.5$. The system was set to run at 4Hz., giving each image a processing time of 250ms.

The three scenarios which we show in the video are:

- 1) *Exploring a coral head*: In this trial, we started the robot near a coral head surrounded by monotonous sand. We see that the robot quickly gets attracted towards the coral head, and continues to bounce around over this structure while staying away from sand. We see the effect of curiosity decay variable γ , as the robot is successfully able to return back to the coral head several times after going over the much less interesting sandy regions.
- 2) *Interaction with a diver* Although our goal was to study the robot as it would interact with a fish, due to lack of cooperation with the fish, we were forced to conduct the experiment with a scuba diver instead. We see that as soon as the diver is in robot's view, it is the singular source of curiosity for the robot. We see the robot following the diver around, and hovering over the diver when he has stopped moving.
- 3) *Exploring the ocean floor* In this trial, we started the robot near the ocean floor, which was sparsely populated with sea plants and corals. We see the robot manages to keep its focus on sea life, while not wasting time over sand.

Overall we see that although the behavior of the robot is purely statistically defined, it results in collecting more data about perceptually interesting things in the scene.

VII. CONCLUSION

We have presented the idea of a curiosity driven robot, which plans a path that aims to maximize the information gained in semantic space. We used ROST topic modeling framework to describe the observations semantically in real-time, and then at each time step moved in the direction of the neighboring location with the highest information content. Empirically we showed that the proposed exploration results in collecting significantly more information compared to spatial coverage or brownian motion. In an an underwater scenario,

this corresponds to more observations of rare corals. We implemented the approach on Aqua underwater robot, and found the emergent behavior to resemble an intelligent life form. In the video demonstrating we showed the behavior of the robot in a variety of situations such as: inspecting coral reefs, exploring sea floor, and following divers, without changing any system parameters. The fact that the same technique can be used in many different scenarios to do different tasks is encouraging.

ACKNOWLEDGMENT

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) through the NSERC Canadian Field Robotics Network (NCFRN).

REFERENCES

- [1] J. G. Bellingham and K. Rajan. Robotics in remote and hostile environments. *Science (New York, N.Y.)*, 318(5853):1098–102, Nov. 2007.
- [2] A. Bender, S. B. Williams, and O. Pizarro. Autonomous Exploration of Large-Scale Benthic Environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 390–396, 2013.
- [3] J. Binney, A. Krause, and G. S. Sukhatme. Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research*, 32(8):873–888, July 2013.
- [4] 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)*, 16:106–114, 2004.
- [5] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte. Information based adaptive robotic exploration. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, pages 540–545, 2002.
- [6] 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*, 5(1999):65–72, 2009.
- [7] H. Choset and P. Pignon. Coverage path planning: The boustrophedon cellular decomposition. *Field and Service Robotics*, 1998.
- [8] J. Das, F. Py, T. Maughan, T. O'Reilly, M. Messie, J. Ryan, G. S. Sukhatme, and K. Rajan. Coordinated sampling of dynamic oceanographic features with underwater vehicles and drifters. *The International Journal of Robotics Research*, 31(5):626–646, Apr. 2012.
- [9] B. B. Dey, S. Manjanna, and G. Dudek. Ninja legs: Amphibious one degree of freedom robotic legs. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5622–5628. IEEE, Nov. 2013.
- [10] G. Dudek, P. Giguere, C. Prahacs, S. Saunderson, J. Sattar, L.-A. Torres-Mendez, M. Jenkin, A. German, A. Hogue, A. Ripsman, J. Zacher, E. Milios, H. Liu, P. Zhang, M. Buehler, and C. Georgiades. AQUA: An Amphibious Autonomous Robot. *Computer*, 40(1):46–53, 2007.
- [11] P. Giguere, Y. Girdhar, and G. Dudek. Wide-Speed Autopilot System for a Swimming Hexapod Robot. In *Canadian Conference on Computer and Robot Vision (CRV)*, pages 9–15, 2013.
- [12] Y. Girdhar, P. Giguere, and G. Dudek. Autonomous adaptive exploration using realtime online spatiotemporal topic modeling. *The International Journal of Robotics Research*, Nov. 2013.
- [13] Y. Girdhar, D. Whitney, and G. Dudek. Curiosity Based Exploration for Learning Terrain Models. In *IEEE International Conference on Robotics and Automation*, page 7, Oct. 2014.
- [14] R. Grabowski, P. Khosla, and H. Choset. Autonomous exploration via regions of interest. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, volume 2, pages 1691–1696. IEEE, 2003.
- [15] T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America*, 101(Suppl 1):5228–5235, 2004.
- [16] G. a. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme. Active planning for underwater inspection and the benefit of adaptivity. *The International Journal of Robotics Research*, 32(1):3–18, Nov. 2012.
- [17] T. Kollar and N. Roy. Efficient Optimization of Information-Theoretic Exploration in SLAM. In *AAAI*, volume 8, pages 1369–1375, 2008.

²http://cim.mcgill.ca/~yogesh/publications/crv2014_video.mp4




Fig. 5. Examples of observations showing cells marked with their curiosity score. Red marks the cell with the highest score.

- [18] T. Kollar and N. Roy. Trajectory optimization using reinforcement learning for map exploration. *The International Journal of Robotics Research*, 27(2)(June):175–197, 2008.
- [19] R. Mannadiar and I. Rekleitis. Optimal coverage of a known arbitrary environment. In *2010 IEEE International Conference on Robotics and Automation*, pages 5525–5530. IEEE, May 2010.
- [20] R. Marthiniussen, K. Vestgird, and R. A. Klepaker. HUGIN-AUV concept and operational experiences to date. In *OCEANS '04. MTS/IEEE TECHNO-OCEAN '04*, pages 846–850, 2004.
- [21] N. Plamondon. *Modeling and control of a biomimetic underwater vehicle*. PhD thesis, McGill University, 2010.
- [22] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. ORB : an efficient alternative to SIFT or SURF. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2564—2571, 2011.
- [23] M. Sibenac, W. J. Kirkwood, R. Mcewen, F. Shane, and R. Henthom. Modutar AUV for Routine Deep Water Science Operations. In *OCEANS '02 MTS/IEEE*, pages 167–172, 2001.
- [24] R. Sim, G. Dudek, and N. Roy. Online control policy optimization for minimizing map uncertainty during exploration. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 2, pages 1758–1763. IEEE, 2004.
- [25] J. Sivic and A. Zisserman. Video Google: Efficient Visual Search of Videos. In J. Ponce, M. Hebert, C. Schmid, and A. Zisserman, editors, *Toward Category-Level Object Recognition*, volume 4170 of *Lecture Notes in Computer Science*, pages 127–144. Springer Berlin / Heidelberg, 2006.
- [26] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme. Persistent Ocean Monitoring with Underwater Gliders: Adapting Sampling Resolution. *Journal of Field Robotics*, 28(5):714–741, 2011.
- [27] X. Song, C.-Y. Lin, B. L. Tseng, and M.-T. Sun. Modeling and predicting personal information dissemination behavior. In *Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining - KDD '05*, page 479, New York, New York, USA, Aug. 2005. ACM Press.
- [28] C. Stachniss, G. Grisetti, and W. Burgard. Information Gain-based Exploration Using Rao-Blackwellized Particle Filters. *Robotics: Science and Systems*, 2, 2005.
- [29] M. Varma and A. Zisserman. A Statistical Approach to Texture Classification from Single Images. *International Journal of Computer Vision*, 62(1/2):61–81, Apr. 2005.
- [30] A. Xu, C. Viriyasuthee, and I. Rekleitis. Optimal complete terrain coverage using an Unmanned Aerial Vehicle. In *2011 IEEE International Conference on Robotics and Automation*, pages 2513–2519. IEEE, May 2011.
- [31] B. Yamauchi. 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'*, pages 146–151. IEEE Comput. Soc. Press, 1997.