SMOKEBOT

Mobile Robots with Novel Environmental Sensors for Inspection of Disaster Sites with Low Visibility

Project start: January 1, 2015

Duration: 3.5 years

Due date: month 36 (December 2017)

Lead beneficiary: ORU

Dissemination Level: PUBLIC

Main Authors:

Erik Schaffernicht (ORU)
Victor Hernandez (ORU)
Han Fan (ORU)

Version History:

0.1 initial version, es, Feb 18
0.2 first revision, es,vh,hf, May 18
0.3 second revision, es,vh, Jun 18

Contents

Abstract

This document discusses the generation of the so called *hazard map* for the Horizon 2020 project SmokeBot. SmokeBot's objective is to improve the application of mobile robots in disaster scenarios with low visibility conditions. The purpose of the hazard map is to provide information about dangerous areas in the environment in form of a cost map for the navigation component of the SmokeBot system.

Thereto, this document discusses different options to create the hazard cost map. The map is derived from the environment information (i.e. the grid map generated by the SLAM system T1.5) and the hazard prediction system (T4.2) to which it is closely related.

1. Introduction

The SmokeBot autonomous navigation system (T5.3) uses a cost map to plan and execute the optimal trajectory between navigation points in a given environment. In its simplest form this cost map is a binary occupancy grid that encodes obstacles and free space in the environment. As such an approach often leads undesirable behavior during plan execution (e.g. cutting corners closely), a common extension is to encode the distance to the nearest obstacle in the cost map implementing a potential field approach to navigation.

In the SmokeBot project the robot not only has to deal with impassable objects encoded in typical grid maps, but also with other dangers that might damage the robot or impair sensor systems. Those hazards include hot areas, the presence of harmful gases or the presence of smoke, which are defined as low-level hazards (see Deliverable 4.2 for a full discussion), but also areas that might be subject to sudden harmful changes (e.g. a flash-over), which were defined as high-level hazards in D4.2.

In order to avoid potential dangerous areas as well as obstacles during navigation the information from the obstacle maps and the hazard prediction system are combined in the hazard map. The resulting hazard map is a 2D grid map, which contains a cost value representing the degree of danger in each respective cell. The remainder of this deliverable explores different options to integrate this additional hazard information in such a cost map.

From a practical point of view, the boundaries between the hazard prediction task (T4.2 and the associated Deliverable D4.2) and the hazard mapping task discussed in this document are not as clear and easily separable as envisioned in the project proposal, hence there is a certain overlap and redundancy between those two tasks, which is particularly visible in the handling of spatial and temporal aspects surrounding the hazard prediction and mapping.

2. Basic Concept

In order to facilitate autonomous navigation the information about hazards in the environment (as described in deliverable D4.2) is condensed to a single number, the degree of danger indicator, for each traversable place in the environment.

Thereto, a state estimator for the degree of danger in each cell is required. As a result of D4.2 the information about the different hazards are available as map layers in GDIM (see deliverable D4.1), which takes care of aligning the maps to a common coordinate system and grid structure.

This common representation structure allows combining the information easily into a hazard map on a cell by cell basis[1]. Figure 1 visualizes this basic idea. We will discuss different ways of combining the different layers next.

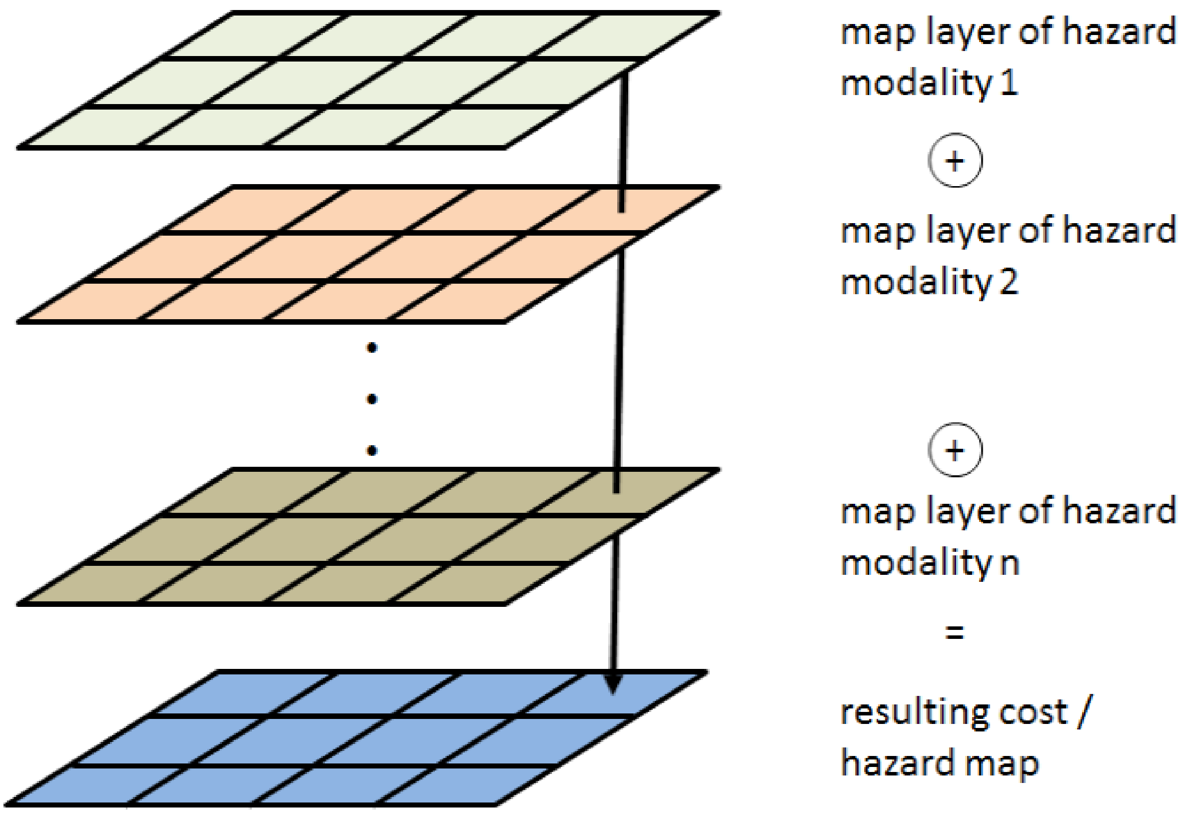

Figure1:Basic concept of generating a hazard cost maps

[1] Obviously, neighboring cells are not independent from each other; hence we will discuss the spatial dependences later in the document.

2.1 Combining the different modalities

The set of all cells $C = \{c_{x,y} \mid \forall x,y\}$ in the hazard cost map is a function of the same cells in the different modalities.

$C_{hazard} = f\,(C_{modal\_1}, C_{modal\_2}, \ldots, C_{modal\_n})$

Without loss of generality, we assume that the value in stored in cells of the modality maps is a continuous number between zero and one, $c_{x,y} = [0\ldots1]$, with 0 representing the absence of a particular hazard and 1 the highest severity of a particular hazard. We will discuss the different hazards defined in D4.2 in the context of severity/dangerousness in Sec. 3 of this document.

The hazard map values are not necessarily restricted to this interval, since they are used in the optimization problem of finding the path with the minimal traveling costs (see deliverable D5.3). The only restriction is that all costs are non-negative.

Maximum operator

This simple approach assumes that any superposition of hazards is being taken care of in the high-level hazard prediction system described in D4.2 and the worst possible hazard has been identified and quantified. Then each cell in the hazard map is just the value of the worst hazard in the modality maps.

$C_{hazard} = max\,(C_{modal\_1}, C_{modal\_2}, \ldots, C_{modal\_n})$

Weighted sum

The assumption made for the first idea can be criticized as unrealistic as it would require to handle all possible co-occurrences of hazards in the previous step, which is not the case. Using the simple assumption that the presence of two different hazards is always more dangerous than the occurrence of the more dangerous single hazard.

$C_{hazard} = w_1C_{modal\_1} + w_2C_{modal\_2} + \ldots + w_nC_{modal\_n}$

Setting $w_1 = w_2 = w_n = 1$ reduces the need for additional parameter settings, which will be chosen heuristically, since it is not feasible to estimate them from data or with experimental trials.

Probabilistic aggregation

So far, the hazard map is created as a single calculation in each time step. A more sophisticated approach would be to consider the hazards as a Markov process, taking into account the current modality maps and the previous hazard map to calculate the new hazard map.

$C_{hazard}(t) = f\,(C_{modal\_1},(t), C_{modal\_2}(t), \ldots, C_{modal\_n}(t), C_{hazard}(t\text{-}1))$

This opens up the possibility of using well-established state-estimators like Kalman filters, treating the problem as a queueing problem and modelling an arrival process using Poisson distributions, or perform event-based modelling and estimating the parameters of a Bernoulli distribution via its conjugate prior, the beta distribution Beta (which could interpreted as counting the number of hazard indications vs. the number of non-indications). A generalization to a multinomial distribution is another possible approach (e.g. for different hazard categories), but then again requires a way to reduce the different modes of the resulting probabilistic distribution to a single number indicating the level of danger.

2.2 Spatial aspects

As mentioned in Sec. 2 of this document, the suggested treatment in a cell-by-aligned-cell way disregards the neighborhood structure in the map as each cell is treated independently from every other cell. This, for obvious reasons, does not reflect reality as the degree of danger is not a discrete value that will show a high variance from cell to cell, nor is such an effect desirable in the context of the optimization to be done for the robot navigation. Hence a smoothing (to reduce variance) and extrapolation (to create a security margin around hazard areas) operation can be employed. Candidate methods include Gaussian Processes [1] and Kernel Estimators like Nadaraya-Watson [2], which then require the definition of the kernel and associated parameters, or Conditional Random Fields [REF], which we do not consider here, do to computational challenges associated with the estimation of the conditional likelihoods and the approximate interference (e.g. by loopy belief propagation).

To a certain extent, spatial aspects are taken into account when creating the hazard modalities in D4.2 or, in the case of gas distribution mapping in the respective module (here D2.7), and there to the neighborhood aspects are already encoded in the respective modality layers.

Hence it is possible to deal with the spatial aspects in different ways:

1. Ignore them on the hazard map level as they are taken care of in the hazard prediction step.
2. Perform the smoothing and extrapolation step on the hazard cost map after combining all the hazard probabilities. The advantage is that there are only few parameters to select, at the cost that you lose the ability to treat the spatial dimension for each hazard type individually (e.g. gas distribution maps already include a similar spatial extrapolation scheme, which should not necessarily extrapolated again).
3. Perform the smoothing and extrapolation step on the individual modality layer before combining them. This allows adjusting the system to the specifics of each hazard, but comes at the cost of a much higher number of parameters to choose (or estimate the hyper-parameters from the data, which is practically often not feasible).

Furthermore, the Gaussian Process (and to some extent similar information can be extracted from CRFs and the NW estimator) provides a confidence measure of the estimate. Yet it becomes a semantic question, whether this matters for creating the cost map for optimization, as it begs the question whether the unknown state of hazards at a position is itself a hazard. Thereto, this information is not immediately useful in the context of robot navigation.

2.3 Other temporal aspects

Besides the Markov assumption approaches discussed before, the system described in deliverable D4.2 uses a finite element model to estimate future hazard states by assuming that changes are governed by diffusion characteristics. As such the hazard modality map layers are available for multiple time steps including future predictions and the described danger estimate can be computed for the predicted modality layers without any changes.

2.4 Difficulties in evaluating the different options

Given those different options discussed before, it proved difficult to find a meaningful evaluation framework. It is inherently hard to define an objective evaluation criterion of what comprises a good hazard map and what comprises a bad map on its own. Attempts to judge its utility for the navigation task is not very fruitful either, as the resulting hazard maps are different in absolute numbers and differences of the cells, but the minimal path costs are in most cases the same. Engineering meaningful scenarios to showcase major differences in behavior were unsuccessful – only the obvious case of long-term exposure to elevated temperatures is an example. Yet, even this example is not realistic as this will not happen in autonomous mode, where to robots goal is getting back into communication range or safety as quickly as possible, whereas in tele operation mode in higher temperature areas the operator is informed about the robots status and makes then his decisions not based on the hazard map which is only used for autonomous navigation.

3. Revisiting the characteristics of hazards

The main challenges to transform the hazard information into a cost map are the different spatial and temporal characteristics of the different hazards as well as the level of danger they present to the robot and the mission.

3.1 Spatial characteristics of hazards

The detection of hazards is often based on in-situ measurements, where the sensor reports the conditions at the position of the sensor itself without knowledge about the spatial extension of the hazard. Yet, in reality the hazards considered are not point hazards but covering areas. Hence, an assumption about the vicinity of the measured position should be made.

In the case of remotely detected hazards either an area can be derived from the actual measurements (e.g. the smoke indication derived from the sensor fusion system) or they result in a limited area hazard (e.g. a hot spot detected by the thermal imaging system) that allows to make similar assumptions as for the in-situ sensors. Partially this is covered by the respective subsystems as e.g. the gas distribution mapping creates a spatial extrapolation already.

3.2 Temporal characteristics of hazards

We distinguish between constant and evolving hazards. Constant hazards are not expected to significantly change over time, while evolving hazards have the potential to change rather dramatically, which often increases the danger of that hazard. Following the distinction between low-level and high-level hazards presented in deliverable D4.2, most of the low-level hazards are generally assumed to not change significantly and are processed as perceived, while the high-level hazards each have a defined set of indication that signal a possible change. An exception to that characterization is the “high ambient temperature” hazard, which is also dependent on the robots internal temperature. If the robot is entering an elevated temperature area in a cool state, this is less problematic compared driving in the same temperature area for a longer period. The feasible operating temperature conditions of SmokeBot are discussed in Deliverable D7.3. In order to model this accumulating temperature hazard, the techniques discussed in Sec 2.1 have to be applied.

3.3 Degree of danger

Any hazard that could potentially damage or destroy the robot has to be avoided, yet it does not make sense to define those areas as not traversable, as a robot may find itself inside suddenly inside such an area and the navigation module should still be able to find its way to safety, even if associated with a high cost. Thereto they are marked with the highest degree of danger (DOD1). Hazards that are only interfering with the sensing capabilities form the second class, which introduce an intermediate cost to navigate through them. After all it is not a problem to cross an area with elevated temperatures, or to drive through a plume of smoke, but if possible it should be avoided nonetheless (DOD2). Hence those two degrees of dangers have to be considered. Please take not that the degree of danger of evolving hazards can change over time.

3.4 Classification of hazards

This section revisits the hazards defined in Deliverable D4.2 and attempts to classify them in terms of the characteristics presented in the previous section.

In the companion Deliverable D4.2 the following low-level hazards where defined.

Hazard Category	Hazard	Spatial and Temporal Characteristics	Degree of Danger
High temperature	Remotely detected hot spot	Spatial: Point or area, effecting the neighborhood Temporal: constant	DOD1
	High ambient temperature	Spatial: point measurement effecting the neighborhood Temporal: evolving	DOD2 to DOD1 depending on the actual temperature and the internal robot temperature
Gas-related	Remotely detected low visibility	Spatial: area calculated from the differences in coverage of the MPR and the laser scanner Temporal: constant	DOD2 the actual value can be dependent on the smoke as indicated in Square Mean Error discussed in D4.2
	Low visibility based on gas measurements	Spatial: point measurement effecting the neighborhood Temporal: constant	DOD2
	Harmful gases (harmful to the robot)	Spatial: point measurements used to calculate a gas distribution map Temporal: constant	DOD2 to DOD1 depending on the gas
Communication problem	Wi-Fi dead spot	Spatial: point measurement effecting the neighborhood Temporal: constant	DOD2 in autonomous mode, DOD1 in tele-operation mode
User defined	No-go area	Spatial: area Temporal: constant	DOD1

D4.2 defined so-called high-level hazards as well. All the high-level hazards, explosion, flashover, backdraft and rollover, present considerable danger to the robot and should be avoided. All of them are area estimates as they require the presence of multiple factors and they are all areas that present considerable danger and are therefore classified as DOD1.

4. Summary

In this deliverable, we discussed ways to transform the hazard information stored in GDIM into a cost map that is used for autonomous robot navigation. The hazard map is itself a grid map that indicates the degree of danger for each of its cells. The challenge of combining all the different hazard information into single number lies in capturing the different threat levels to robot that are presented by the different hazards. Another issue that complicates the creation of a proper hazard map is the handling of spatial and temporal aspects, since developments over time and space are used already to create the hazard prediction of D4.2. In the SmokeBot project we were unable to separate them properly and take a more principled approach that combines the hazard prediction and hazard map creation in unified framework.

A. References

[1] Rasmussen, C.E.; Williams, C.K.I (2006). „Gaussian Processes for Machine Learning." MIT Press. ISBN 0-262-18253-X.

[2] Nadaraya, E. A. (1964). "On Estimating Regression". Theory of Probability and its Applications. 9 (1): 141–142

[3] Monroy, J.G., Blanco, J.L., Gonzalez-Jimenez, J. (2016). "Time-variant gas distribution mapping with obstacle information." Autonomous Robots 40 (1), 1-16