Title	Cognitively Adequate Topological Robot Localization and Mapping
Authors(s)	Corcoran, Padraig, Bertolotto, Michela, Leonard, John J.
Publication date	2014-11
Publication information	Corcoran, Padraig, Michela Bertolotto, and John J. Leonard. “Cognitively Adequate Topological Robot Localization and Mapping.” ACM, November 2014. https://doi.org/10.1145/2676528.2676534.
Conference details	22nd ACM SIGSPATIAL International Workshop on Indoor Spatial Awareness (ISA), Dallas, Texas, USA, 4-7 November, 2014
Publisher	ACM
Item record/more information	http://hdl.handle.net/10197/6126
Publisher's version (DOI)	10.1145/2676528.2676534

Cognitively Adequate Topological Robot Localization and Mapping

Padraig Corcoran
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA.
padraigc@mit.edu

Michela Bertolotto
School of Computer Science and Informatics
University College Dublin
Ireland
michela.bertolotto@ucd.ie

John Leonard
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA.
jleonard@mit.edu

ABSTRACT

Simultaneous Localization and Mapping (SLAM) is a fundamental problem in the field of robotics which concerns mapping an environment or space while simultaneously localizing within this map. Given that one of the major goals of robotics is to perform tasks commonly performed by humans, we argue that SLAM methods should be cognitively adequate; that is, they should model the same properties of a space as the human cognition models. Topological properties are considered the most fundamental of those modelled by the human cognition. Therefore in order to achieve cognitive adequacy such properties must be modelled explicitly. Research in the domain of spatial cognition has demonstrated that the topological properties modelled by the human cognition can be quantified using the Egenhofer Nine-Intersection Model (9-IM). In this work we propose a conceptual SLAM method which models the same properties as the 9-IM. Relative to existing topological SLAM methods, which model a single topological property of connectivity between locations, this method achieves a stronger degree of cognitive adequacy.

Categories and Subject Descriptors

E.1 [**Data Structures**]: Graphs and Networks; I.5.4 [**Artificial Intelligence**]: Pattern Recognition—*applications*

General Terms

Algorithms, Theory

Keywords

Robotics, Localization, Mapping, Nine-Intersection Model

ISA '14, November 04 - 07 2014, Dallas/Fort Worth, TX, USA
Copyright 2014 ACM 978-1-4503-3137-1/14/11 ...$15.00
http://dx.doi.org/10.1145/2676528.2676534

1. INTRODUCTION

Simultaneous Localization and Mapping (SLAM) is a problem in the field of robotics which concerns modelling, or mapping, the geometry of the space within which a robot is located, while simultaneously localizing within this model [26]. In this context geometry may be defined as the properties of a space. SLAM methods can be categorized in term of *what* aspects of a space's geometry they attempt to model. Since localization is defined in terms of a space's geometry, this categorization implies what aspects of localization are modeled.

When designing a SLAM method deciding what properties of a space's geometry to model requires careful consideration. One approach is to model those set of properties which are necessary for an intended application. For example, if the intended application is navigation, then properties relating to connectivity between locations should be modelled. However the disadvantage of this approach is that for each new application a corresponding set of properties must be specified and in turn a new SLAM solution may potentially be needed to infer these properties. A better solution is to define a general set of properties which is pragmatic with respect to many applications [22]. However what constitutes a general set of properties is debatable [5]. A promising solution proposes to specify a set of properties that represents the same aspects of a space which human cognition models. This has led to the use of the term *cognitively adequate* to describe a set of properties which are believed to be an accurate model of these aspects [7]. One of the major goals of robotics is to perform tasks which humans can. Therefore we argue that SLAM methods should aim to model the geometry of a space such that these models are cognitively adequate.

In terms of what properties existing SLAM methods model, a common distinction made between methods is whether they consider the space within which a robot is located to be a metric or topological space and subsequently attempt to model properties of that space [25]. A metric space includes the concept of distance and the geometry of such a space will include properties such as the distance between objects. SLAM methods which consider the space within which a robot is located to be a metric space are known as *metric SLAM* methods and the resulting models are known as metric maps. A topological space includes the concept of neighborhood but does not include the concept of distance. The geometry of a topological space will include properties such as containment and intersection. SLAM methods which consider the space within which a robot is located to be a topological space are known as *topological SLAM* methods and the resulting models are known as topological maps. It has been confirmed by experiments both in psychology and geographical information science that the most fundamental properties of an environment which the human cognition models are topological [13]. Therefore designing SLAM methods which are cognitively adequate with respect to such properties represents an important research goal and is the focus of this paper. Previous works have focused exclusively on the development of topological SLAM methods which model a single property of a topological space. That is, connectivity between locations. In this paper we argue that such SLAM methods cannot be considered cognitively adequate. Toward addressing this issue, we propose a conceptual topological SLAM method. We demonstrate that, under certain assumptions, this method not only models connectivity between locations but achieves cognitive adequacy with respect to topological properties.

The layout of this paper is as follows. Section 2 introduces the concept of a topological space and the basic tools necessary for modelling the geometry of such a space. Section 3 reviews existing topological SLAM methods. In section 4 we introduce the concept of cognitive adequacy with respect to a model and discuss how this can be potentially defined in terms of topology concepts. Given a definition of cognitive adequacy with respect to topological properties, in section 5 we define a conceptual SLAM method which achieves cognitive adequacy. Section 6 presents simulated results which demonstrate this fact. Finally in section 7 we draw some conclusions and discuss some of the future research directions that this work presents.

2. TOPOLOGY OF EUCLIDEAN PLANE AND GRAPH

In this work we are principally concerned with the 2D topological SLAM problem which considers the space within which a robot is located to be the Euclidean plane and in turn considers this plane to be a topological space. In this section we briefly introduce two branches of topology known as *point set topology* and *topological graph theory* which concern modeling the topology of sets and a graph respectively. For greater details please see [20] and [12] respectively.

2.1 Point Set Topology

The topology of the Euclidean plane is known as the *Euclidean* or *usual topology*. An intuitive understanding of this topology can be gained by imagining the Euclidean plane within which the environment is embedded to be a rubber sheet. Now imagine this sheet is arbitrarily stretched but not torn or folded. The transformation induced by such stretching is called a *topological transformation* or *homeomorphism*. If the space in question is considered a topological space, the properties which are invariant under an homeomorphism represent its geometry. We refer to such properties as being *topological properties* of the space [11]. For example, if a point is contained inside a region, this

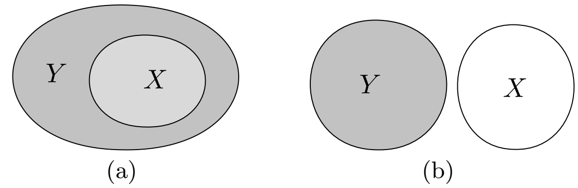

Figure 1: The topological relationships of *inside* and *disjoint* between the sets X and Y are illustrated in (a) and (b) respectively.

property of containment is invariant under any homeomorphism and therefore is a topological property of the space. On the other hand, the area of a region is not invariant under any homeomorphism and therefore is not a topological property of the space. Instead area is a property which is invariant under *Euclidean transformations* and represents the geometry of a metric space. As such, it is referred to as a *metric property.* If two spaces are related by a homeomorphism they have equal topological properties and are called *topologically equivalent.* All spaces which are topologically equivalent belong to a single *topological equivalence class.*

Using point set topology it is possible to express topological properties of the Euclidean plane in terms of the intersection or non-intersection of the boundaries, interiors and exteriors of sets in the plane. Let ∂A, A^o and A' be the boundary, interior and exterior respectively of a set A. Consider the Euclidean plane which contains two sets X and Y illustrated in Fig. 1(a). The topological relationship existing between these sets can be described linguistically as X *inside* Y and, in terms of point set topology, satisfy the following four conditions: $\partial X \cap \partial Y = \varnothing$, $X^o \cap Y^o \neq \varnothing$, $\partial X \cap Y^o \neq \varnothing$ and $X^o \cap \partial Y = \varnothing$. Likewise the topological relationship existing between the two sets in Fig. 1(b) can described linguistically as X *disjoint* Y and, in terms of point set topology, satisfy the following four conditions: $\partial X \cap \partial Y = \varnothing$, $X^o \cap Y^o = \varnothing$, $\partial X \cap Y^o = \varnothing$ and $X^o \cap \partial Y = \varnothing$.

Any metric space has a corresponding topology. This topology is commonly referred to as the topology induced by the metric in question. Consequently, for a given metric space the set of metric properties corresponding to that space is a superset of the topological properties corresponding to the induced topological space. For example containment is both a metric and topological property due to the fact that it is invariant under any Euclidean and homeomorphism transformation respectively. By similar reasoning, area is a metric property but not a topological property.

2.2 Topological Graph Theory

In this work we adopt the following topological definition, as opposed to a combinatorial definition, of a graph [12]. Here the graph consists of a set of points V, called *vertices*, and a set of arcs E, called *edges*, joining pairs of vertices such that each edge is homeomorphic to the closed interval $[0, 1]$. The embedding of a graph in a surface is defined as a continuous one-to-one function from the graph into the surface [16]. Intuitively this can be thought of as a drawing of the graph on the surface such that no edges intersect

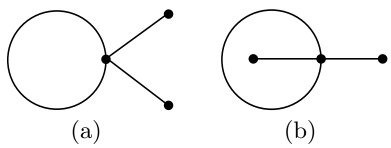

Figure 2: Three planar embeddings of a graph are represented. Vertices are represented by solid circles and edges by solid lines.

except at vertices. Every graph can be embedded in a given surface if and only if that surface has a sufficient number of handles or bridges [12]. A planar graph is a graph that can be embedded in a plane, that is, a surface with genus zero.

If a connected graph is embedded in a plane and cuts the plane along the graph what remains is a set of *faces* F where each is homeomorphic to an open disk. The degree of a face is the number of edges incident to that face. If both 'banks' of an edge belong to the same face, that edge is incident to the face twice. For example, consider the embedding of a graph in the plane represented in Fig. 2(a). This embedding contains two faces; a face of degree one and an *outer* face of degree five.

A connected graph may have multiple embeddings in a surface where these embeddings are not topologically equivalent. For example consider Fig. 2(a) and Fig. 2(b) which represent two different embeddings of the same graph in a plane. One can confirm this by noting that the set of vertices, the sets of edges, the edge incidence and vertex adjacency relations are equal [16]. However these embeddings are not topologically equivalent; that is, a homeomorphism does not exist between them. Next consider Fig. 2(a) and Fig. 2(c) which also represent two different embeddings of the same graph in a plane. However in this case both embeddings are topological equivalent.

The topological equivalence class of 2-connected graphs embedded on a sphere has a corresponding combinatorial representation, known as a *rotation system*, which is defined by the clockwise ordering of adjacent edges around each vertex in the graph (please see Theorem 3.2.2 of [12]). The topological equivalence class of 2-connected graphs embedded in a plane has a corresponding combinatorial representation which is defined by the rotation system of the graph and its outer face. Schnyder [23] proposed an algorithm which given the rotation system and the outer face of a graph embedded in the plane, returns a graph embedded in the plane which belongs to the topological equivalence class as the original graph. This algorithm has a time complexity of $O(n)$. For example, consider the graph embedded in the plane displayed in Fig. 3(a). The rotation system and outer face of this graph were extracted and subsequently supplied as parameters to the algorithm of Schnyder. The resulting graph embedded in the plane is displayed in Fig. 3(b). It is evident that both embedded graphs belong to the same topological equivalence class.

3. TOPOLOGICAL SLAM

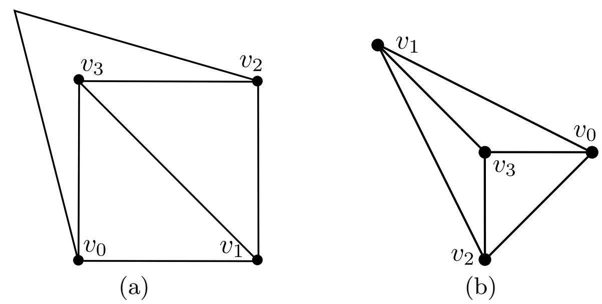

Figure 3: An embedded graph and a corresponding embedded graph belonging to the same topological equivalence class are displayed in (a) and (b) respectively.

As discussed in section 2.1, metric properties are a superset of topological properties. Therefore a metric map contains both metric and topological properties. Apart from not containing any metric properties, a topological map differs from a metric map in the way it represents topological properties. In a metric map, topological properties are generally represented implicitly while in a topological map they are represented explicitly. As will be discussed in section 4, topological properties are extremely important in the context of performing many tasks. Therefore it important to represent such properties explicitly; therein lies the motivation for developing and the goal of topological SLAM methods.

Current topological SLAM methods assume that the space in question consists of a series of paths and can be accurately modeled using a graph. Vertices in such graphs represent points in the Euclidean plane while the existence of an edge between two vertices signifies that the points in question are connected by a direct path. The embedding of the graph in the plane is not modelled and in fact only a single topological property of connectivity between points is modelled. We refer to these methods as *graph-based topological SLAM* methods. Graph-based topological SLAM was originally proposed by Kuipers [15] and it has become a standard paradigm adapted by all other researchers in the field. In the remainder of this section we review existing graph-based topological SLAM methods.

A number of graph-based topological SLAM methods use metric SLAM methods as a platform. These methods are based on the principle that, every metric space has an induced topology and therefore given a metric map it is possible to derive a topological map through a process of abstraction or generalization. We refer to methods based upon this principle as *indirect* methods. The method by [2] first constructs a detailed metric map and from this extracts a corresponding graph based representation. In this representation vertices correspond to convex regions in the metric map while an edge exists between any pair of vertices if the corresponding regions in the metric map touch. Localization is achieved by constructing a local metric map and matching this against the metric map for the entire environment. The method of [3] models the environment using a generalized Voronoi graph (GVG) which can be considered a skeleton like graph of the free space. The authors construct the GVG by analysing successive local metric maps. Localization is achieved by storing a number of metric and topological properties at each vertex which can be matched against.

The drawback of the above indirect approaches is that they compute metric properties which are ultimately abstracted away. A better solution is to only compute the necessary topological properties. We refer to methods based upon this principle as *direct* methods. Appearance based SLAM methods can be considered a direct approach to graph-based topological SLAM. In such methods correspondences between points in a robot's trajectory are determined using appearance features such as visual SIFT features [8]. Edges are constructed between all pairs of vertices successively visited by the robot trajectory. No metric properties of the space are computed. The advantages of using a direct method, as opposed to an indirect method, when performing graph-based topological SLAM is reflected in the scale at which the corresponding methods can perform. The appearance based methods of [8] can perform accurately within environments of a scale greater than one thousand kilometres. There currently exists no metric SLAM method which can perform accurately at such a scale.

A number of methods exist which combine metric and topological SLAM methods [24]. Similar to graph based topological SLAM these methods model the geometry of the space within which a robot is located using a graph representation. However this graph is used to model both metric and topological properties. Vertices contain local metric maps while edges not only represent connectivity but may also represent metric relationships between local maps [1].

4. COGNITIVE ADEQUACY

When designing a SLAM method deciding what properties of a space's geometry to model requires careful consideration. A set of properties are *complete* if they are necessary and sufficient for determining if the geometry of two spaces are equivalent [4]. On the other hand, a set of properties are *incomplete* if they are necessary but not sufficient for determining if the geometry of two spaces are equivalent. Although it is possible to model the geometry of a space using a complete set of properties, the use of an incomplete set is generally more appropriate for a number of reasons. Relative to a complete set of properties, an incomplete set of properties contains less information about the geometry of the space in question. Therefore in the context of robotics, where all information must be inferred from noisy sensor measurements, modelling using an incomplete set represents an easier problem. Also, the information contained in a complete set may be of too fine a resolution than is appropriate for performing a given task. For example consider Fig. 4(a) and 4(b) where each figure displays two sets X and Y embedded in the Euclidean plane. Irrespective of whether we consider the Euclidean plane to be a metric or topological space, the geometry of both figures is not equivalent. For a given task it may not be appropriate to use a set of properties which distinguishes between the geometry of these figures. An example of a set which does not perform such a distinction is a single topological property which indicates if the sets X and Y intersect.

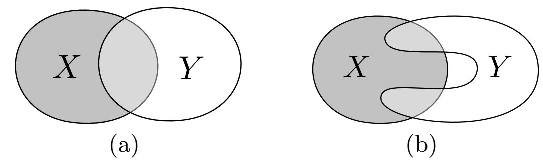

Figure 4: The geometry of (a) and (b) are not equal.

Cohn [6] states that not all sets of properties are equally useful and the actual set used must be relevant to the task being performed. As discussed in the introduction to this paper, it is advantageous to model a set of properties which is pragmatic to many applications. A promising solution towards defining such a set is to model the properties of a space's geometry which human cognition models or conceptualizes. This has led to the use of the term *cognitively adequate* to describe such a set [14]. The degree of cognitively adequacy can vary from *strong* (idealized) to *weak* (conforming to well known ergonomic standards) [14].

What constitutes a cognitively adequate set of properties for the Euclidean plane has been the focus of much research. Toward answering this question Egenhofer [10] defined the Nine-Intersection Model (9-IM). The 9-IM models the geometry of the Euclidean plane by modelling all pairwise relationships between sets embedded in the plane. Each of these pairwise relationships is in turn represented by a set of nine features indicating the intersection or non-intersection of the interiors, boundaries and exteriors of the sets in question. Specifically, for two sets X and Y their relationship is modelled using the following set of nine features:

Table 1: Features in the 9-IM.

$\partial X \cap \partial Y$	$\partial X \cap Y^o$	$\partial X \cap Y'$
$X^o \cap \partial Y$	$X^o \cap Y^o$	$X^o \cap Y'$
$X' \cap \partial Y$	$X' \cap Y^o$	$X' \cap Y'$

All of these nine features are defined in terms of point set topology constructs and therefore the 9-IM models topological properties of the Euclidean plane. Each of these features evaluates to intersection or non-intersection and therefore the 9-IM can distinguish between 512 (2^9) possible topological relationships where each relationship corresponds to a particular topological property. However most of the 512 relationships cannot be realized for connected sets embedded in the Euclidean plane. For a point X and region Y, 3 relationships can be realized. These realizations along with corresponding linguistic representations are illustrated in Fig. 5. For two simply-connected regions X and Y, 8 relationships can be realized. These realizations along with corresponding linguistic representations are illustrated in Fig. 6. It is straight forward to derive the corresponding 9-IM representation of these relationships.

Following the proposal of the 9-IM, a number of studies were performed to determine if, and if so to what degree, does the 9-IM constitute a cognitively adequate set of properties for modelling the geometry of the Euclidean plane. In the case of a metric space containing two connected sets and a single pairwise relationship, through human subjected test-

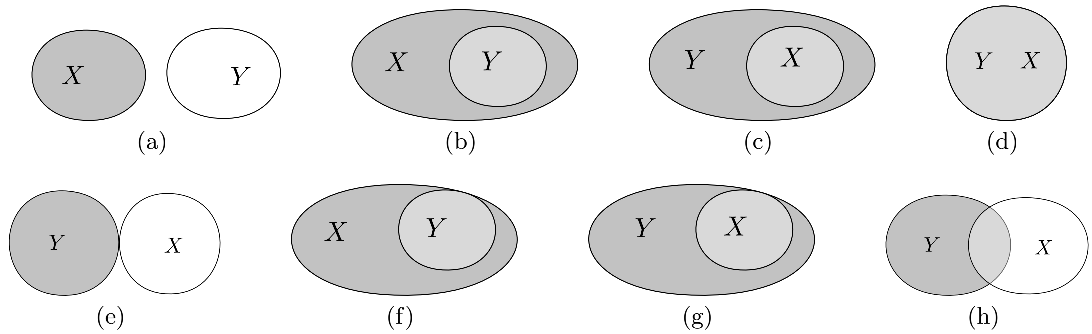

Figure 6: For two regions X and Y embedded in the Euclidean plane, the topological relationships of X *disjoint* Y, X *contains* Y, X *inside* Y, X *equal* Y, X *meet* Y, X *covers* Y, X *covered by* Y and X *overlap* Y, are illustrated in (a), (b), (c), (d), (e), (f), (g), and (h) respectively.

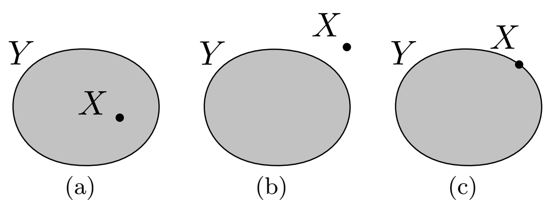

Figure 5: For a point X and a region Y embedded in the Euclidean plane, the topological relationships of Y *contains* X, X *disjoint* Y and X *meet* Y are illustrated in (a), (b) and (c) respectively.

ing, Mark [17] argued the 9-IM exhibits strong cognitive adequacy. This claim was subsequently promoted in later studies.

One of the major goals of robotics research is to ultimately performs tasks which humans can. Therefore we argue that SLAM methods should aim to model the geometry of a space such that the resulting models are cognitively adequate. We refer to any SLAM method which achieves this goal as being a *Cognitively Adequate SLAM* method. Given the above discussion, the importance of modelling topological properties toward achieving this goal is evident. In this work we assume the 9-IM to be cognitively adequate with respect to topological properties. In this context, any SLAM method which models the same properties at the 9-IM can be considered cognitively adequate with respect to topological properties. It is not difficult to see that the topological properties modelled by the 9-IM would be useful in the context of many applications. For example, consider the case where a set X represents a room and a set Y represents a building. In the context of a localization task it would be useful to be able to determine if the room represented by X is *inside* the building represented by Y.

As discussed in section 2.1, metric properties are a superset of topological properties. However, metric SLAM methods generally model topological properties implicitly where these properties must be inferred through a process of abstraction. Therefore they cannot be considered cognitively adequate with respect to topological properties.

5. EXTENDING TOPOLOGICAL SLAM

In section 3 we discussed the fact that topological SLAM is generally defined as the problem of modelling the geometry of the Euclidean plane using a graph while simultaneously localizing within this model. Such a model can be achieved if one can effectively detect loop closures. A loop closure occurs when the robot is exploring an environment and returns to a previously visited location. However the drawback of using such a model is that it only models a single topological property; that is connectivity between points. Therefore such SLAM methods cannot be considered cognitively adequate with respect to topological properties given the definition of this term in the previous section. For example if the Euclidean plane contains two sets corresponding to a region and a point, in order to achieve cognitively adequacy one must be capable of modelling the relationships of *contains*, *disjoint* and *meet* between these sets. Here the region in question may correspond to that enclosed by a cycle in the graph while the point may correspond to a graph vertex.

In this section we propose a conceptual SLAM methodology which achieves cognitive adequacy with respect to topological properties. This method makes three assumptions. Firstly we assume the space within which the robot is located is a graph embedded in the Euclidean plane. This is a reasonable assumption if the space in question corresponds to a street network or a building interior; in fact this assumption is made in most existing topological SLAM methods. Secondly, we assume that one can detect locations corresponding to vertices in the graph and perform loop closures; that is, detect when one has returned to a previously visited vertex. Finally we assume that one has an accurate orientation sensor, such as an gyroscope or heading sensor.

Before modelling the topological properties of the Euclidean plane we must define sets in the plane to model such properties with respect to. All existing works which aim to achieve cognitive adequacy, and were discussed in section 4, assume the sets are known a priori. However in the context of SLAM this is not the case and the sets must be extracted from sensor data. There have been some works which have proposed methods for extracting sets toward modelling topological properties. Derenick et al. [9] defined sets to be regions of space with uniform landmark visibility. Remolina et al. [21] defined sets to be regions to the left and right of different paths through the space. However there exists no standard approach to extracting suitable sets and the approach used generally depends on the application or context. In this work we define the sets in the space to be points corresponding to the graph vertices, lines corresponding to graph edges and regions corresponding to the interior of simple cycles in the graph.

The proposed topological SLAM methods is performed in two steps. In the first step we compute a graph and its embedding in the Euclidean plane such that this graph belongs to the same topological equivalence class as the space being modelled. Two graphs belonging to the same topological equivalence class exhibit equal topological properties. In the second step we exploit this fact by computing all properties necessary, to achieve cognitive adequacy, from the embedded graph computed in the first step. Both these steps are described in detail in the following two subsections.

5.1 Computing topological equivalence class

The aim of this step is to compute a graph and its embedding in the Euclidean plane such that this graph belongs to the same topological equivalence class as the space being modelled. We propose to compute this using the algorithm of Schnyder [23] introduced in section 2.2. However in order to apply this algorithm we must first compute the following three properties of the graph which serve as algorithm parameters: 1) the set of graph vertices and edges, 2) the rotation system of the graph (see section 2.2), and 3) the outer face of the graph. We now describe in turn how each of these is computed.

5.1.1 Vertices and edges

The set of vertices and edges is computed as follows. If the robot encounters a previous unseen vertex it is added to the set of vertices. If the robot visits vertex x followed by vertex y an edge is added between these vertices. The computational time complexity of this step depends on the exploration strategy the robots uses. If an efficient exploration strategy, such as depth first search, is used the time complexity will be $O(|V| + |E|)$.

5.1.2 Rotation system

Let $O(v_i, v_j)$ be the orientation of the edge (v_i, v_j) relative to the vertex v_i. In this work we define $O(v_i, v_j)$ to the orientation measurement in the edge (v_i, v_j) which is next to v_i. For example consider the vertex v_1 of the graph in Fig. 3(a). The orientation of the three edges incident to this vertex, that is $O(v_1, v_0)$, $O(v_1, v_3)$ and $O(v_1, v_2)$, are π, $^{3\pi}/_4$ and $^{\pi}/_2$ respectively. Given the orientations of all edges incident to a given vertex, a sorting of these orientation values is then performed to give the required clockwise ordering. The computational time complexity of this step is $O(|V|)$.

5.1.3 Outer face

The outer face of a connected planar graph embedded in the plane corresponds to a simple cycle $C = (v_j, \ldots, v_{i-1}, v_i, v_{i+1}, \ldots, v_j)$ with the following topological property. If the cycle in question is clockwise the orientation of all edges incident to each vertex v_i in the cycle lie in the closed interval $[O(v_i, v_{i-1}), O(v_i, v_{i+1})]$. On the other hand, if the cycle in question is counter-clockwise the orientation of all edges incident to each vertex v_i in the cycle lie in the closed interval $[O(v_i, v_{i+1}), O(v_i, v_{i-1})]$. Fig. 7 illustrates this property for a single vertex v_i contained in a counter-clockwise cycle. It is evident that if the cycle in question corresponds to the outer face, all edges incident to v_i must have an orientation lying in the closed interval $[O(v_i, v_{i+1}), O(v_i, v_{i-1})]$.

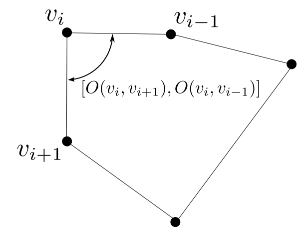

Figure 7: A counter-clockwise cycle containing the vertex v_i.

Using the above topological property we determine the outer face using the following steps. First we extract all simple cycles in the graph using a search procedure with backtracking [18]. The computational time complexity of this step is $O(|V||E|)$. Next for each cycle we determine whether it corresponds to a clockwise or counter-clockwise cycle by examining its turning number which equals the total curvature of the cycle divided by 2π. A turning number of $+1$ corresponds to a counter-clockwise cycle while a turning number of -1 corresponds to a clockwise cycle. We approximate the total curvature of a cycle as the sum of its turning angles which corresponds to the differences in consecutive orientation measurements along the cycle [19]. Finally, by jointly considering the orientation of each cycle and the orientation of all edges incident to vertices in that cycle, the outer face can be determined. The computational time complexity of this step is $O(|E|)$.

Given the necessary Schnyder algorithm parameters computed above, one can now compute a graph and its embedding in the Euclidean plane such that this graph belongs to the same topological equivalence class as the space being modeled. For example consider the graph embedded in the plane displayed in Fig. 8(a). The rotation system and outer face of this graph were extracted and subsequently supplied as parameters to the algorithm of Schnyder. The resulting graph embedded in the plane is displayed in Fig. 8(b). It is evident that both embedded graphs belong to the same topological equivalence class.

5.2 Computing topological properties

In order to achieve cognitive adequacy it is necessary to compute the nine point set topology features contained in the 9-IM model of table 1. This will in turn allow us to compute the topological relationship which exists between the sets in question. In the previous section we computed a graph embedded in the plane belonging to the same topological equivalence class as the space being modelled. Therefore computing the above point set topology features is reduced

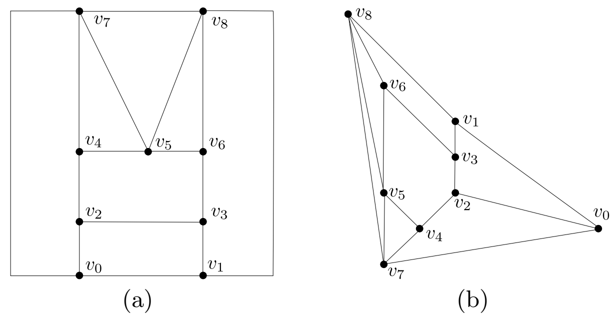

Figure 8: An embedded graph and a corresponding embedded graph belonging to the same topological equivalence class are displayed in (a) and (b) respectively.

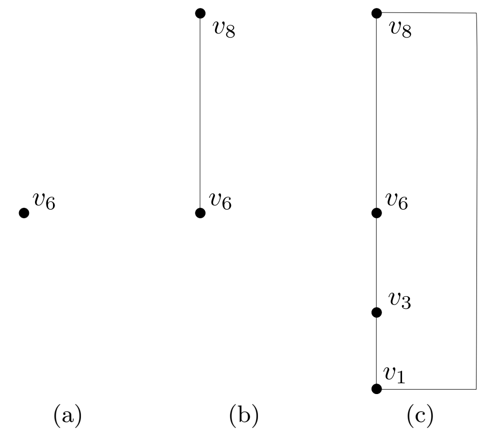

Figure 9: The exploration of the simulated graph of Fig. 8(a) is illustrated.

to a straight forward application of computational geometry techniques to the above graph. To illustrate this point we now demonstrate how to compute two of these features.

As a first example consider the case of computing the feature $X^o \cap Y^o$ where the sets X and Y are both of dimension two; that is, they correspond to regions defined by cycles in the graph. In the context of computational geometry, these sets correspond to polygons and determining if two polygons intersect is a standard technique in this domain. As a second example consider the case of computing the feature $\partial X \cap Y^o$ where the sets X and Y are of dimensions zero and two respectively; that is, they correspond to a vertex and a region defined by a cycle in the graph respectively. In the context of computational geometry, these sets correspond to a point and a polygon respectively. Determining if a point intersects a polygon is a standard technique in this domain and is commonly know as a *point in polygon* query.

Given two sets, their respective dimension and the corresponding set of nine point set topological features of Table 1, the 9-IM defines the topological relationship which exists between the sets in question. For example if both sets are of dimension two, eight unique relationships can exist; these are illustrated in Fig. 5. Consider Fig. 8(a) and the regions defined by the cycles $(v_0, v_1, v_8, v_7, v_0)$ and $(v_2, v_3, v_6, v_5, v_4, v_2)$ which we entitle X and Y respectively. The 9-IM defines the relationship between these sets to be X *contains* Y.

6. RESULTS

In order to demonstrate the proposed SLAM method we created a dataset containing six simulated graphs embedded in the Euclidean plane. The minimum and maximum number of vertices in each graph was four and ten respectively. The minimum and maximum number of edges in each graph was six and sixteen respectively. Fig. 3(a) and Fig. 8(a) display two of the embedded graphs contained in the dataset. The graphs were scaled such that the x and y locations of their vertices lay in the range $[0, 1000]$. The process of a robot exploring the embedded graphs was simulated with control inputs being manually specified. The exploration process was terminated when all vertices were visited and all edges were traversed. An accurate orientation sensor was assumed and this was sampled every unit distance travelled.

To illustrate the exploration process consider the simulated graph of Fig. 8(a). The robot begins its exploration at the vertex v_6 and adds this vertex to its current map of the environment. This map is displayed in Fig. 9(a). The robot next traverses the adjacent edge to vertex v_8. It subsequently adds this vertex and edge to its current map which is displayed in Fig. 9(b). The robot next visits the vertices v_1, followed by v_3, followed by v_6. Upon visiting vertices v_1 and v_3 the robot adds the corresponding vertices and edges to its current map. When the robots visits the vertex v_6 it detects a loop closure and therefore does not add a new vertex to its current map. Instead it adds an edge between the vertices v_1 and v_6 in its current map which is displayed in Fig. 9(c). This process of exploration continues until the entire graph has been explored.

For all simulated graphs in the dataset, the correct number of vertices and the connectivity between these vertices was correctly determined using the above exploration procedure. For each graph the corresponding rotation system and outer face was also correctly determined. In all cases this subsequently allowed the generation of an embedded graph belonging to the same topological equivalence class as the graph being modelled. Fig. 3 and Fig. 8 illustrate this fact for two of the embedded graphs contained in the dataset.

For each of the computed embedded graphs the correct topological relationships, according to the 9-IM, between sets in the space were identified. For example in the context of Fig. 8 it was correctly identified that the topological relationship existing between the set X, corresponding to the region defined by the cycle $(v_0, v_1, v_3, v_2, v_0)$, and the set Y, corresponding to the region defined by the cycle $(v_4, v_5, v_6, v_8, v_7, v_4)$, is X *disjoint* Y. In the context of the same embedded graph, it was also correctly identified that the topological relationship existing between the set X, corresponding to the region defined by the cycle $(v_0, v_1, v_8, v_7, v_0)$, and the set Y, corresponding to the point defined by the vertex v_3, is X *contains* Y. The above results demonstrate that the proposed SLAM method achieves cognitive adequacy with respect to topological properties.

7. CONCLUSIONS

This article makes the argument that SLAM methods should aim to achieve cognitive adequacy through modelling properties of a space which the human cognition models. The most fundamental of these are topological and therefore achieving cognitive adequacy with respect to such properties represents an important research goal. Drawing from research in spatial cognition, we quantify such properties in terms of point set topology. Specifically, we consider the Nine-Intersection Model of Egenhofer. This analysis has identified that current topological SLAM methods only model very limited topological properties; specifically, a single topological property of connectivity between locations. Modelling a richer set of topological properties, and in-turn potentially achieving cognitive adequacy, has many potential robotic applications. Such applications include localization and object manipulation.

ACKNOWLEDGMENT

This research was funded in part by a European Marie Curie International Outgoing Fellowship.

8. REFERENCES

[1] M. Bosse, P. Newman, J. Leonard, and S. Teller. Simultaneous localization and map building in large-scale cyclic environments using the atlas framework. *The International Journal of Robotics Research*, 23(12):1113–1139, 2004.

[2] J. Choi, M. Choi, S. Y. Nam, and W. K. Chung. Autonomous topological modeling of a home environment and topological localization using a sonar grid map. *Autonomous Robots*, 30(4):351–368, 2011.

[3] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (slam): toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17(2):125–137, 2001.

[4] E. Clementini and P. Di Felice. Topological invariants for lines. *IEEE Transactions on Knowledge and Data Engineering*, 10(1):38–54, 1998.

[5] E. Clementini, P. Felice, and P. Oosterom. A small set of formal topological relationships suitable for end-user interaction. In D. Abel and B. Chin Ooi, editors, *Advances in Spatial Databases*, volume 692 of *Lecture Notes in Computer Science*, pages 277–295. Springer Berlin Heidelberg, 1993.

[6] A. G. Cohn and S. M. Hazarika. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae*, 46(1):1–29, 2001.

[7] P. Corcoran, P. Mooney, and M. Bertolotto. Spatial relations using high level concepts. *ISPRS International Journal of Geo-Information*, 1(3):333–350, 2012.

[8] M. Cummins and P. Newman. Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research*, 27(6):647–665, 2008.

[9] J. Derenick, A. Speranzon, and R. Ghrist. Homological sensing for mobile robot localization. In *IEEE international conference on robotics and automation*, pages 572–579, 2013.

[10] M. J. Egenhofer. Reasoning about binary topological relations. In O. Gunther and H.-J. Schek, editors, *Advances in Spatial Databases*, Lecture Notes in Computer Science, pages 141–160. Springer Berlin Heidelberg, 1991.

[11] M. J. Egenhofer and R. D. Franzosa. Point-set topological spatial relations. *International Journal of Geographical Information System*, 5(2):161–174, 1991.

[12] J. L. Gross and T. W. Tucker. *Topological graph theory*. Courier Dover Publications, 2001.

[13] A. Klippel. Spatial information theory meets spatial thinking: Is topology the rosetta stone of spatio-temporal cognition? *Annals of the Association of American Geographers*, 102(6):1310–1328, 2012.

[14] M. Knauff, R. Rauh, and J. Renz. A cognitive assessment of topological spatial relations: Results from an empirical investigation. In *Spatial Information Theory A Theoretical Basis for GIS*, pages 193–206. Springer, 1997.

[15] B. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and autonomous systems*, 8(1):47–63, 1991.

[16] S. K. Lando and A. K. Zvonkin. *Graphs on surfaces and their applications*, volume 141. Springer, 2004.

[17] D. M. Mark and M. J. Egenhofer. Modeling spatial relations between lines and regions: combining formal mathematical models and human subjects testing. *Cartography and geographic information systems*, 21(4):195–212, 1994.

[18] P. Mateti and N. Deo. On algorithms for enumerating all circuits of a graph. *SIAM Journal on Computing*, 5(1):90–99, 1976.

[19] M. E. Mortenson. *Mathematics for computer graphics applications*. Industrial Press Inc., 1999.

[20] J. R. Munkres. *Topology*, volume 2. Prentice Hall Upper Saddle River, 2000.

[21] E. Remolina and B. Kuipers. Towards a general theory of topological maps. *Artificial Intelligence*, 152(1):47–104, 2004.

[22] J. Renz, R. Rauh, and M. Knauff. Towards cognitive adequacy of topological spatial relations. In C. Freksa, C. Habel, W. Brauer, and K. Wender, editors, *Spatial Cognition II*, volume 1849 of *Lecture Notes in Computer Science*, pages 184–197. Springer Berlin Heidelberg, 2000.

[23] W. Schnyder. Embedding planar graphs on the grid. In *Proceedings of the first annual ACM-SIAM symposium on Discrete algorithms*, pages 138–148. Society for Industrial and Applied Mathematics, 1990.

[24] S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.

[25] S. Thrun. *Exploring Artificial Intelligence in the New Millennium*, chapter Robotic mapping: a survey, pages 1–35. Morgan Kaufmann San Mateo, CA, 2002.

[26] S. Thrun, W. Burgard, and D. Fox. *Probabilistic robotics*. MIT press Cambridge, 2005.