

Probabilistic Region Connection Calculus

Codruta Girlea and Eyal Amir¹

Abstract. We present a novel probabilistic model and specification language for spatial relations. Qualitative spatial logics such as RCC are used for representation and reasoning about physical entities. Our probabilistic RCC semantics enables a more expressive representation of spatial relations. We observe that reasoning in this new framework can be hard. We address this difficulty by using a factored representation based on Markov Random Fields.

We formally present the syntax and semantics of a probabilistic RCC. We then use Markov Random Fields to represent our models compactly. Using this representation, we show a basic algorithm for answering queries about the probability of a relation to hold between two entities. Finally, we illustrate the effectiveness of the new approach experimentally over a small set of examples.

1 Introduction

We provide a logic for representing and reasoning about spatial elements, in the presence of uncertainty. Our framework combines a high-level approach based on qualitative spatial reasoning, that avoids the pitfalls and complexities of pixel-level reasoning, with a probabilistic semantics, able to deal with and quantify uncertainty.

Reasoning about space at the pixel level requires too complex computations and does not capture higher-level properties of objects. As a solution, higher-level qualitative calculi have been introduced, such as Region Connection Calculus, or RCC [7]; however, in such calculi there are no shades of gray in representing uncertainty. We take the flexible, high-level approach of qualitative spatial reasoning, RCC-8 in particular, and define probabilistic models.

Using our probabilistic spatial calculus, we are able to answer more accurately questions about the relations between regions: in classic RCC, uncertainty with respect to the base relation that holds between two regions means that some base relations are possible. There is no cue as to which of these relation is more likely. In the worst case, the entire base relation is possible. However, generally, in real world situations, some relations might be more probable than others; using our probabilistic calculus, one can find the probabilities for all the base relations between the two regions and then get, rather than a set of relations, the most probable base relation.

An example of an application for our calculus is recreating a spatial landscape, consisting of all spatial relations that hold between all entities, from a natural language description. The landscape description can be analysed to extract an initial set of spatial relations as the first, incomplete, landscape, and then the most likely complete image can be recreated using inference in probabilistic RCC. The techniques used here could be extended to other spatial formalisms, that are able to capture other meaningful relations between entities. Reconstructing a spatial landscape from text can be useful to answering

deeper understanding queries regarding the text. This kind of queries can nowadays be answered in the context of natural language processing by means of textual entailment [9]. Here, either one uses only lexical cues, which can only lead to a shallow understanding of the text, or one learns to infer deeper, semantic relations implied by the text by training on large corpora of annotated textual entailment pairs. In the latter case, much effort is spent on annotating a corpus and feature engineering. By using qualitative spatial reasoning, one only needs to spend effort in extracting the obvious spatial relations from the text, whereas the deeper understanding queries can be answered by reasoning in the underlying spatial logic.

The paper is structured as follows: first, we present some background notions on RCC. Then, we describe the syntax, semantics and inference for our calculus. Next, we present the MRF-based representation and inference. We then show the results on some examples. Finally, we give an overview of related work and conclude.

2 Background

Qualitative Spatial Reasoning [3] is a term used for any relational reasoning technique for which the objects are spatial entities.

Region Connection Calculus (RCC), introduced by Randell, Cui and Cohn in 1992 [7], is a qualitative spatial calculus used to reason about the relations between regions. The distinction between base relations is made based on either connectedness or the mereological 'part of' relation. The two definitions are equivalent, as the two relations can be defined by means of each other. Given the possible distinctions and additional information considered (e.g., whether the region borders are taken into account or not), the space of possible relations is broken into a set of jointly exhaustive and pairwise disjoint, or JEPD, base relations.

For RCC-8, the base relations are:

- disconnected (DC) - the regions are not connected, i.e. they share no common parts;
- externally connected (EC) - the regions are connected, but their interiors are not;
- partially overlap (PO) - the regions' interiors are connected, but there are regions that are part of either one but not the other
- tangential proper part and its inverse (TPP, TPPI) - one region is a part of the other, but that region is not part of the other's interior (equivalently anything that connects to the inner region also connects to the outer region);
- non-tangential proper part and its inverse (NTPP, NTPPI) - one region is a part of the other one's interior;
- equivalent (EQ) - each region is part of the other.

RCC-8 can be formalized as a relation algebra in the sense of Tarski based on the set algebra over 2^B , where B is $\{DC, EC, PO, EQ, TPP, NTPP, TPPI, NTPPI\}$ (the set of base relations).

¹ University of Illinois at Urbana-Champaign, USA, email: girlea2@illinois.edu

To complete the relation algebra, each relation has a converse (TPP is the converse of TPPI, NTPP is the converse of NTPP, the rest of the base relations are each its own converse), and the composition table for base relations is as shown by Wöfl et al. [12]. EQ is the composition identity.

For RCC-5, the border information is not considered, so EC and PO are coalesced into O (overlap), TPP and NTPP are collapsed into PP (proper part), and analogously for their converse relations. Consequently B is $\{DC, O, EQ, PP, PPI\}$ and the relation algebra is changed accordingly.

3 Probabilistic RCC

Let us consider the following image description:

John's office is on the second floor of the building. Andy's office is across the corridor, right next to the service room. There's a tree right beside Andy's office window. Andy is standing by the window. John realized a couple of minutes ago he needed something from the service room, and thought he'd pass by Andy's office on his way there to exchange a few words.

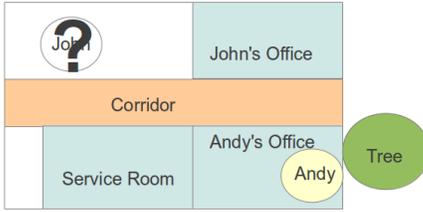


Figure 1. Example of an image

When one reads this description, one builds a mental abstract image (e.g. Figure 1), consisting of spatial relations between entities, and based on this particular image, one can answer questions on what the most likely relative positions of the entities in this world are. Each of these abstract images is similar to a probabilistic RCC model.

In general, in the problem we are trying to solve, we are given a set of regions in a topology, a set of region names or region constants, and a set of spatial constraints on them expressed as a formula. We want to be able to answer queries regarding the probability of certain relations to hold between certain pairs of regions.

3.1 Syntax

In general, the signature of probabilistic RCC is a first-order logic signature of a particular form, containing: a set of constants \mathcal{C} (the region names); and a set of arity 2 relations \mathcal{B} (the base relations).

For RCC-8, $\mathcal{B} = \{DC, EC, PO, EQ, TPP, NTPP, TPPI, NTPPI\}$ and for RCC-5, $\mathcal{B} = \{DC, PO, EQ, PP, PPI\}$.

Two probabilistic RCC-8 (RCC-5) signatures may differ from each other on their set of constants. This leads to the following definition of an RCC-8 signature:

Definition 1 A probabilistic RCC signature is a set of region constants \mathcal{C} .

In the story described above, the signature contains the constants: Andy, John, the corridor, the service room, the tree, and the offices.

Henceforth we will refer to RCC-8 only; the results can easily be applied to RCC-5.

A basic sentence encodes the set of constraints for the problem; this is just a ground FOL sentence. Our example can be encoded as the basic sentence:

$$\begin{aligned} \phi = & TPP(\text{OfficeJohn}, \text{Floor}) \wedge TPP(\text{Corridor}, \text{Floor}) \wedge \\ & TPP(\text{OfficeJohn}, \text{Floor}) \wedge TPP(\text{Corridor}, \text{Floor}) \wedge \\ & EC(\text{OfficeJohn}, \text{Corridor}) \wedge \\ & EC(\text{OfficeAndy}, \text{Corridor}) \wedge \\ & DC(\text{OfficeJohn}, \text{OfficeAndy}) \wedge \\ & EC(\text{ServiceRoom}, \text{Corridor}) \wedge \\ & EC(\text{Tree}, \text{Floor}) \wedge EC(\text{OfficeAndy}, \text{ServiceRoom}) \wedge \\ & EC(\text{Tree}, \text{OfficeAndy}) \wedge TPP(\text{John}, \text{Floor}) \wedge \\ & TPP(\text{Andy}, \text{OfficeAndy}) \wedge \\ & TPP(\text{OfficeJohn}, \text{Floor}) \wedge \\ & TPP(\text{OfficeAndy}, \text{Floor}) \wedge \\ & TPP(\text{ServiceRoom}, \text{Floor}) \wedge \\ & DC(\text{ServiceRoom}, \text{OfficeJohn}) \end{aligned} \quad (1)$$

Definition 2 The basic sentences of probabilistic RCC-8 are defined inductively as follows:

- atoms are of the form $r(a, b)$, where $a, b \in \mathcal{C}$ and $r \in \mathcal{B}$;
- if ϕ and ψ are basic RCC-8 sentences, then $\phi \vee \psi$ and $\phi \wedge \psi$ are also basic RCC-8 sentences;
- if ϕ is a basic RCC-8 sentence, then $\neg\phi$ is also a basic RCC-8 sentence

In our example, a query is on the probability of the 'part-of' relation between John and each of the rooms. In general, the queries we want to be able to answer are about the probability of a relation to hold between two regions. This relation may be either a base relation ('externally connected') or a general relation (a disjunction of base relations). In this case, 'part-of' is the disjunction of 'proper part', 'tangential proper part' and 'non-tangential proper part' relations.

One property of PRCC sentences, that stems from JEPD-ness, namely the fact that the negation of a literal can be rewritten as a positive disjunction, is the following:

Property 1 Any basic sentence of probabilistic RCC-8 can be written as a positive sentence

Next, we define query-type sentences. These are the sentences that express probabilities of relations and, as the name implies, will be used to answer queries. A conditional query-type sentence expresses the probability of a relation given a basic type sentence: this is the kind of sentence that generally encodes a full problem. The semantics of these sentences is defined using the semantics of non-conditional query-type sentences.

In the following, α is the probability we are looking for: $p_\alpha(\forall r \in \mathcal{B}_q r(a, b))$ has the intuitive meaning that the probability that $r(a, b)$ holds is α .

Definition 3 If $0 \leq \alpha \leq 1$, $a, b \in \mathcal{C}$, $\mathcal{B}_q \subset \mathcal{B}$ and ϕ is a basic sentence, then:

- $p_\alpha(\forall r \in \mathcal{B}_q r(a, b))$ is a non-conditional query-type sentence or a query-type atom;
- $p_\alpha(\forall r \in \mathcal{B}_q r(a, b) \mid \phi)$ is a conditional query-type sentence.

Definition 4 A probabilistic RCC-8 sentence is either a basic sentence or a query-type sentence.

3.2 Semantics

In our example, a model is any spatial configuration and assignment of names to elements in the spatial configuration, i.e. which room is Andy's office, that satisfies the set of constraints given.

In general, a model of a PRCC signature will specify the topology, a subset of this topology (the 'working' regions), the set of interpretations of region constants in the 'working' region set and a probability distribution on these interpretations.

Let T be a topology on some universe U and let $X \in \mathcal{R}$ be a closed set in T . In the following, let $Int(X)$ be the interior of X and $\Gamma(X) = X - Int(X)$ be the border of X .

Definition 5 Given an RCC-8 signature \mathcal{C} , a model M of the signature is a structure of the form $M = (U, T, \mathcal{R}, W, P)$, where:

- U is a (possibly infinite) universe of points;
 - T is a topology on U ; the closed regular sets in T are called regions;
 - $\mathcal{R} \subset T$ is a finite set of regions;
 - $W = \{(U_w, w) \mid w : \mathcal{C} \cup \mathcal{B} \rightarrow U_w \uplus (U_w \times U_w)\}$ is a set of possible worlds, where for each possible world w :
 - $U_w = \mathcal{R}$ is the world universe;
 - $w|_{\mathcal{C}} : \mathcal{C} \rightarrow U_w$ is an interpretation of constant symbols as regions;
 - $w|_{\mathcal{B}} : \mathcal{B} \rightarrow U_w \times U_w$ is an interpretation of base relation symbols
- and the interpretation of base relation symbols $w|_{\mathcal{B}}$ is such that:
- $\forall X, Y \in U_w, w(DC)(X, Y)$ iff $X \cap Y = \emptyset$;
 - $\forall X, Y \in U_w, w(EC)(X, Y)$ iff $Int(X) \cap Int(Y) = \emptyset$ and $X \cap Y \neq \emptyset$;
 - $\forall X, Y \in U_w, w(PO)(X, Y)$ iff $Int(X) \cap Int(Y) \neq \emptyset$ and $X \not\subseteq Y$ and $Y \not\subseteq X$;
 - $\forall X, Y \in U_w, w(EQ)(X, Y)$ iff $X = Y$;
 - $\forall X, Y \in U_w, w(TPP)(X, Y)$ iff $X \subsetneq Y$ and $X \not\subseteq Int(Y)$;
 - $\forall X, Y \in U_w, w(TPPI)(X, Y)$ iff $w(TPP)(Y, X)$;
 - $\forall X, Y \in U_w, w(NTPP)(X, Y)$ iff $X \subseteq Int(Y)$;
 - $\forall X, Y \in U_w, w(NTPPI)(X, Y)$ iff $w(NTPP)(Y, X)$.
- $P : W \rightarrow [0, 1]$ (with $\sum_{w \in W} P(w) = 1$) is a probability distribution over the set of interpretations.

These properties of interpretation functions also ensure that the set $w(\mathcal{B})$ forms a partition over $U_w \times U_w$, or in other words the relations in $w(\mathcal{B})$ are jointly exhaustive and pairwise disjoint (JEPD).

In the rest of the paper, we will assume the topological space of the model fixed. The interpretation of base relations in this space will be the same for all models so we will omit both the topology and the interpretation of relations from the definition of a model as implied. Moreover, for all models we will have the set of interpretations to be the entire set of functions from \mathcal{C} to \mathcal{R} , so W will be completely defined by \mathcal{R} and can thus be omitted as well (restrictions to a subset of this \mathbb{I} can be made by forcing the probability of the missing interpretation functions to 0).

For basic sentences, sentence satisfaction is defined for every possible world, inductively on the structure of the sentence, as in any fragment of FOL. A sentence is satisfied if it is satisfied in every world that has a non-zero probability.

Definition 6 Given model $M = (\mathcal{R}, W, P)$, the satisfaction of a basic formula in a possible world $w \in W$ is defined inductively as:

- $w \models r(a, b)$ iff $(w(a), w(b)) \in w(r)$;
- $w \models \phi \wedge \psi$ iff $w \models \phi$ and $w \models \psi$;
- $w \models \neg \phi$ iff $w \not\models \phi$;
- $w \models \phi \vee \psi$ iff $w \models \neg(\neg \phi \wedge \neg \psi)$;

We say a model $M = (\mathcal{R}, W, P)$ satisfies a basic formula ϕ and write $M \models \phi$ iff $w \models \phi$ for all $w \in W$ with $P(w) > 0$.

Next, we will show how to answer queries, given a model and a set of constraints. The intuition is that, when we are presented with a new piece of information about the world, we constrain our model of the world so as to discard all interpretations that are not consistent with the new piece of information. The model we end up with is what we will call the restriction of a model via a basic-type sentence. Restricting the model via a sentence lowers to 0 the probabilities of all the interpretations that do not satisfy the sentence, and scales the other probabilities such that they still sum to 1.

Definition 7 Let ϕ be a basic formula and $M = (\mathcal{R}, W, P)$ a probabilistic RCC-8 model; then we can define the restriction of M via ϕ as $M|_{\phi} = (\mathcal{R}, W, P|_{\phi})$, where:

- $P|_{\phi}(w) = P(w) \cdot \frac{1}{Z(\phi)}$ if $w \models \phi$;
- $P|_{\phi}(w) = 0$ if $w \not\models \phi$

and $Z(\phi) = \sum_{w \models \phi} P(w)$ is the normalization constant.

Thus $M|_{\phi}$ is intuitively the largest submodel of M that satisfies ϕ .

In order to answer the query given a set of constraints, we restrict the model in order for it to satisfy the set of constraints, and then we sum the probabilities of the interpretations that satisfy the query. So, the satisfaction of a query-type sentence by a model M is defined as follows:

Definition 8 Given model $M = (\mathcal{R}, W, P)$, basic sentence ϕ , $a, b \in \mathcal{C}$, $\mathcal{B}_q \subset \mathcal{B}$ and $0 \leq \alpha \leq 1$, the satisfaction of query-type sentence $p_{\alpha}(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$ is defined as:

- $M \models p_{\alpha}(\bigvee_{r \in \mathcal{B}_q} r(a, b))$ iff $\sum_{w \models \bigvee_{r \in \mathcal{B}_q} r(a, b)} P(w) = \alpha$;
- $M \models p_{\alpha}(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$ iff $M|_{\phi} \models p_{\alpha}(\bigvee_{r \in \mathcal{B}_q} r(a, b))$.

It is worth noting that we are really not interested in what exactly the interpretations of constant symbols in a possible world look like, but in their relative position. So we can restrict our attention to equivalence classes of possible worlds, under the equivalence relation \simeq given by the set of base RCC relations that hold in these worlds:

$$w_1 \simeq w_2 \quad \text{iff} \quad \text{for each pair } a, b \in \mathcal{C} \text{ and for each } r \in \mathcal{B} \\ w_1 \models r(a, b) \Leftrightarrow w_2 \models r(a, b) \quad (2)$$

This will be particularly useful when introducing the factored representation.

3.3 Inference in Probabilistic RCC

Using definitions 6 and 8, we can derive the following alternative condition for the satisfaction of conditional query-type sentences - $M \models p_{\alpha}(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$:

$$\alpha = \frac{\sum_{w \models \phi \text{ and } w \models \bigvee_{r \in \mathcal{B}_q} r(a, b)} P(w)}{\sum_{w \models \phi} P(w)} \quad (3)$$

It is straightforward to implement an algorithm that finds α using this formula. If N is the size of ϕ , R is the number of regions and C is the number of constant symbols in the signature, this algorithm would require $O(R^{C+1})$ space and $O(N \cdot R^C)$ time.

Notice that this algorithm requires us to know P , the probability distribution over possible worlds. If we don't know it, the proper probability distribution to use is the one with the maximum entropy, according to the principle of maximum entropy. The set of possible worlds being a discrete and finite domain, the maximum entropy probability distribution is the uniform probability distribution.

Using this observation and the equation (3) derived in the beginning of the previous section, we can compute α in $p_\alpha(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$ as:

$$\alpha = \frac{|\{w \in W \mid w \models \phi \text{ and } w \models \bigvee_{r \in \mathcal{B}_q} r(a, b)\}|}{|\{w \in W \mid w \models \phi\}|} \quad (4)$$

4 Factored Representation of PRCC

Given a signature \mathcal{C} and model $M = (\mathcal{R}, W, P)$, for each pair of distinct constant symbols $a, b \in \mathcal{C}$, let $X_B^{a,b}$ be the random variable encoding the base relation that holds between the regions named by a and b . Then, the probability distribution P over possible worlds induces a joint probability P_B distribution over $\{X_B^{a,b}\}_{a,b \in \mathcal{C}, a \neq b}$:

$$P_B(X_B^{p_1} = r_1, \dots, X_B^{p_N} = r_N) = \sum_{w \models \bigwedge_{1 \leq i \leq N} r_i(p_i)} P(w) \quad (5)$$

where $N = \binom{\mathcal{C}}{2}$, $\{p_1, \dots, p_N\} = \{\{a, b\} \in \mathcal{C} \mid a \neq b\}$ and $r_i \in \mathcal{B}$ for $1 \leq i \leq N$.

If we consider the model consisting of equivalence classes of possible worlds, we can recover the probability distribution over these equivalence classes from the joint probability P_B , as:

$$P([w]_{r_1(p_1), \dots, r_N(p_N)}) = P_B(X_B^{p_1} = r_1, \dots, X_B^{p_N} = r_N) \quad (6)$$

where $[w]_{r_1(p_1), \dots, r_N(p_N)} = \{w \in W \mid w \models r_1(p_1) \wedge \dots \wedge r_N(p_N)\}$.

Therefore, reasoning in probabilistic RCC can be reduced to reasoning with such joint probability distributions:

Theorem 1 *Given model $M = (\mathcal{R}, W, P)$, basic sentence ϕ , expressed as a conjunction of atoms, $a, b \in \mathcal{C}$, $r \in \mathcal{B}$ and $0 \leq \alpha \leq 1$, the satisfaction of query-type sentence $p_\alpha(r(a, b) \mid \phi)$ can be computed as follows:*

- $M \models p_\alpha(\bigvee_{r \in \mathcal{B}_q} r(a, b))$ iff $P_B(X_B^{a,b} = r) = \alpha$;
- $M \models p_\alpha(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$ iff $P_B(X_B^{a,b} = r \mid \phi) = \alpha$;

Furthermore, since for every $X_B^{a,b}$, the events $X_B^{a,b} = r$ and $X_B^{a,b} = r'$ are disjoint for every $r \neq r' \in \mathcal{B}$:

Corollary 1 *Given model $M = (\mathcal{R}, W, P)$, basic sentence ϕ , expressed as a conjunction of atoms, $a, b \in \mathcal{C}$, $\mathcal{B}_q \subset \mathcal{B}$ and $0 \leq \alpha \leq 1$, the satisfaction of query-type sentence $p_\alpha(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$ can be computed as follows:*

- $M \models p_\alpha(\bigvee_{r \in \mathcal{B}_q} r(a, b))$ iff $\sum_{r \in \mathcal{B}_q} P_B(X_B^{a,b} = r) = \alpha$;
- $M \models p_\alpha(\bigvee_{r \in \mathcal{B}_q} r(a, b) \mid \phi)$ iff $\sum_{r \in \mathcal{B}_q} P_B(X_B^{a,b} = r \mid \phi) = \alpha$.

4.1 Markov Random Fields

A Markov random field (MRF) is a compact representation of a joint probability distribution by means of an undirected graph describing conditional independence. More specifically, given a joint probability distribution P over random variables X_1, X_2, \dots, X_M , an MRF consists of the following:

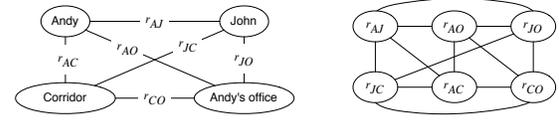


Figure 2. Compact representation of a model

- an undirected graph with vertices X_1, X_2, \dots, X_M , such that:
 - any two non-adjacent random variables are conditionally independent given all the others (*pairwise Markov property*)
 - a random variable is conditionally independent of all the others given its neighbours (*local Markov property*)
 - any two sets of random variables are conditionally independent given a third set that separates any path between the two (*global Markov property*)
- a set of factors ϕ_k associated to the cliques of the graph (over which k ranges). The joint probability distribution is then:

$$P(X_1, X_2, \dots, X_M) = \frac{1}{Z} \prod_k \phi_k(\bar{X}_k) \quad (7)$$

where \bar{X}_k is the set of variables in clique k and Z is a normalization factor called the *partition function*

Going back to PRCC, we can represent the joint probability distribution P_B in a compact way, as an MRF, using the observation that the base relation that holds between two regions named a and b is independent of any other relation that holds in the world, given the relations that hold between region named a and any other region, and the relations that hold between any other region and region named b , and these relations' duals:

$$I(X_B^{a,b}, X_B^{c_i, c_j} \mid X_B^{a,c} \cup X_B^{b,c} \cup X_B^{c,a} \cup X_B^{c,b}) \quad (8)$$

where $X_B^{a,c} = \{X_B^{a,c} \mid c \in \mathcal{C}\}$ and likewise $X_B^{c,a} = \{X_B^{c,a} \mid c \in \mathcal{C}\}$. For example, if we know where John is with respect to Andy, all the rooms on the current floor, and the current floor, and all the spatial relations that hold between the tree and Andy, the floor and all the rooms on the floor, then we don't need to know what spatial relation holds between the service room and the corridor in order to find the relation that holds between John and the tree.

This observation does not hold for all PRCC models, and we will only be able to use this compact representation for those models that do have this property. Intuitively, this is the case if we don't have any prior knowledge of the space of regions \mathcal{R} , and indeed in this case, using the naïve algorithm described in the previous section is infeasible. So we will require, in the rest of the section, that the following property holds for the models for which we give the compact representation:

Property 2 *Let $M = (\mathcal{R}, W, P)$ be a PRCC model. We say that M is unstructured if:*

- P is the uniform probability distribution, and
- for every sentence ϕ (on the same signature), if ϕ is satisfiable in RCC-8, then there is $w \in W$ such that $w \models \phi$

The MRF representation for unstructured models is illustrated on a simple example in Figure 2, for the case where we restrict our attention to the subsignature consisting of only Andy, John, the corridor

and Andy’s office. Notice that if (8) holds, then we only have edges between nodes that share a symbol:

Lemma 1 *If Property 2 holds for a model M , then in the MRF representation of M , there is an edge between nodes X_B^p and X_B^q if pairs p, q share a constant symbol, i.e. $|p \cap q| = 1$.*

This lemma leads to the following theorem:

Theorem 2 *Let \mathcal{C} be a probabilistic RCC signature and let $P_B(X_B^{p_1} = r_1, \dots, X_B^{p_N} = r_N)$ be the probability distribution over the base relations that hold between the interpretations of every two constant regions, given an unstructured model M . Then, in the MRF representation of P_B , the largest clique has size $C - 1$.*

Intuitively, every node $X_B^{a,b}$ ($a \neq b \in \mathcal{C}$) in the MRF representation is connected to two cliques of size $C - 1$: one containing all the pairs that share symbol a , and one containing all the pairs that share symbol b . Other cliques that appear in the MRF are triangles representing the relations that hold between any three regions. The interactions represented by those latter cliques stem from the constraints imposed by RCC relation composition.

Let $\bar{X}_B^{a,a'} = (X_B^{a,b})_{b \neq a, b \in \mathcal{C}}$ be the tuple containing the nodes in the clique sharing symbol a , and let $X_B^{a,b,c} = X_B^{a,b}, X_B^{b,c}, X_B^{c,a}$. For an unstructured model, one can have any combination of base relations between a region and all the other regions, i.e., we can assume $\phi(\bar{X}_B^{a,a'})$ a constant and therefore ignore it in the factorization. Therefore the probability distribution can be written as:

$$P_B(X_B^{p_1} = r_1, \dots, X_B^{p_N} = r_N) = \frac{1}{Z_B} \prod \phi_{a,b,c}(X_B^{a,b,c}) \quad (9)$$

4.2 Inference in the Factored Models

In the following we will assume we know the factors $\phi_{a,b,c}(X_B^{a,b}, X_B^{b,c}, X_B^{c,a})$ in the joint probability distribution. We can infer the probability α of $\bigvee_{r \in \mathcal{B}_q} r(a, b)$ as the sum of probabilities of each $r(a, b)$, given an evidence $\phi = r_1(a_1, b_1) \wedge \dots \wedge r_k(a_k, b_k)$:

$$\begin{aligned} \alpha &= \sum_{r \in \mathcal{B}_q} P(r(a, b) \mid \phi) \\ &= \frac{\sum_{r \in \mathcal{B}_q} P(r(a, b), r_1(a_1, b_1), \dots, r_k(a_k, b_k))}{\sum_{r \in \mathcal{B}} P(r(a, b), r_1(a_1, b_1), \dots, r_k(a_k, b_k))} \quad (10) \end{aligned}$$

using any inference method in the corresponding MRF.

If we further assume $\phi_{a,b,c}(X_B^{a,b}, X_B^{b,c}, X_B^{c,a}) = w_{a,b,c} f_{a,b,c}(X_B^{a,b}, X_B^{b,c}, X_B^{c,a})$, where the value of the feature $f_{a,b,c}$ is 1 if the configuration specified by the relations between a, b and c is possible and 0 otherwise, we can use any MRF learning algorithm to infer the set of weights $\{w_{a,b,c}\}_{a,b,c}$. We intend to explore this direction in the future, for the current work we assume that all the factors are known, or all the weights are 1.

Note that, although all the cliques have size at most $C - 1$, variable elimination can lead to factors of greater size. Since every node links two cliques of size $C - 1$, eliminating the first variable produces a clique of size $2(C - 2)$. Eliminating further variables increases by at least $C - 2$ the size of the largest clique, therefore in the course of running the algorithm, the largest clique may reach size $O(N \cdot (C - 2))$. Since $N = \binom{C}{2}$, we get the following theorem:

Theorem 3 *Variable Elimination for the factored model of PRCC has a time complexity of $O(2^{C^3})$.*

Approximate inference methods such as loopy belief propagation or sampling are beyond the scope of this paper.

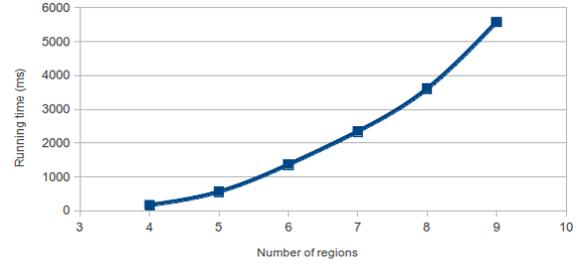


Figure 3. The running time of answering a query as a function of the number of region symbols in the signature

5 Results on Some Examples

We implemented our framework, using the MRF representation for the joint probability distribution. We used variable elimination as the inference algorithm. We experimented with answering queries on our running example (Figure 1) with slight modifications. Figure 3 shows the running time as a function of the number of regions, and Figure 4 shows the dependence on the number of constraints.

We got that John is most likely to be in his office, with probability 0.226, and that Andy is standing by the tree with probability 0.38.

We also investigated a simple story, where John is in his office, and Andy enters the office (Andy partially overlaps both the office and the corridor). In this case, Andy meets John (the disjunction of base relations EC and PO) with probability 0.435.

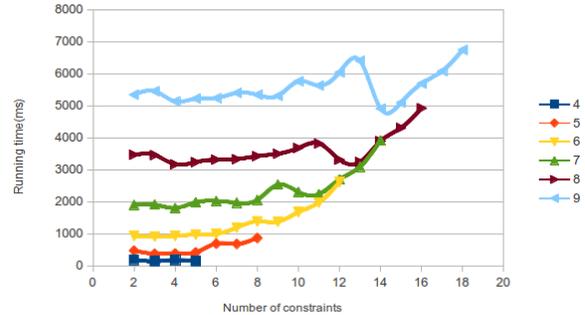


Figure 4. The running time of answering a query as a function of the number of atoms in the evidence, for different numbers of region constants

6 Related Work

Another way to do probabilistic reasoning in RCC is to use the language of Markov Logic Networks [8]. This amounts to representing PRCC as an MLN built from an axiomatization of classic RCC, such as the original axiomatization [7]. All constraints imposed by the axiomatization are considered hard, therefore the sentences in the MLN will have infinite weight. What we do here is to give probabilistic RCC an individuality of its own, with its own well-defined syntax and semantics. Furthermore, we encode the PRCC models directly as Markov Random Fields, taking advantage of the particular independence assumptions that stem from the spatial domain.

We will further discuss other related approaches to represent and reason about uncertainty in Region Connection Calculus.

Cohn et al. [2] address the problem of reasoning in the presence of vague topological information, more specifically in the case where the regions have vague boundaries. They introduce the 'egg-yolk' representation, where each region is divided into its crisp, certain subregion (the 'yolk') and a surrounding vague part (the 'white'). The intuition is that the actual region lies anywhere within the borders of the 'white' and necessarily covers the 'yolk'. In this work there is no quantification for the degree of uncertainty.

Schockaert et al. [11] [10] also deal with vague regions and add quantifiable uncertainty. Rather than work in a probabilistic setting, as in our approach, or divide each region, as in the previous approach, they develop a framework based on fuzzy logic. They take the 'connected' relation to mean the degree to which regions are connected, not a crisp truth value as in the classical RCC. With this, they redefine the entire set of base relations of RCC and subsequently the RCC framework. In contrast, we keep the classic logical framework of RCC and give it a probabilistic semantics.

In order to deal with uncertainty regarding regions, Bittner et al. [1] represent approximate regions by relating them to a frame of reference consisting of a set of unit regions. The definition of approximation makes qualitative distinctions based on the coverage of those unit regions. They then define an approximate region as a set of regions with the same approximation. With this definition, they rewrite the RCC framework to work with approximate regions.

All of these lines of work look at dealing with or quantifying vagueness rather than quantifying the likelihood of relations.

Probabilistic logic programming, or PLP [6], resembles our work mainly in the way they define the satisfaction of probabilistic sentences. One major difference is that, in PLP, each probabilistic formula (representing the probability of a conditional event) is assigned a probability *interval* - we are reasoning over single probabilities, not probability intervals. Another important difference is that any sentence in PLP is a probabilistic sentence; in our case, only the queries are probabilistic, whereas the knowledge base consists only of sentences expressed in classic RCC.

A maximum entropy semantics has also been defined for PLP [13]; that definition is based on the notion of degree of satisfaction. Since we do not use probabilistic sentences in our knowledge base, our maximum entropy model is much more simple.

Kontchakov et al. [5] have proved the sensitivity to the underlying topological space of the complexity of reasoning in a superset of RCC-8, enriched with a unary connectedness predicate and Boolean functions over regions. They also prove NP-completeness of satisfiability for the calculus enriched with connectedness only, as well as EXPTIME-hardness of the full superset. Further results [4] prove reasoning in the 2D Euclidean space RE-hard for the case when Boolean functions are allowed over regions. We do not make assumptions on the underlying topological space, and we do not talk about Boolean functions over regions.

7 Conclusions and Further Work

We showed the syntax and semantics of probabilistic RCC-8. We showed how to represent the models of this calculus compactly, by using Markov Random Fields to model the joint probability distribution over spatial relations. We then used this framework to answer queries regarding the probabilities of relations between regions, given a set of constraints, on a small set of examples.

One problem we don't address is how to handle disjunctive evidence. One way to look at this, is that, writing the evidence in disjunctive normal form, every clause serves as evidence for a possible

abstract image of the world. We would then want to combine the probabilities of base relations that result from each of these possible images. One could take an optimistic approach and use the maximum of these probabilities, for every query, but this does not accurately represent the probability distribution encoded by the model. We will explore ways to look at disjunctive evidence in future work.

As another line of future work, we want to explore more efficient algorithms for the compact representation, possibly sampling, and experiment with real world examples, including learning the weights from a larger description of the world. We would also like to explore the idea of allowing quantification over PRCC sentences and using types of regions to derive meaningful and more general representations. We believe the approach holds promise for recreating a spatial scene from a natural language description, so another line of future work would be exploring this possibility.

ACKNOWLEDGEMENTS

This work was supported by NSF IIS grants 09-17123, 09-68552 and grant NSF EAR 09-43627 EA.

REFERENCES

- [1] Thomas Bittner and John G. Stell, 'Approximate qualitative spatial reasoning', *Spatial Cognition and Computation*, **2**, 435–466, (May 2001).
- [2] A.G. Cohn and N.M. Gotts, 'The 'egg-yolk' representation of regions with indeterminate boundaries', pp. 171–187. Francis Taylor, (1995).
- [3] Anthony G. Cohn, 'Qualitative spatial representation and reasoning techniques', in *Proceedings of the 21st Annual German Conference on Artificial Intelligence: Advances in Artificial Intelligence*, pp. 1–30, London, UK, (1997). Springer-Verlag.
- [4] Roman Kontchakov, Yavor Nenov, Ian Pratt-Hartmann, and Michael Zakharyashev, 'On the decidability of connectedness constraints in 2d and 3d euclidean spaces', *CoRR*, **abs/1104.0219**, (2011).
- [5] Roman Kontchakov, Ian Pratt-Hartmann, and Michael Zakharyashev, 'Interpreting topological logics over euclidean spaces.', in *KR*, eds., Fangzhen Lin, Ulrike Sattler, and Miroslaw Truszczyński. AAAI Press, (2010).
- [6] Thomas Lukasiewicz, 'Probabilistic logic programming', in *In Proc. of the 13th European Conf. on Artificial Intelligence (ECAI-98)*, pp. 388–392. J. Wiley & Sons, (1998).
- [7] David A. Randell, Zhan Cui, and Anthony Cohn, 'A Spatial Logic Based on Regions and Connection', in *KR'92. Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*, eds., Bernhard Nebel, Charles Rich, and William Swartout, 165–176, Morgan Kaufmann, San Mateo, California, (1992).
- [8] Matthew Richardson and Pedro Domingos, 'Markov logic networks', *Mach. Learn.*, **62**, 107–136, (February 2006).
- [9] M. Sammons, V. Vydiswaran, and D. Roth, 'Ask not what textual entailment can do for you...', in *ACL*, Uppsala, Sweden, (7 2010). Association for Computational Linguistics.
- [10] Steven Schockaert, Martine De Cock, Chris Cornelis, and Etienne E. Kerre, 'Fuzzy region connection calculus: Representing vague topological information', *International Journal of Approximate Reasoning*, **48**(1), 314 – 331, (2008). Special Section: Perception Based Data Mining and Decision Support Systems.
- [11] Steven Schockaert and Martine De Cock, 'Reasoning about vague topological information', in *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management, CIKM '07*, pp. 593–602, New York, NY, USA, (2007). ACM.
- [12] Stefan Wöflf, Till Mossakowski, and Lutz Schröder, 'Qualitative constraint calculi: Heterogeneous verification of composition tables.', in *FLAIRS Conference*, eds., David Wilson and Geoff Sutcliffe, pp. 665–. AAAI Press, (2007).
- [13] Anbu Yue, Weiru Liu, and Anthony Hunter, 'Measuring the ignorance and degree of satisfaction for answering queries in imprecise probabilistic logic programs', in *Proceedings of the 2nd international conference on Scalable Uncertainty Management, SUM '08*, pp. 386–400, Berlin, Heidelberg, (2008). Springer-Verlag.