A Framework for Applying Qualitative Spatial and Temporal Reasoning

(blind review) (blind review) (blind review)

ABSTRACT
Many application domains make extensive use of spatial and temporal information. However, the numerical approaches employed by most software tools have limitations, particularly when information is vague or incomplete. To address this, alternative qualitative spatial and temporal reasoning (QSTR) methods have been developed, yet few applications have made significant use of these techniques. In response to this we are developing a framework that will support the application of QSTR by allowing software developers to create custom qualitative modelling systems. In this paper we compare our framework to more standard toolbox approaches for QSTR application support. We present fundamental principles of qualitative modelling, and demonstrate, using an architectural lighting example, how these principles provide a basis for creating qualitative modelling systems that incorporate domain knowledge.

Categories and Subject Descriptors
I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods – *relation systems, representation languages*

General Terms
Design, Theory

Keywords
qualitative spatial and temporal reasoning, application framework, artificial intelligence

1. INTRODUCTION
Spatial and temporal information is a central component in many application domains, including geographic information systems when considering proximity constraints or pedestrian and vehicular traffic flow modelling, project scheduling when reasoning about uncertain temporal information, construction IT, astronomy, robotics, and so on. Standard software tools used to assist in these fields often rely on numerical approaches for representing (or modelling) and processing (or reasoning about) this spatial and temporal information. However, numerical methods have serious limitations when information is vague or incomplete, which are often inherent characteristics of spatial and temporal data available. To illustrate this, consider a numerical ray tracing tool used by architects to simulate light being emitted from a source, and reflected and absorbed by surfaces e.g. [1]. The aim of the simulation is to provide estimates for light intensity across surfaces in a room, which are used by the architect to determine whether a proposed lighting configuration meets necessary criteria.

The first limitation is that, if information is incomplete numerical methods breakdown. If an equation required by the ray tracing simulator has unknown variables, the equation can not be resolved, i.e.:

$$x = 5 + ? + 3$$

Because one of the values in the equation is unknown, the value of x must also remain unknown: x=?. Incompleteness can arise due to limitations of measuring equipment, or simply as a feature of the early stages of a design project, e.g. the exact materials may not be specified, certain dimensions may not be available, the client may still be deciding on whether or not to have an open-plan living space, and so on.

Secondly, even when information about the uncertainty is available, such as range restrictions or probability distributions, numerical methods can become highly sensitive to small changes in the initial parameters, making the results of an individual simulation run unreliable. For example, in Figure 1 the dimensions of a protruding wall are incomplete, but the length is known to lie between 1.5m and 2.5m. However, at 2.1m the wall begins to occlude light from entering a portion of the room. Thus, only a small change in wall length, from 2.0m to 2.1m, results in a drastically different light intensity distribution across surfaces in the room. In more complex simulation runs with many uncertain parameters, sensitivity can be a very significant problem.

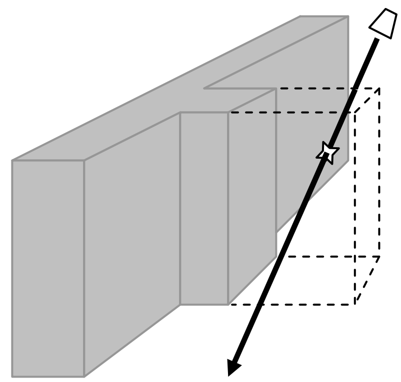

Figure 1. example of numerical sensitivity, where a small numerical change in wall length can abruptly begin to block light (arrow), leading to very different simulation results.

Finally, many concepts which people use in everyday life are impossible to directly encode in numerical equations, for example, the architect's client may require a room to appear 'dramatic' and 'sophisticated'. There are no numerical units for 'drama', and no numerical formulae that directly translate light intensity measurements and furniture configurations into a measurement for 'sophistication'. Thus, standard software tools provide very limited support, if any, for these tasks.

These numerical method limitations have motivated the development of alternative approaches for representing and reasoning about coarse qualitative concepts, called qualitative spatial and temporal reasoning (QSTR) [2,3]. For example, rather than specifying a numerical path for navigation such as "walk 105.6m, turn 44^o, walk 55.3m, ...", a qualitative solution would be "walk a *small distance* down Symonds St, *turn left* into Alfred St,...". Allen's temporal interval calculus [4] has been particularly influential. The calculus specifies thirteen different qualitative relations that can exist between two time intervals, such as before, overlaps, and during (see Figure 2). An inference mechanism is provided so that, given relations for time intervals (t1, t2) and relations for (t2, t3), the relations that can hold between (t1, t3) are defined, allowing reasoning about networks of time intervals. Two examples follow, where • is the composition operator:

t1 before t2 • t2 contains t3 = t1 before t3

t1 overlaps t2 • t2 during t3 = t1 (overlaps, or during, or starts) t3

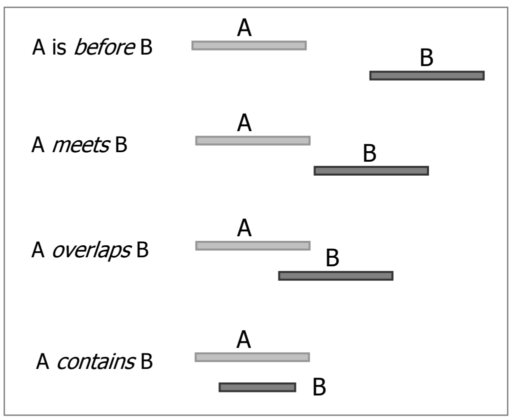

Figure 2. extract of Allen's [4] qualitative relations between temporal intervals

Despite the growing number of specialised qualitative methods being developed, relatively few substantial applications based on qualitative spatial and temporal reasoning have emerged [3,5]. We are addressing this issue by developing a framework that supports the application of qualitative spatial and temporal reasoning techniques. The framework will allow software developers to create completely customised qualitative modelling systems, rather than being restricted to existing 'off-the-shelf' qualitative tools. Section 2 provides an overview of the framework in the context of other related approaches. Section 3 presents our three fundamental principles of qualitative modelling. Section 4 uses an architectural lighting example to illustrate how the three principles can be applied. Section 5 discusses further modelling, verification, and implementation issues as part of our future research.

2. FRAMEWORK OVERVIEW

The lack of qualitative spatial and temporal reasoning applications has prompted the recent development of QSTR software libraries and toolboxes such as SparQ [5] that provide qualitative method implementations such as Allen's interval calculus, in a standard framework with a uniform interface (see Figure 3 (a)).

However, in many cases an existing qualitative formalism will not directly meet the needs of a task. Firstly, qualitative terms are vague, and rely on context for disambiguation and meaning, e.g. "near" can be a function of financial cost, time travelled, or distance covered, and can vary depending on the user being a pedestrian, a car owner, or a city planner [6]. Secondly, the qualitative concepts being modelled must be relevant for the task – it is difficult to gauge the usefulness of a qualitative term a priori without knowing the requirements of the task.

Thus, qualitative formalisms often need to be extended or modified, by adding or removing qualitative relations, or changing some part of an inference rule. A serious difficulty is that, in general, it is difficult to predict the resulting change in reasoning properties such as soundness, completeness, and computational complexity [7] – the modified reasoning engine may be intractable and logically faulty.

Rather than adopting a purely toolbox approach, where software developers are presented with readymade plug-in components, we present a framework that acts as a practical guide in creating customised qualitative modelling systems (see Figure 3 (b)), aimed at software developers who have little or no experience with the qualitative reasoning literature. Note that the framework may recommend a toolbox approach where appropriate. The framework currently defines fundamental principles of qualitative information, e.g. answering questions like:

- what is a quality?
- how do qualities work?
- how do you reason with qualities?

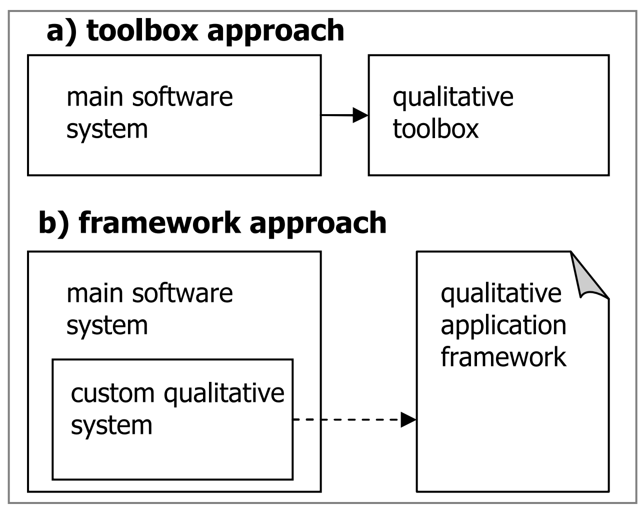

Figure 3. different approaches to applying QSTR: boxes represent software systems, page represents body of research (not software), normal arrow indicates software dependency, dotted arrow indicates design dependency only. (a) main software refers to a library of pre-made qualitative methods (b) user creates custom qualitative system as part of main software based on our framework (e.g. in the form of design guidelines and theoretical information)

These principles provide a basis for constructing custom qualitative modelling systems. The framework will also provide guidelines for efficient software implementation and techniques for verifying the customised logic of the qualitative system.

3. QUALITATIVE MODELLING AND REASONING

This section describes our three fundamental principles of qualitative modelling, which provide a foundation for developing custom qualitative modelling systems.

The first principle relates to modelling using qualitative information. Qualitative approaches require the definition of "qualities" used to describe objects and their relationships. The way in which a quality describes the world can be determined from the term's use in everyday language: a "quality" is an inherent or distinguishing property [8] (where 'property' can be generalised to n-ary relations). Thus, the function of introducing these qualitative relations is to provide a distinction between the objects that are being modelled. For example, "eared" (having ears) is not a common everyday qualitative term for describing people (at least in my experience), most likely because it would fail to provide useful distinctions within an everyday context, i.e. most people fall into the same category ("tall", "doesn't wear glasses", etc. are more common). The simplest way that a qualitative relation can introduce a distinction is if the relation introduces a single point of difference, where a relation either (a) holds, or (b) does not hold. This leads to the first qualitative modelling principle: a "quality" provides a distinction between objects – those that have the quality, and those that do not.

The second principle relates to reasoning using qualitative information. For most tasks there is indefiniteness about exactly which relations hold for a given object in a dataset. Reasoning reduces this indefiniteness by making inferences based on initial premise information and, by doing so, can be used to solve a problem or accomplish some task. For example, the qualitative relations "near" and "far" can not hold for the same pair of objects at the same time (mutually exclusive). If premise information states that "near" holds, then according to the constraint, "far" must not hold. A more complex inference rule is as follows, where a, b and c are intervals of time:

"if **a** happened *before* **b**, and **c** happened *during* **b**,
then **a** must also have happened *before* **c**"

A schedule is illustrated in Figure 4 for which this rule holds. Inference rules can also refer to qualitative relations at different levels of abstraction, for example:

"if a **room** has *warm*, *bright* ambient illumination
then it will evoke a *relaxing* impression"

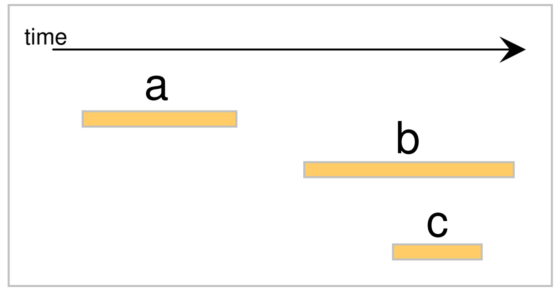

Figure 4. a, b, and c are time intervals. Schedule illustrates that a is before c, given that a is before b, and c is during b.

This leads to the second qualitative modelling principle: certain qualities can not hold for the same object at the same time – these are domain constraints that provide the basis for inference and reasoning. A great advantage of qualitative reasoning is that, given only vague or incomplete qualitative information at any of the represented levels of abstraction, a qualitative approach will infer as much as possible according to distinctions within the model. In contrast, a numerical approach requires information to be expressed as unit values, and processing capability is limited if the quantities are not available.

The final principle relates to tasks that can be performed by a qualitative modelling system. The group of objects for which some combination of qualitative relations hold represents a "class" or "type" of object, e.g. the expression that "all roads far from the city, and overlooking the lake are idyllic" is defining the class of idyllic roads according to qualitative relations that must hold. All tasks accomplished by a qualitative approach are some form of object classification, including inference (when premise information is incomplete), instantiation (e.g. determining concrete numerical dates for a schedule so that they are consistent with the qualitative relations), and querying (e.g. "find cafes near the art gallery on a quiet street"). This provides the final qualitative modelling principle: the purpose of qualitative reasoning is ultimately object classification - object classes are defined by whether certain qualities hold and do not hold; qualitative reasoning solves tasks by determining the class of objects. We have now identified the exact limit of a qualitative approach – if the required task can not be accomplished by a form of object classification, then a qualitative modelling system will not be of any assistance.

These three principles provide the necessary background for developing a qualitative modelling system.

4. ARCHITURAL LIGHTING EXAMPLE

This section illustrates the application of our three principles in the domain of architectural lighting, by describing the design of a software tool that we have developed. The tool assists architects by analysing electrical lighting installations and reporting on the broad subjective impressions that will be evoked. The architect will use the tool to quickly explore various designs during the early stages of a project.

4.1 Background to Architectural Lighting

An architect is not only concerned with ensuring that their buildings satisfy well-defined technical specifications, but also

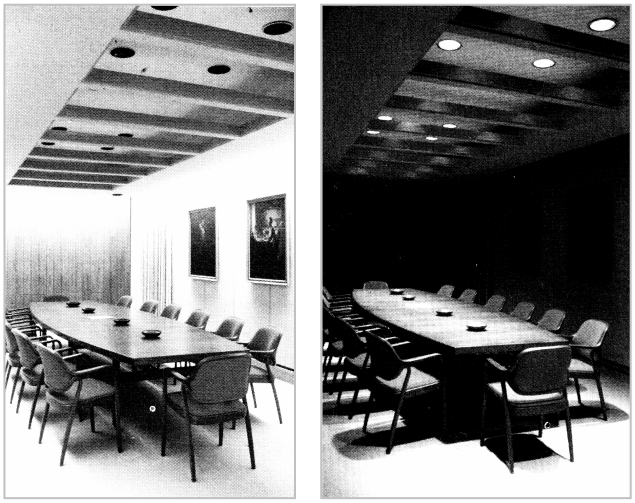

1) neutral clarity
2) spacious

1) generally hazy, quiet
2) strong confinement

Figure 5. two scenes eliciting very different subjective responses through a change in lighting (original research and images by Flynn [11,12], reproduced from [13] fig. 4.22 and 4.23, section 4.5.4, p.64)

eliciting some emotional, subjective response to an environment [9]. For example, a client may require the impression or atmosphere of an environment to be dramatic and sophisticated. Within the discipline of architectural lighting it is now accepted that a wide range of subjective impressions, such as relaxation, excitement, intimacy, and spaciousness can be achieved simply by varying parameters of the lighting installation [10]. In particular, Flynn has performed studies on the relationships between metric lighting parameters (luminances and luminous patterns) and subjective response [11,12]. The aim was to establish basic guidelines on how to influence a range of nonvisual effects with a lighting scheme, resulting in the identification of the following five key impressions (see Figure 5):

(i) visual clarity, referring to a person's subjective impression of how clearly or distinctly interior details, objects, and other people's features appear, ranging from clear to hazy,

(ii) spaciousness, referring to the apparent volume of a space, ranging from spacious to cramped,

(iii) relaxation, referring to the apparent work intensity, ranging from relaxed to tense,

(iv) intimacy, referring to the feeling of privacy in a space, ranging from private to public, and

(v) preference or pleasantness, referring to the subjective evaluation of the lighting environment, ranging from like to dislike.

The above subjective responses are elicited by a number of intermediate qualitative lighting conditions; these relationships are shown in Table 1. For example, to create a sense of relaxation indirect luminaires could be selectively placed around the periphery (e.g. wall sconces or accent lighting on wall art decorations), complemented with direct low intensity incandescent lamps placed over the area of occupancy [13].

4.2 Designing the Qualitative Modelling System

We will now apply the qualitative modelling principles for the software tool's design.

The tool will take rooms and electric lighting designs, and then generate a report on the subjective impression that the rooms will evoke. The third principle states that qualitative approaches can only perform tasks requiring some form of object classification. In this system the objective is to determine the class of subjective impression that a room object falls under.

The first principle states that qualities are used to make relevant distinctions between objects. Each of Flynn's five subjective impressions provide two broad categories for describing an environment, making perfect candidates for providing relevant distinctions between rooms - we define the qualitative unary relations: *clear, hazy, spacious, cramped, relaxed, tense, private, public, like,* and *dislike*. We can encode some basic structure in these relations by applying the second principle of constraining relations; each pair is mutually exclusive (e.g. a room can never be both clear and hazy).

The studies conducted by Flynn (summarised in Table 1) indicate further constraints with other qualitative relations, e.g. "low ambient illumination in the area of occupancy, with bright perimeter emphasis will evoke intimacy". These provide key

Table 1. lighting conditions (rows) required to elicit the desired subjective impression (columns). Based on research by Flynn [11], and adapted from [13] (pp. 61-72), and [14] (pp. 118-119).

	Clarity	Spaciousness	Relaxation	Intimacy	Pleasantness
Ambient Illumination			☑ (bright)	☑ (low in occupancy area)	
Room colour temperature	☑ (cool)		☑ (warm)		☑ (warm)
Perimeter emphasis	☑ (some)	☑ (uniform)	☑ (nonuniform)	☑ (high brightness)	☑
Work surface illumination	☑ (bright, uniform)	☑ (bright, uniform, central)			

inference rules for our qualitative modelling system, using the second principle. Relations such as *bright ambient illumination,* and *warm colour temperature* are added to the system, and constrained according to Table 1.

At this point the system has relations representing relevant output. Now the input or premise information will be considered. The input lighting configurations are metric descriptions, and thus, applying the second principle, we can refer to research that defines our relations in terms of metric ranges, e.g. the qualitative appearance of different colour temperatures given in Table 2 is widely agreed upon [10,15]. If research is unavailable (e.g. when is a table considered to be in the 'centre' of a room?) the relations can become necessary premises specified by the user, possibly with a default coarse definition, which becomes a modelling assumption.

Table 2. values of a metric called correlated colour temperature (CCT) and their qualitative colour appearance

Colour appearance	CCT
cool (bluish white)	≥ 5000 K
intermediate (white)	< 5000 K; ≥ 3300 K
warm(yellowish white)	< 3300 K

Finally, standard spatial relations for orientation, distance, and topology are used to infer the state of more abstract relations such as 'uniform perimeter emphasis' (refer Table 1). For example, qualitative *light source direction* is a relationship indicating whether a light source is directed towards or away from a surface. The key to defining these qualitative spatial relations comes from domain knowledge (as argued in Section 2) – at an early stage of design, the basic initial decision to apply accent lighting to a particular art piece is far more important than details such as precisely orienting a light source to minimise spillage. The type of light that the architect has chosen (e.g. spot light vs. diffuse light) is a qualitatively significant indicator of its intended use. Thus, the type of light provides a basis for approximating the light beam: the beam shape of a 'directed' light source (such as a spot light or a small aimable light) is represented as a line, so that the source is only directed at surfaces to which it is purposefully aimed, regardless of beam width. Alternatively, if a light source is 'diffuse' then it is directed at every surface in the room (see Figure 6).

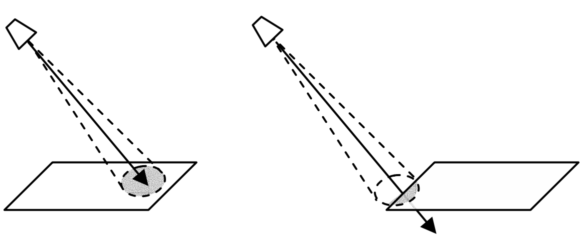

Figure 6. approximating the beam of a 'directed' light source as a line shape, to qualitatively determine whether the source is directed *at* or *away* from a surface. If the line intersects the surface (left) the directed value is *at*, otherwise (right) the directed value is *away* regardless of beam width.

The resulting system is then verified according to results of the studies by Flynn (verification details of this particular system are in [16]). This example illustrates that the application of the three main principles, along with appropriate domain knowledge, are sufficient for developing effective qualitative modelling systems. A more formal set of guidelines are currently being developed around these principles.

5. FUTURE WORK

A number of issues relating to modelling, verification, and implementation, are the focus of current and future research.

Firstly, while the core of a qualitative concept is well-defined, the boundary between neighbouring concepts may be vague or incomplete [17,18]. In some cases it is necessary to directly model this vagueness by referring to the membership that an object has in a relation, e.g. using fuzzy logic or rough sets.

Secondly, giving a user freedom to completely customise a qualitative modelling system brings about the risk of faulty logic, poorly chosen qualitative relations, and so on. Measures to quantify the efficiency and effectiveness of a system are being explored, along with techniques such as test case verification as used in software testing (e.g. verification against the Flynn studies in Section 4).

Finally, we are developing a straightforward method for implementing any qualitative modelling system. The principles in Section 3 have an underlying set theoretical basis. For example, the qualitative concept that "**houses** *far away from* the **city** with *lake side views* are *relaxing*" can be interpreted as specifying three sets of houses $\textit{far from}_{\text{city}}$, $\textit{adjacent}_{\text{lake}}$, and *relaxing*, with a relationship that:

$$\textit{far from}_{city} \cap \textit{adjacent}_{\text{lake}} \subseteq \textit{relaxing}$$

This leads to a natural and simple implementation in relational databases, which also have an underlying set theoretic definition. This provides an ideal platform for implementation, as many software developers are very familiar with relational databases.

6. CONCLUSIONS

Qualitative spatial and temporal reasoning (QSTR) techniques can address the limitations of purely numerical approaches. We are developing a framework to support the application of qualitative approaches by allowing software developers to create custom qualitative modelling systems, in contrast to other software toolbox approaches. We have presented three main principles of qualitative modelling and reasoning:

- a "quality" provides a distinction between objects – those that have the quality, and those that do not
- certain qualities can not hold for the same object at the same time – these are domain constraints that provide the basis for inference and reasoning
- the purpose of qualitative reasoning is ultimately object classification - object classes are defined by whether certain qualities hold and do not hold; qualitative reasoning solves tasks by determining the class of objects

Using an example from architectural lighting, we have illustrated how a user can develop their own custom qualitative modelling system by combining the above principles with domain knowledge. Current research is focused on further issues relating to modelling, verification, and implementation.

7. ACKNOWLEDGMENTS

(blind review)

8. REFERENCES

[1] Ward GJ, The RADIANCE lighting simulation and rendering system. Proceedings of the 21st annual conference on Computer Graphics and Interactive Techniques SIGGRAPH, ACM Press New York, NY, USA;1994

[2] Forbus KD, Qualitative reasoning. In: Tucker AB, editor. CRC Handbook of Computer Science and Engineering, CRC Press; 1996, p. 715-733.

[3] Cohn A, Hazarika, SM, Qualitative spatial representation and reasoning: An overview. Fundmenta Informaticae 2001; 46,1-2:2-32.

[4] Allen JF, Maintaining knowledge about temporal intervals. Communications of the ACM 1983; 26,11:832-843.

[5] Dylla F, Frommberger L, Wallgrun JO, Wolte D, SparQ: A toolbox for qualitative spatial representation and reasoning. Proceedings of the Workshop on Qualitative Constraint Calculi:Application and Integration, Workshop at KI 2006; 79-90.

[6] (blind review)

[7] Ligozat G, Mitra D, Condotta J, Spatial and temporal reasoning: beyond Allen's calculus. AI Communications 2004; 17,4:223-233.

[8] Concise Oxford English Dictionary. 11th ed. Soanes CS, Stevenson A, editors. Oxford University Press; 2006.

[9] Jay PA, Review: Subjective criteria for lighting design. Lighting Research and Technology 2002; 34, 2:87-99.

[10] Cuttle C, Lighting by Design. Elsevier Science; 2003.

[11] Flynn JE, A study of subjective responses to low energy and nonuniform lighting systems. Lighting design & application 1977; 7,2:6-15.

[12] Flynn JE, Spencer TJ, Martyniuk O, Hendrick C, Interim study of procedures for investigating the effect of light on impression and behaviour. Journal of the Illuminating Engineering Society 1973; 3,2:87-94.

[13] Steffy G, Architectural lighting design. John Wiley & Sons, Inc. New York; 2002.

[14] Egan MD, Concepts in architectural lighting. McGraw-Hill; 1983.

[15] Sanders MS McCormick EJ, Human Factors in Engineering and Design. 7th ed. McGraw-Hill, 1993.

[16] (blind review)

[17] Cohn A, Gotts NM, The 'Egg-Yolk' representation of regions with indeterminate boundaries. Proceedings of GISDATA Specialist Meeting on Geographical Objects with Undetermined Boundaries, Taylor & Francis; 1995:171-187.

[18] Guesgen HW, Fuzzifying spatial relations. Applying soft computing in defining spatial relations, Physica-Verlag, Heidelberg, Germany; 2002, p. 1-16.