In: Singh, M., Travé-Massuyès, L. (eds.) Proc. of the IMACS Workshop on Decision Support Systems and Qualitative Reasoning, 181-187, North-Holland, Amsterdam 1991.

# Conceptual Neighborhood and its role in temporal and spatial reasoning

Christian Freksa

Institut für Informatik
Technische Universität München
Arcisstr. 21, 8000 Munich 2, Germany
(freksa@informatik.tu-muenchen.de)

An extension of Allen's approach to interval-based temporal reasoning is presented. The new method allows for temporal and spatial reasoning on the basis of incomplete or imprecise knowledge of the kind that is available from inference and perception processes. The central idea of the representation method is the structuring of knowledge according to the *conceptual neighborhood* of temporal and spatial relations. This representation allows for integration of coarse and fine knowledge. Logical reasoning on the basis of such knowledge therefore takes place within a unified scheme. The method presented not only is more efficient than Allen's method, it also is more 'cognitively adequate' in comparison with previous approaches.

## 1. Interval-Based Temporal Reasoning à la Allen

In his popular paper on maintaining knowledge about temporal intervals, James Allen [1] proposes to represent temporal relations between intervals to describe qualitative relations between events. For reasoning about relations between events, Allen develops rules which allow for deriving knowledge about certain relations between events from knowledge about temporal relations between related events.

### 1.1 Thirteen Qualitative Temporal Relations

The basis for Allen's interval algebra is the set of thirteen mutually exclusive qualitative relations between two intervals by which the temporal relationship between any two events can be unambiguously described. These relations are depicted in Fig. 1. The set consists of seven basic relations and their inverses; in case of the 'equal' relation the basic and inverse relations are identical. In the pictorial example, the events 'X' and 'Y' are denoted by contiguous sequences of the characters 'X' and 'Y', respectively.

Temporal reasoning is done by deriving the set of possible relations between two events which have known relations to a third event. For this purpose, all 13*13 possible compositions of two relations are listed in a table. These compositions may result in unique relations or in a set of several alternative possibilities.

| Temporal Relation | Basic Symbol | Inverse Symbol | Pictorial Example |
|---|---|---|---|
| X *before* Y | < | > | XXX    YYY |
| X *equal* Y | = | = | XXX<br>YYY |
| X *meets* Y | m | mi | XXXYYY |
| X *overlaps* Y | o | oi | XXX<br>    YYY |
| X *during* Y | d | di |   XXX<br>YYYYYY |
| X *starts* Y | s | si | XXX<br>YYYYY |
| X *finishes* Y | f | fi |     XXX<br>YYYYY |

Fig. 1: The 13 possible relations between two intervals

1.2 Examples for Interval-Based Inferences

1) If we know that event A immediately precedes event B (A *meets* B) and event C takes place during event B (C *during* B), we can derive that event A takes place *before* event C:

```
A meets B          A m B
C during B         B di C AAAABBBB
---------------    ---------     CC
A before C         A < C
```

2) If we know that A takes place *during* B and C *starts* B, we can derive that either A takes place *during* C or A *finishes* C, or A is *overlapped-by* C, or A i*s-met-by* C, or A takes place *after* C; none of the remaining eight relations may hold between events A and C:

```
A during B          A d B
C starts B          B si C
---------------     ---------
A {during,          A {d, f, oi, mi, >} C
finishes,
is-overlapped-by,                  AA
is-met-by,                    BBBBBBBBB
after} C                 >    C
                         mi   CC
                         oi   CCC
                         f    CCCC
                         d    CCCCC
```

1.3 Criticism of Allen's Approach

Allen's reasoning is carried out by look-up in the exhaustive composition table. This is an efficient method but the table does not reflect an understanding of a physical or logical structure underlying the inferences. As a consequence, it is difficult to extend the representation scheme in such a way that it is robust against variations or small errors in the input knowledge: if the input cannot be relied on, the output becomes completely unpredictable.

Although Allen's reasoning scheme can represent incomplete knowledge in form of disjunctions, the inference mechanism does not really process such knowledge; instead, incomplete knowledge is complemented into an exhaustive set of completely specified alternatives; each of these pieces of complete knowledge then is processed individually.

Incomplete knowledge, however, is omnipresent in temporal and spatial reasoning: typically, the knowledge that is available for reasoning is incomplete to start with; but even if it is complete, after only one inference step we have to face incomplete knowledge (previous section, example 2).

From a cognitive perspective, Allen's inference scheme has a rather unplausible property: the less knowledge is given to the procedure (due to incomplete knowledge), the more complex is its representation and the more processing has to be done. This is due to the fact that incomplete knowledge is treated as a set of completely specified alternatives.

Allen's inference scheme obeyes the laws of logics but ignores some useful laws of the physics of time which we have incorporated in our approach.

## 2. THE REPRESENTATION OF INCOMPLETE KNOWLEDGE

There are two ways of dealing with incomplete knowledge corresponding to a bottom-up and a top-down view of the world, respectively. The bottom-up view suggests that incomplete knowledge is due to omission of specifications; thus it can be completed by considering the set of possible augmentations. The top-down view suggests that incomplete knowledge is due to possible distinctions of details which are not made; thus, by ignoring details, we can deal with coarse knowledge.

Each of those views is associated with a corresponding knowledge representation and processing philosophy. Both approaches eventually yield the same results; but they use different paths of differing length.

### 2.1 Disjunctions of Completed Knowledge

The use of disjunctions of complete knowledge to represent and process incomplete knowledge is a classical propositional approach which is used by Allen: incomplete knowledge is augmented into a disjunction of all alternatively possible complete propositions; each of the alternatives suggested by the disjunctions is then treated as complete knowledge. The results obtained by the individual complete reasoning processes are again combined into disjunctions reflecting the set of actually possible outcomes.

This approach is restricted to situations in which all possible completions of knowledge are foreseeable. For complexity reasons, the number of alternatives considered must be small.

## 2.2 Abstraction from Details

This approach is suggested by models of perception: if I see something for the first time at some distance, I may not see all the details of the object nor may I be able to consider all possible instances of details which I am missing; nevertheless I will be able to make certain observations and inferences about the object. The reason why this is possible is that many observations and inferences are independent of the missing details; knowledge of the details merely would allow for a refinement of the observations and inferences and would not require their correction.

With this view of the world we organize knowledge hierarchically according to the level of detail which is available: less knowledge corresponds to the higher levels and more knowledge to the lower levels in this organization. On any level certain inferences can be drawn. These inferences can be expressed in terms of knowledge represented on the same level of detail or of a higher or lower level. One advantage of this approach is that inferences that can be drawn on a higher level subsume several (possibly many) corresponding inferences on lower levels. If additional knowledge about details becomes available, the inferences are refined. Knowledge is treated more like a painting than as a text.

A prerequisite for employing this approach is monotonicity of the reasoning processes involved. This means, that inferences carried out on the basis of coarser knowledge must remain valid when additional knowledge becomes available. Differently stated this means that the 'picture' painted by the coarse knowledge always must be complete in the sense that there are no gaps on this level which could be filled by additional knowledge; ambiguous situations 'melt' into a single representation. The picture may be incomplete in the sense, that additional knowledge resolves the painting more finely, however. This monotonicity property can be established for temporal and spatial knowledge as will be seen in the following sections of the paper.

# 3. Conceptual Neighborhood

Although the examples given in the present paper are taken from the temporal domain, the statements about neighborhood apply to the spatial domain as well [2]. The basic approach to exploiting monotonicity for temporal and spatial reasoning is guided by the observation that conceptually neighboring relations between events have similar behavior [3]. In order to be more precise, we make the following

Definition:
Two relations between pairs of events are *conceptual neighbors* if they can be directly transformed into one another by continuous deformation (i.e., shortening or lengthening) of the events.

Example:
The relations *before* (<) and *meets* (m) are conceptual neighbors since a temporal extension of the earlier event may cause a direct transition from the relation *before* to the relation *meets:*

```
A < B      AAAAA   BBBBB

A m B      AAAAAAABBBBB
```

The relations *before* (<) and *overlaps* (o) are not conceptual neighbors, since a transition between those relations must go through the relation *meets:*

```
A o B      AAAAAAAAA
                BBBBB
```

## 3.1 Neighborhood Structure of the 13 Relations

The thirteen qualitatively different pairs of intervals (depicted by a dumbbell-shaped line and a rectangle)

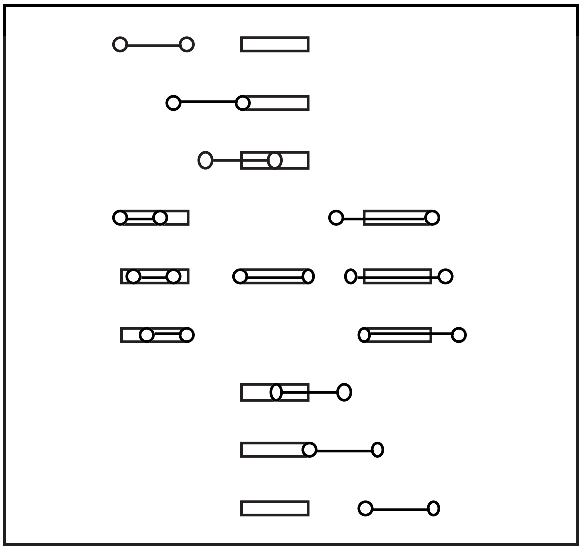

Fig. 2: The 13 relations arranged according to conceptual neighborhood

are arranged in Fig. 2 in such a way, that continuous transformation of the corresponding events only will result in transitions between spatially neighboring pairs. Since the resulting neighborhood structure will play a major role in the following sections, we will develop a neighborhood-oriented symbolism for referring to the relations within the structure.

### 3.2 Iconization of the Neighborhood Structure

The neighborhood structure obtained by continuous deformation of intervals is shown in Fig. 3. The images of the intervals from Fig. 2 are replaced here by circles containing the symbolic abbreviations of the names of the corresponding relations as given in Fig. 1. The neighborhood relations are depicted by solid lines:

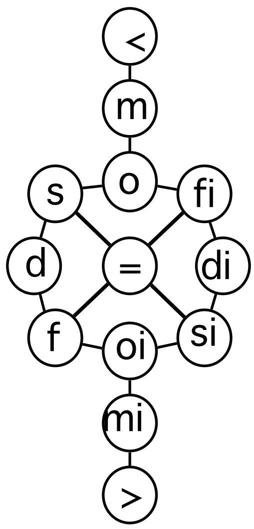


Fig 3: Symbolized relation neighborhood structure

The topological arrangement of Fig. 3 is used as generic structure for symbolizing disjunctions of interval relations by means of icons. The individual relations are symbolized as follows:

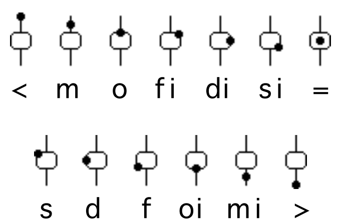


Any of the $2^{13}$-1 disjunctions of one or more individual relations can be symbolized by superposition of icons as shown in the following examples.

The icon

corresponds to the disjunction of the relations <, m, o, s, d and the icon

corresponds to the disjunction of all thirteen relations (which means that no constraint on the relationship between the intervals is given).

### 3.3 Linearization of the Neighborhood Structure

In effect, we now have a system of 2-dimensional symbols for representing disjunctions of temporal relations. We use these symbols for representing abstract temporal and spatial relationships. We will exploit the fact that the icons can be viewed both as logic symbols and as topological images of temporal or spatial relationships.

The advantage of the topolog ical view of the icons is that we obtain information on the neighborhood of relations without explicitly reasoning about neighborhood.

In order to visualize the reasoning procedure within a 2-dimensional graphical domain, we will modify the topological arrangement in such a way that only a subset of the actual neighborhood relations is reflected. Specifically, we will neglect the neighborhood between the relations o and s, oi and si, f and =, = and fi, for this purpose. This yields a linear neighborhood structure which consists of the following sequence of the thirteen relations: <, m, o, fi, di, si, =, s, d, f, oi, mi, >.

We use this ordering of relations for creating a neighborhood-oriented variation of Allen’s composition table for reasoning about temporal or spatial intervals. The resulting table is depicted in Fig. 4.

### 3.5 Neighborhood-Oriented Composition Table

The composition table consists of 13 rows and 13 columns arranged according to the linear neighborhood structure developed in the previous section. The entries of the table refer to the disjunctions of relations which may hold under composition.

The significance of the neighborhood-oriented arrangement of the relation composition table stems from the fact that certain properties of temporal and spatial structures are preserved which are not represented in Allen’s scheme.

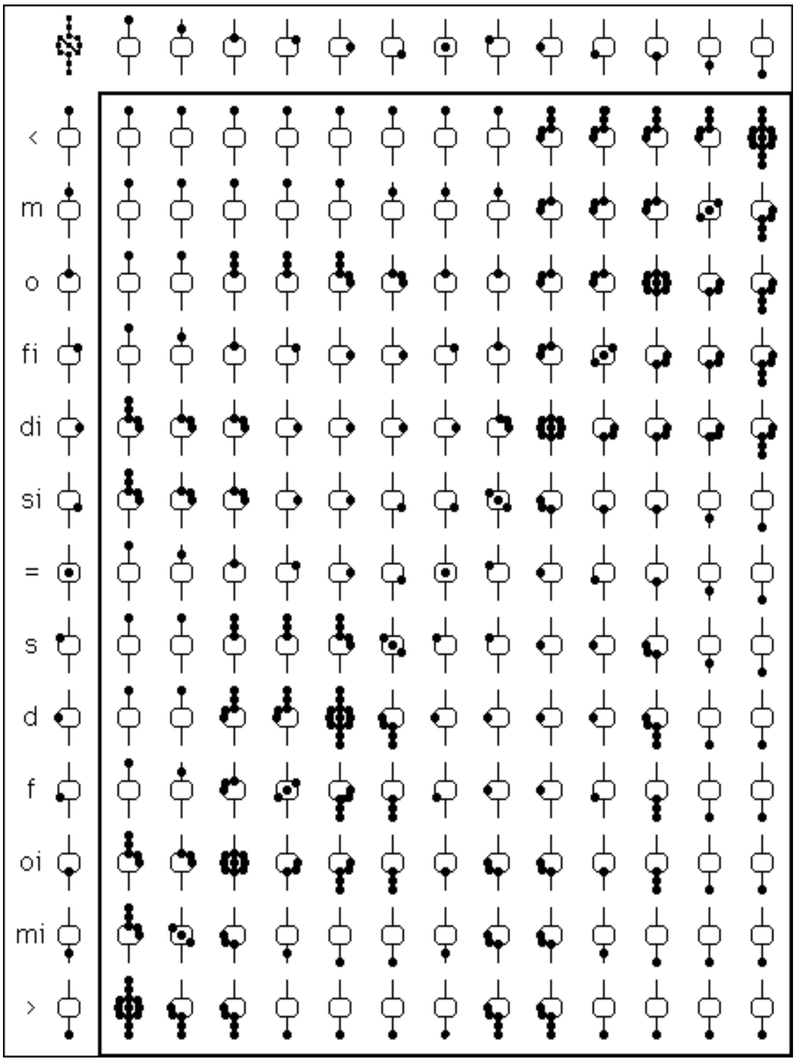


Fig. 4: Neighborhood preserving relation composition table.

Some important properties of the composition table are listed below:

1) Most compositions (= temporal or spatial inferences) result in disjunctions of several alternative relations; some compositions result in unambiguous relations.

2) The relations within a disjunction always form a *conceptual neighborhood*, i.e., they are connected via conceptual neighbors.

3) In many cases, a transition to neighboring initial conditions results in the identical conclusion or in a subset or superset of an inference neighborhood.

4) In no case, a transition among neighboring initial conditions results in a jump between non-neighboring conclusions.

5) Only a small subset of 27 of the possible conceptual neighborhoods corresponds to actually possible inferences.

6) The inference table shows many symmetries which may be utilized in the inference process.

7) All these properties hold for the complete 4-dimensional composition space which we would get without linearization of the neighborhood structure

## 4. COARSE CODING AND NEIGHBORHOOD-BASED REASONING

The benign monotonicity properties of the compositions in the neighborhood-based framework allow us to represent the inference knowledge of the composition table on a coarser level: we construct neighborhoods of initial conditions in such a way that we obtain each of the original 13 relations by conjoining neighborhoods. This corresponds to a more abstract level of representing the basic knowledge.

### 4.1 Condensed Composition Table

We choose the following neighborhoods of relations as 'primitives' of our representation:

The first and last two relations are trivial neighborhoods consisting of only a single relation each. This is due to the fact that they only have one or two neighbors. The remaining nine of the thirteen relations each have at least three neighbors.

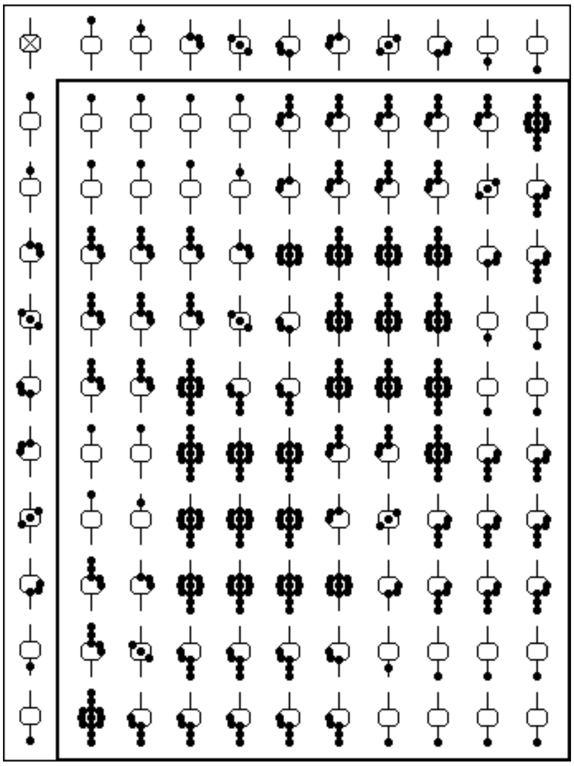

Fig. 5 Inference table for coarse reasoning

The condensed inference table is obtained by superimposing rows and columns of the original table. Formally, this corresponds to a logical OR-operation. With the condensed table, reasoning can be done in terms of entire neighborhoods – instead of single relations – at once. This, of course, is helpful only if the initial conditions for the reasoning process are given in terms of these neighborhoods. This is indeed frequently the case, both from initially incomplete knowledge and from previous inference steps, as has been shown in detail in [3].

More surprisingly, the neighborhood-based coarse composition table can be used for fine reasoning on the basis of the original thirteen interval relations: this is done by accessing the rows and columns of the table whose neighborhoods contain the individual relations involved. The largest neighborhood which is contained in all icons found at the resulting intersections represent the conclusion of the inference procedure. Formally, this is the conjunction of all disjunctions involved. Due to favorable orthogonality properties, no information gets lost in this process.

Accordingly, coarse and fine initial conditions can be combined.

### 4.2 Examples for Neighborhood-Based Reasoning

1) Coarse reasoning

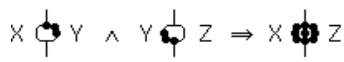


From the initial conditions describing the relations between events X and Y and between the events Y and Z, the relation between events X and Z is derived.

2) Fine reasoning

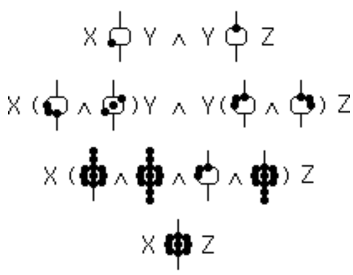


The initial conditions are expressed in terms of conjunctions of neighborhoods; conclusions are drawn by coarse reasoning; the intersection of the resulting neighborhoods yield the final result.

## 5. CONCLUSIONS AND OUTLOOK

We have sketched an approach to temporal and spatial reasoning based on Allen's interval calculus for qualitative temporal reasoning [1]. Our approach augments Allen's calculus by the notion of *conceptual neighborhood* which is elaborated in [3]. This notion supplies structures for abstracting knowledge and for higher-level reasoning.

The use of conceptual neighborhoods is motivated by physical properties which are considered essential for most cognitive operations in the temporal and spatial domains. The incorporation of neighborhood structure allows for a representation of generalized knowledge which turns out to be extremely useful both from a cognitive and a computational perspective. This generalized knowledge enables efficient higher-level reasoning in terms of cognitively meaningful concepts [3].

The higher-level reasoning processes exhibit a much more cognitively adequate behavior in that they treat incomplete knowledge as coarse knowledge rather than as underspecified knowledge; as a consequence, the required 'computational effort' decreases with decreasing knowledge

Another consequence of the incorporation of neighborhood information is the new role of uncertainty: due to the monotonicity properties of the temporal and spatial domains, neighboring initial conditions result at worst in neighboring consequences; thus, small uncertainties in the initial conditions do not cause drastically wrong conclusions. The neighborhood concept also allows for a cognitively more adequate treatment of fuzzy knowledge [4].

On a different dimension, there exists a favorable complexity-theoretical result which applies to our neighborhood structure. Vilain and Kautz [5] have shown that computing the consequences of temporal assertions in Allen's framework is computationally intractable. Nökel [6] has investigated a fragment of Allen's full algebra, namely the algebra of convex relations, which is closely related to the 'neighborhood algebra' presented here. He showed that in the algebra of convex relations global consistency can be verified in polynomial time. This result is expected to hold at least for the specific neigborhood structure used in the present paper.

The approach presented is suitable both for sequential and parallel processing. In fact, it is possible to utilize much more structure in the composition table (Fig. 4): the table of originally 169 entries can be compressed down to 7 entries by exploiting a variety of symmetries [3]. This result is particularly useful for dealing with higher dimensions [7] and for hardware realizations of temporal or spatial inference engines.

## ACKNOWLEDGEMENTS

The research reported in this paper was supported by the Deutsche Forschungsgemeinschaft and by Siemens AG. I also acknowledge discussions with Wilfried Brauer, Jerry Feldman, Daniel Hernández, Kerstin Schill, and Kai Zimmermann.

## REFERENCES

[1] Allen, J., Maintaining knowledge about temporal intervals. CACM 26 (1983), 11, 832-843.

[2] Freksa, C., Qualitative spatial reasoning. Appears in: Mark, D.M., Frank, A.U. (eds.) Cognitive and linguistic aspects of geographic space, Kluwer, Dordrecht.

[3] Freksa, C., Temporal reasoning based on semi-intervals. ICSI TR-90-016, International Computer Science Institute, Berkeley 1990.

[4] Freksa, C. *in preparation*

[5] Vilain, M., Kautz, H., Constraint propagation algorithms for temporal reasoning. Proc. AAAI-86, 377-382.

[6] Nökel, K., Convex relations between time intervals. 5. Österr. Artificial-Intelligence-Tagung, 298-302, Springer-Verlag, Berlin 1989.

[7] Güsgen, H.W., Spatial reasoning based on Allen's temporal locic. ICSI TR-89-045, International Computer Science Institute, Berkeley 1989.