

Generating Spatial Descriptions for Cross-modal References

Peter Wazinski

SFB 314, Department of Computer Science
University of Saarbrücken
D-6600 Saarbrücken, Germany
email: wazinski@cs.uni-sb.de

Abstract

We present a localisation component that supports the generation of cross-modal deictic expressions in the knowledge-based presentation system WIP. We deal with relative localisations (e.g., “The object to the left of object X.”), absolute localisations (e.g., “The object in the upper left part of the picture.”) and corner localisations (e.g., “The object in the lower right corner of the picture”). In addition, we distinguish two localisation granularities, one less detailed (e.g., “the object to the left of object X.”) and one more detailed (e.g., “the object above and to the left of object X.”). We consider corner localisations to be similar to absolute localisations and in turn absolute localisations to be specialisations of relative localisations. This allows us to compute all three localisation types with one generic localisation procedure. As elementary localisations are derived from previously computed composite localisations, we can cope with both localisation granularities in a computationally efficient way. Based on these primary localisation procedures, we discuss how objects can be localised among several other objects. Finally we introduce group localisations (e.g., “The object to left of the group of other objects.”) and show how to deal with them.

1 Introduction

The increasing amount of information to be communicated to users of complex technical systems nowadays makes it necessary to find new ways to present information. Neither the variety of all possible presentation situations can be anticipated nor it is further adequate to present the required information in a single communication mode, such as either text or graphics. Therefore, the automatic generation of multimodal presentations tailored to the individual user has become necessary. Current research projects in artificial intelligence like SAGE ([Roth *et al.*, 1990]), FN/ANDD ([Marks and Reiter, 1990]), COMET ([Feiner and McKeown, 1990])

and WIP ([Wahlster *et al.*, 1991a]) reflect the growing interest in this topic.

For the knowledge-based presentation system WIP, the task is the generation of a multimodal document according to the formal description of the communicative intent of the planned presentation and a set of generation parameters. The current scenario for WIP is the generation of instructions for using an espresso-machine. A typical fragment of an instruction manual for an espresso machine is shown in figure 1.




Figure 1: Fragment from an instruction manual

Cross-modal deictic expressions, e.g., “the lid” or “the knob in the middle,” help to establish the coreferentiality between the entities mentioned in the text and shown in the picture as well ([Wahlster *et al.*, 1991b]). The use of spatial relationships such as “the knob in the middle” simplifies the generation of referring expressions that have to identify a particular object in a picture. Obviously these spatial relationships cannot be computed in advance because they depend on the projection parameters for the picture, e.g., the viewpoint, which in turn themselves depend on the communicative intent of the document to be planned¹.

The localisation component described in this paper was developed in order to support the generation of cross-modal deictic referring expressions. All procedures are fully implemented and were recently integrated into the first WIP prototype. They are coded in Commor

¹ Even if the projection parameters are constant, it is not feasible to compute all possible relative localisations from a combinatoric point of view.

Lisp and run under Genera 8.0 on a MacIvory. A testbed called LOC-SYS was also developed: it allows the convenient generation and manipulation of rectangle scenes like the examples given in this paper.

Before we describe the methods which underlie the various localisation procedures, in the following section we present our views about localisation phenomena and introduce the terminology used in the rest of this paper.

2 Object Localisation

A lot of work has been done on 'object localisation' and its linguistic complement, 'spatial prepositions'. Wunderlich/Herweg ([Wunderlich, 1982], [Wunderlich and Herweg, forthcoming]) and Herskovits ([Herskovits, 1985]) provide linguistic approaches to the semantics of spatial prepositions. NL-systems like NAOS ([Neumann and Novak, 1986]), HAM-RPM ([Hahn *et al.*, 1980]), SWYSS ([Hußmann and Scheffé, 1984]) and CITYTOUR ([André *et al.*, 1985],[André *et al.*, 1986]) address various issues regarding computational aspects. Schirra ([Schirra, to appear 1992]) and Habel/Pribbenow ([Habel and Pribbenow, 1988],[Pribbenow, 1990]) also incorporate relevant work from cognitive psychology.

In our approach, we concentrate on the requirements for localising objects in pictures. We assume that the user can see the picture containing the objects to be localised and we do not deal with the problem of anticipating possibly wrong visualisations of the user in the case he/she cannot see the picture. We do not deal with possible intrinsic orientations of depicted objects (c.f. [Retz-Schmidt, 1988]) and assume the deictic reference frame of a common viewer (c.f. figure 5). Together with every localisation, we compute a so-called applicability degree from the interval [0..1]. The applicability degree is not only used to generate linguistic hedges (c.f. [Lakoff, 1972]) as in SWYSS or CITYTOUR, but also for selecting the 'best' localisation from a set of alternatives. The localisations computed on our system are two-dimensional localisations in the sense that they are based on the 2D-projection of a picture and not on its possible 3D-representation. In the rest of this section we will describe the localisation phenomena we take into account and introduce our terminology.

2.1 Relative and absolute localisations

The objects shown in part A of figure 2 can be localised as follows:




Figure 2: Localising objects in a picture

- (1) "Object A is on the right side of the picture."

- (2) "Object B is in the lower part of the picture."
- (3) "Object A is to the right of Object B."
- (4) "Object B is below Object A."

Sentences (1) and (2) are considered to contain **absolute localisations**: an object is localised by stating its *absolute* position in the picture. Sentences (3) and (4) are examples of **relative localisations**: an object is localised by stating its position *relative* to another object. The object to be located will be called the **primary object** (LO for short). The object that serves as reference for locating the primary object is called **reference object** (REFO for short).

How can we explain the similarity between absolute and relative localisations, between "on the right side of the picture" and "to the right of Object B"? Our hypothesis is:

Absolute localisations are specialisations of relative localisations in the sense that for absolute localisations the center of the picture functions as an implicit reference object.

Part B of figure 2 shows how the absolute localisation of part A can be explained as a relative localisation by assuming a circle-shaped center: "Object A is on the right side of the picture." is equivalent to "Object A is to the right of the center of the picture."

2.2 Elementary and composite localisations

Whereas the unambiguous localisations of the objects in figure 2 could be achieved by naming either the horizontal ("on the right side", "to the right of") or vertical relation ("in the lower part", "below"), figure 3 shows a situation in which it is necessary to give both the horizontal and vertical position of the object with respect to the reference object:




Figure 3: Elementary and composite localisations

In part A of figure 3, it is sufficient to describe object C as the object "to the right of" or "above" object A. But in part B, both descriptions would be ambiguous, because "to the right of" or "above" could refer to object D or B respectively. The only possibility to localise C unambiguously is to describe it as being "above and to the right" of A.

Localisations where either the horizontal or vertical relation is given will be called **elementary localisations**. If both relations are stated together, we will call it a **composite localisation**.

The localisation types introduced so far — absolute vs. relative and elementary vs. composite — are orthogonal. Therefore, an absolute or a relative localisation can be further subcategorized as being an elementary or a composite localisation.

Composite localisations cannot always be applied, e.g., in figure 2 object B cannot be localised as “the object in the lower left part of the picture.” Criteria for the applicability of composite localisations will not be examined further in this paper as this would lead to more complex questions, e.g., whether an object can be localised at all. A detailed discussion of these problems is given in [Wazinski, 1991].

2.3 The construction of the horizontal and vertical reference frame

One important feature of the localisation procedures is the division of the horizontal and vertical reference frame into three parts. The reason for this are ‘center’-localisations as shown in figure 4:




Figure 4: Center localisations

In all pictures, object A can be localised as the object “in the center of the picture.” In order to integrate this observation with the elementary vs. composite distinction we divided the horizontal and vertical dimension into three parts: ‘top’, ‘horizontal center’ and ‘bottom’ and ‘left’, ‘vertical center’ and ‘right’ respectively (c.f. figure 5). Under these conditions the ‘center’-localisation in the left part of figure 4 can be analysed as a composite (‘vertical center’, ‘horizontal center’)-localisation. For the picture in the middle it is an elementary ‘vertical center’-localisation and for the right one an elementary ‘horizontal center’-localisation. When transforming these different localisations into a surface string they all become the same: “in the center of the picture.”

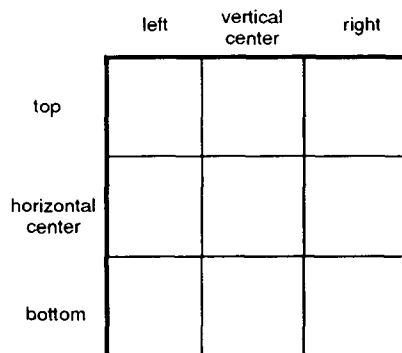


Figure 5: Horizontal and vertical reference frame

Figure 6 shows that it is also useful to adopt the this partition scheme for relative localisations: B would usually be described as the object “to the right of A” and C as the object “above and to the right of A.” With respect to the partition scheme a (‘right’, ‘top’)-localisation can be applied to C and a (‘right’, ‘horizontal center’)-localisation to B. The former matches exactly with the surface string. The latter can be matched with

“to the right of A” by assuming that the ‘center’-part of a composite localisation is a special part of a composite localisation that does not appear at the linguistic level.




Figure 6: Center localisations and relative localisations

2.4 Corner Localisations

An additional localisation type that can be used to localise objects in pictures is the **corner localisation**: if an object is placed in one of the four corner regions of the picture it can be localised as, e.g., “the object in the left upper corner of the picture.”

The difference between absolute composite localisations and corner localisations is illustrated in figure 7: While object B can be localised as being “in the lower right corner of the picture” it is not possible to use a corner localisation for A. In that case, only “in the left upper part of the picture” could be used. Therefore, we consider corner localisations to be more precise than absolute composite localisations, i.e., the applicability of a corner localisation implies the applicability of the corresponding absolute composite localisation but not vice versa.




Figure 7: Corner localisations vs. absolute composite localisations

3 Basic Localisation Procedures

In this section we present matrix-oriented localisation procedures for absolute and relative localisations. As mentioned in section 2.2, both the horizontal and vertical relation of the primary object are given in case of a composite localisation. This suggests that composite localisations are composed of elementary localisations. The procedures presented here, though, behave differently: for the sake of efficiency they compute the composite localisations first and derive the elementary localisations from these previously computed localisation results.

3.1 Absolute localisations

We approximate the center of the picture with a rectangle whose horizontal and vertical extension is one third of the horizontal and vertical extension of the picture. Figure 8 shows the construction of the horizontal and vertical reference system according to the rectangular center region.




Figure 8: The construction of the horizontal and vertical reference system

Before describing the evaluation function for composite localisations, we give a few definitions:

- The horizontal reference system is abbreviated by $XLOC = \{\text{left, x-center, right}\}$, the vertical one by $YLOC = \{\text{top, y-center, bottom}\}$. Composite localisations are denoted by $CLOC = XLOC \times YLOC$. Both reference systems together are described with $ULOC = XLOC \cup YLOC$.
- The constant **CENTER** denotes the center rectangle of a given picture.
- **POLY** denotes the set of all polygons that can appear in a picture. For given polygons P_1 and P_2 the associative and commutative operator \cap , $_ \cap _ : POLY \times POLY \mapsto POLY$ computes the intersection polygon. The empty polygon is denoted by P_\emptyset . The following holds: $\forall P \in POLY : P_\emptyset \cap P = P \cap P_\emptyset = P_\emptyset$.
- The function **PR** (Partial Rectangle), $PR : CLOC \times POLY \mapsto POLY$, computes the rectangle corresponding to a given composite localisation and the rectangle partition of the picture induced by a given polygon. For example $PR((\text{left, top}), \text{CENTER})$ computes the upper left rectangle according to the partition scheme shown in figure 8.
- \mathfrak{R} denotes the set of the real numbers. Given a polygon P , the function f , $f : POLY \mapsto \mathfrak{R}$ computes the area of a polygon. It is $f(P_\emptyset) = 0$.

The applicability degree of a composite localisation evaluates how good the position of the object in question is described by that particular localisation. We define the applicability degree as the portion of the area of the object that lies in the rectangle of the picture that corresponds to the composite localisation and the rectangle partition of that picture. Thus we can define A_c as follows:

$$A_c : CLOC \times POLY \mapsto \mathfrak{R}$$

$$A_c(l, LO) = \frac{f(p)}{f(LO)}$$

with $p = PR(l, \text{CENTER}) \cap LO$

For object **LO** in figure 9, the above definition yields the following results:

$$\begin{aligned} A_c((\text{left, top}), LO) &= 1/12, \\ A_c((\text{x-center, top}), LO) &= 1/6, \\ A_c((\text{left, y-center}), LO) &= 1/4, \\ A_c((\text{x-center, y-center}), LO) &= 1/2. \end{aligned}$$

For all other $l \in CLOC$ we have $A_c(l, LO) = 0$ because $f(p) = f(P_\emptyset) = 0$.




Figure 9: Computing absolute localisations

For elementary localisations we adopted an analogous definition: the applicability degree A_e of an elementary localisation is determined by the portion of the area of the object that lies in the corresponding row or column of the picture. As already mentioned at the beginning of this section we can write A_e in terms of A_c :

$$A_e^x : XLOC \times POLY \mapsto \mathfrak{R}$$

$$A_e^x(l_x, LO) = \sum_{l_y \in YLOC} A_c((l_x, l_y), LO)$$

$$A_e^y : YLOC \times POLY \mapsto \mathfrak{R}$$

$$A_e^y(l_y, LO) = \sum_{l_x \in XLOC} A_c((l_x, l_y), LO)$$

$$A_e : ULOC \times POLY \mapsto \mathfrak{R}$$

$$A_e(l, LO) = \begin{cases} A_e^x(l, LO) & \text{if } l \in XLOC \\ A_e^y(l, LO) & \text{if } l \in YLOC \end{cases}$$

A_e^x and A_e^y compute the applicability for the horizontal and vertical dimension by summing up the applicability degrees of the corresponding composite localisations. They are combined in A_e order to have a function that is defined on both dimensions, i.e., $ULOC$.

With respect to figure 9 we get:

$$\begin{aligned} A_e(\text{top}, LO) &= A_c((\text{left, top}), LO) + A_c((\text{x-center, top}), LO) = 1/4, \\ A_e(\text{y-center}, LO) &= A_c((\text{left, y-center}), LO) + A_c((\text{x-center, y-center}), LO) = 3/4, \\ A_e(\text{left}, LO) &= A_c((\text{left, top}), LO) + A_c((\text{left, y-center}), LO) = 1/3 \text{ and} \\ A_e(\text{x-center}, LO) &= A_c((\text{x-center, top}), LO) + A_c((\text{x-center, y-center}), LO) = 2/3. \end{aligned}$$

As argued in paragraph 2.4 corner localisations are similar to composite ('left'/'right', 'top'/'bottom')-localisations, but less general. This property can be modelled by corner regions that are smaller than the corner regions for absolute localisations. In turn, these corner regions correspond to a larger center as shown in figure 10. Thus we can compute corner localisations just by changing the size of the center.




Figure 10: The relation between corner and center regions

Instead of $1/3$ as for absolute localisations we take $4/5$ of the horizontal and vertical extension of the picture for the extended center.

3.2 Relative localisations

The localisation procedure for relative localisations is similar to the one for absolute localisations. One major difference is that now the construction of the horizontal and vertical reference frame is done with respect to a given reference object and not to the implicit assumed center of the picture (c.f. figure 11). The second difference concerns the computation of the applicability degree: for relative localisations, not only the portion of an area is taken into account, but also the distance between the primary object and the reference object.




Figure 11: The construction of the reference frame for relative localisations

The basic idea for the evaluation of the distance between primary object and reference object is adopted from the CITYTOUR system: first we compute the center of gravity for the primary object. Then we determine its coordinates with respect to the reference system established by the reference object. Finally we use these coordinates for the computation of the applicability degree. Figure 12 illustrates the various factors that affect the applicability of an 'above'-localisation:

1. The applicability degree decreases with an increasing vertical distance. In Part A of figure 12 the ap-

plicability degree for " P_1 is above REFO" is greater than for " P_2 is above REFO."

2. The applicability degree decreases with an increasing horizontal distance. In Part B the applicability degree for " P_3 is above REFO" is greater than for " P_4 is above REFO."
3. If the horizontal and vertical distances increase by the same amount, then the applicability degree decreases more with the increasing horizontal distance than with the increasing vertical distance. This is shown in Part C: the applicability degree for " P_6 is above REFO" is greater than for " P_7 is above REFO", although the vertical distance between P_5 and P_6 and the horizontal distance between P_5 and P_7 are equal.




Figure 12: Evaluating the distance of a point

Let $eval$ denote the function that evaluates the distance between a point and a rectangle according to the criteria mentioned above. Let further $POINT$ denote the set of all points within a picture and $RECT \subseteq POLY$ the set of all rectangular polygons. Then the signature of $eval$ can be written as²:

$$eval : CLOC \times POINT \times RECT \mapsto \mathfrak{R}$$

Now we are almost able to define the function A_c , which computes the applicability degree of a composite localisation. Let $CG, CG : POLY \mapsto POINT$, compute the center of gravity for a polygon and let further $SR, SR : POLY \mapsto RECT$, compute the smallest surrounding rectangle for a polygon. Then the applicability degree A_c of a composite localisation can be defined as:

$$A_c : CLOC \times POLY \times POLY \mapsto \mathfrak{R}$$

$$A_c(l, LO, REFO) = w \cdot eval(l, CG(p), SR(REFO))$$

$$\text{with } p = PR(l, REFO) \cap LO$$

$$w = \frac{f(p)}{f(LO)}$$

²In reality $eval$ is slightly more complicated because it maps into $\mathfrak{R} \times \mathfrak{R}$ and not only into \mathfrak{R} . The reason for this is that the different evaluation of increasing vertical and horizontal distances can result in different evaluations for points to which both a horizontal or vertical localisation can be applied. E.g., P_7 in figure 12 would get a different evaluation for an 'above'- than for a 'right of'-localisation. Therefore, these two values would be grouped to a tuple. For the computation of an elementary localisation $l \in YLOC$ we would sum up the first component of the tuple. If $l \in YLOC$, we take the second component. We abstract from this detail in order to make the principle of the procedure clearer.

p is the part of the primary object that lies in the rectangle corresponding to the composite localisation l . The factor w weighs the result of $eval$ according to the portion of the area of the primary object that lies in the rectangle corresponding to l .

Now the definition of A_e , the applicability degree for an elementary localisation, can be given in terms of A_c again:

$$A_e^x : \text{XLOC} \times \text{POLY} \times \text{POLY} \mapsto \mathfrak{R}$$

$$A_e^x(l_x, \text{LO}, \text{REFO}) = \sum_{l_y \in \text{YLOC}} A_c((l_x, l_y), \text{LO}, \text{REFO})$$

$$A_e^y : \text{YLOC} \times \text{POLY} \times \text{POLY} \mapsto \mathfrak{R}$$

$$A_e^y(l_y, \text{LO}, \text{REFO}) = \sum_{l_x \in \text{XLOC}} A_c((l_x, l_y), \text{LO}, \text{REFO})$$

$$A_e : \text{ULOC} \times \text{POLY} \times \text{POLY} \mapsto \mathfrak{R}$$

$$A_e(l, \text{LO}, \text{REFO}) = \begin{cases} A_e^x(l, \text{LO}, \text{REFO}) & \text{if } l \in \text{XLOC} \\ A_e^y(l, \text{LO}, \text{REFO}) & \text{if } l \in \text{YLOC} \end{cases}$$

This means that the applicability degree A_e for a primary object LO is the sum of the composite localisations for the corresponding row or column of the reference frame.

For figure 13 we get the following results:

$$\begin{aligned} & A_c((x\text{-center}, \text{top}), \text{LO}, \text{REFO}) \\ &= \frac{1}{3} \text{eval}((x\text{-center}, \text{top}), P_1, \text{SR}(\text{REFO})) \\ &= \frac{1}{3} * 0.7 = 0.23 \end{aligned}$$

$$\begin{aligned} & A_c((\text{right}, \text{top}), \text{LO}, \text{REFO}) \\ &= \frac{2}{3} \text{eval}((\text{right}, \text{top}), P_2, \text{SR}(\text{REFO})) \\ &= \frac{2}{3} * 0.65 = 0.43 \end{aligned}$$

$$\begin{aligned} & A_c(l, \text{LO}, \text{REFO}) \\ &= 0 \text{ as for all other } l \in \text{CLOC} : w = \frac{f(p)}{f(\text{LO})} = 0 \end{aligned}$$

$$\begin{aligned} & A_e(\text{top}, \text{LO}, \text{REFO}) \\ &= A_c((x\text{-center}, \text{top}), \text{LO}, \text{REFO}) + \\ & A_c((\text{right}, \text{top}), \text{LO}, \text{REFO}) = 0.66 \end{aligned}$$

$$\begin{aligned} & A_e(\text{right}, \text{LO}, \text{REFO}) \\ &= A_c((\text{right}, \text{top}), \text{LO}, \text{REFO}) = 0.43 \end{aligned}$$

$$\begin{aligned} & A_e(x\text{-center}, \text{LO}, \text{REFO}) \\ &= A_c((x\text{-center}, \text{top}), \text{LO}, \text{REFO}) = 0.23 \end{aligned}$$

4 A generic localisation procedure for absolute and relative localisations

The similarities between the localisation procedures discussed in the previous section allow us to design one generic localisation procedure that can be specialised to a procedure for absolute, relative or corner localisations.

Given the primary object LO and the reference object REFO the first step is to determine the 3×3 matrix M^\cap , which contains the intersection polygons of LO and the partial rectangles in the picture with respect to REFO. For relative localisations, REFO varies, for absolute localisations and corner localisations the parameter is set to either the normal or the extended center area (c.f. section 3.1). Thus, for $x \in \text{XLOC}$, $y \in \text{YLOC}$ we compute

$$M_{x,y}^\cap = \text{PR}((x, y), \text{REFO}) \cap \text{LO}$$




Figure 13: Computing relative localisations

The second step is the computation of the evaluation matrix M^A , which contains the applicability degrees of the composite localisations. The computation requires a function E , $E : \text{POLY} \times \text{POLY} \times \text{POLY} \mapsto \mathfrak{R}$. E corresponds exactly to the function A_c for absolute and relative localisations in section 3.1 and 3.2. The only difference results from the previous computation of M^\cap : the subexpression $p = \text{PR}((x, y), \text{REFO}) \cap \text{LO}$ is factored from A_c and therefore computed only once.

$$M_{x,y}^A = E(M_{x,y}^\cap, \text{LO}, \text{REFO})$$

The third step is the computation of the elementary localisations. The vector \vec{X} contains the evaluations of the horizontal localisations and \vec{Y} the ones for the vertical localisations:

$$\begin{aligned} \vec{X}_x &= \sum_{y \in \text{YLOC}} M_{x,y}^A \\ \vec{Y}_y &= \sum_{x \in \text{XLOC}} M_{x,y}^A \end{aligned}$$

This means that we have $\vec{X}_l = A_c(l)$ for $l \in \text{XLOC}$ and $\vec{Y}_l = A_e(l)$ for $l \in \text{YLOC}$.

Finally, we can determine the best composite and elementary localisation and their applicability degrees by computing the maximum value of M^A and \vec{X} or \vec{Y} respectively.

For figure 13 we get

$$M^A = \begin{pmatrix} 0 & 0.23 & 0.43 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$\vec{X} = (0 \ 0.23 \ 0.43)$ and $\vec{Y} = (0.66 \ 0 \ 0)$. The best composite localisation is “(right, top)” with applicability degree 0.43. The best elementary localisation is “top” with applicability degree 0.66.

5 Localising objects in a complex scene

In the previous sections we considered pictures with a minimal number of objects. In order to deal with more complex object configurations the localisation procedures presented above have to be extended. The new task is no longer “Localise LO with respect to REFO!” but “Given a set of REFO candidates, choose the best one for LO!”

In order to reduce the search space for REFO candidates, first a kind of ‘between’-test is applied to the set of possible reference objects. The idea behind this test is that an exclusion procedure based on simple geometric overlapping tests can be performed more efficiently than a comparison of applicability degrees that have to be computed by the rather complex localisation procedures. An example is given in figure 14: When searching for a suitable reference object for object A in figure 14, object D would be ruled out because object B is found in the ‘between’-area of A and D.

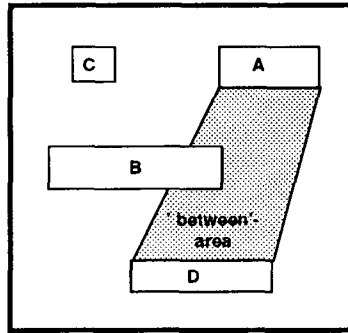


Figure 14: Search space reduction for complex object configurations

The determination of the best reference object raises the problem of ambiguity. Not only is the applicability degree of a localisation important, but also whether the use of the reference object would result in an ambiguous localisation. In that case, a different reference object has to be chosen. If all possible localisations are ambiguous, then the particular object cannot be localised at all. E.g., in Part A of figure 15 object D could be localised as being either ‘above A’ or ‘to the right of C.’ But the first localisation is ambiguous because both, C and D, are ‘above A.’




Figure 15: Ambiguous reference objects

With respect to elementary and composite localisations we distinguish three cases of ambiguity:

1. In Part A of figure 15, the localisation of object C or D would be ambiguous with respect to A because for both objects the composite localisations, (x-center, top), are equal.
2. In Part B a composite localisation cannot be applied to object D (neither ‘D is above and to the right of A’ nor ‘D is immediately above A’ are adequate) and its elementary localisation, ‘top’, is part of the composite localisation, (x-center, top), of object C.
3. In Part C a composite localisation can be applied

neither to C nor to D and their elementary localisations, ‘top’, are equal.

6 Localising Groups of Objects

Control knobs and switches are often grouped together in a control panel in order to provide for easier operation of technical devices. Moreover spatially adjacent objects can also be grouped as one perceptual unit according to the ‘law of the good gestalt’ in Gestalt psychology ([Murch and Woodworth, 1978]). Thus the possibility to generate localisations with respect to a given group structure is necessary for the ‘naturalness’ of a localisation. Besides this, group localisations are also useful if the objects in the immediate neighbourhood of the primary object have exactly the same properties (c.f. [Wahlster *et al.*, 1978]). In this case, the primary object can be localised with respect to its group and has not to be localised with respect to the whole scene, which could have resulted in an ambiguous localisation.

For our localisation procedures this means that groups can function as a reference object as well as a primary object. In addition, objects can be localised absolutely with respect to the group they are contained in. In figure 16 object B would be localised as the object ‘to the right of the triangles.’ Vice versa we can say ‘The triangles to the left of object B’ and we can localise object A as being ‘the upper left of the triangles that are to the left of B.’




Figure 16: Group localisations

The last example also illustrates the hierarchical character of group localisations: An object can be localised absolutely within a group. This group might be localised again within a surrounding group or — if there is none — this group can be localised relatively with respect to another (group of) object(s).

The algorithm for group localisations cannot detect group hierarchies. Instead it expects a tree representation of the group hierarchy as an input. The output consists of two parts: According to the depth of the group tree the algorithm computes a chain of absolute localisations. In addition the outermost surrounding group of the primary object is localised relatively to an optional (group of) reference object(s).

7 Conclusions

We have introduced a unifying approach for absolute relative and corner localisations of objects in pictures. In addition, the use of a special partition scheme for the reference frame of a preposition allows us to deal with two different localisation granularities for absolute and relative localisations. By defining the evaluation function

for elementary localisations in terms of the evaluation functions for the corresponding composite localisations, we have been able to design one procedure that handles all three localisation types and both localisation granularities efficiently. Furthermore, we have given a solution to the problem of localising an object within a complex configuration on the basis of this localisation procedure. Finally, we have shown how our system deals with group localisations.

References

- [André *et al.*, 1985] E. André, G. Bosch, G. Herzog, and T. Rist. CITYTOUR – Ein natürlichsprachliches Anfragesystem zur Evaluierung räumlicher Präpositionen. Abschlußbericht des Fortgeschrittenenpraktikums, Department of Computer Science, University of Saarbrücken, 1985.
- [André *et al.*, 1986] E. André, G. Bosch, G. Herzog, and T. Rist. Characterising Trajectories of Moving Objects using Natural Language Path Descriptions. In *Proc. of the 7th ECAI*, pages 1–8, 1986.
- [Feiner and McKeown, 1990] S. K. Feiner and K. R. McKeown. Coordinating Text and Graphics in Explanation Generation. In *Proc. 8th AAAI*, pages 442–449, 1990.
- [Habel and Pribbenow, 1988] C. Habel and S. Pribbenow. Gebietskonstituierende Prozesse. LILOG-Report 18, IBM Germany, 1988.
- [Hahn *et al.*, 1980] W. v. Hahn, W. Hoepfner, A. Jameson, and W. Wahlster. The Anatomy of the Natural Language Dialogue System HAM-RPM. In L. Bolc, editor, *Natural Language Based Computer Systems*, pages 119–254. München: Hanser, 1980.
- [Herskovits, 1985] A. Herskovits. Semantics and Pragmatics of Locative Expressions. *Cognitive Science*, 9:341–378, 1985.
- [Hußmann and Scheffe, 1984] M. Hußmann and P. Scheffe. The Design of SWYSS, a Dialogue System for Scene Analysis. In L. Bolc, editor, *Natural Language Communication with Pictorial Information Processing*. München: Hanser McMillan, 1984.
- [Lakoff, 1972] G. Lakoff. Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts. In J.N. Levi and G.C. Phares, editors, *Papers from the 8th regional Meeting of the Chicago Linguistics Society*, pages 183–228. University of Chicago, Chicago, IL, 1972.
- [Marks and Reiter, 1990] J. Marks and E. Reiter. Avoiding Unwanted Conversational Implicatures in Text and Graphics. In *Proc. 8th AAAI*, pages 450–455, 1990.
- [Murch and Woodworth, 1978] G.M. Murch and G.L. Woodworth. *Wahrnehmung*. Stuttgart: Kohlhammer, 1978.
- [Neumann and Novak, 1986] B. Neumann and H.-J. Novak. NAOS: Ein System zur natürlichsprachlichen Beschreibung zeitveränderlicher Szenen. *Informatik Forschung und Entwicklung*, pages 83–92, 1986.
- [Pribbenow, 1990] S. Pribbenow. Interaktion von propositionalen und bildhaften Repräsentationen. In C. Habel and C. Freksa, editors, *Repräsentation und Verarbeitung räumlichen Wissens*, pages 156–174. Berlin: Springer, 1990.
- [Retz-Schmidt, 1988] G. Retz-Schmidt. Various Views on Spatial Prepositions. *AI Magazine*, 9(2):95–105, 1988.
- [Roth *et al.*, 1990] S. Roth, J. Mattis, and X. Mesnard. Graphics and Natural Language as Components of Automatic Explanation. In J. W. Sullivan and S. W. Tyler, editors, *Intelligent User Interfaces*, pages 207–239. Reading, MA: Addison Wesley, 1990.
- [Schirra, to appear 1992] J. Schirra. A Contribution to the Reference Semantics of Spatial Prepositions: The Visualization Problem and its Solution in VITRA. In *Proceedings of the IAI Workshop "On the Semantics of Prepositions in Natural Language Processing"*. Mouton, de Gruyter, to appear 1992. Also available as Technical Report 75, SFB 314, Department of Computer Science, University of Saarbrücken.
- [Wahlster *et al.*, 1978] W. Wahlster, A. Jameson, and W. Hoepfner. Glancing, Referring and Explaining in the Dialogue System HAM-RPM. *American Journal of Computer Linguistics, Microfiche 77*, pages 53–67, 1978.
- [Wahlster *et al.*, 1991a] W. Wahlster, E. André, S. Bandyopadhyay, W. Graf, and T. Rist. WIP: The Coordinated Generation of Multimodal Presentations from a Common Representation. In O. Stock, J. Slack, and A. Ortony, editors, *Computational Theories of Communication and their Applications*. Berlin: Springer, 1991.
- [Wahlster *et al.*, 1991b] W. Wahlster, E. André, W. Graf, and T. Rist. Designing Illustrated Text: How Language Production is Influenced by Text and Graphics. In *Proc. 5th Conf. of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 8–14, 1991.
- [Wazinski, 1991] P. Wazinski. Objektlokalisierung in graphischen Darstellungen. Master's thesis, Universität Koblenz-Landau, Abt. Koblenz/DFKI Saarbrücken, 1991.
- [Wunderlich and Herweg, forthcoming] D. Wunderlich and M. Herweg. Lokale und Direktionale. In A. v. Stechow and D. Wunderlich, editors, *Handbuch der Semantik*. Königstein Ts.: Athenäum Verlag, forthcoming.
- [Wunderlich, 1982] D. Wunderlich. Sprache und Raum. *Studium Linguistik*, 12:1–19, 1982.