On Spatial Reference Frames in Qualitative Motion Representation

Alexandra Musto (musto@in.tum.de)
Institut für Informatik
Technische Universität München

Abstract. In the discussion of spatial representations in technical as well as in biological systems, frames of reference are of great importance. In this paper, we focus on spatial frames of reference in motion representation and stress that we have to introduce a distinction between frames of reference in measurement and in representation. We discuss some problems that arise in qualitative motion representation when measuring the course of motion in the egocentric and the allocentric frame of reference. Then we present a method to automatically generate route descriptions from these qualitative representations where granularities and frames of references can be mixed.

Keywords: Qualitative Reasoning, Spatial Reasoning, Motion Representation

1. Introduction

When talking about spatial representations, we have also to talk about reference frames. Especially in motion representation, a careful look at in which reference frame we actually operate has to be taken. Imagine the case of locomotion: one could be tempted to say that the only reference frame for locomotion would be the egocentric one. But is this also the case for a co-pilot at the rallye Granada-Dakar, who navigates with a GPS system?

To make things even more complicated, we want to point out that in motion representation we have to distinguish two different levels at which different reference frames may be used: The lower level is the level of *measurement* of the motion event, the level above is the level of *representation*. Imagine again you were a participant in the rallye Granada-Dakar and have measured your course of motion of the day's leg through a GPS-Track, that is in an allocentric frame of reference. In analyzing it you could easily transform this track into a representation in the egocentric frame of reference, and make a statement like "In this village we made a mistake and turned to the right instead of turning to the left". In this case, a motion event that was measured in the allocentric frame of reference is represented in the egocentric one. So, in motion representation, it is not only important which frame of reference is used, but also on which level it is used. However, there are some difficulties in collecting qualitative data in the egocentric frame of reference and in changing reference frames on the different levels that we want to discuss in further detail.

From a cognitive point of view, Tversky (1993) advocates that people's spatial mental models also use only two basic perspectives: Locating elements relative to one another from a point of view or locating an element to a higher order environmental feature or reference frame. The first corresponds to an egocentric frame of reference and a route perspective, the second to an allocentric frame of reference and a survey perspective. Regarding these two basic perspectives, the mental models people use seem to be more abstract than either, allowing switches of perspective and inference in both of them. Furthermore, people often mix both perspectives when producing descriptions of the environment. We discuss some suggestions on this mix of reference frames for motion representation on technical systems on different granularities.

2. Definitions

First, we want to clarify some notions important with respect to motion representation. We make two assumptions that seem reasonable in the context of (loco-)motion: We deal with objects that have an intrinsic front side and normally move forward with respect to their front side (and not to the side like, e.g., a crab).

Most of the following definitions are in accordance with (Klatzky, 1998). See also figure 1 for illustration.

The axis of orientation is the perpendicular with respect to the front side of the object, directed against this front side.

The heading or *orientation* of an object in space is the angle between some directed reference axis external to the object ($\rightarrow$ *allocentric* heading) and the object's axis of orientation.

The bearing from point A to B is the angle between the directed reference axis and a line from A to B. In an $\rightarrow$ *egocentric* frame of reference, the reference direction is ego's heading.

The direction of locomotion at a position $n-1$ is the angle between the heading of ego at position $n-1$ and position n. This is the egocentric bearing of position n with respect to ego and its heading at position $n-1$. The direction at position $n-1$ equals to the new heading of the object at position n.

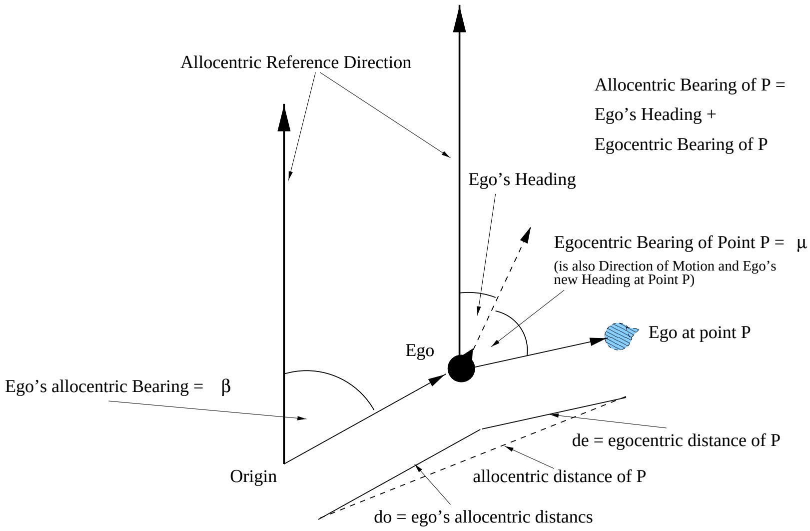

Figure 1. Illustration of basic concepts

Frames of reference Different notions for reference frames can be found in the literature, e.g. egocentric and allocentric like in (Klatzky, 1998; Brewer and Pears, 1993)[1], or intrinsic, extrinsic and deictic like in (Clementini et al., 1997). In this paper, the following terms are understood to mean:

- allocentric: In this frame of reference a fixed coordinate system is imposed by some external factors like in geographic space. A point P in the allocentric representation has coordinates (d_o, β), where d_o is the distance from the origin and β is the bearing from the origin, defined with respect to the reference direction (see Klatzky (1998)).
- egocentric: There is no external coordinate system. A point P in the egocentric representation has coordinates (d_e, μ), where d_e is the egocentric distance of the point and μ is its egocentric bearing (see Klatzky (1998)). Egocentric bearing is defined with respect to an intrinsic axis of orientation that is imposed by ego's physical configuration. In the case of locomotion, we take this frame of reference with us with every step we take and its allocentric heading, bearing and distance changes constantly.
- extrinsic: synonymous with allocentric frame of reference.

[1] The latter, however, argue that the better names for these reference frames according to their use in the literature would be body-centered and environment-centered.

- intrinsic: The coordinate system is determined by some inherent characteristics of the reference object, like its topology, size, or shape. E.g. a church yard has an intrinsic front — the side where the church stands. The egocentric FoR is a special case of an intrinsic FoR, where the reference object is Ego. On the other hand, regarding Ego in the curchyard, the intrinsic frame of reference is a special case of an allocentric frame of reference. Therefore, the intrinsic reference frame works like the allocentric reference frame with its own reference direction. That is, a point P in the intrinsic representation has coordinates (d_{io}, β_i), where d_{io} is the distance from the intrinsic origin and β_i is the bearing from the intrinsic origin, defined with respect to the intrinsic reference direction. The intrinsic frame of reference can be embedded in a bigger, allocentric FoR. Then it has its own allocentric bearing β_a. See fig. 2 for illustration. The meaning of this for motion representation will be discussed in section 4.
- deictic: A frame of reference is imposed by an external observer. This includes problems like distortion through perspective and locomotion of the observer.

In our approach in section 4, we use three different frames of reference, namely egocentric, allocentric and intrinsic, as defined above. This distinction happens to correspond with the one Levinson (1996) made of relative (in our terminology: egocentric), intrinsic, and absolute (in our terminology: allocentric) reference frames.

In the context of qualitative motion representation we feel the need to introduce another distinction with respect to the use of reference frames: the distinction between the use of a special frame of reference in **measurement** vs. the use of a frame of reference in **representation**.[2] They form two independent layers in the task of motion representation, where appropriate frames of reference can be chosen:

We can measure motion data from an egocentric point of view without an allocentric, fixed coordinate system, e.g. if we count steps and memorize the angles when turning. Then the frame of reference in measurement is egocentric. Then we can represent this data also in an egocentric frame of reference; the

[2] Actually this distinction is conceptually important also in quantitative motion representation, but there egocentric measurement is not so problematic since angles are not mapped onto intervals. Nevertheless, due to the lack of an external coordinate system, egocentric measurement is alway erroneous, since small errors accumulate. This should be kept in mind when dealing with egocentrically measured data.

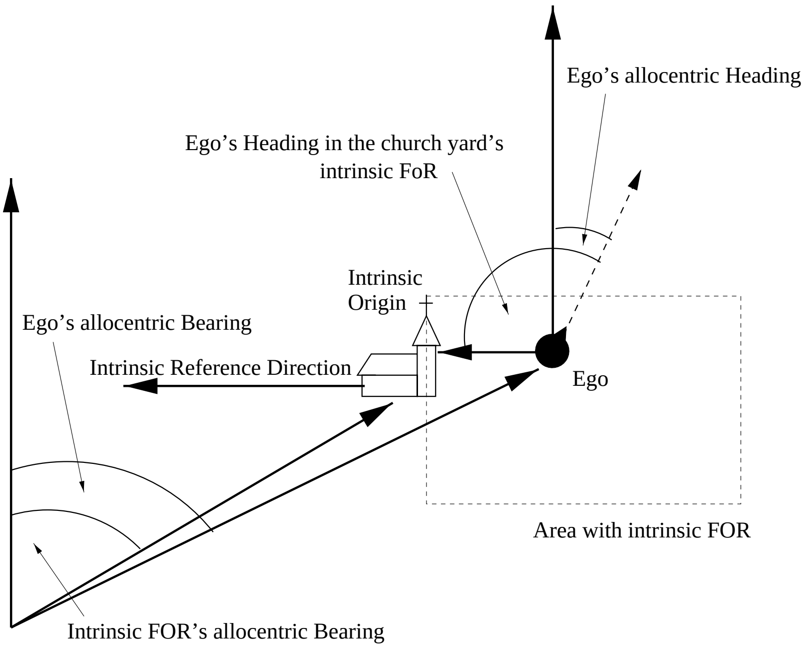

Figure 2. Interplay of an intrinsic and an allocentric Frame of Reference

frame of reference in representation is then egocentric, too. Or we can try to map the egocentrically measured data into a global coordinate system and represent the course of motion in an allocentric frame of reference in representation.

On the other hand, we can measure motion data from an allocentric point of view with a fixed coordinate system, e.g. when we store the GPS-Track of a certain route we traveled. We then are able to transform the measured data into an egocentric frame of reference, e.g. for a route description. Then, the farme of reference in measurement is allocentric, and the frame of reference in representation is egocentric. This is e.g. done in (Musto et al., 1998).

Discretizations of Space and Time Since space and time are inherently continuous domains, we have to discretize space and time somehow to achieve a qualitative representation. For the representation of positional change we need only two components, namely distance and direction. These have to be discretized into intervals, e.g. like in the figures 3 and 4.

Time can be discretized using a fixed scan rate, which seems reasonable for technical as well as for biological systems (see (Pöppel

Figure 3. A discretization of the distance domain

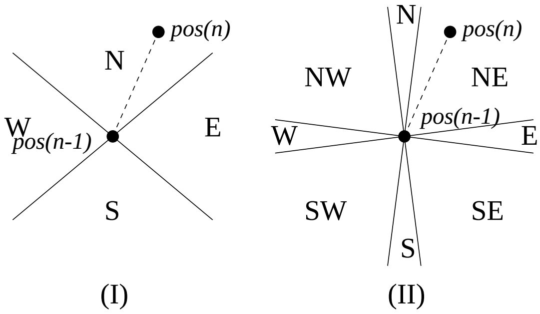

Figure 4. Two discretization of the direction domain

and Schill, 1995)). Another possibility is an event-driven discretization of time. We won't discuss this further, since this paper deals with spatial reference frames.

Qualitative Motion Vectors (QMV'S) are vectors that describe the motion of an object from position $n-1$, measured at measurement point $n-1$, to position n, measured at measurement point n, where the vector components are some qualitative descriptions of, e.g., direction, distance and speed of the object when moving from point $n-1$ to point n.

3. Frames of Reference in Measurement

There are two major problems in the description of movement by qualitative intervals, no matter whether in the egocentric or allocentric frame of reference in measurement.

1. Two similar movement vectors, e. g. (18cm, 43°) and (20cm, 47°), may result in totally different QMV's. E. g., corresponding to figure 4 (I), the direction of the first would be expressed by `north`, the direction of the second by `east`. With distances the same situation arises when the border between two distance intervals lies between 18 and 20.

2. Since many directions are mapped into the same direction intervals, small changes in direction between two movement vectors may not be noticed at all.

How serious each of these problems is depends on different items: The first problem is serious if we want to judge over similarity, since very similar courses of motion might result in very unsimilar representations. The second problem is serious for measuring locomotion without falling back on an extrinsic coordinate system, which we now want to discuss further.

In (Musto et al., 1998), a qualitative representation of the course of motion in an allocentric frame of reference in measurement and, first, allocentric frame of reference in representation by means of qualitative motions vectors (QMV's) is described. The formalization is done relatively straight forward:

A course of motion is measured with a fixed scan rate and qualitatively represented by only two components, namely distance and direction. Speed can be derived from the distance the moving object covered in a single scan cycle. Space is discretized sharply in the domain of distance and direction like in figure 3 and 4(I), following the suggestions of (Clementini et al., 1997). So, a course of motion is represented as sequence of qualitative motion vectors, e.g.

`<close east>`^{5} `<close north>`^{2} `<close west>`^{3} `<close south>`^{1} `<medium-dist south>`^{1} `<medium-dist east>`^{1}.[3]

The indices indicate for how many scan cycles no change in distance and direction occured.

If the course of motion is measured in an allocentric frame of reference in measurement like here, switches between allocentric and egocentric frame of reference in representation in the domain of direction are possible without loss of information. A representation in the egocentric frame of reference in representation of the above QMV Sequence reads:

`<close forward>`^{5} `<close left>`^{2} `<close left>`^{3} `<close left>`^{1} `<medium-dist forward>`^{1} `<medium-dist left>`^{1}.

Problem (2) is no real problem in an allocentric frame of reference in measurement, since, if small changes accumulate to a big change, this is eventually noticed when the border to another direction region is crossed.

Unfortunately, the same is not true when using an egocentric frame of reference in measurement. Since there is no external, fixed,

[3] Mapping of distance combined with the counter-information (e.g. 5 times "close" yields "medium-dist") and speed in qualitative intervals yields a representation like this: `<medium-dist east slow> <close north slow> <close west slow> <close south slow> <medium-dist south medium-vel> <medium-dist east medium-vel>` (vel = velocity).

and relatively unchanging frame of reference when measuring locomotion, the direction gridlock has to be newly aligned in each scan cycle, depending on the new intrinsic orientation of the moving person or robot. Therefore, small changes in the direction in each scan cycle might be never noticed, but may accumulate to a rather big change in many scan cycles[4]. E.g., if we measure direction under these circumstances with a coarse discretization like in figure 4 (I), and we make only small changes in direction in each scan cycle (e.g. 20°), we would never notice, even if we had made a whole loop in the end (cf. figure 5). The only changes in direction we would notice would be very sharp ones in one single scan cycle. This leads to the dissatisfactory situation that the representation of the spatial path of a course of motion at a given scan rate depends greatly on the speed of the moving object.

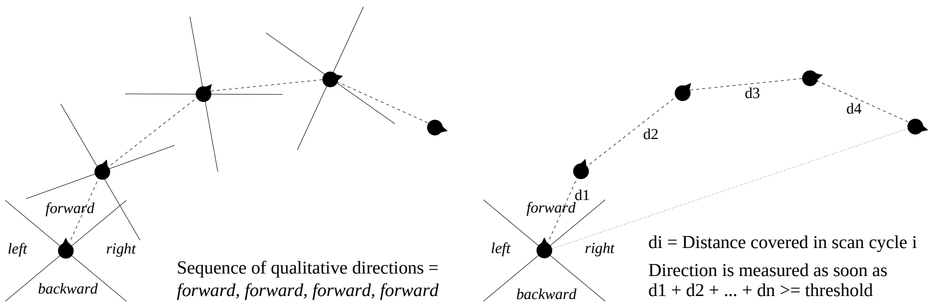

Figure 5. Left: Small changes in direction may be problematic – Right: Solution

So, for the case of locomotion with its egocentric frame of reference in measurement we have to take into account these special circumstances when we want to create an adequate qualitative representation.

We can solve this problem by measuring changes in direction not in every scan cycle, but only after a certain distance was covered (see figure 5). This is the reverse paradigm to (Musto et al., 1998), where time is fixed and distance is variable in the measurement. Here, we fix the distance and therefore have to measure the time needed to cover the fixed distance in order to derive speed. This makes sure that the resulting representation doesn't depend on the speed of the motion.

Following (Musto et al., 1998), we can then represent locomotion as a sequence of egoQMV's, measured in the egocentric frame of reference in measurement and represented in the egocentric frame of reference in representation. At first, an egoQMV consists

[4] This is true for humans, too. When we are in an environment with no landmarks like a desert, we tend to walk in counterclockwise circles while believing to walk in a fixed allocentric direction.

of the components Direction D, e.g. {`forward, backward, left, right`} and Number of time cycles t needed to cover the fixed distance. A counter indicates for how many measurement cycles no change in direction and speed[5] occured: `<t, D>`^{i}. Thus a course of motion is represented as a sequence of egoQMVs, for example:

<6 forward>^{5} <6 left>^{2} <6 left>^{3} <6 left>^{1} <4 forward>^{6}
<4 right>^{6}.

We can transform this representation into an allocentric frame of reference in representation if we know the allocentric heading and distance of the object when starting the motion or simply take the starting point and heading as origin for the allocentric coordinate system. Actually, if we represent distance and speed too in discrete, qualitative values, and assume that the allocentric heading at the starting point is east, we get the same sequence as in the previous example (if the measurement distance falls into the qualitative interval "close"):

<medium-dist east slow> <close north slow>
<close west slow> <close south slow>
<medium-dist south medium-vel>
<medium-dist east medium-vel>[6]

We see that we can get to the same representation of the motion event at the representational level regardless of the frame of reference we chose for the measurement level. Nevertheless we have to take into account that we possibly make big mistakes in this transformation, since we rotate the direction grid constantly when mapping the direction of motion onto the qualitative interval in the egocentric frame of reference in measurement, whereas in the allocentric frame of reference in measurement this grid is fixed.

[5] Since we measure after a fixed distance, t is equivalent to speed: the larger the number, the slower the robot is moving.

[6] The distance information is computed from the counter combined with the measurement distance: in the first vector, the moving object went 5 times the measurement distance forward $\rightarrow$ `medium-dist`. The speed information is directly derived from t: the moving object needed 6 time cycles to cover the measurement distance, which is a `slow` motion.

4. Granularity and Frames of Reference

A QMV representation is a qualitative representation of a course of motion, but still on a relatively fine granularity. Furthermore, it does not take into account the environment the motion takes place in.

In most cases where motion descriptions occur, however, they are of a coarser granularity and landmarks are of importance. Consider, for example, a route description: This is a description of a course of motion in a certain environment, with landmarks, and in a coarse granularity. It abstracts from the fine structure of the course of motion one may perform while following the route, when, e.g. crossing the street for window shopping or giving way to other pedestrians, and focusses on the coarse structure in a mostly egocentric frame of reference (in representation). Route descriptions can be segmented into pieces that belong mainly into four categories: start point, reorientation (direction), path/progression, and end point (see, e.g., (Tversky and Lee, 1998)). However, Tversky and Lee (1998) also found that people often give additional information like extra landmarks (not only at turning points), cardinal directions, shape of the path between landmarks. "This information, while not essential, may be important for keeping the traveler confidently on track".

So, a route description should not only consist of landmarks with turning directives, but also mention the coarse shape of the course of motion between landmarks and regions to cross, and even cardinal directions, when possible. Regions can serve as a kind of extended landmarks, like in "Then you cross Goetheplatz". The motion performed inside the regions to cross is irrelevant and not mentioned. But: If we zoom in these regions, i.e. expand the course of motion here, a switch from egocentric to intrinsic frame of reference is possible. The region may define its own frame of reference with an intrinsic front/backside, etc. E.g. if our course of motion crosses a church yard, the side where the church is defines the intrinsic front of the region. If someone were to find something on this place, a route description would probably switch from a coarse to a finer granularity at the point the place is entered, and switch from an egocentric to an intrinsic perspective by using the church yard's intrisic frame of reference; or define the motion in an egocentric perspective, but not relative to the last motion, but to the intrinsic frame of reference.

How could we automatically obtain such a coarser route description from a QMV sequence?

- **Incorporating Landmarks** First of all, we have to incorporate landmarks in the QMV sequence. The easiest and most exact way to associate landmarks with a QMV sequence is to incorporate them in the QMV's themselves: a QMV then not only represents direction, distance and speed of the motion, but also what landmarks were passed. The landmarks then get a temporal extension; according to their spatial extension and the speed of the motion while passing the landmark. For this purpose, it doesn't matter whether the frame of reference in measurement is egocentric or allocentric. To ensure efficient reasoning, the landmarks should not only be given a name, but it should also be memorized on which side the landmark was passed and whether the landmark is a region with its own intrinsic frame of reference.

- **Generalization** As mentioned before, the fine structure of a course of motion is irrelevant for giving a route description. Therefore, when we want to generate a route description from recorded (loco-)motion data, we should generalize this data to obtain the relevant coarse structure. Generalization algorithms for QMV Sequences and on-the-fly-generalization for numeric motion data are described in (Musto et al., 1998; Musto et al., 1999b; Musto et al., 1999a).

- **Segmentation and classification in motion shapes** As also mentioned before, a route description is not a mere sequence of landmarks with turning directives at the landmarks. The shape of the path between the landmarks can serve as a landmark itself (e.g. in a description like "Go up the street until you reach the japanese restaurant. Turn to the left there and follow the street. The street *makes a sharp turn to the left*, which you follow. After this turn, take the next street at the right.") Therefore, the shape a QMV sequence describes between the landmarks that serve as turning points should be extracted. This can be done with the segmentation and classification algorithm described in (Musto et al., 1998). In this process we additionally memorize the QMV sequences from which the shapes were generated to make switches between different granularities possible.

This algorithm has several consequences:

Multi-purpose landmarks Segmentation and classification of the QMV sequence in a sequence of motion shapes takes place independently from the landmarks. That is because landmarks can not only be used to indicate places in the route where some directional change has to be made, but can serve also for control purposes: "You pass three traffic lights before you get to a travel agency. There you turn left".

Landmarks with spatial and temporal extension Since landmarks get a temporal extension, not only landmarks are associated to QMVs, but it is possible that a whole sequence of motion shapes is associated to a single landmark. This is the case when the landmark is a larger region like, e.g. the English Garden in Munich. To these landmarks an intrinsic frame of reference can be associated.

Mixed granularities and frames of reference Routes can be described in mixed granularities and mixed frames of reference. When the course of motion is measured in an allocentric frame of reference, conversion in the egocentric frame of reference in representation is easy. In fact, in our software, we make this conversion automatically and have the egocentric and allocentric data available in the QMV sequence at any time. When we measure in the egocentric FoR, we can convert in the allocentric frame of reference in representation too if we know Ego's allocentric heading when beginning the motion (and accept of making mistakes due to errors in measurement). The allocentric distance is irrelevant, since we only deal with positional change, not with absolute position. So we can switch at any place of our QMV sequence between allocentric and egocentric frame of reference. This corresponds to a route description like "Go east until you reach the mall; turn left there".

We can also give a route description in different granularities. We can, e.g., give the description in the coarse granularity of motion shapes and landmarks and switch to the finer granularity of QMV sequence when we get close to the goal.

Last but not least, we can mix the egocentric perspective of the motion shape/landmark representation with the intrinsic frame of reference of a region: "You go to the left, down the street to the traffic light and then to the right until you enter the church yard. The butcher is on the church's left side". To this end, we have to know the allocentric heading of the intrinsic frame of reference.

Example:

For the sake of simplicity, we leave out the step of generalization. Different generalization algorithms are described in detail in (Musto et al., 1998; Musto et al., 1999b; Musto et al., 1999a). So assume that we have measured the QMV sequence corresponding to the dashed track in figure 6 in the allocentric frame of reference which generalizes to the following sequence (corresponding to the solid track in figure 6):

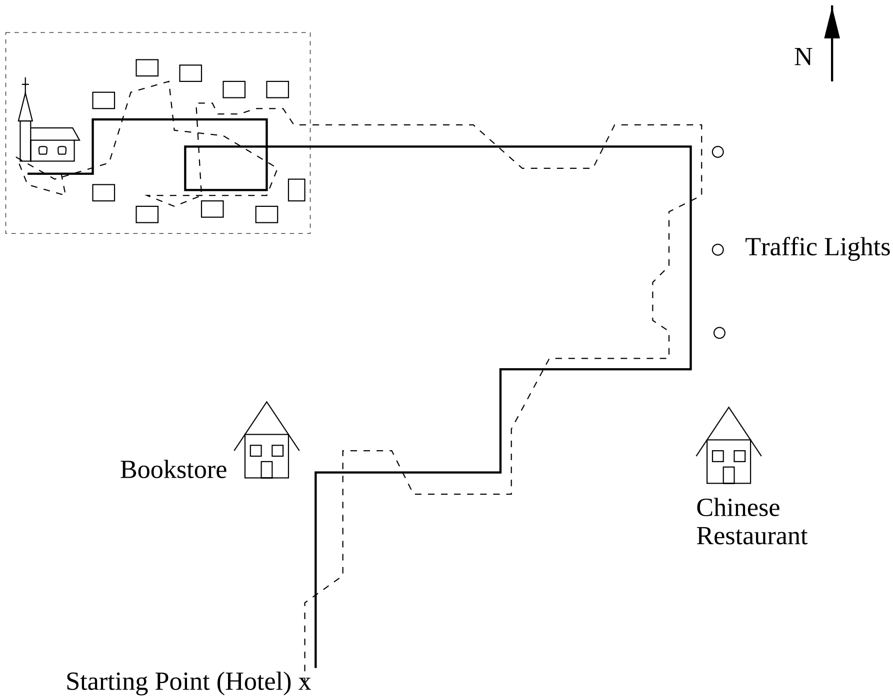

Figure 6. Example track with landmarks and generalization

```
<medium-dist north slow [s,lm1]>
<medium-dist east very-slow [lm1,lm0]>
<medium-dist north slow [lm0]>
<medium-dist east slow [lm2]>
<medium-dist north very-slow [lm2,lm3,lm4,lm5]>
<very-far west very-slow [lm5,e]>
<close south very-slow [e]><close east very-slow [e]>
<close north slow [e]> <medium-dist west slow [e]>
<close south slow [e]><close west very-slow [e]>
```

Since the landmark are something we assume to know, we identify them only by numbers. “0” means that there is no landmark, “s” and “e” are starting and end points. We see that landmark e must be a region since there is a longer subsequence associated with it.

During generalization, some QMV’s are combined to a single one. That’s the reason why some QMV’s have several landmarks associated.

We automatically compute egocentric directions:

```
<medium-dist north forward slow [s,lm1]>
<medium-dist east right  very-slow  [lm1,lm0]>
<medium-dist north left slow [lm0]>
<medium-dist east right slow [lm2]>
<medium-dist north left very-slow [lm2,lm3,lm4,lm5]>
<very-far west left very-slow [lm5,e]>
<close south left very-slow [e]>
<close east left very-slow [e]>
<close north left slow [e]>
<medium-dist west left slow [e]>
<close south left slow [e]>
<close west right very-slow [e]>
```

Now, we can start the segmentation and classification algorithm, which produces motion shapes and the correspronding list of landmarks passed during the motion was performed:

Start:	straight-line
Bookstore:	right-turn
	s-curve-left
	straight-line
Chinese Restaurant:	left-turn
Traffic Lights No. 1, 2:	straight-line
Traffic Light No. 3:	left-turn
Church Yard:	straight-line
Church Yard:	loop-left
Church Yard:	s-curve-left

At a coarser granularity, we can ignore the motion performed inside landmark e (maybe when the region of landmark e is reached, everything is obvious, and no turning directives have to be given, only a description like "then you can already see the church") and give only this description:

Start:	straight-line
Bookstore:	right-turn
	s-curve-left
	straight-line
Chinese Restaurant:	left-turn
Traffic Lights No. 1, 2:	straight-line
Traffic Light No. 3:	left-turn
Enter Church Yard:	straight-line

However, since we have memorized the QMV sequence associated with each motion shape, we could also expand the description at any place and switch to a finer granularity:

Start:	straight-line
Bookstore:	right-turn
	s-curve-left
	straight-line
Chinese Restaurant:	left-turn
Traffic Lights No. 1, 2:	straight-line
Traffic Light No. 3:	left-turn
Enter Church Yard:	straight-line
Inside Church Yard:	`<close south left very-slow [e]>`
	`<close east left very-slow [e]>`
	`<close north left slow [e]>`
	`<medium-dist west left slow [e]>`
	`<close south left slow [e]>`
	`<close west right very-slow [e]>`

Knowing the intrinsic frame of reference of landmark e, the allocentric directions can be transformed into intrinsic ones. Since landmark e is a churchyard with the church at the west side, intrinsic directions for the above given QMV subsequence associated with landmark e read:

```
<close left-intrinsic left very-slow [e]>
<close backward-intrinsic left very-slow [e]>
<close right-intrinsic left slow [e]>
<medium-dist forward-intrinsic left slow [e]>
<close left-intrinsic left slow [e]>
<close forward-intrinsic right very-slow [e]
```

5. Conclusion

When talking about spatial frames of reference in motion representation, we have to make a distinction between the frame of reference in which the course of motion is measured and the frame of reference in which the course of motion is represented.

We can conclude that measuring a course of motion in the egocentric or the allocentric frame of reference raises different problems. In general, measuring motion in an allocentric frame of reference in measurement is less problematic because of the unchanging coordinate system. Nevertheless, egocentrically measured data can be represented in an allocentric frame of reference when we accept the transformation to be possibly erroneous.

If landmarks are incorporated in the measured QMV sequences, route descriptions can be automatically computed which can be used at different and even mixed granularities and with different and mixed frames of references in representation. Possible frames of reference are the egocentric, allocentric, and, in landmarks with extension, the local intrinsic frame of reference.

There are, however, still some problems in the current approach: the assignment of landmarks to the appropriate motion shapes is too coarse — we do not know whether the landmark was passed at the beginning or the end of the shape, or in the middle. This is because through generalization this information is lost. We can remedy this by memorizing whether the landmark was encountered in the beginning, end, or middle of the motion shape. Furthermore, we should also memorize whether the landmarks were passed at the right or left side or wheter they are in front of or at the back of the moving object. This depends a little bit on the sonsory apparatus the moving object/person has — whether there are any sensors at all at the backside, for example.

Another problem is that our segmentation algorithm, that is independent drom the landmarks, is focussing much on the shape of the course of motion, which seems unnecessary when there is a landmark at the turning point. Then we don't need to segment the course of motion in curves as smooth as possible, but rather into straight lines with sharp turns at the landmark. Therefore, we will develop a second segmentation algorithm to be used at landmarks at turning points.

References

Brewer, B. and J. Pears: 1993, 'Introduction: Frames of Reference'. In: N. Eilan, R. McCarthy, and B. Brewer (eds.): *Spatial Representation. Problems in Philosophy and Psychology.* Blackwell Publishers.

Clementini, E., P. Di Felice, and D. Hernández: 1997, 'Qualitative representation of positional information'. *Artificial Intelligence* **95**(2), 317–356.

Freksa, C., C. Habel, and K. Wender (eds.): 1998, *Spatial Cognition. An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, Vol. 1404 of *Lecture Notes in Artificial Intelligence.* Springer.

Klatzky, R. L.: 1998, 'Allocentric and egocentric spatial representations: Definitions, distinctions and interconnections'. in (Freksa et al., 1998).

Levinson, S. C.: 1996, 'Frames of reference and Molyneux's question'. In: P. Bloom, M. A. Peterson, and N. Nadel (eds.): *Language and Space.* Cambridge, MA: MIT Press.

Musto, A., K. Stein, A. Eisenkolb, and T. Röfer: 1999a, 'Qualitative and Quantitative Representations of Locomotion and their Application in Robot Navigation'. In: *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99). To appear.*

Musto, A., K. Stein, A. Eisenkolb, K. Schill, and W. Brauer: 1998, 'Generalization, Segmentation and Classification of Qualitative Motion Data'. In: H. Prade (ed.): *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI-98).* pp. 180–185.

Musto, A., K. Stein, K. Schill, A. Eisenkolb, and W. Brauer: 1999b, 'Qualitative Motion Representation in Egocentric and Allocentric Frames of Reference'. In: *Proceedings of the International Conference on Spatial Information Theory COSIT-99.* Stade, Germany. To appear.

Pöppel, E. and K. Schill: 1995, 'Time Perception: Problems of Representation and Processing'. In: M. A. Arbib (ed.): *The Handbook of Brain Theory and Neural Networks.* The MIT Press, pp. 987–990.

Tversky, B.: 1993, 'Cognitive Maps, Cognitive Collages, and Spatial Mental Models'. In: A. U. Frank and I. Campari (eds.): *Spatial Information Theory. A Theoretical Basis for GIS. European Conference, COSIT'93*, Vol. 716 of *Lecture Notes in Computer Science.* Berlin,Heidelberg,New York, pp. 14–24.

Tversky, B. and P. U. Lee: 1998, 'How Space Structures Language'. in (Freksa et al., 1998).

Address for Offprints:
Alexandra Musto
Institut für Informatik der Technischen Universität München
D-80290 München, Germany

Tel: +49-89-289-28480
Fax: +49-89-289-28483
e-mail: musto@in.tum.de