Using Commonsensical Cardinal Directions to Describe Bordering Objects

Gregory Kritzman
School of Computing and Information Science,
University of Maine
Orono, Maine
gregory.kritzman@maine.edu

Torsten Hahmann
School of Computing and Information Science,
University of Maine
Orono, Maine
torsten@spatial.maine.edu

ABSTRACT

Geographic information systems (GIS) customarily encode spatial information using geometric objects (points, polylines and polygons) and their locations. But people frequently use qualitative relations, such as topological relations (e.g., connection or overlap) or cardinal direction relations (e.g. North or Southeast), to describe spatial scenes. While topological relations have been integrated into modern GIS, direction relations have remained isolated from GIS and are not available for user interaction. Instead, a user must visually infer them from map depictions.

This work uses the problem of generating and interpreting cardinal direction labels that describe the direction between a region and its surrounding neighbors (e.g., all neighbors of a US state) to identify principles for computing more qualitative descriptions of directions that are intuitive to people and to correctly interpret descriptions commonly used by people. This is a step towards bridging the qualitative-quantitative divide between spatial information systems and human conceptualizations of space.

CCS CONCEPTS

• **Information systems → Geographic information systems**; Multimedia and multimodal retrieval; • **Computing methodologies → Spatial and physical reasoning**; *Cognitive science*; Natural language generation; Ontology engineering;

KEYWORDS

cardinal directions, qualitative spatial reasoning, spatial relations

ACM Reference Format:
Gregory Kritzman and Torsten Hahmann. 2018. Using Commonsensical Cardinal Directions to Describe Bordering Objects. In *26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '18), November 6–9, 2018, Seattle, WA, USA*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3274895.3274910

1 INTRODUCTION

Spatial information systems (SIS) typically encode the location of objects geometrically in an absolute coordinate system and users interact with that information via map displays. Map displays, however, place an often unnecessary high cognitive burden on the user because much of the presented metric details, such as precise distances, lengths or sizes of areas, must be filtered out by a human observer who is only interested in, for example, the general direction between two objects. To communicate such high-level spatial knowledge, people rely frequently on *qualitative spatial relations*, including topological and parthood relations (e.g., are they at all in contact?, do they overlap?) or cardinal directions (e.g., which general direction is one relative to the other?) between spatial objects rather than pinpointing the objects' precise locations or quantifying the direction between them (e.g., measuring some angular direction). For example, to identify an unnamed object, a qualitative description of its location relative to other objects (e.g., "the highway southeast of Bangor" or "the county south of Aroostook and east of Piscataquis") often suffices. To make non-visual human interaction with spatial information less burdensome for people, we aim to generate more intuitive *qualitative spatial descriptions* from geometric map sources. These descriptions use the kind of spatial relations people are accustomed to in everyday conversation and that do not require access to a detailed viewable map. In the longer term, we aim to extend our approach to also *interpret* qualitative descriptions relative to maps sources to allow more user-friendly spatial querying of maps.

While much progress has been made on defining and computing intuitive topological and mereological (i.e., parthood) relations from geometric data using the 9-intersection matrix [4] and the Region-Connection Calculus [3], equally intuitive sets of directional relations – in particular cardinal direction relations such as North (N), Northwest (NW), or North-Northwest (NNW) – are still elusive. We will show that multiple competing conceptualizations for the same set of cardinal direction relations exist, so that one static framework may not work in general. For example, the interpretation of "North" varies widely in the size of the angular acceptance range: it could cover the Northern half of an object (i.e., be interpreted as "not South") or cover only the Northern quarter of an object (i.e., be interpreted as "neither East, South, nor West"), with the interpretation depending on the specific spatial context. As another example, saying that to get to Bangor, you need to drive North from Waterville, is taken to mean that one should take the direction that is approximately North (as compared to South) on the Interstate (I95) connecting the two towns, though in absolute geographical terms Bangor lies more to the East than to the North of Waterville. But since I95 is divided into the two complementary directions North and South, North is more appropriate than South.

SIGSPATIAL '18, November 6–9, 2018, Seattle, WA, USA

ACM ISBN 978-1-4503-5889-7/18/11.
https://doi.org/10.1145/3274895.3274910

The specific objective of this paper is to develop a method for automatically deciding which cardinal direction terms best describe the directions of all of a 2D region's neighbors, for example, for all the counties (or any other administrative region) that surround Somerset County in Maine (or any other administrative region) as illustrated in Figure 2. Here, it would be reasonable to describe either of Piscataquis (PS) and Penobscot (PN) county as being to the East, but in the presence of both, it might be more appropriate to describe Penobscot as being to the Southeast to distinguish between the two county's direction from Somerset. Such contextualization is also useful for interpreting descriptions, for example references to "the county east of Somerset" are then resolved to more likely refer to Piscataquis rather than Penobscot.

Compared to prior work, such as [1, 8], our approach is novel in that it automatically adapts the directions' granularity (when to use NE or NW rather than N? When to use NNE?) and is able to mix cardinal direction terms of different granularities. Our work's novelty also lies in identifying first principles for computing a single descriptive direction label for each neighbor such that the direction label is unique among all of a region's neighbors to the extent possible. While previous approaches have always related objects as if they exist in isolation, we *consider all neighbors of an object at once* to decide on the best direction label for each of them. Where the neighboring objects are spaced out fairly equally, our approach yields descriptions comparable to prior approaches. However, in cases where many targets are close together, we are often still able to generate unique direction labels whereas prior approaches would describe multiple neighbors using the same direction, thus rendering them indistinguishable from the description alone. Moreover, prior work solely describes the direction between the "bulk" of two objects, rather than of the boundary one shares with the other (e.g., "Piscataquis borders Somerset to the East") as we attempt.

2 BACKGROUND & RELATED WORK

Qualitative direction relations are either *relative* or *absolute*. Relative relations (e.g., [7, 13, 15, 23]) construct an internal reference system either using the intrinsic orientation of an object (e.g., the "front") or via a third reference object. Objects in this system are generally described as being 'in front of', 'behind', 'left of', or 'right of' a reference object. Absolute relations [6, 8, 9, 17, 19, 21] rely on some external reference system (i.e., coordinate system) that is independent of the orientation of the considered object. In geographic domains, objects described using *absolute directions* typically employ so-called cardinal directions, such as being 'north of', 'east of', or 'south of' a reference object. We are interested in absolute directions in order to provide a birds-eye description of the direction between objects.

In both absolute and relative directional systems, the reference and/or target objects are typically simplified. One common simplification uses the minimum bounding box of the reference object [2, 8, 9, 11, 12, 16, 20, 21], around which the space is divided using rectangular projections. An alternative (see, e.g., [5, 10, 19]) represents the reference object as a point, e.g. its center of mass, and establishes a set of conical projections radiating from that point. By adjusting the size of the cones, the granularity of the computed relations can be adjusted [10, 19], whereas refining the granularity of rectangular projections often leads to drastic increases in cognitive complexity [12]. Hybrid approaches (e.g., [18, 22]) combine the bounding box simplification of the reference object with the conical projections originally developed for the point-based simplifications.

Our work differs from studies of direction relations in qualitative spatial reasoning (QSR) [6, 7, 14, 15, 19, 23] by our focus on extracting and, eventually, interpreting qualitative directions *relative to a geometrically encoded map*. Thus, we aim to connect qualitative and quantitative spatial representations rather than perform purely qualitative reasoning on a given set of qualitative direction facts.

3 PRINCIPLES AND APPROACH

We use a hybrid approach that combines conical projections with the reference object's original shape. Rather than simplifying the reference object to a point or a bounding box, we simplify each target object to the segment of boundary it shares with the reference object. More precisely, we partition the reference object's boundary and thus avoid introducing overlap between target objects to help generate unique labels for each target later on. The chosen simplification allows us to generate more concise descriptions of a reference object's boundary and immediately surrounding area than other approaches, though with two drawbacks: (1) it is limited to adjacent objects (i.e. objects that have a common boundary), and (2) it requires full knowledge about how the reference region's boundary is segmented by neighboring objects.

One of the main challenges we face in producing and interpreting human-comprehensible qualitative direction descriptions is the use of multiple competing spatial conceptualizations by people. Cognitive science literature suggests that people adjust level of detail (i.e., granularity) as necessary or convenient [10]. The Star calculus [19] is one of the only previous approaches that allows adjusting the granularity of directions by specifying how often to subdivide sectors. However, it does not give any guidance on how to choose the right granularity and when to vary the granularity among the neighbors of a reference object. We aim to adjust the size according to the arrangement of the neighbors: in directions where multiple neighboring regions are close together a fine granularity is used whereas in directions were neighbors are further apart, the granularity is decreased. For example, relative to Somerset county, we will describe Waldo as being in the SSE eighth whereas Piscataquis is in the Eastern (E) half (cf. Fig. 2). Generally, we want to use the coarsest granularity that still uniquely describes a target object as coarser labels are often cognitively simpler – especially sectors of eights are less commonly found in human descriptions.

3.1 Conceptualizations of Directions

As a first step, we have compiled the set of eight common directional conceptualizations of different granularity shown in Fig. 1, wherein all sectors ("cones") are of equal size and that use only labels of length 1 (N, E, etc.), 2 (NE, SE, etc.) or 3 (NNE, ENE, etc.) as these labels are arguably the only ones people would use in everyday descriptions. We denote the used sectors using a granularity index of 8 for the 16 eighths (NNE_8, NE_8, ENE_8, E_8, ESE_8, SE_8, SSE_8, S_8, SSW_8, SW_8, WSW_8, W_8, WNW_8, NW_8, NNW_8, N_8), 4 for the eight quarters (NE_4, E_4, SE_4, S_4, SW_4, W_4, NW_4, N_4), and 2 for the eight halves (NE_2, E_2, SE_2, S_2, SW_2, W_2, NW_2, N_2). Note, however, that

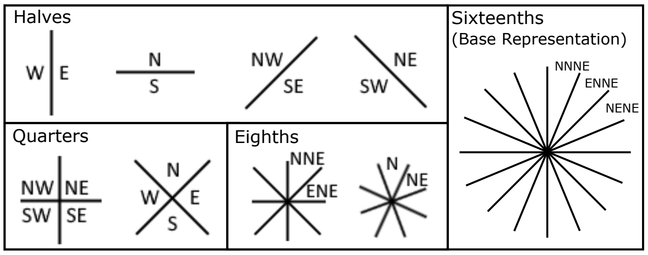

Figure 1: The eight conceptualizations of cardinal directions that yield the eighth, quarter and half sectors. The base representation using sixteenths is shown on the right.

Table 1: The sixteenths from the base representation and their aggregation into coarser eighths, quarters and halves with the curly brackets showing which sixteenths are covered by the coarser sectors. *Note the table shows only the first 10 sixteenths and aggregations into larger sectors that include only those sectors. The empty cells will "wrap around" at the end of the table because of the circular nature of the relationship between sectors.*

Sixteenth	Eighths		Quarters		Halves	
(internal)	len. 3	len. 2 & 1	len. 2	len. 1	len. 2	len. 1
NNNE$_{16}$	NNE$_8$					
ENNE$_{16}$		NE$_8$	NE$_4$			
NENE$_{16}$	ENE$_8$					
EENE$_{16}$		E$_8$		E$_4$	E$_2$	
EESE$_{16}$	ESE$_8$					
SESE$_{16}$		SE$_8$	SE$_4$			SE$_2$
ESSE$_{16}$	SSE$_8$					
SSSE$_{16}$		S$_8$		S$_4$		
SSSW$_{16}$	SSW$_8$					
WSSW$_{16}$						
…						

our approach extends to more fine-grained conceptualizations or to an entirely different way of labeling sectors, e.g., one that uses egocentric direction label such as front, left, right, and rear, though the base representation must be adjusted accordingly.

3.2 The Base Direction Representation

To help compute human-comprehensible direction labels, we have devised a finer-grained *base representation* consisting of 16 "sixteenths" (shown on the right in Fig. 1) such that all sectors from the eight conceptualizations are sums of sets of sixteenths, as outlined in Table 1. To generate a single unique label for each target region then requires finding a way to *generalize* the sixteenths to coarser sectors while avoiding so-called *congruent labels* that only differ in their index (e.g., E_8, E_4 and E_2 are all congruent).

Computing the base representation is the first step in our two-step approach. We start by determining the reference region's centroid (using ArcPy's centroid function) and create the cone-shaped sectors (represented as triangles) for all sixteenths from the centroid. Each sixteenths is then intersected with every boundary segment that represents a neighboring target object to determine

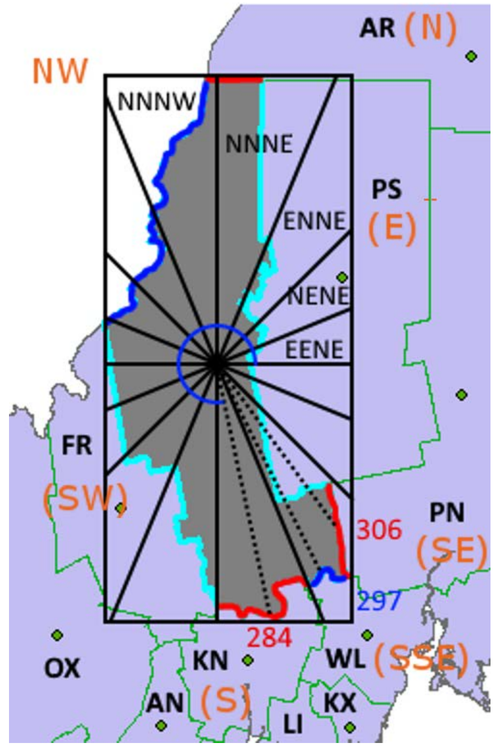

Figure 2: Visualization of our approach using Somerset County in Maine as example reference region (dark grey). Its boundary is partitioned into the segments (shown in red, turquoise and dark blue) it shares with its neighboring counties Aroostook (AR), Piscataquis (PS), Penobscot (PN), Waldo (WL), Kennebec (KN), Franklin (FR), and the border with Quebec, Canada. These boundary segments are intersected with the 16 sectors emanating from Somerset's centroid to form the base representation. In addition, the angle of the midpoint of each boundary segment is stored.

the set of overlapping sixteenths for each target object. For example, Piscataquis overlaps Somerset in seven sixteenths: NNNE, ENNE, NENE, EENE, EESE, SESE, ESSE, whereas Penobscot overlaps Somerset only in the ESSE sixteenth. This produces a fine-grained, declarative abstraction of the directions of the neighboring regions that is mostly qualitative except that it also retains the angle of the midpoint of each boundary segment for later tie-breaking in case two or more target regions overlap exactly the same set of sectors. Other geometric and quantitative information from the underlying map, such as shape, distances, or area sizes, is completely discarded.

3.3 Computing Single and Unique Directions

For each target object, we aggregate adjacent sectors into coarser ones to identify the coarsest labels that exactly describe the target. The aggregation combines two adjacent sixteenths into an eighth, multiple overlapping eighths (e.g., NNE$_8$, NE$_8$, ENE$_8$) into a quarter (e.g., NE) if possible[1], and three overlapping quarters into a half.

[1]For example, NE$_8$, ENE$_8$, and E$_8$ cannot be aggregated into a single quarter.

For example, the seven sixteenths for describing Piscataquis' direction from Somerset result in the sectors SE_8, E_4, and NE_4. These aggregations are lossless in that the resulting coarser sectors cover precisely the set of sixteenths from the base representation.

While aggregation produces *coarser* and *fewer* labels, it does not necessarily reduce them to a single label nor does it ensure uniqueness. The second step aims to identify the coarsest sector for each target such that there is no conflict across targets. Rather than picking one of the sectors produced by the aggregation step, all sectors of a target region are overapproximated by one or two sectors of the next coarser granularity. For example, Piscataquis' three labels SE_8, E_4, and NE_4 are approximated by a half (because E_4 and NE_4 are included as quarters), with E_2 being the half that maximally covers the three sectors. In this specific example a single sector is produced and no other target uses a *congruent label*, thus "E" is used to uniquely describe Piscataquis' direction relative to Somerset. Two kinds of complications can arise: (1) multiple sectors can be good overapproximations of a target's aggregated sectors; and/or (2) two or more targets share an overapproximated sector. Both issues are, in part, addressed by an elimination process that deletes the worst sector among all targets. If that leaves a target with a single sector, it is chosen (and deleted from all other targets), unless the sector is congruent to an already assigned one. Conflicts also arise when deletion leaves two or more target regions with only congruent sectors. All conflicts are resolved – to the extent possible – by *refining* the sectors via the addition of neighboring or finer-grained labels as additional choices. In the exceptional case when this broadening does not resolve the conflict, two or more target regions are assigned the same label. This only happens when more than 16 target regions exist (i.e., more than the total number of available labels of length 1 to 3) or more than 5 target regions overlap the same quarter (i.e., more than the total number of available labels available for that area).

4 CONCLUSIONS AND FUTURE WORK

We have presented an approach for computing human-comprehensible direction descriptions that bridges the divide between *precise geometric, yet implicit* (i.e., not explicitly queryable) *computational representations* of directions, which are less suitable for human interaction, and *explicit* (e.g., using cardinal direction relations, such as North, between pairs of objects) *qualitative representations of directions*, which people understand more intuitively but which are inherently more vague and ambiguous. We have utilized three principles to minimize the produced description's cognitive complexity: (1) using a single label for each neighbor, (2) preferring the coarsest and thus least complex yet descriptive labels, and (3) assigning unique labels to each neighbor to the maximum extent possible. We developed a mechanism for working with and mixing direction labels of various granularities and from different directional conceptualizations, thus employing the direction labels with nuanced meanings that adapt to the specific spatial context. We specifically proposed how to automatically adjust the granularity based on how far neighboring regions are spaced apart, thereby implicitly selecting the conceptualizations that is most appropriate for each neighbor. Full cognitive evaluation, and comparison to other approaches are still outstanding.

In the presented setting, we require all target regions to border the reference region. Thus, the produced direction labels are refinements of the topological "borders" relation between 2D regions, such as "Piscataquis county borders Somerset county to the East." But in the future, we plan to generalize this approach to objects that not necessarily partition the border of the reference object (e.g., towns scattered around a lake or in- and outlets of a lake).

Acknowledgement. This material is based upon work supported by the National Science Foundation under Grant Number III-1565811.

REFERENCES

[1] K. Buchin, V. Kusters, B. Speckmann, F. Staals, and B. Vasilescu. 2011. A splitting line model for directional relations. In *Intern. Conf. on Advances in Geographic Information Systems (ACM SIGSPATIAL)*. ACM, 142–151.
[2] Juan Chen, Dayou Liu, Haiyang Jia, and Changhai Zhang. 2007. Cardinal direction relations in 3D space. In *International Conference on Knowledge Science, Engineering and Management*. Springer, 623–629.
[3] Anthony G. Cohn, Brandon Bennett, John M. Gooday, and Nicholas M. Gotts. 1997. Qualitative spatial representation and reasoning with the Region Connection Calculus. *GeoInformatica* 1 (1997), 275–316.
[4] Max J. Egenhofer and Robert D. Franzosa. 1991. Point-set topological spatial relations. *Int. J. Geogr. Inf. Sci.* 5, 2 (1991), 161–174.
[5] Andrew U Frank. 1992. Qualitative spatial reasoning about distances and directions in geographic space. *J. Vis. Lang. Comput.* 3, 4 (1992), 343–371.
[6] Andrew U Frank. 1996. Qualitative spatial reasoning: Cardinal directions as an example. *Int. J. Geogr. Inf. Sci.* 10, 3 (1996), 269–290.
[7] Christian Freksa. 1992. Using orientation information for qualitative spatial reasoning. In *Theories and methods of spatio-temporal reasoning in geographic space*. Springer, 162–178.
[8] R Goyal and Max J Egenhofer. 1997. The direction-relation matrix: A representation for directions relations between extended spatial objects. *USGIS annual retreat* 3 (1997), 95–102.
[9] Roop Goyal and Max J. Egenhofer. 2000. Consistent queries over cardinal directions across different levels of detail. In *Int. Workshop on Database and Expert Systems Applications*. IEEE, 876–880.
[10] Alexander Klippel, Jan Oliver Wallgrün, Jinlong Yang, and Kevin Sparks. 2015. Intuitive Direction Concepts. *Baltic International Yearbook of Cognition, Logic and Communication* 10, 1 (2015), 3.
[11] Lars Kulik, Carola Eschenbach, Christopher Habel, and Hedda Rahel Schmidtke. 2002. A graded approach to directions between extended objects. In *Int. Conf. on Geographic Inf. Sci. (GIScience-02)*. Springer, 119–131.
[12] Yohei Kurata and Hui Shi. 2009. Toward heterogeneous cardinal direction calculus. In *German Conf. Artif. Intell. (KI-09) (LNCS 5803)*. Springer, 452–459.
[13] G. Ligozat. 1998. Reasoning about Cardinal Directions. *J. Vis. Lang. Comput.* 9, 3 (1998), 23–44.
[14] Weiming Liu, Xiaotong Zhang, Sanjiang Li, and Mingsheng Ying. 2010. Reasoning about cardinal directions between extended objects. *Artif. Intell.* 174, 12 (2010), 951 – 983. https://doi.org/10.1016/j.artint.2010.05.006
[15] Till Mossakowski and Reinhard Moratz. 2012. Qualitative reasoning about relative direction of oriented points. *Artif. Intell.* 180 (2012), 34 – 45. https://doi.org/10.1016/j.artint.2011.10.003
[16] Isabel Navarrete and Guido Sciavicco. 2006. Spatial reasoning with rectangular cardinal direction relations. In *Workshop on Spatial and Temporal Reasoning at ECAI-06*. 1–10.
[17] Dimitris Papadias, Max Egenhofer, and Jayant Sharma. 1996. Hierarchical reasoning about direction relations. In *ACM Workshop on Advances in Geographic Information Systems (ACM GIS 1996)*. ACM, 105–112.
[18] D. Peuquet and Z. Ci-Xiang. 1987. An algorithm to determine the directional relationship between arbitrarily-shaped polygons in the plane. *Pattern Recognition* 20, 1 (1987), 65–74.
[19] Jochen Renz and Debasis Mitra. 2004. Qualitative direction calculi with arbitrary granularity. In *Pacific Rim Int. Conf. on Artif. Intell. (PRICAI-04)*. 65–74.
[20] Markus Schneider, Tao Chen, Ganesh Viswanathan, and Wenjie Yuan. 2012. Cardinal directions between complex regions. *ACM Transactions on Database Systems (TODS)* 37, 2 (2012), 8.
[21] Spiros Skiadopoulos and Manolis Koubarakis. 2004. Composing cardinal direction relations. *Artif. Intell.* 152, 2 (2004), 143–171.
[22] Spiros Skiadopoulos, Nikos Sarkas, Timos Sellis, and Manolis Koubarakis. 2007. A family of directional relation models for extended objects. *IEEE Transactions on knowledge and data engineering* 19, 8 (2007), 1116–1130.
[23] Diedrich Wolter and Jae Hee Lee. 2010. Qualitative reasoning with directional relations. *Artif. Intell.* 174, 18 (2010), 1498–1507.