Evaluating Field Crisping Methods for Representing Spatial Prepositions

Mark M. Hall
School of Computer Science
Cardiff University
5 The Parade, Roath
Cardiff CF24 3AA, United Kingdom
M.M.Hall@cs.cardiff.ac.uk

Christopher B. Jones
School of Computer Science
Cardiff University
5 The Parade, Roath
Cardiff CF24 3AA, United Kingdom
C.B.Jones@cs.cardiff.ac.uk

ABSTRACT

There is a need for GIR systems to interpret the vague aspects of spatial language. Here we describe an initial approach towards evaluating crisp realisations of a field-based model of the use of the spatial preposition "near", based on evidence of usage of the term in image captions.

Categories and Subject Descriptors

J.2 [**Physical Sciences and Engineering**]: [Earth and atmospheric sciences]

General Terms

Languages, Experimentation, Verification

Keywords

Spatial language, kernel density estimate, active contours, evaluation

1. INTRODUCTION

Many aspects of spatial language concerned with real world features and with the relationships between them are essentially vague. While this vagueness is managed quite effectively in natural language communication between people, there are currently only very limited facilities for interpreting such language when used to communicate with computers. This is adequate for many professional applications of GIS technology, but in an application such as geographic information retrieval (GIR) it would be desirable to be able to make "intelligent" interpretations of vague spatial language. This is the case, for example, when using vague spatial prepositions such as "north-of" and "near". This paper presents an experimental evaluation of an empirical model for the spatial preposition "near", but the methodology is applicable to any spatial preposition.

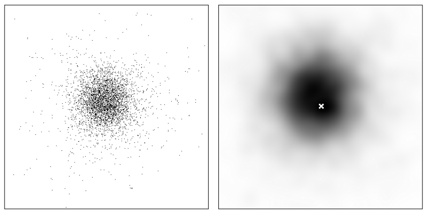

Figure 1: Training set of actual uses for "near", common centre indicated by the cross. KDE based on these uses using an axis-aligned normal kernel.

2. FIELD MODEL AND CRISPING

In order to exploit the knowledge in the spatial prepositions they need to be quantitatively modelled. We advocate a field based approach that allows for a continuous representation of the area that the spatial preposition applies to. The field representation in this paper is derived from a set of 4600 actual uses of the preposition "near" from photographic image caption data provided by the Geograph project[1] (fig. 1, the cross indicates the combined reference locations to which the other points were regarded as "near"), unlike previous work based direct acquisition from people [2]. Based on these uses a kernel density estimate (KDE) with a two-dimensional normal kernel derives the continuous field from the point data, normalised to the 0 to 1 scale (fig. 1).

For GIR purposes the field is applied to a query such as "Hotels near Cardiff" by centring the field on the location of the place in the query (here "Cardiff") and then using the field values to determine which documents apply to the query. For this a crisp boundary needs to be determined to decide whether a point belongs or does not belong to the region defined by the query. Two crisping methods have been investigated, based on the *contourLines* algorithm (as defined by the `R` statistical package) and on an active contour method. For the simple contouring algorithm 11 contours at 0.1 confidence steps between 0 and 1 were created (fig. 2).

Active contours are energy minimising splines that were initially developed for finding outlines in images [1]. To use them to to create contours we define two energies that maintain the active contour's shape and the spatial preposition field influence. A third energy is used to control how far the active contour contracts and this energy is increased in regular steps to create 10 contours (fig. 2).

[1] `http://www.geograph.org.uk`

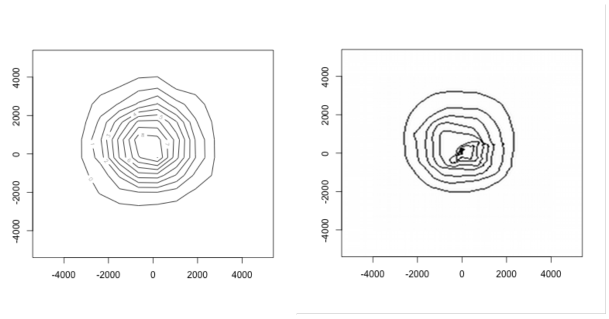

Figure 2: Contours created by the `R` *contour* function, in 0.1 steps between 0 and 1 (left). Contours created using active contouring in 10 contraction energy steps (right). The outermost active contour is not shown, as it is outside the core 8x8km region.

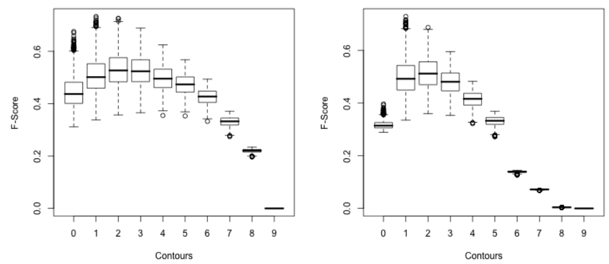

Figure 3: Boxplots for the F-scores by contour number for both the `R` *contour* (left) and active contour (right) methods. Visible is the similar distribution around the high F-score values and the large drop for the active contour results between contours 5 and 6.

3. CONTOUR EVALUATION

To evaluate the contours a quality metric needs to be defined. One possible quality metric is to interpret the contours as classifiers for a "near" / "not-near" classification task and then use standard classification task *precision*, *recall* and *F-score* metrics to evaluate the contours.

For this it is necessary to create a test set of "near" and "not-near" labelled points around a reference location. The set of "near" labelled points is taken from the original Geograph data by randomly selecting 500 points, which were then removed from the training set. Defining a set of "not-near" points is not as easy, as very few image captions use "not near <reference location>" which would make it possible to directly build up the "not-near" test set, but it is possible to create a simulation of "not-near" by calculating the points "near" alternative locations scattered around the reference location and then labelling these as being "not-near" the reference location. Five such locations are randomly placed around the original centre at distances between 1.5 and 5km, these distances being informed by an analysis of distances between toponyms in the Geograph caption data. The "near" test set is then used to create the points "near" these alternative locations, which are then labelled as "not-near" the original reference location.

As the "not-near" test-set creation is random within certain bounds, a 1000 iteration Monte-Carlo simulation was run to negate the influence of the more extreme placements of the alternative locations. In each iteration a new "not-near" test-set is created using the method described above and precision, recall and F-score calculated for each contour. The combined results of all Monte-Carlo runs are shown as box-plots per contour in figure 3.

3.1 Results

Analysis of the F-scores (fig. 3) indicates that the contours with the highest F-scores deliver comparable results for both methods with F-scores of 0.53 (simple contouring) and 0.51 (active contour). This is also supported by a visual analysis (fig. 2) of the highest scoring contours (contour 2 for both methods) which cover roughly the same area. Active contours can thus provide results that are of comparable quality to a simple crisping approach.

The main difference between the two crisping methods is that the active contouring method provides a clearer differentiation between contour results that provide a good and bad representation of the area "near". F-scores for those contours that are too large (contour 0) or too small (contours 6, 7, 8, 9) are very significantly lower than for the acceptable contours, whereas in the simple contouring method there are no such clear jumps between the mean F-scores of adjacent contours.

4. CONCLUSIONS

The results indicate that active contour based contour creation delivers results that are as good as those derived using a simple contouring algorithm. The advantage of active contours is that the algorithm provides a pronounced distinction between those contours that are not good representations and those that are acceptable representations, making it easier to filter degenerate contours. Another advantage is that further influences such as a hard boundary defined by a shoreline can easily be added, further improving the quality of the results.

The contours presented here have no inherent interpretation except for being arbitrary, crisp "near"/"not-near" decision points. Future work will focus on is how the contours are evaluated and interpreted by humans. Results of that evaluation can then be linked to the results presented here to provide an automatic contour evaluation that is grounded in human spatial interpretation.

5. ACKNOWLEDGEMENTS

We would like to gratefully acknowledge contributors to Geograph British Isles (see `http://www.geograph.org.uk/credits/2007-02-24`, whose work is made available under the following Creative Commons Attribution-ShareAlike 2.5 Licence (http://creativecommons.org/licenses/by-sa/2.5/).

This material is based upon work supported by the European Community in the TRIPOD (FP6 cr n°045335) project.

6. REFERENCES

[1] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[2] M. Worboys. Nearness relations in environmental space. *International Journal of Geographic Information Science*, 15(7):633–651, 2001.