Supervaluation Semantics for an Inland Water Feature Ontology

Paulo Santos* and **Brandon Bennett**

{psantos, brandon}@comp.leeds.ac.uk

School of Computing

University of Leeds, Leeds, UK

Georgios Sakellariou

georgios.sakellariou@ic.ac.uk

Electrical and Electronics Engineering

Imperial College, London, UK

Abstract

This paper describes an ontology for inland water features built using formal concept analysis and supervaluation semantics. The first is used to generate a complete lattice of the water domain, whereas supervaluation semantics is used to model the variability of the concepts in terms of threshold parameters. We also present an algorithm for a mechanism of individuation and classification of water features, from snapshots of river networks, according to the proposed ontology.

1 Introduction

There are at least three main aspects of an ontology for Geographical Information Systems (GIS): the ontology at the *data level*, whereby features of how raw data (either from sensor images or geographic surveys) are expressed in terms of numeric an symbolic information; the ontology at the *expert level*, that incorporates scientific knowledge about geographical objects; and the ontology at the *user level*, which should in some way specify the meanings of the vague and ambiguous geographic terminology which is used to describe the world in natural language. The present work is an investigation of an ontology of the latter kind for inland hydrographic features.

The ontology of inland water features proposed in this work is based on the *Waters* 'lexical field' presented as an example in [Ganter and Wille, 1999]. However, we shall analyse the domain in greater detail. In particular we shall explicitly model the intrinsic vagueness and ambiguity in water feature terms (such as *pond and lake*) by means of a variant of *supervaluation semantics* as proposed in [Fine, 1975; Bennett, 2001a; 2001b]. The methods developed in this work, however, are not restricted to the water feature domain, but can be applied to any ontology for GIS or, indeed, any complex domain containing vague concepts.

We argue in this work that *formal concept analysis* (FCA) and *supervaluation semantics* complement each other for building ontologies of naïve geography. On the one hand, FCA provides constructive tools for extracting complete lattices for any particular domain. Therefore, formal aspects of

the ontology can be assessed using the tools of lattice theory; moreover, there are also automatic methods for building lattices from a set of terms. Therefore, the problem of going from elements of a domain to the formal structure of this domain is solved by formal concept analysis. On the other hand, FCA does not handle changes in the interpretation of terms in the ontology due to intrinsic vagueness in the meaning of the corresponding concept. Supervaluation semantics provides rigorous methods for handling such changes of 'standpoint' while providing the machinery to make *reliable inferences* involving vague concepts. This ability to derive the reliable consequences of vague knowledge is the main advantage of supervaluation semantics over fuzzy logic [Elkan, 1993].
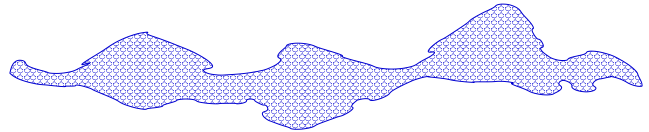

Figure 1: Vagueness in water features: *three lakes or a meandering river?*

Figure 1 shows an example of the kind of vagueness we are going to be dealing with in this work. This figure may depict either three lakes connected by channels or a meandering river stretch. The decision of one interpretation from another depends on the *standpoint* of the observer, i.e., on how the set of parameters distinguishing rivers from lakes are determined by the observer's judgement. Therefore, not only is vagueness pervasive to the feature definitions but also on how the features are individuated. It is worth pointing out also that this work concentrates on *sorites* vagueness, which is an indeterminacy in the boundaries of applicability of a term. This can be distinguished from conceptual *ambiguity* where a term has several qualitatively different interpretations.

The ontology of water features and the supervaluation semantics are encoded into a prototype implementation whose purpose is precisely to individuate qualitatively distinct features from a snapshot of a river network. Each individual feature is, thus, classified according to the definitions in the ontology, taking into account parameters handled by the supervaluation semantics.

This paper is organised as follows: Section 2 overviews some basic definitions of formal concept analysis and intro-

---

*Currently at Centro Universitário da FEI, São Paulo, Brazil.

duces the conceptual domain of inland water features (Section 2.2). Supervaluation semantics is discussed in Section 3. Section 4 describes our prototype implementation and Section 6 concludes this work.

2 Formal Concept Analysis

Formal concept analysis (FCA) [Ganter and Wille, 1999] is a powerful tool for formalising conceptual domains by means of algebraic structures such that the set of concepts can be represented as complete lattices.

2.1 Overview

The basic ideas underlying formal concept analysis are those of "formal context" and "formal concept" that are introduced below[1]. Informally, a *context* refers to the formal structure that characterises a conceptual domain and a *concept* within a context is composed of the ontology terms and their attributes. *Terms* in an ontology are called *objects* in formal concept analysis terminology, in this work we use these words interchangeably.

Formal contexts and concepts
A context K is a triple $K \equiv \langle G, M, I \rangle$ in which G is the set of terms in the context K, M is the set of attributes of the context and I is an *incidence relation* between G and M. Let's assume $g \in G$ and $m \in M$, then gIm is read as " the term g has attribute m".

In order to obtain the set of all attributes m ($m \in M$) common to a particular subset of objects, A, a *derivation operator* is defined as follows:

$$A' = \{ m \in M \mid gIm \ \text{ for all } \ g \in A \}.$$

This operator is overloaded to refer to the set of all objects $g \in G$ that have all of their attributes in a particular subset of M. Formally:

$$B' = \{ g \in G \mid gIm \text{ for all } \ m \in B \}.$$

Given a context $\langle G, M, I \rangle$, a *concept* is identified with a pair (A, B), where $A \subseteq G$, $B \subseteq M$, such that $A' = B$ and $B' = A$. The sets A and B are called, respectively, the *extent* and the *intent* of the concept (A, B).

Given two concepts $c_1 = (A_1, B_1)$ and $c_2 = (A_2, B_2)$, we say that c_1 is a subconcept of c_2 (equivalently, c_2 is a superconcept of c_1), or simply $(A_1, B_1) \leq (A_2, B_2)$, if $A_1 \subseteq A_2$. The set of all concepts of a particular context, ordered in this way, forms a complete lattice in which infimum and supremum are, respectively [2]:

$$\bigwedge_{t \in T} (A_t, B_t) = \left(\bigcap_{t \in T} A_t, \left(\bigcup_{t \in T} B_t \right)'' \right),$$

and

$$\bigvee_{t \in T} (A_t, B_t) = \left(\left(\bigcup_{t \in T} A_t \right)'', \bigcap_{t \in T} B_t \right).$$

---

[1]For simplicity we are going to drop the word formal in "formal context" and "formal concept" throughout this paper.

[2]As proved by the *basic theorem of concept lattice* [Ganter and Wille, 1999].

Many-valued contexts and scaling
Until now we have been describing the basic machinery of formal concept analysis without taking into account values of attributes with respect to objects. In fact, the definition of context described above only deals with cases of one-valued attributes where objects either have a particular attribute or not. However, in order to deal with the many-valued attributes found in the water feature domain (such as "flow" and "depth"), many-valued contexts should be taken into account.

A many-valued context is a quadruple $\langle G, M, W, I \rangle$ where G is a set of objects, M is a set of many-valued attributes, W are the attributes values and I is a ternary relation of $G \times M \times W$.

In order to represent many-valued contexts as complete lattices, they have to be reduced to one-valued contexts. This is done by means of partitioning each many-value attribute into a finite number of one-value attributes that are distinguished from one another by a set of thresholds. This partitioning process is named *scaling* and the set of partitions of an attribute is called the *conceptual scale* of that attribute in FCA terminology. A scaling can be understood as an interpretation of each of the attributes of a many-valued context by means of a *context of attributes*. A scale for a many-valued attribute m is a one-valued context $Sm = \langle Gm, Mm, Im \rangle$ with $m(G) \subseteq Gm$. In practise, every many-valued attribute m of a context is replaced by the one-valued scale attributes relative to $m(g)$, for any object $g \in G$.

2.2 The context of inland water features

In this section we propose a context for inland water recalling the *lexical field waters* introduced in [Ganter and Wille, 1999]. The present context is constituted of a set of objects (or water features): G = {river, stream, canal, reservoir, lake, marsh, pond} whose set of attributes is M={linearity, size, origin, flow, depth}.

The scale assumed for the linearity attribute partitions its values into *linear* and *non linear* according to a threshold on a *linearity coefficient*. The linearity coefficient assumed in this work is defined as the ratio between the diameters of the minimal bounding circle and the maximal internal circle w.r.t. the feature's border.

In an ideal world, where lakes have the shape of discs and rivers the shape of lines, rivers and lakes could be identified as features whose linearity ratio are, respectively, different from and equal to one. However, this idealistic distinction does not hold for lakes that are elongated or for wide river stretches, for instance. In order to cope with this issue, we assume a threshold on the linearity ratio (called the *linearity threshold*) whereby if the ratio of a water feature is greater or equal than this threshold, this feature is linear. It will be considered non-linear otherwise. This is in fact a simplification of the actual threshold used in the implementation presented in Section 4. The actual definition, however, needs concepts that will be introduced later on this paper.

Size is partitioned into four categories: *large linear*, *small linear*, *large non-linear* and *small non-linear*. The linear sizes are determined according to a threshold on the maximum length of linear features (*maximum linear length thresh-*

old), whereas the non-linear sizes take into account thresholds on both: the *maximum non-linear length* and the *maximum non-linear width*[3]. The distinct size thresholds for linear and non-linear features is due to the intrinsic dissimilarity between judgement of sizes regarding objects in these two classes (for instance, a *small lake* does not necessarily have the same dimensions as a *small river*).

Flow is scaled as *flowing* or *stagnant* given a threshold on the average flow of the feature. Likewise, depth is partitioned into *deep* and *shallow* according to a threshold on depth. Finally, origin is partitioned into *natural* and *artificial*.

All of the scales used in this work partition the many valued attributes into mutually exclusive partitions, so called *nominal scales* [Ganter and Wille, 1999]. Other possibilities for scales could also be considered in this context depending on the level of granularity required for the attributes in the application domain. In this work, however, we are not going to get into details about the diverse possibilities of scales.
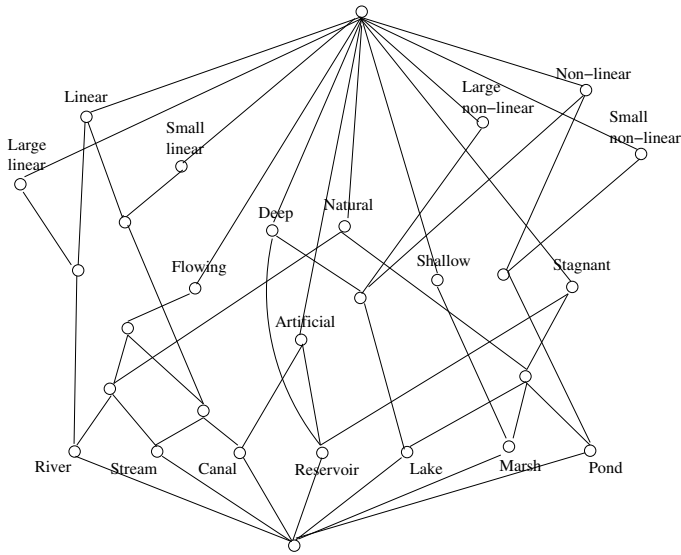

Figure 2: Concept lattice of inland water features.

The context of inland water features can, thus, be represented as the lattice shown in Figure 2.

Extracting definitions from the water feature formal context

The lattice in Figure 2 can be easily translated into definitions that could be handled by the supervaluation semantics, as discussed in Section 3 below.

In order to present the definitions of water features we first introduce the following functions on features. Let X be a water feature, then $linearity\_ratio(X)$, $average\_flow(X)$, $average\_depth(X)$, $max\_length(X)$ and $max\_width(X)$ are mappings from the feature X to real values representing, respectively, its linearity ratio, average flow, average depth, maximum length and maximum width. With these functions,

---

[3]It is worth pointing out that, for clarity, we maintained *linear* and *non-linear* along with *large (small) linear/non-linear*, even though the related attributes are never used separately in this work.

and appropriate thresholds, it is now possible to define the attributes of a inland water feature (F), as presented in Definition 1[4].

Definition 1 (Relations of the domain water feature)

- $Non\_linear(F) \equiv$
  $$linearity\_ratio(F) < linearity\_thresh$$
- $Linear(F) \equiv$
  $$linearity\_ratio(F) \geq linearity\_thresh$$
- $Large\_linear(F) \equiv linear(F) \wedge$
  $$max\_length(F) \geq linear\_length\_thresh$$
- $Small\_linear(F) \equiv linear(F) \wedge$
  $$max\_length(F) < linear\_length\_thresh$$
- $Large\_non\_linear(F) \equiv linear(F) \wedge$
  $$max\_length(F) \geq nonlinear\_length\_thresh \wedge$$
  $$max\_width(F) \geq nonlinear\_width\_thresh$$
- $Small\_non\_linear(F) \equiv non\_linear(F) \wedge$
  $$maximum\_length(F) < nonlinear\_length\_thresh$$
  $$\wedge maximum\_width(F) < nonlinear\_width\_thresh$$
- $Flowing(F) \equiv average\_flow(F) \geq flow\_thresh$
- $Stagnant(F) \equiv average\_flow(F) < flow\_thresh$
- $Deep(F) \equiv average\_depth(F) \geq depth\_thresh$
- $Shallow(F) \equiv average\_depth(F) < depth\_thresh$
- $Natural(F) \equiv \neg \, Artificial(F)$

In this paper we do not stipulate the actual thresholds for Definition 1. Future research may use machine or user surveying to determine appropriate default values corresponding to an intuitive classification.

We now can define the inland water features, as presented in Definition 2.

Definition 2 (Inland Water Features) *Let each of the objects of the inland water feature context be a predicate symbol. The terms of the ontology are defined as follows.*

$$
\begin{aligned}
\{River\}' &\equiv \{Linear, Large\_linear, \\
&\quad Flowing, Natural\} \\
\{Stream\}' &\equiv \{Linear, Small\_linear, \\
&\quad Flowing, Natural\} \\
\{Canal\}' &\equiv \{Linear, Small\_linear, \\
&\quad Flowing, Artificial\} \\
\{Reservoir\}' &\equiv \{Stagnant, Artificial, \\
&\quad Deep\} \\
\{Marsh\}' &\equiv \{Stagnant, Natural, Shallow\} \\
\{Lake\}' &\equiv \{Non\_linear, Large\_non\_linear, \\
&\quad Stagnant, Natural\} \\
\{Pond\}' &\equiv \{Non\_linear, Small\_non\_linear, \\
&\quad Stagnant, Natural\}
\end{aligned}
$$

The next section introduces the supervaluation semantics and discusses how it can handle variations in the meaning of the definitions above.

---

[4]In this paper predicate symbols start with an uppercase letter whereas functions are written with lower case letters.

3 Supervaluation Semantics

The formalisation of the inland water feature domain presented above used thresholds to define partitions on the values of attributes. However the implications of varying the values of these thresholds is not accounted for by the formal concept analysis. In fact, distinct sets of values for the thresholds may result in distinct (sometimes incompatible) definitions of the context objects. It is desirable to view such variations as *standpoints* on the vague definitions of features. We assume the supervaluation semantics [Fine, 1975] to provide a formal framework for standpoints on feature definition.

Supervaluation semantics views a vague language as a set of distinct precise versions of itself. Each of these versions is called a *precisification*. Formally, each precisification p is identified with a particular interpretation I_p of the language. A *supervaluation model* is defined as a set of precisifications. Therefore, given a supervaluation model Υ we can talk about propositions that are unequivocally true (i.e. are true in *every* interpretation $I_p \in \Upsilon$) and propositions that are *in some sense* true (i.e. are true in *some* interpretation $I_p \in \Upsilon$).

In this work the scales applied to attributes of the inland water feature domain compose one particular precisification. Therefore, the definitions of the inland water feature domain may vary according to the multiple meanings that the vague concepts might assume, or multiple values that the thresholds can assume. For instance, the two possible interpretations for the scenario in Figure 1 correspond to two distinct precisifications whereby the particular linearity thresholds dictate whether the scenario depicts three lakes or a meandering river. In fact, the apparatus of supervaluation semantics allows for logical relationships between vague concepts to be represented by quantifying over the (possibly infinite) space of precisifications [Bennett, 2001a].

Moreover, attributes in formal concepts can be used as predicates in the bodies of Horn clauses and the context objects as arguments in the clause heads of logic-based axioms. A set of such clauses can be fed in a system allowing logic reasoning about elements of the ontology. In other words, the methods of generating concept lattices provided by the formal concept analysis can be adapted to generate logical axioms that could be used within a logic-based reasoning system encoding supervaluation semantics.

The next section presents a prototype implementation of these ideas.

4 Implementation

In this section we discuss a prototype implementation of the inland water feature ontology with supervaluation semantics. The purpose of this prototype is to automatically individuate and classify different geographical features within two-dimensional images of river networks (e.g. river stretches, lakes and ponds). In a general way, the process of individuation and classification of water features involves the construction of the medial axis (or skeleton) of the river network and further translate it into a data structure. This data structure provides the basis for feature segmentation/classification. The elements involved in this process are described below.

4.1 Individuation and classification of features

By individuation in the context of water features we mean the process of identifying qualitative distinct stretches of river networks. In this work, this distinction is accomplished in three phases. First, the medial axis (or skeleton) of a river network is obtained by an off-the-shelf image processing method. Second, the axis thus obtained is processed into a data structure from which qualitative distinct features from the network can be identified by observing changes in the width along the axis. Finally, the distinct stretches are classified according to the definitions proposed in Sections 2.2 and 3 above.

Skeletonisation and data structure

The medial axis (or *skeleton*) [Blum, 1973] is commonly defined as the locus of the diameters of all maximal discs enclosed in a Jordan curve[5] that touch at least two points of the curve's boundary. The best known application of skeletonisation is in the field of image processing, where the skeletons of object's images are used as an abstraction of the two dimensional shape of the objects depicted, as surveyed in [Loncaric, 1998]. In fact, a linear algorithm has been proposed that computes the medial axis of simple polygons[6] in linear time [Chin *et al.*, 1995].

The idea of representing river networks by their medial axis is not new. To the best of our knowledge the first work that attempts such endeavour is [McAllister and Snoeyink, 2000], where the skeletons of river networks are used to identity opposite points in river banks, to connect the centreline of the networks with original river bank data, to calculate the areas of rivers and, finally, to automatically obtain the ordering of tributaries along the network. The problem of using the skeletons for the individuation and classification of water features (such as the identification of rivers and lakes in a single river network as proposed in the present paper), however, was only hinted in that work. We believe that the tasks of segmentation and classification of such features cannot be accomplished without a precise and consistent body of definitions such as those presented in sections 2.2 and 3.

In this work, we use a divergence-based method [Dimitrov *et al.*, 2000] as an off-the-shelf method for extracting the medial axis of river network snapshots. This method is based on considering a 2D region, representing the image of a shape boundary, as gradient fields of the region's distance transform. The skeletal points are, thus, identified as singularities in the *net outward flux* of the vector field relative to the region where each pair of points in opposite borders shrinks to a point.

Skeletons produced by the divergence-based method are transformed into an 8-connected list whereby each cell is subdivided into 9 registers. These registers encode each skeletal point in its middle register and pointers to any of its 8 cardinal neighbours. It is worth pointing out that, along with the skeletal point, the middle register also encodes the distance from this point to the nearest point in the shape's boundary (thus providing a measurement of the width of the water feature).

---

[5] A Jordan curve is any closed curve that does not intersect itself.

[6] A simple polygon is a single closed polygonal chain that do not intersect itself [de Berg *et al.*, 1997].

The task now is to sweep the 8-connected list, segmenting the river-network skeleton into stretches. These stretches are, thus, classified according to appropriate thresholds based on the water-feature ontology.

Integrated Classification and Individuation

Points where the skeleton branches and points where there are abrupt changes in the width of a river network (according to a threshold) are used to provide a means of obtaining a *proto-segmentation* of the data into *candidate water features*. The initial set of segments may further be concatenated or split according to the definitions in the ontology. This provides, as a result, the segmentation of the river network into water features, according to formal definitions.

In order to summarise these ideas in an algorithm we first need to introduce the ideas of *branching point*, *sample stretch* and *nexus point*.

A branching point is a discontinuity in the orientation of the boundary of a shape[7]. In a skeleton encoded as a 8-connected list, branching points are identified as a cell with more than two neighbours.

A *sample stretch* of a skeleton at a given point p is an ordered set of skeletal points with p as the mid-point and which extends a certain number of points either side of p. The *sample stretch ratio* is the ratio of the length of the sample stretch to the width of the water body at the sample point (i.e. the width at the mid-point of the stretch). We have found that a reasonable segmentation can be obtained by taking the sample stretch ratio as 1, so that the stretch has the same length as the with of the river.

We call a *nexus point* a point where two or more linear chains in the medial skeleton meet.

We can then define an *(approximately) linear sample stretch* of water as one where the minimum and maximum width along the stretch lie within a certain variation threshold relative to the width at the mid-point. More specifically a sample stretch s is linear according to a given *stretch linearity ratio threshold* (SLR_thresh) just in case:

$$(mid\_width(s)/min\_width(s) < SLR\_thresh) \wedge \quad (1)$$
$$(max\_width(s)/mid\_width(s) < SLR\_thresh)$$

The lake-river segmentation and classification algorithm can be summarised as follows:

> ALGORITHM: Lake-river segmentation and classification
> INPUT: A bitmap depicting a river network
> OUTPUT: A set of segments of the river network classified according to the water-feature ontology
>
> - first, a matrix output by the divergence-based skeletonisation is converted to a set of Prolog facts – one for each point in the skeleton. These identify the coordinates of the skeleton point and its distance from the nearest edge point of the water body;

- Prolog predicates are defined that test whether a point is an end point, a nexus point or a mid point of a chain. This is simply done by counting the number of points of the surrounding 8 pixes that are also skeleton points;
- For each chain of mid points the algorithm checks that the point is on a chain that extends for a distance in each direction that is at least equal to the width at that point;
- a list of points corresponding to the sample stretch is collected;
- Maximum and minimum widths along the sample stretch are computed;
- The ratio of the maximum and minimum widths to the width of the mid point are computed;
- If the ratio's lie within the given *stretch linearity ratio threshold* (according to statement 1) the point is marked as a point on a linear stretch; otherwise, it is marked as belonging to a non-linear segment.
- In the case of a linear segment, the linear-length/width thresholds are taken into account to decide whether the given segment is a river or a stream; in the case of a non-linear segment, the non-linear size thresholds are used to decide between a lakes or a pond.

Changes in the stretch linearity ratio threshold (and on the linear and non-linear size thresholds) allow for variations of the feature segmentation and classification. This provides a rigorous account of multiple standpoints in the classification of river, lakes, stream and pond features in a river network.

4.2 Prolog Implementation

The algorithm outlined in the previous section has been partially implemented in Prolog.

The following screen dump shows the output for an artificial example which was constructed to illustrate some key problems of lake/river discrimination. Note that the example contains a wide river and a much narrower river that could be interpreted either as a stream or as a river depending of an observer's standpoint. Moreover, the narrow river is broken by a small lake (or pond) and in this case the small lake/pond has a maximum diameter which is much smaller than the width of the wide river.

Figure 3 presents an output of the prototype implementation whereby linear stretches are indicated by thick black lines and non-linear parts of medial skeleton are shown as thin gray lines. A new classification, according to a distinct standpoint, is obtained by simply providing a new set of thresholds. Junctions between distinct rivers, however, were misclassified by the prototype, as shown by the junction between the thin and the wide rivers in Figure 3. [McAllister and Snoeyink, 2000] presents a possible solution for this problem, which shall be further exploited within the framework proposed in the present paper.

---

[7]Branching points are, in fact, *first-order shocks* as defined in [Siddiqi and Kimia, 1996].
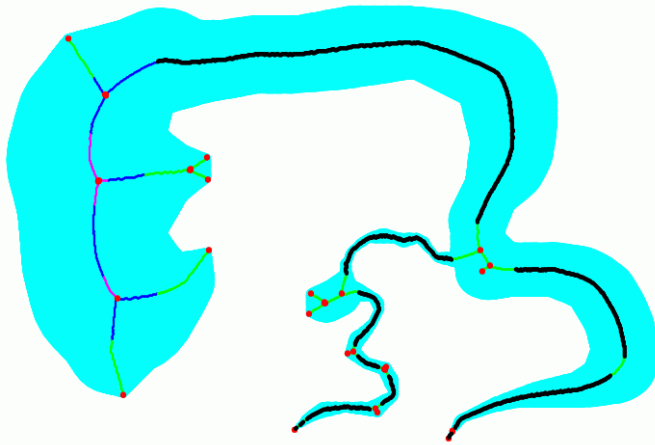
Figure 3: River stretch identification using sample stretches with sample length ratio 1 and sample variance ratio 1.5.

5 Further Work

Several aspects of the proposed ontology require further work to provide a more discriminating and flexible conceptualisation. In particular the articulation of size thresholds on water bodies requires further refinement to identify the relevant metrics. For instance, area and width measures should be distinguished.

More importantly, further work needs to be done on the integration between *individuation* and *classification*, which is only sketched in the present work. It will probably be necessary to introduce a further layer of definitions that alow features to be constructed by aggregating several components. Specifically, we would like to define the way in which a 'river' may be composed from several river stretches which may be joined in complex ways. An adequate treatment of this may require information on water flow to be integrated within the classification.

6 Conclusion

In this paper we presented an ontology for inland water features constructed using formal concept analysis and supervaluation semantics. The former allows a rigorous account of the domain structure by means of a complete lattice of concepts and their attributes. The latter formalises multiple interpretations of ontology terms owing to vagueness in the relative concept. The key for the integration of formal concept analysis and supervaluation semantics is the process of scaling. Scaling partitions many-value attributes of features into one-value ones, which permits a lattice representation of the domain. These partitions are defined by means of thresholds on the attribute values. Supervaluation semantics is used in this work to formalise variations on these thresholds modelling multiple possible interpretations of vague terms. Analysis of the hydrographic domain brings out the fact that vagueness not only affects the definitions of features but also how these features are individuated.

We proposed an algorithm for individuation and classification of water features according to the ontology presented.

The basic idea of this algorithm is to use the medial axis of a river network (extracted via an off-the-shelf divergence-based skeletonisation algorithm) to identify stretches of the network where there is relatively small variation in width. I.e., the process is to look at the variance along the axis checking that the maximum and minimum width values lie within some percentage variation threshold. This threshold can be regarded a supervaluation semantics parameter, and may vary according to any particular interpretation of the feature. This interaction between supervaluation semantics and the actual algorithmic feature individuation is the main contribution of this paper.

Acknowledgements

References

[Bennett, 2001a] B. Bennett. Application of supervaluation semantics to vaguely defined spatial concepts. In D.R. Montello, editor, *Spatial Information Theory: Foundations of Geographic Information Science; Proceedings of COSIT'01*, volume 2205 of *LNCS*, pages 108–123, Morro Bay, 2001. Springer.

[Bennett, 2001b] B. Bennett. What is a forest? on the vagueness of certain geographic concepts. *Topoi*, 20(2):189–201, 2001.

[Blum, 1973] H. Blum. Biological shape and visual science. *Journal of Theoretical Biology*, 38:205–287, 1973.

[Chin *et al.*, 1995] F. Chin, J. Snoeyink, and C.-A. Wang. Finding the medial axis of a simple polygon in linear time. In *Proc. of the 6^{th} Annual International Symposium of Algorithms and Computation*, volume 1004 of *LNCS*, pages 382–391. Springer Verlag, 1995.

[de Berg *et al.*, 1997] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computation Geometry*. Springer Verlag, 1997.

[Dimitrov *et al.*, 2000] Pavel Dimitrov, Carlos Phillips, and Kaleem Siddiqi. Robust and efficient skeletal graphs. In *Proc. CVPR*, June 2000.

[Elkan, 1993] C. Elkan. The paradoxical success of fuzzy logic. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-93)*, pages 698–703, 1993.

[Fine, 1975] K. Fine. Vagueness, truth and logic. *Synthése*, 30:263–300, 1975.

[Ganter and Wille, 1999] B. Ganter and R. Wille. *Formal Concept Analysis*. Springer Verlag, 1999.

[Loncaric, 1998] S. Loncaric. A survey of shape analysis techniques. *Pattern Recognition*, 31(8):983–1001, 1998.

[McAllister and Snoeyink, 2000] Michael McAllister and Jack Snoeyink. Medial axis generalization of river networks. *CaGIS*, 27(2):129–138, 2000.

[Siddiqi and Kimia, 1996] K. Siddiqi and B. Kimia. Toward a shock grammar for recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1996.