Grounding spatial named entities for information extraction and question answering

Jochen L. Leidner **Gail Sinclair** **Bonnie Webber**
School of Informatics
University of Edinburgh
2 Buccleuch Place
Edinburgh EH8 9LW
Scotland, UK
jochen.leidner@ed.ac.uk, csincla1@inf.ed.ac.uk, bonnie@inf.ed.ac.uk

Abstract

The task of named entity annotation of unseen text has recently been successfully automated with near-human performance.
But the full task involves more than annotation, i.e. identifying the scope of each (continuous) text span and its class (such as place name). It also involves grounding the named entity (i.e. establishing its denotation with respect to the world or a model). The latter aspect has so far been neglected.

In this paper, we show how geo-spatial named entities can be grounded using geographic coordinates, and how the results can be visualized using off-the-shelf software. We use this to compare a "textual surrogate" of a newspaper story, with a "visual surrogate" based on geographic coordinates.

1 Introduction

The task of named entity annotation of unseen text has recently been successfully automated, achieving near-human performance using machine learning (Zheng and Su, 2002). But many applications also require **grounding** – i.e., associating each classified text span with a referent in the world or some model thereof. The current paper discusses spatial grounding of named entities that may be referentially ambiguous, using a *minimality heuristic* that is informed by external geographic knowledge sources. We then apply these ideas to the creation of "visual surrogates" for news articles.

This paper is structured as follows: Section 2 discusses how spatial named entities can be grounded and how this interacts with their extraction and applications. Section 3 describes a geo-spatial resolution algorithm. Section 4 shows how maps can be automatically constructed from named-entity tagged newswire text using resolved place names, hence introducing a new, graphical document surrogate. Section 5 deals with the usefulness of grounded named entities for question answering. Section 6 presents some related work, and Section 7 concludes this paper.

2 Spatial Grounding

Gazetteers are large lists of names of geographic entities, usually enriched with further information, such as their class (e.g., town, river, dam, etc.), their size, and their location (i.e. with respect to some relative or absolute coordinate system such as longitude and latitude).

Appendix A identifies some publicly available sources. *UN-LOCODE* is the official gazetteer by the United Nations; it is also freely available from the UNECE Web site[1] and contains more than 36 000 locations in 234 countries (UNECE, 1998). The *Alexandria Gazetteer* (Smith et al., 1996; Frew et al., 1998) is another database of geographical entities, including both their coordinates and relationships such as: `in-state-of`, `in-province-of`, `in-county-of`, `in-country-of`, `in-region-of`, `part-of` and `formerly-known-as`.

To date, Named Entity Recognition (NER) has only used gazetteers as evidence that a text span could be some kind of place name (LOCATION), even though their finite nature makes lists of names of limited use for classification (Mikheev et al., 1999). Here we use them for spatial grounding – relating linguistic entities of subtype LOCATION (Grishman and Sundheim, 1998) to their real-world counterparts.

"World Atlases" and the gazetteers that index them are not the only resources than can be used for grounding spatial terms. In biomedicine, there are are several brain atlases of different species, using various different techniques, and focussing on both normal and disease state; as well as a digital atlas of the human body

[1] http://www.unece.org/cefact/locode/service/main.htm

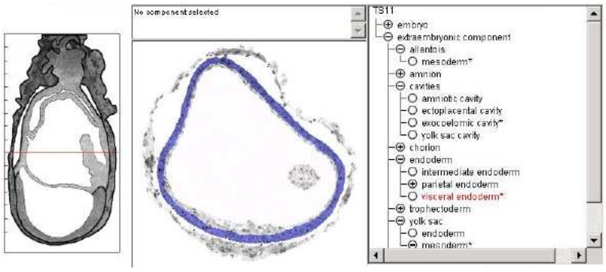

Figure 1: Grounding an XML Ontology in Voxels: The Mouse Atlas (Baldock et al., 1999).

based on data from the Visible Human project. Such atlases and the nomenclatures that label their parts, provide an important resource for biomedical research and clinical diagnosis. For example, the *Mouse Atlas* (Ringwald et al., 1994) comprises a sequence of 3D (volumetric) reconstructions of the mouse embryo in each of its 26 Theiler States of development. Indexing it is an `part-of` hierarchy of anatomical terms (such as `embryo.organsystem.cardiovascularsystem-.heart.atrium`), called the Mouse Anatomical Nomenclature (MAN). Each term is mapped to one or more sets of adjacent voxels[2] that constitute the term's denotation in the embryo. Figure 1 illustrate this linkage (using 2D cross-sections) in the EMAGE database.[3]

Just as one might find it useful for information extraction or question answering to ground grographic terms found in previously unseen text, one may also find it useful to ground anatomical terms in previously unseen text. One example of this would be in providing support for the curation of the Gene Expression Database (GXD).[4] This support could come in the form of a named entity recognizer for anatomical parts in text, with grounding against the Mouse Atlas, using the gazetteer-like information in the MAN.

So what is the relationship between a place name gazetteer like UN-LOCODE and the Mouse Atlas? The MAN is structured in a similar `part-of` hierarchy to that of geographical locations:

USA	embryo
California	organ system
San Mateo County	cardiovascular system
Redwood City	heart

Because both gazetteers like UN-LOCODE and biomedical atlases like the Mouse Atlas provide spatial grounding for linguistic terms (Figure 2), both can be used to reason about spatio-temporal settings of a discourse, for instance, to resolve referential ambiguity.

[2] Pixels are points in the 2D plane $\langle$x, y$\rangle$; voxels are 3D generalizations of pixels $\langle x, y, z \rangle$.

[3] http://genex.hgu.mrc.ac.uk/Emage/database/intro.html

[4] http://genex.hgu.mrc.ac.uk/Resources/GXDQuery1/

3 Place-Name Resolution for Information Extraction

There are many places that share the same (*Berlin, Germany* ∼ *Berlin, WI, USA*) or similar names (*York, UK* ∼ *New York, USA*), usually because historically, the founders of a new town had given it a similar or the same name as the place they emigrated from.

When ambiguous place names are used in conversation or in text, it is usually clear to the hearer what specific referent is intended. First, speaker and hearer usually share some extra-linguistic context and implicitly adhere to Grice's Cooperative Principle and the "maxims" that follow, which require a speaker to provide more identifying information about a location that the recipient is believed to be unfamiliar with. Secondly, linguistic context can provide clues: an accident report on the road between Perth and Dundee promotes an interpretation of Perth in Scotland, while an accident on the road between Perth and Freemantle promotes an interpretation of Perth in Western Australia. Computers, which are bound to select referents algorithmically, can exploit linguistic context more easily than extra-linguistic context, but even the use of linguistic context requires (as always) some subtle heuristic reasoning.

Grounding place names mentioned in a text can support effective visualization – for instance, in a multimedia document surrogate that contains textual, video and map elements (e.g. in a question answering scenario), where we want to ensure that the video shows the region and the map is centered around the places mentioned.

To make use of linguistic context in resolving ambiguous place names, we apply two different *minimality heuristics* (Gardent and Webber, 2001). The first we borrow (slightly modified) from work in automatic word sense disambiguation (Gale et al., 1992), calling it "one referent per discourse". It assumes that a place name mentioned in a discourse refers to the same location throughout the discourse, just as a word is assumed to be used in the same one sense throughout the discourse. Neither is logically necessary, and hence both are simply *interpretational biases*.

The second minimality heuristic assumes that, in cases where there is more than one place name mentioned in some span of text, the smallest region that is able to ground the whole set is the one that gives them their interpretation.[5] This can be used to resolve referential ambiguity by proximity: i.e., not only is the place name *Berlin* taken to denote the same Berlin throughout a discourse unless mentioned otherwise,[6] but so does a *Potsdam* men-

[5] Probably the smaller the span, the more often valid will this heuristic be.

[6] This paper is a rare exception due to its meta-linguistic nature.

Gazetteer	Named Entity	Grounding	Real World
UN-LOCODE	`Novosibirsk`	$\langle 55°02' \text{ N}; 82°55' \text{ E}\rangle$ (longitude/latitude)	(Novosibirsk)
Mouse Atlas	`atrium`	$\{\langle 345, 2, 345\rangle; \ldots\}$ (set of voxels)	(part of the heart)

Figure 2: Comparison between Spatial Grounding in UN-LOCODE and the Mouse Atlas.

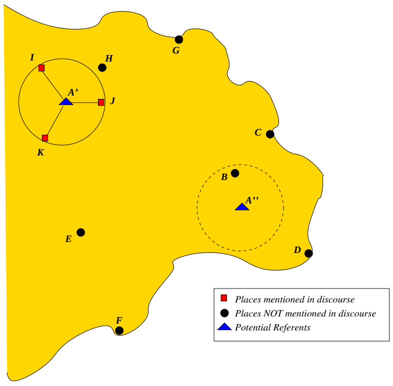

Figure 3: A Place-Name Resolution Method.

tioned together with a *Berlin* uniquely select the capital of Germany as the likely referent from the set of all candidate Berlins.[7]

To illustrate this "spatial minimality" heuristic, consider Figure 3: Assume that a mention of place A in a text could either refer to A' or A''. If the text also contains terms that ground unambiguously to I, J, and K, we assume the referent of A is A' rather than A'' because the former leads to a smaller spatial context.

To use this "spatial minimality" heuristic, we start by extracting all place names using a named entity recognizer. We then look up the "confusion set" of potential referents for each place name, e.g. for *Berlin*: { Berlin, FRG (German capital); Berlin, WI, USA; Berlin, NJ, USA; Berlin, CT, USA; Berlin, NH, USA; Berlin, GA, USA; Berlin, IL, USA; Berlin, NY, USA; Berlin, ND, USA; Berlin, NJ, USA }. Each member of the set of potential referents is associated with its spatial coordinates (longitude/latitude), using a gazetteer. We then compute the cross-product of all the confusion sets. (Each member of the cross-product contains one potential referent for each place name, along with its spatial coordinates.) For each member of the cross-product, we compute the area of the minimal polygon bounding all the potential referents, and select as the intended interpretation, the one with the smallest area.[8] The resulting behaviour is shown in Figure 4: depending on contextually mentioned other places, a different Berlin is selected.

{ *Berlin*; Potsdam } $\mapsto$ Berlin, FRG (Germany)
{ Fairburn; *Berlin* } $\mapsto$ Berlin, WI, USA
{ West *Berlin*; Bishops; Dicktown } $\mapsto$ Berlin, NJ, USA
{ Kensington; *Berlin*; New Britain } $\mapsto$ Berlin, CT, USA
{ Copperville; *Berlin*; Gorham } $\mapsto$ Berlin, NH, USA
{ Moultrie; *Berlin* } $\mapsto$ Berlin, GA, USA
{ *Berlin*; Prouty } $\mapsto$ Berlin, IL, USA
{ *Berlin*; Berlin Center; Cherryplain } $\mapsto$ Berlin, NY, USA
{ Medberry; *Berlin* } $\mapsto$ Berlin, ND, USA

Figure 4: Spatial Reference Resolution Using Spatial Minimality.

The value of this heuristic needs to be assessed quantitatively against various types of text.

In resolving anatomical designators in text, we may employ a variation of the spatial minimality heuristic, based on the fact that no listing will ever be complete with respect to all the existing or new-minted synonyms for anatomical terms.

When grounding the anatomical terms in the text

> *In subsequent stages until birth, cytokeratin 8 continues to be expressed in embryonic taste buds distributed in punctuate patterns at regular intervals along rows that are symmetrically located on both sides of the* ***median sulcus*** *in the dorsal anterior developing tongue.*

we find no *median sulcus* within the MAN, only *alveolar sulcus*, *optic sulcus*, *pre-otic sulcus*, *sulcus limitans* and *sulcus terminalis*. We just assume that all anatomical terms refer to previously recognized anatomical entities, just as we assume that all geographic terms refer to existing geographic entities and not, for example, some new town called "Berlin" or "London" that is not yet in the gazetteer. Hence *median sulcus* is assumed to be a synonym for one of the five *sulci* given in the MAN. At this point, we can invoke the spatial minimality heuristic, looking for the minimal bounding space that includes *tongue* and one of the five *sulci*, here yielding "sulcus terminalis". Thus the spatial minimality heuristic is here

[7] despite the fact that most places named *Berlin* are in the United States

[8] One can approximate this, either by computing the sum of pairwise point-point distances, or symbolically, using a hierarchical gazetteer's relations, such as `in-region-of`.

used with other assumptions to resolve missing or previously unseen terms.

4 Visualization of Geo-Spatial Aspects in Narrative

The usefulness of visual representations to convey information is widely recognized (cf. (Larkin and Simon, 1987)). Here, we use the grounding of named entities in news stories to create a visual surrogate that represents their "spatial aboutness".

Two news stories were selected from online newspapers on the same day (2003-02-21): one story (Appendix B) reports the tragic death of a baby from London in a Glasgow hospital despite flying it to a Glasgow specialist hospital in the Royal aircraft (BBC News, 2003), and the other report (Appendix C) describes the search of the Californian police for a pregnant women from Modesto, CA, who has disappeared (The Mercury News, 2003).

We use the term "surrogate" to refer to a partial view of a text (e.g. (Leidner, 2002)). Figure 5 shows a textual surrogate in the form of all place names found in a text: an analyst who wants to get a quick overview about the locations involved in some item of news reportage, to decide its local interest or relevance, might find such a surrogate helpful, although the source would still have to be skim-read.

Story A

... Scotland ... Tooting ... London ... Glasgow ... London ... Glasgow ... Northolt ... Glasgow ... Britain ... Prestwick ... Tooting ... Glasgow ... UK ... (Glasgow) ...

Story B

Modesto ... (Southern California) ... (Modesto) ... Los Angeles ... Sacramento ... Berkeley (Marina) ... Fresno ... Oakland ... Modesto ... Los Angeles ... Southern California ... Modesto ... Southern California ... New York ... Long Island ...

Figure 5: A Textual Geo-Spatial Document Surrogate for the Stories in Appendices B and C.

We now compare this "baseline" textual surrogate to a graphical map representation that draws on the algorithm introduced before. Our simple visualisation method comprises the following components (Figure 6): an (open-domain) news item is fed into **locotagger**, our simple named entity tagger for place names based on UN-LOCODE.[9] It recognises location names, resolves multireferential place names and looks up the coordinates:

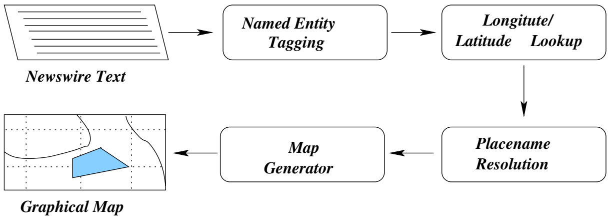

Figure 6: System Architecture.

```
Scott, more than a dozen news crews
from <ENAMEX type="LOCATION" longitude=
"-118.25" latitude="34.05">Los
Angeles</ENAMEX> to <ENAMEX type=
"LOCATION" longitude="121.5"
latitude="38.583333">Sacramento
</ENAMEX> camped out front.
```

From the text we obtain a vector of types of all spatial named entities with their frequency of occurrence in the text:

$$\begin{pmatrix} UK: & 1 \\ Scotland: & 1 \\ Tooting: & 2 \\ London: & 2 \\ Glasgow: & 5 \\ Northolt: & 2 \\ Prestwick: & 1 \\ Britain: & 1 \end{pmatrix} \quad \begin{pmatrix} Modesto: & 3 \\ SouthernCalifornia: & 2 \\ LosAngeles: & 2 \\ Sacramento: & 1 \\ Berkeley: & 1 \\ Fresno: & 1 \\ Oakland: & 1 \\ NewYork: & 1 \\ LongIsland: & 1 \end{pmatrix}$$

For simplicity, we drop those that correspond to regions (which are represented by sets of points) and feed the remaining list of point coordinates (corresonding to villages and cities) into a map generator to generate a Mercator projection of the geographical area that includes all the points plus 10% of the surrounding area. For this, The Generic Map Tools (GMT)[10] were used, in this case via HTTP.[11]

Figure 7 shows the resulting map for the story in Appendix B[12] Figure 9 shows the map for the story in Appendix C. Clearly, such a visual surrogate is superior with respect to comprehension time than the textual surrogate presented before. It is interesting so see what happens if we leave out the final paragraph for the map creation (Figure 8): we obtain a "zoomed-in" version of the map. This turns out to be the case for many stories and is due to the convention of news reportage to close a report with linking the narrative to similar events in order to present the event in a wider context.

[9] For the experiment reported here, we also used data from http://www.astro.com/cgi/aq.cgi?lang=e.

[10] http://gmt.soest.hawaii.edu/

[11] http://stellwagen.er.usgs.gov/mapit/

[12] Place names that are too small to be contained in the gazetteer (Tooting/Northold) are ignored.

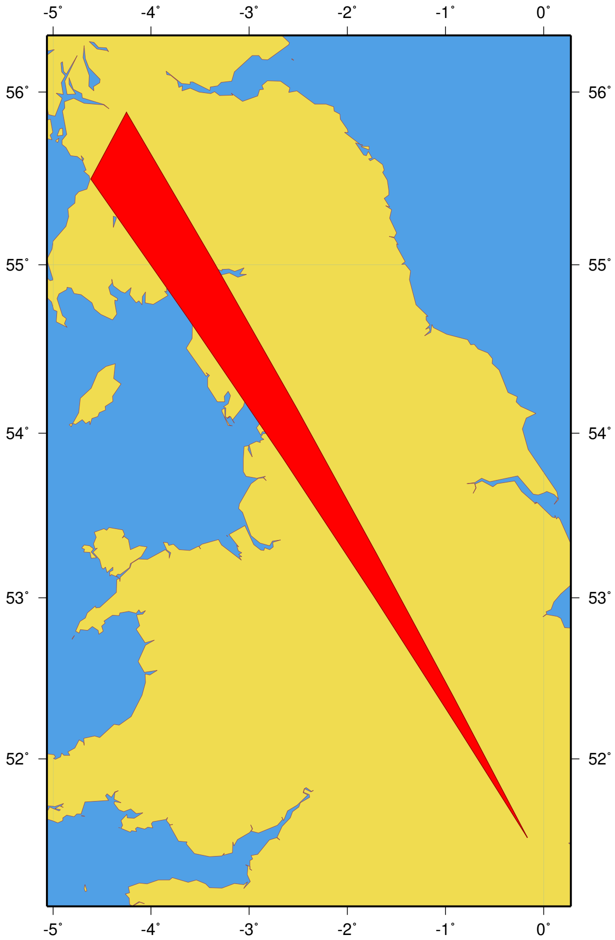

Figure 7: Automatic Visualization of Story B: A Baby Flown from London to Glasgow for Medical Treatment Dies there.

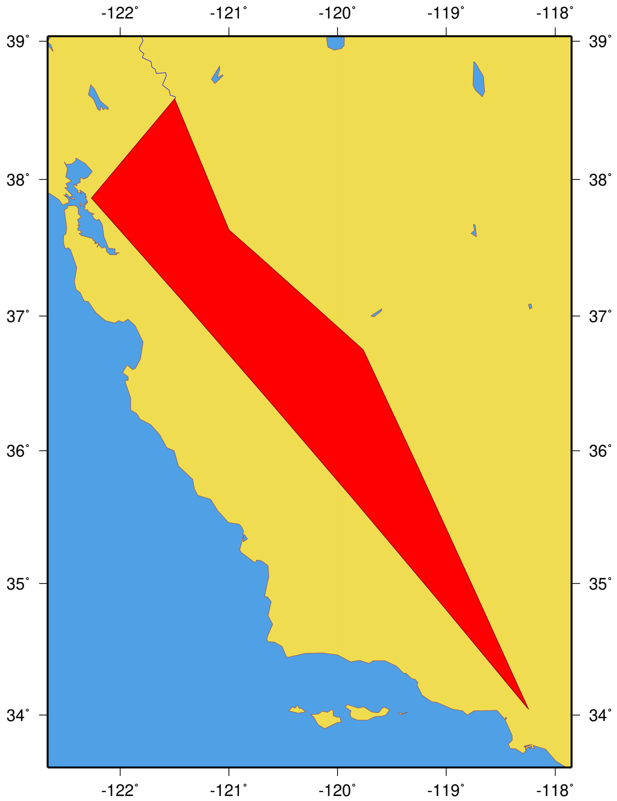

Figure 8: Automatic Visualization of Story C: A Pregnant Woman is Missing in Modeno, CA (Local View; Final Paragraph Excluded).

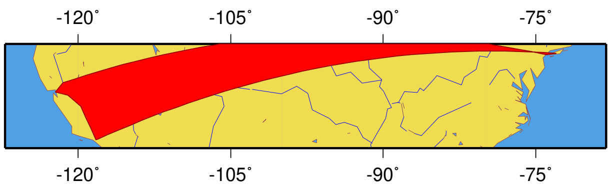

Figure 9: Story C: The Final Paragraph Places the Event in Context (Global View; Complete Story).

(Shanon, 1979) discusses how the granularity of the answers to *where*-questions depends on the reference points of speaker and listener (*Where is the Empire State Building?* – (a) *In New York*, (b) *In the U.S.A*, (c) *On 34th Street and 3rd Avenue*); the map generation task depends on such levels of granularity in the sense that to create a useful map, entities that belong to the same level of granularity or scale should be marked (e.g. city–city rather than village–continent).

5 Question Answering

Using grounding knowledge in gazetteers also enables us to answer questions in natural language more effectively:

1. *What is **X**?*
   Q: *What is Kangiqsualujjuaq?*
   A: *Kangiqsualujjuaq is a place approximately 1500 kilometers north of Montreal, Canada.* (For some place names, many humans cannot tell they refer to places.)

2. *Where is **X**?*
   Q: *Where is Cannes located?*
   A: [should yield a surrogate based on textual descriptions generated from the gazetteer relations:
   X `is-type` Y, X `part-of` Z and the coordinates, plus a map as generated above, with additional images, e.g. from satellites or picture search engines as available.]

3. *What **X** is **Y** part of?*
   Q: *What is 'Bad Bergzabern' part of?*
   A: *Bad Bergzabern is part of the Federal Republic of Germany.*
   Q: *Is Andorra la Vella part of Spain?*
   A: *No, Andorra la Vella belongs to Andorra.*

4. *How far is **X** from **Y**?*
   Q: *How far is Cambridge from London?*
   A: *The distance between London, England, United Kingdom and Cambridge, England, United Kingdom is 79 km (49 miles or 43 nautical miles).*
   Note here that the spatial minimality heuristic resolves Cambridge and London to places in the UK rather than, say, London, Ontario, Canada and Cambridge, Mass., USA. However the answer makes clear the precise question being adressed, so the user can follup up with a different question if this was not what he or she intended.

Since sophisticated gazetteers are available, answering such questions should not be based on textual extraction from Internet sources, but be based on the gazetteers directly, which reduces noise.

6 Related Work

Whereas a lot of work has been done in the area of spatial databases (Shekhar et al., 1999; Freytag et al., 2000; Paredaens, 1995; Paredaens and Kuijpers, 1998), little attention seems to have been given so far to the problem of generating maps from text or annotating maps with information elicited from unstructured documents.

The work presented here perhaps most closely resembles that of (Mackinlay, 1986; Casner, 1990; Roth and Hefley, 1993) who describe systems that generate multimedial representations, potentially including maps, but from a formalized knowledge source rather than unstructured text, and (Waldinger et al., 2002), who describe Geo-Logica, a system that can transform requests to see a place (formulated in natural language) into a three-dimensional aerial view.

7 Summary and Conclusion

Rich gazetteers can be helpful in determining unique geospatial referents and have many applications.

Starting with a description of how spatial named entities can be grounded in the physical world using gazetteer knowledge in two different domains, we have described a heuristic method to resolve place names using such gazetteers.

We have then applied our method to the task of drawing maps from text automatically. These (geo-)graphical document surrogates give an overview about *where* a reported news event takes place. We do not know of previous attempts to create geographical maps automatically from unconstrained newswire text.

Many venues for further research can be conceived, for instance, a finer granularity of analysis could reveal more named entity classes (including e.g. brigdges and other artifacts) and their relative ortientation. Maps as created here could also be used to produce an animation from multiple documents (with different time-stamps) on the same topic to illustrate how the importance of places shifts over time, which could aid analysts in their timely identification of regions of crisis.

Acknowledgments. The authors are grateful to the United Nations Economic Commission for Europe (UN-ECE) for creating UN-LOCODE, to Johanna Moore, Michael Piotrowski, IBM UK, The National e-Science Centre, Edinburgh, Richard Baldock (MRC Genetics Unit) and the Edwina Gazetteer Project for discussions and support. The three anonymous reviewers helped to improve the quality of this paper. We would also like to acknowledge the financial support for the first author of the German Academic Exchange Service (DAAD) under grant D/02/01831, of Linguit GmbH (research contract UK-2002/2), and the School of Informatics, University of Edinburgh.

References

Richard A. Baldock, Chrisophe Dubreuil, Bill Hill, and Duncan Davidson. 1999. The Edinburgh Mouse Atlas: Basic structure and informatics. In S. Levotsky, editor, *Bioinformatics Databases and Systems*, pages 102–115. Kluwer Academic Press.

BBC News. 2003. Royal Mercy Flight Baby Dies. Online (Accessed Friday 2003-02-21).

S. M. Casner. 1990. *A Task-Analytic Approach to the Automated Design of Information Graphics*. Ph.D. thesis, University of Pittsburgh.

J. Frew, M. Freeston, N. Freitas, L. L. Hill, G. Janee, K. Lovette, R. Nideffer, T. R. Smith, and Q. Zheng. 1998. The Alexandria Digital Library architecture. In *Research and Advanced Technology for Digital Libraries: Second European Conference (ECDL'98), Lecture Notes in Computer Science, vol. 1513*, pages 19–23.

Johann C. Freytag, M. Flasza, and Michael Stillger. 2000. Implementing geospatial operations in an object-relational database system. In *Statistical and Scientific Database Management*, pages 209–219.

W. Gale, K. Church, and D. Yarowsky. 1992. One sense per discourse. In *Proceedings of the Fourth DARPA Speech and Natural Language Workshop*, pages 233–237.

Claire Gardent and Bonnie Webber. 2001. Towards the use of automated reasoning in discourse disambiguation. *Journal of Logic Language and Information*, 10:487–509.

R. Grishman and B. Sundheim, editors. 1998. *Proceedings of the Sixth Message Understanding Conference (MUC-7)*. Morgan Kaufmann.

J. Larkin and H. Simon. 1987. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11:65–99.

Jochen L. Leidner. 2002. *Discovery of Artefacts in Scientific Prose Using Machine Learning*. M.Phil. dissertation, University of Cambridge.

Jock D. Mackinlay. 1986. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5:110–141.

Andrei Mikheev, Marc Moens, and Claire Grover. 1999. Named entity recognition without gazetteers. In *Proceedings of the Annual Meeting of the European Chapter of the Association for Computational Linguistics (EACL'99)*, Bergen, Norway.

Jan Paredaens and Bart Kuijpers. 1998. Data models and query languages for spatial databases. *Data Knowledge Engineering*, 25(1-2):29–53.

Jan Paredaens. 1995. Spatial databases, the final frontier. In G. Gottlob and M. Y. Vardi, editors, *Proceedings of the 5th International Conference on Database Theory (ICDT), Lecture Notes in Computer Science*, volume 893, pages 14–32. Springer-Verlag.

M. Ringwald, R. A. Baldock, J. Bard, M. H. Kaufman, J. T. Eppig, J. E. Richardson, J. H. Nadeau, and Davidson D. 1994. A database for mouse development. *Science*, 265:2033–2034.

S. Roth and W. Hefley. 1993. Intelligent multimedia presentation systems: Research and principles. In Mark Maybury, editor, *Intelligent Multimedia Interfaces*, pages 13–58. AAAI Press.

Benny Shanon. 1979. Where questions. In *Proceedings of the 17th Annual Meeting of the Association for Computational Linguistics*.

Shashi Shekhar, Sanjay Chawla, Siva Ravada, Andrew Fetterer, Xuan Liu, and Chang tien Lu. 1999. Spatial databases – Accomplishments and research needs. *Knowledge and Data Engineering*, 11(1):45–55.

T. Smith, D. Andresen, L. Carver, R. Dolin, C. Fischer, J. Frew, M. Goodchild, O. Ibarra, R. Kemp, R. Kothuri, M. Larsgaard, B. Manjunath, D. Nebert, J. Simpson, A. Wells, T. Yang, and Q. Zheng. 1996. A digital library for geographically referenced materials. *IEEE Computer*, 29:54–60.

The Mercury News. 2003. News Crews Wait and Watch as Police Search Home of Missing Woman. Online (Accessed Friday 2003-02-21).

UNECE. 1998. LOCODE – Code for Trade and Transport Locations. Technical Report and UNECE Recommendation 16, United Nations Economic Commission for Europe.

R. Waldinger, M. Reddy, C. Culy, J. Hobbs, and J. Dungan. 2002. Deductive response to geographic queries. In *GIScience 2002*, Boulder, CO.

GuoDong Zheng and Jian Su. 2002. Named entity tagging using an HMM-based chunk tagger. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 209–219, Philadelphia.

A Some Online Gazetteer Sites

Alexandria Gazetteer
http://www.alexandria.ucsb.edu/gazetteer
GIS WWW Resource List
http://www.geo.ed.ac.uk/home/giswww.html
Gazetteer of Planetary Nomenclature
http://wwwflag.wr.usgs.gov/USGSFlag/Space/nomen/nomen.html
Getty Thesaurus of Geographic Names
http://shiva.pub.getty.edu/tgn_browser/
NIMA GEOnet Names
http://164.214.2.59/gns/html/index.html
Place Name Servers
http://www.asu.edu/lib/hayden/govdocs/maps/geogname.htm
U.S. Census Gazetteer
http://www.census.gov/cgi-bin/gazetteer/
U.S. Geological Survey Geographic Names IS
http://www-nmd.usgs.gov/www/gnis/

All URLs cited as of Friday 2003-02-21.

B Story *Royal Mercy Flight Baby Dies*

BBC NEWS UK EDITION

Last Updated: Thursday, 20 February, 2003, 11:36 GMT

Email this to a friend | Printable version

Royal mercy flight baby dies

A baby who was flown to Scotland for specialist treatment on an aircraft normally used to carry the Royal Family has died in hospital.

The royal aircraft made the flight

Baby Khola was born at St George's Hospital in Tooting, south London, last month with an undisclosed rare condition.

She was transferred to the Royal Hospital for Sick Children at Glasgow's Yorkhill Hospital on 23 January.

However, Khola - who was just over a month old - died on Monday.

"The last few days have been very sad for Khola and her parents," said a spokesman for the hospital.

"Our thoughts are with them at this tragic time."

> **"We wish to thank everyone who prayed for our daughter"**
>
> Khola's parents

Khola's parents, who wish to remain anonymous, thanked the Queen for allowing the use of the jet.

"We would like to say thank you for all the care Khola has received," they said in a statement.

"We wish to thank the hospital staff in Glasgow and London, the team who transported Khola, and the Queen for allowing the use of her plane. We also wish to thank everyone who prayed for our daughter."

Khola was taken to Yorkhill Hospital An RAF Hercules had initially been scheduled to take the baby to Glasgow last month. However, it burst a tyre as it landed at Northolt.

It was then decided to fly the baby and the medical team to Glasgow in an aircraft from the Royal Squadron.

The four-engine British Aerospace 146 was the same aircraft used for the Queen's Jubilee Tour of Britain last year.

A Royal Navy helicopter at Prestwick was mobilised and flown in bad weather to RAF Northolt where the specialist medical team on board was transferred to St George's Hospital in Tooting.

The baby was stabilised in an incubator throughout the night before the transfer to Glasgow. Doctors at Yorkhill said Khola had been in need of respiratory support.

Four hospitals in the UK can provide extracorporeal membrane oxygenation (ECMO) for breathing difficulties, but Glasgow was the only hospital able to accommodate Khola.

The process involves taking the baby's blood out of the body, through an artificial lung and pumping it back in, allowing the baby's lungs to rest. □

C Story *News Crews Wait and Watch as Police Search Home of Missing Woman*

Posted on Thu, Feb. 20, 2003 — The Mercury News

News crews wait and watch as police search home of missing woman

SCORES OF REPORTERS HOPE FOR MORSELS OF INFORMATION

By Lori Aratani and Linda Goldston
Mercury News

MODESTO - The search for Laci Peterson took on a carnival air Wednesday as dozens of reporters and 18 satellite trucks staked out the missing woman's home while police completed a two-day search for evidence.

Associated Press

Laci Peterson.

The big news of the day? Police measured Laci and Scott Peterson's driveway.

They declined to explain why.

But there was so much jostling for tidbits that police strung yellow crime-scene tape around the home to keep reporters back.

The media circus was only the latest in a long line that have sprung up around high-profile criminal cases, where the news is scant but the demand for information – even the smallest morsel – soars off the charts. "People have been interested in this type of story since pre-biblical times," said Joe Saltzman, associate dean at the Annenberg School for Communications at the University of Southern California. "We live in a global village where everyone is our neighbor, and this is an interesting story about a man whose pregnant wife is kidnapped."

As Modesto police completed their second day combing through the home that missing mother-to-be Laci Peterson shared with her husband, Scott, more than a dozen news crews from Los Angeles to Sacramento camped out front. Many of the reporters huddled under a tent hastily erected by one reporter as it rained and hailed.

Detective Doug Ridenour said Wednesday that investigators had completed their work, removing about 95 bags of evidence. On Tuesday, Amy Rocha, the sister of the missing woman, accompanied detectives into the modest three-bedroom home for about two hours. But on Wednesday, investigators worked alone – completing their task about 5 p.m. They also planned to do a walk-through with Scott Peterson.

"We've been able to accomplish a lot," Ridenour told reporters at the Peterson house. "But here at this point we just don't have the significant evidence we need to find Laci or to move in another direction."

Ridenour would not reveal when detectives would analyze what they took away or when they would release a detailed description of what they found.

He also repeated a familiar refrain: Scott Peterson is not a suspect, but he has not been ruled out as one. Suspicion has hovered over Peterson since his wife vanished Dec. 24, a day he said he was fishing at the Berkeley Marina 85 miles away.

[...]

Kornberg, who is best known for representing a man the public loved to hate—Joey Buttafuoco, who had a torrid affair with Long Island teenager Amy Fisher – said the only way to end such sideshows is to convert the American system of justice "to the English system, where you announce the arrest of an individual and then you announce the verdict when it's over." □