UNCERTAINTY IN GIS SUPPORTED ROAD EXTRACTION

George Vosselman
Faculty of Geodetic Engineering
Delft University of Technology
Thijsseweg 11, NL-2629 JA Delft
The Netherlands
e-mail: g.vosselman@geo.tudelft.nl

Commision III, Working Group 3

KEY WORDS: Uncertainty, GIS Updating, Change_Detection, Image Interpretation, Statistics, Modeling

ABSTRACT

Most image interpretation methods suffer from a lack of knowledge about the scene contents. In the case of updating road databases by interpretation of aerial images, the outdated database is a valuable source of knowledge. Both the old database and the generic, often heuristic, models of the scene objects contain uncertain data. This paper deals with different ways of representing this uncertainty and combining data from different uncertain knowledge sources. It is shown how uncertainty affects the quality of the road extraction and how modeling uncertainty can contribute to a better result.

1 INTRODUCTION

After one or two decades of digital mapping many countries are completing or already have completed extensive digital topographic databases, both in small and large scales. The updating of these databases will therefore be one of the major topics for many national mapping organisations. Although the efforts for updating are expected to be significantly lower compared to those of the initial database creation, there is a widespread interest in developing tools for (semi-)automatic mapping in digital images to further reduce the costs.

Contrary to initial expectations, the interpretation of aerial images for the purpose of mapping roads or houses has shown to be an extremely difficult task to automate. Simple schemes of thresholding, edge detection and grouping are clearly insufficient. One has come to realize that a human operator exploits an enormous amount of knowledge to interpret images and that an extensive knowledge base will be necessary to even partially solve tasks like understanding the complex aerial images.

Several ways are being explored to incorporate more knowledge into the image interpretation. A first one is to use the knowledge base of a human operator. By using the computer's speed for low level image processing and recognition of simple patterns and leaving the more difficult interpretation tasks to the operator, interactive algorithms are a very attractive way to speed up the mapping process. Further improvement should be possible if one can provide computer algorithms with detailed specific and generic models of the objects and their interrelationships that are encountered in aerial images. How to represent this knowledge is one of the major research issues. A third important source of knowledge is, of course, the database with the objects that have been mapped at a previous occasion and now need to be updated. The outdated databases not only outline many objects that only need to be verified instead of detected and measured, but also supply the context within which new objects may be found.

Since the modeling of the scene contents is an extremely difficult task, it is clear that a considerable degree of automation in mapping can only be achieved by combining most of the knowledge available. As most of the future mapping projects will be concerned with updating, the outdated databases should be considered a very important source of knowledge.

All sources of knowledge will, however, contain errors. These errors will influence the interpretation process. In order to assess the quality of the image interpretation results we will therefore need to describe the quality of the information for each individual source. Furthermore it is necessary to describe all steps in the interpretation process such that the quality of the different sources can be propagated to the final result.

This paper is about all kinds of uncertainties that are encountered during the updating of a topographical database by interpreting digital aerial images. In particular, we will focus on the extraction of roads using the context of an outdated road database.

Representation of uncertainty in data and processing of uncertain data are relevant to a wide range of disciplines. Especially in the field of artificial intelligence, many research efforts have been devoted to these topics [Kanal et al., 1986]. More recently, user requirements and standardisation efforts have also lead to better descriptions for uncertainty in (or quality of) geographical information [Guptill and Morrison, 1995].

After briefly describing the data and processing steps needed for the updating of road maps (section 2), we therefore first give an overview of representations of uncertainty as they have been developed in different disciplines (section 3). In section 4 some examples are given of how these uncertainty descriptions relate to the data used for road map updating.

The problem of image interpretation is a reasoning problem in which many sources of evidence need to be combined with rules and heuristics in order to generate the most likely explanation of the image. In section 5 methods for processing these uncertain data are reviewed. This also includes the propagation of uncertainty towards the final interpretation result. Section 6 then describes how these methods are being applied or could have been applied to the extraction of roads from aerial imagery. The last section summarizes the findings and outlines the work that remains to be done to assess the uncertainty in updated road databases.

2 ROAD MAP UPDATING

Updating (road) maps is usually considered a two step procedure: first it is verified whether the roads in the old database are still present in the imagery and, second, the new roads are mapped and roads that no longer exist are removed from the database. Looking closer at each of these two steps, many aspects can be discerned.

Verification involves the comparison of the representation of a road in a database (GIS) with the road appearance in the image. The common GIS is vector oriented and also may be generalised. In order to make the two object representations comparable, features are usually extracted from the image. The hypothesis that the features of the image and the GIS originate from the same road is subjected to a test. This test will require information about the uncertainty of the two object descriptions that may be obtained through error propagation in the feature extraction (and generalisation) process. Already this first updating step includes many aspects of uncertainty. A further analysis will be given in section 4.

Mapping new roads is even more complex than the verification step. If the hypothesis of the verification step is rejected, we only conclude that something has changed. The kind of change, e.g. a new road exit, a new fly-over, an extra lane, a new pavement, or a removed road part, is yet unknown. Using the features in the image, the context of the old GIS, knowledge related to all processing steps of both the image and GIS features, and generic knowledge about road networks, hypotheses have to be generated about the type of change. Some of these hypotheses (new exit, new fly-over) may, if accepted, lead to the detection of new parts of the road network. After this detection the road sides will need to be outlined and the consistency with the already known part of the road network needs to be verified [Gunst and Hartog, 1994]. As with the initial verification, all data used in the second step of the updating process can be erroneous and affect the quality of the final interpretation result.

Although very complex, updating of road databases may still be considered a little easier than updating maps with e.g. houses. It seems fair to assume that new roads are always connected to the roads in the old database. Therefore, junctions of the new roads with the old network should be detectable in the verification step. This gives a strong indication about where to look for new roads.

3 UNCERTAINTY IN DATA

When photogrammetrists talk about uncertainties they usually do so in terms of standard deviations of Gaussian distributed variables. Yet, when dealing with image interpretation tasks it soon becomes clear that many aspects of uncertainty cannot be described in those terms.

In this section we will first review the most popular ways of representing uncertainty in data. Many of the newer concepts have been developed in the AI literature and are related to reasoning problems.

In the last part of this section quality descriptions of (GIS) data are discussed. It shows that there is a variety of aspects of data quality that all affect the uncertainty about the correctness of the data.

3.1 Representations of uncertainty

- Probabilities

The best known representations of uncertainty are, of course, probabilities (and probability densities). In the Bayesian formalism, there are three basic axioms of probability theory regarding the belief measure that is attached to propositions [Pearl, 1988, Fine, 1973]:

$$0 \leq \mathsf{P}(A) \leq 1$$

$$\mathsf{P}(\text{Sure proposition}) = 1$$

$$\mathsf{P}(A \vee B) = \mathsf{P}(A) + \mathsf{P}(B) \quad \text{if } A \text{ and } B \text{ are mutually exclusive.}$$

From these axioms it also follows that:

$$\begin{aligned} \mathsf{P}(A) + \mathsf{P}(\neg A) &= \mathsf{P}(A \vee \neg A) \\ &= \mathsf{P}(\text{Sure proposition}) = 1 \end{aligned}$$

I.e., a proposition and its negation must be assigned a total belief of 1.

Bayes introduced the concept of conditionalisation of the belief in a proposition A by the knowledge or context B

$$\mathsf{P}(A|B) = \frac{\mathsf{P}(A, B)}{\mathsf{P}(B)}$$

These conditional probabilities play a very important role in all kinds of reasoning processes.

- Probabilistic networks

  Probabilistic networks [Pearl, 1988] also use the above defined probabilities and are therefore not a different way of expressing uncertainty. Instead they are useful graphical representations of the dependencies between propositions. The nodes of these graphs are the propositions and the links between the nodes show the dependencies. Two basic network types are often encountered:

  + Markov networks: A Markov network is an undirected graph. The links of this graph represent symmetrical probabilistic dependencies.
  + Bayesian networks: A Bayesian network is a directed graph. The arrows of this graph represent causal influences between the propositions. The Bayesian network may not contain cycles.

  Bayesian networks are very attractive for reasoning problems, since they directly show the lines along which the reasoning has to take place.

- Information theoretic measures

  The information $\mathsf{I}(A)$ of a proposition A, on the one hand is directly related to the above defined probability $\mathsf{P}(A)$ by

$$\mathsf{I}(A) = -\log \mathsf{P}(A)$$

and is interpreted as an amount (in bits) of surprise or uncertainty. Hence, I(Sure proposition) = 0. On the other hand, it is also motivated by research on communication theory. The amount of information of a proposition is the number of bits required to encode the proposition with an optimal coding scheme [Shannon and Weaver, 1949, Blahut, 1987].

Popular concepts from this theory are the mutual information

$$\mathsf{I}(A; B) = \log \frac{\mathsf{P}(A|B)}{\mathsf{P}(A)} = \log \frac{\mathsf{P}(B|A)}{\mathsf{P}(B)}$$

and the minimum description length (MDL) principle.

- Certainty factors

  Certainty factors were introduced in the famous MYCIN programme, a rule based expert system for medical applications [Buchanan and Shortliffe, 1984]. A certainty factor (CF) is based on a measure of increased belief (MB) and a measure of increased disbelief (MD) in some proposition A given the fact B.

$$\mathsf{MB}(A,B) = \begin{cases} 1 & \text{if } \mathsf{P}(A)=1 \\ \frac{\max[\mathsf{P}(A|B),\mathsf{P}(A)]-\mathsf{P}(A)}{1-\mathsf{P}(A)} & \text{otherwise} \end{cases}$$

$$\mathsf{MD}(A,B) = \begin{cases} 1 & \text{if } \mathsf{P}(A)=0 \\ \frac{\mathsf{P}(A)-\min[\mathsf{P}(A|B),\mathsf{P}(A)]}{\mathsf{P}(A)} & \text{otherwise} \end{cases}$$

$$\mathsf{CF}(A,B) = \mathsf{MB}(A,B) - \mathsf{MD}(A,B)$$

  Though still often used in expert systems, it is recognised [Weichselberger and Pöhlmann, 1990, p. 64] that the usage of certainty factors can be very misleading as it emphasizes the difference $\mathsf{P}(A|B) - \mathsf{P}(A)$, but almost disregards the prior $\mathsf{P}(A)$ itself.

- Dempster-Shafer theory

  One of the major criticisms on Bayesian probabilistic reasoning is the requirement for a complete probabilistic model. And indeed, the determination of prior and conditional probabilities for all relationships is often very difficult and cumbersome. The Dempster-Shafer theory allows a partially complete model by assigning probabilities to sets of propositions [Shafer, 1975]. Let $\Omega = \{X, Y, Z\}$ be an exhaustive set of mutually exclusive propositions (called the frame of discernment). A mass function m assigns probabilities to all subsets of Ω, such that the sum of these probabilities equals 1. For instance, one may know confidence that $\mathsf{m}(X) = 0.2$, but may have no further information. Then, $\mathsf{m}(X \vee Y \vee Z) = \mathsf{m}(\Omega) = 0.8$. Thus, for the remaining 80%, no preference is given to Y, Z or even X.

  The belief $\mathsf{Bel}(S)$ in some subset of Ω is defined as the sum of the probabilities assigned to subsets of S. I.e. in the example, $\mathsf{Bel}(X) = \mathsf{m}(X) = 0.2$, $\mathsf{Bel}(X \vee Y) = \mathsf{m}(X \vee Y) + \mathsf{m}(X) + \mathsf{m}(Y) = 0.2$, etc. Note that always $\mathsf{Bel}(\Omega) = 1$.

  Similarly, a plausibility $\mathsf{Pl}(S)$ is defined as 1 - $\mathsf{Bel}(\neg S)$. I.e. in the example, $\mathsf{Pl}(X) = 1 - \mathsf{Bel}(Y \vee Z) = 1 - \mathsf{m}(Y \vee Z) - \mathsf{m}(Y) - \mathsf{m}(Z) = 1$, $\mathsf{Pl}(Y \vee Z) = 1 - \mathsf{Bel}(X) = 1 - \mathsf{m}(X) = 0.8$, etc. Thus, the belief and plausibility functions define lower and upper bounds for the probabilities of the propositions. The difference $\mathsf{Pl}(S)$ - $\mathsf{Bel}(S)$ expresses our lack of knowledge in the probabilistic model and is often referred to as the ignorance. Hence, in contrast to Bayesian probability theory, $\mathsf{Bel}(S) + \mathsf{Bel}(\neg S)$ do not need to sum up to unity.

- Probabilistic logic

  Probabilistic logic (also called interval probability theory), like Dempster-Shafer theory is used to compute the bounds for the space of all probability assignments that are consistent with the available specifications [Nilsson, 1986, Weichselberger and Pöhlmann, 1990]. Thus, if $\Omega = \{X, Y, Z\}$ and $\mathsf{m}(X) = 0.2$, like in the above example, one can conclude that $\mathsf{P}(X) \geq 0.2$, and, since $\mathsf{P}(X) + \mathsf{P}(Y) + \mathsf{P}(Z) = 1$, it follows that $\mathsf{P}(Y) + \mathsf{P}(Z) \leq 0.8$. Major differences with the Dempster-Shafer theory arise when new sources of evidence are added to the available knowledge (discussed in section 5).

- Possibilities

  Possibility theory was developed based on the idea of fuzzy sets [Zadeh, 1978]. In set theory an object is either member of a set or not. In fuzzy sets an object can be a member to a certain degree, called the possibility. A typical example is whether a person belongs to the set of young persons. At the age of 1 this is definitely true (possibility 1). At the age of 100 the possibility will be 0. Somewhere in between there is a transition, but there is no specific age at which a person is suddenly no longer young. I.e. the dividing line between the class of young and "not young" persons is vague.

  There are many similarities between possibilities and probabilities. Cheeseman [1984], e.g. argues that the possibility in the above example can be considered as a conditional probability P(young|age).

3.2 Aspects of data quality

In the report of the Spatial Data Quality Commission of the ICA [Guptill and Morrison, 1995], the quality of GIS data is described in terms of lineage, positional and attribute accuracy, completeness, consistency, semantic accuracy and temporal accuracy.

The lineage is the history record of a database. It contains information about the time of observation (photo flight), the methods of data capturing and data processing, and, most important, the purpose of data collection.

Positional and attribute accuracy are well-known elements of data quality and can well be described by probability (density) distributions and derived quantities, like standard deviations.

Completeness is often understood as data completeness and expressed by probabilities of omission and commission errors. Another aspect, raised in [Guptill and Morrison, 1995], is model completeness which should indicate whether the data model used at the time of data acquisition contains the features and attributes sufficient to solve the tasks of a particular application (that may or may not have been foreseen at the time of data acquisition).

Assuming that the data model is consistent, data inconsistency is caused by errors in redundant data. Therefore, there is a close link to attribute accuracy. Since, consistency always involves two or more features or attributes, conditional probabilities are in particularly suitable.

Semantic accuracy has to do with the meaning of objects and attributes. E.g., does an object or attribute contain the information we think they do? Semantic accuracy is very difficult to characterize. Especially for image interpretation tasks, it is, however, of major importance, since the interpretation process is supposed to attach a meaning to the objects in the image.

Finally, accuracy of temporal information is often related to consistency of a database during the actions of an update or the accuracy of a time attribute. These aspects were already covered. A new aspect may be information about the rate of change that is to be expected. This can be used to roughly estimate the completeness of a database at some point in time which may be valuable information for the verification step.

4 UNCERTAINTY IN DATA FOR ROAD EXTRACTION

On a first sight one might think that we will have to deal with the precision of the GIS data and the precision of the extracted image features. Yet, from the above aspects of data quality it should have become clear that there are many more causes for uncertainty.

4.1 GIS data

In fact, positional accuracy of GIS data is only a minor source of uncertainty. The available data usually permits to outline the sides of the road surface or the middle of the road within a few pixels in the image. The (very few) results on updating road networks by image interpretation are, however, much worse. Thus, there must be other sources of uncertainty.

Like positional accuracy, attribute accuracy is usually very high in comparison to the quality of the interpretation results.

Consistency of the GIS data may be considered a more important factor. Inconsistent data will yield conflicting evidence to some hypotheses and thereby can mislead the reasoning process. Of special interest in GIS data are the topological relations between the features. Egenhofer and Franzosa [1991] classified eight different topological relationships between two two-dimensional regions (like meet, overlap, disjoint, etc.). Winter [1994] argues that such relationships between regions can not be considered as certain, due to positional inaccuracies, however small they are. E.g., even due to the smallest possible error, two regions that actually meet may be classified as disjoint or overlapping. Other changes in topology are less likely. E.g., if one region is actually contained in another, it is unlikely that they will be classified as disjoint. Winter [1994] therefore derives conditional probabilities of topological relations between regions in a GIS, given their true topology. These probabilities very well model the uncertainty in topological relations. A reasoning process can now take into account the confidence that has to be given to some relationship and does not have to accept all relations as correct.

Semantic errors in the GIS can also have a large impact on the image interpretation. Suppose that, according to the data model, a road database contains the roadsides. This definition of the data still allows several interpretations. E.g., does the road include the sidewalk, or the shoulder? An incorrect interpretation of the data model can clearly lead to a large number of errors in the verification step. If the data model is ambiguous, the verification step should comprise hypotheses for each of the different interpretations in order to find the correct one.

Since the purpose of the image interpretation is in updating the road database, it is obvious that the data completeness of the GIS is significantly lower than what can be expected for an up to date GIS. If available, a rate of change may be used to calculate the expected data completeness at the time of updating. This number can then be compared to the results of the verification step.

The information in a GIS is clearly insufficient to automatically solve the interpretation of the aerial images. In this sense, the model completeness of the GIS for the task of automatic updating is very low. Image interpretation requires much richer descriptions of the objects than a few vectors in a GIS. This leads to the problem of object modeling.

4.2 Object models

Describing roads in generic models such that these models contain sufficient information to recognize all kinds of roads is an extremely difficult task. Yet, humans have no problems in recognizing the roads in figure 1 despite the large variety in shape, size, scale, and pavement.

Figure 1: Variety of road appearances in aerial imagery.

Gunst [1996], after [Garnesson et al., 1990] describes a road model in terms of geometry, radiometry, topology, functionality and context. Many attempts to describe a road only use geometrical and radiometrical properties. E.g., a road is defined as two parallel edges that include an elongated homogeneous area. According to this definition, a side walk, a single traffic lane, a river, a dike, a beach, and probably many other objects also can be classified as a road. Some improvements can be made by including colour or texture information, but a good result can only be expected if the context of a road is also considered.

Cars on a road, houses and trees alongside the road, shadows of fly-overs, junctions with other roads, road markings, and traffic signs are all clues that help an operator to identify a road. The description of all these objects may in turn require some other context information. The question then rises how extended the context of a road should be and how detailed each of the objects needs to be described. The answer is not known, but it is clear that a lack of modeled knowledge about the objects and their context is a major source for the uncertainty in the outcome of image interpretation procedures. Instead of finding support in the presence of cars, houses, road markings, etc., most road detection schemes consider these objects as noise which leads to detection failures.

4.3 Image data

The image data itself and the feature extraction process are also sources of uncertainty. In the imaging process uncertainties are introduced by the sensor noise and the imaging circumstances. Due to a different perspective or changed (weather) conditions object appearances may change drastically and thereby systematically affect the number and shape of the extracted features.

When propagating the image noise to the parameters of extracted features, the assumed noise level is usually taken much higher than the sensor noise (which is almost neglectable). This higher noise level is required to account for small violations of the image models used in the feature extraction algorithms. E.g. many edge extraction operators assume ideal straight step edges with constant grey values on both sides of the edge. When extracting the side of a road, small grey level variations due to structures in the concrete or clumps of grass are ignored and (incorrectly) considered as noise. Such incomplete or simplifying image models give rise to a substantial amount of uncertainty in the extracted features.

Due to the complexity of feature extraction a straightforward error propagation is often not possible. In those cases extensive experiments are required on either simulated [Fuchs et al., 1994] or real [Vosselman, 1992] imagery in order to capture the stochastic properties of the feature extraction process. Transition matrices with conditional probabilities have proven to be adequate for describing the uncertainty. Once extraction probabilities of some basic features are known, some probabilities of detecting more complex features can be derived theoretically. E.g., Fuchs et al. [1994] determine the probability of detecting line junctions by propagating the probability of detecting edge pixels.

In the previous paragraph it was argued that many road models are to poor for a successful recognition. This recognition is based on a comparision between object models and image features. Like for the objects, it is largely unknown how to describe an image such that the description is suitable for interpretation purposes. Many feature extraction processes do not preserve the information that would be very helpful for interpretation and thus complicate the high level reasoning.

5 PROCESSING UNCERTAIN DATA

Image interpretation tasks have to combine several knowledge sources. To assess the final quality the uncertainty in the knowledge sources needs to be propagated. Related to the different methods of representing uncertainty (section 3),

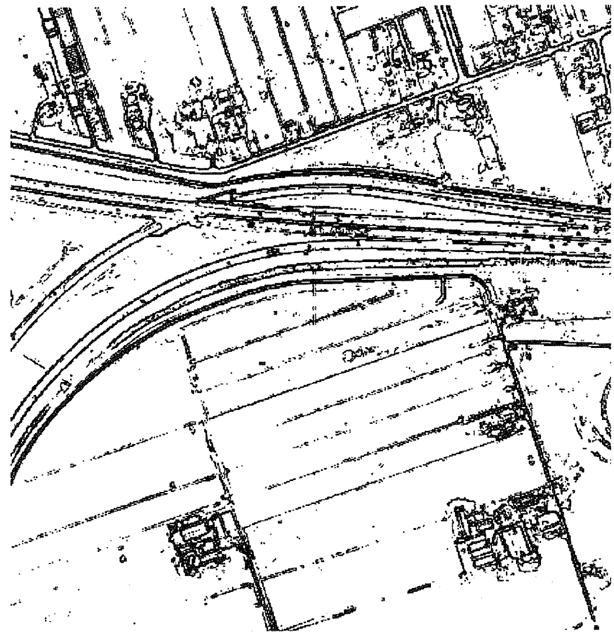

Figure 2: Edges do not contain sufficient information to distinct roads from other linear features.

several techniques for combining uncertain knowledge and propagating uncertainty have been developed.

- Probabilities

  Most computations with probabilities are in some way related to Bayes' theorem

  $$\mathrm{P}(A|B) = \frac{\mathrm{P}(B|A)\mathrm{P}(A)}{\mathrm{P}(B)}$$

  in which the probability of the event A, given that B has been observed is derived. Beside prior probabilities $\mathrm{P}(A)$ and $\mathrm{P}(B)$, also the conditional probability of observing B in case of the event A has to be known. This conditional probability corresponds to the stochastic model used in adjustments, i.e. the assumption of a Gaussian distribution with a certain standard deviation. Error propagation with Bayes' theorem or least squares adjustments of linearized models are very common in photogrammetric calculations, but still find little attention when dealing with GIS data. Heuvelink et al. [1989] and Goodchild and Gopal [1989] give a few examples of error propagation in GIS.

- Probabilistic networks

  Associated with the links of a probabilistic network are conditional probabilities. The probability of each proposition (node) may depend on the probability of several neighbouring nodes. So-called relaxation methods update the probability of a proposition by using the probabilities at the adjacent nodes together with the conditional probabilities [Rosenfeld et al., 1976]. In its simplest form, the probability of proposition A is derived from neighbouring propositions $B_1 \dots B_n$ by

  $$\mathrm{P}(A) = \sum_{i=1}^{n} [\mathrm{P}(A|B_i)\mathrm{P}(B_i) + \mathrm{P}(A|\neg B_i)\mathrm{P}(\neg B_i)]/n$$

In Markov networks this approach may lead to problems, since the probability at a node A that has been derived with the above formula, is used in a later stage to recompute the probability at one (or more) of its neighbouring nodes. Pearl [1988, p. 149] gives a nice example of this kind of circular reasoning:

"Imagine that a processor F, representing the event Fire, communicates asynchronously with a second processor S, representing the event Smoke. At time t1, some evidence (e.g. the distant sound of a fire engine) gives a slight confirmation to F, thus causing the probability of Fire to increase from P(F) to P1(F). At a later time, t2, processor S may decide to interrogate F; upon finding P1(F), it revises the probability of Smoke from P(S) to P2(S) in natural anticipation of smoke. Still later, at t3, processor F is activated, and upon finding an increased belief P2(S) in Smoke, it increases P1(F) to an even higher value, P3(F). This feedback process may continue indefinitely, the two processors drawing steady mutual reinforcement void of any empirical basis, until eventually the two propositions, Fire and Smoke, appear to be firmly believed."

This kind problem can be solved by keeping track of the source of evidence. However, this involves a more complex algorithm, such that the advantages of local asynchronous probability updates are lost.

- Certainty factors

  Certainty factors CF1$(A, B1)$ and CF2$(A, B2)$ arising from two observations $B1$ and $B2$ are used to derive a combined certainty factor with [Buchanan and Shortliffe, 1984]

$$\mathrm{CF} = \begin{cases} \mathrm{CF1} + \mathrm{CF2} - \mathrm{CF1} \cdot \mathrm{CF2} & \text{if } \mathrm{CF1}, \mathrm{CF2} > 0 \\ \mathrm{CF1} + \mathrm{CF2} + \mathrm{CF1} \cdot \mathrm{CF2} & \text{if } \mathrm{CF1}, \mathrm{CF2} < 0 \\ \frac{\mathrm{CF1} + \mathrm{CF2}}{1 - \min(|\mathrm{CF1}|, |\mathrm{CF2}|)} & \text{otherwise} \end{cases}$$

  Whereas single certainty factors can already be misleading, combined certainty factors are even more dangerous, since any correlation between observations is neglected.

- Dempster-Shafer theory

  Given two sources of evidence, the mass functions m1 and m2 such that the combined probability of a subset S, m1 + m2(S), is the sum of the joint probabilities of all combinations of two subsets (T_i, U_j) which intersection equals S. This sum is normalised by the sum of the joint probabilities of all combinations of two subsets which intersection is not an empty set. This normalisation is required in order to take out the so-called weight of conflicting evidence.

$$\mathrm{m1{+}m2}(S) = \frac{\sum_{\{i,j|T_i \cap U_j = S\}} \mathrm{m1}(T_i)\mathrm{m2}(U_j)}{\sum_{\{i,j|T_i \cap U_j \neq \emptyset\}} \mathrm{m1}(T_i)\mathrm{m2}(U_j)}$$

  This update formula shows many resemblances to combining evidence from two independent sources with Bayesian probability theory. The Dempster-Shafer update formula is, however, controversial. Especially in case of incomplete probabilistic models, i.e. Bel(S) + Bel$(\neg S) < 1$, it may lead to curious results (see e.g. [Pearl, 1988, p. 447]).

- Probabilistic logic

  As more evidence becomes available, the theory of probabilistic logic will use this information to further constrain the space of all possible probability assignments until the probabilistic model. In this way results remain consistent with Bayesian probabilistic methods. Pearl [1988] therefore concludes that in case of analysis problems with incomplete probability models probabilistic logic should be preferred above Dempster-Shafer theory.

- Possibilities

  Possibilities of set membership are typically updated with

$$\begin{aligned} \mathrm{poss}(A \wedge B) &= \min(\mathrm{poss}(A), \mathrm{poss}(B)) \\ \mathrm{poss}(A \vee B) &= \max(\mathrm{poss}(A), \mathrm{poss}(B)) \end{aligned}$$

  These update rules are only equivalent to probabilistic rules when A and B are completely dependent, i.e. $A \rightarrow B$ or $B \rightarrow A$. But if, e.g, A and B are mutually exclusive, it is clear that P$(A \wedge B)$ should be zero [Cheeseman, 1984].

6 UNCERTAINTY IN EXTRACTING ROADS

Surprisingly, only a very few publications deal with automatic updating of road maps. The usage of an old road database as a valuable source of knowledge still is very uncommon. Many more papers have been published on road extraction to build up a database from scratch. Most of these publications, however, pay very little or no attention to the uncertainty in the extracted roads. It seems that, like in many areas of image understanding, the results are too poor to seriously consider to describe their quality.

In this section we will again make a distinction between the verification and the detection step in the updating process. For both steps several presented results will be shown and it will be discussed how the uncertainty in these steps was dealt with or could have been dealt with.

6.1 Verification

Four examples are discussed that compare the contents of an aerial image with roads in a database. The first two are aimed at verification. The goal of the last two papers was the location of a road junction. However, the same strategy might have been used for verification as well.

Gunst and Hartog [1994] and Gunst [1996] discuss the advantages of a knowledge based interpretation strategy for updating road maps. The existence of an old road in the new image is verified by submitting the cross correlation between grey value profiles of road cross sections and an artificial road profile to a statistical test. If the cross correlation is lower than a threshold, a change is hypothesized. Problems arise with (larger) cars and overhanging trees alongside the road. Since the road model does not contain any knowledge about possibly occluding objects, many false alarms result. Hence, the uncertainty about the correctness of the verification results are mainly due to insufficient modeling of the road's context.

Baumgartner et al. [1996] compare extracted linear features to the road sides in a vector-based GIS. Checks are performed on parallelism, straightness and symmetry. With some effort in error analysis of the feature extraction process, conditional

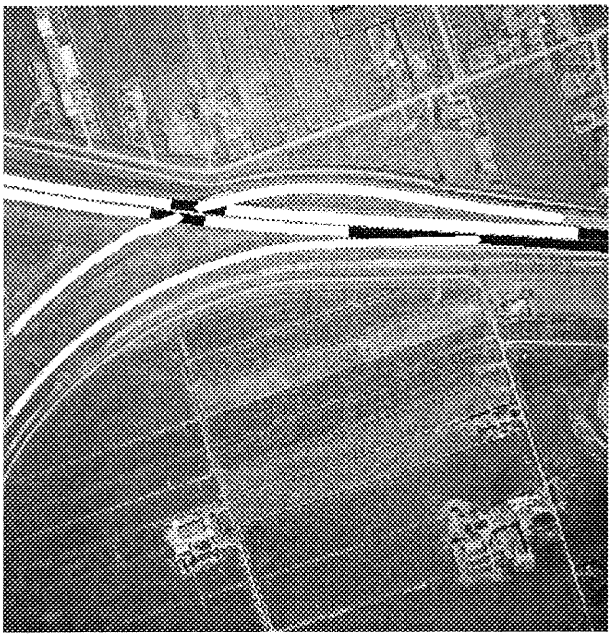

Figure 3: Verification and detection of junctions [Gunst, 1996]. In black are detected changes. Starting from the middle of the changed area, new roads are detected and followed.

probabilities may be derived that can give these tests a statistical basis. The question, however, remains whether a test on linear features can be sufficient to confirm that the road has not been changed. In other words, do extracted edges that match in position with the roadsides in the GIS indeed represent the sides of a road, or can they be due to some other linear structure? This question is important for automatic updating, since a false verification will not result in a search for changes in the road network and therefore cause omission errors in the updated database.

Nevatia and Price [1982] use a relaxation labeling technique to match a structural description of a map to a structural description extracted from an aerial image. The nodes in such a graph-like description do not represent propositions that are either true or false, but the different assignments of map features that can be made to a certain feature of the image. In [Price, 1985] different relaxation schemes for updating the likelihoods of assignments are compared. Although some schemes are called probabilistic, the final likelihoods of the assignments cannot be considered as probabilities. Since much of the evidence found in the labelings at the neighbouring nodes is used several times in the updating process, many labeling "probabilities" converge to either zero or one. The mapping between the descriptions that follow from the likely assignments is, however, very useful for hypothesizing the correct correspondences.

Haala and Vosselman [1992] use relational matching to match the structural descriptions of a map and an image. Their evaluation measure is based on the mutual information between the two descriptions and is derived from conditional probabilities that were obtained in many feature extraction experiments. Using the same probabilities the distribution of the mutual information was also derived. Thus, it is possible to define a statistical test on the amount of mutual information of the best mapping found. Since the structural descriptions contain many attributes of features and their relations, the probability of accepting a wrong match is fairly low. Although the acceptance test was well defined, the number of performed matching experiments was too low to draw conclusions about its applicability.

Summarizing, one may conclude that in many verification tests it is possible to base these tests on a statistical analysis. However, verification errors are likely to happen in case of unmodeled occluding objects or poor object descriptions.

6.2 Detection and measurement

Gunst and Hartog [1994] and Gunst [1996] consider the case of detecting and mapping new exit roads and fly-overs. After the verification steps several locations with a possible change are marked. In these areas goal-directed segmentation algorithms try to detect parts of other roads of the junction. Once these have been found, the new road parts are classified as either an exit road or a fly-over. The decision is based on the values of a few attributes (e.g. the angle between road elements) and knowledge of the road design rules (e.g. exits only have a small angle with the main road). Probability distributions of the attribute values are not used in this test. They could, however, have been used to show the uncertainty in the classification. In case of a high uncertainty it would then be useful to consider multiple hypotheses about the kind of junction. In the current implementation no alternative classifications are considered.

Cleynenbreugel et al. [1990] suggest to use many more layers of a GIS for the detection of new roads. Except for old roads, other information like land cover, DEM's, and hydrological information can also be helpful. Since land cover will discriminate between urban and rural terrain, expectations for the shape of road networks can be tuned to these classes. DEM's can be used to derive slope maps that constrain the possible directions of roads. It may even be possible to derive probability distributions of the road direction at some point given the terrain slope at that point. Finally, hydrological information (position of rivers, lakes, etc.) is also useful. Roads in mountainous areas are often parallel to rivers and have as few bridges (=construction costs) as possible. Roads in the middle of lakes are very unlikely. Many of these heuristics are valuable for image interpretation and will indeed be used by human operators.

McKeown and Denlinger [1988] use profile matching for tracking roads. Road trackers are usually initialized at positions indicated by an operator. In case of updating road networks, it can, however, also be done automatically at those positions where new junctions have been found. Vosselman and de Knecht [1994] use the least squares method for profile matching and Kalman filtering to estimate the position, direction and curvature of the road. This approach enables them to also estimate the precision of the road parameters and to detect failures in the profile matching. Thus the uncertainty in the road extraction is fairly well described. Road trackers, in general, can however only deal with simple roads and will fail at e.g. Y-junctions.

Other methods to outline roads are often based on snakes, deformable templates or dynamic programming. Grün and Agouris [1994] combine the advantages of snakes and least squares template matching by constraining the matching results. Precision estimates are also obtained.

Summarizing, relatively simple tasks like outlining single, and completely visible roads can be performed successfully. If the algorithm can be formulated as a least squares estimation problem, quality descriptions can also be obtained. The more complicated tasks that relate to interpretation of junctions and the road context are far from being solved. It is also clear, however, that much of the available knowledge is not adequately modelled and, therefore, not available to the interpretation process.

7 DISCUSSION

The main issues raised in this paper were the different ways to describe and process uncertain data related to updating road databases. It was also shown that there are many reasons for the poor results in image interpretation.

Many ways to describe and process uncertain data are based on probabilities and Bayes' rule. Several alternatives, like certainty factors, Dempster-Shafer theory, and possibility theory, have been applied successfully in some domains of AI, but may lead to wrong results in other applications. The need for conditional or prior probabilities is often mentioned as a disadvantage of Bayesian probabilistic reasoning, but, in general, the alternative strategies can not fill the gap of such a lack of (modeled) knowledge. Gathering this information remains important. Many claimed advantages of alternatives for probabilities, e.g. possibility theory, can also be realized with a probabilistic approach [Cheeseman, 1984].

Probabilities describe uncertainty, but probability numbers themselves are also uncertain. Especially intuitive heuristics will often only give us a rough idea about some probability number. The endorsement theory [Cohen, 1985] studies how to represent and reason with heuristic knowledge about uncertainty. An interesting analogy can be found between the concept of external reliability in least squares adjustments (the influence of an error in an observation onto the estimated unknowns) and the question how changes in probability distributions affect the outcome of a reasoning process.

Most algorithms for automatic road detection have very low success rates. The context of an old road database contains very useful knowledge to improve this. However, this knowledge is far from sufficient to solve image interpretation tasks. Much additional knowledge concerning the appearances of roads in aerial images and the context of roads will need to be modelled.

Uncertainty plays an important role in this knowledge. Much of our knowledge is heuristic and therefore uncertain. This uncertainty needs to be described in order to properly reason with knowledge. In many cases conditional probabilities will be appropriate, e.g. P(road direction|terrain slope).

Propagation of errors is also necessary. The purpose of error propagation is not only to assess the quality of the final result, but also to value the correctness of intermediate results. The latter motive may be even more important. A sound interpretation can only be made when the quality of all data in all processing steps is known. If the results of some step are found to be uncertain, this knowledge can be used to formulate multiple alternative hypotheses instead of only pursuing the most likely one.

Modelling knowledge and propagating uncertainty are two complicated tasks. A lot of effort will be required to obtain satisfactory results in automatic image interpretation. It therefore seems a good approach to first start with semi-automated methods and gradually increase the interpretation tasks of vision algorithms.

LITERATURE

Baumgartner, A., C. Steger, C. Wiedemann, H. Mayer, W. Eckstein, and H. Ebner [1996]: Road Extraction for Update of GIS from Aerial Imagery: A Two-Phase Two-Resolution Approach. International Archives of Photogrammetry and Remote Sensing, vol. 31, part B3.

Blahut, R.E. [1987]: Principles and Practice of Information Theory. Addison-Wesley, Reading, MA, U.S.A..

Buchanan, B.G. and E.H. Shortliffe (Eds.) [1984]: Rule Based Expert Systems. The MYCIN Experiments of the Stanford Heuristic Programming Project, Addison-Wesley, Reading, MA, U.S.A..

Cheeseman, P. [1986]: Probabilistic vsFuzzy Reasoning. In: Uncertainty in Artificial Intelligence, vol. 4, L.N. Kanal and J.F. Lemmer (Eds.), North-Holland, Amsterdam, pp. 85-102.

Cleynenbreugel, J. van, F. Fierens, P. Suetens and A. Oosterlinck [1990]: Delineating Road Structures on Satellite Images by a GIS-Guided Technique. Photogrammetric Engineering and Remote Sensing, vol. 56, pp. 893-898.

Cohen, P.R. [1985]: Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach. Research Notes in Artificial Intelligence, Pitman Publishing, London.

Egenhofer, M.J. and R.D. Franzosa [1991]: Point-set Topological Spatial Relations. International Journal of Geographical Information Systems, vol. 5, pp. 161-174.

Fine, T. [1973]: Theories of Probability. Academic Press, New York.

Fuchs, C., F. Lang, and W. Förstner [1994]: On the Noise and Scale Behaviour of Relational Descriptions. International Archives of Photogrammetry and Remote Sensing, vol. 30, part 3/1, pp. 257-264.

Garnesson, Ph., G. Giraudon, and Montesinos [1990]: An Image Analysis System, Application for Aerial Image Interpretation. Proceedings 10th IAPR International Conference on Pattern Recognition, Atlantic City, vol. 1, pp. 210-212.

Goedhart, B. [1994]: A General Framework for Compositional Reasoning with Uncertainty. PhD-thesis, Delft University of Technology.

Goodchild, M. and S. Gopal (Eds.) [1989]: The Accuracy of Spatial Databases. Taylor & Francis, London.

Grün A., and P. Agouris [1994]: Linear Feature Extraction by Least Squares Template Matching Constrained by Internal Shape Forces. International Archives of Photogrammetry and Remote Sensing, vol. 30, part 3/1, pp. 316-323.

Grün A. and H. Li [1994]: Semi-automatic Road Extraction by Dynamic Programming. International Archives of Photogrammetry and Remote Sensing, vol. 30, part 3/1, pp. 324-332.

Gunst, M.E. de, and J.E. den Hartog [1994]: Knowledge-Based Updating of Maps by Interpretation of Aerial Images. Proceedings 12th International Conference on Pattern Recognition.

Gunst, M.E. de [1996]: Knowledge Based Updating of Road

Maps from Aerial Images. PhD-thesis, Delft University of Technology, to appear.

Guptill, S.C. and J.L. Morrison (Eds.) [1995]: Elements of Spatial Data Quality. Elsevier Science, Kidlington, U.K.

Haala, N. and G. Vosselman [1992]: Recognition of road and river patterns by relational matching. International Archives of Photogrammetry and Remote Sensing, vol. 29, part B3, pp. 969-975.

Heuvelink, G.B.M., P.A. Burrough and A. Stein [1989]: Propagation of errors in spatial modeling with GIS. International Journal of Geographical Information Systems, vol. 3, pp.303-322.

Heuvelink, G.B.M. [1993]: Error propagation in quantative spatial modeling. PhD-thesis, University of Utrecht.

Hummel, R.A. and S. Zucker [1983]: On the Foundations of Relaxation Labeling Processes. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 5, pp. 267-287.

Kanal, L.N. et al. [1986-]: Series on "Uncertainty in Artificial Intelligence". North-Holland, Amsterdam.

Maître, H., I. Bloch, H. Moissinac, and C. Gouinaud [of aerial images and maps for the interpretation of urban scenes. Proceedings Ascona Workshop on Automatic Extraction of Man-Made Objects from Aerial and Space Images, Birkhäuser Verlag, pp. 297-306.

McKeown D.M. and J. Denlinger [1988]: Cooperative Methods for Road Tracing in Aerial Imagery. IEEE Proceedings on Computer Vision and Pattern Recognition, Ann Arbor, MI, pp. 662-672.

Nevatia, R. and K.E. Price [1982]: Locating Structures in Aerial Images. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 4, no. 5, pp. 476-484.

Nilsson, N.J. [1986]: Probabilistic Logic. Artificial Intelligence, vol. 28, no. 1, pp. 71-87.

Pearl, J. [1988]: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers, San Mateo, CA, U.S.A..

Price, K.E. [1985]: Relaxation Matching Technique - A Comparison. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 7, no. 5, pp. 617-623.

Rosenfeld, A., R.A. Hummel and S.W. Zucker [1976]: Scene Labeling by Relaxation Operations. IEEE Transactions on Systems, Man, and Cybernetics, vol. 6, no. 6, pp. 420-433.

Shafer, G. [1976]: A Mathematical Theory of Evidence. Princeton University Press, Princeton, NJ, U.S.A..

Shannon, C.E. and W. Weaver [1949]: The Mathematical Theory of Communication. University of Illinois Press.

Vosselman, G. [1992]: Relational Matching. Lecture Notes in Computer Science, vol. 628, Springer Verlag.

Vosselman, G. and J. de Knecht [1995]: Road tracing by profile matching and Kalman filtering. Proceedings Ascona Workshop on Automatic Extraction of Man-Made Objects from Aerial and Space Images, Birkhäuser Verlag, pp. 255-264.

Waltz, D. [1975]: Generating semantic descriptions from drawings of scenes with shadows. In Winston, P.H. (Ed.) The Psychology of Computer Vision, McGraw-Hill, New York, pp. 19-92.

Weichselberger, K. and S. Phlmann [1990]: A Methodology for Uncertainty in Knowledge-Based Systems. Lecture Notes in Artificial Intelligence, vol. 419, Springer Verlag, Berlin.

Winter, S. [1994]: Uncertainty of topological relations in GIS. International Archives of Photogrammetry and Remote Sensing, vol. 30, part 3/2, pp. 924- 930.

Zadeh, L.A. [1978]: Fuzzy Sets as a Basis for a Theory of Possibility. In: Fuzzy Sets and Systems, vol. 1.

Zadeh, L.A. [1984]: Is Probability Theory Sufficient for Dealing with Uncertainty in AI: A Negative View. In: Uncertainty in Artificial Intelligence, vol. 4, L.N. Kanal and J.F. Lemmer (Eds.), North-Holland, Amsterdam, pp. 103-116.