Working Group: Explanation

Hidir Aras[1], Norbert Fuhr[2], Seung-won Hwang[3], Ander de Keijzer[4], Friederike Klan[5], Hans-J. Lenz[6], Tom Matthé[7], Heinz Schweppe[6], Mirco Stern[8], Guy De Tré[7]

[1] University of Bremen, Germany
[2] University of Duisburg-Essen, Germany
[3] POSTECH, Korea
[4] University of Twente, Netherlands
[5] University of Jena, Germany
[6] Free University of Berlin, Germany
[/] Ghent University, Belgium
[8] University of Karlsruhe, Germany

This working group addressed the issue of explaining the results of an uncertainty information system to a user. For that, we structured the problem along three major queries: why, what, and how.

Why?

First, there is the question why an information system should provide any explanation at all – e. g. in contrast to a standard information system, which usually does not comprise an explanation component. There is obviously an urgent need for transparency from the user's side: Since the system cannot provide answers with certainty, users might want to understand the system's inference process (which might affect their interpretation of the system's results). Furthermore, there might be legal requirements for transparency (e. g. explain the system's reasoning to a customer who was rejected a loan application).

Second, one might ask why a system should provide any uncertain answer at all – instead of restricting to the certain answers. In some applications, there may not be any certain answers for the specific query, or the domain is uncertain per se (e. g. information retrieval). Even if there are some certain answers, they may not be sufficient, and the user also would like to see the uncertain ones.

What?

Once the system has to provide an explanation, there is the question which aspects of the derivation of the answer should be explained. Most important, the possible sources of uncertainty should be made transparent to the user: There may be uncertain data stored as such in the database, and/or uncertainty may be due to the model that derived the answer (e. g. in classification tasks). In addition to the

sources of uncertainty, also the derivation graph of the result should be presented. This issue has been addressed by the working group dealing with lineage.

Explanations should be provided at different levels of granularity – these may be domain-dependent. Also, explanations should be user-dependent, subject to the user's expertise as well as taking privacy issues into account.

How?

Finally, our group addressed the issue of how to present uncertainty.

Here we distinguished several methods:

1. Probability estimates can be presented as point values, as a pair of mean and variance, or most detailed in the form of a histogram.

2. In some applications, users may only be interested in answers that exceed a certain threshold probability (e. g. in spam filters)

3. There are also applications where numeric values may be of little benefit for a user, since the application bears some intrinsic fuzzyness, like e. g. in similarity search and fuzzy joins.

4. Giving only a ranking of answers, without any quantification, is also a method used frequently (e. g. in information retrieval applications).

5. Finally, listing the possible worlds may be especially appropriate when there are no answers at all (due to disjointness of the underlying probabilistic events).

Given these five possibilities, we saw the following major application areas:

- In business applications, probabilities are most desirable, since they can be used for computing the expected gain or loss.

- Thresholds are appropriate when users are interested in a binary classification only.

- Ranking allows users to inspect answers in the order of decreasing likelihood.

- Possible worlds and lineage should be provided for users asking why a specific tuple was either retrieved or not retrieved.

Summary

Explanation is an important function of uncertainty information systems. The type of explanation given is heavily domain-dependent, and may also be user-spe-

cific. As an alternative to presenting probability value for each result tuple, other methods of indicating uncertainty may be more appropriate in many situations.

Literature

Joseph Y. Halpern, Judea Pearl: Causes and Explanations: A Structural-Model Approach: Part 1: Causes. UAI 2001: 194-202.

Joseph Y. Halpern, Judea Pearl: Causes and Explanations: A Structural-Model Approach - Part II: Explanations. IJCAI 2001: 27-34