EnviroInfo 2013: Environmental Informatics and Renewable Energies
Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

Involving the Expert in the Delivery of Environmental Information from the Web

Leo Wanner[1], Harald Bosch[2], Stefanos Vrochidis[3], Nadjet Bouayad-Agha[4], Gerard Casamayor[4], Lasse Johansson[5], Ari Karppinen[5], Anastasia Moumtzidou[3], Ioannis Kompatsiaris[3], Thomas Ertl[2]

Abstract

Automatic provision of accurate user need- and profile-tailored environmental information is of increasing demand. However, it is a challenging task with many facets. Among them is the difficulty to compile and cast into a formal representation all the expert knowledge needed to accurately and exhaustively acquire, assess and process the data in order to be able to reliably judge their relevance to the user and to produce an adequate summary and recommendations. Studies in Human-Computer Interaction show that both the satisfaction of the users with an application and the objective performance of a service increases if the users (in particular, experts) are assigned an active role in the system. Based on this insight, we propose a largely interactive environmental information acquisition and generation framework, which has been realized in the FP7 project "Personalized Environmental Service Configuration and Delivery Orchestration" (PESCaDO). The PESCaDO service involves the experts in four central tasks: (i) determination of criteria for the search of environmental nodes in the web; (ii) assessment of the relevance of the identified nodes; (iii) assessment of the quality of the data provided by the nodes, and (iv) selection of the content to be communicated to the user. Quantitative evaluations and user trials show that the performance of the system is good and the satisfaction with the service is high.

1. Introduction

Automatic provision of accurate user need- and profile-tailored environmental information is a challenging task with many facets. It involves the interpretation of the query of the user and of her profile in the light of the given environmental conditions, determination of the content that is relevant to the user in the given context and, finally, presentation of this content in the desirable format (e.g., text in the language of the preference of the user). The task becomes even more complex if the environmental information is first to be searched for and distilled from the web. In general, services that address this task can be interpreted as "expert systems" (Jackson, 1998) in that they attempt to model the expert knowledge and context-dependent decisions that have been made explicit before, for instance, in interviews, in order to use them both in a fully automatic mode. However, the experience teaches us that it is difficult to (i) obtain from experts all the necessary knowledge – be it because they are reluctant to reveal this knowledge or because some (less common) contexts are omitted; (ii) represent the knowledge and the decision making process such that for each query of the user the most adequate and relevant information is provided. The consequence is that often neither the expert nor the end user judge the performance of the service sufficiently good to use it on a regular basis.

[1] Catalan Institute for Research and Advanced Studies (ICREA) and DTIC, Universitat Pompeu Fabra
email: leo.wanner@upf.edu
[2] VIS, Universität Stuttgart
[3] ITI, Centre for Research and Technology Hellas
[4] DTIC, Universitat Pompeu Fabra
[5] Finnish Meteorological Institute

Studies in Human-Computer Interaction (HCI) and, more specifically, in its younger branch Visual Analytics (VA) show that the satisfaction of the users with an application, as well as its objective performance (i.e. system-oriented evaluation) can be considerably improved by assigning the user an active role in the system (see, among others, Card & Pirolli, 2005; Keim et al., 2010). Based on this insight, we propose a largely interactive environmental information generation framework, which is realized by the PESCaDO[6] service. The goal of the PESCaDO service is to provide end user need- and profile-tailored meteorological and air quality information from the web in the language of the preference of the user (Wanner et al., 2011). The end users interact with the service via a graphical interface. The interface allows them to register, personalize their profile, submit information requests and receive the produced information.

In order to be able to deliver appropriate information to the user, PESCaDO performs the following tasks: 1. search of environmental nodes in the web that provide potentially relevant environmental information in the corresponding region; 2. filtering the encountered nodes with respect to their relevance; 3. extraction of the data from the textual and image material of the websites; 4. assessment of the quality of the data provided by the websites of the relevant nodes and, when necessary, data fusion; 5. incorporation of the data into ontologies and their semantic interpretation; 6. selection of the content appropriate for the communication to the user; 7. automatic generation of a bulletin with the selected content in the language of preference of the user. The intervention of the expert seems particularly crucial with respect to the tasks 1, 2, 4, and 6. In what follows, we first discuss how this intervention is realized in PESCaDO and then assess its role in the current and future PESCaDO services.

2. Intervention of the expert in the provision of environmental information

In general, the intervention of the expert can be evaluative, corrective, or active. In the case of evaluative intervention, the expert provides feedback on the correctness (satisfaction, etc.) of the intermediate output; in the case of corrective intervention, she corrects the intermediate output; and in the case of active intervention, she makes choices that guide (or at least, influence) the further processing of the system. What kind of intervention is chosen for a given task depends on (i) the task itself; (ii) the realization of the task; (iii) the assumptions made with respect to the profile of the intervening expert. Thus, if the task consists in the interactive validation of measured time series, the intervention by default will be evaluative. However, a corrective dimension can be added to it in that an expert with a solid experience in time series from the domain in question can be further asked to correct the outliers. If the task consists in training an algorithm for detecting striking or irregular data, the task is *per se* also evaluative, but can be further extended by an active dimension if an interactive learning strategy is chosen. And so on.

In what follows, we sketch how the expert intervention is realized in PESCaDO for the four tasks mentioned in Section 1.

2.1 Determination of the criteria for the search of environmental websites in the web

The search for environmental nodes in the web is considered in PESCaDO to be a domain-specific search problem. Towards the solution of this problem, we propose a novel search framework which builds on top of a general purpose search engine, as well as on query formulation and filtering techniques based on classification (Moumtzidou et al, 2012). Specifically, we employ the Yahoo! BOSS API[7] as a general purpose search engine, while the queries are formulated by combining terms of the PESCaDO ontology[8], keyword

[6] PESCaDO stands for "Personalized Environmental Service Configuration and Delivery Orchestration". The development of the PESCaDO service has been funded by the European Commission under the contract number FP7-ICT-248594.

[7] http://developer.yahoo.com/search/boss/

[8] http://pescado-project.eu/ontology.php

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

spices (i.e. domain specific terms extracted using supervised machine learning) and geographical terms. The created queries are submitted to the general purpose search engine. The filtering (i.e. distinguishing relevant results from irrelevant ones) is performed using binary classification based on textual features (textual concept extraction and bag-of-words model) and Support Vector Machines SVMs. The expert is able to directly modify the automatically formulated queries by adjusting the keyword spices and introducing new domain-specific terms. In addition, the expert selects the geographical area for which environmental nodes are to be discovered; cf. Figure 1. That is, the intervention of the expert is of corrective and active nature.

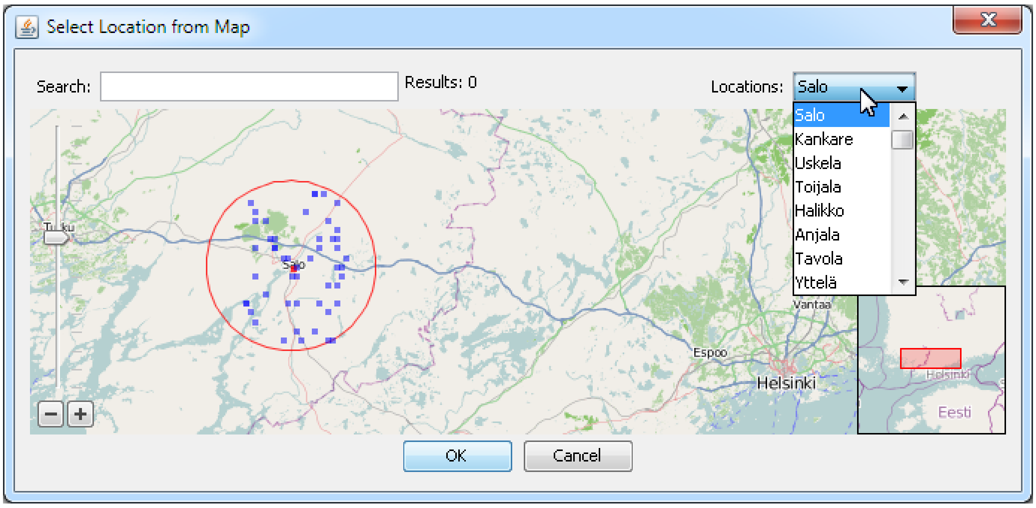

Figure 1
Geographical area selection

2.2 Assessment of the relevance of the identified sites

Once the search engine comes up with some site hits, their relevance must be determined. This is done in PESCaDO using interactive Support Vector Machine (SVM)-based classification (Vrochidis et al., 2012). The user is involved directly in the core of the classification process via a relevance feedback-based visualization module that she uses in the corrective mode. The initial training set is constructed manually and the SVMs are trained with textual features. To support interactive classification, we provide a tool (see Figure 2 immediately below) that shows the websites' feature vectors in relation to a visual abstraction of the classifier to let them assess the classifier's performance interactively.

The most prominent component (Main View – Figure 2a) is the visual abstraction of the decision boundary, which is presented as a straight vertical white space separating the non-relevant region on the left from the relevant region on the right. The Cluster View (Figure 2b) provides in clusters the results based on the distance from the decision boundaries. The Important Terms View (Figure 2d) shows to the user how the classifier "perceives" the websites (i.e. show the most prominent features of the support vectors). If multiple sites are selected, the websites preview in the low left corner (Figure 2c) shows the titles of the websites as a list. The aggregated content of multiple sites can be inspected with the Term Lens (circle with terms to its right – Figure 2f). Using this visualization tool the user is capable of relabeling the websites providing in that way explicit feedback (i.e. which website is relevant). After some labeling actions, the users can select to retrain the main SVM model. During the retraining procedure, the layout in the Cluster and in the Main view are recomputed, the bar chart with the new feature weights is updated, and new results are generated.

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

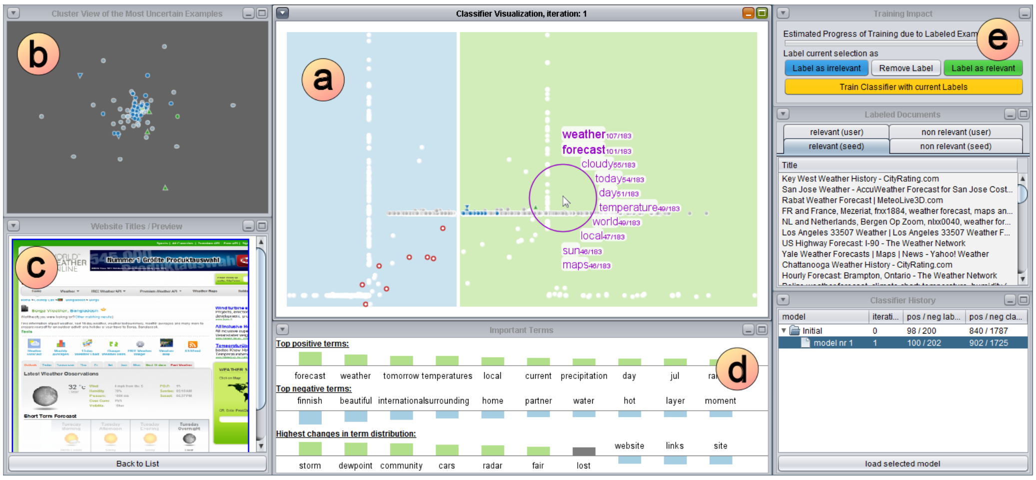

Figure 2a-e
Interactive classification based on visualization layer

2.3 Assessment of the quality of the data

The quality of the data distilled from the sites of the relevant environmental nodes is assessed by applying data uncertainty and confidence metrics. The derivation of these metrics has been an important task in PESCaDO (Epitropou et al., 2012; Johansson et al., 2012). One way to obtain them is the interaction with the expert. Our interaction strategy for this task is based on a set of criteria and a database of artifacts[9] to be evaluated. The strategy can be separated into the following steps: (i) select representative artifacts or artifact groups; (ii) transform the criteria values into confidence values for the confidence metric; (iii) weight and combine the resulting values; (iv) review the confidence metrics on larger scales; and (v) adjust and override the metrics by a faceted metric determination mechanism. This process is depicted in Figure 3, where steps two and three are condensed into the metric itself.

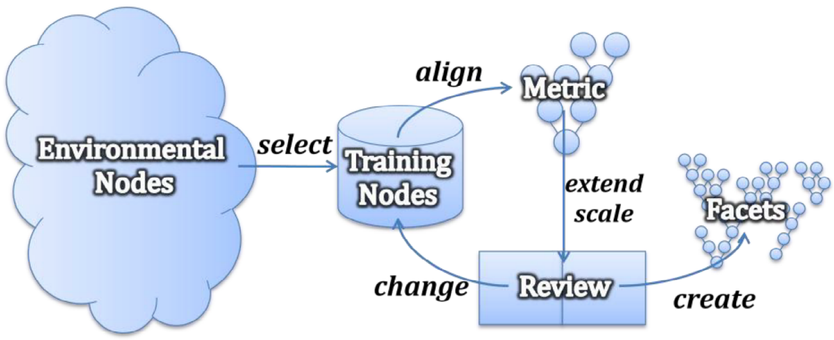

Figure 3
Overview of the metric determination strategy

The outputs of the strategy are graph structures containing transformation rules to receive metric scores from criteria values, as well as tree structures that describe how – and with which weights – the scores are combined into a final score, and also a description of which of these structures are suitable for which kind

[9]We refer to the environmental nodes, websites, and web services that are being evaluated with the metrics as *artifacts*.

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

of environmental node. Together, these outputs can be used to calculate a single score for a given environmental node.

Figure 4 shows examples of the elements that are related to the determination of the confidence metric. The elements for the uncertainty metric in these steps are analogous. This helps to avoid confusion on whether we refer to nodes as parts of a graph structure or environmental nodes and also abstracts from the concrete type of environmental node.

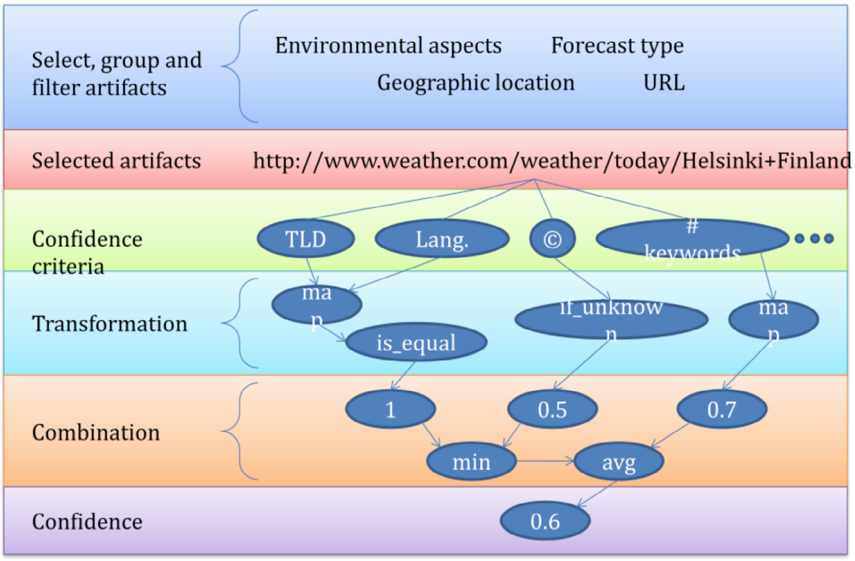

Figure 4
Examples of the elements of the metric determination process

Each step of the metric determination process is carried out in interaction with the expert, who is thus involved in a highly active way. Consider, for illustration, the corresponding interface for the transformation of the criteria into confidence values in Figure 5. Available criteria and operator nodes can be added to the workspace from a list and removed via a delete button. Nodes depicted in the interface in green (in Figure 5, these are “TDL” and “Language”) show that a value is known for the selected artifacts. The nodes depicted in red (in Figure 5: #Keywords) show missing values.

In the example in Figure 5, the top level domain is mapped to a language, which is compared to the language of the artifact. The result is a metric score of 1.0 because they are equal. The number of keywords is missing; therefore a default score of 0.25 is used.

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

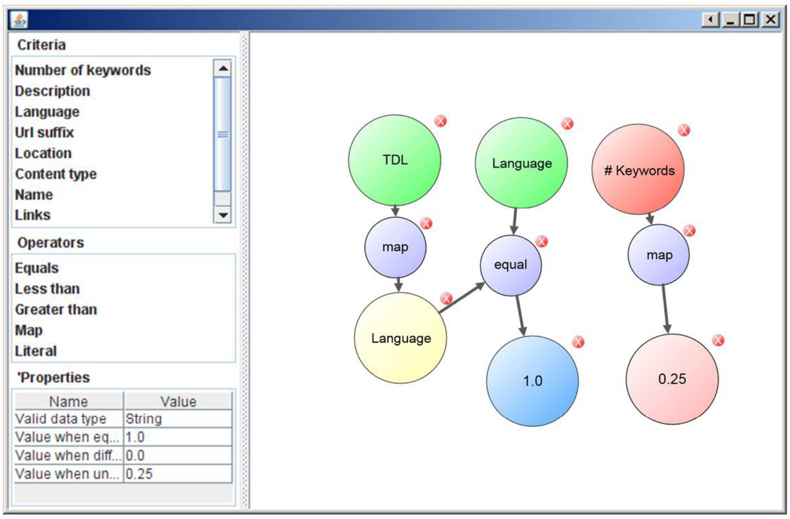

Figure 5
Mapping criteria onto confidence values

2.4 Selection of the content to be communicated in the bulletin

Expert users of the personalized environmental system expect to be able to exert control over the selection of the content to be communicated in a bulletin rather than have this content determined fully automatically. The content selection functionality supports two modes of operation, an automatic selection of content and an interactive content selection strategy for expert users. The raw content in PESCaDO's Knowledge Base (KB) cannot be presented directly to the user because it is difficult to comprehend for users not familiar with Semantic Web representations in which the KB in PESCaDO is encoded. Besides, there is potentially a large amount of content added to the KB following an end user request. For this reason, the interactive content selection functionality is supported by a natural language-based user interface integrated in the main project UI (see Figure 6) and a machine learning mechanism.

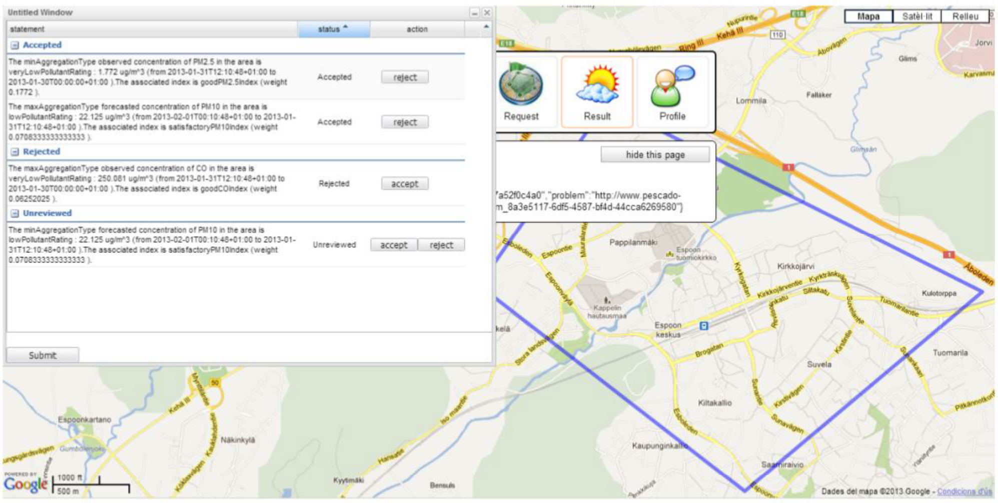

Figure 6
UI of the interactive content selection mode

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

The user interface renders the content in an accessible way by using natural language, while the learning mechanism (based on reinforcement learning methods) aims at predicting what content administrative users would select in each scenario, thus simplifying the task of browsing all the content in the KB.

The interactive selection task is formulated as a sequential decision process in which the content in the KB is divided into subsets using content templates. Each subset is verbalized as a short text. The expert is prompted with one of the short texts at a time and given the choice between accepting and rejecting the text (and, therefore, the contents it conveys). In other words, the expert user is involved in this task in an evaluative mode.

Once the user considers she has seen enough of the content, a high-quality text is generated from the selected data. Unlike the generation of the final report, the verbalization of content during the interactive selection process does not rely on fully-fledged text generation. Instead, it uses fill-in text templates, one per each content template used by the system to select contents from the KB; cf. a number of examples:

- The general Air Quality Index is very poor. This is due to a very high concentration of CO. This is due to a very high concentration of O3. This is due to a high concentration of NO2.
- The concentration of carbon monoxide (CO) reached a maximum of 180 micrograms per meter cubed (107 minimum). The index for this measurement is very poor and the rating is very high. Warning: Carbon monoxide causes oxygen deficiency, which is most harmful to persons suffering from cardiovascular and pulmonary diseases, or anaemia, and also to the elderly, pregnant women and infants.
- A concentration of 180 micrograms of carbon monoxide (CO) exceeded the warning threshold of 160 micrograms in Suomelina at 09:00h 03/09/2011.
- The forecasted index of carbon monoxide (CO) will range from poor to fair. Warning: Adverse health effects are possible on sensitive individuals.
- The maximum wind speed measured was 9.7 meters per second (strong wind speed) and the minimum was 2.3 meters per second (weak wind speed). Recommendation: Umbrella use becomes difficult. Empty plastic garbage cans tip over.

3. Discussion and Future Work

Interactive solving of complex tasks is an increasingly popular methodology across different scientific fields and different types of applications. The motivation may differ from field to field and from application to application. For some of the applications, the involvement of the expert or the end user is opted for due to the fact that the system's processing and decision strategies require knowledge not encoded in the system – such that external guidance is required. For other applications, their involvement can be justified by less objective motives: it is well-known that the acceptance of an application by users (be they experts or naïve end users) rises if they are given the possibility to intervene.

In the case of the process of the delivery of environmental information from the web, the involvement of the experts proved to be very positive from the perspective of both the performance of the service and its acceptance. Thus, the quantitative and qualitative evaluation carried out for the modules that handle the tasks in which the experts were involved shows that their intervention led to a high accuracy of the outcome of the modules. Test trials in which expert and naïve end users carried out hands-on experiments confirmed that the involvement of experts in the processing loop was appropriate.

The PESCaDO service had been projected for the use in a series of applications – including internet-based personalized delivery of environmental information for citizens, support of municipal administration in environmental decision making, validation of the quality of environmental web-based service nodes,

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5

etc. The involvement of the expert in the central tasks dealt with by the service is decisive for achieving the required quality for these applications.

There are several aspects of the interactive side of the PESCaDO service that can be further improved in the future. For instance, an integrated interface for all types of system–expert interaction would ease the communication of the system with the user and be more in line with the principles of Human-Computer Interaction software design. Furthermore, there is some evidence that the involvement of the end user in some of the tasks (as, e.g., exclusion of some environmental nodes from consideration or giving preference to some of the environmental nodes of the preference of the user) would further increase the performance and acceptance of PESCaDO.

Bibliography

Card, S. and P. Pirolli (2005): Sensemaking Processes of Intelligence Analysts and Possible Leverage Points as Identified Through Cognitive Task Analysis, in: Proceedings of the International Conference on Intelligence Analysis

Epitropou V., Johansson L., Karatzas K., Bassoukos A., Karppinen A., Kukkonen J. and Haakana M. (2012): Fusion of Environmental Information for the Delivery of Orchestrated Services for the Atmospheric Environment in the PESCaDO Project, in: R. Seppelt, A.A. Voinov, S. Lange, D. Bankamp (eds.). International Environmental Modelling and Software Society (iEMSs), 2012 International Congress on Environmental Modelling and Software, Managing Resources of a Limited Planet, Sixth Biennial Meeting, Leipzig, Germany

Johansson, L., M. Haakana, V. Tarvainen, and A. Karppinen (2012): Fusion of forecasted and measured air quality data using landscape profiling and historical performance database as a part of an online PESCaDO system, in: Proceedings of the 9th International Symposium on Advanced Environmental Monitoring and Modeling (ISAEMM), Helsinki

Keim, D.A., J. Kohlhammer, G. Ellis and F. Mansmann (eds.) (2010): Mastering the Information Age. Solving Problems with Visual Analytics, Eurographics Association, Goslar

Jackson, P. (1998): Introduction to Expert Systems, Reading: Addison Wesley

Moumtzidou, A., Vrochidis, S., Tonelli, S., Kompatsiaris, I., and Pianta E. (2012): Discovery of Environmental Nodes in the Web, in: Proceedings of the 5th IRF Conference, Austria, Vienna

Vrochidis, S., Bosch, H., Moumtzidou, A., Heimerl, F., Ertl, T., Kompatsiaris, I. (2012): An Environmental Search Engine based on Interactive Visual Classification, in: Proceedings of the ACM International Workshop on Multimedia Analysis for Ecological Data (MAED 2012), Nara, Japan

Wanner, L., S. Vrochidis, S. Tonelli, J. Moßgraber, H. Bosch, A. Karppinen, M. Myllynen, M. Rospocher, N. Bouayad-Agha, U. Bügel, G. Casamayor, T. Ertl, I. Kompatsiaris, T. Koskentalo, S. Mille, A. Moumtzidou, E. Pianta, H. Saggion, L. Serafini, and V. Tarvainen (2011): Building an environmental information system for personalized content delivery, in: Proceedings of the 9th IFIP WG 5.11 International Symposium on Environmental Software Systems – Frameworks of eEnvironment (ISESS 2011), Brno, Czech Republic, pp. 169–176

Copyright 2013 Shaker Verlag, Aachen, ISBN: 978-3-8440-1676-5