PEGOV 2014:
2nd International Workshop on Personalization in eGovernment Services and Applications

Nikolaos Loutas[1], Fedelucio Narducci[2], Adegboyega Ojo[3], Matteo Palmonari[2], Cécile Paris[4], and Giovanni Semeraro[5]

[1] PwC, Belgium
nikolaos.loutas@pwc.be
[2] Department of Informatics, Systems and Communication
University of Milano-Bicocca, Italy
surname@disco.unimib.it
[3] Insight Centre for Data Analytics, National University of Ireland, Galway, Ireland
adegboyega.ojo@deri.org
[4] CSIRO - Computational Informatics, Australia
Cecile.Paris@csiro.au
[5] University of Bari "Aldo Moro", Italy giovanni.semeraro@uniba.it

1 Preface

User modeling and personalization have proved to play a strategic role in adapting the behavior of intelligent systems to the specific characteristics of their users. The systems exploiting user-profiling strategies are able to give a different answer to requests derived by different users. In this context, it is important to define effective and accurate techniques for extracting the users characteristics, interests and preferences to be used in the next steps of the personalization pipeline (adaptation, recommendation, etc.).

In parallel, e-Government (e-Gov) has, in recent times, transformed interactions between governments, citizens and other stakeholders in the society. Public services and public sector information can now be delivered electronically through Web portals and mobile apps. In this new context, innovative solutions that are better tailored to citizens' needs can facilitate better access to e-Gov services and reduce the red tape that usually characterizes the provisioning of public services. In addition, government and citizens can better engage with each other. Following these initiatives, governments are now looking at providing personalized services, often grouped in life-events and business episodes. Another interesting development related to governments is the recent push towards more openness of public sector information, with an emphasis on opening up government data, which presents new application areas and opportunities for personalization. This trend has specifically created the need for personalized access to Open Government Data (OGD) predominantly by means of visualizations and faceted browsers. It has also given rise to opportunities for improved decision making, as well as recommendation and personalization of e-Gov services.

Personalization methods and user modeling techniques have been applied successfully in several domains (e.g., e-commerce), and personalization is being extensively studied in domains characterized from the digital-object consumption (e-commerce, news, music, video recommendations, etc.). However, personalization in the e-Gov domain is still in its infancy, at least in production systems. One reason for this is the difficulty in the government domain to obtain some of the information required for personalization, due to personal data privacy and confidentiality constraints. For example, while in an e-commerce portal, one can quite easily obtain user preferences, in the e-Gov domain, the concept of preference itself is difficult to define, and confidentiality, ethical and privacy concerns play a more important role. Significant ethical issues could prevent governments from leveraging the direct access they have to citizen information, and undeclared or automatic user profiling could be considered intrusive and undesirable.

The main goal of this workshop is to stimulate the discussion around problems, challenges and research directions about personalization in e-Gov, with a dual focus on both services and OGD. Following up on the successful inaugural edition of the UMAP-PEGOV workshop of last year, we aim with this workshop to stimulate further interest of the scientific and business communities on the aforementioned issues to move towards more user-aware and adaptive services in the e-Gov domain by means of personalization methods. We are particularly interested in the role of user modeling and profiling in e-Gov service design and delivery, as well as in the provision and access to OGD, considering related issues like privacy, security and multilingualism. This workshop builds on other workshops concerned with e-Government and citizen engagement (cf: Gandrup Borchorst et al., 2011; Loutas et al., 2013).

The original topics of interest listed on the call for paper for the workshop included:

- Motivation, benefits and roadblocks of personalization in e-Gov
- Recommender systems for integrated public sector information
- Approaches for the personalization of inclusive and interactive e-Gov services
- User and context awareness in personalization of e-Gov services
- Multilingual e-Gov services
- Adaptation, personalization and recommendation models and goals in e-Gov
- Personalized access to (Linked) Open Government Data (OGD)
- Visualization and faceted search for (Linked) OGD
- Big Data analytics for user profiling in e-Gov
- User, group and family modeling in e-Gov
- Mining of user behavior, opinion mining, and sentiment analysis in e-Gov
- User preference measurement and econometric indicators
- Applications for subjective well-being and happiness assessment
- Persistence, removal, and update of citizen profiles
- Semantic techniques for user profiling and personalization in e-Gov
- Ethical issues, including privacy, in e-Gov and OGD
- Usability of e-Gov applications, covering both services and OGD
- Evaluation of personalized services in e-Gov

– Applications of personalization methods in e-Gov
– Communities and social networks in participatory e-Gov
– Citizen-centered service design and modeling

‘ We received 8 submissions; we accepted five long papers, two short papers and rejected one paper. Each submission was reviewed by at least two PC members (none of the chairs has been involved in the review process). Two of the accepted papers were ultimately withdrawn by their authors. These proceedings thus include four long papers and one short paper. The selected papers show the wide variety of topics and issues that arise in personalization for e-Government, from mining Twitter data to better understand the citizens to explicit crowd sourcing, to the personalization of the information space, documents and services, and, finally, the assessment of value of government services as perceived by citizens.

For several years now, social media, Twitter in particular, has been exploited to assist in a variety of government decisions, including: emergency management (e.g., Kireyev et al., 2009; Heinzelman and Waters 2010; Sakaki et al., 2012; Bruns et al., 2012; Cameron et al., 2012; Griffen et al., 2012; Yin et al., 2012; Karimi et al., 2013; Kreiner et al., 2013; Power et al., 2013); health-related matters (e.g., Leaman et al., 2010; Chee et al., 2011; Liu et al., 2011; Bian et al., 2012); and citizens engagement, i.e. or more generally to listening to citizens feedback to improve government services by and understanding citizens behaviours and concerns in order to improve government services by (e.g., Loutas et al., 2011; Paris and Wan, 2011; Loutas et al., 2012; Wan and Paris, 2014). In their paper, Villena-Román continue this trend to make use of social media to informe public entities. Their paper is entitled “TweetAlert: Semantic Analytics in Social Networks for Citizen Opinion Mining in the City of the Future”. It describes work done in the framework of the Ciudad 2020 Spanish national R&D project, specifically a system that aims to analyze tweets in real time, annotating them with information such as their topic and sentiment, and presenting them to an end user (such as a government agency employee) in a variety of visualisations.

While mining social media to obtain information about a number of topics related to citizens and government matters can already be seen as crowd sourcing, explicit crowd sourcing (i.e., actively asking citizens to contribute to a specific site with their comments) is being exploited by an increasing number of government agencies to manage participary decision making (e.g., Public Sphere 2, 2009; Karamanou et al., 2011; Lee et al., 2011) and obtain information on various topics, such as road conditions (e.g., RAVC Pot Hole Patrol). Liliana Ardissono and her colleagues present a proposal for the management of 3D Community Maps which provide a virtual representation of a locality, also enabling users to contribute to policy making. Their paper, entitled “Community Mapping for Participatory decision-Making Processes” presents in particular an analysis of user requirements and personalization issues for this type of application. In particular, it suggests the selection of specific information from the community maps based on individual interests. This is of course the personalization of the information space for individual users, albeit in a specific context.

The personalization of information for individual users in the e-Government domain is further explored from two different perspectives: one of retrieval, following work such as (Chen and Sycara, 1998; Sugiyama et al., 2004; Gauch et al., 2007). In their paper entitled "Personalization of Parliamentary Document Retrieval using different user profiles", Vicente-López and his colleagues performed a comparative study of several content-based user profile representations to support citizens' access to the Records of Parliament Proceedings in Andalucìa (Spain). They show that personalization helps citizens find information relevant for their needs. In a different line of research, based on the work on the production of personalized documents, where "one-size-fits-all is replaced by one document for one user, as in (Colineau et al, 2012a & b), Penadés and her colleagues propose a method that enables the generation of personalized documents in domains with high variability and with high levels of reuse. Their paper, "Product Line-based customization of e-Government documents" presents their approach (called Document Product Lines, or DPL), together with the principles that underlie it and a case study.

These proceedings finish with a paper by Torsello et al., entitled "A fuzzy model for service value assessment". While it is assumed that personalization in e-government is beneficial to businesses and citizens (Baldassarre et al., 2013; Palmonari et al., 2008), to avoid information overload and ensure citizens receive the appropriate services with the least amount of effort, this proposition must to be verified. In particular, the personalization usually comes at a cost, for example the disclosure of information on the part of the citizen and the construction of a profile to be kept by government agencies. Some evaluations have already been performed (e.g., Colineau et al., 2012a). In their work, Torsello et al. present a model for the assessment of the service value of government services. As in other models which propose to balance costs and benefits (e.g., Paris et al., 2009), Torsello et al. define a service value as a trade-off between benefits and sacrifices, in this case as perceived by citizens according to their experiences using the services. Their model proposes the user of fuzzy concepts to reflect the fact that human perceptions are subjectives.

We hope the workshop will stimulate discussions around the presented papers, the invited talk and the following questions:

1. How can personalization methods support the design of services and applications, which better adapt to the different roles that a citizen/business plays when interacting with public administrations?
2. Which user characteristics (demographic, cultural, family, etc.) can influence the design and delivery of e-Gov services as well as the access and reuse of OGD?
3. How can citizens be involved in the design of adaptive e-Gov platforms and services?
4. To what extent are the general techniques adopted for user modeling and profiling in different domains suitable for modeling the citizen characteristics?

5. Can semantic models and ontologies support the representation of prototypical users in order to identify categories of citizens based on different characteristics?
6. Can novel methods for socio-economic analysis based on mobile applications be used for driving the personalization of access to OGD and e-Gov services?
7. How can personalization methods improve the access to OGD, e.g., with proper visualizations, faceted browsers, and/or suitable recommendations?
8. Can semantic models and ontologies support the representation of prototypical users in order to identify categories of citizens based on different characteristics?
9. How can ethical issues and privacy be addressed to positively influence trust in personalized e-Gov services?
10. Would personalization methods be favorably accepted and desired by citizens?

2 Workshop Chairs

- **Nikolaos Loutas**, PwC, Belgium.
  Nikolaos is manager at PwC's Technology Consulting practice, involved mainly in projects on interoperability of trans-European ICT solutions, data and software products. Nikolaos specialises in semantic aspects of interoperability, through the application of Semantic Web technologies and Linked Data. He has deep insights into open semantic standards, such as the Asset Description Metadata Schema, the e-Government Core Vocabularies and the DCAT Application Profile for data portals in Europe. Nikolaos is currently driving the Open Data Support project of DG CONNECT, which aims at facilitating the access of citizens and business to Open Government Data published by governments across Europe. Before joining PwC, Nikolaos had been working for leading EU research centers. He has published more than 55 papers and reports in the field of Semantic Web in international journals, conferences and books.
- **Fedelucio Narducci**, University of Milano-Bicocca, Italy. Swap Research Group, University of Bari Aldo Moro, Italy
  Fedelucio Narducci is research assistant at University of Milan-Bicocca, Department of Informatics, Systems and Communication. He is also member of the SWAP (Semantic Web Access and Personalization) research group of University of Bari Aldo Moro. His primary research interests lie in the areas of machine learning, content-based recommender systems, user modeling, and personalization. From April 2012 he is working for the SMART (Services & Meta-services for smART eGovernment) project whose goal is to define models, methodologies, languages for planning, production and delivery of services characterized by optimal social value, value of use, and value of exchange. He served as Co-chair of Pegov 2013. Fedelucio was reviewer and co-reviewer for international conferences and journals on the topics of recommender system, user modeling and personalization. He is also author of several papers in international conferences and journals.

– **Adegboyega Ojo**, INSIGHT Center for Data Analytics, National University of Ireland, Galway
Adegboyega Ojo is a Research Fellow and leads the E-Government Group at The INSIGHT Center for Data Analytics, National University of Ireland, Galway; Republic of Ireland. His research focuses on how to drive innovations in government organizations through the applications of Semantic Web, Linked Open Data and Collaboration technologies. His current portfolio of research and development projects is funded under the Seventh Framework Programme of the European Commission. Before his current role, he worked as Academic Program Officer, Research Fellow and Post-doctoral Fellow at the Center for Electronic Governance, United Nations University International Institute for Software Technology (UNU). At UNU, his work benefitted several governments including Macao, Korea, Mongolia, Colombia, Cameroon and Nigeria. He has published widely in the areas of Strategies, Architecture and Standards, e-Participation, Open Governance and Open Data. He obtained his PhD at the University of Lagos, Nigeria (1998), where he was appointed Senior Lecturer and Associate Professor in Computer Science in 2003 and 2012 respectively. He is also Adjunct Lecturer at the National University of Ireland, Galway.
– **Matteo Palmonari**, University of Milano-Bicocca, Italy
Matteo Palmonari is an assistant professor in the Department of Informatics, Systems and Communication at the University of Milan-Bicocca. His research interests include semantic matchmaking, information quality, knowledge representation, and ontologies for the semantic web; several of his research have been applied to service modeling, service matchmaking and e-Government applications. He has been a visiting postdoc and a visiting assistant professor with the ADVIS Laboratory, University of Illinois at Chicago. He has published more than 40 papers in international journals and conferences.
– **Cécile Paris**, CSIRO, Computational Informatics, Australia
Dr Cécile Paris is a Science Leader at the CSIRO Computational Informatics, Sydney, Australia. Dr Paris also holds Adjunct Professorships at Macquarie University (Sydney) and the ANU (Australian National University, Canberra, Australia). Dr Paris leads the research in Language and Social Media. Dr Paris received her B.A. degree in Computer Science from The University of California at Berkeley, USA, and her Masters and PhD degrees from Columbia University, New York, USA. Her PhD was one User Modeling and Natural Language Generation. Her main research interests lie in the areas of personalized information delivery and language technology. She has been involved in e-Government for over 5 years, and her current work includes tailored delivery for Public Administration, online communities and social media in the context of e-Government. Dr Paris co-organised the workshop on Government and Citizen Engagement at the Communities and Technology conference in 2011. In 2011, she was an invited speaker at the 2nd (Australian) Public Officer Digital Media Forum, and at the 7th Annual AIMIA Digital Summit (AIMIA is the Australian Interactive Me-

dia Association). She was a keynote speaker at the (Australian) Emergency Management New and Emerging Technologies Forum in October 2013 and at the National Medicine Symposium in May 2014. Dr Paris has authored over 250 referred technical articles at international journals and conferences. She is currently the chair of CHISIG, the Computer Human Interaction Special Interest Group of the Human Factors and Ergonomics Society of Australia.

– **Giovanni Semeraro**, SWAP Research Group, University of Bari Aldo Moro, Italy
He is associate professor of computer science at the University of Bari Aldo Moro and leads the Semantic Web Access and Personalization Research Group Antonio Bello. His research interests include AI, recommender systems, user modeling, personalization, intelligent information retrieval, semantic and social computing, the Semantic Web, natural language processing, and machine learning. He received his M.Sc. degree in computer science from the University of Bari. He served as General Co-chair of UMAP 2013, IIR 2013, SemExp 2012, IIR 2012, IIA 2008, AI*IA 2008, SWAP 2007, CILC 2006, and as Program Co-chair of Decisions@RecSys 2013 & 2012, DART 2013, 2012 & 2011, RSmeetDB@DEXA 2013 & 2012, SeRSy@RecSys 2013 & SeRSy@ISWC 2012, DEMRA@UMAP 2011, SPIM@ISWC 2011, EC-Web 2010, SWAP 2010, Web Mining 2.0@ECML/PKDD 2007, ISMIS 2006, Web-Mine@ECML/PKDD 2006, IEA-AIE 2005, He is co-author of more than 300 papers published in journals, international conferences and workshops.

References

1. Baldassarre, C., Cremaschi, M. And Palmonari, M. (2013). Bridging the Gap between Citizens and Local Administrations with Knowledge-Based Service Bundle Recommendations. In Morvan, F., Tjoa, A.M., & Wagner, R. (Eds.): 24th International Workshop on Database and Expert Systems Applications, DEXA 2013. IEEE 2013: 157-161
2. Bian, H., Topaloglu, U. And Yu, F. (2012). Towards large-scale Twitter mining for drug-related adverse events. In Yang, C. C, Chen, H., Wactlar, H., Combi, C. &Tang, X. (program chairs). Proceedings of the 2012 International Workshop on Smart Health and Wellbeing (SHB 12). New York, NY, USA: ACM.
3. Bruns, A., Burgess, J., Crawford, K. And Shaw, F. (2012). #qldfloods and @QPSMedia: Crisis Communication on Twitter in the 2011 South East Queensland Floods. Brisbane, QLD, Australia: ARC Centre of Excellence for Creative Industries and Innovation.
4. Cameron, M., Power, R., Robinson, B. And Yin, J. (2012). Emergency Situation Awareness from Twitter for Crisis Management. In Mille, A., Gandon, F., Misselis, J., Rabinvich, M. &Staab, S. (general and program chairs). Proceedings of the 21st International Conference Companion on World Wide Web (WWW 12 Companion). New York, NY, USA: ACM.
5. Chee, B.W., Berlin, R. And Schatz, B. (2011). Predicting adverse drug events from personal health messages. AMIA Annu Symp Proc. 2011. p. 217?226.
6. Chen, L. And Sycara, K. (1998). A Personal Agent for Browsing and Searching. In Proceedings of the 2nd International Conference on Autonomous Agents, 132-139.

7. Cuidad 2020: http://innprontaciudad2020.es
8. Colineau, N., Paris, C. And Vander Linden, K. (2012a). An evaluation of tailored Web materials for Public Administration. In the Proceedings of the 23rd ACM Conference on Hypertext and Social Media, Adaptive Hypertext and Hypermedia track, June 25-28, 2012, Milwaukee, WI. USA, p.265-274.
9. Colineau, N., Paris, C. and Vander Linden, K. (2012b). Government to citizen communications: From generic to tailored documents in public administration. In Information Polity, 17(2):177-193, Special Issue on Open Government and Public Participation, Editors: L. F. Luna-Reyes, S. Ae Chun and V. Atluri. IOS press.
10. Gandrup Borchorst, N., Bdker, S., Colineau, N., Korn, M. and Paris, C. (2011). Proceedings of the C&T 2011 Workshop on Government and Citizen Engagement. June 30, Brisbane. Publisher: IISI - International Institute for Socio-Informatics (Volume 8, Issue 2, 2011, ISBN 1861-4280).
11. Griffen, G., Jones, R. And Paris, C. (2012). Strategic implications of social media for emergency management. In Clark, M. &Griffen, G. (eds.). Next Generation Disaster and Security Management. Canberra, ACT, Australia: The Australian Security Research Council.
12. Grauch, S., Speretta, M., Chandramouli, A. And Micarelli, A. (2007). User profiles for personalized information access. The Adaptive Web: Methods and Strategies of Web Personalization. LNCS 4321, 54-89, 2007.
13. Heinzelman, J. and Waters, C. (2010). Crowdsourcing Crisis Information in Disaster-affected Haiti. United States Institute of Peace Special Report. Washington DC, USA: United States Institute of Peace.
14. Karamanou A., Loutas N., Tarabanis K. (2011) ArgVis: Structuring Political Deliberations Using Innovative Visualisation Technologies. In Proceedings of 3rd ePart Conference, Delft, The Netherlands.
15. Karimi, S., Yin, J. and Paris, C. (2013). Classifying microblogs for disasters. In Culpepper, S., Zuccon, G. &Sitbon, L. (eds.). Proceedings of the 18th Australasian Document Computing Symposium (ADCS 13). New York, NY, USA: ACM.
16. Kireyev, K., Palen, L. And Anderson, K.M., (2009). Applications of topics models to analysis of disaster-related Twitter data. In. NIPS 2009 Workshop on Applications for Topic Models: Text and Beyond. Amherst, MA, USA, December 2009.
17. Kreiner, K., Immonen, A. And Suominen, H. (2013). Crisis management knowledge from social media. In Culpepper, S., Zuccon, G. &Sitbon, L. (eds.). Proceedings of the 18th Australasian Document Computing Symposium (ADCS 13). New York, NY, USA: ACM.
18. Leaman, R., Wojtulewicz, L., Sullivan, R., Sakariah, A., Yang, J. and Gonzalez, G. (2010). Towards Internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks. In. Proceedings of the 2010 Workshop on Biomedical Natural Language Processing (BioNLP 10). Stroudsburg, PA, USA: ACL.
19. Lee D., Loutas N., Sanchez Nielsen E., Mogulkoc E., Lacigova O. (2011) Inform-Consult-Empower: A Three-tiered Approach to eParticipation. In Proceedings of 3rd ePart Conference, Delft, The Netherlands.
20. Liu, J., Li, A. And Seneff, S. (2011). Automatic drug side effect discovery from online patient-submitted reviews: focus on statin drugs. In. Proceedings of the First International Conference on Advances in Information Mining and Management (IMMM 11). Barcelona, Spain, October 2011.
21. Loutas, N., Narducci, F., Palmonari M. and Paris, C. (2013). Proceedings of PE-GOV 2013: 1st International Workshop on Personalization in eGovernment Services

and Applications. Held in conjunction with UMAP (User Modelling, Adaptation and Personalization), Rome, June 2013.
22. Loutas N., Peristeras V., Zeginis D., Tarabanis K. (2012). The Semantic Service Search Engine (S3E). In Journal of Intelligent Information Systems, vol. 38 (3), pp. 645-668.
23. Loutas N., Peristeras V., Tarabanis K. (2011). A citizen-centric semantically-enhanced governmental portal. In Electronic Government, An International Journal, vol. 8 (4), pp. 363-384.
24. Palmonari, M., Viscusi, G. And Batini., C. (2008). A semantic repository approach to improve the government to business relationship. Data Knowl. Eng. 65(3): 485-511 (2008)
25. Paris, C., Colineau, N. and Wilkinson, R. (2009). Evaluating Web-Based Information Systems from a Cost-Benefit Perspective: A Case Study. In the Proceedings of Hypertext 2009, the 20th ACM Conference on Hypertext and Hypermedia, June 29th—July 1st, Torino, Italy. ISBN: 978-1-60558-486-7. PP 353-355. CSIRO Number: 09/051.
26. Paris, C., and Wan, S. (2011). Listening to the community: social media monitoring tasks for improving government services. In Tan, D., Begole, B. &Kellogg, W. A. (general and program chairs). Proceedings of the HI 11 Extended Abstracts on Human Factors in Computing Systems (CHI EA 11). New York, NY, USA: ACM.
27. Power, R., Robinson, B. And Radcliffe, D. (2013). Finding fires with Twitter. In Karimi S &Verspoor K. Proceedings of the 2013 Australasian Language Technology Association Workshop (ALTA 13), Brisbane, QLD, Australia, December 2013.
28. Public Sphere 2: Government 2.0 Briefing Paper (2009) http://www.katelundy.com.au/wp-content/uploads/Public-Sphere-2-Government-2.pdf (retrieved March 29, 2010).
29. RAVC Pot Hole Patrol: https://www.facebook.com/RACVPotholePatrol. (Accessed May 31, 2014).
30. Sakaki, T., Okazaki, M. And Matsuo, Y. (2010). Earthquake shakes Twitter users: real-time event detection by social sensors. In Rappa, M., Jones, P., Freire, J. &Chakrabarti, J. (general and program chairs). Proceedings of the 19th International Conference Companion on World Wide Web (WWW 10 Companion). New York, NY, USA : ACM.
31. Sugiyama, K., Hatano, K. And Yoshikawa, M. (2004). Adaptive web search based on user profile constructed without any effort from users. In Proceedings of the 13th International Conference on the World Wide Web (WWW 2004), 675-684.
32. Yin, J., Lampert, A., Camern, M., Robinson, B. And Power, R. (2012). Using social media to enhance emergency situation awareness. IEEE Intelligent Systems 27 (6). p. 52?59.
33. Wan, S. And Paris, C. (2014). Improving government services with social media feedback. In Proceedings of the the 2014 International Conference on Intelligent User Interfaces (IUI 2014). Haifa, Israel, Feb 24-27, 2014.

TweetAlert: Semantic Analytics in Social Networks for Citizen Opinion Mining in the City of the Future

Julio Villena-Román[1,2], Adrián Luna-Cobos[1,3],
José Carlos González-Cristóbal[3,1]

[1] DAEDALUS - Data, Decisions and Language, S.A.
[2] Universidad Carlos III de Madrid
[3] Universidad Politécnica de Madrid
{jvillena,aluna}@daedalus.es, josecarlos.gonzalez@upm.es

Abstract. In this paper a highly configurable, real-time analysis system to automatically record, analyze and visualize high level aggregated information of user interventions in Twitter is described. The system is designed to provide public entities with a powerful tool to rapidly and easily understand what the citizen behavior trends are, what their opinion about city services, events, etc. is, and also may used as a primary alert system that may improve the efficiency of emergency systems. The citizen is here observed as a proactive city sensor capable of generating huge amounts of very rich, high-level and valuable data through social media platforms, which, after properly processed, summarized and annotated, allows city administrators to better understand citizen necessities. The architecture and component blocks are described and some key details of the design, implementation and scenarios of application are discussed.

Keywords: Semantic analytics, social networks, citizen, opinion, topics, classification, ontology, events, alerts, big data, city console.

1 Introduction

With the recent success and proliferation of mobile devices, the democratization of Internet accessibility and the possibility of meta-information, such as user location, user profile and demographics, etc., the vastly amount of data that is being generated has very rapidly grown. This unstructured source of data is already being used in multiple fields like sociology, advertising, etc. and may also be used to improve public administration services and functionality, as a new version of e-Gov application. User interventions in social networks often contains agreement, disagreement or comments about city services, city administrators, events in the city, etc. However, these data are not really useful unless some semantic processing or data mining technique is applied in order to automatically distinguish between relevant and not relevant information and provide a higher level of abstraction.

This work has been developed in the framework of Ciudad 2020 [1] Spanish national R&D project, which aims to achieve improvements in areas such as energetic efficiency, Internet of the Future, Internet of Things, human behaviour, environmental sustainability and mobility and transport, in order to design the City of the Future. The project proposes a new city model designed for the citizen *–ad civitates civis–* that aims to include citizen reality into the city decisions.

Usually, the final objective of the government decisions is the citizen welfare. However, it is not always an easy task for the administration services to quickly identify the most important facts that their citizens are facing, to correctly scale them regarding their relative importance according to what citizens think about them, or just to be quick enough to recognize recent issues that may suddenly appear. In such cases, citizen opinion mining will be a key factor to identify and later solve such concerns. Therefore, the citizen is observed here from a dual point of view: on the one hand as the main user of the services that the city offers, and on the other hand, as a proactive city sensor capable of generating huge amounts of data through social media platforms. The citizen sensor is an innovative way to capture high-level heterogeneous information, very descriptive and with great value, especially when considering aggregations. If the city administrators get to properly analyze such vast amount of data coming from Social Media, they will be able to better know trends, generate hypotheses over urban behaviour models in order to improve municipal management policies, bringing them closer to the actual reality of the citizens, thus, turning them into real actors within management mechanisms of smart cities.

In such process of data understanding and mining, technologies to analyze natural language allow to semantically analyze citizen interventions in social media such as Twitter. The aim of our system is to provide city promoters with a powerful tool to rapidly and easily understand what the citizen behavior trends are, what their opinion about city services, events, etc. is, and finally to provide them a primary alert system that may improve the efficiency of emergency systems. In the same way, but applied to a smaller scale as what we propose here, the system could be used to track public services Twitter profiles' and collect user opinion about e.g., e-Gov sites or applications, allowing them to act more quickly to possible lacks of usability, services failures, etc.

The rest of the paper presents the system description and architecture, and further explores the details of each block that composes the system. Finally, a discussion and future work section with insights to improve the system currently in the development branch are presented.

2 System Architecture

We present a highly configurable, real-time analysis system to automatically record, analyze and visualize high level aggregated information of user interventions in Twitter that may be used by public entities to better understand citizen necessities. The system is composed by four main components, shown in Figure 1.

The central component is the ***datawarehouse***, the core information repository that is able to store the high volume of data that the system manages and also provides advanced search functionality to be able to exploit the information. The system is based on Elasticsearch [2], which is a flexible and powerful open source, distributed, real-time search and analytics engine. Its distributed capabilities and the fact that it scales very good when the system grows were key factors in the selection of this architecture. Elasticsearch runs on top of Apache Lucene, so it offers quite complex search capabilities and a scalable and high-performance environment.

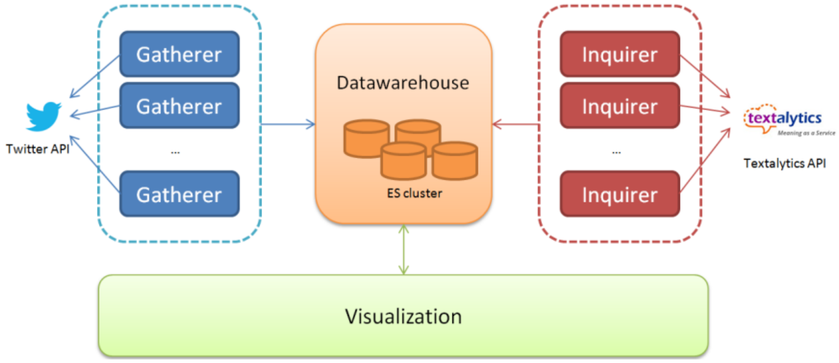

Fig. 1. System architecture

The second component is composed by a set of concurrent ***gatherer*** processes, which query the Twitter APIs [3] to collect tweets regarding to certain filters. The configuration file defines the query parameters to the Twitter streaming API, allowing to filter tweets by a list of user identifiers, a list of keywords to track (terms, hashtags) and/or a set of geographical bounding boxes to restrict the search.

The third component is composed of a set of concurrent ***inquirer*** processes, whose task is to annotate the messages using several of our Textalytics Core APIs [4]. The system is deployed to use the text classification API using two specific models specially designed for this business case (SocialMedia and CitizenSensor, described later), the topics extraction API, which extracts topics such as entities, concepts, money, URI expressions, etc., the sentiment analysis API, which extracts sentiment polarity and also subjectivity and irony indications, and finally, the user demographics API, which currently returns the gender, age and type of the author of the tweet.

Specifically, for each tweet, the system tries to identify the thematic area of the message (energy, transport, economy, politics, social interests…), concepts mentioned (city services, weather...), events to which the text refers (cultural events, soccer matches...), special alert situations (road accidents, fires, street violence, security issues…), and the specific location of the user (a building, means of transport...). This analysis is complemented by an analysis of the sentiment polarity of the message: very positive, positive, negative, very negative and neutral.

An example of an annotated tweet is shown in Figure 2, where a Twitter user alerts from a crash in a public tunnel of the city of Madrid that needed of the presence of the firemen. The system correctly detects that the issue is located in a *public road*, classifies the message in the topic of *Security* in the Citizen Sensor ontology (under *Concepts>Services>Security*) and as *Disasters and accidents* in the general Social Media ontology. Furthermore, it finds out that the entity *Calderón* (soccer stadium nearby) appears in the sentence and also several concepts: *accident*, *tunnel*, *closed*, *firefighter*, *exit*, *lane*, etc. Finally, it detects that it is an objective, non ironic comment with negative polarity written out by a male aged in the range of 35 to 65 years.

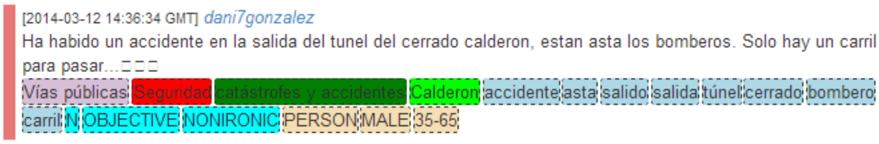

Fig. 2. Example of a tweet annotated by the system

The semantic annotation task is the highest time consuming task and constitutes most of the times the bottleneck of the system. The inquirer processes annotate the unprocessed messages in descending order of insertion time, so that the most recent information is available first to be able to react to early alerts. If the input rate of messages being indexed in the system is higher than the multi-threaded annotation rate, it is still not possible to access the high-level annotations on real time, but once this peak situation is reversed and the system manages to annotate at a higher rate than the indexing of new documents, it will start annotating the rest unprocessed documents.

Finally, the ***visualization component*** is used to exploit the annotated data. Several widgets have been developed to present the data, either just for query and reporting or also for data analytics purposes. These visualization modules can be specifically adapted to better match the city needs.

The datawarehouse and the gatherers are obviously language independent, but the inquirer components are strongly dependent on language lexicons and models. Although the text classification engine is itself language independent, classification models (consisting of training text and rules) are developed for a specific language. The topic extraction engine relies on Part-of-Speech and parsing modules specifically designed to build a sentence syntactic tree in a given language. Moreover, the sentiment analysis engine makes use of that syntactic tree and also depends on a lexicon containing polarity units and modifiers for a given language. The user demographics engine is the only module where no information in a given language is used for creating the model. Our initial business case is deployed to analyze data in Spanish, but modules exist for other languages: English, French, Italian, Portuguese and Catalan.

3 Semantic Annotation

Much effort has been invested in the semantic annotation task, specifically focusing on this scenario, tuned to properly deal with the special singularities of this kind of text snippets (tweets) that usually contain misspellings, emoticons, typographic symbols, letter/number homophones, shortenings, contractions, etc.

The inquirer provides several levels of analyses to classify the text with respect to several (customizable) categories of specifically-defined ontologies, identify topics, perform a demographics analysis to get the user age range, gender and whether he/she is a person or an organization, and sentiment analysis of polarity and subjectivity.

All modules have been exhaustively tested and successfully evaluated in various scenarios, both separately and also integrating two or more modules, in actual systems currently in production, and also in different national and international evaluation workshops such as SEPLN [5], CLEF [6] [7], NTCIR [8] and SemEval [9].

3.1 Text Classification

Another semantic annotation dimension is obtained with an automatic text classification [10] according to pre-established categories defined in a model. The algorithm used [11] [12] combines statistical classification with rule-based filtering, which allows to obtain a high degree of precision for very different environments. Two ontologies were specially designed for this system including concepts and situations that we find relevant to this particular problem; however, the system allows building particular ontologies and classification models for each scenario.

The Social Media ontology defines the general topic classification of the tweet, and contains the first-level categories shown in Figure 3a. The Citizen Sensor ontology, shown in Figure 3b, focuses on features considering the citizen as a sensor.

Category
Art and culture
Police and justice
Disasters and accidents
Economy and treasury
Education
Environment, meteorology and energy
Health
Social issues
Work
Tourism, trips and people movement
Leisure, lifestyle and free time
Politics
Religion and creed
Science and technology
Sport
Disturbances, conflicts and war
Greetings and thankfulness

Category	Subcategory
010000 Locations	010100 Government 010200 Education 010300 Justice 010400 Health service 010500 Culture 010600 Sport 010700 Religion 010800 Commerce 010900 Hotel industry 011000 Outdoor 011100 Means of transportation 011200 Accommodation 011300 Social Institution 011400 Leisure Centre 011500 Workplace
020000 Events	020100 Demonstration 020200 Sport events 020300 Conference and convention 020400 Cultural events 020500 Celebrations
030000 Disasters	030100 Fire 030200 Explosion 030300 Landslide 030400 Avalanche 030500 Flooding 030600 Drought 030700 Earthquake 030800 Seaquake 030900 Epidemic / Plague 031000 Toxic discharge

Category	Subcategory
040000 Accidents	040100 Car crashes 040200 Plane crashes 040300 Maritime accidents 040400 Rail accidents 040500 Nuclear accident
050000 Criminal acts	050100 Robbery 050200 Aggression 050300 Harassment 050400 Rape 050500 Abuse / Mistreat 050600 Kidnapping / Disappearances 050700 Gunfire 050800 Murder 050900 Attempt 051000 Drug trafficking 051100 Intelectual property 051200 Prostitution and Pederasty
060000 Medical emergencies	060100 Pregnancy and childbirth 060200 Decease 060300 Suicide 060400 Drug addiction 060500 Infarction 060600 Asphyxiation 060700 Intoxication 060800 Injury 060900 Burn 061000 Fainting 061100 Epilectic seizure 061200 Electrocution

Category	Subcategory
070000 Weather warnings	070200 Snowfall / Frost 070300 Heatwave 070400 Cold snap 070500 Storm 070600 Hurricane / Tornado 070700 Wind
080000 Incidents	080100 Traffic congestion 080200 Public road damage 080300 Interruption of supplies 080400 Caution at the beach
090000 Concepts	090100 Lighting 090200 Signposting 090300 Air-conditioning 090400 Supplies 090500 Services 090600 Acoustic environment 090700 Odoriferous environment 090800 Environment 090900 Quality of life

Fig. 3. a) Social Media ontology; b) Citizen Sensor ontology (1^{st} and 2^{nd} level categories)

3.2 Topics Extraction

Topics extraction process is carried out by combining a number of complex natural language processing techniques that allow obtaining morphological, syntactic and semantic analyses of a text and using them to identify different types of significant elements. In short, the text is first divided into paragraphs, sentences and tokens, and then each token is lemmatized and tagged with its Part-Of-Speech. A rule-based parser in a series of sequential steps creates the sentence syntactic tree, detecting and tagging the existing coordinated and subordinated clauses, word groups and dependencies among them, and also recognizing named entities and concepts, based on both language resources and also language dependent heuristics (such as `[Mr.|Sir|Dr.]+NAME=>PERSON`). This process also carries out a disambiguation step for the morphosyntactic and semantic information of each token and also anaphora detection and resolution for sentence interlinking.

Currently the system is able to identify (allowing word inflections, variants and synonyms) the following topic categories: named entities (people, organizations, places, etc.), concepts (significant keywords in the text), time expressions, money expressions and URIs.

3.3 Sentiment Analysis

The system also includes functionality to perform a detailed multilingual sentiment analysis of texts from different sources. The text provided is analyzed to determine if it expresses a positive/negative/neutral sentiment polarity. First [6], the local polarity of the different sentences in the text is identified and the relationship among them is evaluated, resulting in a global polarity value for the whole text. Besides polarity at sentence and global level, natural language processing techniques also detect the polarity associated to both entities and concepts in the text (aspect-based polarity).

Moreover, although perhaps not very useful in this context, the sentiment analysis module can also detect if the text processed is subjective or objective and if it contains irony marks, both at global and sentence level, giving the user additional information about the reliability of the polarity obtained from the sentiment analysis.

3.4 User Demographics

The user demographics analysis module extracts some important demographics (type, gender, age) for a given Twitter user. State-of-the-art information extraction and text classification algorithms are used to guess those facts from his/her login, name and profile description, based on n-grams model, developed using Weka [13].

4 Visualization

The visualization is a web interface that allows to easily building complex queries in a structured way, enabling, thus, a versatile filtering of the data and high level visualization with the aim to provide the final user with a highly aggregated and condensed information at a first sight. The system is designed to provide both real time analysis and backtracking of previously stored data.

The system console is created defining several elements called widgets, in such a way that the template may be changed between different user cases (different cities and their particular needs) to adapt the system to each community.

Some of the components make use of the Highcharts JavaScript library [14] to create intuitive and interactive charts, OpenLayers [15] to display maps and geoposition information, as well as self-customized components. The user interface makes use of the capabilities of Elasticsearch, allowing the user to create their own queries by filtering on the semantic tags and aggregating information using its Facets API.

An example of an analysis dashboard using some of the built widgets is shown in next figures. Figure 4 shows filter capabilities, analysis of total number of tweets and alerts as well as last minute tracked tweets and alerts, a timeline tracking the number of tweets and alerts per minute, as well as the number of positive and negative ones.

Fig. 4. Dashboard with filters, statistics and timelines

Figure 5 presents several pie charts with user statistics (number of users by age range and gender), global sentiment polarity, and a list of the most frequent alerts, locations and events.

Fig. 5. Dashboard with user demographics, sentiment polarity, alerts and events

The console also displays a map with the locations of the alerts that contained this information and also includes the semantically annotated tweets that match the filtering criteria (Figure 6).

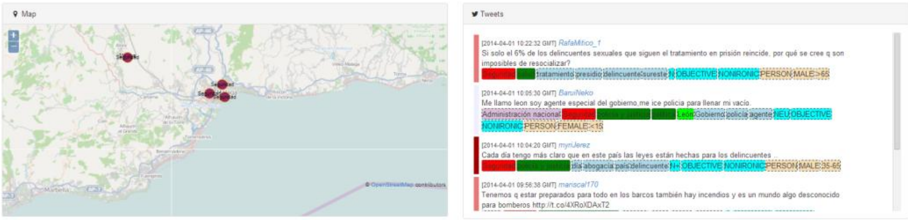

Fig. 6. Dashboard with tweets and map

Last, Figure 7 shows some widgets with tag clouds listing the most relevant (by number of appearance) topics, entities, concepts and hashtags.

Fig. 7. Dashboard with tag clouds of topics, entities, concepts and hashtags

5 Discussion and Future Work

In this work a real time semantic annotation engine for Twitter data with datawarehouse capabilities and a search engine for backtracking and later data analytics has been described. The system allows community promoters to more quickly react to specific events that may happen (catastrophes, accidents, traffic congestion, etc.), react to people feelings and detect which initiatives are more likely to be improving quality of life for their citizens, to detect the topics that are worrying the citizens... Thus, it will increase the degree of engagement of the smart cities that use the system with their citizens.

Currently the system in beta-testing process, adapting the interfaces, fine-tuning the different modules and removing noise in the annotations. The system will be deployed in different scenarios in a short or medium term. There are several business cases under negotiation. The first scenario is to build a city console for a local administration to be able to analyze in real-time the behavior and topics of interest of the citizens, with two components: a private console, internal for the city services, and a public console, a dashboard with attractive, summarized, non-confidential information to be projected or displayed at selected public locations of the city (town hall, libraries, museums) or even in a LED video wall in a populous square in downtown, to engage citizens with these technologies and also promotion. The second scenario is to focus on emergencies services, providing early detection of security-related issues.

Regarding the technology, the storage capabilities of the system allow not only to analyze real time data, giving a snapshot of the current city state, but also to apply data mining algorithms to the stored data in order to better understand particularities of the population, clustering and profiling of the different groups that form the city environment, compare the singularities of the different detected clusters, etc. Currently, steps to further explore this path are being taken: city mobility analysis (how, when, why people move from one place to another), relevant topics analyzed at neighbourhood level, city reputation and brand personality, etc.

Finally, the same approach that has being used analyzing Twitter data will be used with other sources of information. The gatherer will be extended to capture data from other sources like other social network like Facebook, LinkedIn (in smart-city related groups), Tuenti, or other social sites such as YouTube, Flickr, Pinterest, etc. In addition, we want to better adapt our core models for NLP to the special features that Social Networks language introduce.

Acknowledgements. This work has been supported by several Spanish R&D projects: Ciudad2020: Hacia un nuevo modelo de ciudad inteligente sostenible (INNPRONTA IPT-20111006), MA2VICMR: Improving the access, analysis and visibility of the multilingual and multimedia information in web for the Region of Madrid (S2009/TIC-1542) and MULTIMEDICA: Multilingual Information Extraction in Health domain and application to scientific and informative documents (TIN2010-20644-C03-01).

References

1. Ciudad 2020 - Hacia un nuevo modelo de ciudad inteligente sostenible. Website. `http://innprontaciudad2020.es`.
2. Elasticsearch.org. Open Source Distributed Real Time Search & Analytics. `http://www.elasticsearch.org`
3. Twitter REST API v1.1. `https://dev.twitter.com/docs/api/1.1`
4. Textalytics API. `http://textalytics.com`
5. Díaz Esteban, A., I. Alegría, and J. Villena-Román (eds). Proceedings of the TASS workshop at SEPLN 2013. *Actas del XXIX Congreso de la Sociedad Española de Procesamiento de Lenguaje Natural*. IV Congreso Español de Informática. 17-20 September 2013, Madrid, Spain.
6. Villena-Román, J., S. Lana-Serrano, C. Moreno-García, J. García-Morera, and J.C. González-Cristóbal. 2012. DAEDALUS at RepLab 2012: Polarity Classification and Filtering on Twitter Data. *CLEF 2012 Labs and Workshop Notebook Papers*, Rome, Italy, September 2012.
7. Villena-Román, J., and S. Lana-Serrano. MIRACLE at VideoCLEF 2008: Topic Identification and Keyframe Extraction in Dual Language Videos. 2009. *Evaluating Systems for Multilingual and Multimodal Information Access. 9th Workshop of the Cross-Language*

Evaluation Forum, CLEF 2008, Aarhus, Denmark, September 17-19, 2008, Revised Selected Papers. Carol Peters et al. (Eds.). Lecture Notes in Computer Science, Vol. 5706, 2009.

8. Villena-Román, J., S. Lana-Serrano, and J.C. González-Cristóbal. 2008. *MIRACLE at NTCIR-7 MOAT: First Experiments on Multilingual Opinion Analysis*. 7th NTCIR Workshop Meeting. Evaluation of Information Access Technologies: Information Retrieval, Question Answering and Cross-Lingual Information Access. Tokio, Japón, December 2008.
9. Villena-Román, J., J. García-Morera, and J.C. González-Cristóbal. 2014. Daedalus at SemEval-2014 Task 9: Comparing Approaches for Sentiment Analysis in Twitter. *Proceedings of the 8th International Workshop on Semantic Evaluation*, SemEval'14, Dublin, Ireland, 2014 (to be published).
10. Sebastiani, F. 2002. Machine learning in automated text categorization. *ACM Computing Surveys*, 34(1), pp 1–47.
11. Villena-Román, J., S. Collada-Pérez, S. Lana-Serrano, and J.C. González-Cristóbal. 2011. Método híbrido para categorización de texto basado en aprendizaje y reglas. *Procesamiento del Lenguaje Natural*, Vol. 46, 2011, pp. 35-42.
12. Villena-Román, J., S. Collada-Pérez, S. Lana-Serrano, and J.C. González-Cristóbal. 2011. Hybrid Approach Combining Machine Learning and a Rule-Based Expert System for Text Categorization. *Proceedings of the 24th International Florida Artificial Intelligence Research Society Conference (FLAIRS-11)*, May 18-20, 2011, Palm Beach, Florida, USA. AAAI Press 2011.
13. Hall, M., E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I.H. Witten. 2009. The WEKA Data Mining Software: An Update. *SIGKDD Explorations*, Volume 11, Issue 1.
14. Highcharts - Interactive JavaScript charts for your webpage. JavaScript library website. `http://www.highcharts.com`
15. Openlayers. `http://openlayers.org`

Community Mapping for Participatory Decision-Making Processes*

Liliana Ardissono(1), Maurizio Lucenteforte(1), Adriano Savoca(1), and Angioletta Voghera(2)

(1) Dipartimento di Informatica, Università di Torino, Italy
liliana.ardissono@unito.it, maurizio.lucenteforte@unito.it,
WWW home page: www.di.unito.it/˜liliana,
http://di.unito.it/lucente
(2) DIST, Politecnico di Torino, Italy
angioletta.voghera@polito.it,
WWW home page:
www.dist.polito.it/personale/scheda/(nominativo)/angioletta.voghera

Abstract. Community mapping is being increasingly used to support crowdsourcing in participatory decision-making processes but the collection of feedback is usually carried out in textual form. We describe a proposal for the management of 3D Community Maps which provide a virtual representation of the territory and enable users to contribute to policy making by sharing different types of contributions, such as comments, documents and 3D models. The paper discusses user requirements and personalization aspects in this type of application.

Keywords: Participatory GIS, 3D User Interfaces, Community Maps, Citizen-centered service design and modeling.

1 Introduction

Various crowdsourcing platforms support the management of participatory decision-making processes by enabling the population to provide feedback and proposals about public policies. For this purpose, most of them enable the collection of textual feedback using standard bi-dimensional maps; e.g., see CrowdMap (https://crowdmap.com/). In a few projects a virtual representation of the territory is offered to enable the visualization of possible scenarios and anticipated effects of the planned actions but the virtual environments cannot be modified by their visitors.

Enabling graphical feedback is important to enrich the communication capabilities within a discussion group as well as to extend the kind of contributions which can be provided by people. For instance, consider the design of a cycling path in a town: several revision proposals (e.g., to correct dangerous curves) could be easily represented by means of a sketch in the town map, possibly enriched with textual explanations of

* This work was partially supported by the University of Torino, grant "Supporto intelligente e flessibile all'esecuzione di attività collaborative complesse". We thank Giuseppe Scaramuzzino for his work in the first version of GroupCollaborate2.

the reasons for proposing the changes. However, in order to do that, intuitive tools are needed which can be used by heterogeneous users; e.g., domain experts vs. generic users, or people with different levels of familiarity with computer science.

As a first step in this direction, we developed GroupCollaborate2, a prototype participatory GIS which enables the on-line sharing and editing of geo-localized documents and 3D models in public and private focus groups. The system enables users to share content and collaboratively edit it, as well as to discuss in order to collectively design proposals for the re-development of a territory.

GroupCollaborate2 does not currently offer advanced adaptation features but it enables users to access a personal view of the shared information space by offering customizable filters on the data items to be presented. In this way, it enables people to focus on the portions of the information space they deem to be relevant.

In the rest of this paper Section 2 introduces participatory processes and Participatory GIS. Section 3 discusses user requirements and Section 4 presents our work. Sections 5 and 6 position it in the related research and close the paper.

2 Background

The consolidated practices developed at a national or European level in the participation field are communication, animation, consultation and empowerment [1]:

- Communication is an information activity about public choices.
- Animation has two goals: (i) informing a large social context about the state of the art of a decision process or about the implementation of a program; (ii) favoring the birth of virtuous behaviors among citizens.
- Consultation involves the community in territorial transformation processes for requirement elicitation purposes.
- Empowerment can be interpreted as the enhancement of individual/group abilities. It represents the citizens participation by promoting the auto-organization abilities of people in the social, economic, cultural and territorial fields of action.

Participatory processes are related to empowerment: they are based on a bottom-up decision-making model which promotes the contribution of the population (involving representatives of stakeholders) to public policy development by expressing needs, proposals and feedback with the aim of raising the Public Administration's awareness of the priorities to be addressed and of reaching consensus on the actions to be carried out.

One of the challenges to be addressed in processes related to territorial policies is how to represent space and its inherent relations in a way that is both informative and involving for lay people. As most urban studies data are found in map forms, visualization capacity (by employing mapping services) and the capacity to describe and represent values that people attach to places [2] are critical. Participatory GIS, which support various forms of community participation, are thus emerging as promising tools to overcome traditional barriers to public involvement in decision making processes and policy making with a spatial dimension [3, 4].

Participatory GIS also offer great opportunities to enhance traditional forms of community planning, such as Community Maps drawing. This is a way to represent people's view of a certain area by gathering and presenting site-specific data, to understand differences in perception and to identify which values people attach to places or elements of their living space. Recent experiences demonstrate that using GIS and other ICT tools can significantly add value to community mapping, allowing to uncover individual and collective neighborhood definitions [5] and to highlight local issues, planning priorities and needs, or identify development sites. In this way communities add information to the map themselves and act as sensors in their local environment [6].

3 Crowdsourcing Support in a 3D Environment

3.1 User Requirements

Even though our work concerns supporting different phases of participatory decision-making (e.g., see [7]), the current paper focuses on crowdsourcing. Domain experts (mainly urbanists) helped us to identify relevant requirements for the design of a Participatory GIS and they stressed the importance of centering the interaction model around the concept of Community Map, which provides an immediate and intuitive representation of indigenous spatial knowledge and needs.

The idea is that of using the community map both as a shared information source and as communication mean which enables group members to interact with each others and to access the shared information items. Thus, the members of a group should have access to a dedicated map for entering or searching for comments, documents and proposals (including drawings and 3D models) relative to the associated geographical area. In this way, the community map becomes a shared, dynamic information source which all stakeholders can use to overview and discuss the existing proposals as well as to collaboratively revise them. A main requirement is that the map provides 3D simulations of the environment and of the planned changes in order to offer a realistic representation of the intended effects of the proposed actions. Furthermore, the map should be the only reference for accessing proposals and feedback, thus supporting an integrated access to information. Obviously, other general requirements have to be considered. For instance, the availability of territorial data for the management of the community maps, the diffusion of the enabling technology for using the Participatory GIS (which should be accessible in internet from standard computers and browsers) and the provision of an intuitive user interface which supports a smooth interaction with the system without requiring technical skills.

We designed the user interface and the functions offered by GroupCollaborate2 by taking the above listed requirements into account; see the next section. As far as the basic requirements are concerned, we designed the user interface of our system in a user-centered way, directly involving domain experts and generic users who are not familiar with Computer Science. Moreover, for the representation of the community maps we decided to exploit Google Maps (and Google Earth), whose GIS layers describe geographical areas with reasonable precision and which can be easily enriched with layers displaying additional geo-referenced data. Google Maps are also largely accessible because they only require the downloading of the Earth plug-in in the user's browser to navigate the virtual environment and edit it. Furthermore, we based the authentication in the collaboration support platform on Google accounts, which are pretty diffused.

3.2 Sample Scenario

Let's suppose that, in the context of its future transformation plans, the Public Administration (PA) of Bruino city (IT) starts a participatory process to collect the population's needs and priorities and to solicit proposals regarding a peripheral area to be re-developed. In order to collect suggestions, PA opens a discussion using a crowdsourcing platform such as GroupCollaborate2 and invites citizens and stakeholders (e.g., local associations) to join a focus group devoted it. The interaction with the platform is based on a Web-based Community Map of the town focused on the area under discussion, which enables users to (i) retrieve detailed information about the existing infrastructures, such as the presence of a road; (ii) upload comments, pictures and drawings sketching users proposals; (iii) edit or annotate user-generated content (e.g., project plans) to provide suggestions and comments; (iv) participate in geo-referenced discussion threads.

PA monitors the discussion space and contributes with opinions, suggestions and evaluations. A set of hypotheses for setting up or updating spaces are developed in collaboration with the population and are stored as a basis for a subsequent evaluation phase, aimed at selecting the most promising solutions through deliberation.

4 Our System

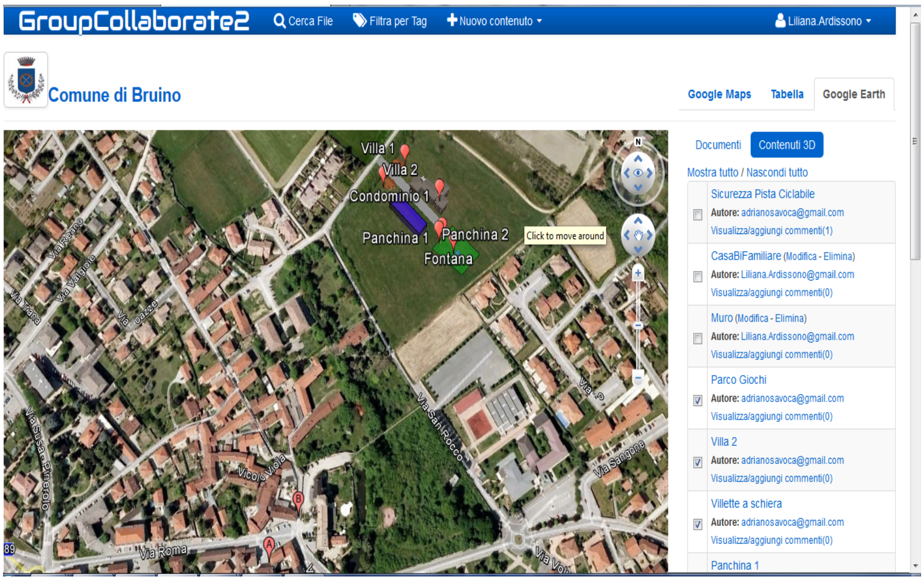

Fig. 1. Community map displaying geo-referenced documents (markers labeled with letters in the bottom part of the map) and 3D models (located in the upper part of the map).

GroupCollaborate2 contributes to the management of focus groups such as the one described in our sample scenario by supporting discussion and document sharing in public and private groups. The system offers a user interface page devoted to the management of focus groups (users subscription, etc.). Moreover, it supports e-mail communication.

For crowdsourcing purposes each group is associated with a map representing the entry point to the shared information items. The map can be visualized as bi-dimensional or three-dimensional and is populated with the shared content. Figure 1 (in Italian), described below, shows the layout of the 3D community map for the Bruino group.

4.1 Information Visualization and Sharing Support

The community map enables users to share and collaboratively edit (possibly) geo-referenced objects of various types, including documents and drawings: 3D models are visualized as shapes and documents are identified by markers; e.g., see A and B in the central area at the bottom of the map in Figure 1. Moreover, each marker/3D model can be clicked to view its metadata (author, title and description; e.g., see the tooltip of "Villa 1" in Figure 2) or to open the associated document for reading/editing purposes.

The "Nuovo contenuto" (new content) link at the top of the page can be used to (i) create or upload a new document; (ii) upload a 3D model from a repository (e.g., a KMZ model), or (iii) draft a new 3D item by means of an editor which allows to sketch broken lines and polygons, select color and height of items, move, orientate and resize them, and set metadata. The right portion of the page displays items in a checkable list which allows the user to further select the elements to be shown in the map and to interact with them: e.g., open/edit documents or zoom the map on 3D objects (through the title), view the list of associated comments or add new ones ("Visualizza / aggiungi commenti"). The list also includes documents which cannot be visualized in the map because they are not geo-referenced.

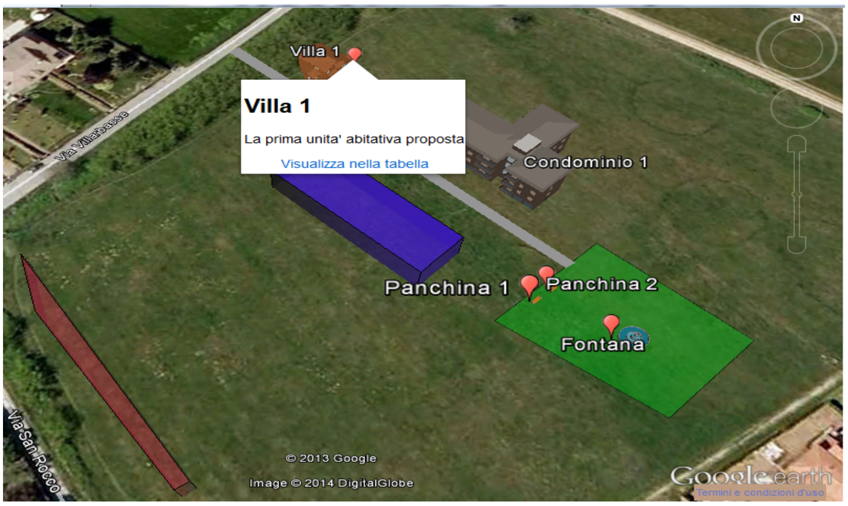

Fig. 2. Zoom on a portion the community map displaying 3D models.

Figure 2 focuses on the 3D items uploaded in the map of Figure 1: “Villa 1”, “Villa 2” (hidden by the tooltip of “Villa 1”) and “Condominio 1” have been uploaded as 3D models; a red polygon representing a wall and a blue polygon representing a building have been drafted using the editor. The map also includes a green area drafted to delimit a playground with benches (“Panchina 1” and “Panchina 2”) and fountain (“Fontana”).

4.2 Information Search and Filtering Support

GroupCollaborate supports hierarchical tag-based item classification and filtering of documents and 3D objects:

- We introduced tag categories to manage different points of view on information: e.g., within a group, it can be useful to distinguish different types of content such as the masterplan of a project or general documentation. Tag categories can be thus defined to support different perspectives in the search for information. Moreover, for each tag category, a set of tags can be introduced to classify specific items in a folksonomy. All group members can define tag categories and individual tags by means of a dialog box supporting the hierarchical exploration of the tag system.
- The “Filtra per Tag” (filter by tag) link at the top of the map (see Figure 1) supports item search: it enables the user to choose a set of tags to be jointly used for selecting shared documents and 3D models. The search results are shown in the map, which visualizes the information satisfying the current filtering criterion.

The system also enables users to search documents by content (i.e., by words included in the documents) and by document name (see “Cerca File” - search file - link at the top of the community map). All these functions support the dynamic generation of community maps reflecting particular viewpoints; moreover, they allow reducing the amount of displayed information depending on users’ interests.

4.3 User Feedback on the System and Personalization Issues

We carried out a preliminary test of the system with a few domain experts and generic users who are not familiar with 3D environments and/or Computer Science but who are internet users, as this is the target expected to use a Web-based Participatory GIS. Users appreciated the functions offered by the system as they proved to effectively support both the introduction of elements to be shared in a focus group and the discussion activity on such elements, based on comments. The tag-based filter was considered very important to reduce the amount of information displayed in the community map, providing particular viewpoints on data.

Domain experts stressed the need to moderate user-generated content (something which is possible as group administrators). Moreover, they suggested to introduce new functions, such as:

- The introduction of subgroups to support the activation of specific discussion lines; e.g., among domain experts only. Subgroup modeling has two facets: (i) the management of restricted information access; (ii) the possibility that a user has multiple associated personae, and thus user profiles, in order to support her/his operations in different contexts.

– The management of user privacy in order to add anonymous contributions. This is particularly important if the system is extended to support voting, in which case also a trusted user authentication is mandatory; see [8].
– The introduction of additional communication channels, e.g., to enable the sharing of multimodal content such as voice comments and videos in the community map.
– The possibility for each user to access personal views on content based on concept selection; e.g., only scholastic buildings, or sport and leisure facilities. This aspect opens research paths on data representation (to classify content by concepts), user modeling (to understand the user's interests) and manual/automatic maps adaptation to derive personal views focused on specific interests.

From the usability viewpoint GroupCollaborate2 should comply with universal accessibility guidelines to support users having different abilities and using different browsers. In the current version the user interface is developed in HTML5 and has a neat layout to address the basic W3C accessibility guidelines. It is however clear that various features of the user interface could be extended to support different types of interaction. For instance, a library of pre-defined shapes and 3D models could be offered for dragging and dropping them in the map. Moreover, sophisticated tools might be proposed to draft complex polygons to represent buildings and architectonic elements having non-trivial shapes; e.g., see [9]. However it must be assessed whether technical users, who might want to produce such shapes, prefer to use specialized tools for this purpose.

5 Related Work

Social networks and Web 2.0 technologies are very used for supporting participation but they typically only collect textual feedback from users. For instance, urbanAPI (http://www.urbanapi.eu/) enhances participatory urban planning by coupling interactive information visualization for the presentation of project plans and of policy simulation results with the management of polls to elicit feedback about territorial policies. Similarly, WE-GOV (www.wegov-project.eu/), NOMAD (http://www.nomad-project.eu/) and FUPOL (http://www.fupol.eu/) deal with both communication and crowdsourcing but they collect people's feedback using textual interaction.

A few Participatory GIS projects enrich communication with 3D information in Virtual Reality (VR) or Augmented Reality (AR) environments describing proposed redevelopment scenarios. For instance, LIVE+GOV (http://liveandgov.eu/) combines AR and VR with social networks to allow internet users upload and receive geo-localized information about buildings and locations in a city, as well as participate in polls and discussions. Moreover, Min Stad (http://minstad.goteborg.se/minstad/index.do - a portal for the City of Goteborg), and partially PlanYourPlace [8], integrate GIS with social networks enabling users to upload 3D contents and to publish comments.

With respect to such works, GroupCollaborate2 lacks the support to deliberation provided by polls, which can be integrated with limited effort. Moreover, it has no direct connection to external social networks because it directly manages communication, integrating it with a richer type of crowdsourcing where users can share and collaboratively edit heterogeneous types of contents. Furthermore, our system supports thematic discussion groups, with consequent information hiding, w.r.t. to a generic upload and visualization of information, the same for everybody. Finally, it supports the selection of the information to be visualized in the Community Maps through content-based and tag-based filtering, thus enabling their projection on different dimensions reflecting individual interests.

6 Conclusions

This paper presented GroupCollaborate2, a prototype Participatory GIS which enables the on-line management of discussion groups for participatory decision making. The described system enables users to share various types of digital content, including 3D objects, and to discuss it by interacting with a bi/three-dimensional Community Map supporting tag-based and content-based information search. The interaction with domain experts (planners) and generic users highlighted avenues for the adoption of personalization strategies, aimed at improving the interaction with the system and the access to information. For technical details about GroupCollaborate2 see [10].

References

1. Ciaffi, D., Mela, A.: La partecipazione. Dimensioni, spazi, strumenti. Carrocci, Roma (2006)
2. Brown, G., Raymond, C.: The relationship between place attachment and landscape values: Toward mapping place attachment. Applied Geography **27**(2) (2007) 89–111
3. Bugs, G.: Assessment of online PPGIS study cases in urban planning. In: ICCSA'12 Proc. of 12th Int. Conf. on Computational Science and Its Applications,LNCS n. 7333. Springer (2012) 477–490
4. Brown, G., Weber, D.: Measuring change in place values using public participation GIS (PPGIS). Applied Geography **34** (2012) 316–324
5. Coulton, C., Chan, T., Mikelbank, K.: Finding place in community change initiatives: Using GIS to uncover resident perceptions of their neighborhoods. Journal of Community Practice **19**(1) (2011) 10–28
6. Ellul, C., Francis, L., Haklay, M.: Engaging with local communities: A review of three years of community mapping. In: Urban and Regional Data Management, UDMS Annual 2011 - Proc. of Urban Data Management Society Symp. 2011, Delft, Netherland (2012) 165–177
7. Ardissono, L., Voghera, A.: Participatory decision making: a step towards e-planning for smart cities. In: Planning Support Tools: Policy Analysis, Implementation and Evaluation. Franco Angeli (2012) 1338–1349
8. Hunter, A., Steiniger, S., Sandalack, B., Liang, S., Kattan, L., Shalaby, A., Alaniz Uribe, F., Bliss-Taylor, C., Martinson, R.: Planyourplace a geospatial infrastructure for sustainable community planning. Int. J. of Geomatics and Spatial Analysis **22**(2) (2012) 223–253
9. Isikdag, U., Zlatanova, S.: Interactive modelling of buildings in google earth: A 3d tool for urban planning. In: Developments in 3D Geo-Information Sciences Lecture Notes in Geoinformation and Cartography, Berlin Heidelberg New York (2000) 52–70
10. Ardissono, L., Lucenteforte, M., Savoca, A., Voghera, A.: GroupCollaborate2: Interactive community mapping. In: Proceedings of UMAP 2014 - Posters and Demos, Aalborg, Denmark, CEUR (2014) to appear

Personalization of Parliamentary Document Retrieval using different User Profiles

Eduardo Vicente-López, Luis M. de Campos, Juan M. Fernández-Luna, and Juan F. Huete

Departamento de Ciencias de la Computación e Inteligencia Artificial, E.T.S.I.I.T., CITIC-UGR, Universidad de Granada, 18071-Granada, Spain
{evicente,lci,jmfluna,jhg}@decsai.ugr.es

Abstract. Owing to the information overload we are faced with nowadays, personalization approaches are becoming almost a must, in order to provide relevant information for users. These personalization techniques retrieve results closer to the user interests and preferences, by using the information stored in the user profile. We have carried out a comparative study between six different user profile representation approaches, based on the content of the documents of the Andalusian Parliament, obtaining quite good personalized performance results and some interesting conclusions about the goodnesses of these content-based approaches.
Keywords. User Profiles, Personalization, Information Retrieval, Parliamentary Documents

1 Introduction

Over the last few years, the amount of digital information is rising exponentially [10], so its access is everyday more difficult. The use of Information Retrieval Systems (IRS) has become essential to find relevant information within this huge bunch of data. The use of such systems in e-Government will help to deliver the needed information to citizens, representing a particular, and important, application of IR techniques.

This paper is framed within our collaboration with the regional Parliament of Andalusia (Spain). Particularly, we have built an IRS [6] to enhance the access of the citizens to the Records of Parliamentary Proceedings, called *Seda* (`http://irutai2.ugr.es/SEDA`), taking the most of the internal structure of such documents and founding our search engine on XML retrieval. Among the proceedings, the textual transcriptions of the working Committees can be found, considering policy issues, conducting inquiries and producing reports on a range of matters. Each Committee is devoted to a wider topic of interest as agriculture, education or economy (the number of committees and the covered topics varies between terms of office – for example, nowadays there are 11 different committees). Each of these records (or documents) contains an average of 5.6 initiatives, which present a detailed discussion of the members of the Parliament about a specific issue. In turn, each of these initiatives is tagged with one or more subjects extracted from the EUROVOC thesaurus[1], being manually assigned by parliamentary documentalists as the best representation of its content.

The citizens can search for a piece of information by submitting a query to the system. Although in the last years our system has been providing quite good results, it offers the same output for a given query, independently of the user, since it only considers the query keywords as the representation of the user information needs. This issue is well-known as the *'one size fits all'* paradigm.

If we join the continuous increase of data, with the users tendency to formulate short and ambiguous queries [19], a new approach is required, in which the user context, and not only the query, is considered as an important part within the retrieval process. Personalization [2, 3, 22] is this possible solution, and hot arising research area, whose main objective is to retrieve results closer to the users, in order to better satisfy their specific information needs.

Any personalization process has three main different stages: 1) to acquire and represent the user interests and preferences in the *user profile*, 2) to exploit the best as possible the user profile information within the retrieval process, and 3) to evaluate the whole personalization process. We may consider some additional issues, such as privacy in the personal data collection and management process [13], or different ways to present the personalized results [1], with the intention of presenting this information to the user in the most easy and intuitive way.

It seems quite obvious that the personalization process expected performance depends on the quality of the user profile information. For this reason, in this article we have focused on the analysis of different ways to build user profiles. Concretely, we shall focus on content-based user profiles which are frequently used in *contextual evaluation* environments, such as [18]. Additionally, this kind of profiles could be ideal for the introduction of personalization in privacy constrained environments. We have concretely faced this problem with the Andalusian Parliament, where the members of the Parliament do not allow any personal data recollection of themselves nor the citizens. In this way, the parliamentary IRS could integrate personalization techniques to improve its retrieval performance and user satisfaction, only giving the user the possibility to choose with which of the simulated profiles he/she is more alike.

To build these user profiles, we have to analyse the content of the Records. We take the advantage of a pre-classified collection, i.e. each document belongs to one committee representing a different area of interest or category, in which the future users could be interested in. As a first approach, we develop a user profile based on the EUROVOC thesaurus subjects, manually assigned by the Andalusian Parliament documentalists to each initiative discussed in a committee session. Secondly, we build a user profile only based on terms (independently of where they appear in the document), and thirdly we have configured a hybrid user profile composed of both terms and subjects. Finally, we evaluate the use of each alternative in order to find the most appropriate in terms of retrieval performance.

[1] http://eurovoc.europa.eu/

While these content-based simulated user profiles could be considered as lacking *'reality'*, since they do not represent real users, they are a valid approach [9, 18] for possible users interested in some areas of interest[2]. If we join the recent rise of personalized systems, together with the fact that their evaluation through user studies is rather complicated (due to the large required resources, such as, access to real users, time, money or even the needed infrastructure for their implementation), we consider particularly important to test and improve the quality of content-based user profiles.

The rest of the article is organized as follows: Section 2 presents a review of the state-of-the-art in user profiles. Section 3 shows how the profiles are built, based on subjects, terms and a combination of both of them at the same time. How these profiles are used in conjunction with the user query, the experimental design and the evaluation is described in Section 4. Finally, the last section of the paper shows the conclusions and proposals for further research.

2 User Profiles Literature

The quality of personalized results will highly depend on the user profile quality and how it is exploited in the retrieval process. Hence, the user profile building process is one of the most important steps to obtain good personalized results, but at the same time very difficult, since user interests and preferences are difficult to be captured and they also change over time [14, 17].

The three most important steps in the user modeling-user profile building process, according to [11] are the following : 1) acquisition of user information, 2) user profile representation, and 3) user profile update. We will focus on the second step, since this article main goal is to make a comparative study between different user profile representations performance.

The three main user profile representation approaches are: a set of weighted keywords, semantic networks, and a set of weighted concepts:

a) Weighted keywords: it is the most common user profile representation. They may be automatically learned from user visited documents or directly given by the user. The keyword weights show the importance of each keyword within the profile. Examples of this approach are [20], where they build three different user profiles based on relevance feedback and implicit information, user browsing history, and a modified collaborative filtering. Other examples are [16], where they learn the user profiles from the user visited web pages, based on the well-know *tf*idf* approach, and [8], where they build user profiles formed by a vector of keywords for each user area of interest.

b) Semantic networks: in order to handle the keyword user profile polysemy problem, a weighted semantic network is included, in which each node represents

[2] In this situation the user might also opt for the inclusion of several terms in the query describing the committee content, terms that could be difficult to select for a citizen, appearing also a query drift problem. Or, otherwise, opting for filtering out the documents which do not belong to the committee, but in this case there might be relevant results which are not shown to the user (around 25% in our studies).

Table 1. Examples of the three proposed user profiles, for the 'agriculture and livestock' area of interest (unstemmed and translated into English).

sProf	$s = \{$ *0.216*"agriculture aid" 0.127*"agricultural policy" 0.098*"agricultural production" 0.098*"oily" 0.095*"food industry" 0.091*"fishing" 0.083*"oil" 0.075*"huelva province" ...* $\}$
tProf	$t = \{$ `0.007*agriculture 0.007*sector 0.004*fishing 0.004*agrarian 0.004*production 0.003*aid 0.003*farmer 0.002*product 0.002*rural 0.001*oil ...` $\}$
stProf	$s_1 =$ *0.216*"agriculture aid"* $t_{s1} = \{$`0.007*aid 0.006*sector 0.006*agriculture 0.005*farmer ...`$\}$ $s_2 =$ *0.127*"agricultural policy"* $t_{s2} = \{$`0.009*agriculture 0.007*agrarian 0.006*production ...` $\}$ ⋮ ⋮

a concept. For example, in [15] a filtering interface is created to personalize the results from the Altavista search engine. Another semantic network example is [18], where a personalized search system with ontology based user profiles is presented. These user profiles are built assigning scores to user interests, implicitly derived from concepts of the ODP ontology. Since the user interests are dynamic, a propagation algorithm is used to keep these interests updated.

c) Weighted concepts: they are similar to the semantic networks, since they also have conceptual nodes and relations between them, but in this case, the nodes are represented by abstract topics of interest for the user instead of terms. But, at the same time, they are also similar to the weighted keyword user profiles, since they are usually represented as vectors of weighted concepts. Nonetheless, in the last few years is common to use a hierarchical representation of concepts, usually derived from a taxonomy, thesaurus, or a reference ontology, instead of using concepts with no structure, allowing a much richer representation. An example of this approach is [21], where using concepts from the ODP ontology first three levels, they build user profiles based on the user browsing history. Another example is [4], where they show three different ways to use ODP: first, as a semantic support to find relations between concepts; second, identifying some ODP structure parts relevant to the user; and third, the user directly choose the ODP concepts he/she is interested in. After that, they study how to exploit these three user profiles, with personalization techniques based on query modification and re-ranking.

3 User Profiles Building Process

Due to the frequent important restriction concerning collecting user personal information, and additionally to the difficulty to have accurate and updated user profiles, we have decided to build simulated user profiles based on content. Particularly we focus on the information available in the transcriptions of the working Committees, where much of the work of the Parliament takes place. Thus, assuming that those topics in a given committee might represent the interests (preferences) of the citizens, we analyse its content to learn the profile. In this paper we will explore three different types of user profiles, see Table 1 for an example:

- *sProf*: This first approach, based on the initiative subjects, can be considered as a *weighted concept* profile, since these subjects represent abstract topics of interest for the user but not terms. They are represented as vectors of weighted concepts, without any structure. Concepts profiles main assets are their robustness to vocabulary variations and a less requirement of user feedback. These characteristics and the fact that the subjects are manually selected by experts in the document collection, as the best content representation for the parliamentary initiatives, were the reasons which made us to start with this approach to learn the user profiles.
- *tProf*: The second profile approach, based on the collection terms, can be considered as a *weighted keyword* profile, since the terms themselves are the items which represent the user interests. These profiles are the easiest to build, but they need to have many terms to accurately define a user interest. These profiles are also less understandable for users than those based on concepts, since their interests are much easily mapped with concepts than with isolated terms. But at the same time, the terms let a more fine-grained representation of the collection content.
- *stProf*: The third profile approach, based on subjects and terms, is an hybrid approach among the weighted concept and weighted keyword profiles, keeping concept abstraction but enriched by the terms fine-grained contribution. To build this profile we learn the most representative terms for each collection subject. Thus, this new profile now contains two levels: the first, with the subjects which represent the profile, and the second formed by the terms which represent each first level subject.

We now show the way we select the elements of each type of profile. Let X represent either a subject in the case of *sProf* or a term in the case of *tProf*, and let Y represent a profile. Then we define $f^{+}(X,Y)$ as the frequency of X in documents belonging to any area(s) of interest which form the profile Y; $f^{+}(Y)$ is the number of elements (either subjects for *sProf* or terms for *tProf*) within Y; $f^{-}(X,Y)$ and $f^{-}(Y)$ are respectively the frequency of X and the number of elements in documents outside the profile Y. For the *stProf* profiles, X represents a term and Y represents a subject, $f^{+}(X,Y)$ being in this case the frequency of X within initiatives classified by the subject Y and $f^{+}(Y)$ the total number of terms within these initiatives; $f^{-}(X,Y)$ and $f^{-}(Y)$ have in this case the obvious meaning. We then define the relevance of X with respect to Y, $R(X,Y)$ as

$$R(X,Y) = \frac{f^{+}(X,Y)}{f^{+}(Y)} - \frac{f^{-}(X,Y)}{f^{-}(Y)}$$

i.e., the normalized frequency of X within Y minus the normalized frequency of X outside Y. If the final value is $R(X,Y) \leq 0$, it means that X is more frequent outside than within Y, so it is not representative of Y and we do not consider it. However, if the final value is $R(X,Y) > 0$, this means that X represents Y at a certain degree, so we keep it. All the retained elements are sorted in decreasing order of relevance to form the profile. In the case of the *stProf* profile we first calculate the list of subjects and next the list of terms associated to each subject.

Table 2. Final *sProf* and *tProf* user profiles using $exp[Subj|Terms] = 5$ and $maxNorm = 0.66$.

sProf	*0.66*"agriculture aid" 0.388*"agricultural policy" 0.299*"agricultural production" 0.299*"oily" 0.290*"food industry"*
tProf	`0.66*agriculture 0.647*sector 0.401*fishing 0.399*agrarian 0.398*production`

4 Evaluation Framework and Results

This section shows the components of the used evaluation framework, how we have used the previous user profiles, and the obtained results and conclusions.

The evaluation framework is composed by the following components: a document collection formed by 658 Committee Sessions from the sixth and seventh Andalusian Parliament terms of office, marked up in XML (containing 432,575 retrievable structural units); an heterogeneous set of 23 queries formulated by real users of the document collection; the search engine is Garnata [5]; the relevance assessments were obtained from a carried out user study, which involved 31 users, with a total of 126 evaluation triplets (user, query, profile), i.e., the relevance assessments provided by a given user, evaluating a given query under a given profile (considering each user chose the user profile closer to his/her interests - none of the user profile representations discussed in this article was provided to the user, but a brief general description of its expected content); the NDCG evaluation metric (Normalized Discounted Cumulative Gain) [12], with some special considerations due to the structured nature of the documents; and the personalization techniques are *NQE, HRR, SRR, IRR, NQE+m, HRR+m, SRR+m, IRR+m, CAS and CAS-or*, which represent a highly heterogeneous set of personalization techniques. You can see [7] for a more detailed explanation about any of these evaluation components.

Using the user profiles. We are going to explain how the profiles have been used in the experimentation.

1) *sProf and tProf:* The use of subject-based and term-based profiles is quite simple. It basically involves taking the top-n relevant subjects (*expSubj*) or terms (*expTerms*), with $n = 5, 10, 20, 40$. Once we have these first *expSubj* or *expTerms*, we normalize (proportionally) their weights in such a way that the maximum value (*maxNorm*) is: 0.33, 0.66, 0.99. The combination of *expSubj* or *expTerms* with *maxNorm*, gives us a total number of 12 different weighted subject or term sets, to provide to each personalization technique. Check Table 2 to see an example of these final user profiles from Table 1.

2) *stProf:* Its use is somewhat more complicated. In principle, the process should be to get the first *expSubj* profile subjects (again $expSubj = 5, 10, 20, 40$), and for each of these subjects to get the first *expTerms* terms, with values *expTerms*$= 1, 5, 10$. Each term weight will be multiplied by its corresponding subject weight. Thus the terms, which will be the ones finally used by the personalization techniques, will already incorporate in their weights the influence of their subjects.

Table 3. Final *stProf* user profile using $expSubj = 2$, $expTerms = 3$ (to make it more clear and short), and $maxNorm = 0.66$.

add	0.66*agriculture 0.421*aid 0.372*sector 0.237*agrarian 0.219*production
max	0.66*aid 0.583*sector 0.548*agriculture 0.371*agrarian 0.344*production
addFill	0.66*agriculture 0.421*aid 0.372*sector 0.307*farmer 0.237*agrarian 0.219*production
maxFill	0.66*aid 0.583*sector 0.548*agriculture 0.482*farmer 0.371*agrarian 0.344*production

But we find a problem in the previous process: when joining the different terms associated to different subjects, some of these terms are repeated (several subjects have terms in common, as `agriculture` in the example of Table 1). Since having repeated terms with different weights makes no sense, we consider the following approaches to fix the weights of these terms:

a) Add weights (add): collapse the repeated terms into one with a weight equals to the addition of the individuals weights.

b) Maximum between weights (max): we keep the repeated term with highest weight, removing the others.

c) Add weights, filling terms (addFill): same as *add*, but each time a term is deleted from a subject, the next one in the list is included until having *expTerms* terms for each subject.

d) Maximum between weights, filling terms (maxFill): same as *addFill*, but using maximum instead of sum.

The first two approaches involve that we do not always obtain the same number of terms for the personalization techniques, as it happens with the last two approaches. It should be noted that, in the last two approaches the filling process starts from the last *expSubj* subject, since we want more information from the most profile representative subject, i.e., the first *expSubj* subject. At the end of this process, the final terms will be also normalized with a maximum normalization value (*maxNorm*), with values: 0.33, 0.66, 0.99. The combination of the *expSubj*, *expTerms* and *maxNorm* gives us a total number of 36 different weighted term sets to provide to each personalization technique. Check Table 3 to see an example of these user profiles from Table 1.

4.1 Results

This section shows the results of the different experiments carried out, under the previous evaluation framework and considering the six proposed user profiles.

In Table 4, each cell represents the maximum NDCG value among the (12 or 36) different profile configurations, for each possible user profile and personalization technique. These values are the averaged NDCG values from the carried out user study 126 evaluation triplets.

Firstly, we can see that personalization helps the user to find relevant information, since in all cases we obtain significant improvements with respect to the non-personalized IRS performance (NDCG = 0.388). Depending on the profile and the personalization technique these improvements range from 50% to 80%.

The two main conclusions drawn from Table 4 are: 1) the best personalization technique is clearly *HRR+m*, and 2) the best user profile approach is *tProf* except in two cases, in which *stProf_maxFill* profile is better. Considering the last conclusion, we may assume that most of the times the best user profile approach to use is the simpler *tProf*, instead of the more complicated *stProf_maxFill*.

With respect to the used number of subjects or terms and the normalization value, the first row in Table 5 shows which profile configuration maximizes performance. We can see that the best user profile configuration is very homogeneous, independently of the user profile approach. This best user profile configuration is formed by $exp[Subj|Terms] = 40$, $expTerms = 10$ (in *stProf* profiles), and $maxNorm = 0.99$. Thus, it seems that the best user profiles are those including a rather high number of subjects and/or terms with high associated weights.

Also, the last two rows of Table 5 show the average and standard deviation for all the proposed user profile approaches. We observe that the highest average value is achieved by the *tProf* approach, with a low deviation value. Meanwhile, the lowest deviation value is achieved by the *sProf* approach, but with a much lower average value than *tProf*. Considering the four *stProf* approaches, we may observe a gradual decrease and increase in the average and deviation values, respectively, following the order of these profiles in the table. This situation indicates that within these user profiles, the further to the right in the table, they achieve more disparate personalization results (higher and lower), so more attention need to be paid to the selection of the right user profile configuration. The fact of having the maximum experimental evaluation performance with *stProf_maxFill* approach confirm this last conclusion.

Considering all the results, could it be concluded that we stand up for the *tProf* profile? Not necessarily, from a user perspective and considering not very small profiles, a *stProf* profile is much easier to understand than a *tProf* profile, since abstract concepts contain more semantics than isolated terms. It is also true that the *stProf* profile with two levels (concepts and terms) could be exploited by a given personalization technique to improve its performance, e.g., easily selecting parts of the user profile which suit more to the query (particularly helpful for heterogeneous profiles). Thus, depending on the application and the used personalization technique, a trade-off decision between pure performance or more expressiveness of the user profile must be taken. Additionally, from a cost

Table 4. NDCG maximum values from the 12-36 possible 'user profile-personalization technique' configurations. Original (non-personalized) NDCG value: 0.388. '*' character shows the best user profile approach for each personalization technique, and '+' character shows the best personalization technique for a given user profile approach.

	NQE	HRR	SRR	IRR	NQE+m	HRR+m	SRR+m	IRR+m	CAS	CAS-or
sProf	0.588	0.603	0.577	0.572	0.632	0.645^{+}	0.552	0.552	0.552	0.578
tProf	0.634*	0.652*	0.625*	0.620*	0.678	0.696^{+}	0.597*	0.597*	0.675*	0.668*
stProf_add	0.610	0.626	0.605	0.603	0.673	0.685^{+}	0.580	0.580	0.659	0.662
stProf_max	0.601	0.624	0.603	0.600	0.681	0.694^{+}	0.584	0.584	0.660	0.662
stProf_addFill	0.626	0.634	0.615	0.611	0.674	0.693^{+}	0.585	0.585	0.660	0.659
stProf_maxFill	0.612	0.633	0.606	0.602	0.683*	0.701^{+*}	0.587	0.587	0.658	0.660

Table 5. Best user profile configuration (*exp[Subj—Terms]-[expTerms]-maxNorm*), average and deviation, for each user profile approach.

	sProf	**tProf**	**stProf_add**	**stProf_max**	**stProf_addFill**	**stProf_maxFill**
Prof. conf. (max)	40-0.99	40-0.99	40-10-0.99	40-10-0.99	40-10-0.99	40-10-0.99
NDCG average	0.543	**0.602**	0.575	0.572	0.571	0.565
NDCG deviation	**0.032**	0.034	0.039	0.042	0.043	0.045

perspective these user profiles are not very demanding, since they only change with the inclusion of new documents, which does not happen very often.

5 Conclusions

In this work, we have presented 6 different user profile representation approaches based on content. Firstly, we focused on the subjects from a thesaurus, which are manually assigned to the initiatives in the documents by documentalists. These subjects are considered as concepts for the user profile based on *subjects*. Then, we did not take into account any other information than simply the document terms, to build the user profile based on *terms*. And finally, we proposed a hybrid approach between the two previous approaches (with four variations), having a two level user profile representation, where the first level is represented by *subjects* and the second level by the *terms* representing these subjects.

We have performed evaluation experiments including ten different personalization techniques and a wide range of user profile configurations, for all the proposed user profile approaches. We have obtained very good results, which in the best case reach up to 80.67% of improvement, with respect to the original non-personalized model. Additionally, we have demonstrated that most of the times the use of a simple user profile based on terms is enough to get good personalized results. Anyway, having a user profile with some structure and abstract concepts may help both, users to better understand their own profiles, and also some personalization techniques which may exploit this richer representation.

As future work, we would like to develop some personalization techniques to exploit the hierarchy of the proposed user profiles based on subjects and terms, and to use them to include personalization in privacy constrained environments.

Acknowledgements. Paper supported by the Spanish Andalusian "Consejería de Innovación, Ciencia y Empresa" and the "Ministerio de Ciencia e Innovación" under the projects P09-TIC-4526 and TIN2011-28538-C02-02, respectively.

References

1. J. Ahn, P. Brusilovsky, D. He, J. Grady, and Q. Li. Personalized Web Exploration with Task Models. Proceedings of the 17th International Conference on World Wide Web (WWW'08), 1-10, 2008.
2. J. Allan. Challenges in information retrieval and language modelling. In Report of a workshop held at the Center for Intelligent Information Retrieval, University of Massachusetts, Amherst, 2002.

3. N.J. Belkin. Some(what) Grand Challenges for Information Retrieval. SIGIR Forum 42(1): 47-54, 2008.
4. S. Calegari, and G. Pasi. Ontology-Based Information Behaviour to Improve Web Search. Future Internet 2(4): 533-558, 2010.
5. L.M. de Campos, J.M. Fernández-Luna, J.F. Huete, and A.E. Romero. Garnata: An information retrieval system for structured documents based on probabilistic graphical models. Proceedings of the 11th Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, 1024-1031, 2006.
6. L.M. de Campos, J.M. Fernández-Luna, J.F. Huete, C.J. Martin-Dancausa, C. Tur-Vigil and A. Tagua. An Integrated System for Managing the Andalusian Parliament's Digital Library. Program: Electronic Library and Information Systems 43(2): 121-139. 2009.
7. L.M. de Campos, J.M. Fernández-Luna, J.F. Huete, and E. Vicente-López. Using personalization to improve XML retrieval. IEEE Transactions on Knowledge and Data Engineering, to appear. DOI: 10.1109/TKDE.2013.75.
8. L. Chen, and K. Sycara. A Personal Agent for Browsing and Searching. Proceedings of the 2nd International Conference on Autonomous Agents, 132-139, 1998.
9. M. Daoud, L. Tamine, and M. Boughanem. Learning user interests for session-based personalized search. In ACM Information Interaction in conteXt (IIiX'08), 57-64, 2008.
10. J.F. Gantz. The Diverse and Exploding Digital Universe. IDC white paper (2): 1-16, 2008.
11. S. Gauch, M. Speretta, A. Chandramouli, and A. Micarelli. User profiles for personalized information access. The Adaptive Web: Methods and Strategies of Web Personalization, LNCS 4321, 54-89, 2007.
12. K. Jarvelin, and J. Kekalainen. Cumulative gain-based evaluation of IR techniques. ACM Transactions on Information Systems 20(4): 422-446, 2002.
13. A. Kobsa. Privacy-Enhanced Web Personalization. The Adaptive Web: Methods and Strategies of Web Personalization, LNCS 4321, 628-670, 2007.
14. E.L.C. Law, T. Klobučar, and M. Pipan. User effect in evaluating personalized information retrieval systems. LNCS, 4227, 257-271, 2006.
15. A. Micarelli, and F. Sciarrone. Anatomy and Empirical Evaluation of an Adaptive Web-Based Information Filtering System. User Modeling and User-Adapted Interaction 14(2-3): 159-200, 2004.
16. A. Moukas. Amalthaea: Information Discovery And Filtering Using A Multiagent Evolving Ecosystem. Applied Artificial Intelligence 11(5): 437-457, 1997.
17. G. Pasi. Issues in Personalizing Information Retrieval. IEEE Intelligent Informatics Bulletin 11(1): 3-7, 2010.
18. A. Sieg, B. Mobasher, and R. Burke. Web search personalization with ontological user profiles. Proceedings of the 16th ACM International Conference on Information and Knowledge Management, 525-534, 2007.
19. A. Spink, S. Ozmutlu, H.C. Ozmutlu, and B.J. Jansen. US versus european web searching trends. Proceedings of the 25th ACM SIGIR conference on research and development in information retrieval, 32-38, 2002.
20. K. Sugiyama, K. Hatano, and M. Yoshikawa. Adaptive web search based on user profile constructed without any effort from users. Proceedings of the 13th International Conference on World Wide Web, 675-684, 2004.
21. J. Trajkova, and S. Gauch. Improving Ontology-Based User Profiles. Proceedings of the RIAO 2004 Conference, 380-390, 2004.
22. G. Uchyigit, and M.Y. Ma. Personalization Techniques and Recommender Systems. World Scientific, 2008.

Product Line-based customization of e-Government documents

Mª Carmen Penadés[1], Pau Martí[1], José H. Canós[1], Abel Gómez[2]
[1] ISSI - DSIC, Universitat Politècnica de Valencia
46022-Valencia, Spain
{mpenades, pmarti, jhcanos}@dsic.upv.es
[2] AtlanMod, Mines-Nantes – Inria - Lina
44307 Nantes, France
abel.gomez-llana@inria.fr

Abstract. Content personalization has been one of the major trends in recent Document Engineering Research. The "one document for n users" paradigm is being replaced by the "one user, one document" model, where the content to be delivered to a particular user is generated by some means. This is a very promising approach for e-Government, where personalized government services, including document generation, are more and more required by users. In this paper, we introduce a method to the generation of personalized documents called Document Product Lines (DPL). DPL allows generating content in domains with high variability and with high levels of reuse. We describe the basic principles underlying DPL and show its application to the e-Government field using the personalized tax statement as case study.

Keywords: Document generation, personalized e-Government services, Software Product Lines, Document Product Line.

1 Introduction and motivation

Within the high diversity of activities placed under the e-Government umbrella, document generation and delivery are key activities required by most processes. Although content management issues – that is, document classification, organization, storage and retrieval – are well solved thanks to the advances on Digital Libraries research of last decades, current content generation solutions are far from effective. Many tools are in the marketplace supporting document-related processes within e-Government environments [1, 2]. Most of them provide tools to generate documents in different ways (from scratch, from templates, via copy/paste, drag&drop…), or workflow-like utilities to automate document circulation and publishing processes. Others provide document signing utilities, and content aggregators. However, there are many issues that still need to be addressed to improve the efficiency and quality of public e-services.

In this paper, we focus on personalization aspects of document generation, with the belief that adapting public documentation to the actual needs of citizens is the best way to keep them involved in the community's governance. Customization is especially relevant in multi-national agencies such as the UN or the EU, where a given document (e.g. an EU-wide law or a UNESCO white paper) must be produced in a high variety of languages, but those parts of a document which are not language-sensitive can be reused in all the instances. Similarly, some regulations may apply to only a subset of the member countries, making no sense to be included otherwise. This is a variation of the "one user, one document" principle as the way to maximize one reader's satisfaction by providing him/her with just the information he/she needs. We face the problem of e-government document customization as one challenge that remains unsolved in a satisfactory way. Current solutions to personalization are hard-coded in Web-based applications, or require high knowledge of low level languages such as XML. We intend to raise the abstraction level, providing tools and methods close to the application domain, hiding the internals to the final users.

Our solution borrows principles and techniques of Product Line Engineering [3] and is based on the definition of families of documents rather than one single document. A family groups a number of documents that share content in some sections while differ in others. Creating a particular document means selecting the appropriate content for that document and having tools for the automatic generation of the document from a set of content components stored in a repository. Our method, called *Document Product Lines (DPL)* [4], is supported by a tool that allows the definition of the family, the management of content components, and the generation of customized documents.

We will use an example of what we call "*sparse documents"* by analogy with Matrix Calculus (http://en.wikipedia.org/wiki/Sparse_matrix). A sparse matrix is such that most of its cells are set to zero. A *sparse document* is one whose content has been designed to a broad audience and, as a consequence, many parts of it are of no interest to a specific user. Examples of *sparse documents* are laws in several languages (one of whose readers is only interested in his/her own language), large emergency plans (with documentation relevant to different rescue teams), or tax statement forms that include sections for different types of economic activities. To illustrate the potential of DPL in the field of e-Government, we have chosen the tax statement completion as our case study. We show how the definition of a family of tax forms avoids users to find irrelevant sections in their statements, simplifying this way their completion since only those sections that are relevant for a particular taxpayer are included.

The paper is structured as follows. Section 2 outlines the DPL approach to support the generation of customized documents. Section 3 describes the tax statement as case study and how DPL is used to generate the personal income tax. Section 4 summarizes related works. Finally, Section 5 presents the conclusions.

2 An overview of DPL

DPL [4] applies product line engineering principles to the generation of documents in domains with high content variability and reuse. Central to DPL is the notion of family of documents. By family we mean a set of documents that share some common, mandatory parts while differ in other, optional parts. Every member of the family is built by assembling a set of content components.

DPL is structured as a two-stage process. First, in the *Domain Engineering* stage, a domain expert defines the characteristics of a family of documents in terms of features. There are two types of features: content features, which represent parts of content that will eventually be components of some document of the family, and technology features, that will allow to specify the different ways a given content feature can be rendered. For instance, a content feature like “street map” could be rendered as a high resolution image, or as an URL linking to some mapping service.

After the family feature model has been defined, every content feature must be linked to a content component. Such pieces are called *InfoElements*, and their granularity varies from very simple items (words or phrases) to large documents. As a rule, an *InfoElements* represents a self-contained reusable piece of content and, as such, can be used to build different documents in the same or even in different document families. DPL assumes the existence of a repository where document components are stored and organized for reuse. The repository is explored to find existing *InfoElements*. If no *InfoElements* in the repository can satisfy a requirement specified in the feature model, a new one should either be developed or retrieved from other repositories. New *InfoElements* development requires the availability of a library of applications to create and/or modify different types of them. Finally, the document product line is generated, that is, a process that specifies how the *InfoElements* are integrated according to the different relationships defined between content features, and between content and technology features.

The second stage, *Application Engineering,* exploits the document product line defined and supports the generation of customized documents. First, the user selects the variability points that will be included in the particular document (an instance of the family) to create a configuration. At that point, an automatic process generates the final document by assembling the different *InfoElements*. The output of the product line is user-defined; we can obtain documents in different formats (PDF, HTML…).

DPLFW [5] is an implementation of DPL developed following Model Driven Engineering principles. The main elements are shown in Fig. 1. The Feature Editor is used to characterize the variability of the domain as a document feature model. The Repository contains the *InfoElements* that will be reused in the generation of the

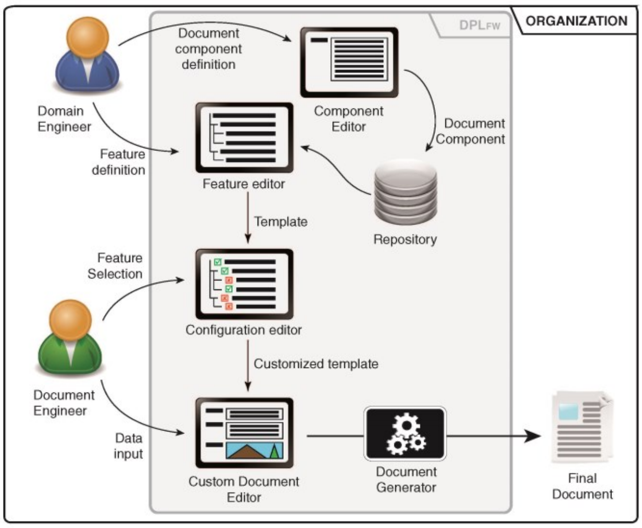

Fig. 1. An overview of the DPLfw

document. The Component Editor is used to create new *InfoElements* and add them to the Repository. All of the above elements support the Domain Engineering stage of DPL. The remaining elements are related to the Application Engineering subprocess. The Configuration Editor supports the selection of variability points. A given configuration is used to generate the Custom Document Editor. We talk about editor since, in some cases; an *InfoElements* may not be complete since some user-provided data is required. For instance, as we will see in our example, a tax statement is a set of forms that a final user has to fill in a given order. Then, the document generated will act as the editor of the true final version of the statement, which will be produced by the user following the logic of the forms generated with DPL. Finally, the Document Generator integrates these components to obtain a fully instantiated document generated in a specific format.

2.1 Variability supported in DPL

DPL supports two types of variability. On one hand, content variability is supported by the configuration editor, which allows selecting optional parts that are added to the mandatory part of the family to generate a new member. On the other hand, variable data support refers to the ability to use partially instantiated document components that contain placeholders that are replaced with values at the mail-merge style when the document is generated or, even, when a user completes the document.

The DPL document feature metamodel describes a document family as a set of document features and these are distinguished as content features (*ContentDocumentFeature or CDF*) and technological features (*TechnologyDocumentFeature or TDF*). A *CDF* can be associated to one or more *TDF*. As in a classical feature model, the features can be mandatory, optional or alternative, and may be related to other features by cross-like relationships such as "requires" or "excludes". For instance, "*f1* requires *f2*" means that every document containing the feature *f1* must also include the feature *f2*. Relationships are important to define document variability because all of these constraints bring the possibility to define that features appear in the whole or only in a group of documents of a document family. Each CDF need an *InfoElement* to provide their content which is visualized using a Disseminator [6].

Variable data in *InfoElement* is represented with the *VariableAttribute* class. The scope of a variable may be local to an *InfoElement*, or global to a feature model (that is, to the entire document). Finally, the *CriterionAttribute* is used to search and retrieve *InfoElements* when a custom document is being instantiated.

3 A case study: Customizing Tax Statement Documents

To illustrate how DPL works to generate customized documents, we use a practical case study of the e-Gov domain, namely the Personal Income Tax Statement in Spain. According to the Spanish regulations, all individuals who obtained incomes over a given threshold in Spain are obliged to file the Personal Income Tax yearly. The Spanish Tax Agency has published every year a set of tax document templates that taxpayers had to fill either by hand or using a Web interface. For the average taxpayer, the template was a large, sparse document that included pages for a variety of cases, most of which were relevant to a subset of the taxpayers (e.g. those owning stock options), being blank in the remaining cases. In the most usual case, more than a 40% of the tax statement was not filled. In the paper version, this meant an unnecessary waste of paper; in the Web version, a number of "Skip" buttons had to be pressed; in both cases, the statement became much more complicated for taxpayers, who had sometimes to interpret the non-trivial meaning of some sections, resulting in errors in the statement that, sometimes, resulted in fines from the Tax Agency. From 2014 the Tax Statement will only be filled electronically. Also, some customization has been included in the Web application that allows hiding unnecessary content to the user; however, the customization is part of the application logic. Using the DPL approach, things work differently. Although, as in the case of the Web application, the taxpayer selects those sections which are relevant to his/her status, such a selection is not part of the tax processing application, which makes changes much easier since the application would not need to be changed.

3.1 Specifying the Personal Income Tax document family

The first step in DPL is the definition of the family of tax statements by specifying the commonalities and variation points as a document feature model (recall Fig. 1). Using the feature model editor, a domain expert can build a document feature model like the one shown in Fig. 2. There, personal data (*Taxpayer* CDF), marital status (*Marital Status* CDF), income to declare (*Income* CDF), deductions to apply (*Deductions* CDF) and returns obtained according to law (*Tax Returns* CDF) are defined as mandatory content features (labelled with an exclamation point). On the other hand, the number of children (*Children* CDF) and the residence of the taxpayer (*Resident*

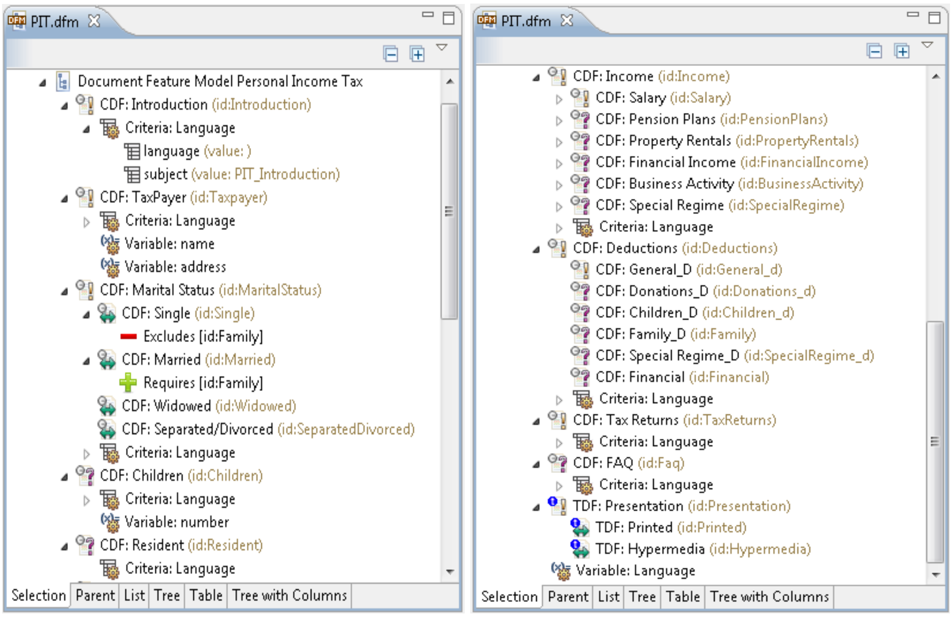

Fig. 2. Document Feature Model of the Personal Income Tax Statement family

CDF) are modeled as optional CDF (labeled with either question mark or a double-head arrow in case of optional or alternative content features, respectively).

A taxpayer must declare some of the following incomes: salary derived from labour relationships (*Salary* CDF), pension plans (*Pension Plans* CDF), rentals from estate property (*Property Rentals* CDF), financial interest or dividends (*Financial Income* CDF), professional fees derived from independent activities (*Business Activity* CDF), and the special tax regime, if any (*Special Regime* CDF). Only the salary is a mandatory CDF, being the remaining ones optional. Similarly, the deductions to apply to the tax amount are represented as children of the *Deductions* CDF. The model is completed with "requires" or "excludes" relationships between CDFs that are used to comply with the rules established by law. For instance, in Fig. 3, the content feature *Single* is exclusive with regard to the *Family* feature.

There is a global variable in the model of Fig. 2, namely *Language*. It is being used to select the final language of the statement. This illustrates another application of DPL to support variable content in multilingual e-government applications. The variable declaration is complemented with its inclusion as search criteria in the model. Notice that some of the CDFs have a child named *Criteria: Language.* This means, on one hand that the actual content for the CDF will be selected at document generation time according to the value of the *Language* variable (see section 3.3).

3.2 Managing the Personal Income Tax content assets

After creating the document feature model, CDFs must be linked to *InfoElements* in the repository. The repository provides facilities to manage *InfoElements* (i.e. creating, deleting and updating) and retrieving them using metadata-based searches. Furthermore, *InfoElements* can be organized hierarchically using folders and resources (a special container which only may contain *InfoElements*). Fig. 3.a shows a view of the DPLFW's repository manager. Notice that there are several folders

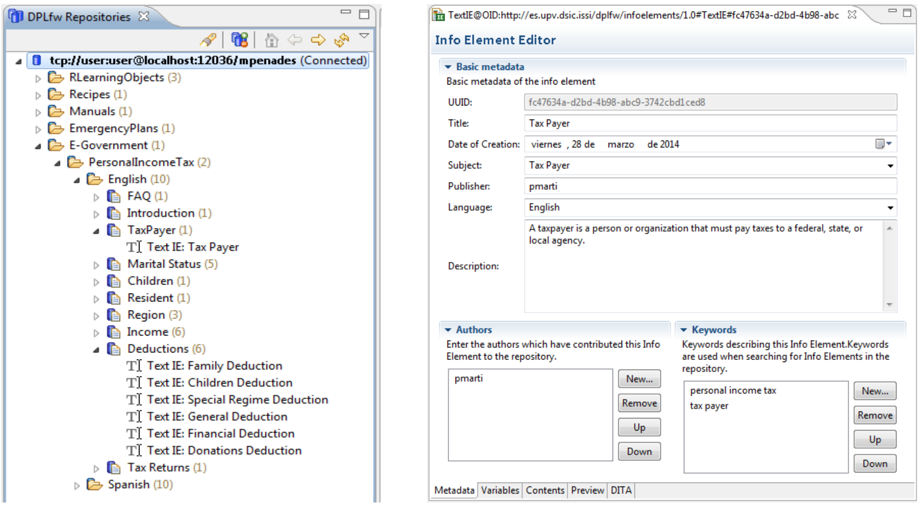

(a) **(b)**

Fig. 3. a. DPL Repository Explorer **b.** Describing the *Tax Payer InfoElement* with metadata

available, one per each domain we are working on, although a DPLFW user can opt for alternative organizations (e.g. storing *InfoElements* for every domain in a single folder). There is a folder named *E-Government* whose content is partially shown in the figure. The leaves in the tree correspond to *InfoElements*, which contain the actual content that will be linked to the CDFs in the model.

When the content assets are not stored in the repository, we need to create a new *InfoElement* using the DPLFW's *InfoElement* Editor. There are several aspects of an *InfoElement* that must be defined, which can be selected by clicking on the tabs situated in the bottom left area of the editor window. Fig. 4.b shows the descriptive metadata capture window when editing the *Tax Payer InfoElement* inspired in the Dublin Core Metadata Set. To edit the content on an *InfoElement*, the contents tab in the editor provides a rich text editor (see Fig. 4.b). This editor allows defining variable data which will be instantiated in the configuration of a member of the document family. In this instance, the *Tax Payer InfoElement* contents two previously defined variables, namely name and address (see Fig. 4.a).

3.2 Configuration of the tax statement

Once the document family has been defined as a document feature model and a Personal Income Tax repository is available, the DPL process is ready to generate customized tax statements. This is performed in two steps. The first one is known as document configuration. A configuration consists of set features that have been selected according to the variability constraints defined by the document feature model; mandatory features are always selected. Figure 5 shows two different configurations of the tax statement. The first one, shown in Fig 5a, corresponds to the case of a woman, married and with two children whose only income is her salary. The language selection of the statement is Spanish. When the married and children features are selected, the deduction features related to family and children are automatically selected according to the document feature model. If the woman has no other deductions, such as donations, the selection of features has finished. At this point, if a given feature (or the full document) has variable attributes, a value must be introduced for each one. For instance, the values of name and address for *TaxPayer* feature have to be introduced, as well as the salary's. Additionally, if a feature has a criterion attribute, the *InfoElement* will be searched and retrieved according to its

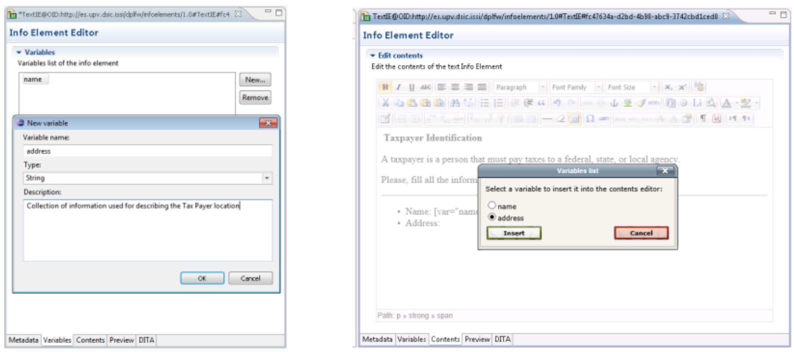

(a) **(b)**

Fig. 4. a. Defining variable data on the *Tax Payer InfoElement* **b.** Editing their content

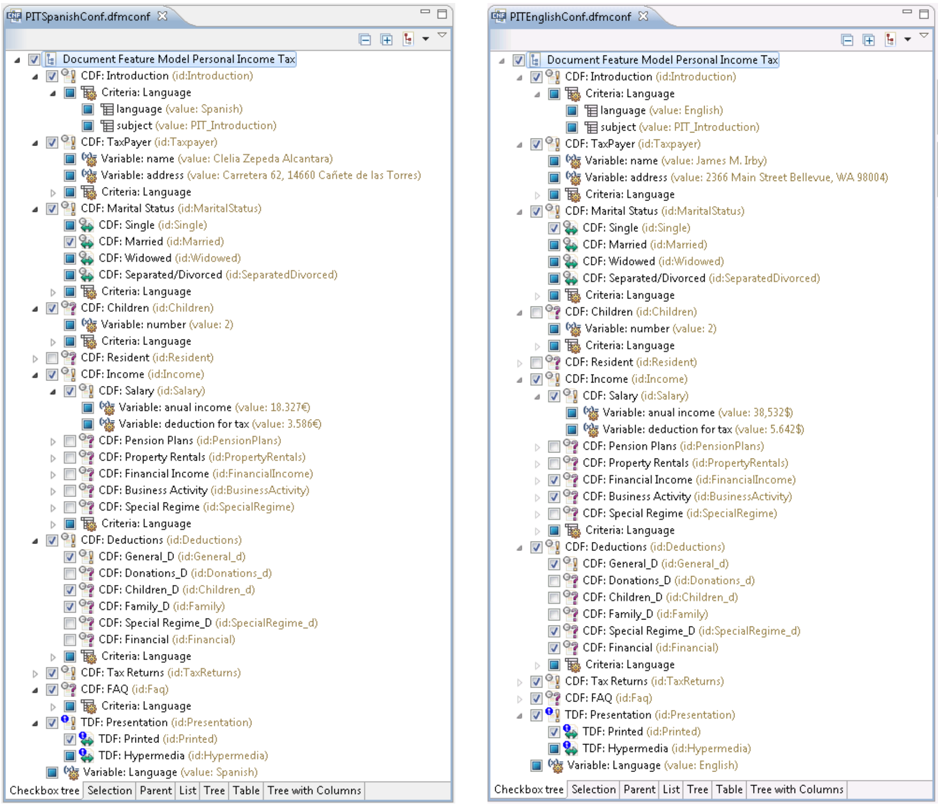

(a) (b)

Fig. 5. a. Spanish Personal Tax configuration **b.** English Personal Tax configuration

value. Next, the customized PIT document can be generated as a printed document. Another tax statement configuration is shown in Fig. 5b. In this case, the language selected is English, and the configuration corresponds to a man who has a business activity as manager and financial income, so that the deductions marked are special regime and financial.

3.4 Customized tax statement generation

When a configuration is finished, an automatic process generates a map describing the structure of the Personal Income Tax Statement configured. The map, along with the *InfoElements* retrieved, is used to generate the customized tax statement. Fig. 6 shows a screenshot of the final Personal Income Taxes obtained.

4 Related work

Some proposals have been presented to deal with the variability in the document generation in the last years. For example, proposals on Variable Data Printing (VDP) use XML and their associated technologies to implicitly represent the variability in

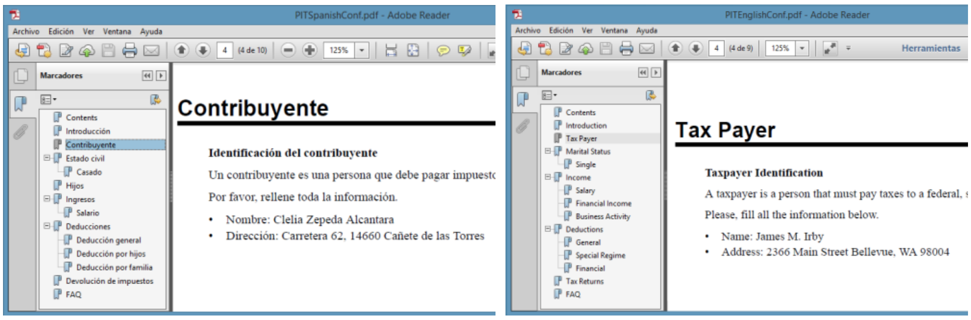

Fig. 6. The Personal Income Taxes generated

documents [7, 8]. In these cases, the variability model is usually represented by the transformations rules and the links defined among a set of XML-based document fragments. The customized document is generated using XSLT or XPath, and a high knowledge about XML is required. The VDP approaches are authoring/editing tools to generate the final XML document and not provide methodological guidance to the document engineer.

More recent proposals such as [9, 10] have explored a product line approach to model the variability explicitly. A feature model identifies the variability points from a domain-oriented perspective and is the basis to support the generation of the customized document. These approaches emphasize the definition of a document generation process based in SPL principles. However, DPLFW supports variable data and variable content in the document generation process, whereas other proposals do not support the variable data ([9]) or they do not generate the final document ([10]).

Some commercial solutions, such as HotDocs (www.hotdocs.com) or Exari (www.exari.com), provide a workflow-based suite tool for document automation. The former has a document generator server which takes a web-based *interview* and generates a document from it. The latter focuses on the domain of contract automation, and it defines a contract lifecycle to improve and automate the document generation. These solutions have similarities to DPLFW since all of them generate a final customized document, but their starting points are different – i.e. DPLFW uses feature models to identify variability points in a document family, and enforces content reuse at domain level following a product line approach.

Finally, there are some proposals that are tailored to customization of e-government documents. In [11], an adaptive hypermedia application reasons about concepts and conditions to produce web materials. The input is a generic document published in the Web by the Public Administration, and using semantic web technologies, a customization of this document is produced (containing the relevant information for a specific user only). In this scenario, the DPL approach would behave as follows: The Public Administration produces a set of content document (the *InfoElements*) and DPLFW generates the customized documents on demand. We also find some domain-oriented proposals within the e-government field: software solutions in the tax statement domain – such as TurboTax, UDoTaxes, StudioTax, QuickTax Online, etc. – create a full tax form for each user. DPLFW also generates a full tax assessment according to the configuration phase and the variable data introduced by the user.

5. Conclusions and Futher Work

The European eGovernment Action Plan 2011-2015 is among a series of initiatives that appeared all around the world to enforce the growth of e-Government. The ultimate goals are, among others, to increase the citizenship awareness by easing the access to the information, and making citizens closer to the governance by means of better public participation mechanisms. Quoting: *"Increasing effective eGovernment means that services are designed around users' needs and provide flexible and personalised ways of interacting and performing transactions with public Administrations"*. Similar initiatives have been launched in other countries such as the USA (http://www.usa. gov/) or India (http://www.indiaegov.org/), to name a few. All the above goals have document personalization at the core of any solution.

In this paper, we have shown how the DPL approach may provide the flexibility that current tools lack in terms of customization and reuse. DPLFW allows the generation of customized documents without writing a single line of code. We switch from programming to feature modelling in a domain-oriented language, hiding the internals of the document generation processes. We have used DPL in domains other than e-Government, such as emergency management, e-learning and cooking recipes generation. All of them share the same personalization and reuse requirements, and show that a general solution is more flexible than domain-oriented tools under some circumstances. We are working on the improvement of the DPLFW to support new facilities like external the document configuration process in a web application to improve the user interface.

ACKNOWLEDGMENTS. The work of M.C. Penadés and J.H. Canós is partially funded under grant TIPEx (TIN2010-19859-C03-03).

References

1. e-Gov Office. http://www.newgensoft.com/products/egovoffice
2. eGov-Suite. FabaSoft. http://www.egov-suite.com/en/dms/create.html
3. Pohl, K., B Böckle, G. and van der Linden, F., Software Product Line Engineering – Foundations, Principles, and Techniques. Springer, Berlin, Heidelberg, New York. (2005)
4. Penadés, M., Canós, J.H., Borges, M.R., Llavador, M. Document product lines: variability-driven document generation. In Proc. 10th ACM DocEng '10, NY, USA, 2010.
5. Gómez, A., Penadés, M., Canós, J.H., Borges, M. and Llavador, M. A framework for variable content document generation with multiple actors. Information and Software Technology. http://dx.doi.org/10.1016/j.infsof2013.12.006. In Press (2014)
6. Kahn, R.,Wilensky, R. A Framework for Distributed Digital Object Services. Handle: cnri.dlib/tn95 - 01. http://www.cnri.reston.va.us /k - v.html. (1995)
7. Lumley, J., Gimson R., Rees, O., A Framework for Structure, Layout & Function in Documents. In Proc. 5th ACM Symposium on DocEng (2005)
8. Quint, V., Vatton, I., Editing with Style. Proc 7th ACM Symposium on DocEng (2007)
9. Rabiser, R., Heider, W., Elsner, C., Lehofer, M., Grünbacher, P. and Schwanninger, C. A flexible approach for generating product-specific documents in product lines. In Proc. SPLC'10 , LNCS vol 6287, Springer. (2010)
10. Pichler, C. Huemer, C. Feature modeling for business document models. Proc. ACM SPLC'11., LNCS, Vo. 2, New York, NY, USA. (2011)
11. Colineau, N, Paris, C, Linder, K.V. An evaluation of tailored Web materials for Public Administration. In Proc. ACM HT'12. New York, NY, USA. (2012)

A Fuzzy Model for Service Value Assessment

M. Alessandra Torsello, Leo Iaquinta, and Marco Comerio

University of Milano-Bicocca - Dep. of Informatics, Systems and Communication
Viale Sarca 336, 20126 Milano (Italy)
{torsello, iaquinta, comerio}@disco.unimib.it

Abstract. E-Government (e-Gov) is becoming an important means to produce value for citizens by using innovative technologies in the delivery of more advanced, efficient and personalized services. In this paper, we propose a model for the assessment of the service value defined as the trade-off between benefits and sacrifices perceived by citizens according to their service usage experience. Since human perceptions are subjective and uncertain in nature, the model proposes the use of fuzzy concepts to effectively represent and handle data under uncertain conditions. The suitability of the proposed model is shown by its application on a case study in the e-Gov domain.

Keywords: service value, value assessment, fuzzy evaluation, service evaluation.

1 Introduction

E-Gov creates opportunities to provide added value services to citizens, increasing efficiency and reducing costs. The concept of service value is of interest for researchers and practitioners as it represents one of the aspects that mainly affect user behavior and satisfaction level. In addition, the analysis of service value could provide useful insights for service personalization.

In literature, service value (referred also as *service value in use*) is typically defined as the trade-off between the benefits which users receive by using the service and the sacrifices that they bear in order to get that service [13, 9]. Many works define service value in use as the overall assessment of the utility of a service based on the perception of what is received and what is given [10]. In many studies, a variety of value components that determine benefits and sacrifices for users have been identified. Such components are related to aspects characterizing the service such as its functional properties (FPs), non-functional properties (NFPs) and qualities (Qs). Specifically, FPs concern the functionalities realized by the service, NFPs (e.g., price and payment method) define how the service performs its functionalities, and Qs are aspects characterizing the efficiency and effectiveness of the service. By gathering perceptions of users about such components, it is possible to evaluate the overall benefits and sacrifices for users who use a considered service and thus to assess the value in use of that service. The assessment of service value may be considered a key element for the provision of personalized services taking into account perceptions and satisfaction level of user groups. Such groups identify user segments including users with similar characteristics which perceive the service in analogous manner. To these segments, organizations may offer more personalized services able to better satisfy the peculiar needs of users.

In this paper, we propose a model for the assessment of the service value in use. The model determines benefits and sacrifices for users by taking into account how a set of service aspects are perceived. The gathered perceptions, being the result of the human thought, have an extremely subjective and uncertain nature. To better capture the uncertainty and the subjectivity that permeate the evaluation process of a service, the model exploits concepts of fuzzy logic. Broadly speaking, the model represents user perceptions in terms of fuzzy sets and it defines a set of fuzzy rules that express the relationship between the actual values of aspects and the benefit/sacrifice as perceived by users. The value in use of the considered service is determined by the fuzzy inference of such rules properly combined with the relevance of each aspect.

The paper is organized as follows. Sec. 2 discusses related works. Sec. 3 describes the proposed model for service value assessment. Sec. 4 reports the results obtained by applying the model to a case study. Finally, Sec. 5 draws conclusions and outlines future works.

2 Related Work

Service value assessment typically involves perceptions expressed by users on service characteristics. In the evaluation processes based on user perceptions, different works propose classical statistical approaches [2, 7]. However, such approaches may result ineffective due to the subjective and uncertain nature of human perceptions. Fuzzy logic is proposed in several research works to handle imprecise knowledge typical in human reasoning. In particular, fuzzy logic has been used for service quality evaluation. For instance, a fuzzy set approach has been proposed in [4] where the customer subjective opinions and the weight of considered factors are described by fuzzy linguistic scales. Each linguistic term is represented by a fuzzy number. The approach considers the importance of each factor and computes the overall fuzzy ratings of all alternatives by using fuzzy number operations. In [3] the authors presented a fuzzy multi-attribute decision-making approach for evaluating dynamically the service quality. Here fuzzy numbers are used to solve the ambiguity of concepts that are associated with human subjective judgments vaguely measured with linguistic terms. In [6] a method based on triangular fuzzy numbers is proposed to measure perceived service quality. The discrepancy between consumer perceptions and expectations is evaluated as the intersection area between two triangular fuzzy numbers.

In previously mentioned works, fuzzy logic has been mainly used to model and process user perceptions in evaluation processes. However, many works exploit fuzzy logic for its ability to build fuzzy inference systems, i.e. models able to simulate the reasoning of a human expert when he has to take decisions in environments characterized by uncertainty and imprecision. Fuzzy systems are gaining widespread acceptance in service quality assessment. In [8] a hybrid fuzzy expert system is applied to investigate service quality in the academic library. The system is a combination of four fuzzy expert systems: three systems that work in parallel to evaluate three different aspects of libraries and the fourth system which determines the library service quality. In [1] a fuzzy method for evaluating user perception of the security level on social networking sites is presented. Inputs to the system are fuzzy sets representing linguistic variables for information security evaluation. A set of fuzzy rules is built based on the intuitive knowledge of the relationships between the variables.

In this paper, we propose a service value assessment model that exploits fuzzy concepts both to represent user perceptions and to evaluate benefits/sacrifices deriving from the use of the service. Specifically, fuzzy numbers are used to model perceptions. Moreover, a number of fuzzy systems is defined to express relationships among aspects and related benefits/sacrifices. Service value is estimated as the trade-off of benefits and sacrifices determined for all selected aspects by considering their relative importance.

3 The proposed model for service value assessment

To assess the value in use of a generic service s for a set of users U, our model comprises the following steps detailed hereafter:

1. Selection of the service aspects to consider as value components;
2. Gathering of user perceptions about the selected aspects;
3. Evaluation of benefits/sacrifices for each selected aspect;
4. Determination of the weights corresponding to the aspects;
5. Evaluation of the overall benefits and the overall sacrifices for users;
6. Determination of the service value in use.

1. Selection of the service aspects In our model, service value arises from the evaluation of some aspects characterizing the same service related to FPs, NFPs and Qs. Each aspect is a value component and it can represent either benefits or sacrifices for users. For example, the price that a user pays to obtain the service is always considered as a sacrifice that he has to bear and when the price is very low, the sacrifice becomes null. On the contrary, the quality of a service is always considered as a benefit for the user and when the quality is very low, the benefit becomes null. Referring to the service considered in the experimental activity of our model, the following aspects have been selected by distinguishing among benefits/sacrifices and specifying (in brackets) their typology:

- **benefits**: Delivery time (NFP), Transparency (NFP), Fulfillment of user needs (FP), and Overall quality (Q);
- **sacrifices**: Price (NFP), and Request time (NFP).

2. Gathering of user perceptions about the aspects The evaluation of each selected aspect is performed by exploiting the perceptions of users that express judges about their service usage experience. To gather user perceptions, different ways could be employed such as interviews, questionnaires, focus groups, etc. In this work, perceptions are gathered by requiring users to fill a questionnaire that includes questions about each considered aspect. Users express their perception by choosing one of the levels among those included in Likert scales with a odd number of levels labeled by linguistic terms. The choice of adopting scales with linguistic terms is essentially due to the fact that users express their perceptions in more natural way in words rather than by numeric values. In addition, for each aspect related to NFPs, users are required to specify the actual value experimented in their experience as well as a range of values retained acceptable for that aspect. Such values are useful to quantify the benefit/sacrifice determined by the aspect. The considered aspects do not have all the same importance and different users may ascribe them different relevance degrees. In order to determine the aspect importance, the questionnaire includes questions requiring users to sort the aspects according to their relevance.

3. Evaluation of benefits/sacrifices for each aspect User perceptions are processed to quantify benefits/sacrifices related to each considered aspect. User perceptions are subjective and uncertain. To better handle the imprecision and the subjectivity of the gathered perceptions, our model exploits the ability of fuzzy logic to represent and process information under uncertain conditions.

First of all, each linguistic term of the evaluation scales adopted in the questionnaire is represented by a triangular fuzzy number (TFN), i.e. a fuzzy set with a triangular shape. Such kind of representation helps to deal with the imprecision inherent verbal perceptions. Fuzzy numbers, with respect to the use of crisp values, allow to better capture the subjectivity of the judges expressed by users. Moreover, since fuzzy numbers are based on real values, the information contained in them may be better handled, explored and mathematically exploited. According to [12], the semantic of scale linguistic terms is defined by fuzzy numbers as shown in Fig. 1 depicting a scale with 7 linguistic terms, namely N (None), VL (Very Low), L (Low), M (Medium), H (High), VH (Very High) e P (Perfect).

To determine benefits and sacrifices related to each aspect, two different procedures are performed according to the aspect typology: a first procedure concerns aspects related to FPs and Qs, and a second procedure concerns NFPs.

Procedure for FPs and Qs As concerns such kind of aspects, perceptions expressed in terms of fuzzy numbers are aggregated by the fuzzy average operator [5] in order to evaluate the corresponding benefit value represented by a TFN. Finally, this is defuzzified to obtain the corresponding crisp numeric value. Different defuzzification processes could be adopted. This work uses the "center of the area" method [11] that substantially associates to the TFN the abscissa of the geometric center of the area under its membership function.

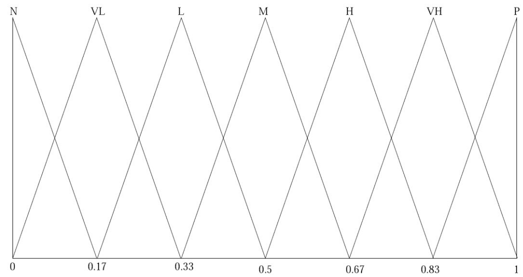

Fig. 1. A scale of 7 terms with its semantics

Procedure for NFPs As concerns NFPs, perceptions of value ranges retained acceptable by users are determined. Such value ranges are processed to derive the benefit or the sacrifice to be associated with actual values assumed by an NFP. To do this, for each aspect, a set of fuzzy rules is defined starting from the gathered data relying on the knowledge of the domain expert. For instance, one rule for the price aspect is defined in the following form:

$$\text{IF } price \text{ is CHEAP THEN } sacrifice \text{ is LOW}$$

To derive such rules, fuzzy sets on input and output variables in the antecedent and the consequent of each rule have to be defined. Fuzzy sets on input variables are defined by analyzing the gathered data. In particular, perception levels in the Likert scales are merged to obtain three fuzzy sets corresponding to positive, neutral and negative perceptions. Neutral perceptions coincide with the middle scale level. Positive and negative perceptions refer respectively to the highest and the lowest scale levels. Membership functions are defined by exploiting the frequency of values on which users have expressed the respective perception level. The membership degree of a value to the respective fuzzy set is proportional to the corresponding frequency. As an example, the three fuzzy sets CHEAP, FAIR, and EXPENSIVE defined for the price input variable of a sample service are depicted in Fig. 2(a). On the output variables (benefit or sacrifice according to the considered aspect), three fuzzy sets are defined, namely LOW, MEDIUM, and HIGH, as shown in Fig. 2(b). Once fuzzy rule set is defined for the considered aspect, the benefit/sacrifice in correspondence of the value taken by that aspect is determined by the inference process of such fuzzy rules [5].

4. Determination of the aspect weights Our model associates a numeric value ranging from 0 to 1 to each aspect representing the relevance degree assigned by users. Weights are determined by considering how each user filling the questionnaire orders the service aspects in terms of the conferred relevance. Precisely, as a first step, based on the specified position, a score is assigned to each aspect. In our model, the score assignment is inspired to competition ranking where each ranked position is associated with a numeric value that reflects the

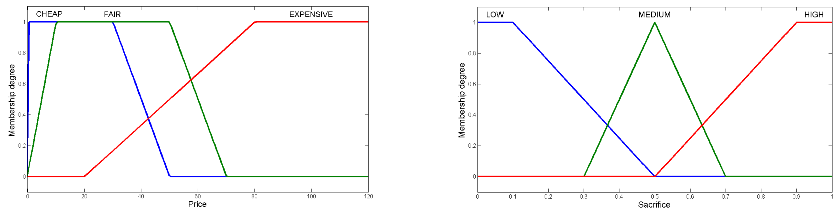

(a) (b)

Fig. 2. Membership functions for price (a) and sacrifice (b) of a sample service

relationships between a set of competitors according to their final position in the ranked list. Next, aspects are split up into two sets representing benefits and sacrifices and the total score for each set is computed. Successively, the weight of each single aspect is calculated as the ratio between the score assigned in the first step and the total score of the set it is member of. In this way, the sum of benefit weights and the sum of sacrifice weights are separately equal to 1. Finally, the weight of each aspect is estimated as the average of weights of that aspect calculated for all users.

5. Evaluation of the overall benefits and sacrifices Weights and values of benefit/sacrifice for the selected aspects are combined to determine the value of the overall benefits and sacrifices for all users. The overall benefits are calculated by the weighted average of the benefit values related to the aspects that determine some kind of benefit for users, as follows:

$$ob_{s,U} = wb_1 * ben^1_{s,U} + ... + wb_{N_b} * ben^{N_b}_{s,U} \quad (1)$$

where N_b is the number of aspects classified as benefits, wb_i, $i = 1, ..., N_b$, represent the weights for those aspects, $ben^i_{s,U}$, $i = 1, ..., N_b$, are the respective benefit values. Analogously, the overall sacrifices are calculated by the weighted average of the sacrifice values related to the aspects that determine some kind of sacrifice for users, as follows:

$$os_{s,U} = ws_1 * sac^1_{s,U} + ... + ws_{N_s} * sac^{N_s}_{s,U} \quad (2)$$

where N_s is the number of aspects classified as sacrifices, ws_j, $j = 1, ..., N_s$, represent the weights for those aspects, $sac^j_{s,U}$, $i = 1, ..., N_s$, are the respective sacrifice values.

6. Determination of the service value in use In our model the value in use of a service s perceived by a set of users U is computed as the ratio between the overall benefits $ob_{s,U}$ (as in eq. 1) and the overall sacrifices $os_{s,U}$ (as in eq. 2) as follows:

$$v_{s,U} = \frac{ob_{s,U}}{os_{s,U}} \quad (3)$$

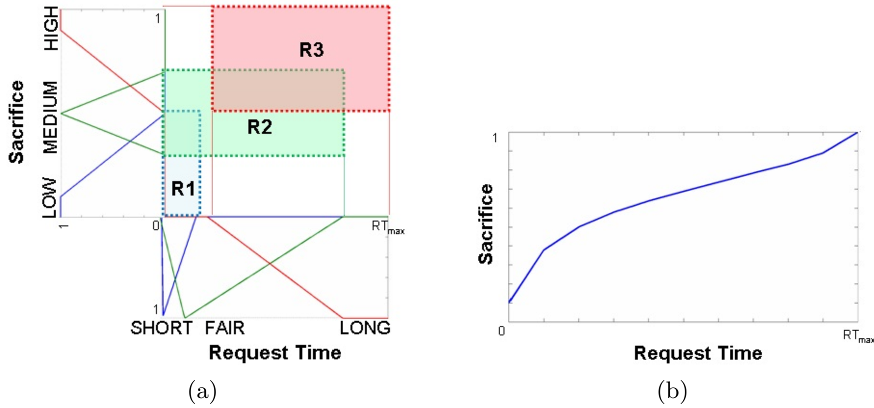

Fig. 3. Rules (a) and sacrifice curve (b) for request time

A value in use equal to 1 is obtained when perceived sacrifices are fully balanced by perceived benefits. A value greater than 1 means that perceived benefits overcome perceived sacrifices. On the contrary, a value in use less than 1 indicates that, for that service, sacrifices are perceived by users as weightier than perceived benefits.

4 A case study

The proposed model for service value assessment was applied on a case study to verify its suitability. In the SMART (Services and Meta-services for smART eGovernment) project, we analyzed services needed for entrepreneurs who want to open public businesses. Among these, in this work, we focus on the results obtained from the value analysis performed on the service of Internet connectivity provision to public businesses being one of the most experienced services.

The first step was consisted in selecting the aspects to consider as value components. For the considered service, the aspects listed in Sec. 3 were selected. The price aspect was distinguished into *activation price* and *monthly price* to indicate respectively the fee paid by the entrepreneurs when the service provision starts and the fee paid each month for the provision. In this way, a total number of 7 aspects were selected, that are *activation price*, *monthly price*, *request time*, *delivery time*, *transparency*, *fulfillment of user needs*, and *overall quality*.

In the second step of our model, user perceptions were gathered by questionnaires investigating the usage experience of about 10 services useful to open public businesses in Italy such as café and Bed & Breakfast. At the end of this step, a total number of 102 questionnaires filled by entrepreneurs were collected. However, each entrepreneur was asked to answer questions about at most three experienced services. Thus, for the Internet connectivity service considered in the case study, perceptions expressed by about 20 entrepreneurs on the aspects previously selected were gathered.

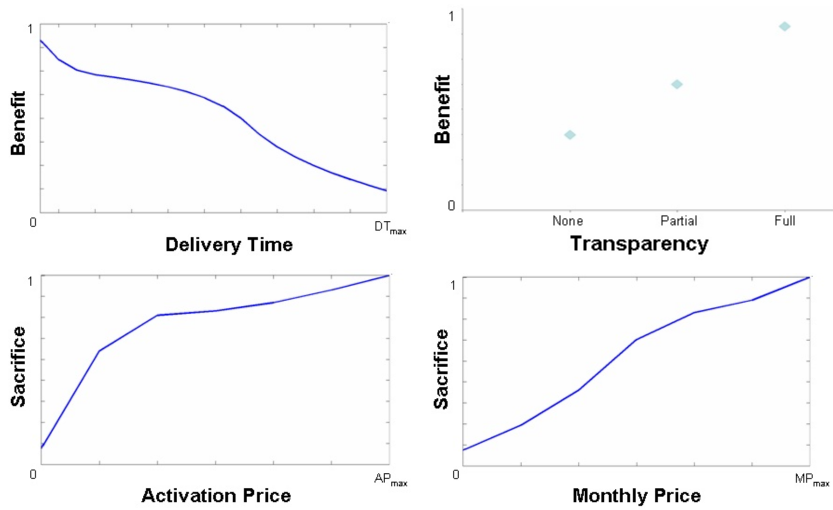

Fig. 4. Benefits/sacrifices for delivery time, transparency, activation/monthly prices

In the third step, benefits and sacrifices related to each aspect were obtained by applying the appropriate procedure as described in Sec. 3. Precisely, as concerns NFPs, for each aspect a set of fuzzy rules was defined starting from the gathered perceptions. As an example, the fuzzy rule set derived for the *request time* aspect is shown in Fig. 3(a). As it can be seen, three rules have allowed to cover the gathered perceptions and to establish the relationships among *request time* and *sacrifice* values. The derived relationships are synthesized in the following rules:

R1: IF *request time* is SHORT THEN *sacrifice* is LOW
R2: IF *request time* is FAIR THEN *sacrifice* is MEDIUM
R3: IF *request time* is LONG THEN *sacrifice* is HIGH

Fig. 3(b) shows the sacrifice curve obtained by the inference of rules for all possible request time values. In Fig. 4, we show the trend of benefits and sacrifices obtained for the other considered aspects related to NFPs. Particularly, the transparency values are not in a continuous range (like for the other considered NFPs), but 3 different levels were considered, namely *None*, *Partial*, and *Full*, corresponding to the different degrees with which entrepreneurs receive information about the provision progress of the requested service.

As concerns FPs and Qs, perceptions expressed in terms of TFNs were aggregated by the fuzzy average operator. The TFNs obtained as a result were $\tilde{F} = (0.47, 0.64, 0.80)$ and $\tilde{Q} = (0.42, 0.67, 0.91)$ for fulfillment of user needs and overall quality, respectively. Such TFNs were defuzzified into crisp values obtaining 0.64 and 0.67, respectively.

In the next model step, weights of all the aspects were calculated as explained in Sec. 3 and the obtained values are reported in the third column of Table 1. To

Table 1. An example of value in use assessment for a sample service

Aspect	Contract term	Weight	Benefit	Sacrifice	Value in use
Activation price	98 €	0.20		0.81	
Monthly price	45 €	0.60		0.83	
Request time	1 day	0.20		0.38	
Delivery time	7 days	0.46	0.63		
Transparency	Partial	0.18	0.30		
Fulfillment of user needs		0.18	0.64		
Overall quality		0.18	0.67		
Overall			**0.58**	**0.74**	**0.78**

apply the last two steps of the model, we need to consider the values included in the contract of a specific service. Thus, by considering a service for Internet connectivity provision characterized by the contract terms reported in the second column of Table 1, benefits and sacrifices of each aspect can be computed by exploiting the curves derived for NFPs and defuzzified values derived for FPs and Qs. Such values are reported in the same Table 1. By computing the weighted average of benefits and sacrifices previously obtained, the overall benefits and sacrifices were determined having respectively 0.58 and 0.74.

As a final step, the value in use of the considered service was derived as the ratio of overall benefits and sacrifices obtaining a value equal to 0.78.

The performed value analysis points out that entrepreneurs perceive the monthly price as the most relevant aspect (weight 0.60) and the related perceived sacrifice is also the highest (0.83). On the other hand, the second most important aspect (weight 0.46) is the delivery time for which entrepreneurs perceive a quite high benefit (0.63). In correspondence of all the other aspects, entrepreneurs associate a low relevance degree (weight 0.18 or 0.20) and, consequently, these components weakly affect the assessed value. Thus the most important sacrifice perceived for the monthly price is mainly alleviated by the benefit perceived for the delivery time. As a consequence, the resulting final value in use (0.78) is quite near to 1 (that corresponds to the situation in which benefits balance sacrifices) despite the most relevant sacrifice referred to the monthly price. Such kind of analysis can be useful whenever a user has to choose among different available services. In fact, in such situations, indicators about perceived benefits and sacrifices may provide helpful cues to support users in selecting the most valuable and personalized service.

5 Conclusions

In this paper a model for assessing the service value in use has been proposed. An accurate analysis of service value may offer useful hints to provide personalized services able to meet the peculiar needs of users. Service value is quantified as the trade-off of benefits and sacrifices perceived by users when using the service. To better model the subjectivity inherent any perception-based evaluation process, the model benefits from the ability of fuzzy logic to handle information under uncertain conditions. Results obtained by the experimental activity carried out on a service of the e-Gov domain encourage the application of the model on other services and perceptions of a wider number of users to definitely assess the effectiveness of the model. As future work, starting from a larger collection of user perceptions, the automatic derivation of membership functions and rules could be an interesting point to address together with the possibility to automatically refine the rules by a proper learning process.

Acknowledgment
The work presented in this paper has been partially supported by the Italian PON project PON01 00861 SMART (Services and Meta-services for smART eGovernment).

References

1. E. Afful-Dadzie and A. Vesely. User perception of security on social networking sites using fuzzy logic. *British Journal of Applied Science & Technology*, 3(4):714–734, 2013.
2. C. Bai and C. Ye. Construct a public service quality evaluation model: evidence from chinese public service industry. In *Proc. of Services Systems and Services Management (ICSSSM'05)*, volume 1, pages 508–513, 2005.
3. J.M. Bentez, J.C. Martn, and C. Romn. Using fuzzy number for measuring quality of service in the hotel industry. *Tourism Management*, 28(2):544–555, 2007.
4. C.T. Chen. Applying linguistic decision-making method to deal with service quality evaluation problems. *International Journal Uncertainty Fuzziness Knowledge-Based Systems*, 9:103–114, 2001.
5. G. Chen and T.T. Pham. *Introduction to Fuzzy Sets, Fuzzy Logic, and Fuzzy Control Systems*. CRC Press, 2000.
6. C.-J. Chien and H.-H. Tsai. Using fuzzy numbers to evaluate perceived service quality. *Fuzzy Sets and Systems*, 116(2):289–300, 2000.
7. P. Monari, M. Bini, and D. Piccolo. *Statistical Methods for the Evaluation of Educational Services and Quality of Products*. Contributions to Statistics. Physica-Verlag HD, 2009.
8. A.A. Pourahmad, M. Neshat, and A. Baghi. Service quality assessment in the academic library: Use of hybrid fuzzy expert system. *African Journal of Business Management*, 6(46):11511–11529, 2012.
9. R.A. Spreng. The impact of perceived value on consumer satisfaction. *Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior*, 6:50–55, 1993.
10. W. Ulaga and S. Chacour. Measuring customer-perceived value in business markets: A prerequisite for marketing strategy development and implementation. *Industrial Marketing Management*, 30(6):525–540, 2001.
11. W. van Leekwijck and E.E. Kerre. Defuzzification: criteria and classification. *Fuzzy Sets and Systems*, 108(2):159–178, 1999.
12. M.-M. Yu and J.-S. Yao. Modification of concordance analysis based on statistical data and ranking of level $1-\alpha$ fuzzy numbers. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 09(03):313–340, 2001.
13. V.A. Zeithaml. Consumer perceptions of price, quality, and value: A means-end model and synthesis of evidence. *Journal of Marketing*, 52(3):2–22, 1988.