



A Semantically Enhanced UPnP Control Point for Sharing Multimedia Content

The Universal Plug and Play (UPnP) protocol lets users share multimedia content across devices and display it in multimedia renderers in home-based LANs. However, UPnP doesn't support sharing this content between homes (for example, with friends, relatives, or workmates), and content can be difficult to find. A UPnP extension provides uniform access to multimedia content belonging to different homes, such as content from a user's Facebook friends. In addition, an extended UPnP control point lets users add semantic annotations to their multimedia resources and link them to the linked data cloud to facilitate their discovery by others.

**Mariano Rico
and Oscar Corcho**
Universidad Politécnica de Madrid

**Víctor Méndez
and José Manuel
Gómez-Pérez**
*Intelligent Software Components
(iSOCO), Madrid*

We live in a multimedia-centric world in which users share all types of professional and user-generated multimedia resources and want to render them anywhere, anytime, and in a large set of heterogeneous devices. Users share most of these multimedia resources through Web 2.0 sites (such as Flickr and YouTube), where they upload, annotate, comment on, find, and download various types of multimedia content while controlling access to sensitive content for trusted colleagues, friends, or relatives. Demand for better multimedia management tools is therefore increasing not only on these online sites, but also in home environments. Such tools let users share their own (often sensitive or private) multimedia content with trusted individuals without having to upload it on external sites and,

most importantly, without losing the copyright.

In home environments, the Universal Plug and Play (UPnP) protocol,¹ defined by the UPnP Forum and supported by many device manufacturers, lets heterogeneous devices share multimedia content without having to copy material. Multimedia content stored in UPnP media servers can be streamed to any UPnP media renderer located in the same home using software applications running in PCs or on set-top boxes – the *UPnP control points*.

However, users also demand access to their multimedia resources outside their home environments so they can listen to a music track on a friend's computer, for example, or show photos or videos from a recent visit with relatives. Although UPnP is a first step toward the provision of multimedia

Table 1. Comparison of relevant multimedia home sharing technologies currently available.

Feature	iTunes	Spotify	W7HG ^a	UPnP	UpGrid
Access outside LAN	No	Yes	No	No	Yes
Authentication	Yes	Yes	Yes	Yes	Yes
Streaming	Yes	Yes	Yes	Yes	Yes
Annotation	Yes ^b	No	No	No	Yes ^c
Search	Yes	Yes	Yes	No	Yes

Notes: (a) Windows 7 Home Group technology, (b) tag-based, and (c) tag- and ontology-based.

management and sharing tools, some important limitations prevent its use in sharing content across homes:

- Its use is restricted to LANs. Previous work has explored extending UPnP outside LANs,² and the UPnP Remote Access Committee is developing standards to let UPnP media servers, renderers, and control points outside the home network interact with devices inside a LAN. However, this work offers only partial or difficult-to-use solutions. For example, a network connection with a remote home is usually through ad hoc virtual private networks, which require complicated handling of network address translation.
- UPnP doesn't provide native authentication and authorization mechanisms, and ongoing efforts only partially describe extensions to handle this. Security and privacy are crucial in a multihome environment.
- Most of the navigation through UPnP devices is done by directory hierarchy browsing, and no advanced search capabilities exist. Consequently, UPnP-enabled systems aren't very usable, especially given a huge amount of heterogeneous multimedia resources.

Furthermore, because the metadata associated with these resources are often general metadata stored in MPEG-7 descriptors or similar formats, it's difficult to find resources using searches such as, "give me the photos or videos from where my cousins were on holiday."

To address these issues, we extended the standard UPnP protocol to handle multihome settings. We also developed a software architecture for creating URIs that supports the concept of "extended home" or "network of interconnected homes" as well as a UPnP control point with semantically enabled annotation and search capabilities. Metadata plays a

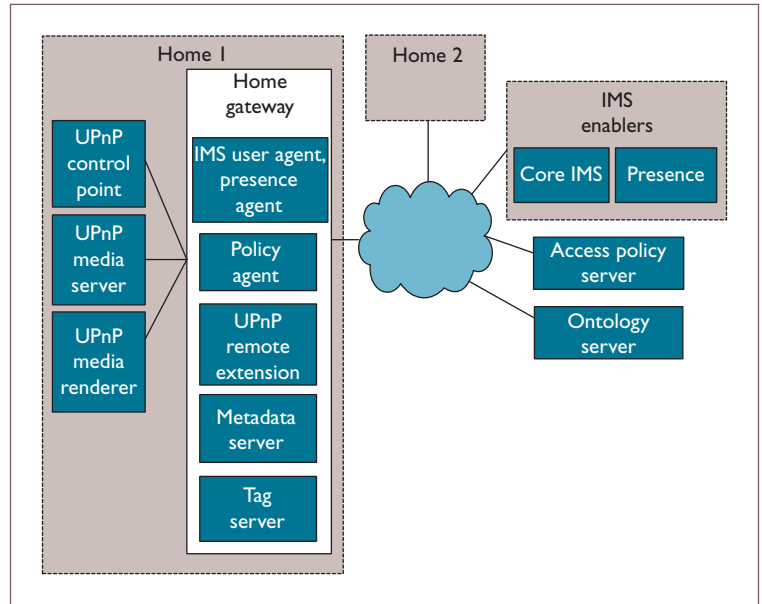


Figure 1. Simplified UpGrid architecture. Homes 1 and 2 are Internet-connected homes with local components (not shown in Home 2). The elements on the right are the centralized components.

key role in our extension, following metadata-management principles proposed in advanced metadata-management architectures. Ontology-based annotations³⁻⁵ are linked to the Web of linked data (<http://linkeddata.org>) to provide existing information as much as possible.

Other initiatives for sharing multimedia content across homes exist, as Table 1 shows. However, these systems are based on downloading resources or are limited to users in a LAN, and they don't use semantic annotation and search.

System Architecture

As Figure 1 illustrates, the UpGrid architecture contains a set of distributed components (located in each home) and a set of centralized components (deployed by the network provider). The distributed components include

those deployed in the home gateway and those that support UPnP-enabled functions in UPnP-based environments.

The following home UPnP-enabled components are usually found in any UPnP-enabled environment:

- *The UPnP media server* (such as a hard disk) stores and shares media content. It serves the content to be rendered by media renderers, and can be browsed from any control point.
- *UPnP media renderers* (such as TVs or digital photo frames) deliver content coming from media servers, and are controlled from control points.
- Standard control points (such as Windows Media Player and Cidero) let the user browse through media server content and send content to media renderers.

We extended the control point to support our extensions. The *extended UPnP control point* can also establish UPnP sessions with remote homes to allow content search and to annotate multimedia resources using tags and ontologies.

The home gateway connects devices in the home to the Internet. Home gateways include the following components:

- The *metadata server* stores ontology-based annotations created by local users. Annotations can refer to local or remote multimedia resources. This component follows the Semantic Open Grid Services Architecture (S-OGSA) for metadata management⁶ and allows ontology-based searches over the entire set of ontology-based metadata available in the participating systems (as long as the authorization restrictions allow access to it).
- The *tag server* stores user-defined tags related to multimedia resources.
- The *IMS user agent* establishes IP multimedia subsystem (IMS) remote sessions between homes and updates the home presence status, using the functions provided by a centralized IMS enabler.
- The *UPnP remote extension* extends UPnP visibility to all the homes participating in the remote UPnP session. That is, it allows authorized remote devices or multimedia resources to act as if they belonged to the same LAN.
- The *access policy agent* acts as the policy enforcement point (PEP) for authorization requests.

The UPnP remote extension uses the agent to provide different visibility configurations.

The centralized components include the following:

- An *IMS enabler* is an architectural framework for delivering IP multimedia services. The enablers that we've created let home gateways register in the IMS network, establish secure IMS sessions, and dynamically update each home's presence.
- The *access policy server* provides authorization services with different levels of granularity: homes, devices, multimedia resources, and metadata. The access policy server also facilitates tracking and control of UPnP session requests, acting as the policy decision point (PDP) for authorization requests.
- The *ontology server* stores the ontologies that will be used for annotations, including multimedia ontologies (such as COMM), and general ontologies about persons, locations, events, and so on. Any online ontology server can replace this component.

Because this architecture is not sufficient to achieve our goal, we extend the UPnP specification.

UPnP Extensions

Our first set of extensions aims to expand the standard UPnP protocols for managing users and creating UPnP sessions beyond the LAN.

Unlike other approaches that handle authorization at a granularity of homes, our approach considers that a home might have a set of users, which can be associated with roles, and applies role-based authorization mechanisms for granting or denying access to resources through the access policy server.

When a user logs in to our control point, the system provides a personalized view of the user's contacts, UPnP devices, and multimedia content, regardless of whether they're local or remote. Actions allowed for each user, such as adding new contacts and remote homes, or authorizing other users to view their local content, can also be personalized. Finally, a Facebook plug-in integrated in the control point facilitates user management.

After users log in, they can establish a set of UPnP remote sessions with each of the available

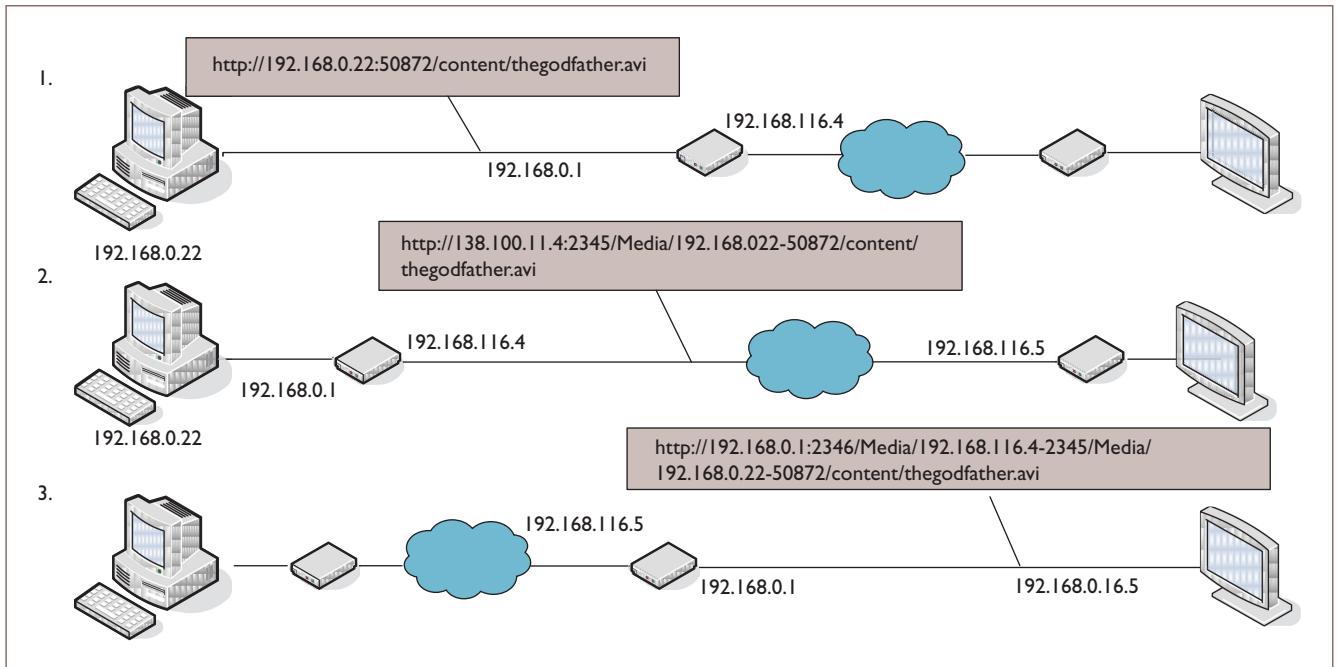


Figure 2. Sample interaction and resource URI transformation across homes. A user sends a request to view a movie from a computer in a connected home using the UPnP remote extension.

remote homes that they're authorized to access. UPnP remote sessions are established on top of IMS sessions, providing the necessary levels of security and quality of service. To start a UPnP session that includes a remote home, the IMS user agent installed in the home gateway starts an IMS session. The destination home's IMS user agent constantly listens for incoming IMS sessions. When the destination IMS user agent receives an incoming IMS session, it sends an authorization request to the access policy agent.

Once the agent discovers a device, it can invoke action methods such as browsing or searching (for media servers) and can receive events. The UPnP remote extension also proxies these methods using the standard UPnP control protocol (SOAP over HTTP), and includes a component that tracks the devices connected from a remote home so it can send them notifications when a remote home is disconnected.

Global Identification and Location

The standard UPnP protocol proposes using URIs for identifying resources available in UPnP-enabled devices as well as actions that users can invoke on those resources. However, these URIs are addressed from a local network, and aren't directly accessible from outside. In fact, when a UPnP server is asked for a URI, it returns the resource's physical address using a local IP address.

Our resource-identification proposal is based on the virtualization of resources available across homes through the use of global URIs that are independent of UPnP sessions and homes. This facilitates resource delivery and management, as well as the linking and publication of their corresponding metadata according to linked data principles.

Home gateways modify local URIs assigned to resources inside a home network for every request or response that leaves the local network. This modification includes a prologue to the original URI with the home gateway public IP address and the port where it is listening, and a string that indicates the type of message being sent (for example, a media request or event message). This last modification informs the UPnP remote extension about the type of message it's handling, so it can direct the message to the appropriate handler, and this speeds processing. The only exception to this modification is when URIs are already global public URIs (for example, when we access a media server available in a device that's connected to the Internet and has a permanent public IP address).

Figure 2 shows an example of these modifications and the process for retrieving a multimedia resource from a remote home. Imagine we're connected to a friend's home and we want to access the movie "The Godfather," which is

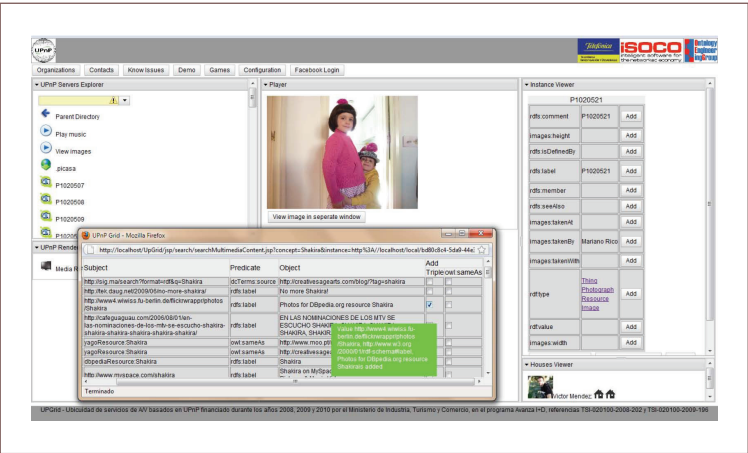


Figure 3. UpGrid control point user interface. The main window shows how a resource (a picture in this case) is rendered, with the semantic data shown in the Instances Viewer panel. Users can add additional semantic information (RDF) to the resource in the lower window. Videos are also available at www.youtube.com/results?search_query=UPnPGrid.

available in AVI format. To do this, we'd take the following steps:

1. We send a request to our friend's computer asking for the movie's location. Because the computer is working as a UPnP media server, it returns `http://192.168.0.22:50872/content/thegodfather.avi`.
2. The response arrives at the UPnP remote extension of our friend's network, which reads it and modifies the URI by adding the public IP address of our home network: `http://138.100.11.4:2345/Media/192.168.0.22-50872/content/thegodfather.avi`. This is the public global URI that will be used across home networks and to link with the metadata stored in metadata servers.
3. When this package arrives at our UPnP remote extension, we add its private IP address, so it can be handled by UPnP-enabled services as if it were a local resource: `http://192.168.0.1:2346/Media/192.168.116.4-2345/Media/192.168.0.22-50872/content/thegodfather.avi`.

Our local renderer then receives a local URI (the one generated in step 3) that our home gateway can convert into a physical location. If we decide to play this movie, the UPnP remote extension will extract the IP address and port of our proxy from this local URI, along with additional information about the message type, and will send the request to the address specified in the rest of the URI path.

Ontology-Based Annotation (Metadata) Management

Ontology-based annotations (generally referred to as metadata) of multimedia resources available in local or remote home environments are stored in local metadata servers, giving users complete control over all their metadata content. Thanks to the aforementioned URI extensions and the ability to access remote content, these annotations can seamlessly refer to local or remote multimedia resources. Figure 3 shows a snapshot of the user interface.

Metadata Generation

The enhanced control point's user interface lets users browse through local or remote accessible multimedia resources, and provides their ontology-based annotations. During the annotation process, users can add RDF triples either directly or through searches over the Web of linked data. For instance, when annotating a song from CD, the user can search for information about the song, CD, artist, and so on, querying various SPARQL end points (for example, DBPedia, MusicBrainz, and BBC), and select triples from the results that allow linking these descriptions to existing URIs or copying triples into the metadata being generated.

Metadata Management

We consider metadata as a first-class entity in our architecture, similar to multimedia resources and devices, with its own lifetime properties and authorization mechanisms, and with strong links to the multimedia resources it describes. This is supported by the semantic bindings concept, as defined in S-OGSA,⁶ for the robust management of this information. In addition to its inherently distributed nature, the use of such architecture supports the easy disclosure, disposal, modification, and archiving of metadata, even when the devices holding the multimedia resources being described are unavailable.

Given metadata management's distributed nature in this context, we use SPARQL-DQP (Distributed Query Processing),⁷ an extension of the OGSA-DQP system⁸ that supports queries over distributed metadata servers. The UpGrid system publishes this metadata using a linked data approach. Our home gateway transforms any external content request requiring an RDF view of a specific resource into a query to the corresponding local metadata server, which provides the set of annotations it has about that specific resource, local or remote.

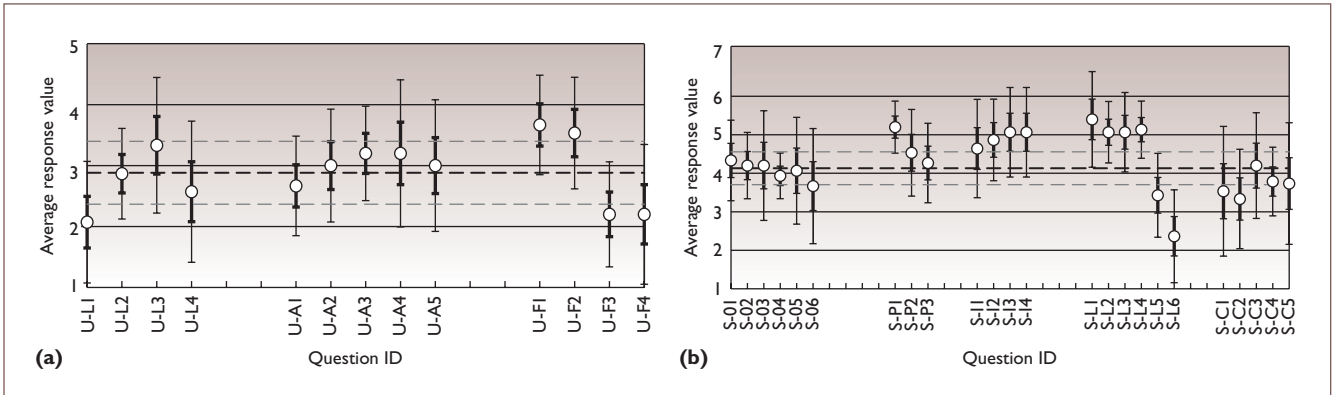


Figure 4. Evaluation results for the (a) usability and (b) satisfaction parts of the questionnaire. Circles represent the mean value of users' responses. Thin bars show the standard deviation. Thick bars show the mean's 90 percent confidence interval. Dark dotted lines show the average value of the responses, and light dotted lines show standard deviation bounds.

UPnP User and Social Networks

The UpGrid control point has a strong connection with social networks (such as Facebook), which allows an easier bootstrapping of the system in terms of user and authorization rights management. The control point lets users establish and end sessions with authorized homes. Whenever we want to connect with an authorized home, our UpGrid control point sends a Session Initiation Protocol (SIP) call to the remote home gateway and, once accepted, has access to the authorized resources.

The UpGrid control point also provides search functionalities using tags and ontology-based annotations, and an alert service associated with these search functionalities that informs the user of any changes to the ontology-based metadata available in our local metadata server or any other authorized remote metadata server according to a given SPARQL query.

UpGrid Control Point Evaluation

We selected 15 participants from our institutions to evaluate the UpGrid control point. Although small, the group represents the potential UpGrid control point users. We created a practical use case and presented it to each participant. The use case covered most of the UpGrid application's functionality, such as annotation of images and audio, semantic search in the available repositories, and management of friends' multimedia resources. Most participants took about 30 minutes to complete the use case.

Next, we gave each participant a detailed questionnaire (available at <http://sites.google.com/site/upgridquestionnaire>). This questionnaire


had 49 questions focusing on three aspects: the user's skills in UPnP and Web technologies, the control point's usability, and the user's satisfaction with the control point user interface. Most of these questions were based on standard questionnaires.⁹

We measured the control point's usability using the "Practical Heuristics for Usability Evaluation" test.¹⁰ This test includes 13 questions with responses ranging from 1 (bad) to 5 (good), providing a useful measure of the user's perception of the control point's usability. Figure 4a shows the test results, including the average value assigned by the participants to the questions related to usability, as well as the average value and deviation bounds.

We based the analysis of usability on the average value assigned in this test, which was 2.88, with a standard deviation of 0.52. This shows medium usability values for the entire range of participants, although the most skilled users assigned slightly lower usability values. A possible explanation is that skilled users are more demanding, which results in slightly lower evaluations. These results confirm that UpGrid is a usable tool even for less skilled users.

To evaluate user satisfaction with the control point user interface, we created a slightly modified version of the "User Interface Satisfaction" test.¹¹ The standard version includes 27 questions, but we reduced this number to 25 because of overlaps with the usability test. Valid responses to these questions were positive integers ranging from 0 (not satisfied at all) to 7 (completely satisfied). Figure 4b shows the results. The average value for user satisfaction was 4.13 with a standard deviation of 0.43.

The results show the dependency between the user's satisfaction and his or her skills. User satisfaction depends on user skills in the same way that usability does — that is, more skilled users assign a slightly lower value to satisfaction.

Our next steps will focus on providing a packaged product that can be easily installed in home gateways using a range of operating systems and browsers. The availability of annotations in the Web of data to facilitate user-based annotation and linking, and the wider adoption of multimedia-based ontologies, will help determine our approach's success. 

Acknowledgments

The Spanish government funded this work under the Avanza program (TSI-020100-2008-202, TSI-020100-2009-0196). We thank the Telefónica R&D staff for supporting this project.

References

1. *UPnP Device Architecture Version 1.0*, UPnP Forum, June 2000; www.upnp.org.
2. *Group Services and System Aspects, IP Multimedia Subsystem, Stage 2*, release 6, technical specification, project 3GPP TS23.228, Third Generation Partnership (3GPP), 2005; www.3gpp.org.
3. A. Chakravarthy, F. Ciravegna, and V. Lanfranchi, "Cross-Media Document Annotation and Enrichment," *Proc. 1st Semantic Web Authoring and Annotation Workshop (SAAW 06)*, CEUR Workshop Proceedings, 2006; <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-209>.
4. S. Bloehdorn et al., "Semantic Annotation of Images and Videos for Multimedia Analysis," *Proc. 2nd European Semantic Web Conf. (ESWC 05)*, LNCS 3532, Springer, 2005, pp. 592–607.
5. K. Petridis et al., "M-OntoMat-Annotizer: Image Annotation Linking Ontologies and Multimedia Low-Level Features," *Proc. 10th Int'l Conf. Knowledge-Based Intelligent Information and Engineering Systems (KES 06)*, LNCS 4253, Springer, 2006, pp. 633–640.
6. Ó. Corcho et al., "An Overview of S-OGSA: A Reference Semantic Grid Architecture," *J. Web Semantics: Science, Services, and Agents on the World Wide Web*, vol. 4, no. 2, 2006, pp. 102–115.
7. C. Buil Aranda, M. Arenas, and O. Corcho, "Semantics and Optimization of the SPARQL 1.1 Federation Extension," *Proc. Extended Semantic Web Conf. (ESWC 11)*, LNCS 6643, part II, Springer, 2011, pp. 1–15.
8. S. Lynden et al., "The Design and Implementation of OGSA-DQP: A Service-Based Distributed Query Processor," *Future Generation Computer Systems*, vol. 25, no. 3, 2009, pp. 224–236.
9. G. Perlman, "Web-Based User Interface Evaluation with Questionnaires," <http://oldwww.acm.org/perlman/question.html>.
10. G. Perlman, "Practical Usability Evaluation," *Extended Abstracts on Human Factors in Computing Systems (CHI 97)*, ACM Press, pp. 168–169.
11. J.P. Chin, V.A. Diehl, and K.L. Norman, "Development of an Instrument Measuring User Satisfaction of the Human-Computer Interface," *Proc. ACM Conf. Human Factors in Computing Systems (CHI 88)*, ACM Press, 1988, pp. 213–218.

Mariano Rico is a teaching assistant in the computer science department at the Universidad Autónoma de Madrid and a collaborator in the Ontology Engineering Group at the Universidad Politécnica de Madrid. His research interests include Semantic Web technologies, virtual worlds, user interfaces, and natural language processing. Rico has a PhD in computer science from the Universidad Autónoma de Madrid. Contact him at mariano.rico@uam.es.

Oscar Corcho is an associate professor in the artificial intelligence department at the Universidad Politécnica de Madrid, where he's a member of the Ontology Engineering Group. His research activities focus on semantic e-Science and real-world Internet, although he also conducts research in the more general areas of Semantic Web and ontological engineering. Corcho has a PhD in artificial intelligence from the Universidad Politécnica de Madrid. Contact him at ocorcho@fi.upm.es.

Víctor Méndez is a research fellow at Intelligent Software Components (iSOCO) in Madrid. His research interests include Semantic Web applied to multimedia, opinion mining, reputation, and provenance. Méndez has a BSc in computer engineering from the Carlos III University of Madrid. Contact him at vmendez@isoco.com.

Jose Manuel Gomez-Perez is the director of R&D at Intelligent Software Components (iSOCO) in Madrid. His research aims to support users in creating, sharing, and accessing knowledge and spans knowledge acquisition, provenance analysis, intelligent information access, and their applications. Gomez-Perez has a PhD in computer science from the Universidad Politécnica de Madrid. Contact him at jmgomez@isoco.com.