On Reusing Other People's Experiences

Enric Plaza

The Web is a vibrant environment for innovation in computer science, AI, and social interaction; these innovations come in such great number and speed that it is almost impossible to follow them. This paper will focus on some emerging aspects on the web that are a great opportunity and challenge for AI, specifically the large amount of records of experiences about the world that individual people share in the Web. I discuss a new approach that, instead of focusing on improving information access in the web, aims at supporting people to reuse other people's experiences, recorded in the web, in order to take more informed actions in the world.

1 Introduction

We habitually reuse other people's experiences in our daily activities. Since large amounts of these experiences are now recorded on the web, we try to use the web to find relevant information which can help us to take more informed decisions. The current web, however, is based on a metaphor where resources are "documents" upon which we perform a "search" in order to find one (or a few) relevant documents. In this paper I surmise that there is a different class of problems to which this *information access* paradigm is unsatisfactory, namely the task of reusing other people's experiences, recorded on the web, in order *to improve our decisions and actions on the world*. Moreover, I will propose that this presents a challenging opportunity for AI, and specifically for Case-based Reasoning: reusing the large amount of *experiences* that individual people share in the Web.

These recorded experiences —ranging from client reports about hotels they have visited to short explanations on how to do certain things— are searched for and reused by thousands of people every day. These experiences can be found in forums and blogs, in normal web pages and in specialized services like Question-Answer websites. Nevertheless, they are treated as *documents*, not as experiences. That is to say, they are represented, organized, analyzed, and retrieved as any other document. The main purpose of this paper is to postulate that there is a special kind of content, namely *experiences*, that provides a specific form of knowledge, *experiential knowledge*, and that they should be represented, organized, analyzed, and retrieved in accordance to their nature. Moreover, the paper will provide some food for thought by proposing some ideas on the conditions required and the techniques suitable to build systems capable of reusing experiential knowledge provided by other people in specific domains.

Notice that the focus of this paper is not on improving *information access* in the web, as it is the case for research developing better search/retrieve techniques. In other words, my approach is not about *finding something* (in the Web) —it's about *doing something* (in the world), taking an action in the world, and the Web is merely used to take a more informed decision or action. For this reason my emphasis is on *reusing* the experiential knowledge provided by others for the actual purposes of a final user. Moreover, this reuse implies that the content a user is interested in finding depends crucially on the actual purposes of that user. Therefore, the goals of the search/retrieve techniques will switch from finding a few items to display to assessing a large amount of (potentially) relevant items that have a bearing on the decision or action to be taken by the user. There is an approach in AI that has singularly and enduringly dealt with the analysis and reuse of experiences, as the next section discusses.

2 The Case for Experiences

Case-based reasoning (CBR) may be understood, first and foremost, as learning to solve problems (or take decisions) from past experiences. More specifically, past experiences are represented in the form of a collection of *cases*, where a case (*situation1*, *outcome1*) is to be understood as knowing that in the past, when what is described in *situation1* held, then the *outcome1* (that may be a consequence or a decision) also happened. Thus, a case is a statement (at some level of description) of a fact observed or experienced in the world. Additionally, CBR systems use case-based inference (also called analogy and similarity-based inference) based on the assumption that when a new *situation2* is similar to an old *situation1* then we can plausibly predict that an *outcome2* similar to *outcome1* is correct.

The representation of cases, situations and outcomes may be very different across domains (from k-NN classification to case-based planning); but they have in common that they present the knowledge of an observed factual situation: e.g. "this is a good hotel because my stay was very agreeable", or "I did this sequence of actions (this plan), in this situation, and I achieved that goal". Although there are no "cases" as such on the web there is a huge amount of this kind of *practical knowledge* present today in the web. This kind of practical knowledge coming from the direct observation or experiences of people is what I will hereto refer as *experience(s)*.

The technological challenge is how to represent, organize, and reuse experiential content. I surmise that the first step to address this challenge is to recognize that there is such a thing as "experiential content," and not merely hyperlinked documents. The way content is organized nowadays is a network of documents, and possibly in the next future, *annotated documents* (using ontology-defined concepts or folksonomy-based tags).

Moreover, the way users work with web content is what I'll call *Search & Browse* (*S&B*). The web users typically *first* need

to use a search engine to find a "resource," this search engine may be an external (e.g. Google or Yahoo to find a website or a page) or an internal (e.g. search inside a forum for the posts that may talk about the topic of interest). *Next*, the users need to browse a (sometimes disturbingly) large collection of "found items," perform a cursory read of them to filter out those blatantly irrelevant, then carefully read the rest (while eliminating those subtly irrelevant) to isolate the relevant content. Finally, the users have to *reuse* the relevant content, that may be dispersed in dozens of pages in different websites; notice that there is no support for the users' task and they simply use "copy & paste" to aggregate the information found or print all those pages and then aggregate that information manually.

A specific example may be useful to illustrate this scenario. Let us consider the task of deciding which hotel to book and consider the existing experiential content of previous hotel clients that describe their good and bad experiences after staying in those hotels. Let us say there are H hotels in the intended destination, W websites with hotel-related experiential content, and each hotel in each website has on average C client reports: a user to be well informed would need to *Search & Browse*, on average, $H \times W \times C$ user-contributed experience items. This is a huge amount of valuable information, but ineffective if it is to be manually processed —as it is the case in the *S&B* paradigm where there is no support for the task the users want to carry out, and for which reason they have performed a search in the first place.

Certainly, the users are capable of cutting down the work by filtering out information: by selecting a few websites (equivalent to performing a sampling operation $w = sample(W)$), by reducing the eligible hotels using hard constraints like "3-star hotels only" (a filtering operation $h = filter(H)$), and finally reading only a subset of all client reports (a sampling operation $c = sample(C)$), the workload is reduced to examining $h \times w \times c$ client reports. Notice that there is no computer support to perform a good sampling of websites or client reports: the users have no way to know if they acquire a *good sample* of the documents —simply having this kind of support automated would improve both the user workload and the output quality.

Moreover, the real task for the users starts now, and is also unsupported: they have to aggregate for each hotel in h the information coming from a number of around $w \times c$ client reports —e.g. determining pros and cons for each hotel according to the majority opinion of those reports, and finally deciding on the hotel that better fit their interests. Clearly, the *S&B* paradigm does not support this process and the users end up making a less informed decision. However, AI could be used to support this decision, and I'm referring not only to data mining or recommender systems, but to a reinterpretation of Case-based Reasoning that would allow us to support users in using experiential knowledge provided by a community of practice.

3 Reusing Other People's Experiences

Considering again the example of the hotel selection task, we can easily substitute the *Search & Browse* process by *Retrieve & Reuse* processes of CBR as follows: (1) the Retrieve process searches for client reports of hotels close to the declared interests of a user and selects a subset of them; then (2) the Reuse process analyzes the retrieved client reports in order to aggregate the information about pros and cons of each hotel, and finally produces a ranking of hotels taking into account both the user's interests and the pros and cons of each hotel. This mapping is sound, in the sense that both Retrieve and Reuse processes follow the ideas in [5]: (1) given a problem (a specific task to be achieved) the Retrieve process selects the subset of cases (experiential knowledge) most similar (or relevant) to that problem, while (2) the Reuse process combines, in some specific way, the (experiential) content of those retrieved cases (and possibly using some domain-specific knowledge as well) in order to achieve a solution for that problem (that specific task to be achieved).

This rather abstract mapping allows us to determine what a CBR approach to experiential reuse in the web adds to the *S&B* paradigm: the definition of a user-defined task to be achieved. Indeed, only when a problem (a specific task to be achieved) is posited then a *Retrieve & Reuse* approach can be used.

Let us return to the hotel selection example again. Clearly the kind of hotel the user is interested in depends on the type of travel: e.g., whether it's in a one-night business trip or a leisure week-long travel, the pro and con factors that are important may vary from one kind of travel to another. For instance, the significance of whether the hotel staff is categorized as friendly (in pros) or unfriendly (in cons) depends on the trip: a friendly/unfriendly staff might not important in a one-night business trip while it might be quite important on a leisure week-long travel. This matching between the hotel client reports and the user interests would be performed inside the Reuse process, e.g. preferring those hotels with a clear majority of client reports stating a friendly staff and the other factors significant for the user. This is precisely the work the human user has to do now, without any support, while examining $h \times w \times c$ client reports.

Nevertheless, there are differences from the traditional CBR approach with respect to a *Retrieve & Reuse* approach to use the experiential knowledge of other people. These differences stem from tacit hypotheses used in CBR or implicit assumptions built from practice in building CBR systems. A first implicit assumption is that the Retrieve process will select one case (or a small number of cases) on which the Reuse process will work upon. As the hotel scenario shows, this is not the best option when dealing with experiential knowledge coming from a (potentially large) number of people. In the hotel scenario the role of the Retrieve process is to select, among a huge number of client reports, a sufficient number of reports about hotels that are relevant for the specific task requested by a user.

Since the Reuse process needs to aggregate information from disparate sources in order to avoid noisy data, the sample of data has to be large enough so that aggregation methods like averages or weighted averages are meaningful. That is to say, in the hotel scenario, the role of the Retrieve process may be to select the hotels relevant for the task at hand within some given ranges, for instance, of price and location, and then gathering all their relevant client reports. Additionally, the Retrieve process could perform an additional filtering of client reports based on their age, client reputation, etc. Then, given this sizable sample of people's reports on their experiences, the Reuse process may be able to aggregate, from the evidence of disparate sources, the likelihood that one or a few hotels are the most adequate for the particular interests of a user's travel.

The robustness of using experiential knowledge originating

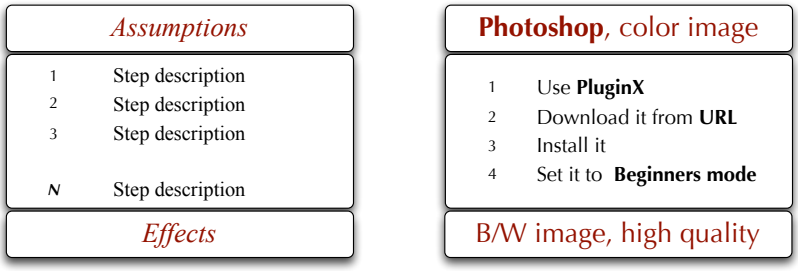

Figure 1: Semi-structured form of experience for *How-To* tasks.

from multiple sources has been studied in several scientific fields. In Machine Learning, the *ensemble effect* states that using an ensemble of learning systems reduces always the error when compared to any single learning system. The only requirements for the ensemble effect to take place is that the prediction of individual learning systems is better than random and that their errors are not correlated with one another [6]. Similar properties have been characterized in Social Choice Theory, where the Condorcet Jury Theorem provides a similar property for taking average measures like voting in a jury [8]. Communities of practice on the web have been known to show a similar effect, a fact popularized in James Surowiecki's book *The Wisdom of the Crowds* — where similar conditions are prescribed in order to insure the emerging effect of wiser decisions or predictions by aggregating information from a crowd of people.

Therefore, a challenge for applying the *Retrieve & Reuse* approach sketched here is to enlarge the core ideas of CBR, namely reasoning and learning from past experience, to a scenario where experiential knowledge comes in large numbers and originates from multiple individual sources; these issues require that we incorporate aggregation measures that obtain the desired *ensemble effect* into the *Retrieve & Reuse* processes. There are other CBR assumptions that need to be challenged to develop systems that reuse experiential knowledge on the web, and I'll summarily address a few of them later in the paper.

The next two sections deal with the semantics and structure of experiences, two challenging issues that need to be addressed in order to reuse experiential knowledge on the web.

4 Semantics and Experience

There are two main approaches to web semantics, namely (1) the top-down approach and (2) the bottom-up approach. Both approaches are suitable to be used in a CBR-like approach to reasoning from experiential knowledge on the web:

1. the semantic web uses ontologies expressed in description logics (specifically the OWL language[1]), which is compatible with the research line on knowledge-intensive CBR systems development using description logics;
2. Textual CBR [9] has worked on a bottom-up and hybrid approaches to semantics in cases expressed as text, which is compatible with the current research goals of folksonomies and web text mining.

Since both semantic approaches, or a combination of top-down and bottom-up approaches, are suitable for a CBR-like approach to reuse web experiential knowledge, the challenges are basically the same of any other web-based system developed using Artificial Intelligence techniques. Moreover, since the applicability and utility of either semantic approach may vary for different application domains, it is an empirical issue to determine when and how these semantic approaches will be useful. In this sense, the approach to reuse web experiential knowledge I'm sketching here would be neutral on these semantic debates, trying to find a suitable trade-off for a particular application domain and to keep up with the new developments in web semantics.

However, the focus on user-contributed experiential knowledge poses some practical constraints. The first one is that the form in which experiences are expressed has to feel easy and natural to the people integrating a community of practice; otherwise, very few content will be contributed by them. This constraint seems to bias experience representation towards text-based content, but this again depends on the specific community of practice and their particular application domain. Ontology-based approaches require a highly structured representation of content, but technical communities of practice (e.g. medicine, engineering) may accept this approach if they find useful the services provided.

For other users in general a text-based approach seems more suitable, but the text needs not be completely free: we should be able to provide semi-structured templates where the users can textually enter their experiences. This idea leads us to the second challenging issue I'd like to discuss: the structure of experiential knowledge.

5 Forms of Experience

An important issue about experiential knowledge on the web, as mentioned before in section 2, is that *cases* as such are not already present on the web. Recalling the hotel selection example we can see that there is no collection of cases of the form $(situation, outcome)$; instead we had *records of individual experiences* in the form of client reports. That is to say, we have a collection of *situations* without the *outcome*. For the task at hand, selecting a hotel, it is tempting to conceive of the *outcome* as the selected hotel: this is true for the system outputting a recommended hotel but it is not applicable to the client reports. A case in the standard sense would be a pair where a *situation* would describe the interests, preferences and constraints of a user and an *outcome* would be a hotel satisfying (most of) them. However, the client reports do not directly specify the persons interests, preferences and constraints; it is an *account*

[1] See an overview of OWL at http://www.w3.org/TR/owl-features.

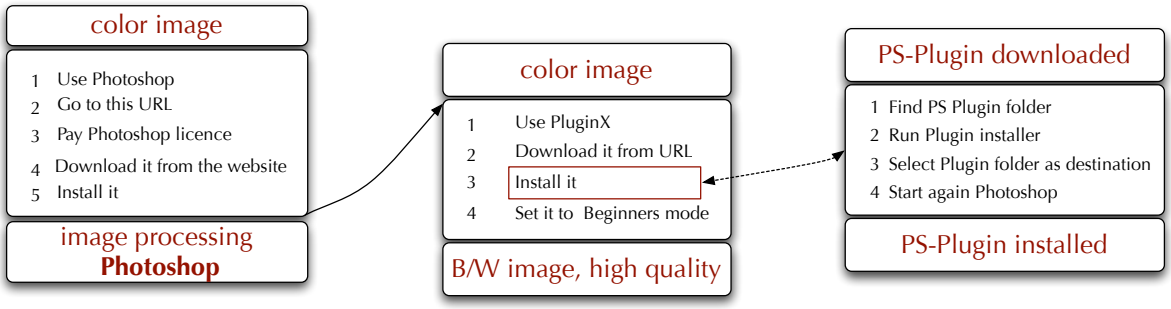

Figure 2: Reusing experiential knowledge by combining *How-To*s.

of an experience that may have been positive or negative (or something in between). Nevertheless, as I tried to show in the hotel scenario, some of this information is implicit and can be extracted: the analysis of the client records in terms of pro and con factors is a way to uncover the tacit interests and preferences of the users giving an account of their experiences.

There may be other ways to uncover the important factors in experiential accounts, since this *pros and cons analysis* is just an example. This leads us to the core issue in this approach: How many different *forms of experience* are there? Do we need to develop a new form or structure of experience for every new application domain? This circumstance could make impractical to apply this approach on the web at large. If not, is there a small collection of *forms of experience* that could be characterized and reused? Which are they and how to find them? I really have no answer in advance, since it is essentially an empirical matter to be settled only after trying to develop systems that reuse experiential knowledge on the web. I may suggest some hypotheses, though, as to how to proceed for developing systems that reuse experiential knowledge on the web.

The first hypothesis is trying to characterize a *form of experience* for each class of task commonly known in CBR systems: e.g. classification, regression, planning and configuration. These tasks are classically differentiated by the *form of the solution*: single-item solutions for classification, sequences for planning and complex graph structures for configuration.

Seems reasonable to assume that the differences on the solution structures of these tasks imply that the corresponding experiential knowledge would also be structurally different. However, each class of task may have a sufficient degree of internal coherence to allow the development of experience-reuse systems applicable inside a class of tasks. For instance, the method of analyzing pros and cons in hotel client reports could be used, in principle, to other application domains whose task is a form of classification: e.g. selecting a digital camera, or selecting a B/W plugin for Photoshop. Moreover, other different techniques to reuse experiential knowledge for classification tasks could be developed. Again, only empirical evidence will determine whether the hypothesis suggested here is correct or not.

As a further example, let us consider *planning* in the context of experiential knowledge on the web. Since a plan is just a way to achieve some effect or goal performing a series of steps, it is easy to see that they are pervasive on the web, although they are not called "plans": sometimes they are called *How-To*s, but most times they are just descriptions of how to do something in a few steps. Forums are websites where a large number of *How-To*s can be found. For instance, forums store numerous records of "question and answer" pairs that may be interpreted as problems and their solution-plan. A specific forum, like one devoted to digital photography, has both a community of practice and a shared vocabulary of terms (e.g. B/W image), verbs (actions) and proper nouns (e.g. "Photoshop"). A typical scenario is when a user asks how to perform some effect on an image and the answer is a plan of the form "*assuming you have Photoshop, you should download this PluginX from this URL, install it and then set it up in the beginner mode; you'll obtain B/W image with an already good quality*." Forums organize this content in a structure based on questions and answers, and thus we are expected to use *Search & Browse* to find and reuse this experience. Capturing this experiential knowledge from free text using NLP techniques may be feasible, but also costly.

Another option is to design some semi-structured representation for this form of experience that, if stored on a website (substituting the questions and answers structure), would facilitate the analysis, retrieval and reuse of *How-To* knowledge. As a further elaboration of this scenario, consider a possible semi-structured template for *How-To* experiential knowledge as that shown in Figure 1. The semi-structured template clearly separates plan preconditions (*Assumptions*), plan goals (*Effects*) and each one of the *Steps* or actions of the plan. Albeit text processing is still necessary, the previous example on *PluginX* shown at the right hand side of Fig. 1 is now more easily analyzed for the purposes of its reuse. Recall that the final user will be able to understand and perform this *How-To*, we only need enough structure to (1) allow a user to express the problem she wants to solve, e.g. "I have Photoshop and I want to transform a color image into a high quality B/W image," and (2) *recognize* that the *How-To* in Fig. 1 is a way to solve that problem.

Moreover, accessing a large repository of *How-To*s would also enable forms of *case-based plan adaptation*. Consider the situation where the user has the same goal but she does not have Photoshop. Figure 2 shows how a new plan can be generated by concatenating two *How-To*s: the first plan is one for acquiring Photoshop, while the second plan is that of Fig. 1, that uses a Plugin to achieve a B/W image. Since the effect of the first *How-To* is having Photoshop, now the second plan can be safely used since the Photoshop assumption is now satisfied. Another form of adaptation is expanding a step, that is in fact a subplan, into its component sub-steps. Fig. 2 shows that Step 3 "Install Plugin" is not an atomic action, but may be expanded

into 4 steps because another *How-To* can be found whose goal is to install Photoshop plugins. This form of plan adaptation should be feasible whenever we have a large repository of plan-like *How-To*s, and it is also akin to the currently fashionable idea of "mash-ups"[2] on the web. Moreover, a similar technique for case-based plan adaptation has been recently suggested for real-time planning when a large repository of plans is available —i.e. new plans are generated by combining existing plans [7].

Planning by reusing, adapting, and combining user-contributed plans can be applied to a large number domains —from *How-To*s and other forms of methods to itineraries and route sheets— as long as a large repository of "action sequences" is available. The fact that these plans have already been tried out by someone, and were successful, gives us a further hindsight. Moreover, the *ensemble effect* can be used on a large repository: when several methods or plans are found to achieve the same result then aggregation techniques like voting can be used to determine the one that is considered more reliable (at least inside a community of practice).

Therefore, the hypothesis put forward in this section is that several forms of experience could be defined with sufficient internal coherence so that it is feasible and practical to build systems for reusing recorded experiences of other people. The next section discusses the overall organization of such systems.

6 The EDIR Cycle

These ideas can be integrated into a process model called the EDIR cycle, shown in Fig. 3. The EDIR cycle consists of four processes: *Express*, *Discover*, *Interpret*, and *Reuse*. They should be understood as interrelated processes, not as sequential phases or causally dependent steps.

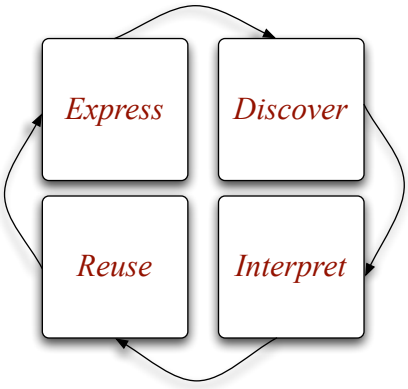

Figure 3: The EDIR cycle for reusing other people's experiences.

1) Express: This process addresses the different ways in which experiences can be expressed by a contributing user inside a community of practice. Free, semi-structured and ontology-based templates for specific forms of experience and application domains need to be developed and tested; the research goal is finding a trade-off that (a) allows sufficient structuring of the expressed experiences for automated analysis and (b) feels as a natural and unobtrusive way to express experiences for the users in a community of practice.

2) Discovery: This process addresses the different ways in which specific experiential content is recognised and retrieved as possibly relevant to a given query posed by a system user. The research goal is determining how to extend existing retrieval techniques to work on experiential content integrating semantic web and/or bottom-up semantic analysis. The conditions under which the Discovery process has to work requires a fast and possibly shallow analysis of large quantities experiential reports; the expected output is a moderately-sized collection of experiences that are (likely) relevant to the current query.

3) Interpret: This process addresses the different ways to build semantic interpretations of the discovered experiences. The semantics are only assumed to hold inside a community of practice. This interpretations can be understood as a more in-depth analysis of the experiences selected by the Discovery process using the semantic model of the community of practice and the available domain knowledge. Several transformations, among others, make part of the Interpret process: (a) eliminating a subset of discovered experiences as non-relevant; (b) transforming discovered experiences into a new canonical representation; (c) translating discovered experiences into a canonical vocabulary coherent with the one used to build the queries of the final users. These or other transformations may be used in a particular system, but the final outcome is a collection of canonical experience descriptions supplied to the Reuse process.

4) Reuse This process addresses the different ways in which the experiential content provided by the Interpret process is used to achieve the goals of a user as described in a particular query. Reuse techniques from CBR may need to be revised or extended in order to be applicable in this context (e.g. case-based adaptation) but also new methods that rely on the nature of large repositories of human experience should be developed (e.g. methods based on the *ensemble effect*). Moreover there may be different modalities of experience reuse: from automated experience reuse (yielding to the user the complete solution provided by reusing experiential knowledge) to the opposite extreme where the user receives directly a small ranked selection of relevant and reliable experiences. Intermediate modalities may perform part of the reuse process automatically while supporting the user in reuse finalization.

The EDIR cycle is a process model, so the relationship of the four processes needs not be sequential in any implementation of the model. Clearly, during an interaction with the final user to elucidate the requirements of her enquiry several discovery and interpretation processes may be launched and their results used to help the user narrow her options or widen her constraints.

Finally, let's compare the EDIR approach with the current *Search & Browse* approach. The main difference is that the EDIR approach requires a *query*: a description of the kind of result needed by the user —a definition of the *problem to be solved*. Only with a query it is possible to *reuse* experiences, since the Reuse process employs methods that try to satisfy the requirements of the current query using a collection of selected experiences. A second but important difference concerns the form and organization of content. The *Search & Browse* approach assumes the existence of just hyperlinked documents: even when some structure is present (e.g. question-answer structures in forums) —this structure is not exploited to improve the

[2] A *mash-up* refers to a web application that combines data from more than one source into a new integrated service.

results. The EDIR approach intends to characterize a particular kind of content, experiential knowledge, and it is thus concerned on how to adequately express, represent, organize, analyze, and retrieve this content.

7 Discussion

This paper is about current and future challenges on reasoning from experience. As such, I've dispensed with some formalities of the typical structure and content of scientific papers. There is no state of the art on the semantic web, natural language processing and text mining applied to the web, but this is because they are orthogonal to the purposes of this paper: they can be applied, and they mostly are applied inside the *S&B* approach; but they could also be used in an EDIR approach to experiential knowledge reuse.

The purpose of the paper is not presenting a specific contribution but a series of ideas intended to set in motion a discussion on how to apply AI techniques, in general, and CBR techniques, in particular, to the ever-growing World-Wide Web. The main idea presented for discussion is whether there is, or is useful to conceive of, experiential knowledge on the web. I've not given a formal definition of *experience*, but my use of the term is close to the common sense meaning, and the examples presented should be enough for grasping its meaning. I found worthy of attention that trying to apply CBR ideas like reuse of past experience to the web, I've been compelled to abandon a straightforward notion of *case*. CBR can not be directly applied to the web, since there are no ready-made cases preexisting on the web. However, if we understand CBR as ways of reusing past experience, we can generalize these core ideas in CBR and investigate how could we possibly reuse the experiences that people are already providing on the web.

The EDIR cycle is simply a way to organize the different issues and challenges to be addressed in developing systems for reusing experiential knowledge on the web. As such, it is a tool for helping to start thinking and debating about how to build systems that reuse experience, and should probably be left aside when enough progress is made that shows how to proceed.

The bottom-line is that we commonly reuse other people's experiences in our daily activities. Since large amounts of these experiences are now recorded on the web, we try to (re)use them in order to take more informed actions in the world. However, we end up with tons of "content items," but no support to filter, analyze, and reuse them. The *S&B* paradigm was built upon a platform for "web content browsing," with the idea of helping people to find an item or a few items. This is based on the metaphor of finding a book in a digital library, and is therefore based on the idea of searching over documents. Basically, this comes down to offering support for tasks of *information access*. We need a new approach, like the one I'm proposing here, for supporting tasks where people use web content for taking actions on the world at large; tasks in which information access is only a part of the whole process. The approach proposed here is that a specific form of content, experiential knowledge, can be fruitfully used to provide support for users to take more informed decisions, and AI can provide the techniques that will eventually support reusing other people's experiences.

Acknowledgements This paper has benefited from long, engaging discussions of the author with many people, specially Agnar Aamodt, Josep-Lluís Arcos, Paolo Avesani, Klaus-Dieter Althoff, Ralph Bergman, Susan Craw, and Nirmalie Wiratunga; all errors and/or misconceptions, nevetheless, are the author's sole responsibility.

This research has been partially supported by the Project MID-CBR (TIN2006-15140-C03-01).

References

[1] Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. Scientific American Magazine (2001).

[2] Shadbolt, N., Hall, W., Berners-Lee, T.: The Semantic Web Revisited. IEEE Intelligent Systems **21**(3) (2006) 96–101.

[3] Davies, J.: Semantic Web Technologies: Trends and Research in Ontology-based Systems. Wiley (2006).

[4] Benkler, Y.: The Wealth of Networks. How Social Production Transforms Markets and Freedom. Yale University Press (2006).

[5] Aamodt, A., Plaza, E.: Case-based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. Artificial Intelligence Communications **7**(1) (1994) 39–59.

[6] Perrone, M.P., Cooper, L.N.: When Networks Disagree: Ensemble Methods for Hybrid Neural Networks. In: Artificial Neural Networks for Speech and Vision, 126–142. Chapman-Hall (1993).

[7] Sugandh, N., Ontañón, S., Ram, A.: Real-Time Plan Adaptation for Case-Based Planning in Real-Time Strategy Games. In: K.D. Althoff, R. Bergman, M. Minor, and A. Hanft (Eds), Advances in Case-Based Reasoning, ECCBR 2008, LNCS Vol. **5239** (2008) 533-547. Springer Verlag.

[8] Sunstein, C.R.: Group Judgments: Deliberation, Statistical Means, and Information Markets. New York University Law Review **80** (2005) 962–1049.

[9] Weber, R.O., Ashley, K.D., Brüninghaus, S.: Textual Case-based Reasoning. Knowledge Engineering Review **20** (2005) 255–260.

Contact

Prof. Enric Plaza
IIIA, Artificial Intelligence Research Institute
CSIC, Spanish Council for Scientific Research
IIIA-CSIC, Campus UAB, 08193 Bellaterra, Catalonia (Spain)
Tel.: +34 935 809 570 ; Email: enric@iiia.csic.es
Web: http://www.iiia.csic.es/∼enric

Bild

Enric Plaza is Research Professor at the IIIA-CSIC and Head of the Learning Systems Department (LSD). His research has spanned different areas of Artificial Intelligence, including knowledge acquisition and validation for expert systems, case-based reasoning, machine learning, and multiagent systems. He is author or co-author of more than 150 scientific papers in journals and conferences and has worked in more than 20 AI research projects. He is an ECCAI Fellow and was the first President of the Catalan Association for Artificial Intelligence (ACIA) from 1994 to 1998 as well as Director of the ACIA Bulletin from 1998 to 2002.