

# On Reusing Other People's Experiences

Eric 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 im-

plies 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 fu-