Universidade Federal Fluminense

Bruno Norberto da Silva

KA-CAPTCHA: An Opportunity for Knowledge Acquisition on the Web

Niterói
2007

KA-CAPTCHA: An Opportunity for Knowledge Acquisition on the Web

BRUNO NORBERTO DA SILVA

Dissertação de Mestrado submetida ao Programa de Pós-Graduação em Computação da Universidade Federal Fluminense como requisito parcial para a obtenção do título de Mestre.
Área de concentração: Otimização Combinatória e Inteligência Artificial.

Orientador:
Ana Cristina Bicharra Garcia

Universidade Federal Fluminense

Niterói
2007

KA-CAPTCHA: An Opportunity for Knowledge Acquisition on the Web

BRUNO NORBERTO DA SILVA

Dissertação de Mestrado submetida ao Programa de Pós-Graduação em Computação da Universidade Federal Fluminense como requisito parcial para a obtenção do título de Mestre.
Área de concentração: Otimização Combinatória e Inteligência Artificial.

Aprovada por:

Prof. Ana Cristina Bicharra Garcia / IC-UFF (Presidente)

Prof. Bianca Zadrozny / IC-UFF

Prof. Clarisse Sieckenius de Souza / DI-PUC-Rio

Niterói, Outubro de 2007

To my very special partner who made this all possible and enjoyable.
K

Acknowledgements

I thank my advisor Ana Cristina for all the support and guidance during these years. I believe this work could not have been completed with your sound advice and thoughtful criticism. I am very thankful for your faith on my potential and for acting precisely on (and only on) my many flaws. The freedom to let me find and develop my own thesis topic was paramount for a stubborn student like me to succeed in this journey.

I must also thank a number of friends and colleagues who accompanied me during these years. The ADDLabs team (including former members) definitely helped me not only with supportive criticism, but also with the dearest friendship. In particular, I must thank Adriana Franco and Antonio Segaloto for all the logistical help, without which this work would never have reached this final state.

I must express gratitude to my thesis committee members who contributed with valuable comments to improve the quality of this work. I also thank Professors Eduardo Laber and, especially, Vinicius Carrasco, who were extremely kind to let me sit and watch their courses at PUC-Rio as a special student.

Additionally, I have been very fortunately to collect a number of very valuable friends during all stages of my life. Ever since primary school until college and later, I know there are a number of people I can count on in the toughest moments.

Finally, the one who made this all worthwhile. She, who *eyes me like a Pisces when I am weak*, who lately has demonstrated the greatest will to overcome every barrier and to support both my professional and personal plans in life, who continuously amazes me by how much she resembles myself, but also how much she completes me. Forever be this wonderful and spontaneous little person who is truly my soul mate. I'm your greatest fan Lucia Helena.

Resumo da Tese apresentada à UFF como parte dos requisitos necessários para a obtenção do grau de Mestre em Computação (M.Sc.)

KA-CAPTCHA: An Opportunity for Knowledge Acquisition on the Web

Bruno Norberto da Silva

Orientadora: Ana Cristina Bicharra Garcia

Programa: Pós-Graduação em Computação

Qualquer usuário da Web é um voluntário em potencial para um procedimento de aquisição de conhecimento, contudo um grande desafio é convencer estes voluntários a colaborar altruisticamente, isto é, convence-los a doar seu tempo sem que recebam qualquer benefício imediato. De maneira geral, a tarefa de alinhar interesses individuais com os sociais requer incentivos artificiais por parte de um projetista. Neste trabalho, apresentamos uma nova ferramenta para aquisição de conhecimento da Web que por estar embutida em aplicativos indispensáveis para usuários, induz comportamentos cooperativos na saga de elicitação de semântica, especialmente para imagens. Nossa proposta se baseia na extensão dos mecanismos CAPTCHA para permitir aquisição de conhecimento de maneira invisível durante tarefas rotineiras da Web. Conseqüentemente, esta funcionalidade se distingue das abordagens anteriores por não exigir nenhum esforço extra por parte dos voluntários, além do esforço já solicitado pelo CAPTCHA.

Duas aplicações deste mecanismo são apresentadas neste trabalho. Na primeira, coletamos conhecimento sobre o domínio de rotulamento de imagens, em conjunto com resultados experimentais que sugerem sua viabilidade. Outra aplicação é apresentada visando extrair transcrições de materiais com conteúdo auditório presentes na Web.

Abstract of Thesis proposed to UFF in partial fulfillment of the requirements for the degree of Master of Science (M.Sc.)

KA-CAPTCHA: An Opportunity for Knowledge Acquisition on the Web

Bruno Norberto da Silva

Advisor: Ana Cristina Bicharra Garcia

Program: Pós-Graduação em Computação

Every Web user is a potential knowledge contributor, but it is a challenge to make them devote their time contributing to some purpose. In order to align individual with social interests, I will present an extension of the CAPTCHA Web resource protection application to embed knowledge elicitation within the users' main task of accessing a Web resource. Consequently, unlike previous knowledge acquisition approaches, no extra effort is expected from users since they are already willing to use a CAPTCHA to perform some particular task.

We present two applications of this mechanism: one that collects image labels from Web users, where experimental results suggest the feasibility of this approach, and another that collects textual transcriptions to auditory media.

Key-words

1. Inteligência artificial

2. Aquisição de conhecimento (Sistemas especialistas)

3. Web semântica

List of Figures

List of Tables

Summary

Chapter 1. Introduction

Knowledge acquisition is extremely important for the development of intelligent systems and its contributions spans fields as diverse as medicine, engineering and financial economics. However, the task of eliciting knowledge from a large population presents a great challenge to researchers. Major difficulties like combining conflicting points of view and motivating self-interested contributors arise when performing a large-scale knowledge acquisition task.

1.1 Unsolved Artificial Intelligence as a Security Mechanism – Motivating scenario

In his seminal work Computing Machinery and Intelligence (Turing, 1950), Alan Turing introduced his Turing Test, a method for measuring progress in the task of building intelligent machines. This test takes places whenever a human referee tries to distinguish a machine from a human being by exchanging nothing but textual messages with both agents. Ideally, the interaction between the referee and subjects should be long enough to allow a conscious decision by the referee, and thus artificial intelligence would be achieved whenever the referee could not distinguish the agents.

Since its creation, the Turing Test has also been used in applications other than measuring machine intelligence. For instance, the need for protecting computational resources from massive attacks and the need to impede non-human access to these resources have motivated the development of Turing-based methods for security problems. Additionally, a large volume of transactions involving these same resources has pushed the demand for methods to legitimate the human side of this computation in an automated way.

A major problem arises in these new applications of the Turing Test due to the significantly short interaction between the referee and the malicious computational agents. Because of adverse and restricted conditions, most of the communication protocols underlying these computations cannot afford to enable a rich communication between those involved in this distributed procedure. Consequently, machines might be able to pass as human agents not due to a behavior to that of a human being, but because of the diminished realm of communication required for the interaction and the small period of time the agents interact. The problem leads to a demand for an automated mechanism that is able to tell humans and computers apart, with as restricted an interaction as possible.

The most successful effort in addressing this issue is a mechanism known as CAPTCHA, which stands for *Completely Automated Public Turing Test to Tell Humans and Computers apart* (von Ahn et Al. 2004). A CAPTCHA is a test that challenges an agent to solve an open Artificial Intelligence problem. Consequently, humans are supposed to easily pass a CAPTCHA test but current computers should find it virtually impossible to succeed.

Examples of successful applications of CAPTCHA include access to public databases (Plataforma Lattes), which might be flooded by denial of service attacks; email account registrations (Yahoo! Mail), which cannot be freely accessed due to risk that spammers will be able to spread even more unsolicited email messages; and SMS message sending services (Mundo Oi), whose unrestricted access could extend the distribution of unsolicited messages to mobile telephone devices.

The fact that a CAPTCHA test is based on an unsolved Artificial Intelligence problem fits well with its required property of distinguishing humans from computers, and popular CAPTCHAs have relied heavily on the problems of distorted character and speech recognition. Regrettably, since CAPTCHAs are employed as protectors of several important resources on the World Wide Web, the inherent incentive for hackers to develop novel computational tools that break a CAPTCHA test has been a persistent concern (Mori and Malik, 2003). This demands a constant effort in the design of either very difficult instances of CAPTCHAs, or totally different new tests based on subfields of Artificial Intelligence where there is still a significant gap between the state-of-the-art computational agents and human performance.

Unfortunately, because of the restricted environment where CAPTCHAs are to be employed, an instance of the test that is difficult enough for computers is frequently also an obstacle (if not a barrier) for legitimate users as well (Chellapilla et al. 2005). Therefore, a CAPTCHA based on different AI problems would be of a greater utility.

1.2 Unsolved Artificial Intelligence as a downside of the Semantic Web – A missed opportunity

Once focusing on the design of new CAPTCHA tests, one can take advantage of the fact that the set of open problems in Artificial Intelligence that are available for the design of new CAPTCHAs is huge. This provides great freedom for designers of new CAPTCHA tests, but it naturally matches several unrealized efforts that aim at enhancing human experience with computers. Many of the AI fields that currently serve as potential grounds for the development of new CAPTCHA tests are also the object of massive efforts at enhancing the potential of the World Wide Web. One appealing archetype of these efforts is currently the Semantic Web, which envisions a system where intelligent agents are to replace the Web's existing reactive architecture and will be able to actually understand and reason about the Web's content.

Since current reasoning methods still fail to satisfy the ambitious goals of the Semantic Web, several endeavors have attempted to make up for the unskilled machine with knowledge collected from humans. These efforts can be classified into two categories: 1) explicit knowledge acquisition, in which information is usually collected from a group of trained individuals (Lenat, 1995; Feigenbaum, 1991; Miller, 1995; Yokoi, 1995) hired to include semantics to data, and implicit knowledge acquisition, constituted basically of volunteers who provide knowledge as a side effect of some other task (von Ahn, 2006; Open Mind; Chklovski, 2003). Each of these approaches has their pros and cons, however they all underperform as they fail to elicit as much knowledge as people would be able to report. Possibly the most compelling issue in eliciting knowledge from humans is the difficulty of aligning individual interests on how to allocate their time with the interests of any knowledge acquisition mechanism that needs people to continuous feed their data with semantics, since people often finds the task of contributing to some common good to be boring and exhaustive, and will not immediately benefit from the contribution.

Recent investigations attempted to overcome this inefficiency of knowledge acquisition mechanisms by working out the psychological aspect of the interaction (Rashid et al. 2006; von Ahn and Dabbish, 2004; von Ahn et al, 2006). However, all previous approaches to this problem have not fulfilled one or more of the following desiderata we consider important for a knowledge acquisition instrument:

- No specific training → it should be possible and straightforward for an altruistic volunteer to collaborate with an existing knowledge repository.
- Incentive compatibility → in order to perform an efficient knowledge extraction, it is essential for a mechanism to align the incentives of the volunteering party with those of the knowledge collector. A mechanism should not depend on altruism to succeed.
- Generality → it is highly desirable for mechanisms to gather knowledge from a wide and diverse range of people, and not restrict its application to small niches.
- Vanishing invasiveness → to facilitate the collection of a significant amount of contributions, a knowledge acquisition mechanism should introduce as little overhead as possible into its user's agenda.

Therefore, it is clear the need for novel methods of knowledge acquisition in order to satisfy requirements such as those of projects like the Semantic Web. As soon as we detect this demand, we are left with a quest for a model of knowledge acquisition that satisfies the properties previously mentioned.

1.3 Recycling a wasted effort – the CAPTCHA Mechanism as an Opportunity for Knowledge Acquisition on the Web

A careful inspection of the CAPTCHA mechanism that was previously described reveals that an opportunity to satisfy these highlighted desiderata might exist if one could embed knowledge acquisition into CAPTCHA's interaction. Because CAPTCHAs present a test that asks users the answer to an instance of an unsolved Artificial Intelligence problem, one could conceive a test which asks its users for a solution to some useful unsolved AI problem, instead of some throwaway question as done by current CAPTCHAs on the Web. This scenario offers a good opportunity for knowledge acquisition because:

- CAPTCHAs have become very popular on the Web due to their simplicity and efficiency. Hence, their model of interaction is already familiar to almost everyone who interacts with protected Web resources.

- CAPTCHA tests represent a unique moment when both sides of the computation have their interests aligned: a test taker is interested in truthfully answering a query because she wants to access some Web resource. And the Web resource administrator might be interested in querying a test taker because he is interested in collecting knowledge from a diverse pool of subjects.
- A CAPTCHA test is independent of the Web resource it is protecting. Consequently, the very same mechanism might be employed to protect many different resources, and consequently reach a wide and diverse audience on the Web.
- It is widely accepted that Web resources cannot persist without some form of protection, and CAPTCHA has become a very successful security provider. Since the CAPTCHA interaction is decidedly present on people's experience on the Web, a knowledge acquisition mechanism that takes advantage of this existing burden would pose no additional effort on either sides of the task. The effort this approach expects from users is already being exerted with the current CAPTCHA design, but only in a fruitless way.

1.4 Illustrative example of using a CAPTCHA for performing Knowledge Acquisition

The problem of assigning valid textual labels to images is of extreme importance to Artificial Intelligence, and is also representative of the number of issues resolved using knowledge acquisition. Since current advancements in computer vision fail to satisfy the needs of applications that make use of pictorial knowledge on the Web, the elicitation of this information from humans has become a promising approach to this problem.

Because the task of describing the content of images is an unsolved Artificial Intelligence problem, it qualifies to be used by a CAPTCHA test. And indeed, the Pix CAPTCHA (The CAPTCHA Project) takes advantage of this lack of advancement in computer vision techniques to guarantee security for the Web. This CAPTCHA maintains a database with a large collection of images, each associated with one or more textual labels. A test is generated by selecting one label at random, picking four images related to that label, distorting them at random, and them asking the an agent to indicate, from a list of all existing labels, which one relates to all four images (Figure 1).

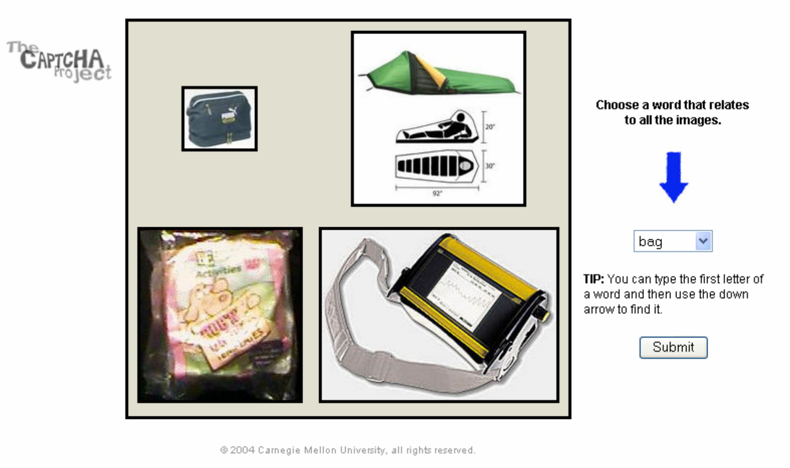

Choose a word that relates to all the images.

TIP: You can type the first letter of a word and then use the down arrow to find it.

© 2004 Carnegie Mellon University, all rights reserved.

Figure 1. Pix, a CAPTCHA in which users are expected to identify a common feature between four images and point it out from a list of labels.

The intuition behind this CAPTCHA is that computers are not able to identify image's content and then reason to find a common aspect between the four images and relate that aspect to a textual label. However, since the Pix database is public and of very limited size, the images must be displayed with such distortion that a computational agent will not be able to relate the presented pictures to one of the public images (because if it could, an automated solution to this test would be trivial). An implementation of this test with such level of distortion might be too difficult for humans to succeed as well, making Pix a very difficult instance of CAPTCHA.

However, this CAPTCHA can serve as a guiding inspiration for our purpose of creating a CAPTCHA that performs knowledge acquisition. Pix works using a small knowledge base of labels and images, and accomplishes its security guarantee by using the human processing power to assign labels to images. Yet, it is clear that the task performed by humans in this task results in an answer that the Pix CAPTCHA is already aware of.

Nevertheless, as was previously illustrated, the problem of assigning valid labels to images is a difficult and still unsolved problem of Artificial Intelligence. Consequently, if we adapted Pix to assign new labels to the images present in its knowledge base (but never at the expense of its security primitives) than a significant tool for pictorial knowledge acquisition could be developed. The design of such mechanism will be presented in Chapter 4.

1.5 Problem statement

In this dissertation, we address the problem of automatically acquiring knowledge for building the Semantic web. More specifically, we focus on the problem of designing a mechanism for knowledge acquisition that is able to collect semantic information from humans.

1.6 Working hypothesis

The CAPTCHA scenario poses a good opportunity for knowledge acquisition because people are facing hard tests and willingly giving their best shots to solve them. The problem lies in designing a CAPTCHA test that impedes computational agents to access unauthorized resources at the same that elicits some knowledge from users.

One universal assumption regarding an agent's behavior is individual rationality (Mas-Colell et al, 1995, Chapter 1), which states that individuals opt for strategies in order to maximize their own utilities, in detriment of a socially optimal policy. This applied to the scenario considered in this thesis reveals a significant flaw in several knowledge acquisition mechanisms in the literature. Most of them ignore the individual rationality property and depends heavily on altruist to succeed (Chklovski, 2003; Singh, 2002). Other mechanisms attempt to attract the interests of knowledge contributors by psychological means (e.g. entertainment), however, as we argued above, this is in detriment of the knowledge collected from volunteers, who bears the risk of ending up too biased.

After considering this argument, and along with the observation that every knowledge acquisition mechanism will require some cognitive effort from the individual who is collaborating with the process, we suppose that knowledge elicitation improves when included within another interactive process, which serves as ground for our solution.

We propose KA-CAPTCHA, which stands for *Knowledge Acquisition CAPTCHA,* as a novel approach to knowledge acquisition that is embedded into the existing CAPTCHA security mechanism that is already widely disseminated over the Web. Our purpose is to take advantage of the time window users must devote to a CAPTCHA procedure to prove being an authentic user and reach their objectives.

The CAPTCHA test should be elaborated to incorporate the users' input as semantics to entities available on the Web, where different perspectives on existing information are incrementally included. Since users are highly motivated to pass the CAPTCHA test it is feasible to believe that the acquired information is true.

KA-CAPTCHA can be considered a designed mechanism to align the users' need to obtain a service to the Web designers' needs to elicit semantics to better index information and consequently to improve information retrieval on the Web. Hence, our contribution adds to the state-of-the-art by providing a mechanism that is general, incentive-compatible, with a potentially large and diverse audience and which introduces little overhead to the Web experience. Additionally, the KA-CAPTCHA allows for the design of specific instances that collects knowledge from a diversified body of users, and therefore to collect conflicting points of views. However, this is a burden of a specific KA-CAPTCHA designer.

We claim that 1) we can design a CAPTCHA that imposes the same cognitive cost to users, but that utilize the users' effort to gather useful information, and 2) the elicited information improves precision and recall in the Web.

1.7 Underlying premises

The scope of this thesis restricts our contribution by imposing two premises to our solution. A profound study of the consequences of dropping any of such premises is a possible direction for future work, as indicated in Chapter 5.

One basic premise is that the CAPTCHAs presented in this thesis that provide knowledge acquisition are not significantly more difficult than traditional existing CAPTCHAs. More concretely, we assume that the difference in difficulty (if any) is not large enough to detriment its application on an environment like the Web where designers compete for user's attention.

Secondly, we will assume that cultural differences among our CAPTCHA's population will not be in detriment of our method's success in distinguishing humans from machines. We believe that pragmatic means (e.g. HTTP cookies) suffice for addressing this issue, although clearly a solution that dispenses such resources are ideal and are object of future research.

1.8 Research method and evaluation metrics

We implemented the KA-CAPTCHA method for the domain of picture description elicitation. Controlled experiments were developed in which 143 university students from 5 different undergraduate and graduate courses used KA-CAPTCHA to have access to their grades. Initially, we measured the number of CAPTCHA trials they needed for reaching the information. Then we measured the precision and recall of our method and compared it to Google's image search. Also, we demonstrate how our solution can be configured for any level of security against random attacks from computer agents.

1.9 Dissertation outline

In the next Chapter, we will review previous efforts on knowledge acquisition, knowledge acquisition from the Web and the existing CAPTCHA mechanism. Chapter 3 introduces our KA-CAPTCHA approach and some of its applications while Chapter 4 presents an empirical evaluation of these applications. Chapter 5 ends this thesis with conclusions and directions for future research.

Chapter 2. Relevant previous work

The field of knowledge acquisition has a rich history of successful efforts. This chapter briefly reviews some of these past endeavors, with a bias towards the literature that contributed to the formalization of our work. In addition, we review the CAPTCHA literature, which serves as grounds to our contribution to the field of knowledge acquisition.

2.1 Knowledge Acquisition

When discussing the work of Alan Turing cited in the beginning of the previous Chapter, Edward Feigenbaum (Feigenbaum, 1996) argues for the importance of knowledge acquisition in the quest for systems that demonstrate human-like intelligence. He claims that the tremendous success of knowledge-based systems (Buchanan and Shortliffe, 1984; Lederberg 1987) led to the belief that is was the specific knowledge contained in a system that counts the most toward its intelligent understanding and behavior. A representative argument of this *knowledge-is-power* hypothesis (Feigenbaum, 1991) states that "*reasoning processes of an intelligent system, being general and therefore weak, are not the source of power that leads to high levels of competence in behavior. (...) In the absence of knowledge, reasoning won't help*" (Kurzweil, 1990, *Knowledge Processing –From File Servers to Knowledge Servers, pp. 1*).

Perhaps the earliest hint that a knowledge base of data collected from humans was to play a key part in the development of intelligent systems was the DENDRAL project (Lederberg, 1987). DENDRAL was a system comprised of different components that enabled the formulation and evaluation of hypotheses that led to the identification of organic chemical molecules. However, the distinguishing element of this system was its knowledge base, which contained the representation of specific knowledge that was relevant to the domain of organic molecules. This information was codified into the system by knowledge engineers who organized all data with the help of professional specialists. The success of the DENDRAL system led to the belief that the development of a large knowledge base could be a strong indication of the success of some intelligent systems.

A subsequent attempt strongly backed-up the knowledge-is-power hypothesis. The MYCIN program (Buchanan and Shortliffe, 1984) was an expert system developed in the early 1970's which achieved a remarkable success at diagnosing infectious diseases and recommending effective antibiotics. Again, the core element of the MYCIN architecture was a knowledge base comprised of about 500 domain specific inference rules, which combined with a very simple backward chaining reasoning method to achieve a rate of success that was greater than that of most physicians who were not specialists in diagnosing infectious diseases.

These attempts, together with other subsequent endeavors, have demonstrated the practical importance of expert systems in particular, and Artificial Intelligence in general. Applications of these systems reached fields as diverse as medicine, engineering, financial services and biology. This tremendous success of knowledge-based systems led to the belief that is was the specific knowledge contained in a system that counts the most toward its intelligent understanding and behavior.

Currently, many different areas of Artificial Intelligence make extensive use of knowledge collected from humans, including machine translation (Jurafksy & Martin, 2000), speech recognition (O'Hara et al. 2004) and reasoning (Panton et al. 2006). Unfortunately, the task of building a broad-ranging knowledge base has shown to be a very difficult problem, something which is now known as the *knowledge acquisition bottleneck* (Feigenbaum, 1991). Many are the problems that arise from the procedures applied to collect and combine information from diverging (possibly conflicting) sources of knowledge, and also from the need to verify this information for completeness and accuracy.

This drawback of knowledge acquisition methods seemed even more disappointing with the advent of the Web, as it served to push the demand for tools that demonstrate intelligent behavior. The Web scenario, in which information is stored in a distributed fashion and the amount of available data lies beyond human processing power, presents a necessity of assistive technologies that efficiently handles users' needs.

At the same time the Web presents new challenges, it also introduces novel opportunities for knowledge acquisition, and consequently the development of innovative intelligent systems. Because of its diverse human audience and the potentially enormous reach of its applications, the Web can be considered as promising grounds for large-scale knowledge acquisition efforts.

We generally categorize recent knowledge acquisition endeavors in two groups. One of them is explicit knowledge acquisition, in which people consciously contribute with knowledge descriptions to systems. In this case, the incentives to users for devoting their time to this task are either monetary (users are paid to contribute with knowledge to the system) or altruistic (users are to cooperate because they somehow feel rewarded with their contribution *per se*).

A more recent approach has been to develop implicit knowledge acquisition tools. According to this approach, the system presents a tangential task that culminates with knowledge elicitation from the user. In this method, the incentives offered to users are based on the tangential task introduced by the system. Consequently the system will succeed in gathering information without imposing any extra cost to users except the performance of the tangential task. A more fruitful knowledge acquisition will take place the more users will be interested in performing the tangential task.

2.1.1 Explicit knowledge acquisition

Explicit approaches represent the more traditional approach to the problem of knowledge acquisition, where information is elicited from assistants who are aware of their knowledge contribution. Many of its past efforts have been responsible for the establishment of the field and for the demonstration of its importance to Artificial Intelligence in general.

One of the most important approaches to explicit knowledge acquisition is the CYC project (Lenat, 1995). CYC was also one of the earliest efforts on building a large-scale knowledge repository, and its main goal was to build a vast ontology (Gruber, 1993) of common sense knowledge that could serve any AI system. By common sense knowledge we refer to knowledge about the world that every human being is supposed to know. CYC developers were motivated by the idea that the foremost drawback in practical AI agents is their ignorance of elementary information about the world, and an ontology of common sense knowledge is to play a key part in the development of major AI endeavors like natural language understanding and machine learning.

The CYC ontology has been constructed by a small group of trained knowledge engineers who manually encode small pieces of common sense knowledge and who design inference procedures to generate new assertions into the knowledge base using a formal predicate logic and predicate calculus.

Although innovative, the CYC project has not gone without much criticism. For instance, its ambitious goal of serving any AI system is believed to be a serious challenge to a manually-created knowledge base, and thus the coverage of the CYC ontology is still regretfully small in several domains. Also, even the encoding of common sense knowledge has shown to be a difficult and conflicting task, and consequently the information that is eventually encoded in the CYC ontology might be the beliefs of a senior engineer and not what is agreed by the whole group (e.g. "*Is the paint on the wall a part of the wall or not? CYC engineers had a great big debate about it, with some taking the position that once dry it is*" (Chklovski, 2007)). Additionally, there is also concern about the variations suffered by what is known as common sense knowledge, i.e. what the general population might assume as true today might be regarded as false in the future, and vice-verse. This instability of common knowledge might suddenly outdate the CYC ontology.

Even though still far away from its declared goal, the CYC project has already shown successful in some areas. The semantic information generated in the CYC ontology has been fruitful in applications involving natural language processing (Curtis, Cabral and Baxter, 2006), reasoning (Panton et al. 2006) and speech recognition (O`Hara et al. 2004).

Another interesting explicit Knowledge Acquisition approach is The Open Mind initiative (Open Mind, 2007), which is an effort comprised of several programs, each of them quite similar in nature to CYC. However, unlike the CYC methodology, Open Mind attempts not to control the number or nature of knowledge contributors, which hopefully will be numerous and diverse as they can freely volunteer from the Web to collaborate in one or more Open Mind programs.

Each Open Mind program serves a very specific purpose and collects some predefined category of knowledge from altruistic volunteers much like the well-known open source software development paradigm (Open Source Initiative, 2007). One successful Open

Mind program is Open Mind Common Sense (Singh et al, 2002), whose objective is very similar to CYC's, namely to generate a large ontology of common sense knowledge, however in a distributed fashion, collecting information from an uncontrolled population from the Web.

In the Open Mind Common Sense program, a system generates a scenario that leads its users to contribute with information to some knowledge base. In one possible scenario, the system might present a hypothetical event like *Bob bought some milk* and then ask volunteers to *write up to five things that someone should already know in order to fully understand the event*. Hopefully, the system will be able to gather the new information supplied by the user with previous submissions to generate a large repository of common sense knowledge.

New data provided by volunteers is archived in the system's database as plain English sentences, exactly as submitted by users. This enables unskilled users to provide richer pieces of information, but demands for an extra task of parsing the knowledge base in order to allow further processing.

Although rather similar, one could highlight advantages of both the Open Mind initiative and the CYC project over each other. For instance, while the Open Mind initiative has the potential to acquire a much more diverse knowledge base, since its information source is distributed among the whole Web, the CYC project provides much stronger guarantees against malicious or ill-informed volunteers who might provide information that is either false or not generally regarded as true. However, both endeavors underperform with respect to the amount of information that is collected from volunteers (Chklovski and Gil, 2005). It seems clear that the absence of incentives for people to contribute with information into these systems leads to a very inefficient use of the most powerful resource in this architecture: the human cognitive power. This drawback is the major appeal for the introduction of implicit approaches to knowledge acquisition, as we present next.

2.1.2 Implicit knowledge acquisition

In order to remedy the inefficiency of previous large-scale knowledge acquisition endeavors, recent advances in the field have invested on the incentives provided to volunteers when participating in knowledge acquisition efforts. The Human Computation paradigm (von Ahn, 2005) fits in this class, and can be described as an attempt to make good use of the human processing power when people perform peripheral ordinary tasks on the Web.

One major ordinary task that has been use as grounds for this Human Computation approach is online game playing on the Web (von Ahn, 2006). In this case, it is shown how one could design mechanisms in the form of entertainment games that try to collect useful information from players who are entertained by the game interaction. This approach was motivated by the large amount of time that is spent by people who seek entertainment by playing games on the Web, and it has proven successful in some domains, especially those related to image recognition problems. We describe next the ESP Game and Peakaboom, two games that have been designed to collect useful information and try to solve computer vision problems at the same time that they entertain the player who is unaware of her volunteering role.

The ESP game (von Ahn and Dabbish, 2004) is a two-player game which aims at assigning descriptive textual labels to images. In this game, two players are anonymously paired over the Web and presented with the same image (Figure 2). They are then challenged to guess whatever their partner is typing, and are only rewarded with some amount of points if both type the exact same label.

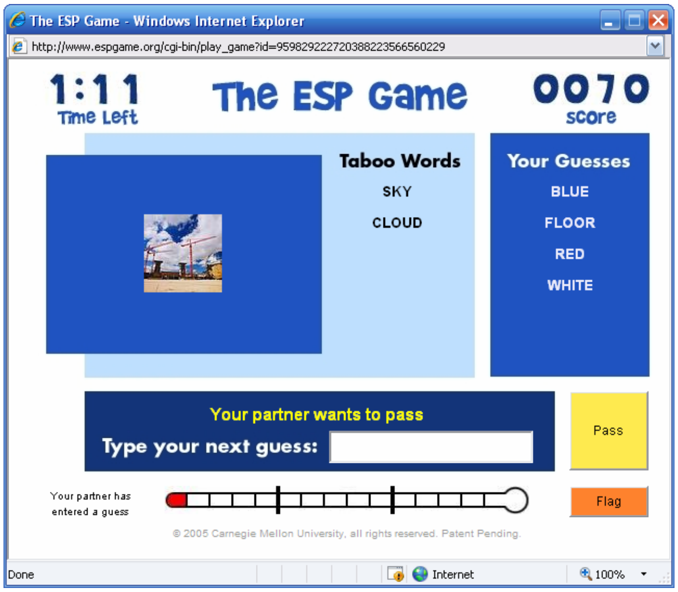

The ESP Game - Windows Internet Explorer

http://www.espgame.org/cgi-bin/play_game?id=959829222720388223566560229

Your partner wants to pass

Your partner has entered a guess

© 2005 Carnegie Mellon University, all rights reserved. Patent Pending.

Figure 2. The ESP game.

Since the identity of her partner is anonymous, it turns out that the only rational strategy a player could employ to succeed in this game is to type something related to the image she knows both players can see. And so, whenever the two players agree on a label, it is reasonable to expect that this label somehow relates to the image presented. Therefore, the system attaches agreed labels to the image and presents users with a new image for another round of the game.

In order to keep users motivated to type good and diverse labels to an image, some precautions are implemented by this game's designers. For example, whenever two players agree on a label to an image, this label enters a list of taboo words related to the respective image. This is done in order to avoid attaching the same label to an image twice. Also, there is a limited period of time when partners must label the maximum number of images they can, so decreasing the chance that players might agree on some random label they might type.

The output of the ESP game is a large collection of label-image relations which can be useful to applications like image searching on the Web, or as training data to machine learning algorithms. The former case is quite straightforward because a textual query could be compared to the textual labels in order to seek valid results. The later, however, presents some further complication, since the ESP game fails to justify how a given label might describe the content of an image. This problem is attacked in the Peekaboom game, which is discussed in the next Section.

Despite the fact that the ESP game has shown success in the task of assigning valid textual descriptions to images, it still does a poor job when serving as training data to machine learning algorithm. This is mainly because, even though it presents a valid label to an image, it fails to indicate which part of an image is represented by the label. For example, in Figure 3, the ESP game could indicate that a good label to this image might be the word *baby*, however, it would never be able to indicate which part of the image relates to this label, or more specifically, which pixels represent a baby in this image.

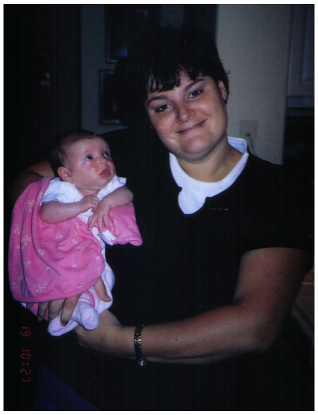

Figure 3. An image related to the label *baby*.

In order to remedy this drawback, a new game has been designed that attempts to provide such information. This game is called Peekaboom (von Ahn, Liu and Blum, 2006), and it is a two-player game to be played on the Web (Figure 4). In each Peekaboom iteration, the *Boom* player observes a pair label-image and must help the *Peek* player identify the label associated with the current image.

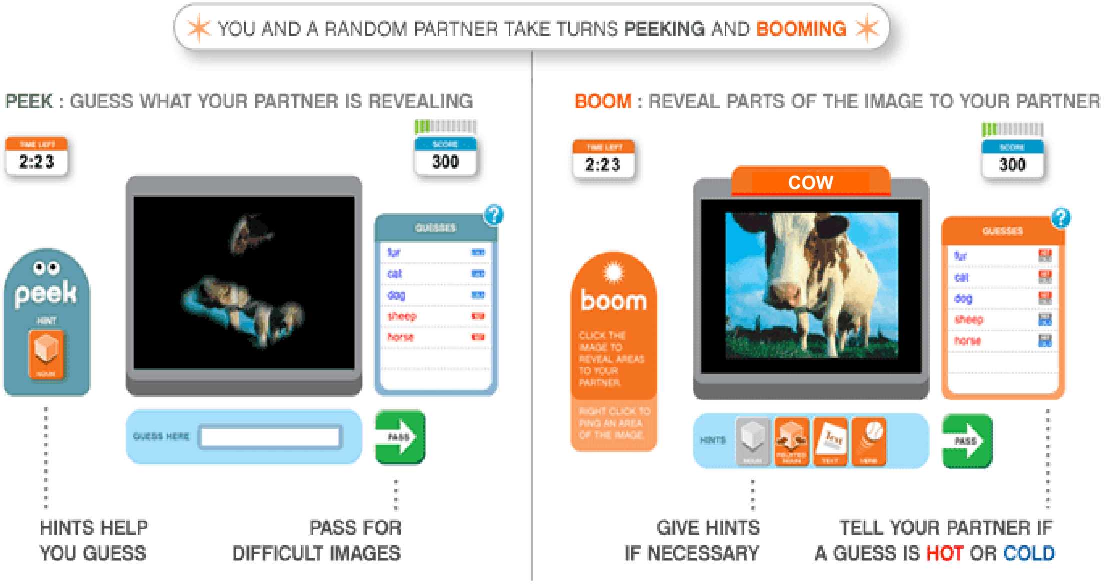
YOU AND A RANDOM PARTNER TAKE TURNS PEEKING AND BOOMING
PEEK : GUESS WHAT YOUR PARTNER IS REVEALING
TIME LEFT
2:23
SCORE
300
peek
HINT
GUESSES
fur
cat
dog
sheep
horse
GUESS HERE
PASS
HINTS HELP YOU GUESS
PASS FOR DIFFICULT IMAGES
BOOM : REVEAL PARTS OF THE IMAGE TO YOUR PARTNER
TIME LEFT
2:23
SCORE
300
COW
boom
CLICK THE IMAGE TO REVEAL AREAS TO YOUR PARTNER
RIGHT CLICK TO PING AN AREA OF THE IMAGE
GUESSES
fur
cat
dog
sheep
horse
HINTS
PASS
GIVE HINTS IF NECESSARY
TELL YOUR PARTNER IF A GUESS IS HOT OR COLD

Figure 4. The Peekaboom game (image from (von Ahn, Liu and Blum, 2006))

The only communication between both players happens when *Boom* has a chance of showing *Peek* some small portion of the image. *Peek* can see no information except the small fragments of the image that are revealed by *Boom*. *Boom* then gains more points the smaller the portion of the image that is revealed to *Peek*. Once the *Peek* player correctly guesses the desired label, the system is able to infer which parts of the images actually relate to the label collected from the ESP game.

The general game approach to knowledge elicitation from users has proven considerably more efficient than previous approaches presented in this Section. It is able to collect much more information because of the psychological incentives underlying the game structure (i.e. the entertaining aspect of the interaction contributes to the addictiveness of players). However, knowledge collected by these means might suffer from some problems. First, and foremost, such a resulting knowledge base would be terribly biased, since this entertainment incentive does not easily generalize to a wide range of different people (for example, a game that entertains a little girl hardly ever will catch the attention of her parents). Second, since every small detail of this system (such as the color of the user interface, as well as its sounds and animation reflecting user's action) is crucial for its success (the amount of time people spend playing this game), it is extremely difficult to adapt a game designed to some specific target-audience to some other public.

In Chapter 3, we will introduce a new approach towards knowledge acquisition which attempts to remedy these negative aspects of the gaming approach and still presents the incentive-compatibility that is lacking in the explicit knowledge-acquisition literature.

2.2 CAPTCHA

CAPTCHA stands for Completely Automated Public Turing Test to Tell Computers and Humans Apart (von Ahn, Blum and Langford, 2004). A CAPTCHA works by generating tests that humans are expected to easily pass, but current computers will find it virtually impossible to succeed. A CAPTCHA is useful whenever there is need for protection of some resource against automated attacks. Since most Web resources need some sort of protection, CAPTCHA has become a very popular gadget in Web resources' access procedures.

Figure 5 represents the CAPTCHA architecture, where a Web Service is protected from direct exploitation from illegitimate users, and therefore protected from abusive misuse. In order to prove being a rightful agent, a human must face a test presented by the CAPTCHA controller.

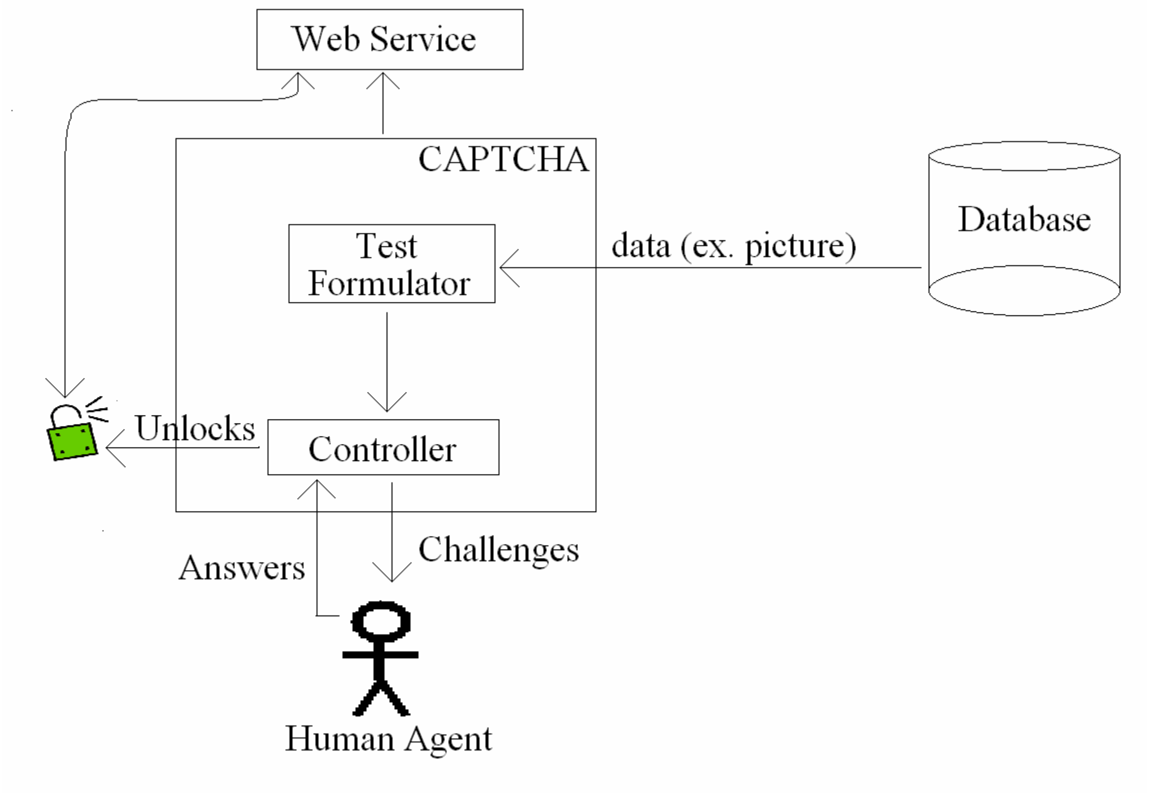

Figure 5. The CAPTCHA architecture.

A CAPTCHA has some additional properties, including its independence from any privacy of data in order to achieve its purpose. In other words, a CAPTCHA test must be unwelcoming for computational agents even if these agents have access to the data used to generate the test. Another property of CAPTCHA tests is to be based on an unsolved Artificial Intelligence problem. This requirement has two purposes: first, it served to guarantee that most humans pass the test; and secondly, it guarantees that the design of every CAPTCHA test presents advancement to either computational security or to Artificial Intelligence. Once a CAPTCHA is developed, either it is never broken (and a major security issue of the Web is forever solved), or it is broken by a computational agent and the Artificial Intelligence field is advanced by this novel agent (since the CAPTCHA was based on an unsolved AI problem).

There are numerous successful applications of CAPTCHA on different Web resources, including email account registrations (Yahoo! Mail), which cannot be freely accessed due to risk that spammers will be able to spread even more unsolicited email messages, public database queries (Plataforma Lattes), which might be flooded by denial of service attacks,

and SMS message sending services (Mundo Oi) whose unrestricted access could extend the distribution of unsolicited messages to mobile telephone devices..
We present next the two most popular instances of CAPTCHA, Gimpy and Sounds.

2.2.1 Gimpy

The oldest and, arguably, most popular CAPTCHA test is Gimpy (The CAPTCHA Project) (Figure 6). It works by first selecting at random some characters of the Latin alphabet, then generating a randomly distorted picture containing these letters and finally presenting this image to some agent, who is then challenged to identify which letters are represented in the image.

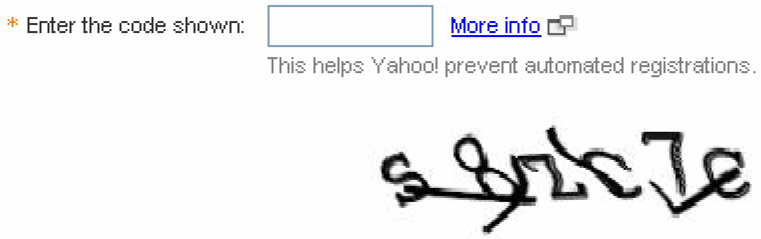

* Enter the code shown:

This helps Yahoo! prevent automated registrations.

Figure 6. An implementation of the Gimpy CAPTCHA from (Yahoo! Mail).

The random distortion introduced in the image is to assure the inability of computers to pass this test, since the Gimpy test is based on the assumption that computers are not capable of recognizing distorted text as well as humans. Indeed, this was the case when this CAPTCHA was developed. However, the incentive presented by the widespread use of this test on the Web has been enough for the development of computational tools that can pass this test with reasonable success (Mori and Malik, 2003), turning Gimpy into a broken CAPTCHA. Current implementations of this test on the Web maintain security guarantees by designing very difficult instances of this test. Nevertheless, it is worthy to note that a Gimpy test that would block access to the state-of-the-art computational agents would be a very strong barrier to humans as well (Chellapilla et al, 2005).

2.2.2 ARTiFACIAL

A significant departure from traditional distorted-characters CAPTCHAs is presented in (Rui and Liu, 2003). This work's premise is that the singular most easily identifiable object by humans is a human face. Therefore, in their CAPTCHA test they create a distorted image of a human face (Figure 8) and ask the user to identify some of its parts like a mouth, the eyes or a nose. The distortion is generated using a basic image of a human face and a 3D model of a generic head (Figure 7). Consequently, the computer can precisely generate a CAPTCHA test and correctly identify correct and incorrect answers to the test.

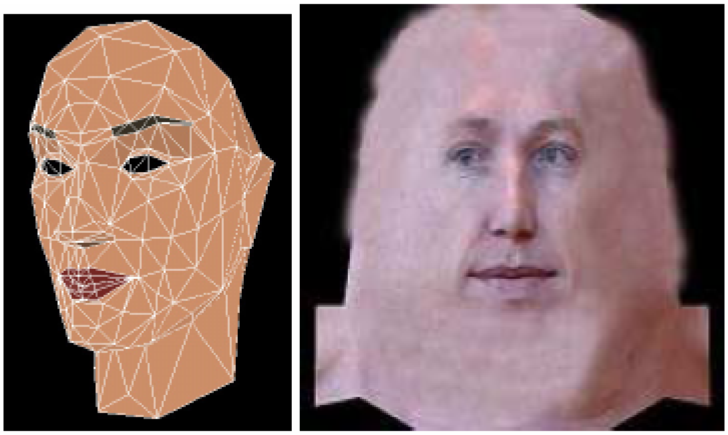

Figure 7. A generic 3D model of a human head, along with a picture of a human face, serves to generate the ARTiFACIAL CAPTCHA (Rui and Liu, 2003).

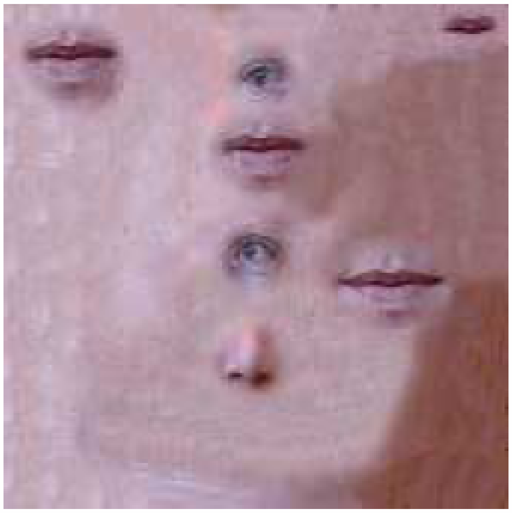

Figure 8. The ARTiFACIAL CAPTCHA: users are challenged to identify parts of the human face, such as eyes, nose and the mouth (Rui and Liu, 2003).

2.2.3 Pix

Pix (Figure 1, Chapter 1) is another example of CAPTCHA. It maintains as database a large collection of images, each associated with one or more textual labels. A test is generated by selecting one label at random, picking four images related to that label, distorting them, and finally asking the challenged agent to indicate, from a list of all existing labels, which one relates to all four images.

The intuition behind this CAPTCHA is that computers are not able to identify image's content and then reason to find a common aspect between the four images and relate that aspect to a textual label. However, since the Pix database is public, the images must be displayed with such distortion that a computational agent will not be able to relate the presented pictures to one of the public images (because if it could, an automated solution to this test would be trivial). An implementation of this test should be attentive to this level of distortion, because an exaggeration might be the source of barriers to humans to succeed in this test.

2.2.4 Sounds

An alternative to Gimpy is the Sounds CAPTCHA (The CAPTCHA Project). This CAPTCHA's design is very similar to Gimpy, but it is based on the inability of computers of recognizing human speech. It works by sampling some predefined number of words, then synthesizing them into a single audio file, and finally distorting it with random noisy sounds.

This type of CAPTCHA is of special interest to visually-impaired people, but many challenges make it a significantly harder test than Gimpy for humans to solve. First, this test is language-specific, which means that in order for someone to correctly pass this test, she must be fluent in the language from which the original words were collected. Also, even completely unsighted people face difficulties related to the speed and volume of the rendered output, which cannot be softened due to security problems.

2.2.5 A Combination of the *Pix* and *Sounds* CAPTCHAs

In a recent work, (Homan et al. 2007) claimed that a test that is based on a combination of audio and visual information represents a significant contribution towards the accessibility of CAPTCHAs. Figure 9 presents a prototype of their solution, where the Pix and Sounds CAPTCHAs are combined. In this test, the user is asked to identify which animal is depicted in a picture and reproduced in a sound clip. Therefore, visually-impaired users could resort to the audio clip to identify the answer, while the visuallygifted ones could follow the Pix-CAPTCHA interaction and simply watch the picture for an answer.

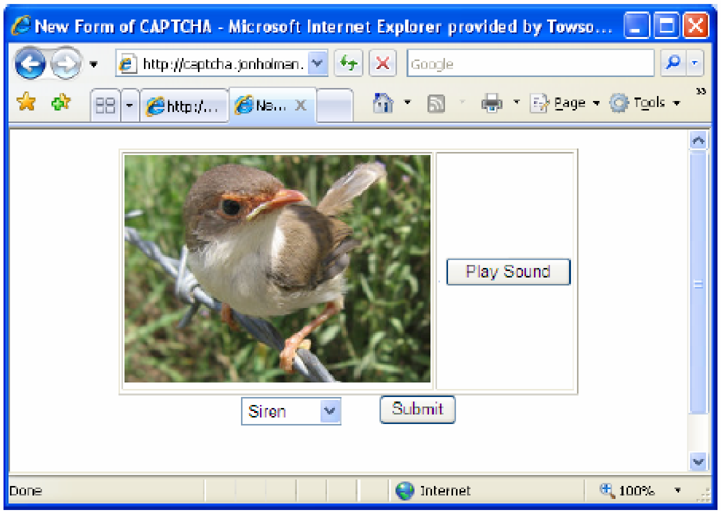
New Form of CAPTCHA - Microsoft Internet Explorer provided by Towso...
http://captcha.jonholman.
Google
http:/...
Ne...
Page
Tools
Play Sound
Siren
Submit
Done
Internet
100%

Figure 9. A Combination of the *Pix* and *Sounds* CAPTCHAs (Holman et al. 2007).

It is important to stress that the current CAPTCHA test stands merely as an overhead to Web services while it does not aggregate value to the domain, making the cognitive effort exerted by humans when taking the tests an unproductive experience.

This scenario poses a good opportunity for knowledge acquisition because people are facing hard tests and willingly giving their best shots to solve them. The problem lies in designing a CAPTCHA test that impedes computational agents to access unauthorized resources at the same that elicits some knowledge from users. In the next Chapter we present an architecture that enables CAPTCHAs to extract knowledge from humans.

Chapter 3. KA-CAPTCHA

In this chapter, we present the KA-CAPTCHA model, a mechanism that is able to elicit knowledge from Web users while performing a CAPTCHA interaction which is common in the Web. We advocate that the KA-CAPTCHA approach has several potential advantages over previous efforts, including the following:

- Unlike the case of knowledge collected from volunteers, the KA-CAPTCHA approach does not rely on altruism from contributors to succeed. Our approach might succeed because people need to use a CAPTCHA to fulfill their agenda (i.e. to access some resource or to perform some task on the Web).
- Unlike the case of entertainment, our method does not require a constant redesign of the user interface in order to keep it attractive to people, as is the case with games. Again, people will not use our CAPTCHA to get distracted, surprised or amused. In fact, they would rather not use a CAPTCHA if they could.
- Because our method's appeal lies in the resource protected by the CAPTCHA and not on the CAPTCHA itself, one could save design and implementation time by applying the very same CAPTCHA mechanism to a number of different services, reaching a wider audience that could contribute with knowledge of a broad nature. We believe this is not the case with entertainment.
- Moreover, the effort our approach expects from our users is already being exerted with the current CAPTCHA design, but only in a very inefficient scenario. Current instances of CAPTCHAs have become so difficult for humans (Chellapilla et al. 2005) that it is a waste not to recycle this considerable effort into something positive. Therefore, our method does not impose any extra costs for volunteers.

As was reviewed in the last chapter, the CAPTCHA architecture expects some set of data that is to be employed in the development of a test that should only be solvable by human beings. As depicted in Figure 10, we model a combination of several of such possibly CAPTCHA-feeding databases as a subset of the knowledge that is collectable from humans, and also as a subset of the Web.

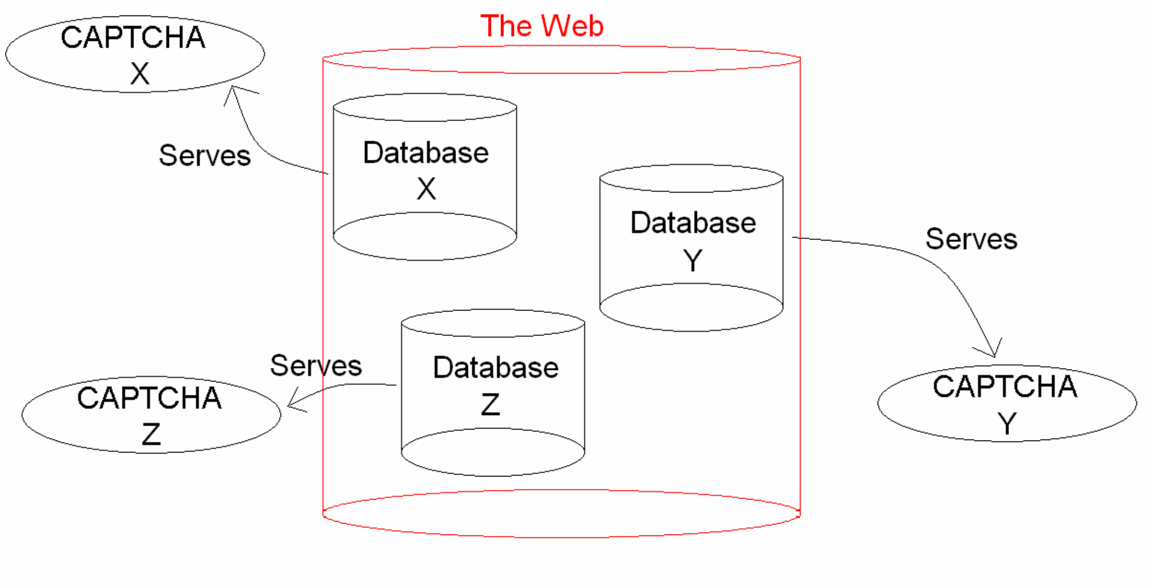

Figure 10. The current architecture of CAPTCHA makes use of a database that can be viewed as a subset of the Web.

Since the usual CAPTCHA model demands for a knowledge base of reliable data, much information from the Web is not adequate to CAPTCHA mechanisms, either because it is poorly annotated or because its semantic content is not reliably documented. Nevertheless, while this is the reason why a considerable subset of the Web is unable to play a part in the development of novel CAPTCHA tests, we envision that the human response to the CAPTCHA interaction might contribute to this poorly annotated data repository.

As is depicted in Figure 5, the CAPTCHA mechanism determines that a challenged agent should solve some kind of test in order to prove being an authentic user. Therefore, the CAPTCHA interaction forces Web users to spend a considerable amount of time in an effort that only annotates that small subset of the Web that is already well annotated. In other words, the CAPTCHA mechanism is also a redundancy generator.

This redundancy is especially inappropriate because one of the greatest shortcomings in annotating the whole Web is the difficulty in convincing people to provide semantics to Web objects. The task is just as onerous as individually-unattractive, since one may not directly benefit from the effort recently exerted. Therefore, the period of time Web users spend solving CAPTCHA tests is very valuable to the development of a Semantic Web.

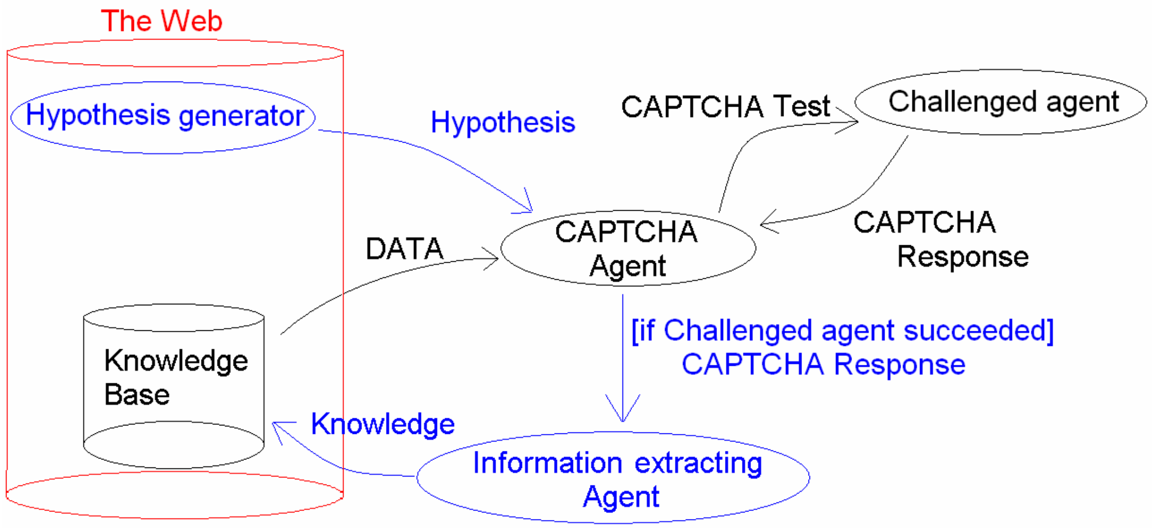

Figure 11. The KA-CAPTCHA architecture, represented by an extended CAPTCHA mechanism that enables knowledge acquisition.

In order to make an efficient use of the human cognitive task employed during the CAPTCHA interaction, we designed a mechanism that benefits from this opportunity called KA-CAPTCHA, which stands for Knowledge Acquisition CAPTCHA. This model works as informally depicted in Figure 11, where all content in black represents the usual CAPTCHA architecture and the blue components are the additions that enable knowledge acquisition.

In the usual CAPTCHA mechanism, an agent generates a test by retrieving data from a public knowledge base. Our extension attempts to collect new knowledge from users by combining what was collected from the usual CAPTCHA knowledge base with data gathered from other sources on the Web, here represented by a Hypothesis generator. The data that is retrieved from the Hypothesis generator represents the content of the Web that is still poorly annotated and whose semantics the CAPTCHA designer believes can be provided by a legitimate user who is solving the CAPTCHA test.

The CAPTCHA agent then generates a test with a combination of data retrieved from the Web and forwards it to some Challenged agent, who reacts with a response to this test. The CAPTCHA agent then reviews the response and grants access to the protected resource conditional on the correctness of the user’s response; i.e. the user gains access to the Web resource if, and only if, her response to the CAPTCHA test is correct. Moreover, if the user submitted a valid response to the CAPTCHA test, the system will enjoy the valid response to the unknown questions that were included in the original CAPTCHA test. It does so by forwarding the user's response to an Information extracting agent (Figure 11) who will analyze the response looking for new knowledge to feed into the original knowledge base. The success of the knowledge extraction from the test's response is directly correlated with the design of the CAPTCHA test and the information retrieval from the Web. Better tests should enable better knowledge acquisition at latter stages, i.e. quantitatively better meaning enable more information being elicited from users, and qualitatively better, resulting in a sound knowledge base.

Figure 12 depicts a formal representation of this mechanism, where the traditional CAPTCHA architecture represented in Figure 5 is extended to allow the recycling of the user's response to feed a database of semantic information.

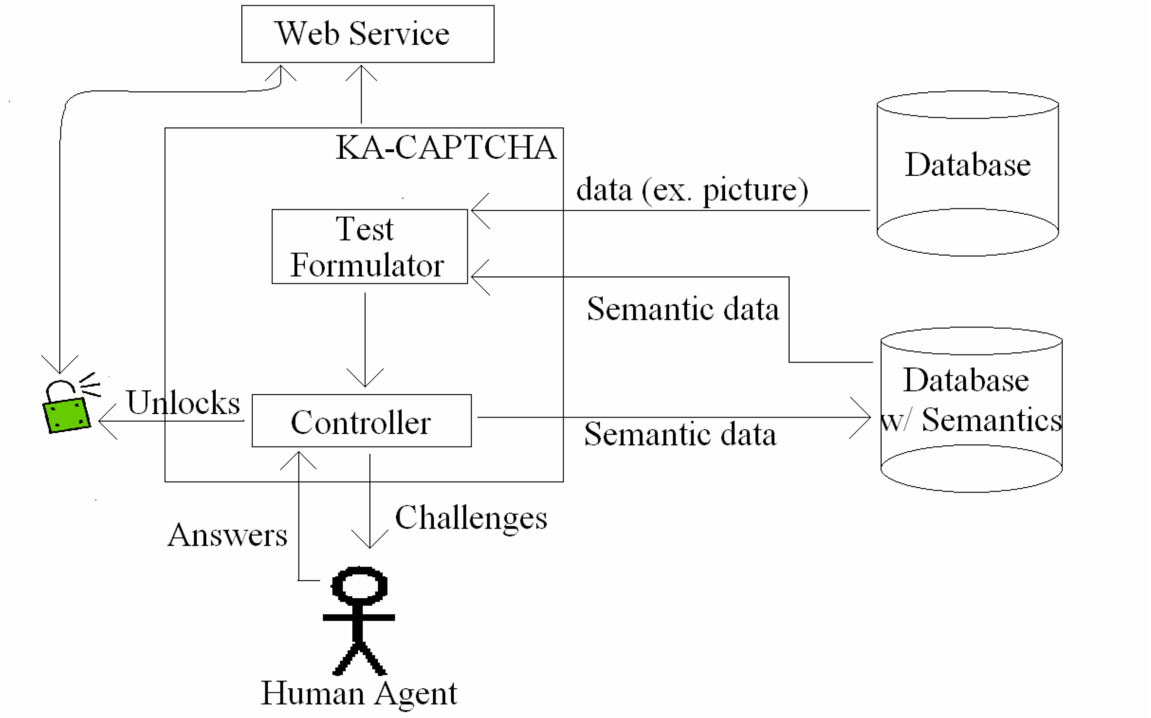

Figure 12. A formal representation of the KA-CAPTCHA mechanism.

In order to assure the soundness of the knowledge elicited from users, the KA-CAPTCHA mechanism associates two measures to its underlying knowledge representation. First, the KA-CAPTCHA requires an ontology (Gruber, 1993) to represent the knowledge collected from users and to correspond to the schema of the CAPTCHA database. Second, our mechanism assigns to each semantic relation in this ontology a Support and a Confidence rank, which will enable a distinction between reliable information and noisy data.

The Support rank associated with each semantic relation in the ontology specifies the number of times a user has indicated that the corresponding relation in the CAPTCHA database is sound. Hence, a high Support rank indicates that a large number of users have recognized the respective relation as sound. A small Support rank would indicate that the relation is not recognized by a large volume of users.

The Confidence rank attempts to measure the probability that a CAPTCHA user will recognize the respective relation as true. It is defined as the ratio between the Support rank and the number of times the relation has been questioned to users, and its value ranges from 0 (when the Support of the relation is 0 as well, i.e. no user has ever recognized the relation as true) to 1 (when every user who has been questioned about the validity of the relation has also identified it as sound).

Figure 13 illustrates the use of these ranks with a sketch of a pictorial knowledge ontology. In this example, the only existing semantic relation in the ontology is a *describes* relation connecting a label to an image. In this scenario, an agent uses a KA-CAPTCHA test to query a set of users about the relation between two elements of its database, namely the label *Love* and a specific image collected from the Web. After each user's response, the KA-CAPTCHA agent updates the Support and Confidence rank of the relation according to the user's response.

As shown in Figure 13, the Support rank is incremented whenever a user answers the KA-CAPTCHA query, and consequently the Confidence rank is affected by this change as well. After the third iteration in this example, the CAPTCHA agent finds that the Support of the relation between the label *Love* and the respective image equals to 2 and the Confidence equals ⅔ ≅ 0.67.

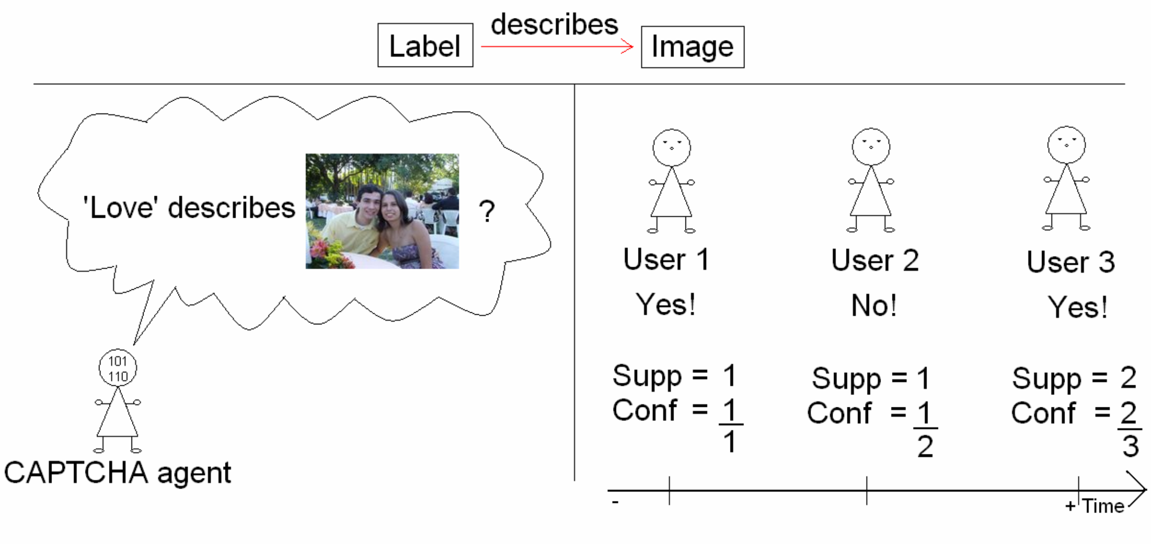

Figure 13. An illustration of the KA-CAPTCHA interaction

In order to interpret the Support and Confidence ranks of the semantic relations in the ontology, the system maintains a set of parameters that indicate the appropriate interpretation of the value of these ranks. Regarding the Support rank (Figure 14, left), the system holds a Support Threshold that indicates the minimum amount of indications needed for a given relation to be considered established. Hence, whenever a given relation's Support is below the Support Threshold, the system will not be confident of the soundness of this relation, independently of the value of its Confidence.

The interpretation of the Confidence rank follows a similar procedure. The system keeps two measures (Figure 14, right), namely a Certainty threshold and a Suspicion threshold.

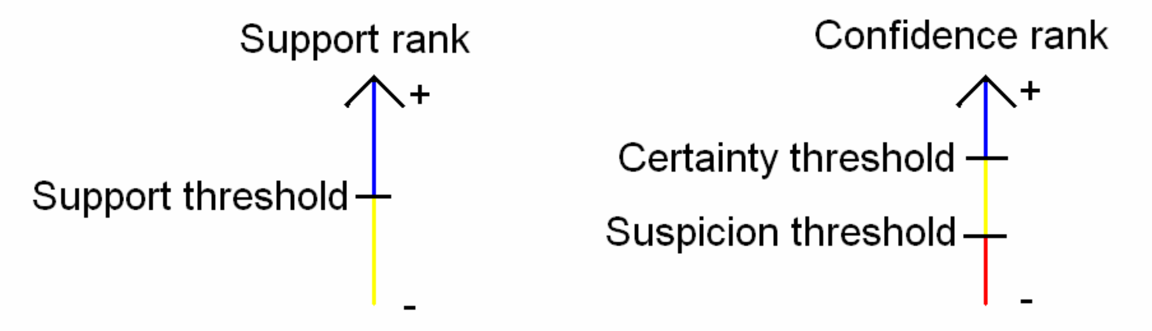

Figure 14. The interpretation of the ranks of each semantic relation in the KA-CAPTCHA ontology.

Whenever the Confidence rank of a certain relation is above the Certainty threshold, it indicates that the respective semantic relation is to be considered sound by the system. Conversely, if a Confidence rank falls below a Suspicion threshold, this makes the system believe that the semantic relation should be considered false. Whenever the Confidence rank is between these two thresholds, this indicates that the system is uncertain about the validity of the relation, independently of the value of its Support rank. Table 1 bellows summarizes the interpretation of these measures, where the ‘+’ sign means the system regards the respective relation as true, ‘–’ means the system regards the relation as false and ‘*?*’ means the relation is yet to be further verified.

The exact value of these parameters should be empirically determined, based on specific properties of the target population like its size and diversity.

Table 1. Interpretation of the ranks attached to each pair label-image in our database.

Confidence rank Support rank	Above Certainty threshold	Between thresholds	Below Suspicion threshold
Below Support threshold	?	?	?
Above Support threshold	+	?	–

With this interpretation of the Support and Confidence ranks, the system is able to automatically generate new CAPTCHA tests based on the knowledge collected in previous iterations. Since a CAPTCHA test must be comprised of a series of questions based on what is encoded in its knowledge base, the KA-CAPTCHA system can generate a CAPTCHA test by selecting those relations whose Support and Confidence ranks are above the Support and Certainty thresholds, respectively. Additionally, knowledge acquisition can be performed by adding to these known relations some relations whose soundness the system is still uncertain of. By mixing these relations together, it is expected that the user will not distinguish between those the system knows the correct answer and those that are still unknown. Consequently, a user will have to answer the whole CAPTCHA test truthfully if she is willing to access the Web resource protected by the system.

After a response to the CAPTCHA test has been submitted by the user, our system must verify its accuracy. If the user correctly recognizes the answer to all of these known relations present in the test, the KA-CAPTCHA can rely on the authenticity of its user.

Furthermore, in case the user succeeds on the test, the KA-CAPTCHA can also take in consideration those relations embedded in the test whose soundness the system wasn't confident of, yet. The Support and Confidence of these relations are to be updated appropriately at future KA-CAPTCHA iterations, in accordance to its future user's response.

Chapter 4. Applications

In this Chapter, we present two instances of the KA-CAPTCHA mechanism described in the previous Chapter. They support the claim that the method is general, feasible for the Web, and is able to satisfy existing demands for knowledge acquisition.

4.1 Application to pictorial knowledge acquisition

Semantic annotation of images is a challenging issue in the field of computer vision. The difficulty in designing computer algorithms that correctly describe the content of visual data lies in satisfying the demand for a method that is general and still generates output with high precision.

This scenario is especially disappointing if one considers the requirement for semantic annotation of images on the Web. The demand for applications that allow textual search queries over image-databases (e.g. Google image search) exposed this lack of proper technology to Web users. Additionally, this condition contributed with Web inaccessibility, since much of the Web content is archived in a visual format, and without automatic generation of proper textual alternatives to these data, visually-impaired users are not able to fully enjoy the Web experience.

In order to deal with the lack of technology for the proper annotation of images, many image-searching mechanisms on the Web have relied on a popular workaround to generate textual descriptions to images. This heuristic consists of the search for textual objects that are related to the images, instead of searching for images directly. Consequently, once a textual Web object (e.g. a hyper-text Web page) has been retrieved according to a given textual query, these systems output a group of images that are related to the Web object (e.g. the images contained in the Web page).

Although this heuristic search is reasonably satisfactory regarding the metric of recall, its precision is poor. This is a direct consequence of the weak semantic association between these two objects (i.e. it is not the case that every image present on a Web page is suitably described by the keywords of the Web page). Therefore, a method that remedies this drawback of image retrieval on the Web would be of great use.

4.1.1 Knowledge modeling

Our goal is to design a CAPTCHA that will enhance the precision of existing image databases on the Web. Therefore, we will model our CAPTCHA's knowledge database as a set of different images and textual labels. Additionally, we will consider links connecting textual labels to image.

Naturally, some of the images in our database will relate to a group of labels because the label might be descriptive of the image's content. We will regard this connection as a d*escribes* relation and will assume that the ontology depicted in Figure 15 serves as a schema of our knowledge base. It comprises two classes of elements (namely *Label* and *Image*) and a *describes* relation between these two classes.

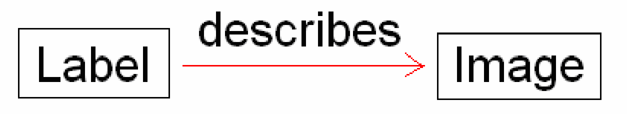

Figure 15. An ontology as a schema for KA-CAPTCHA to elicit pictorial knowledge from Web users.

Once we model our knowledge base as the ontology represented above, the task of retrieving the correct information from our knowledge base is reduced to the problem of unscrambling the appropriate *describes* relations from the bad ones. To support this task, we associate to each relation the *Support* and *Confidence* ranks described in Chapter 3 (Figure 14).

4.1.2 KA-CAPTCHA design

The design of our test addresses the issues of security and knowledge acquisition. In accordance to these principles, we designed some KA-CAPTCHA tests that are similar to that depicted in Figure 16. In this test, the CAPTCHA agent collects at random a label and a set of images from its database; some images that relate to the current label (those whose ranks match a '+' in Table 1), some that do not relate to the label (ranks match '-' in Table 1), and some whose relation to the label is still unknown to the system ('?' in Table 1). Subsequently, these images are randomly distorted in order to block a naïve strategy of searching the CAPTCHA's database and retrieving the relations between the current label and the displayed images. And finally, the images are organized in columns (Figure 16) and the user is challenged to indicate, for each column, which image is described by the current label (i.e. the user is challenged to identify which image's connection to the label is marked with a '+' in the system's database).

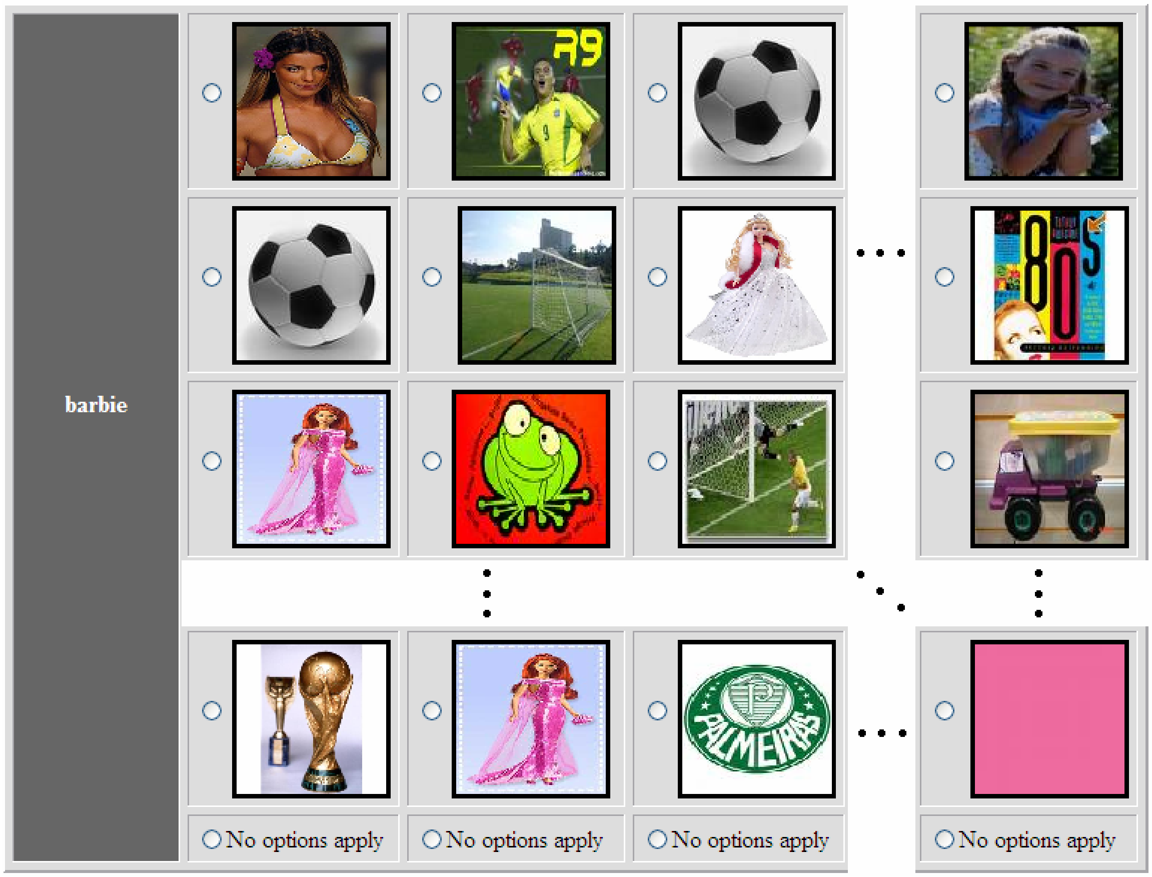

Figure 16. A KA-CAPTCHA to elicit pictorial knowledge from Web users.

After the user submits an answer to this test, our system verifies the option selected in each column. If the user chooses an image whose relation to the relation is recognized as false by the system, then this user is believed to be a computer and is denied access to the resource protected by the CAPTCHA. Otherwise, the user is believed to be a human and is granted access to the Web resource.

In addition to this security inspection, the KA-CAPTCHA is also able to elicit knowledge from the interaction with humans. For this reason, the KA-CAPTCHA randomly chooses one of the columns of this test and inserts in it only image whose relation to the current label is still unknown. This special column is presented to the user as a regular column, as if the user would be judged based on his answer to this column as well. However, since the system is unsure if any of the images on this special column is indeed related to the current label, the system is forced to provide the user with a *No options apply* alternative.

Consequently, since the end user must not distinguish this special column from the others, this *No options apply* option must be available for every column of the test.

After verifying the legitimacy of an agent, if this agent is believed to be a human (i.e. if the agent has correctly recognized the option in every regular column of the test), the KA-CAPTCHA agent will regard any selection that the user has made in the special column to be true and update its knowledge base accordingly, as illustrated in Figure 13. If the human agent has selected an image from the special column, then this is interpreted as the *Yes* answer in Figure 13, and the *Support* and *Confidence* of the pair label-selected image will increase. If the user has chosen *No options apply* in the special column, then the system will interpret this as a *No* answer (from Figure 13) for every image in this column. Consequently, every image in this column will have their *Support* increased and their *Confidence* decreased.

The total number of columns in this test is variable, and depends on the number of images retrieved from the knowledge base. Ideally, there should be enough columns in this test to make a random attack look unattractive (i.e. to make the probability of passing this test by choosing an answer at random be sufficiently small). Since the number of images per column (*#i*) and the number of columns per test (*#c*) are variable, if we let *M* be the maximum acceptable probability of a success by chance, we can choose any *#i* and *#c* that satisfies the equation

$$M\ (\#i+1)^{\#c-1} \leq 1.$$

In addition to the test depicted in Figure 16, a KA-CAPTCHA designer can present users with different tests. A more obvious variation asks users to indicate which image does **not** relate to the label. This variation is faster at pruning away bad relations, in detriment of highlighting good ones. Also, Figure 17 shows a slightly different test, where an image is the main focus of the interaction, and the user is challenged to identify which label describes the image (instead of being asked which image is described by the label).

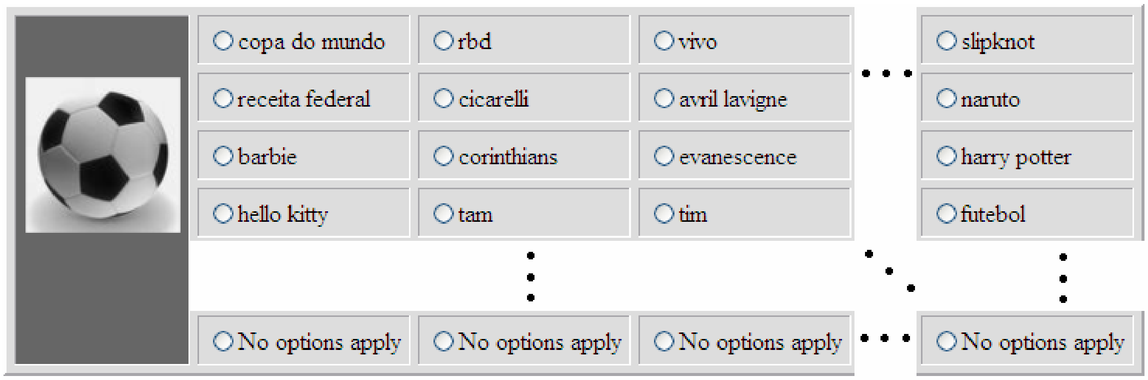

Figure 17. A different pictorial knowledge KA-CAPTCHA, where the test is focused on a single image and its different potential labels.

4.1.3 Empirical evaluation

In order to evaluate our method, we implemented KA-CAPTCHA for the domain of picture description elicitation. This controlled experiment was developed in 3 stages. In total, 147 volunteers took part in the experiments collecting data, interacting with KA-CAPTCHA or evaluating the resulting knowledge base.

During the experiment, we measured the number of KA-CAPTCHA trials each volunteer needed before proving being a human. Then we measured the precision and recall of our resulting knowledge base and compared it to Google's image search.

The first stage of the experiment served to generate this knowledge base. A second phase represents the actual use of KA-CAPTCHA by volunteers and consequent refinement of its knowledge base, and a final stage was designed to evaluate the data extracted from steps 1 and 2. There were no overlap of volunteers, i.e. each participant took part in one, and only one, stage of the experiment.

Previous preliminary experiments served to indicate adequate values for our method's parameters to this sample. During the final experiments, we used a Certainty threshold of 0.8, Suspicion threshold of 0.2 and a Support threshold of 2. We used 5 images per column, with 4 columns per test.

Stage 1

Participants. Two volunteers from Brazil, ages 24 and 25 years old, took part in this stage for free. One was a male engineering graduate student and the other a female literature undergraduate. Participants had no previous knowledge of this work or of the participation of each other.

Procedure. At first, we collected the 15 most popular queries made to Google from Brazil during the month of September 2006 – data is available in the Google Press Center. Then, we presented each query as a label and asked our volunteers to retrieve from the Web, for each label, 10 images they believe related to the label and 10 images they believe did not relate to the label. Volunteer 1 managed to collect 101 images, while Volunteer 2 collected only 63. We labeled each group of images as Knowledge base 1 and Knowledge base 2, respectively.

Stage 2

Participants. 143 Brazilian computer science, physics and engineering graduate and undergraduate students indirectly took part in this stage, for free. Their respective professors, all belonging to our Computer Science department, supported this study by providing us their students' test grades. No participant had any knowledge of the research described in this paper.

Procedure. In order to attract users to our CAPTCHA, we designed a small application where students would be able to consult their grades online in a private way. In our Computer Science department, student's grades used to be published on a piece of paper attached to a wall, with no attempt to preserve grades' confidentiality. With the advent of our system, professors were able to appropriately address student privacy. We collected students' grades twice. Initially, in the middle of the semester when students accessed the system to view their grades to mid-term tests (when we used Knowledge base 1), and then at the end of the semester, when the system was accessed for information about final grades and class status, during which time we used Knowledge base 2. Students were told by their instructor that the only way they could know their test score was through our system, but weren't notified of the presence of the CAPTCHA, how to solve it or what its objective was. All information they had about the system was its Web address and instructions in its interface (which was the Portuguese equivalent of "*Select, for each column on the right, the alternative that is best described by the label on the left*").

Stage 3

Participants. Two male volunteers from Brazil, 24 years old, took part in this final stage for free. Both were undergraduate students, one studying law and the other veterinary medicine. Both had no knowledge of the research underneath the experiment.

Procedure. In this stage, participants were presented with all the labels and images collected from stage 1 and were asked to indicate, for each image, which labels describe the image. They were instructed to assign any number of labels they thought described the current image, including none if that was the case.

Our first evaluation regards the ability of our users to pass the KA-CAPTCHA test. Figure 18 presents a graph indicating how many attempts our users have made before passing the CAPTCHA (0 means the user passed the test at first attempt, 6 means the user failed the test 6 times before finally succeeding). Eventually, every user passed the KA-CAPTCHA test.

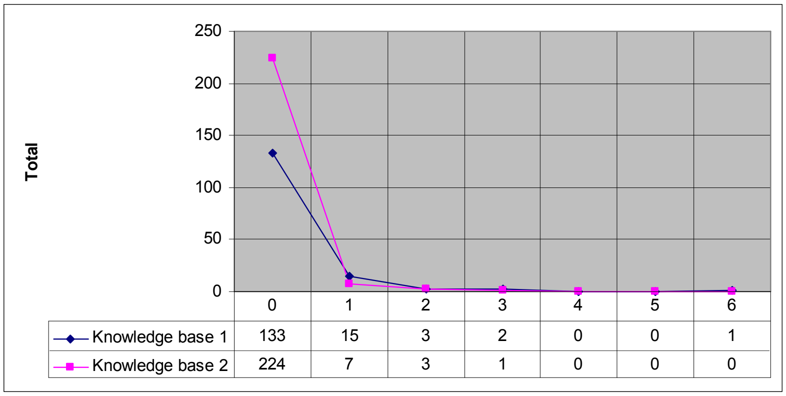

	0	1	2	3	4	5	6
Knowledge base 1	133	15	3	2	0	0	1
Knowledge base 2	224	7	3	1	0	0	0

Figure 18. Number of attempts made by volunteers before passing the KA-CAPTCHA.

Access to our CAPTCHA using Knowledge base 1 resulted in users passing the test in 81.48% of the attempts. Knowledge base 2 was used in a moment when users were more familiar with the system, as now they achieved the higher rate of success of 93.08%.

After verifying our user's ability to pass the KA-CAPTCHA test, our final evaluation concerned the quality of the information elicited from users. Figure 19 shows a representation of these results, where each bar represents the final status of each knowledge base.

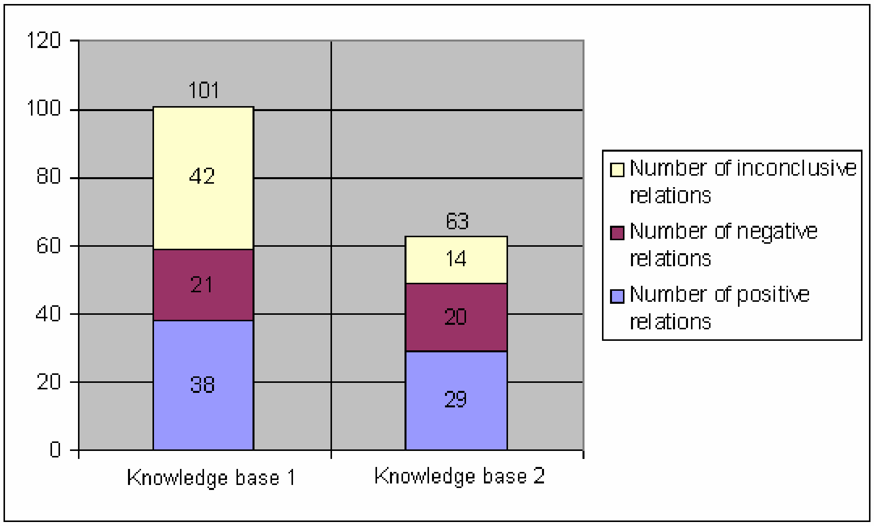

Figure 19. A representation of the knowledge elicited by our KA-CAPTCHA from volunteers.

Knowledge base 1 was formed by 101 relations between images and labels. After all students had consulted the system for their mid-semester grades, we observed that 38 relations were regarded by the system as true, 21 had been considered false and 42 were still inconclusive (the system couldn't yet decide whether they were true or false). Comparing this knowledge collected by our CAPTCHA to data collected from volunteers of stage 3, we noticed two false-negative relations (i.e. relations the CAPTCHA regarded as false but later volunteers regarded as true) but no false-positives (i.e. all relations the CAPTCHA regarded as true were agreed by our volunteers). This results in a precision of (108 – 2) / 108 = 0.98148.

Knowledge base 2 was initially composed of 63 relations. After the experiments, 29 of these were considered valid relations, 20 were considered false and 14 were still inconclusive. When comparing this data to the information extracted from stage 3, we observed two false-negatives but no false-positives.

In order to better evaluate these results to the state-of-the-art, we compared our knowledge base to Google's Image database. Google's data is constructed based on two methods. The first is the heuristic described in Section 4.1, where a query for an image is first applied to a textual database of Web pages, and then the images associated with this Web page are retrieved. Also, Google applies a variation of the ESP game called Google Image Labeler in order to refine its search results.

In order to compare both method's precision and recall, we asked a volunteer from our previous experiment to indicate, for each label in our database, which of the first 180 search results from Google were correct (i.e. which results were a plausible search result for the respective label). Table 2 shows the results of this experiments, indicating and overall precision of 0.54307 among all queries made to Google's database.

Table 2. An evaluation of the search results from Google's image database.

	Correct results	Incorrect results	Precision
Avril Lavigne	117	63	0.65
Barbie	100	80	0.55556
Cicarelli	96	84	0.53333
Copa do mundo	69	111	0.38333
Corinthians	105	75	0.58333
Evanescence	149	31	0.82778
Harry Potter	126	54	0.7
Hello Kitty	154	26	0.85556
Naruto	131	49	0.72778
Rbd	119	61	0.66111
Receita Federal	34	146	0.18889
Slipknot	141	39	0.78333
Tam	55	125	0.30556
Tim	29	151	0.16111
Vivo	25	125	0.16667
TOTAL	1450	1220	0.54307

When comparing this result with our method's precision of 0.98148, we observe the clear contribution that the KA-CAPTCHA method presents to the state of the art image search heuristics. However, the experiments demonstrated that the KA-CAPTCHA recall achieved the low rate of 108 / 164 = 0.65853. This was expected since our method ran for a relatively small time window (during a couple of weeks), and additionally, we argue that our method's appeal is very robust against a decline of interest from users, therefore this small level of recall could be compensated with a constant rate of users accessing our KA-CAPTCHA.

4.2 Application to acquisition of textual transcriptions

A careful inspection on the Web reveals a constant increase in the distribution of audiovisual content like personal videos (www.youtube.com, video.google.com) and university lectures or seminars (mitworld.mit.edu, itunes.stanford.edu). However, a large percentage of the audience of these services need access to it in circumstances where audio resources cannot be relied on (students may want to study lectures in libraries) or have little value (foreign individuals might be unable to understand the spoken language or hearing-impaired users won't be able to understand the audio content). Consequently, this technology will never reach its full potential unless transcriptions accompany its auditory component.

Unfortunately, the generation of transcriptions is not an easy task. State-of-the-art speech recognition systems are still less than satisfactory in natural scenarios (Munteanu et al. 2006), and the cost of hiring professionals to manually transcribe the data would make this approach economically infeasible. Nevertheless, it is important to note that the visually-impaired could perform this transcription with relatively ease. Existing CAPTCHA tests (The CAPTCHA Project) already challenge users to perform a similar task, but they only ask for a transcription of meaningless pieces of noisy sound. We could then recycle this thrown-away effort and turn it into a knowledge acquisition mechanism.

4.2.1 Knowledge modeling

The modeling of a KA-CAPTCHA to acquire textual transcriptions to videos is very similar to the pictorial example of the previous section. Here, we will consider the existence of two entities called words and transcriptions. The only semantic relation we will model is a *part of* relation (Figure 20), indicating that a given word is part of a correct transcription being made to a video.

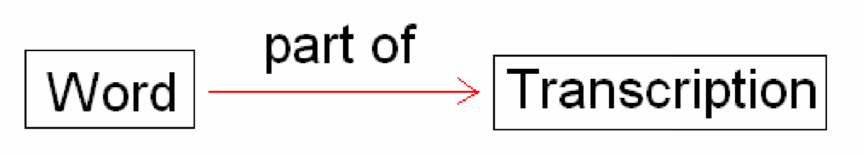

Figure 20. A taxonomy of the schema used by KA-CAPTCHA to elicit textual transcriptions to videos from Web users.

Once we model our knowledge base as the taxonomy represented above, the task of generating valid textual transcriptions reduces to the task of identifying which Word-Transcription relations are correct. Again, this is done assigning the Support and Threshold ranks described in Chapter 3 to each *part of* relation.

4.2.2 KA-CAPTCHA design

The test we design here should collect textual transcriptions to online videos. Therefore, in this CAPTCHA test we will present users with a segment of audio data and ask for a transcription of what the user has heard.

In order to present a reasonable challenge to humans, we break each video from our database into various segments of fixed length. The exact size of each of these segments should be empirically determined. A too-short segment would make it easy for computers to pass our test. A lengthier fragment will complicate computers' success but could also become a considerable obstacle for humans, who might lack the patience, attention or memory to transcribe the entire passage.

After the user submits a response to this CAPTCHA test, we evaluate the submitted response and compare it to the current transcription present in the KA-CAPTCHA knowledge base. This current transcription is retrieved by joining all the words associated with the transcription that have its Support and Confidence above their respective minimum thresholds.

When confronting the current transcription to the transcription submitted by the user, the credibility of users decreases the more words (s)he fails to transcribe and the more certain the system was of these missing words (ω_2 in Figure 21). A user succeeds the CAPTCHA if her/his credibility lies above a *credibility threshold*, when the system is able to grant access to the user to some Web resource.

In case of user success, the system also performs the knowledge acquisition procedure. In this step, the beliefs in each of the words that belonged to the user's response and which the system is not yet confident about are increased by a unit (ω_4 in Figure 21).

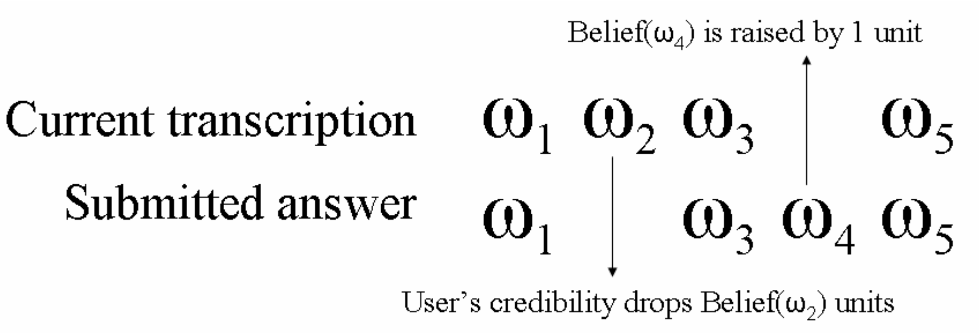
Belief(ω4) is raised by 1 unit
Current transcription
ω1 ω2 ω3 ω5
Submitted answer
ω1 ω3 ω4 ω5
User's credibility drops Belief(ω2) units

Figure 21. Interpretation of a user response.

Chapter 5. Conclusions

5.1 Contributions

This thesis presented KA-CAPTCHA, a new approach to knowledge acquisition on the Web. We took advantage of the existing and widely-disseminated CAPTCHA model and embedded a knowledge elicitation component underneath the existing security mechanism. This approach allows for costless acquisition of information because the effort we expect from our knowledge contributors is similar to that already being exerted in usual CAPTCHA mechanisms.

In support to our contribution, we designed two KA-CAPTCHA tests whose interaction results in knowledge acquisition from Web users. Experimental results suggest the feasibility of the KA-CAPTCHA approach in a pictorial knowledge scenario and its good performance compared to existing image retrieval systems like Google Image Search.

5.2 Limitations

While the KA-CAPTCHA model was designed to collect knowledge from a wide audience on the Web, a final analysis of the experiments reported on Chapter 4 revealed some limitations of our method. With regards to the KA-CAPTCHA presented in Section 4.1, a potential inefficacy of this application is the lack of consideration to cultural factors that might play an important role in the efficacy of our method. Unlike the Gimpy CAPTCHA (Section 2.2.1) or Sounds (2.2.4), which rely on human perception to distinguish legitimate users, our KA-CAPTCHA focused on the acquisition of semantic information of pictorial data, and therefore presents a strong dependency of cultural aspects of human cognition. Consequently, our method introduces the possibility of access segregation not because a legitimate user does not possess the attributes that lack to a machine, but because he/she fails to share a sufficiently similar cultural background as previous KA-CAPTCHA users.

One other inherent limitation of our approach is its need of an initial knowledge base to allow the security verification of earlier test takers. Because our method verifies the legitimacy of incoming users with the same knowledge base that grows with the subsequent knowledge acquisition step, it becomes impossible for our KA-CAPTCHA to develop an early autonomous mechanism that is relieved of an initial amount of knowledge. On the other hand, while this is indeed a strict requirement, we argue that it is not a severe limitation to our approach's application because this initial amount of information can be constructed using a set of heuristic functions. The pictorial KA-CAPTCHA took advantage of existing heuristics for labeling images on the Web, while for the generation of textual transcriptions to videos we relied on the speech recognition state-of-the-art.

5.3 Lessons Learned

The development of this work presented a significant contribution to the field of collaborative knowledge acquisition. Nevertheless, several obstacles appeared during this research and played an important role in the final stage of this thesis.

A critical issue appeared with regards to the experiments that were to empirically validate the KA-CAPTCHA that acquires textual transcriptions to auditory media (Section 4.2). The aim of this experiment was to evaluate how frequently the visually-impaired individuals were able to succeed on this CAPTCHA. Consequently, we designed an experiment to be conducted with the help of volunteers from the Benjamin Constant Institute (IBC) for the Blind[1]. Unfortunately, this study faced a number of logistical problems, especially concerning security access between the UFF and the IBC network domains. Consequently, no significant result was available by the time this dissertation was written.

5.4 Future Work

As one direction for future work, we mention the necessity to perform an extended version of this experiment, where the performance of our method can be evaluated on the long run. This evaluation would enable designers to observe the robustness of our method against timely variations in the pictorial perceptions of the population.

Also, as was mentioned in Chapter 1, the premises of this thesis both represent possibilities for future work. First, one could analyze the KA-CAPTCHA usability and compare it to usual CAPTCHA's. Also, the estimation of the effect of heterogeneous

[1] www.ibc.gov.br

users to their success on the KA-CAPTCHA test is of significance. In particular, the case of a KA-CAPTCHA employed for semantic knowledge acquisition seems to be specially concerning because the very attribute that would distinguish legitimate users is culture-dependent, in detriment of sensorial attributes which are more universal. Therefore, the development of a KA-CAPTCHA that does not need context information (i.e. HTTP cookies, or the like) and still extracts semantic knowledge from users would be a challenging yet very interesting research agenda.

One interesting possibility for the extension of this method would be the application of different statistical measures to estimate the validity of ontology relations. For instance, the *lift* statistical measure might be useful to indicate how frequently a population recognizes some information as true given that some other piece of information was identified as truth.

Another direction for future work might be the design of different KA-CAPTCHA instances to knowledge acquisition on the Web. One promising area of investigation might be the elucidation of a Web site's affordance, i.e. the KA-CAPTCHA might be able to elicit information from the Web user about how (s)he interacts with a Web site, and possibly detect errors that might lead to a bad human-computer interaction.

References

Buchanan, B.G. and Shortliffe, E.H. (eds). Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project. Reading, MA: Addison-Wesley, 1984.

Chellapilla, K.; Larson, K.; Simard, P.; Czerwinski, M. 2005. Designing Human Friendly Human Interaction Proofs (HIPs). In *Proceedings of ACM CHI*, 2005, 711-720

Chklovski, T. 2003. LEARNER: A System for Acquiring Commonsense Knowledge by Analogy. In *Proceedings of Second International Conference on Knowledge Capture (K-CAP 03)* 4-12

Chklovski, T.; Gil, Y. 2005. An Analysis of Knowledge Collected from Volunteer Contributors. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*, 564-570. Menlo Park, Calif.: AAAI Press

Chklovski, T. 2007. Personal communication

Curtis, J., Cabral, J., Baxter, D. 2006. On the Application of the Cyc Ontology to Word Sense Disambiguation. In *Proceedings of the Nineteenth International FLAIRS Conference*, pp. 652-657, Melbourne Beach, FL, May 2006.

da Silva, B., Garcia, A. 2007. KA-CAPTCHA: An Opportunity for Knowledge Acquisition on the Web. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, pp. 1322-1327, Menlo Park, Calif.: AAAI Press

Feigenbaum, E. 1991. Expert Systems: Principles and Practice. Stanford University Technical Report. Stanford University. Stanford CA

Feigenbaum, E. 1996. How the “What” Becomes the “How”. Communications of the ACM May 1996/Vol. 39, No. 5 pp. 97-104

Google Press Center. http://www.google.com/press/zeitgeist.html

Gruber, T. 1993. A Translation Approach to Portable Ontology Specifications. Knowledge Acquisition 5(2), pp. 199-220

Holman, J., Lazar, J., Feng, J., D'Arcy, J. 2007. Developing usable CAPTCHAs for blind users. In *Proceedings of the 9th international ACM SIGACCESS conference on Computers and accessibility*, Poster Session, pp. 245-246

Lederberg, J. 1987. How Dendral Was Conceived and Born. In *ACM Symposium on the History of Medical Informatics*. Rockefeller University. New York: National Library of Medicine.

Lenat, D. 1995. CYC: A Large Scale Investment in Knowledge Infrastructure. Communications of the ACM 38(11):33-38

Mas-Colell, A.; Winston, M; Green, J. Microeconomic Theory. Oxford University Press, 1995

Miller, G. 1995. WordNet: A Lexical Database for English. Communications of the ACM 38(11):39-41

Mori, G.; Malik, J. 2003. Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA. In IEEE Computer Vision and Pattern Recognition, 134-141, Madison, Wisconsin, June 2003

Mundo Oi. http://mundooi.oi.com.br/

O'Hara, T., Bertolo, S., Witbrock, M., Aldag, B., Panton, K., Schneider, D., Salay, N., Curtis. J. Inferring Parts of Speech for Lexical Mappings via the Cyc KB. In *Proceedings of the 20th International Conference on Computational Linguistics (COLING-04)*, Geneva, Switzerland, August 2004.

Open Mind. http://www.openmind.org

Open Source Initiative. http://www.opensource.org

Panton, K., Matuszek, C., Lenat, D., Schneider, D., Witbrock, M., Siegel, N., Shepard, B. Common Sense Reasoning – From Cyc to Intelligent Assistant. In Yang Cai and Julio Abascal (eds.), Ambient Intelligence in Everyday Life, pp. 1-31, LNAI 3864, Springer, 2006.

Plataforma Lattes. http://buscatextual.cnpq.br/buscatextual/index.jsp

Rashid, A M., Ling, K., Tassone, R. D., Resnick, P., Kraut, R., Riedl, J. 2006 Motivating Participation by Displaying the Value of Contribution. In *Proceedings of ACM CHI* 2006, 955-958

Rui, Y., Liu, Z. 2003. ARTiFACIAL: automated reverse turing test using FACIAL features. In *Proceedings of the eleventh ACM international conference on Multimedia*, pp. 295 - 298

Singh, P., Lin, T., Mueller, E., Lim, G., Perkins, T., Zhu, W. 2002. Open Mind Common Sense: Knowledge acquisition from the general public. LNCS, 2519:1223-1237

Stork, D. 1999. The Open Mind Initiative. IEEE Expert Systems and Their Applications pp. 16-20

The CAPTCHA Project. http://www.captcha.net

Turing, A. (1950). Computing machinery and intelligence. Mind, 59, 433-460

von Ahn, L., Blum, M., Langford, J. How Lazy Cryptographers do AI. In Communications of the ACM, February 2004. Pages 56-60.

von Ahn, L., Dabbish, L. Labeling Images with a Computer Game. In *ACM Conference on Human Factors in Computing Systems*, CHI 2004. Pages 319-326.

von Ahn, L. 2005. Human Computation. PhD dissertation, Department of Computer Science, Carnegie Mellon University

von Ahn, L. Games With A Purpose. In IEEE Computer Magazine, June 2006. Pages 96-98.

von Ahn, L., Liu, R. and Blum, M. Peekaboom: A Game for Locating Objects in Images. In *ACM Conference on Human Factors in Computing Systems*, CHI 2006. Pages 55-64.

Yahoo! Mail, http://www.yahoo.com

Yokoi, T. 1995. The EDR Electronic Dictionary. Communications of the ACM 38(11):42-44