Biological Solutions for Engineering Problems: A Study in Cross-Domain Textual Case-Based Reasoning

Swaroop S. Vattam and Ashok K. Goel

Design & Intelligence Lab, School of Interactive Computing,
Georgia Institute of Technology Atlanta, GA 30332, USA
{svattam,ashok.goel}@cc.gatech.edu

Abstract. Textual Case-based Reasoning (TCBR) is a powerful paradigm within CBR. Biologically inspired design – the invention of technological systems by analogy to biological systems - presents an opportunity for exploring cross-domain TCBR. Our *in situ* studies of the retrieval task in biologically inspired design identified findability and recognizability of biology articles on the Web relevant to a design problem as major challenges. To address these challenges, we have developed a technique for semantic tagging of biology articles based on Structure-Behavior-Function models of the biological systems described in the article. We have also implemented the technique in an interactive system called Biologue. Controlled experiments with Biologue indicate improvements in both findability and recognizability of useful biology articles. Our work suggests that task-specific but domain-general model-based tagging might be useful for TCBR in support of complex reasoning tasks engaging cross-domain analogies.

1 Introduction

Textual case-based reasoning (TCBR) entails the use of unstructured cases in the form of textual documents (Weber, Ashley & Bruninghaus 2006). TCBR has become especially important with the advent of the Web that provides access to a large number of textual documents containing potential cases. A major question then becomes how do we access the right cases from the Web for a given query or problem? Thus, research into TCBR is closely intertwined with research on information retrieval (IR) and text mining (TM) (Rissland & Daniels 1996). According to Weber, Ashley & Bruninghaus (2006), the major differences between TCBR and IR/TM are that the former (1) is more explicitly interested in supporting complex reasoning, and (2) thus uses task-specific and domain-specific knowledge to access the right case for supporting the reasoning. For example, Burke *et al.* (1997) describe a technique that uses task-specific and domain-specific knowledge to answer FAQ questions in a specific domain; Bruninghaus & Ashley (2001) describe a different technique that too uses task-specific and domain-specific knowledge to access textual cases in the legal domain. Lenz (1998) describes knowledge layers for supporting TCBR; Raghunandan *et al.* (2008) propose evaluation measures for TCBR systems.

S.J. Delany and S. Ontañón (Eds.): ICCBR 2013, LNAI 7969, pp. 343–357, 2013.

In this work, we are interested in a related but slightly different question: how might TCBR work if the target problem and the textual cases are from different domains? That is, we are interested in *cross-domain* TCBR. We have encountered this problem in the context of *biologically inspired design (BID)* – the invention of new technological products, processes and systems by analogy to biological systems. The needed biological knowledge typically is found in the form of unstructured textual documents, typically on the Web. Due to its growing importance, we posit that BID presents a great opportunity for exploiting and exploring cross-domain TCBR.

In general, BID entails all the major tasks of CBR such as retrieval, adaptation, evaluation and storage (Kolodner 1993). In this paper, we focus on the *retrieval* task. Given a target design problem, one of the first tasks in the BID process involves finding the right biological system to emulate in order to generate a design solution. Designers, including expert designers, typically are novices in biology and thus are aware of only a small fraction of the vast space of biological systems. Thus, designers typically rely on external information environments such as the Web for finding biological cases relevant to their design problems. Most biological cases on the Web are available only in unstructured forms such as textual documents. Thus, the retrieval task in BID takes a query in a design domain, such as engineering, as input, and has the goal of returning as output textual documents in the domain of biology that are relevant to the query. The retrieval task is challenging not only because of the unstructured nature of the cases, but also because the retrieval process cannot rely on domain-specific knowledge and conventional techniques for retrieving textual documents lead to poor precision and recall.

Our *in situ* studies of the retrieval task in BID identified findability and recognizability of biology articles relevant to a design problem as major challenges (Vattam & Goel 2011); we define findability and recognizability below. To address these challenges, we have developed a technique for semantic tagging of biology articles based on Structure-Behavior-Function (SBF) models of the biological systems described in the article. We have also implemented the technique in an interactive system called Biologue. Controlled experiments with Biologue indicate improvements in both findability and recognizability of useful biology articles. In this paper, we describe the design and evaluation of Biologue.

2 Background

The growing movement of biologically inspired design (BID) or biomimicry views nature as a large library of sustainable designs that could be a powerful source of technological innovation (e.g., Benyus 1997). Recent examples of BID include the design of wind turbines inspired by the tubercles on the pectoral fins of humpback whales, and fog harvesting devices inspired by the arrangement of hydrophilic and hydrophobic surfaces found on the back of Namibian beetles, etc.

2.1 Related Research

Research on computational methods and tools for supporting BID can be categorized into three broad approaches. The first approach uses digital libraries of functional models of biological systems (Chakrabarti *et al.* 2005). For example, the DANE system provides access to a functionally indexed digital library of SBF models of biological systems (http://dilab.cc.gatech.edu/dane/; Goel *et al.* 2012). The difficulty with this approach is scalability: it takes expertise, time and effort to build such a library.

The second approach uses text mining techniques (Shu 2010), including syntax-level heuristics customized to BID (Chiu & Shu 2007), and enhanced by engineering-to-biology thesaurus (Nagle, Stone & McAdams 2010). Although more scalable than the first approach, this technique could be subject to the usual limitations of keyword-based search; the efficacy of this approach is still being explored.

The third approach uses semantic indexing for accessing biological information. For example, Biomimicry 3.8 Institute's AskNature (Biomimicry 2008) is a popular Web portal that provides access to a functionally indexed digital library of biology articles. Our work on TCBR in BID presented here takes a similar approach. We posit that our approach is more human-centered, emphasizing (1) first gaining a deep understanding of TCBR in BID as it naturally occurs in the real world, (2) grounding our system design in that understanding, and (3) rigorously evaluating our claims using controlled experiments.

This work both builds on and differs from our previous work on case-based design. In earlier work, we grounded the process of case-based design in SBF models of physical systems (Goel, Bhatta, Stroulia 1997), exploited TCBR for understanding design problems stated in natural language (Peterson, Mahesh & Goel 1994), and explored TCBR for acquiring SBF models of everyday devices from textual documents. In more recent work, we have formalized SBF models (Goel, Rugaber & Vattam 2009), conducted *in situ* studies of BID (Vattam, Helms & Goel 2008; Vattam & Goel 2011), and developed digital libraries of SBF models of biological systems in support of BID (Goel *et al.* 2012). In this paper we describe the design and evaluation of Biologue, our interactive system for addressing the findability and recognizability challenges of cross-domain TCBR in the retrieval task of BID.

3 Human-Centered TCBR in the Context of BID

We conducted our *in situ* studies of BID in the context of ME/ISyE/MSE/PTFe/BIOL 4740, a senior-level, interdisciplinary, project-based course at Georgia Tech. Yen *et al.* (2011) describe this course in detail. The two studies described below were conducted in Fall 2006 and Fall 2008 sections of the class, respectively. In these studies we observed a total of ten interdisciplinary teams of designers engaged in open-ended, semester-long BID projects that led to novel *conceptual* designs of technological systems such as the design of a new levee for New Orleans inspired in part by Iron Snail. While details of the studies can be found in other sources (Vattam 2012), here we summarize our findings related to TCBR in BID.

3.1 Characteristics

CBR in BID can be characterized as follows:

- *Cross-domain analogies*: The design problem originates in a design domain such as engineering but the cases for addressing the problem are in the domain of biology.

• *Compound analogies*: A single design solution may often require multiple biological source cases (Vattam, Helms & Goel 2008).
• *Textual cases*: Cases that are retrieved and used by the human designers mostly are found in the form of textual biological articles.
• *Cases distributed across multiple online environments*: Designers use a range of online information environments to seek biological cases, including (1) digital libraries like Web of Science, Google Scholar, ScienceDirect, etc., (2) online encyclopedia like Wikipedia, (3) popular life sciences blog sites like Biology Blog, (4) biomimicry portals like AskNature, and (5) general web search engines like Google.
• *Human-in-the-loop retrieval process*: In our observations, the designers used search results from a design query to formulate new queries for online search in an iterative process of formulating queries, searching online, finding biology articles, reading the articles, formulating new queries and so on. We call this process *interactive analogical retrieval* (Vattam 2012).

3.2 Challenges

We found that designers faced three major challenges in accessing biology articles on the Web relevant to their design problems (Vattam & Goel 2011): *findability*, *recognizabili*ty and *understandability*. These difficulties were encountered irrespective of the specific type of information environment used and made the retrieval process quite inefficient.

Findability**:** designers often went for long periods without finding a single relevant biological case in a retrieval process that typically extended over several weeks and often was tedious and frustrating for the designers. Thus, the relative frequency of encountering relevant articles containing biological cases was very low, suggesting that the match between the retrieval task and the information environment was not very good. A rough calculation suggests that designers spent approximately three person-hours of search time on the Web in order to find a single relevant article.

Recognizability**:** designers were prone to making errors of judgment about the true utility of articles that they encountered in the search process. In almost all online environments, search queries brought back a ranked list of search results (a set of articles). One important aspect of the search process was assessing and selecting promising articles from this list for further consumption. But, this decision had to be made based on *proximal cues* – hyperlink titles and snippets of text that are intended to represent the distal documents. In many instances, designers picked up on low-utility articles, only to realize later that it was not actually very useful (false positives). False positives lead to wasted time and effort (resource cost). Conversely, consider situations where designers might dismiss an article they encounter during the search as having low utility even though it might have contained a potential biological source (false negatives). False negatives represent lost opportunities.

Understandability: Since designers typically are novices in biology, they often have difficulty understanding the biological systems described in the textual documents they retrieve from the Web. While this challenge is covered in detail in other sources (Vattam 2012), here we focus on addressing the findability and recognizability challenges.

4 Addressing the Challenges of TCBR

Let us consider the issue of *findability*. According to the ACME theory of analogy (Holyoak & Thagard 1989), in order to retrieve source cases analogous to a target problem, the retrieval mechanism should simultaneously satisfy three constraints: semantic similarity (the overlap in terms of the number of similar concepts between the target and potential sources), structural similarity (the overlap in terms of the higher-order relationships between the target and potential sources), and pragmatic similarity (the overlap in terms of the pragmatic constraints or goals surrounding the target and potential sources). It is these three constraints acting simultaneously that distinguish analogical retrieval from other kinds of information retrieval mechanisms.

However, keyword-based search mechanisms found in common current online information environments support access to cases based on literal similarity (word-for-word matching), or at most semantic similarity to a limited degree, while ignoring structural and pragmatic similarity. As a result, each attempt at access can contain a large number of spurious articles that contain systems that are superficially similar to the target design as opposed to analogically similar. This results in low precision and recall.

Alternate methods of indexing and accessing biological articles in online environments may help address this challenge. Literature on case-based reasoning suggests guidelines for the alternate method (Kolodner 1993). (1) Indexing at storage time should anticipate the vocabulary the reasoner might use at retrieval time. (2) Indexing should use concepts and relations described at a level of abstraction that is justified from the perspective of the reasoning task.

In our *in situ* studies we found that designers' vocabulary used concepts and relations like functions, structures, physical principles, and operating environments. This vocabulary significantly overlaps with the vocabulary of Structure-Behavior-Function (SBF) models (Goel, Rugaber & Vattam 2009). Briefly, SBF models are a family of conceptual functional models. In SBF models, *Structure* pertains to components of a system; *Behaviors* describe causal processes or mechanisms in the system; and *Functions* specify outcomes of the system. In past work, SBF models have proved to be useful for design (Goel *et al.* 2012), understanding (Helms, Vattam & Goel 2010; Vattam, *et al.* 2011) as well as TCBR (Goel *et al.* 1996; Peterson, Mahesh & Goel 1994). Therefore, we posit that semantically indexing and accessing biology articles using concepts and relations derived form SBF ontology may better address the issue of findability in BID.

Now consider the second issue of *recognizability*. Information foraging theory (Pirolli 2007), which explains human information-seeking behavior in online information environments, claims that navigation towards useful information is guided by perception of *information scent* based on the proximal cues available in these environments. The issue of recognition errors in BID is attributable to the affordances - or lack thereof - of the proximal cues for accurately perceiving the information scent of the biology articles in the information environments.

One way to address this problem is by enhancing the proximal cues with additional information to help designers perceive the analogical similarity between the target design problem and the contents of the biological cases in the textual documents. We posit that enhancing the proximal cues of the distal articles by including visual overviews derived from the SBF models of the biological cases described in the articles may lower the rate of recognition errors.

5 Biologue

Biologue is an interactive information system that embodies the two main claims discussed in the previous section for addressing the challenges of the retrieval task in BID. Biologue represents a social approach to establishing an online corpus of biology articles annotated by their corresponding SBF models. It is based on the principle of social bookmarking (Sen *et al.* 2006) and is aimed to promote the sharing of biology articles in the BID community.

As one posts a reference to an article in Biologue, one can also manually add tags for annotating and organizing that reference. However, instead of keyword-based tags, the semantic tags in Biologue are linked to the ontology and schema of SBF models. As more and more people tag a particular reference, partially-structured SBF models of biological systems emerge in a socially-distributed fashion and get associated with that article. As Figure 1 illustrates, Biologue leverages these models to: (1) index an article and provided access to it based on the SBF schema and ontology, (2) offer visual overviews of the SBF models of biological cases described in the article.

As a use-case scenario, consider a situation where a designer, in the course of her day-to-day work, comes across a relevant online article on rat intestine and how that organ passively transports water across an osmotic gradient. The designer uses Biologue to: (1) bookmark this article in her personal library, (2) enter the article's bibliographic information, (3) tag the article with **`Function:Transport(Water)`**, where Transport **`Is-A`** Function in SBF ontology, and (4) share this tagged article with a teammate. The teammate reads the article and understands that the intestine achieves this function using the *three-chamber* mechanism, which uses a combination of *forward-* and *reverse-osmosis* principles. The teammate then adds a new tag to this article, **`Behavior:Three-Chamber-Method`**, and links it to the **`Function:Transport(Water)`** tag, where Behavior **`Is-A-Part-Of`** Function in SBF ontology. The teammate further adds two mores tags **`Principle:Osmosis`** and **`Principle:Reverse-Osmosis`** and links them to the **`Behavior:Three-Chamber-Method`** tag, where Principle **`Is-A-Part-Of`** Behavior in SBF ontology. Assuming that this article is read, tagged and retagged by many people, a conceptual model of how the rat intestine works emerges through negotiation and gets associated with this article.

Biologue implements an auto-complete feature to encourage tag reuse and minimize proliferation of user-generated tags, and a simple drag-and-drop interface for linking one tag to another and for linking tags to parts of the document.

Now, consider a use-case scenario where some other designer, completely unrelated to the first set who tagged the article, is trying to design a bio-inspired, energy-efficient, seawater desalination technique. This user logs into Biologue and proceeds to search the collection of articles in Biologue for a relevant biological source. Biologue currently allows users to search for articles based on features derived from SBF ontology, including Function, Principle, Structure, and Operating environment.

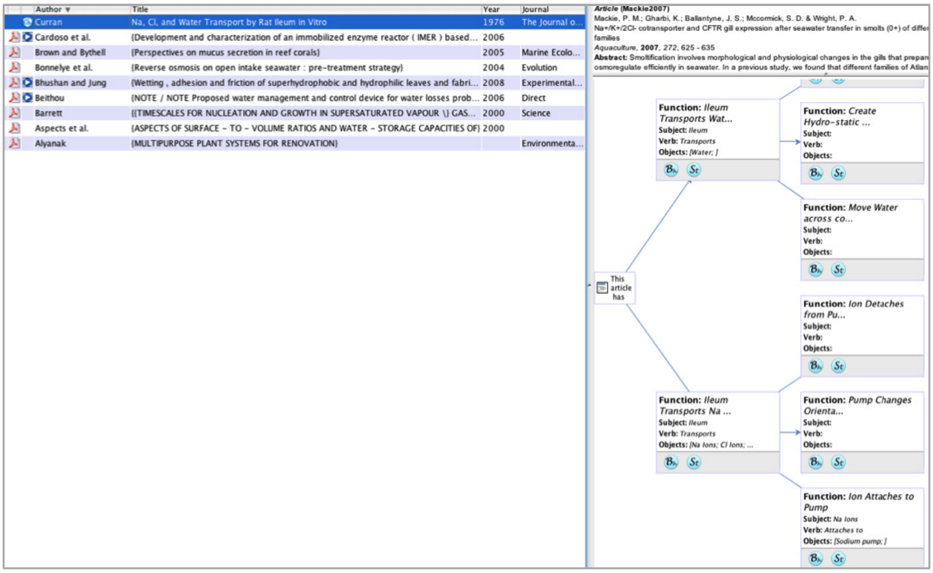

Fig. 1. Search results in Biologue; a selected article, and the visual overview of the SBF model tagging the article

Let us assume that this designer chooses to search the collection based on Function "remove salt from water" and the Principle "reverse osmosis". Because of the partial match across Principle, Biologue returns a non-ranked set of articles that includes the rat intestine article tagged by the first set of users. When this designer clicks on this article in the result set, she can not only view the traditional information that one would expect, but also the SBF model associated with the article as illustrated in Figure 1. Studying this model gives her a gist of how the rat ileum works from a designer's perspective, which also helps her make the analogical connection between the target desalination technology and the rat ileum system and allows her to decide whether it is worth pursuing the article. The model also gives her a high-level conceptual schema that she can refer back and forth to guide her development of understanding of the article. Upon reading the article and understanding how the rat ileum works in detail, she may choose to use this biological system as a source of inspiration and develops a novel desalination technique that removes salt from water by a combination of forward- and reverse-osmosis.

6 Findability Study

Hypothesis: In the context of BID, indexing and accessing biological articles using concepts and relations derived from their corresponding SBF models will lead to higher rate of findability when compared to keyword-based indexing and retrieval.

Procedure: This was a between-subject study conducted in Spring 2011. Sixteen subjects were recruited to participate in this study. A preselected BID challenge was presented to all subjects. The goal was to use Biologue to find as many articles relevant to the challenge as possible (and provide a rationale for their choice) within a stipulated amount of time. Eight subjects comprising the control group were given a version of Biologue where the articles were indexed and accessed using the conventional keyword-based approach. The other eight subjects in the experimental condition were given a version of Biologue where the articles were indexed and accessed using features derived from SBF representation. The performance of the participants on the search task was compared across the two groups.

Materials: The BID challenge involved a technology for solar thermal collectors, and included the design of (1) a bio-inspired reflective panel that could be fitted onto an existing absorber and was capable of dynamically changing its reflectivity, and (2) a bio-inspired feedback control system that regulates the temperature of glycol by regulating the reflexivity of the panel. This design challenge was an authentic problem attempted by a design team in one of the previous offerings of the BID course at Georgia Tech (Yen *et al.* 2011).

Two versions of Biologue were created for this study. In one version, the articles in Biologue's repository were indexed by keywords. Consequently, the articles were accessible only through keyword search. The search panel consisted of a single text box similar to Google and a search and a clear button. In the second version, the articles in Biologue's repository were indexed by concepts and relationships that were part of the SBF models associated with those articles. Biologue's repository had more than 200 articles in it. Fourteen designated articles that were known to be relevant to the target design challenge were included in the repository. If a participant's information seeking was efficient and accurate, then nearly all these 14 articles would be found and reported by the participant.

Data: Three kinds of data were collected for this study. First, the video of each participant's entire search process was captured using screen-capture software. This provided the bulk of the data to understand and analyze the retrieval process of the participants. Second, every article found by the participant in the course of their search was collected, along with his or her stated rationale for selecting that article. Third, participant demographic data was also collected.

Analysis: First, the participants in the two groups were compared to establish the equivalency of the two treatment groups. There was no statistically significant difference between the two groups with respect to participants' gender, their biology background, extent of design experience, extent of interdisciplinary research experience, or extent of use of scholarly articles in their work practices.

Second, the found-article data was analyzed to determine: (1) the total number of articles found by a participant, (2) the number of designated articles within that total number. The rationale provided by the participant for selecting an article was also analyzed to ensure that participants were meaningfully undertaking the task.

Third, the video data obtained for each participant was coded using a coding scheme. The origin of the coding scheme lies in information foraging studies conducted by researchers in the human-information interaction community (Pirolli, 2007). Their coding scheme coded and visualized the behavior of a person engaged in online information activity in the form of *web behavior graphs*. From these web-behavior graphs, collecting interesting statistics about the information seeking behavior in our experiments becomes possible. Two coders independently coded the videos (inter-coder reliability was 87.93%; in those cases where there was no consensus, the coding of the experienced researcher was included).

From the coded video data we were able to derive four retrieval performance characteristics for each participant: (1) *find period* (minutes per article found), (2) *mean between-patch foraging time* (minutes spent searching and navigating), (3) *mean number of regions foraged* (number of hops from one search page to another), and (4) *mean information yield per region* (ratio of actually relevant articles in an information region to all the articles encountered in that region). These four dependent variables, related to the findability issue, indicate the efficiency with which participants were able to search and retrieve the articles that they found.

Results: In the experimental condition, we expected that find period would be lower, that between-patch foraging time would be lower, that number of regions foraged would be smaller, and that information yield would be larger. Table 1 shows the actual results that confirm our predictions.

Table 1. Participant performance on the findability task

Treatment	Find period (avg)	Between-patch foraging (avg-mean)	Num. of regions foraged (avg-mean)	Yield per region (avg-mean)
Control	11.48 mins	11.68 mins	4.3	0.07
Experimental	5.85 mins	2.96 mins	2.45	0.212

The data collected in this study was submitted to a statistical significance test. An *a priori* power analysis suggested a sample size of 27 (power=0.7, d=0.5, α=0.05) for utilizing the t-test. Since our sample size of 16 was less than the recommended number, the data could not be subjected to the t-test. Therefore a non-parametric version of the t-test, known as *Mann-Whitney U* test was used.

There was *no* statistically significant difference between the two groups with respect to the number of articles found. The median *total articles found* in the control and experimental groups were 9 and 7.5 respectively; the distributions in the two groups did not differ significantly (Mann-Whitney U = 23.5, n1 = n2 = 8, P = 0.369 two tailed). Similarly, the median *total designated articles found* in the control and experimental groups were 3 and 3.5 respectively; the distributions in the two groups did not differ significantly (Mann-Whitney U = 27, n1 = n2 = 8, P = 0.595 two tailed).

Although the number of articles found by the two groups was comparable, there was a statistically significant difference between the two groups with respect to the cost incurred to find those articles.

In experimental condition, the average *find period* was 52% less as compared to the control condition and the difference was statistically significant (median1 = 92.5, median2 = 39, Mann-Whitney U = 1.5, n1 = n2 = 8, P = 0.001 two tailed). In the experimental condition, the average *mean between-patch foraging time* was 74.63% less compared to the control condition and the difference was statistically significant (median= 587, median2 = 157, Mann-Whitney U = 2, n1 = n2 = 8, P = 0.002 two tailed). In the experimental condition, the average *mean number of information regions foraged* was 43% less compared to the control group and the difference was statistically significant (median1 = 3.1, median2 = 1.8, Mann-Whitney U = 16, n1 = n2 = 8, P = 0.093 two tailed). Finally, in the experimental condition, the average *mean information yield per region* was 67% more compared to the control group and the difference was statistically significant (median1 = 0.063, median2 = 0.221, Mann-Whitney U = 5, n1 = n2 = 8, P = 0.005 two tailed).

Discussion: The above results suggest that both the treatment groups were similar with respect to the quantity and quality of articles that they found during this task. But, the experimental group took significantly less time and effort compared to the control group. In other words, for a similar output, the cost of retrieval in the experimental group was significantly lower. The differences between the two groups with respect to the four measurements taken together indicate that participants in the experimental condition more frequently encountered relevant biology articles when compared to the control condition. This implies that SBF-based indexing and access to biology articles has greater affordance for dealing with the findability issues, thus validating our proposed hypothesis.

7 Recognizability Study

Hypothesis: In the context of BID, enhancing proximal cues to include visual overviews derived from corresponding SBF models will lead to lower rate of recognition errors when compared to traditional proximal cues that do not include such overviews.

Procedure: This too was a between-subject study conducted in Spring 2011. The same sixteen subjects from Study 1 also participated in this study, but a sufficient time gap was provided between the two studies. A second preselected BID challenge was presented to all subjects. Biologue was used to then present a set of eight biology articles' proximal cues to the participants. The goal was to judge the relevancy of each of the eight biology articles for the given design challenge. The relevancy was reported on a five-point scale; the subjects were also asked to provide a rationale for their ratings. The articles were chosen such that four of them were relevant to the given design challenge and the other four were not relevant.

It is important to note that rather than entire articles, Biologue presented them with just proximal cues associated with articles. Eight subjects in the control group were given a version of Biologue that presented conventional proximal cues (containing information like the title, publication information, and an abstract of the article). The other eights subjects in the experimental condition were given a version of Biologue that presented them proximal cues which were additionally enriched with visual overviews derived from the articles' associated SBF models. All subjects were given a stipulated amount of time to complete the task. Because the researchers knew beforehand which four articles were relevant and which four were not relevant, we were able to calculate the extent of correct and incorrect classifications for each participant. We compared the classification accuracy of the participants across the two groups.

Materials: The BID challenge given to the participants in this study involved the design of a bio-inspired desalination technique such that: (1) the salinity of output fresh water should be fit for human consumption (specifically drinking), and (2) the energy footprint of the new technique must be less than the existing industry-standard techniques. This design challenge was subject to the following simplifying assumptions: (1) the feed water is already filtered and pre-treated to remove all other unwanted contents, leaving designers to deal with only pure saline water, and (2) the design will not actively control for other parameters like pH and alkalinity, free residual chlorine, boron, etc. The subjects were also given information about two existing industry-standard techniques for doing desalination, namely flash distillation method and reverse osmosis method. Some of the energy-related problems associated with the industry-standard techniques were also presented. To sum up, the subjects were given enough information so that a novice could be brought up to speed on the problem and had a rich enough mental model of the problem to be able to read an article and make a determination about its relevancy. They were also tested on their knowledge about this problem before they proceeded to perform the rating task. Again, this was an authentic problem addressed by one of the design teams (Team FORO) in the BID course at Georgia Tech in Fall 2008.

A total of eight biology articles were chosen for this task. These articles were selected from a pool of articles that Team FORO had researched in Fall 2008. Four of those articles were noted by the team as being relevant to solving the problem, and four as being irrelevant and leading to dead ends.

Biologue's repository for the purposes of this study consisted of only those eight articles. Biolgue for this study was instrumented such that as soon as a participant launched it, she would be instantly presented with a list of these eight articles (proximal cues only). This was meant to simulate a snapshot in the information seeking process where the seeker has just entered an information region and then needs to prioritize the order in which these articles would be visited based on the perceived relevance of each article to the target problem.

Two versions of Biologue were created for this study. In the control condition version, participants saw the traditional version of proximal cues, consisting of articles' title, abstract, and publication information. In the second experimental condition version, participants saw the SBF-augmented version of the proximal cues consisting of visual model overviews in addition to the other elements.

To minimize research bias, the primary researchers recruited another researcher to build the SBF models of biological systems discussed in the eight articles. This model builder had not encountered the desalination problem and was not aware of the purpose to which the SBF models would be put to use. Therefore, he could not introduce bias by tailoring the SBF models to match the desalination problem. These SBF models were then entered into Biolgoue and made available as part of the cues in one of the treatment groups.

All participants were required to rate the eight articles on a scale of 1 to 5. This was achieved by asking the participants to take an online survey when they were ready to rate the articles. The survey contained eight questions, one for each article they were required to rate. The rating was couched as a recommendation question: what would be their recommendation for the article to a team doing the desalination project on a scale of 1 to 5, where 1 represented “completely irrelevant (skip reading the article altogether)” and 5 represented “absolutely relevant (mimic the biological system in the paper and you will have solved the problem).” The middle value 3 represented “may be relevant, may not be relevant, can’t say which.”

Data: Participant demographic data was one of the data points used for this study. But the primary data for this study came from the online survey, which contained participants’ article classification data, including the rationale for their classification.

Analysis: Although the participants were the same in the two studies, their distribution across the treatment groups was different. Therefore, a group equivalency test had to be performed in this study as well. We found no statistically significant difference between the two groups with respect to participants’ gender, biology background, design experience, interdisciplinary research experience, and the use of scholarly articles in their everyday work practices.

Participant classification data, which was on a 5-point scale, was converted into a 3-point scale. A value of 1 or 2 was classified as “irrelevant,” a value of 4 or 5 was classified as “relevant,” and a value of 3 was classified as “unclassified.” For each participant and for each article, the participant classification was compared against the actual classification of the article. Based on this comparison, a determination was made as to whether it the classification was correct, false positive, false negative, or null (no classification).

Table 2. Participant performance data on the recognizability task

Treatment		Correct classification	False +ve	False -ve	Undecided
Control	avg	3.75	1.88	0.88	1.50
	stdev	1.04	0.83	0.64	0.93
Experimental	avg	5.50	0.38	0.88	1.25
	stdev	1.31	0.52	0.64	1.16

Results: Table 2 summarizes the data from this study. Again, the data collected in this study was submitted to Mann-Whitney U test. This data shows that in the experimental condition, the *average recognition error* was 41.18% less compared to the control condition and the difference was statistically significant (median1 = 4, median2 = 2.5, Mann-Whitney U = 9, n1 = n2 = 8, P = 0.015 two tailed).

In the experimental condition, the average *false positives* was 79% fewer when compared to the control condition and the difference was statistically significant (median1 = 2, median2 = 0, Mann-Whitney U = 4.5, n1 = n2 = 8, P = 0.003 two tailed). But, in the experimental condition, there was 0% difference in *false negatives*. Finally, in the experimental condition, there was 16.67% fewer *undecided classifications*, but the differences was *not* statistically significant median1 = 1.5, median2 = 1.5, Mann-Whitney U = 28, n1 = n2 = 8, P = 0.66 two tailed).

Discussion: The above results shows that in the context of this study, the group that worked with redesigned proximal cues containing SBF information did significantly better in terms of number of false positives. This difference in false positives heavily contributed towards the difference observed in the total error rate between the two groups. It is not clear why there was no change in the number of false negatives or undecided classifications. More fine-grained studies are required to determine the affordance of proximal cues vis-à-vis the different kinds of recognition errors.

8 Conclusions

We know from past work that TCBR is a powerful paradigm within CBR especially with the advent of the Web (e.g., Bruninghaus & Ashley 2001; Burke *et al.*, 1997; Lenz 1998; Rissland & Daniels 1996; Raghunandan *et al.* 2008; Weber, Ashley & Bruninghaus 2006). However, on one hand complex reasoning tasks such as BID require access to cross-domain analogies, and, on the other, conventional search engines on the Web do not support easy access of cross-domain analogies. This gap between the demand and lack of supply for cases creates both a challenge and an opportunity for TCBR: how to accomplish cross-domain TCBR?

Our *in situ* studies of the retrieval task in BID identified findability and recognizability as two of the main challenges for cross-domain TCBR. There are potentially several approaches that one can use to mitigate the identified challenges, including engineering a structured case-base of biological and engineering systems using a domain-general knowledge representation language, semantically tagging documents in a socially-distributed fashion, using natural language processing to automatically extract the semantic tags from textual documents, using machine learning to learn semantic tags for the biology articles. While in this work we have chosen to use the social semantic approach, in other threads we exploring alternative approaches.

We have developed a technique for model-based semantic tagging of biology articles based on SBF ontology of systems because SBF models have proved to be useful for design, understanding as well as TCBR. We have also implemented, fielded and evaluated the technique in an interactive system called Biologue. From the point of view of TCBR, this is a new approach not only because of the cross-domain nature of the cases, but also because of the social dimension of semantic tagging of the cases that is part of the solution. From the point of view of BID, Biologue represents a new class of technological aids; earlier technologies relied on fully structured knowledge-bases that typically entailed a high cost of knowledge engineering, or employed mostly-syntactic bottom-up natural language processing techniques that often are prone to poor precision and recall.

Our controlled experiments with Biologue indicate that when users do indeed adopt the social semantic approach, the improvements in retrieving cross-domain textual cases are significant. In particular, improvements were noticed with respect to both findability and recognizability issues. In our other work (Helms, Vattam & Goel 2010), we have found that SBF annotations also help design teams better *understand* biological systems.

Technological aids such as Biologue however raise the issue of reducing the chances of serendipitous encounter with fringe information that can sometimes lead to creative design solutions. The additional focus brought in by semantic search can inadvertently have the result that designers do not spend as much time browsing through articles and accidentally finding information that might be useful. Therefore, in practice, a more directed search feature in such a tool should be accompanied by other features that allow users to browse articles using a different set of criteria so that serendipity too is supported.

We know from past work that TCBR uses task-specific and domain dependent knowledge to retrieve and reuse textual cases for complex reasoning tasks (Weber, Ashley & Bruninghaus 2006). Our work indicates that task-specific but domain-general knowledge might be useful for TCBR in support of complex reasoning tasks engaging cross-domain analogies.

Acknowledgements. We are grateful to our research partners Michael Helms and Bryan Wiltgen for many discussions about biologically inspired design. We thank Professor Jeannette Yen, the coordinator of the ME/ISyE/MSE/PTFe/BIOL 4740 class that acts as a teaching and research laboratory for us. We are grateful to the US National Science Foundation that has generously supported this research through an NSF CreativeIT Grant (#0855916) entitled "Computational Tools for Enhancing Creativity in Biologically Inspired Engineering Design."

References

1. Benyus, J.: Biomimicry: Innovation Inspired by Nature. William Morrow (1997)
2. Biomimicry Institute (2008), Ask Nature – The Biomimicry Design Portal, `http://www.asknature.org/`
3. Burke, R., Hammond, K., Kulyukin, V., Lytinen, S., Tomuro, N., Schoenberg, S.: Question answering from frequently-asked questions files: experiences with the FAQ Finder system. AI Magazine 18(1), 57–66 (1997)
4. Brüninghaus, S., Ashley, K.: The role of information extraction for textual CBR. In: Aha, D.W., Watson, I. (eds.) ICCBR 2001. LNCS (LNAI), vol. 2080, pp. 74–89. Springer, Heidelberg (2001)
5. Chakrabarti, A., Sarkar, P., Leelavathamma, B., Nataraju, B.: A functional representation for aiding biomimetic and artificial inspiration of new ideas. AIEDAM 19, 113–132 (2005)

6. Chiu, I., Shu, L.: Biomimetic design through natural language analysis to facilitate cross-domain analysis. AIEDAM 21, 45–59 (2007)
7. Goel, A., Bhatta, S., Stroulia, E.: Kritik: An Early Case-Based Design System. In: Maher, Pu (eds.) Issues and Applications of Case-Based Reasoning in Design, pp. 87–132 (1997)
8. Goel, A., Mahesh, K., Peterson, J., Eiselt, K.: Unification of Language Understanding, Device Comprehension and Knowledge Acquisition. In: Proc. Cognitive Science Meeting 1996 (1996)
9. Goel, A., Rugaber, S., Vattam, S.: Structure, Behavior & Function of Complex Systems: The SBF Modeling Language. AIEDAM 23, 23–35 (2009)
10. Goel, A., Vattam, S., Wiltgen, B., Helms, M.: Cognitive, collaborative, conceptual and creative - Four characteristics of the next generation of knowledge-based CAD systems: A study in biologically inspired design. Computer-Aided Design 44(10), 879–900 (2012)
11. Helms, M., Vattam, S., Goel, A.: The Effects of Functional Modeling on Understanding Complex Biological Systems. In: Proc. 2010 ASME IDETC/CIE, Montreal, Canada (2010)
12. Holyoak, K., Thagard, P.: Analogical Retrieval by Constraint Satisfaction. Cognitive Science 13(3), 295–355 (1989)
13. Kolodner, J.: Case-Based Reasoning. Morgan Kaufmann Publishers, San Mateo (1993)
14. Lenz, M.: Defining knowledge layers for textual case-based reasoning. In: Smyth, B., Cunningham, P. (eds.) EWCBR 1998. LNCS (LNAI), vol. 1488, pp. 298–309. Springer, Heidelberg (1998)
15. Nagel, J., Stone, R., McAdams, D.: An engineering-to-biology thesaurus for engineering design. In: Proc. ASME 2010 IDETC/CIE, Montreal, Canada (2010)
16. Peterson, J., Mahesh, K., Goel, A.: Situating Natural Language Understanding in Experience-Based Design. IJHCS 41, 881–913 (1994)
17. Pirolli, P.: Information foraging theory: Adaptive interaction with information. Oxford University Press, Oxford (2007)
18. Raghunandan, M.A., Wiratunga, N., Chakraborti, S., Massie, S., Khemani, D.: Evaluation Measures for TCBR Systems. In: Althoff, K.-D., Bergmann, R., Minor, M., Hanft, A. (eds.) ECCBR 2008. LNCS (LNAI), vol. 5239, pp. 444–458. Springer, Heidelberg (2008)
19. Rissland, E., Daniels, J.: Using CBR to Drive IR. In: Procs. International Joint Conference on Artificial Intelligence, vol. 14, pp. 400–407 (1995)
20. Sen, S., Lam, S., Rashid, A., Cosley, D., Frankowski, D., Osterhouse, J., Harper, F., Riedl, J.: Tagging, communities, vocabulary, evolution. In: Procs. CSCW 2006, Banff, Canada, pp. 181–190 (2006)
21. Shu, L.H.: A natural-language approach to biomimetic design. AIEDAM 24(4), 483–505 (2010)
22. Vattam, S.: Interactive Analogical Retrieval: Practice, Theory and Technology, Doctoral Dissertation, Georgia Institute of Technology (2012)
23. Vattam, S., Goel, A.: Foraging for inspiration: Understanding and supporting the information seeking practices of biologically inspired designers. In: Proc. ASME DETC Conference on Design Theory and Methods, Washington, DC (August 2011)
24. Vattam, S., Helms, M., Goel, A.: Compound Analogical Design: Interaction Between Problem Decomposition and Analogical Transfer in Biologically Inspired Design. In: Proc. DCC 2008, Atlanta, pp. 377–396. Springer (June 2008)
25. Weber, R., Ashley, K., Bruninghaus, S.: Textual Case-Based Reasoning. Knowledge Engineering Review 20(3), 255–260 (2006)
26. Yen, J., Weissburg, M., Helms, M., Goel, A.: Biologically Inspired Design: ATool for Interdisciplinary Education. In: Bar-Cohen, Y. (ed.) Biomimetics: Nature-Based Innovation. Taylor & Francis (2011)