Want World Domination? Win at Risk!

Matching To-Do Items with How-Tos from the Web

Denny Vrandečić
Karlsruhe Inst. of Techn. /
Univ. of South. California
denny.vrandecic@kit.edu

Yolanda Gil
Information Sciences Inst.
Univ. of South. California
gil@isi.edu

Varun Ratnakar
Information Sciences Inst.
Univ. of South. California
varunr@isi.edu

ABSTRACT

To-Do lists are widely used for personal task management. We propose a novel approach to assist users in managing their To-Dos by matching them to How-To knowledge from the Web. We have implemented a system that, given a To-Do item, provides a number of possibly matching How-Tos, broken down into steps that can be used as new To-Do entries. Our implementation is in the form of a web service that can be easily integrated into existing To-Do applications. This can help users by providing them with an approach to tackle the To-Do by listing smaller, more actionable To-Dos. In this paper we present our implementation, an evaluation of the matching component over two sets of To-Do corpora with very different characteristics, and a discussion of the results.

Author Keywords

to-do, personal information management, tasks, how-to, intelligent assistant

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*Interaction styles*

General Terms

Design

INTRODUCTION

To-Do lists are widely used tools for personal task management. Many To-Do list managers are readily and freely available on the Web or as desktop or mobile applications. Although studies have analyzed To-Do lists, their usage patterns, and have come up with desiderata (e.g. [2, 9]), no system to date has been developed to provide comprehensive, intelligent assistance to the user. Popular methodologies to manage To-Dos, like GTD (Getting Things Done [1]) or the Pomodoro technique [3], suggest to make sure that items on To-Do lists are "actionable", i.e. that they can be done and checked off in a reasonable timeframe. Whereas there are plenty of software solutions supporting the management of To-Do lists in general – making it easy to enter them and check them off, synchronizing them between devices, providing reminders, creating repeated To-Do items, etc. – there is not yet much support to break down a non-actionable item or a project into a set of smaller, more immediately actionable items. At the same time, there is a growing number of Websites offering guides and manuals towards tackling an enormous number of tasks. We often use a search engine to find help with getting a specific task done.

In this paper we present HowToDo,[1] a prototypical implementation of the idea of matching To-Do entries with How-To knowledge gathered from the Web. We have taken samples of 480 To-Do entries from two To-Do corpora available to us, and tagged them manually to see how well our system matched them to the available How-To articles. We follow with a discussion of the results.

Different types of To-Do entries require different types of assistance. Our previous work has been focusing on two of the types. The first type is matching To-Do entries to the capabilities of available agents, like *"schedule meeting with John"* or *"email Catherine about XTime"*. We have developed Beamer, a system using a paraphrasing approach to match To-Do entries and connect them to relevant agents within the TOWEL system [4]. Our analysis of To-Do lists in the CALO project has revealed that 14% of such To-Dos can be matched to agent capabilities in the system [8], and Beamer achieved an accuracy of 88% when matching [7].

HOW-TO KNOWLEDGE ON THE WEB

There are many different resources on the Web to find knowledge on how to perform specific tasks. In order to create an initial set of a relevant number of How-Tos for our How-To repository, we have turned to one of the bigger websites directly aiming at collecting How-To knowledge, wikiHow.[2] wikiHow is a community-built site collecting How-To manuals. It is multi-lingual, even though the English part vastly outnumbers the rest. wikiHow holds currently more than 85,000 articles written collaboratively by more than 200,000 registered users under a Creative Commons license. It has more than 25 million unique visitors per month.

Articles on wikiHow are usually uniformly structured: after a short introduction we find a list of steps, which in turn are

[1] `http://howtodo.isi.edu`
[2] `http://www.wikihow.com`

IUI'11, February 13–16, 2011, Palo Alto, California, USA.

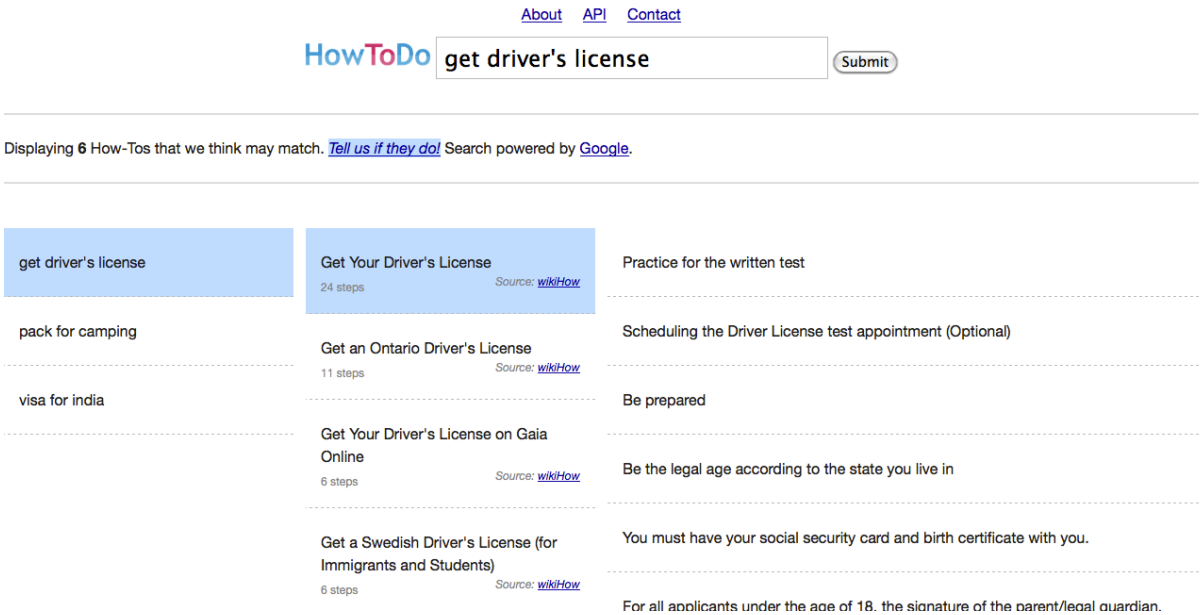

Figure 1. Screenshot of HowToDo, showing How-To data from the Web as a list of possible To-Dos. On the left we see some user-entered To-Do entries, the middle row shows a list of possible matches for the currently selected entry, and the right row displays the steps for the selected match.

often further detailed. We use the steps to offer the create the more detailed subtasks. The articles end with a few general notes, links to external resources, and also a list of related wikiHow articles. Some of the steps themselves may further link to external resources or further wikiHow articles. A How-To usually has half a dozen to a dozen steps.

TURNING HOW-TOS INTO TO-DOS

Users are expected to use the UI of the To-Do application they are already using and are familiar with. We provide a web service that can be integrated into the user interface of a given application without the user having to exit their application. Thus To-Do management applications can enrich themselves with the results from the service provided by HowToDo.

The matching component takes an incoming To-Do entry and retrieves and ranks fitting How-To articles from the How-To repository. The current matching algorithm is based on an web search service. The results are then further filtered and processed. Whereas the resulting steps are often merely an extract from the original How-To article, the search uses the complete full text and the article's link structure in order to calculate the ranking. We provide the HowToDo website for a convenient access to the web service (see Figure 1).

The HowToDo web service is accessible as a RESTful interface providing results in JSON.[3] The query to HowToDo consists of the To-Do entry, and HowToDo returns a ranked list of possibly matching How-Tos and the steps for each of these How-Tos. We imagine an application to let the user select a fitting match, and then to either just read the steps, add them as subentries to the queried entry, or replace the original To-Do item with the selected steps. This allows the user to remain within the context of their To-Do application instead of having to switch to a Web search. We expect a custom-made interface combining To-Dos and How-Tos to be particularly beneficial on mobile devices, which are often used to provide ubiquitous access to a user's To-Do list.

EVALUATING THE HOW-TO MATCHING

We have used two datasets with To-Do entries to perform our evaluation. The first dataset comes from the CALO project [4]. CALO provided the TOWEL interface to enter and manage To-Do entries [4]. The To-Do entries we are using for our evaluation were collected with no automatic support when being entered (like auto-completion or suggestions). All users were participating in the CALO case study and thus part of the same office environment. 2,400 To-Do entries were collected from a dozen users [8].

The second dataset was collected by a Facebook app for managing To-Dos.[5] All Facebook users can log in to the application and create To-Dos there, and, optionally, share the entries with their friends. Our dataset consists of 1,500 To-Do entries by 325 users, of which about 100 were regular users of the application [6].

Some examples from both corpora are given in Figure 2. There were a number of major differences in how the two corpora were collected: the CALO corpus was collected and accessible only when the user was working at their computer. The Facebook data on the other hand was only available from within Facebook and thus in many cases not from the office workplace. Otherwise, Facebook can be accessed

[3] `http://howtodo.isi.edu/apidoc.php`

[4] http://caloproject.sri.com

[5] `http://apps.facebook.com/todo-lists/`

CALO
plan Washington Meeting
new use cases for Mary
Prepare brief for CALO mgmt review
Think about extending John's algorithm
figure out signitures
CALO ui - window should auto resize
Facebook
Buy christmas gifts!
hair cut
Call USAA to discuss mileage
Joanna's birthday party
Find speaker
Buy Dardanelles Gun #2 - Level 75

Figure 2. Examples from the two To-Do corpora used in our evaluation. Note that this is not a random sample but rather selected to emphasize the differences (names have been changed for privacy reasons).

through mobile devices and from anywhere with access to the Web. The users of the Facebook application were self-selected, and the application was published on Facebook and made available open-ended, whereas the CALO data was collected during two annual evaluation periods, each lasting several months.

As a result, the entries in the two corpora exhibit different characteristics as illustrated in Figure 2. Whereas most of the entries in CALO are related to office work, the Facebook data encompasses mostly leisure, study, and general activities. Both corpora may contain extremely obscure entries that can be only understood within a specific context.

Before evaluating the HowToDo system, we formulated our hypotheses as follows: 1) only a small, but still interesting fraction of To-Do items will be matched by HowToDo, 2) the system will perform significantly better on the Facebook data, 3) the system will rather return no How-Tos at all than useless How-Tos, and 4) the overlap between the set of To-Do items that are understood by Beamer and the set of To-Do items that can be assisted with HowToDo is small.

We took samples from both datasets. For the CALO dataset, we took a sample that was previously used for the evaluation of Beamer [7] in order to also evaluate the coverage of Beamer in comparison to the coverage of HowToDo. From the Facebook data we collected 300 random entries (only omitting entries not in English). We processed every entry in the corpus and tagged the results manually by either stating that all suggestions were bad matches, that no suggestions were made, or the rank of the first matching suggestion (sometimes, several suggestions matched).

An overview of the results is given in Figure 3. *Good* gives the percentage of true positives: it means that there was a match between the To-Do entry and one of the top five ranked suggested How-Tos. *Top* indicates the percentage of entries where the top ranked How-To was the correct result (thus *Good* contains *Top*). *Bad* lists the percentage of false positives: it means that the system did return results, but

	CALO	Facebook
Good	16%	36%
Top	*4%*	*24%*
Bad	27%	27%
None	57%	37%
n	*180*	*300*

Figure 3. Matches on the evaluation samples.

none of the top five ranked fitted to the To-Do entry. We have been rather strict: the system often suggested thematically related, interesting How-Tos, but not one actually breaking down the given task (like matching *"Banana cupcake"* with *"Make Vegan Banana Muffins"*). Finally, *None* lists the percentage of entries where no How-To was suggested.

As a further result, we compared the set of To-Do entries tagged in the CALO corpus as having the potential of being automatically assisted by agents, the latter being 14% of the entries [8]. Only 3% of the sample was both amenable to automatic assistance and matched by HowToDo with a fitting How-To. This shows that the two approaches towards assisting the user with their To-Do lists – automatizing the tasks, as suggested within CALO, and matching with How-Tos, as suggested in this work – are highly complementary, pointing towards a hybrid solution for optimal coverage.

DISCUSSION

Our first hypothesis was correct, and the evaluation actually exceeded our expectations in the Facebook case. This is a promising result, but we have to remember that we only evaluated the matching of the To-Do entries to the How-Tos, and not the usefulness of the suggested, individual steps.

The two datasets exhibit very different properties: in the CALO case, the number of false positives outnumbered the good suggestions, and even with the good ones, most of the time the user had to look through a number of suggestions before finding the correct one. In the Facebook case, though, not only the number of matched entries was significantly higher, but also in most cases, when there was a match, the match came up as the top-ranked suggestion. This confirmed our second hypothesis: this approach works much better on the Facebook data.

Even though the number of false positives is comparable, the false positives in the Facebook case were still often considered as interesting to the query, related to the topic at hand, or outright funny (matching *"world domination"* to *"Win at the Game Risk"* or matching *"create task to-do list app"* with a work-saving *"Use Remember the Milk"*). In turn, in the CALO case the false positives seldom had anything to do with the task, often providing an irritating distraction for the task – a consequence much less serious for users using a Facebook app to manage their To-Dos. We assume that this is due to the particular repository of How-Tos we have used. CALO was collected in an office setting, and it contains many references to project names, persons, specific forms, and procedures. Often projects are given names that are also common words – this was not the case in our corpus,

thus leading to the high rate of 57% of entries not matched at all. We assume that if the projects were given names that are also common English words, we would have had a drastic increase in the percentage of false positives. But even without this further complication, our third hypothesis was shown wrong. The number of false positives was very high.

As we have seen the fourth hypothesis was shown correct on the CALO data: Beamer and HowToDo are indeed complementary in their potential.

Another shortcoming of the current system is its lack of parametrization. Whereas wikiHow provides How-Tos like *"Move"* or even *"Move to a Bigger City for a New Job"*, not many cities have a dedicated article on how to move there.[6] Again, a knowledge-based backend could expand the query by recognizing parts of the query that could be regarded as parameters and provide with a more general match, if a specific one is not found.

RELATED WORK

There are numerous implementations of To-Do lists, many of them available online. Two of them are of particular interest with the extension described in this paper. First, Life Balance provides so-called *exchange files* for tasks.[7] They basically provide How-Tos that can loaded into the To-Do list. A similar feature was integrated into our own Facebook app, called *technique*, even providing parametrization. Both cases differ from our approach in that the repository of How-Tos is created manually, and thus has a problem of providing enough How-Tos to entice users to use it. By using wikiHow, we are able to match an interesting number of To-Do items.

Although the work as described in this paper is based on To-Do applications, we are aware that many people use their email inbox as a replacement for a separate To-Do list application [10]. But with systems like RADAR that can identify tasks embedded in emails as well, and then sends these tasks to the To-Do management system [5], we can basically provide a similar assistance as if users were using a To-Do application themselves.

CONCLUSIONS

This paper presents an approach to assist users with their To-Dos by matching To-Do entries with How-To articles automatically and use the latter to support users with their task management. We have implemented the approach and provide the implementation as a web service, so that it can be integrated in external systems. We also provide a website, HowToDo, as a demonstrator of the service, so that the idea and implementation can be easily tested.[8]

We have evaluated the implementation using two different corpora, one from the project CALO, one from a Facebook app. We have found that the system performs very differently on the two corpora. In the Facebook setting, the system is able to provide users with support in one out of three To-Do entries, whereas in the CALO corpus the ratio drops to one out of six.

The current implementation is based on the repository of How-Tos created by the wikiHow community. We eventually envision the web containing a huge collection of How-Tos, from the very specific – like *"deliver furniture to 3544 Jones Av, Apt 27"* – to generic ones, like the ones we find today. These How-Tos can be further annotated with rich metadata, helping to discover and parametrize them, and even connect them to services on the Web so that some parts of the To-Dos can be resolved with a single click. HowToDo is a prototypical implementation of a small part of this puzzle.

ACKNOWLEDGEMENTS

The work presented in this paper is funded by the National Science Foundation (NSF) under Grant number IIS-0948429.

REFERENCES

1. D. Allen. *Getting Things Done: The Art of Stress-Free Productivity*. Penguin, 2001.
2. V. Bellotti, B. Dalal, N. Good, P. Flynn, D. G. Bobrow, and N. Ducheneaut. What a to-do. In *CHI '04*, pages 735–742. ACM, 2004.
3. F. Cirillo. *The Pomodoro Technique*. Lulu, 2006.
4. K. Conley and J. Carpenter. Towel: Towards an intelligent to-do list. In *AAAI 2007 Spring Symp. on Interaction Challenges for Int. Assistants*, 2007.
5. A. Faulring, B. Myers, K. Mohnkern, B. Schmerl, A. Steinfeld, J. Zimmerman, A. Smailagic, J. Hansen, and D. Siewiorek. Agent-assisted task management that reduces email overload. In *IUI '10*, pages 61–70. ACM, 2010.
6. Y. Gil, P. Groth, and V. Ratnakar. Social task networks. In *AAAI Fall Symposium Series on Proactive Assistant Agents*, Arlington, VA, USA, November 2010.
7. Y. Gil and V. Ratnakar. Automating To-Do lists for users. In *AAAI'08*, July 2008.
8. Y. Gil and V. Ratnakar. Towards intelligent assistance for to-do lists. In *IUI'08*, Canary Islands, Spain, January 2008.
9. G. R. Hayes, J. S. Pierce, and G. D. Abowd. Practices for capturing short important thoughts. In *CHI Extended Abstracts*, pages 904–905, 2003.
10. S. Whittaker, V. Bellotti, and J. Gwizdka. Email in personal information management. *Communications of the ACM*, 49(1):68–73, 2006.

[6] Even though *"move to amsterdam"* not only yields the helpful How-To on *"Move to the Netherlands"*, but also the very practical and relevant *"Change the Timezone in Linux"* and *"Talk to a Boy While Bicycling"*.

[7] `http://www.llamagraphics.com/resources/exchange%20files/index.html`

[8] `http://howtodo.isi.edu`