Research Networking Programmes

Short Visit Report: *Transactive Search Evaluation*

Search engines have been the foundation of Internet for the last 2 decades. Recently though, specialized platforms have started to drive a conspicuous volume of Web traffic (e.g., Amazon as the standard-de-facto for product search), showing how a generalist approach to search is becoming more and more inefficient.
In my research, I build novel platforms to serve the information needs revolving around the memories of a user, i.e., memory queries denote queries where the user is trying to recall from his/her past personal experiences. Neither Web search nor structured queries can effectively answer this type of queries, even when supported by Human Computation solutions. Last year we proposed a new approach to answer memory queries that we call Transactive Search: the user requested memory is reconstructed from a group of people by exchanging pieces of personal memories in order to reassemble the overall memory, which is stored in a distributed fashion among members of the group. Experimental results comparing Machine Learning, Social Network, and Human Computation techniques showed that Transactive Search significantly outperforms the effectiveness of existing search approaches for memory queries [1].

Evaluation of Transactive Search has to go beyond the evaluation of traditional Information Access Systems. During my visit at the University of Sheffield (Information School), I worked with my long-time research collaborator Gianluca Demartini. Our goal was to leverage the expertise in Data Science and Psychology offered by the University of Sheffield to understand how to measure human performance and, thus, to improve efficiency of the transactive search queries (i.e., runtime execution, leveraging user profiling and incentives to the group members, etc.) and their expressive power (i.e., support for Boolean queries, point queries, etc.).

[1] Michele Catasta, Alberto Tonon, Djellel Eddine Difallah, Gianluca Demartini, Karl Aberer, Philippe Cudré-Mauroux: TransactiveDB: Tapping into Collective Human Memories. PVLDB 7(14): 1977-1980 (2014).

Background

Some questions can only be answered by humans. A common approach is to post a question openly on the Web and wait for the right person to answer it. However, in an organizational setting social connections can be leveraged to find the person who can answer a question.

During my week in Sheffield, we started to build on top of the idea of transactive memories, that is, the fact that some knowledge is shared across individuals of a community to design and evaluate question routing in an interaction network. We compared proposed approaches in terms of question cost against standard approaches such as asking the question to all members of the community and to machine-support approaches such as expert finding.

Related Work

Transactive Search aims at leveraging transactive memories of a community of people sharing the same experience. It has been tested and compared against machine-based as well as hybrid human-machine and human computation approaches.

In our work we instead focus on point queries who can be answered by one and only one member of the transactive memory group.

Psychology studies showed how the communication process affects the retrieval process in transactive memories and how having an established relations among the transactive memory group members helps in better retrieving information from memory. In our work we leverage individuals with no prior relation but we experimentally compare different levels of knowledge overlap between group members.

A related field of work is also expert finding which proposes machine-based techniques to construct and index expertise profiles of individuals, for example, in a company. On top of such index we can search for expert and, ask the question directly to them. This limits the cost of asking questions to all members of the community but, for some question, it is not very effective as we show in our experiments. Popular expert finding approaches use information retrieval techniques based on language models.

Also related to this is the task of building teams that contain diverse experts able to solve different prediction tasks. As compared to them we rather aim at studying how the crowd can best identify the single person who can answer the query.

In a crowdsourcing setting, expert finding maps to the problem of predicting worker's accuracy on a task in order to assign the task directly to the most accurate worker. For example, Pick-a-Crowd (Difallah et al.) model workers in the crowd by the social media activity to understand skills and interests and assign tasks accordingly. Other works look

at the performance on past tasks to predict the accuracy of workers on future tasks showing how this performs better than randomly assigning a task to a worker.
As compared to these works, we rather let the crowd self-route tasks based on the knowledge they have about other members of the transactive memory group.

In a model for prediction tasks and delegation in a network of humans, authors show by means of simulation how routing rules perform effectively. On the other hand, we run experiments over a crowdsourcing platform showing how human-driven task routing can be efficiently answer point queries.

Lastly, questions and open tasks can be routed to experts within an organization. However, sometimes experts asked for consultation are not able to answer the question. In these cases, they delegate it to someone else within the organization who they believe is able to solve the task. This setting has been studied by Huan Sun et al. (SIGKDD2014) who analyzed how experts behaved and proposed a formal model for it. As compared to this work we rather focus on a crowdsourcing platform setup where knowledge about peers is limited and controlled by the delegation system that we propose. We show how questions that can be answered only by humans over a large digital collection can be efficiently processed by means of crowdsourcing and transactive search.

Task allocation in groups has been studied by "Power in Unity: Forming Teams in Large-scale Community Systems". Authors propose efficient and effective algorithms to assign tasks to groups and deal with the computational complexity of the problem. In our work we instead focus on support the self-organization of a crowd in efficiently routing questions to the right member.

Description of the work carried out during the visit

The first few days have been dedicated to a thorough literature review (as per the section above), plus a few brainstorming sessions to define the problem we wanted to address. While Q&A platforms (both user and Web-content driven) have been thoroughly studied, in our work on "Transactive Point Queries" we want to discover how groups of people can be orchestrated efficiently to answer questions that only one or few in the group can know about.

To this aim, we designed 2 experiments to be run on an anonymous crowdsourcing platform (e.g., CrowdFlower) in order to study our orchestration approaches.

The first experiment will mimick the famous experiment that proved how couples are more efficient to memorize a list of words, as they know the expertise area of the other partner (therefore they tend to remember more about their own expertise area). As "experiment 0", therefore, we will try to perform a similar experiment, but with the anonymous crowd. Each worker will be assigned a certain expertise profile, and we will assess if we can enforce an ephemeral transactive memory in the crowd.

Once this initial experiment succeeds, we will leverage the anonymous crowd to answer semantically-advanced questions on a large text corpus (i.e., questions that cannot be answered with state-of-the-art NLP techniques). Workers will be assigned expertise profiles, and we'll run different experiments to assess how rich the profile should be to strike a balance between the question routing efficiency, and the complexity of the supporting software system/UX.

To this purpose, we already started to work on a public dataset on Reading Comprehension, released by Microsoft:
http://research.microsoft.com/en-us/um/redmond/projects/mctest/

We also designed our crowdsourcing experiment to minimize the amount of money spent, while at the same time texting different "point query" routing strategies in a single run.

Future collaboration with host institution

During my stay in Sheffield, I haven't only strengthened my collaboration ties with Dr. Gianluca Demartini, but I've also had a long meeting with Dr Tom Stafford -- he's a world renowned psychologist (lecturer in Psychology and Cognitive Science at the University of Sheffield and book author) who gave us valuable insights on our study methodology, and showed genuine interest for our research direction.
As a matter of fact, we have already decided to collaborate together on our first publication on "Transactive Point Queries".

Projected publications

After laying the foundation of our paper during the first days, I spent the rest of my short visit to setup the first experiments. Realistically, we will require another couple of months to run all the experiments and write the paper we've currently in mind. As such, we decided to target CHI 2016, one of the top conferences in the "Human Computation" field. The tentative deadline should be September 2015.

Other comments

I benefited greatly from the generosity and flexibility of the ESF short visit program. As a matter of fact, I've already recommended it to all my colleagues and labmates in EPFL. At the same time, I hope I will be able to benefit from it again in the near future, as I'm in a stage of my career when maintaining my collaboration ties is fundamental.