

Pedagogy: a social networking approach to knowledge discovery

Mohammad Raza, Microsoft Research Cambridge, February 2011.

The construction of semantically structured, accurate, and comprehensive knowledge bases remains an important but difficult problem. Acquiring the wealth of human knowledge in a machine-readable format would have invaluable benefits for semantic web services and AI applications in general. In this report I outline a new approach that aims to address this challenge by harnessing the power of social networking.

Background and related work

A number of projects have made significant progress toward the goal of knowledge discovery, and these may be broadly classed into two categories. One approach is to develop methods for extracting knowledge from latent sources that currently exist on the web in human-readable form, such as Wikipedia. Examples of such projects include Yago [14,15], DBPedia [13], True Knowledge [12] and Powerset's "quick facts" [11]. The challenge in this case is to develop sophisticated extraction techniques that can both guarantee accuracy and also deliver comprehensive coverage. However, existing methods need to trade one for the other, and the problem seems to be AI complete in general.

The other approach, which we take here, is based on the notion of *crowd-sourcing*. The idea is to elicit knowledge directly from people, using the distributed power of the internet. Examples of projects directly employing this method include the Open Mind Common Sense project [1,2] and Freebase [16]. Whilst this approach by-passes the need for sophisticated extraction techniques, the main challenge here is to motivate people to input knowledge in a structured format, and to ensure the quality of the acquired knowledge in terms of both accuracy and coverage.

Some steps toward this goal have been taken by the so called *Games With A Purpose* [5,7], such as Verbosity [9] for collecting common sense relations between words, and the ESP game [8,10] for labelling images. These are simple online games with the goal of eliciting knowledge from players as a by-product of game play. In Verbosity, for example, a player is matched with an anonymous partner, and at every round, one player is given a random dictionary word which he must describe to his partner to help him guess it. This is done by giving hints about how the word is related to other dictionary words, using a fixed set of relations such as "is a" or "looks like". Both players get a certain number of points if the partner guesses the word. The aim is to acquire common sense knowledge by collecting semantic relationships between dictionary words in this way.

While having certain advantages over the direct approach, games such as Verbosity have significant short-comings and can even be counter-productive because the goal of the player is different from the goal of knowledge acquisition [6,17]. For example, accuracy is compromised as players often misuse the system to win points, such as giving hints about syntax rather than semantic relationships, or using relations incorrectly to help their partner's guesses. Coverage is also limited, because players tend to input minimal and

obvious information that is needed to guess the given concept. Since players must be anonymous (to prevent cheating), the approach is limited to 'common sense' and cannot deliver expert or encyclopedic knowledge. Apart from the issues in the quality of extracted knowledge, motivation also remains a difficulty: there is little sense of reward or competition, the game is short term and inconsequential, and restricted to anonymous players.

A social networking approach

We propose a novel approach to address the problem of motivation and knowledge quality in crowd-sourcing techniques, which uses the power of social networking to implement a direct method of knowledge acquisition from people. In the style of a Facebook game such as Farmville, we propose a social networking game based on the idea of 'teaching' in a virtual world. Players in a social network are part of a virtual world in which each player is in charge of a number of 'students' or 'apprentices', perhaps in some institution such as a school or other organisation belonging to the player. Students are virtual entities that exist in the game, and exhibit limited intelligence, such as the ability to learn knowledge, ask questions, interact with one another, and be tested about the knowledge that they have been taught. Each player can teach his students about any topics of his own choice, answer their questions, and design tests of varying difficulty to test his own students or the students of his opponents. The goal is to improve one's institution and students by excelling in the quality of education that is provided by one's institution, in terms of the performance of students in tests designed by opposing players, and in terms of the quality of the tests that one has designed.

Various forms of social collaboration are possible e.g. if player's choose to teach each other's students, or pay other's students to teach their own about a specific topic, or design tests together. Motivation would be driven in the context of a virtual economy, very much in the style of Farmville, so that income is generated by the quality of education that one provides, and can be spent on various products or services from a 'marketplace' to improve one's institution or students. For example, if a player's institution has received a good reputation for the quality of their tests, then they may charge a high fee for any student to take their test. In return, the student will get a reputable qualification for passing that test, increasing their odds of being hired to teach in the future.

A concrete instantiation

With the conceptual basis described above, we now outline a minimal concrete instantiation of the game, which may be further developed in various directions.

Knowledge representation. The knowledge base of each student can be represented as a set of RDF triples. For each subject that the student is taught, the player can enter a number of (relation,object) pairs to described how the subject is related to other concepts. For example, a player may choose the subject 'Barack Obama' and enter the tuples

- (president of, America)
- (member of, Democratic Party)
- (alma mater is, Harvard University)
- (born in, 1961)

The aim of the player is to enter as many triples as possible to improve the student's chances on tests designed by fellow players. The size limit of a student's knowledge base may increase with respect to time to simulate natural growth of the student, much in the same way as crops grow in Farmville. This growth can also be more strongly tied to how much the student is taught and how well it performs, to reward players according to their effort.

Questions. There may be two kinds of questions that a student can ask his teacher. The first can arise from interaction with other students belonging to the player's opponents, so that the student may ask his teacher about subjects that his friends know about and he does not. Students will receive reward for introducing subjects to their friends which become popular in the community, in order to encourage the input of diverse and detailed knowledge. The second kind of questions that students may ask arise from what the student has been taught, such as any objects that occur in the RDF triples in the student's knowledge base which do not exist as subjects. For example, after being taught about 'Barack Obama' as above, the student may have the question about what 'Harvard University' is. Other forms of questions may also be developed, such as using machine learning techniques to generate new hypotheses from the student's existing knowledge, using methods such as [3,4].

Tests. For any subject that a student has been taught, such as 'Barack Obama', any player can design a test to assess the student's knowledge on this subject. These tests can have a format as a set of hints about the subject, where the goal of the student is to correctly guess the subject when shown the hints. Hints are in the form of incomplete RDF triples that are (relation,object) pairs. When performing a test, the student matches the given hints with its knowledge base to narrow down the possible subjects, hopefully to one. From the best possible matches, the student chooses a random subject as the answer, and passes the test if it is the intended answer. The matching mechanism is based on string comparisons of words occurring in the hints and the knowledge base, but can become more sophisticated, e.g. with the incorporation of a thesaurus, Wordnet or an existing knowledge base. The matching ability can also be set to improve over time and according to the player's effort to simulate student growth.

The difficulty level of tests depends on the hints: a difficult test will contain fewer and less obvious or less well-known properties of the subjects. The student can take any number of retakes to pass a test, but the number of attempts would affect his performance valuation. At every failed attempt, the student can ask for help from its teacher by revealing some information about the test, such as a single randomly chosen word in the test that was not recognised. The teacher can then teach more knowledge before another retake. For example, after being taught about Barack Obama as above, a student can pass a test containing the hints

(president of, America)
(born in, 1961)

but may not pass a test containing the hints

(president of, America)
(married to, Michelle Obama)

unless the random guess succeeds. On failure, the student may respond by saying that it did not recognize the word "Michelle" in the test, in which case the teacher can teach the

appropriate knowledge before a retake. After a successful attempt, the student can also respond about how sure it was of the answer, e.g. by listing possible guesses it considered.

Anonymous valuation. The quality of tests will be determined by a method of anonymous valuation, to ensure the quality of knowledge and prevent 'cheating' amongst players. For a given test, the system can select a large sample of anonymous students that are not part of the player's friends community and have been taught the test subject. The average pass rate of these students on the test is taken as the measure of the difficulty level or quality of the test. For example, a test with a 0% pass rate is too difficult or not well-designed, and a 100% pass rate means it is too easy.

A student's overall performance measure will also be based on anonymous valuation. The system can choose a large sample of random tests not designed by the player's friends on subjects that the student has been taught, and the performance evaluation will take into account the difficulty level of the test and the number of attempts required by the student to pass. It may also take into account the popularity of the subjects that the student knows about, and how many new subjects the student introduced to fellow students in the community.

Advantages and comparison

We explain some of the important advantages that we expect to gain from the various features of Pedagogy described above, especially in comparison with previous approaches.

- **Knowledge Quality**

- **Accuracy.** Firstly, the accuracy of the acquired knowledge is based on the anonymous valuation aspect. This was also employed by Verbosity, but the advantage here is that it is only used for valuation purposes so that gameplay is not forced to be anonymous. Secondly, the goal of knowledge acquisition directly modelled by the teaching metaphor also alleviates the inaccuracies arising from the misuse problems encountered in Verbosity: in Pedagogy, the goal of the player as a teacher corresponds directly with the goal of knowledge acquisition rather than, for example, making another player guess a word. Hence inaccurate knowledge input will not be advantageous and likely be harmful due to anonymous valuation.
- **Coverage.** Expert and encyclopedic knowledge is supported in Pedagogy because players have the freedom to choose topics of their choice (as in Wikipedia), unlike Verbosity where one is restricted to common sense knowledge forced by anonymity of players. Secondly, the aspect of reward for new and different knowledge in Pedagogy encourages depth and diversity of inputted knowledge.
- **Sophistication.** This refers to the format of the acquired knowledge. Although we have discussed the simple RDF format, this could potentially become more sophisticated (such as n-ary relations, additional fields or type classification) as this may make the game more interesting, e.g. permitting different kinds of tests or questions.

- **Motivation**
 - **Competition.** Unlike previous approaches, there is a strong sense of competition in the game, driven by Farmville-like features such as wealth generation and spending, social status, long-term development of gameplay and gameplay between friends.
 - **Community/Collaboration.** As a social networking game, there is a strong sense of social community presence and collaborative aspects that add entertainment value.
 - **Personal presence.** Gameplay is very personalised, with personal choice of topic and control over the development of one's students and institution.
- **Educational aspect**
 - The design of the game encourages players to seek new and different knowledge about various topics of interest, and to articulate it into a structured and organised form. For example, one may imagine children searching for and extracting detailed knowledge from Wikipedia pages (or even rarer sources) to develop their students or to set difficult tests in competition with their friends. It also brings a new twist to the negative view of education amongst children, as they may enjoy playing the role of the teacher for a change, and learn proactively and positively in the process.

References

- [1] Open Mind Common Sense. www.openmind.media.mit.edu/
- [2] P. Singh, T. Lin, E. Mueller, G. Lim, T. Perkins and W. Zhu. Open Mind Common Sense: knowledge acquisition from the general public. Conference on Ontologies, Databases, and Applications of Semantics for Large Scale Information Systems. 2002.
- [3] Divisi. www.divisi.media.mit.edu/
- [4] R. Speer, C. Havasi, and H. Lieberman. AnalogySpace: reducing the dimensionality of common sense knowledge. AAAI 2008.
- [5] Games with a purpose. www.gwap.com/gwap/
- [6] R. Speer, C. Havasi, H. Surana. Using verbosity: common sense data from games with a purpose. Florida Artificial Intelligence Research Society Conference, 2010.
- [7] M. Yuen, L. Chen and I. King. A survey of human computation systems. Computational science and engineering, 2009.
- [8] L. von Ahn and L. Dabbish. Labelling images with a computer game. Computer Human Interaction, 2004
- [9] L. von Ahn, M. Kedia and M. Blum. Verbosity: a game for collecting common sense facts. Computer Human Interaction, 2006
- [10] Google image labeler. www.images.google.com/imagelabeler.

[11] Powerset. www.powerset.com.

[12] True Knowledge. www.trueknowledge.com.

[13] DBpedia. www.dbpedia.org.

[14] Yago. www.mpi-inf.mpg.de/yago-naga/yago

[15] G. Kasneci, M. Ramanath, F. M. Suchanek, and G. Weikum. The YAGO-NAGA approach to knowledge discovery. SIGMOD Record, 2008

[16] Freebase. www.freebase.com.

[17] S. Robertson, M. Vojnovic and I. Weber. Rethinking the ESP game. Computer Human Interaction, 2009.