Linked Data Games: Simulating Human Association with Linked Data

Jörn Hees, Thomas Roth-Berghofer, and **Andreas Dengel**
Knowledge Management Department,
German Research Center for Artificial Intelligence DFKI GmbH
and Knowledge-Based Systems Group, University of Kaiserslautern,
{firstname.lastname}@dfki.uni-kl.de

Abstract

Teaching machines to understand human communication is one of the central goals of artificial intelligence. Psychological research indicates that human associations are an essential requirement to understand human communication. In this paper the hypothesis is presented that simulating human associations with the help of Linked Data could improve text understanding capabilities of machines. To investigate whether human associations can be simulated with Linked Data, two preliminary problems are identified: (i) A reasonable ground truth for human associations is lacking and (ii) human associations have different strengths while Linked Data treats all triples equally and does not provide edge weights. To overcome these problems, two ideas for web games in accordance with Luis von Ahn's Games with a Purpose are proposed trying to turn the tedious acquisition processes into fun games. The resulting datasets are then to be used for quantitative comparisons of human associations and Linked Data.

1 Introduction

Since its introduction in 2001 the Semantic Web has gained much attention. In recent years, especially the Linked Open Data (LOD) project[1] contributed many large, interlinked and publicly accessible datasets, generating one of the world's largest, distributed knowledge bases. The accumulated amount of Linked Data can already be used to answer astonishingly complex questions (e.g., compiling a list of all musicians who were born in Berlin before 1900) or to provide additional information for selected concepts on a website (e.g., providing a short abstract and a thumbnail when hovering over the name Barack Obama).

In many ways Linked Data reminds of *spreading activation semantic networks* [Collins and Loftus, 1975]. Spreading activation semantic networks are successfully used in psychology to model human associations (e.g., thinking of *Barack Obama*, one will most likely also think of *USA*). Human associations are an important processing ability of our memory allowing us to retrieve related thoughts, traverse from one thought to another and thereby facilitate our way of thinking. They are also crucial for our understanding of everyday communication [Gerrig and Zimbardo, 2010, pp. 240ff] as they help us to build a context and resolve ambiguities. Hence it seems plausible that simulating such associations could help to improve text understanding capabilities of machines.

This work investigates the question if and how it is possible to simulate human associations with the help of Linked Data. A first analysis identifies two problems:

Currently, a quantitative comparison between Linked Data and human associations is impossible due to the lack of a reasonably large ground truth of human associations. Such a ground truth would consist of a large number of association pairs (e.g., (*Barack Obama*, *President of the US*)) collected from many different test persons. Collecting such a ground truth would allow us to answer questions such as: how large is the overlap between Linked Data and human associations? Nevertheless, due to the desired size, the acquisition would be infeasible with traditional approaches, such as paying test persons to record their associations.

The second problem is that human associations have different strengths, while Linked Data treats all triples equally and currently does not provide edge weights. Solving this problem would for example allow us to ask a machine to only show us the 20 strongest associations related to a resource, which in turn could be used to narrow down search spaces, use spreading activation algorithms in a meaningful way, or rank the results by association strengths. While several approaches try to rate triples by heuristics, none of them was compared to a dataset of human association strengths. Nevertheless, the acquisition of such a dataset would require us to assign weights to a very large number of Linked Data triples, which again would be infeasible with traditional approaches.

As both datasets need to be collected in order to investigate if it is possible to simulate human associations with Linked Data, two ideas for games in accordance with Luis von Ahn's Games with a Purpose [von Ahn and Dabbish, 2008] are proposed, turning the tedious process of entering associations or ratings into fun games.

The remainder of this paper is structured as follows: First the current state of the art is presented. Then a brief introduction into human communication is given explaining the important role played by associations in the language understanding process. In the following section human associations are conceptually compared to Linked Data, resulting in the two problems outlined above. For each of these problems a game idea is proposed, preceding the final conclusion and outlook.

2 State of the art

In recent years the research community devoted much attention to the Semantic Web [Berners-Lee *et al.*, 2001],

[1] `http://esw.w3.org/SweoIG/TaskForces/CommunityProjects/LinkingOpenData`

which proposed the vision of sharing information in the web not only to other humans, but also transforming it into a meaningful form for machines. Two of the most famous outcomes of these developments are the Resource Description Framework (RDF)[2] and the Web Ontology Language (OWL)[3]. Still, as manually providing information in these languages is quite cumbersome and seldom results in an immediate benefit for the authors, a few years went by without a reasonable amount of RDF being published on the web. Realizing this, a kind of grassroots development began to publish data which was already available in several centralized, corporate or governmental databases. Quickly guidelines were developed for publishing such data and interlinking it with others. With the ongoing help of volunteers and famous proponents, the so called Linked Open Data project was born, nowadays forming the so called Linked Data Cloud of interlinked datasets, including over 13.1 billion RDF triples as of November 2009. The content of this cloud, often called Linked Data, poses the world's largest distributed knowledge base and can be seen as a first challenge of the vision of the Semantic Web.

As more and more datasets were integrated into this Linked Data Cloud certain centralized points evolved, one of them being DBpedia[4]. The DBpedia team tries to automatically extract structured information from Wikipedia articles. In contrast to many other datasets, DBpedia represents very much knowledge across a large number of domains, which makes it very interesting for tying domains together. At the same time due to the automatic nature of the extraction, DBpedia also introduces a lot of errors into the Linked Data Cloud.

Even long before the Semantic Web started to evolve, psychological research has been very busy in the field of how human beings understand language. While the next sections go more into detail on spreading activation semantic networks [Collins and Loftus, 1975], it shall be mentioned here that they belong to the human semantic memory, which is a part of the so called explicit memory [Baddeley *et al.*, 2009, pp. 113–121].

As mentioned in the introduction, in this paper two games are proposed that try to turn the otherwise unfeasible work into fun games. This is motivated by Luis von Ahn's Games With A Purpose[5], which are part of a field called Human Computation. After the big success of the famous ESP Game [von Ahn and Dabbish, 2004], which turns the tedious process of labeling images into a fun game, and further games, a summary of design principles for Games With A Purpose [von Ahn and Dabbish, 2008] was published.

The proposed game ideas are especially related to the ESP Game, Verbosity [von Ahn *et al.*, 2006], Matchin [Hacker and von Ahn, 2009] and OntoGame [Siorpaes and Hepp, 2007]. Verbosity is a game which turns the widely known Tabu game into a game collecting common-sense facts. The game is of an asymmetric type, which means that there's a describer who gets a word that she has to describe to the guesser. The guesser can see the describer's output and guesses her input. In order to prevent cheating and to direct the collection of common-sense facts, the describer can not enter text freely, but has prepared snippets, which can be completed, such as "is a", "has a", "is the opposite of". If human associations were extracted from the facts collected with this game, they would be strongly biased towards these snippets. Besides this, the results of the game do not seem to be published. Matchin is a game, which presents pairs of images to the players and asks them which one their partner will prefer. The collected amount of relative votings is then used to globally rank the pictures. OntoGame can be seen as kind of the first game of this kind applied to Linked data. Nevertheless, it solves a very different task than the later on proposed games, namely finding out whether a Linked Data resource is an instance or a class and then trying to file them into a taxonomic structure.

[2]`http://www.w3.org/TR/rdf-primer/`

[3]`http://www.w3.org/TR/owl-ref/`

[4]`http://dbpedia.org`

[5]`http://www.gwap.com`

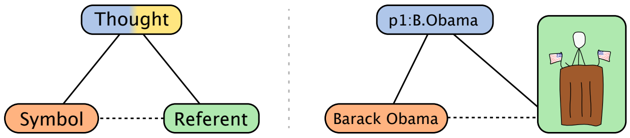

Figure 1: Semiotic Triangle after Ogden & Richards (left) and an example (right)

3 Human Communication

Human communication is the process of transporting information from one human being to another.[6] In such communication we can distinguish between *symbols*, THOUGHTS and `referents` as they can be visualized in the so called semiotic triangle (Figure 1).

This distinction allows us to see one of the main characteristics of human communication: Thoughts heavily depend on the respective person, and we are not able to exchange thoughts directly. A THOUGHT (e.g., P1:B.OBAMA) is someone's mental representation of some `referent` (e.g., `Barack Obama`, the one person with that name, currently being president). Instead of exchanging a thought directly, we are only able to exchange a *symbol* for the thought in written or spoken form (e.g., the two words *Barack Obama*). As a speaker or writer we then hope that the listener or reader of such a symbol finds an own thought that is sufficiently similar to our thought (see Figure 2, P2:B.O.), or creates a new thought for what we described.

Communication is not limited to the exchange of single thoughts, but is all about exchanging information (i.e., the connections between thoughts). As we can not exchange thoughts directly, we can only exchange information by a from of symbolic indirection. From a Semantic Web point of view, information can be expressed as simple statements. Each statement essentially is a simple sentence in the form of (subject, predicate, object)-triples, but instead of the usual symbolic form of a sentence now subject, predicate and object are thoughts. The exchange of the triple (P1:B.OBAMA, P1:BIRTHPLACE, P1:HONOLULU) from person $p1$ to person $p2$ via the described symbolic indirection is visualized in Figure 2. Note that the thoughts of $p1$ and $p2$ are disjoint as they are in different brains, but hopefully they are similar enough, so that $p2$ understands $p1$. Also note that P2:B.O. is more detailed than P1:B.OBAMA, as $p2$ also knows a more specific name.

From this simple example one can identify two problems that can occur in our communication:

- A thought can be referred to by several symbols (e.g., *Barack Hussein Obama II*, *Obama*, *President of the*

[6]We focus on written or spoken communication here.

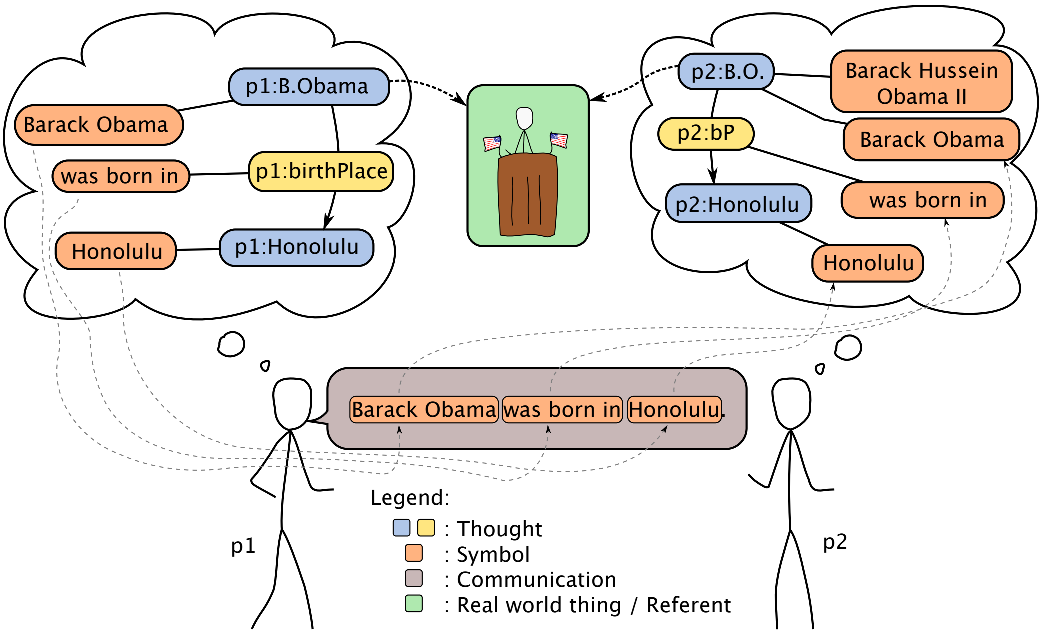

Figure 2: Human to human communication

United States, *President Obama* all could refer to one single thought p2:B.O.).

- A symbol can refer to several thoughts, also known as lexical ambiguity (*Apple* might refer to Apple (the fruit) or AppleInc (the company)).

3.1 Human associations

Now, when we read or hear some symbol (e.g., *Barack Obama*) we instantly connect it with its thought(s) (e.g., p1:B.Obama), but also, we somehow remember related thoughts (e.g., p1:PresidentUSA, p1:USA, ..., p1:Honolulu), called associations.

Associations can be thought of as a graph of thoughts or a so called spreading activation semantic network [Collins and Loftus, 1975], in which each thought is a node and associations between these nodes are edges. An example for such a semantic network is depicted in Figure 3 and shows some possible associations of p1:B.Obama. The length of an edge represents the strength of the association (e.g., p1:B.Obama's association to p1:PresidentUSA is stronger than to p1:Honolulu).

3.2 Context

It is argued that such associations are key requirements for our daily communication [Gerrig and Zimbardo, 2010, pp. 240ff]. Whenever we read or hear a symbol, our brain looks up its connected thought(s). These thoughts and their associations (in the following called associations of the corresponding symbol) are used to generate and adjust a *context* of thoughts that are related to the current communication sequence. The associations of a symbol may either support the current context or oppose it. In the first case our associations reinforce the context and our confidence that we understood what the speaker was telling us and that we are thinking about similar thoughts rises. In the latter case we have an indication that our lookup of the symbol's thought was wrong and we struggle to find a different meaning for what we have heard so far (e.g., in the sentence "Last year the pen was abandoned as it was too dirty for the animals to live in." [Gerrig and Zimbardo, 2010, p. 241] one would first assume that "pen" was a writing instrument, as it does most of the time (so called biased ambiguity), but when reading about animals living in it we suddenly realize that it refers to an enclosure).

Such a context, which emerged from our ongoing integration of associations, can be seen as our condensed set of thoughts related to the current communication. We can also think of this context as a working space, a drastically reduced search space for things that we expect to be related to what the speaker / writer expressed, instead of considering every thought we have in our mind to be potentially related. The context hence allows us to do much faster lookups (e.g., for symbols of thoughts which until now only were mentioned implicitly or for the integration of new information into our memory) and gives us the ability to resolve ambiguities.

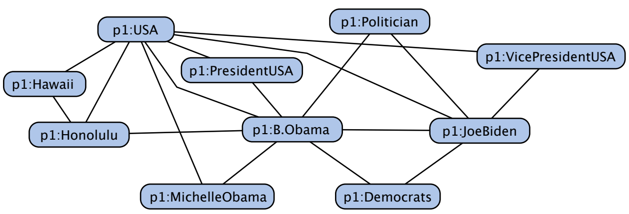

Figure 3: Semantic Network

4 Comparing Linked Data with human associations

Linked Data and spreading activation semantic networks used to model human associations are very similar. Both can be seen as network structures with thoughts as nodes and edges in between, both of huge dimensions.

In contrast to human communication, "thoughts" in Linked Data are meant to be exchanged directly between machines, which means without the symbolic indirection that we humans have to use. In contrast to the human brain the evolving knowledge base is thus distributed and interlinked, which is at the same time one of its largest advantages and disadvantages. For example it allows continuing growth and allows machines to exchange information without reasoning about whether their symbols mean the same things, but at the same time this freedom introduced many errors into Linked Data, as anyone can issue statements about anything or as people use properties in different ways (cf. the use of `owl:sameAs` [Halpin and Hayes, 2010]).

Another difference between Linked Data and human associations is that Linked Data has strongly typed edges, which identify the kind of relation between two concepts very precisely, while human associations do not. Usually it takes us much longer to find a name for the association of two thoughts than it takes us to name the associated thought. Furthermore, human associations do not seem to be limited to a certain type of relation between thoughts, even though it would be interesting to investigate which of these properties we use most frequently.

To be able to investigate questions as the one above, and more general to perform any quantitative comparison between Linked Data and human associations, a large ground truth of human associations is needed, which currently does not exist.

Except from the need for such a ground truth, until now only differences were mentioned, which render Linked Data more specific than human associations and hence would allow to easily simulate human associations with Linked Data by just ignoring details such as edge labels. Still there exists one key feature of human associations, which currently is not part of Linked Data: association strengths. Even though heavily dependent on the person, the current context and task, nearly all humans will agree that they generally associate Barack Obama stronger with USA than with Honolulu. In contrast to this, the facts in the Linked Data cloud are facts in a logical sense. They are assertions, all of the same "truth", none being more valuable than another.

A collection of such association strengths would allow us to ask machines not only to give us all information about an instance (e.g., all 600 triples for Barack Obama), but also to rank this information by association strength (e.g., only presenting the top 10 of these to an end user) and thus, to constrain the number of results. With regard to understanding human texts, this would allow us to propagate activation from thoughts whose symbols directly occur in the text to thoughts occurring only implicitly in a human like way. This in turn, would enable us to narrow down our search space from the whole Linked Data graph to only those thoughts associated with the current communication sequence by an "average human". Another immediate benefit from annotating Linked Data triples with an association strength is a kind of feedback for automated extraction processes such as the one underlying DBpedia. One could investigate, which extraction rules yield high and which ones yield low association strengths, possibly driving an improvement process.

Besides these immediate uses, such a collection of association strengths would also allow us to test whether current common heuristics truly model how we associate thoughts and if they do, they could be used to bootstrap the acquisition of associations strengths. Examples for such heuristics include word co-occurrences on websites or graph intrinsic features such as page rank, betweenness and the like, trying to model how much activation flows from one thought to another.

For further research it is thus proposed to generate both: a human association ground truth as described earlier and a collection of Linked Data triples rated with their association strengths.

5 Linked Data Games

In order for both datasets to be of any use, their size needs to cover a sufficiently large amount of human associations / Linked Data. The larger the amount of data collected, the more meaningful later quantitative comparisons will be. At the same time all data collected (e.g., associations of Barack Obama or whether USA is stronger associated to him than Honolulu), is highly subjective, resulting in a large amount of input from different persons to be needed in order to arrive at a collectively agreed association / rating. Additionally entering a large amount of associations or ratings is an extremely tedious task.

The former constraints render the application of traditional approaches to generate a ground truth, such as paying test persons to record or write down their associations or ratings, infeasible. Nevertheless, there exists a novel approach called Human Computation [von Ahn and Dabbish, 2008] that might be suitable to solve this acquisition problem. The approach suggests to turn problems which are difficult to solve for machines into fun games with atomic decisions which are easily answerable for us humans.

In the following for each of the needed datasets an idea is presented on how the individual problem could be turned into such a game.

5.1 Building a human association ground truth

In order to acquire a human association ground truth one would usually present an item to a participant and ask her to enter all associations she has with this item in free text in a fixed amount of time (e.g., one minute). Then the next item would be shown. In order to make the collected data robust against priming effects, the order of the items usually would be randomized.

In order to turn this tedious process into a game many aspects of the ESP Game [von Ahn and Dabbish, 2004] can be borrowed. In contrast to the ESP Game the items of this game are not web images but Linked Data resources (i.e., their symbols).

In its simplest form the envisioned game called *Associator* is a symmetric two player output agreement web game, where a player starts a game and is randomly grouped with some other player for a short, fixed period of time (e.g., 2 minutes). Both players then play in rounds and can not communicate by other means than described in the following. At the start of each round they see an input item $\mathcal{S}$ (e.g., *Barack Obama*). What makes this game a Linked Data game is that this item is a symbol of a Linked Data resource $\mathcal{R}$ (i.e., there is a triple ($\mathcal{R}$, RDFS:LABEL, $\mathcal{S}$) in the LOD cloud (e.g. (DBPEDIA:BARACK_OBAMA, RDFS:LABEL, *Barack Obama*)). Both players are then asked to enter what their partner will associate with this item. The round will last until any of the outputs of player p_1 matches any of the outputs of player p_2 or until both players decide to pass this item, as they can not agree on any association. All outputs and timings in this process are recorded. In case of a match both players get points depending on the time it took them to find the match and get into the next round. In case of a pass both players get into the next round without getting points.

What this game is going to return behind the scenes is a collection of associations between the input items (i.e., Linked Data resources' symbols) and user entered symbols, which actually represent associated thoughts as we know. The matches recorded in this process are very good candidates for collectively agreed associations of the presented item. The other non matching guesses as well pose associations, but are very susceptible to subjective associations, cheating and the like. Still any recorded round can be used for single player games, for example when there is an unequal amount of players who want to play or when a player quits out of a running game. In this case the current player is not matched with another player who is playing at the same time, but one whose session was recorded earlier. This single player process can then validate additional guesses of the recorded session.

In contrast to the ESP Game a couple of open questions remain trying to use the collected data to build a human association ground truth.

First of all this simplest version of the envisioned game does not include taboo words. Taboo words are a list of words shown to the players in addition to the round's item. Any word in the taboo list can not be used as a match to get into the next round. In the ESP Game taboo words are used to force the players to enter a larger variety of labels, by adding outputs that are agreed on for a certain amount of times to the corresponding input's taboo list. The problem with this approach is that taboo words certainly bias the association process in a way that not only the round's input item primes associations, but also the taboo words (they sort of work as additional inputs). One of the ways to weaken this problem is randomly selecting taboo words out of an actually larger list of taboo words, as it is proposed for the ESP Game. Still it might be interesting to investigate if it is not possible to completely eliminate this kind of biasing, e.g., it is proposed here to try a kind of covered taboo list. In this case the players would only see a list of covered words, indicating that there are taboo words. In case the player enters a word from this list, it blocks a potential match, gets revealed (for this player only) and the player is awarded with some small amount of additional points.

Another open problem is related to the fact that all collected associations actually are symbols for associated thoughts. In order to analyze whether these associations are part of the LOD cloud already, the symbols need to be matched back to Linked Data resources, where possible. This process often is very ambiguous and it is not clear if this symbol to resource resolution can be seamlessly integrated into the Associator game without introducing new distortions. Also it should be investigated how to select the input items, which in the presented form are symbols of randomly selected Linked Data resources.

5.2 Rating Linked Data triples with association strengths

While the previous game idea targets creating a human association ground truth, in this section an idea shall be presented how to turn the process of assigning association strengths to Linked Data triples into a game. A simple traditional approach for this task could show all Linked Data triples related to a previously selected Linked Data resource (e.g., all triples with `dbpedia:Barack_Obama` as subject and/or object) to a test person and ask her to order these triples according to which one of the relations she would have thought of first, when hearing or reading *Barack Obama*. After ordering the list, the next list for another Linked Data resource is shown.

Before presenting the idea to turn this process into a game, two things shall be mentioned: First, as showing a list of URI triples to the end-user is not of much use, the users will always see a symbolic representation of this list. Luckily Linked Data resources are usually labeled with a symbol with the `rdfs:label` datatype property. Second, the outcome of each of these experiments, which is a user centric absolute ranking, is not only highly subjective, but sometimes even unstable for one person, as a lot of relative decisions are involved within this process and the human brain tends to lose track of them. Hence it seems to be better to ask for these atomic relative comparisons of two facts and then use an objective ranking algorithm to generate an absolute ranking out of them. Generating an absolute ranking out of such results can be compared to chess ranking systems, where based on the outcomes of atomic competitions (player p_1 won against p_2), a global ranking is calculated, just that in this case there are no players competing, but facts.

The envisioned game is called *BetterRelations* and borrows many ideas from Matchin [Hacker and von Ahn, 2009], which asks users to compare images against each other and calculates a global ranking from this.

In its simplest form BetterRelations is a symmetric two player decision agreement web game. A player starting the game is randomly matched with some other player either for a predefined amount of time (e.g., two minutes) or a predefined amount of rounds (e.g., 50). At the start of each round both players are presented with one input item, which actually is a Linked Data resource's symbol (e.g., *Barack Obama*), and two facts about this item (e.g., *is president of the USA* and *was born in Honolulu*). Both players are then asked to select the fact that their partner will have thought of first. In case both players agree they are rewarded with points and get into the next round. In case of disagreement, both players get into the next round but do not get points. As in Matchin the amount of points collected in a round rises with the number of decision agreements of both players in a row, punishing a disagreement without actually subtracting points from the user and preventing cheat strategies to gain many points by simply selecting a random fact as fast as possible. In order to prevent another obvious cheating strategy, the facts are presented to both players in randomized order. All actions in this process are recorded, and for example used for a single player mode.

Behind the scenes this game acquires a large amount of relative decisions between facts. Decisions which both players agreed on are especially validated. Disagreements could mean a lot of things and could result from subjective rankings and cheat attempts, or could indicate that the two facts were equally important.

In contrast to Matchin, BetterRelations will not create one globally ranked list for all Linked Data triples, but instead is going to create a list for each Linked Data resource of interest and all facts related to this resource. The ranking algorithm, which transforms the relative ratings into these global ratings hence has to deal with a lot smaller lists than in the case of Matchin.

Also the algorithm shall be able to quickly exclude a large number of erroneous facts, as they occur in Linked Data, from being played again, in order not to bore players with such facts. One possibility would be to provide the players with an explicit button saying “both facts are nonsense”. Nevertheless, it has to be investigated how to include this third choice into the rewarding system without abandoning the effective cheating prevention mechanisms.

Last but not least, BetterRelations is a game of a very exploratory nature. Often it could happen that the player does not recognize the round’s item and might want to read a few sentences, explaining what the resource is about. For many Linked Data resources such a short introduction exists in the form of a `rdfs:comment` or a `dbpedia-owl:abstract`, but it is unclear how this information can be included in the game without interfering with the rating choices. One promising possibility could be to include such information on demand only and to treat the resulting rating in a different way on the server side. Also as it generally could take a few seconds for the user to switch context, it should be investigated if the throughput of the game can be raised by playing a few consecutive rounds about the same item, only changing the facts, without introducing priming effects.

6 Conclusion and Outlook

This paper introduced and motivated the hypothesis that simulating human associations could improve text understanding capabilities of machines. Thanks to Linked Data we have a very large and promising dataset at hand to simulate human associations. Nevertheless, in order to investigate more thoroughly whether Linked Data really is a good dataset to simulate associations, first of all a reasonably large human association ground truth is needed. Also as human associations can be of different strengths, such associations strengths would need to be annotated to many Linked Data triples.

As generating both of these datasets would be infeasible with traditional approaches, for each of them an idea was presented how to turn the acquisition process into a “Game with a Purpose”.

After presenting these ideas the next research steps include implementing several versions of the games and properly evaluating them with respect to the desired data quality, the throughput, average lifetime play and expected contribution as mentioned in [von Ahn and Dabbish, 2008].

The resulting datasets are then to be used to investigate the overlap of human associations and Linked Data, to rate the extraction process of DBpedia and to benchmark several heuristics used to infer non-existent edge weights for Linked Data.

Acknowledgements Special thanks to Benjamin Adrian for the many insightful conversations about this field of research.

References

[Baddeley *et al.*, 2009] Alan Baddeley, Michael W. Eysenck, and Michael C. Anderson. *Memory*. Psychology Press, 2009.

[Berners-Lee *et al.*, 2001] Tim Berners-Lee, James Hendler, and Ora Lassila. The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American*, 284(5):34–43, May 2001.

[Collins and Loftus, 1975] Allan M. Collins and Elizabeth F. Loftus. A Spreading-Activation Theory of Semantic Processing. *Psychological Review*, 82(6):407–428, 1975.

[Gerrig and Zimbardo, 2010] Richard J. Gerrig and Philip G. Zimbardo. *Psychology and Life*. Allyn & Bacon, Pearson, Boston, USA, 19th edition, 2010.

[Hacker and von Ahn, 2009] Severin Hacker and Luis von Ahn. Matchin: Eliciting User Preferences with an Online Game. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1207–1216, Boston, MA, USA, 2009. ACM.

[Halpin and Hayes, 2010] Harry Halpin and Patrick J. Hayes. When owl: sameAs isn’t the Same: An Analysis of Identity Links on the Semantic Web. In Christian Bizer, Tom Heath, Tim Berners-Lee, and Michael Hausenblas, editors, *Proc. of the WWW2010 Workshop on Linked Data on the Web (LDOW)*, Raleigh, USA, 2010. CEUR Workshop Proceedings.

[Siorpaes and Hepp, 2007] Katharina Siorpaes and Martin Hepp. OntoGame: Towards Overcoming the Incentive Bottleneck in Ontology Building. In *Proc. of the 3rd International IFIP Workshop On Semantic Web & Web Semantics (SWWS)*, Vilamoura, Portugal, 2007. Springer, LNCS.

[von Ahn and Dabbish, 2004] Luis von Ahn and Laura Dabbish. Labeling Images with a Computer Game. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 319–326, Vienna, Austria, 2004. ACM New York, NY, USA.

[von Ahn and Dabbish, 2008] Luis von Ahn and Laura Dabbish. Designing games with a purpose. *Communications of the ACM*, 51(8):58–67, August 2008.

[von Ahn *et al.*, 2006] Luis von Ahn, Mihir Kedia, and Manuel Blum. Verbosity: A Game for Collecting Common-Sense Facts. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 75–78, Montréal, Québec, Canada, 2006. ACM, New York, NY, USA.