



## Embedded Human Computation for Knowledge Extraction and Evaluation (uComp)

### D2.2: Progress Monitoring and Quality Control Model

16 December 2015

Version: 1.1

#### Version History

Version	Date	Author	Comments
0.1	01/08/2015	M. Föls	Initial Draft Document
0.2	15/08/2015	A. Scharl	Feedback and Minor Edits
0.3	22/08/2015	F. Fischer	Revision of Language Quiz Section
0.4	04/09/2015	M. Föls	Major Revision
1.0	30/09/2015	A. Scharl	Final Revision and Document Layout
1.1	16/12/2015	M. Föls	Update of Screenshots and Usage Statistics

Dissemination Level: PU – Public

uComp receives the funding support of EPSRC EP/K017896/1, FWF 1097-N23, and ANR-12-CHRI-0003-03, in the framework of the CHIST-ERA ERA-NET program line.

## Table of Contents

[Table of Contents](#)

[Overview](#)

[Applications](#)

[Climate Challenge](#)

[Language Quiz](#)

[Progress Monitoring](#)

[Quality Control](#)

[Incentive Mechanisms](#)

[Summary](#)

[References](#)

[Extended Bibliography](#)

## Overview

This deliverable presents quality control methods for the human computation processes of the uComp research project, including task progress and reliability measures, cheating prevention strategies, and triggers for incentive mechanisms.

Cheating prevention is being addressed through strategies customised to the specifics of knowledge extraction tasks, including (i) economic models to ensure cheaters do no better than break-even; (ii) defensive task design to encourage users to put in genuine efforts to carry out the tasks; and (iii) statistical models for quality and agreement enforcement that identify outliers and compare contributor answers against gold standard data.

Task and progress monitoring during the human computation cycle can also amplify the incentive mechanisms. Monitoring processes can be exposed as REST services through an Application Programming Interface (API), which allows researchers to define alerts or configure automatic behaviour rules.

The HC Framework plays a core part in the uComp project. During the first year, we have designed the framework (e.g., define core task types, agree on an API with partners) and started its implementation by extending MOD's existing application framework for social media platforms. The following system features have been implemented: multi-channel deployment and social logins to ensure that games can be played from mobile devices; viral notification system, allowing the possibility to invite friends (other options to follow); a bridge to CrowdFlower, allowing the publication of game elements as *Human Intelligence Tasks* (HITS) on mechanised labour marketplaces.

In the second year, the API was tested and reviewed within a small team, to ensure the stability of the code, the clarity of the documentation and also the range of the functionality of the API. During the process a small subset of partners used the API to test their use cases. The feedback of the testers was used to address several minor issues and extend the documentation of the API in cases where the initial version was unclear. Since the API of the system is a very central component that cannot be changed easily once it is being used by many people, this process was carefully planned and executed.

At first the game was intended to only use the login capabilities of Facebook, but in order to maximize the number of possible participants (and create a generic login framework for other projects as well), a more comprehensive Single-Sign-On component (using Shibboleth and SimpleSAMLphp) was created. The core SimpleSAMLphp functionality was extended with custom plugins to authenticate via Facebook, Twitter and Google. This should allow all interested players to participate in the game. The authentication framework set the stage for launching two concrete game applications, the the *Language Quiz* and the *Climate Challenge*.

The crowdsourcing engine and the individual game components were continually refined based on the feedback from beta testers and early adopters. Special emphasis was placed on the interaction design - e.g., the scoring and engagement mechanisms, as well as the flexible support of different task types - e.g., questions where the correct answer will never be known, or questions where the correct answer will be determined at a future point in time (in the former case, the points are based on the mean answer of all players; in the latter case, the points are awarded *ex post* once the correct answer is known).

## Applications

### Climate Challenge

The *Climate Challenge* is an online competition in the tradition of games with a purpose that combines practical steps to reduce carbon footprint with predictive tasks to estimate future climate-related conditions. The application is designed to increase environmental literacy and motivate users to adopt more sustainable lifestyles. Its feedback channels include a leaderboard and a visual tool to compare answers of individual players with (i) the average assessments of their direct social network contacts as well as the entire pool of participants, (ii) a selected group of experts, and (iii) real-world observations (Scharl et al., 2015).

The Climate Challenge supports the following task types:

- Ten *multiple choice questions* per month, which are related to the topic of climate science. Since the answers to these questions are known, the players will get immediate feedback on the correctness of their answer. The players will get points for the right answer, and a point penalty for wrong answers.
- *Opinion poll* obtain survey-like results from players - for example, asking them whether they think governments should regulate the release of greenhouse gases, and to what extent depending on assumed cost factors. *Opinion polls* are bonus challenges, there is no fixed amount of monthly tasks. They offer a more open alternative to multiple choice questions - more flexible to use (e.g., in cases where there is no correct answer), and should appeal to a wider audience when tied to current events.
- Fifty *sentiment detection* tasks per month, where the players state whether they perceive certain words or phrases as positive, neutral or negative. Correct answers for this task type are not immediately available, as the assumed correct value will be calculated as an average once a certain task has been evaluated by ten different players. Therefore, points are distributed *ex post* based on how closely a player's assessment matches the average rating by other players.

- A monthly *prediction task*, which is promoted in a cooperation with *Climate Program Office* of the *National Oceanic and Atmospheric Administration* (NOAA), where the players are asked to predict future environmental conditions - e.g., *What percentage of land area in the Northern Hemisphere will have a "White Christmas"?* Three experts will also provide their answer to the monthly question. After the real-world answer is known (typically through actual measurements), the players are being awarded points depending on how closely their answer matches the observed value. A diagram lets them compare their own assessment with the average predictions of climate scientists, the predictions of all players, and the observed value.
- Five *pledges* per month ask participants for feedback on recommendations for reducing energy consumption and making sustainable lifestyle choice. They can declare if they have already adopted the recommendation, rate the environmental impact of this action, and share accepted pledges via their social media channels.

Each task type has a different mechanism to distribute points. For the prediction and sentiment detection tasks, the players will receive the points once the correct answers are available (i.e., at a future point in time). For multiple choice question, players will receive points immediately depending on whether the answer was correct or not. For pledges and opinion polls, players will receive a certain amount of points just for participating, because the task type does not have a right or wrong answer.

Since the main goal of the Climate Challenge is not only to collect answers, but also to raise awareness, the promotion and community aspect of the game has to be treated different than in the Language Quiz (Scharl et al., 2015), which will be presented in the following section. As Sabou et al. (2014) noted there are three different categories of crowdsourcing paradigms, and it is hard to mix the monetary reward system of Crowdfunder with a gameplay-driven and altruistic reward system of the Climate Challenge. Instead, social media channels such as Facebook and Twitter were used to promote the game and create a stable community. This approach is also supported Huberman (2008), observing that ideas shared through friends' activities spread across social networks and can benefit the reach of a crowdsourcing application by creating widespread attention.

Another special aspect of the Climate Challenge is the recurring nature of the tasks. Inevitably, there is a period where a prediction task is "open" for answers, followed by a waiting period until the real-world answer becomes available. Then the next prediction question can be asked. To account for this delay, and as an incentive measure, monthly game rounds were introduced where only a certain number of tasks per task type are being made available. This approach ensures that players will always find new tasks at the beginning of each month, and have an incentive to return to the game to find out about the correct answer to the previous question.

The created community was used as a platform to promote the new round each month, supplemented by e-mail notifications and Facebook announcements (to the approximately 2,500 followers of the Facebook community page) that new results were available. This approach was used as a regular reminder to re-activate players back every month, and to keep them engaged.

Since its launch in March 2015, the Climate Challenge attracted 3,236 unique users as of December 2015. From those, 644 created a user-account in the game leading to a high conversion rate of 20%. Out of those 644 users, 554 became active players. In total the game collected 680 prediction answers, 4,316 multiple choice answers, 15,979 sentiment answers, 2,003 pledge answers and 34 answers to the recently introduced opinion poll.

## Language Quiz

The approach of the *Language Quiz* differs significantly from the *Climate Challenge*, even though both games are built upon the same framework. Instead of creating environmental awareness and promoting sustainable lifestyle choices, the Language Quiz aims to acquire multi-lingual language resources for research purposes.

A flexible and open architecture to support multiple languages, different types of task, and the inclusion of third-party tasks were among the key factors that have guided the development efforts. Conceptual insights gathered from analyzing the results of the *Sentiment Quiz*, a game with a purpose to assess sentiment terms (Scharl et al., 2012), also influenced the design of the *Language Quiz* - although the two applications are based on a completely different technology stack.

The screenshot displays the main interface of the uComp Language Quiz. At the top, there is a navigation bar with 'Play', 'Progress', 'Leaderboard', and 'About' options. The main content area features a 'Sentence Sentiment' task with a question in German: 'Drückt der fett hervorgehobene Text eine negative, neutrale oder positive Meinung aus?'. Below the question is a text snippet with a highlighted sentence: 'Zu wenig, sagen die Umweltschutzorganisationen. Gerade noch erträglich, sagen die Bauern. **Denn der Kampf gegen Schädlinge hat sich in Österreich in den vergangenen Jahren verschärft.** Schädlingsplage aus Amerika Ein fünf Millimeter kleiner schwarz-gelber Käfer sorgt derzeit für Unruhe auf den Äckern: der Westliche Maiswurzelbohrer.' Below the text are four buttons for sentiment selection: red minus, red plus, green plus, and green minus. A section for 'AVAILABLE LANGUAGES' shows buttons for EN (English), DE (German), ES (Spanish), RU (Russian), ZH (Chinese), and CS (Czech). The 'Game Level' section shows a progress bar with a level indicator set to 1, and a 'You' marker at the bottom. On the right side, there is a 'Top Scores December 2015' table and a 'Bonus Points December' section with a value of 0. The bonus section includes a megaphone icon and the text: 'Invite your friends to play! You'll receive a bonus of 5% of their points after they have accepted your invitation.'

Rank	Player Name	Level	Score
1	Crowd Flower	Level: 3	345
2	Reinhard Fischer	Level: 2	265
3	Michael Föls	Level: 1	125
4	Arno Scharl	Level: 1	105
5	Stefan Gindl	Level: 1	25

**Figure 1.** Main interface elements of the uComp Language Quiz<sup>1</sup>

<sup>1</sup> quiz.ucomp.eu

To open the game engine for use by the partners of the project, an API was created which allows approved partners to send tasks to the game. Currently, Language Quiz supports multiple choice and sentiment assessment tasks, on both a term level and a sentence level (the latter serving either stand-alone statements, or highlighted sentences in a context of a whole paragraph). When uploading the task data, the project owner can decide if the tasks should be sent (i) to the game, (ii) to CrowdFlower, or (iii) to a hybrid workflow that uses CrowdFlower to recruit game participants. The flexibility of this approach increases the usefulness of the framework as an evaluation tool, especially when the size of a community and its engagement level are difficult to assess *ex ante*, or if the task is rather repetitive in nature. This is often the case when acquiring language resources, as compared to games with a purpose such as the Climate Challenge, which offers higher intrinsic motivation due to a limited amount of handpicked monthly tasks in a domain that is often characterized by altruistic motivation. Since the *Language Quiz* does not follow such a curated approach, the quality of the questions can vary.

The mechanics of the game were also adapted to support the hybrid scenario. If a user is sent to the game via Crowdflower, he would receive 20 questions. Upon successful completion of the game rounds, the user will receive a code which he can enter in his Crowdflower task to be paid. Since only the highest-rated Crowdflower users were being targeted, the quality of the results were correspondingly good. To filter out cheating players, each Crowdflower user has to answer three gold standard questions. In the case of a wrong answer on a test question, the user does not receive Crowdflower payments, and his answers will be deleted from the game. This ensures a consistent high quality of the obtained results.

The hybrid support of native gaming and paid Crowdflower evaluation means that the task uploaders have the chance to leverage the benefits of both crowdsourcing categories depending on the requirements of the given task. The paid tasks also improve the user experience of the native game participants, since the correct answers and related game point calculations are available sooner. As Sabou et al. (2013) have identified, games with a purpose and mechanized labour have different pros and cons in the categories speed, cost and quality of answers - using the uComp framework, task providers can make the decision for each task individually.

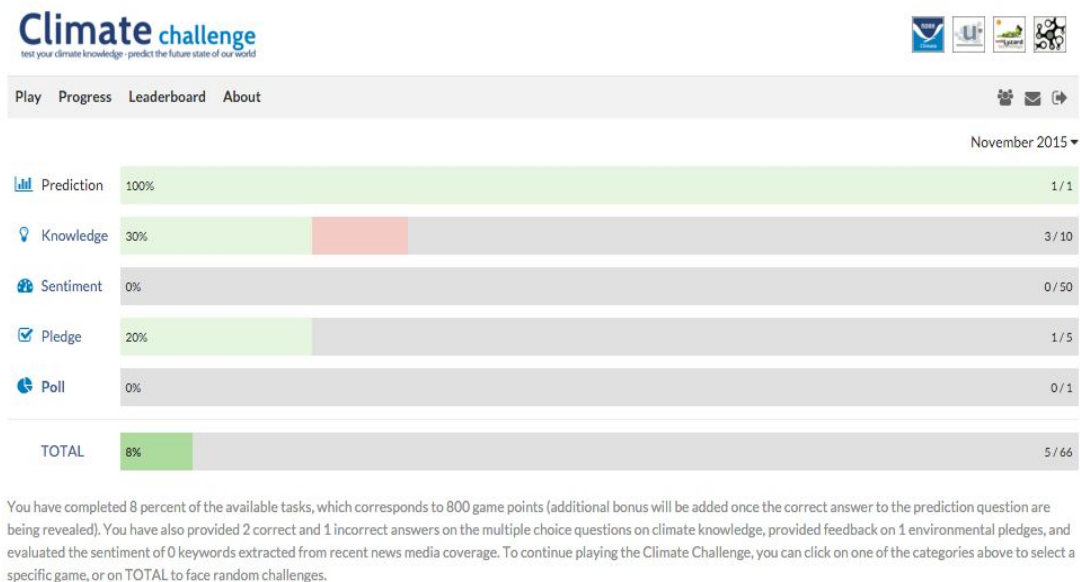
Since its launch in October 2015, the Language Quiz attracted 814 unique users. 579 of these users created a user account, and 525 became active players. Out of the 525 active players, 503 were contacted via the Crowdflower campaign option. The game was also promoted at RANLP'2015 and through social media. 347 users became active players who submitted valid answers, 178 failed the test questions. In total 12,283 valid answers were submitted to the game, 7,459 of those answers where paid via Crowdflower and 4,824 where organic answers from the players.

Further promotion activities by USFD are planned for early 2016, pending ethical approval from the ethics committee. The approval is required since we will be recruiting human participants via the University's volunteer list and also via the University run GATE users community list.

## Progress Monitoring

The progress monitoring serves as an incentive mechanism to engage the players. The uComp framework uses a dynamic level bar below the question box, to show the user at which level he currently is and where other players are on the same level (displayed as a horizontal bar that indicates how far other players are ahead or behind the user). This tool was designed to motivate the user to play more because he can always see how close other players are positioned around him and by playing the game the user can see his avatar bypass the other users to reward him visually for actively playing the game.

Another real-time indicator is the monthly high score table in the right sidebar, which shows the overall top three scores as well as the two participants who rank immediately before and after the player. This tool was designed to visualize the overall position of the player and to show how much distance is between the player and the top players. Like the level indicator, its dynamic updates are intended to encourage the players to keep playing, as they will always be visually rewarded by watching their avatar climb up the leaderboard.

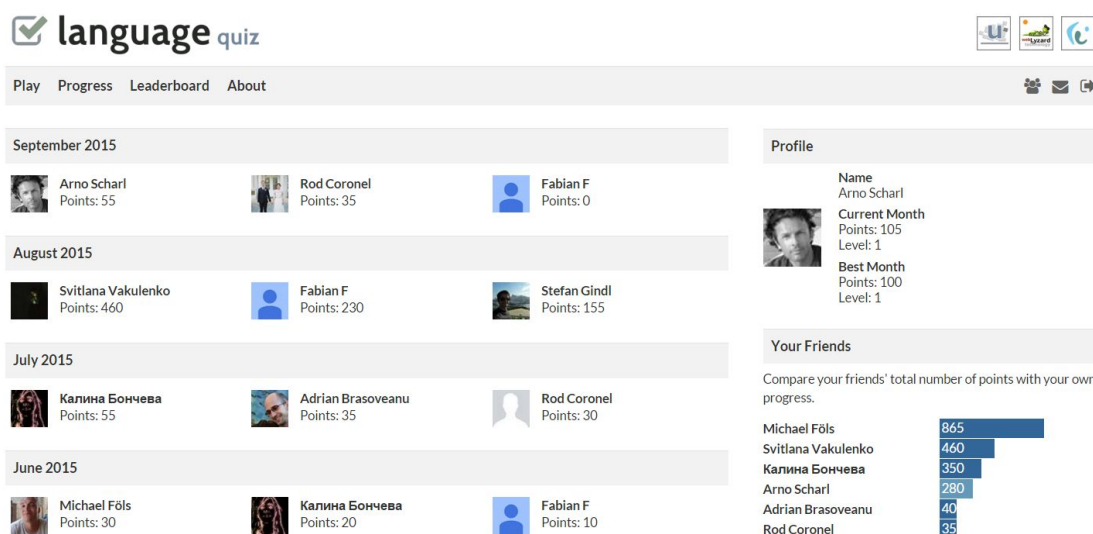


**Figure 2.** Progress page of the Climate Challenge



Another incentive to motivate players are the dynamic progress bars shown in Figure 2, which indicate how many of the monthly tasks in each category have already been solved by the player - e.g., if a player has answered two of the five pledges, the corresponding bar on the progress monitor will be at 40 percent. This feature aims to trigger the user's desire for completion while they have not answered all the monthly questions. At the same time, the component serves as a navigational aid. Users can click on the desired task type to only show questions of this type as long as new ones from the monthly pool of questions are available.

The progress page also includes an archive with historic performance data in order to amplify the motivation to complete the monthly set of questions. With a similar intention, the game offers a leaderboard (see Figure 3) that lists the top three players of each month. This feature complements the progress page archive and the dynamic sidebar elements, designed to support the players' long-term motivation and to grant them "bragging rights" as they become part of the official history of the game.



**Figure 3.** Leaderboard of the Language Quiz

In addition to the progress indicators for players, the Language Quiz as an open platform for trusted third parties to upload their own tasks provides a progress monitoring system for administrators as well (the Climate Challenge with its curated system of monthly tasks does not require such a feature). This system is embedded in the API. After creating a job for the game, the task owner will receive a Job ID. Afterwards he can send a status request with the ID to the game. This works both for Crowdfunder and game tasks. The returned result indicates if the job has been completed; and, if not, how many answers still have to be collected. For game tasks, it is also possible to request the preliminary set of answers collected up until now, even if the job has not yet been completed.

## Quality Control

Measures of quality control vary according to the type of task, and according to the application. Quality control measures in the Climate Challenge, for example, are only required for specific tasks - e.g. the multiple choice questions or the language resource acquisition task for assessing the sentiment of extracted climate change keywords.

### *Climate Challenge*

Task types such as the pledges or the polls, where the answer is a matter of opinion and cannot be considered right or wrong, do not yield quality metrics. Similarly, answers to the prediction questions cannot be immediately evaluated because the real-world answers are not known when the question is asked. The multiple choice questions, by contrast, are instantly evaluated and the user will receive points for right answers and a point penalty for wrong answers. In all four cases it is not possible for players to cheat.

For the sentiment detection tasks in the Climate Challenge (as well as all tasks in the Language Quiz), however, the computation of quality metrics is important. One possible approach would be the evaluation of users based on gold standard questions. Because of the recurring nature of the game and the limited testing ability of the other task types, a cross-validation approach has been chosen (which can be configured via the uComp API) - collecting a total of ten answers per question. The “correct” sentiment will be determined based on the mean of the assessments, and is then used to award game points. The standard deviation of the received answer serves as an indicator of ambiguity, and can be used to decide whether the acquired language resources are valid and should be considered for further processing.

### *Language Quiz*

Cheating prevention becomes essential in the case of the Language Quiz, which offers the option to send tasks to Crowdfunder and/or to recruit game participants via Crowdfunder. The API allows developers to include gold standard data in their set of uploaded tasks. This data is then forwarded to Crowdfunder to include test questions and refuse payment based on the number of failed tasks.

*Hybrid Crowdsourcing Model.* If Crowdfunder is only used as a promotional tool to recruit players for a certain task, there are two systems in place to prevent cheating:

- Users have to insert a code in the Crowdfunder submission form, which they will receive after having answered 20 questions in the game. This code is a complex set of characters and special characters, which cannot be guessed.

- Crowdfunder allows to specify a minimum time that it usually takes to complete a task (filtering out users who do not provide a genuine answer); e.g. if the minimum time is set to two minutes, but the answer is received within 30 seconds, it can be assumed that the user has not really tried to solve the problem.

*Test Questions.* If Crowdfunder is used to generate the results (i.e. the uComp API only serving as an interface between the task owner and the crowdsourcing marketplace), every player has to correctly answer three test questions as part of the batch of 20 questions that need to be answered to receive payment. This flags users who try to cheat, or those who lack the expertise for a specific task type (e.g. a lack of language skills). Answers of players who fail one test question will be disregarded, and they will not receive the payment. This system will be explained in advance to ensure that players pay attention to the quality of their answers.

*Player Selection.* The Crowdfunder marketplace distinguishes three levels of players, depending on the average quality of their answers. If tasks are made available to all levels, they will be solved more quickly at reduced accuracy. The quality requirements and additional complexity introduced by redirecting participants to the game interface suggests to only consider Level 3 players to solve the tasks, except in very time-critical situations. Those players will always be more careful when solving a task, as they risk their Level 3 status if they fail too many test questions.

## Incentive Mechanisms

Crowdsourcing strategies can be classified into three categories (Sabou et al. 2014): *mechanised labour* offering financial rewards, *games with a purpose* with gameplay incentives, and approaches based on *altruistic work*. The Language Quiz uses the first two strategies, while the Climate Challenge relies on gameplay features and altruistic incentives.

*Ranking System.* The main incentive mechanisms in the Climate Challenge are traditional gameplay mechanisms, which are used to engage the users: (i) a real-time leaderboard that shows how they bypass other players and rank higher while playing the game; (ii) a leaderboard with historic data on the top three monthly players, and (iii) a leveling system to motivate users to reach higher levels and surpass their previous record.

*Sharing of Game Content.* Players can also invite their friends to participate in the game, in general or specific to certain tasks - players can use Facebook, Twitter, Google+ or E-Mail to notify others when they accept a pledge. A strong incentive to share are bonus points that players receive when their friends follow their invitation.

*Prizes and Bragging Rights.* Each month, the best players receive small prizes such as coffee mugs or t-shirts, which on purpose do not represent a significant monetary value. But even with this physical reward system in place, we consider the in-game mechanics and the achieved bragging rights as the most important incentives. The game relates to climate change as an important environmental topic, where players can demonstrate their knowledge vis-à-vis other players and friends. Previous work has shown that games that tackle a specific domain will benefit from the intrinsic motivation of supporting a good cause (Rafelsberger and Scharl 2009).

*Notification System.* Another mechanism to motivate players are push notifications via the Facebook API and/or E-Mail notifications, have proven to be a very effective tool to (re-)engage players at the start of a new monthly game round. Users who have signed up are notified once new answers to their open questions are available. Notifications are only sent out once or twice each month, to minimize the risk of losing subscribers. Weekly updates only go to players who play regularly, the other players will just receive rare reminders if they want to rejoin the game or other major announcements such as the final game round of the year. Players have the option to unsubscribe from the notifications and not receive any E-Mail or Facebook notifications at all.

*Multilinguality.* For the Language Quiz, a major benefit of using Crowdfunder as a promotional tool is the ability to recruit native speakers in many different languages (Sabou et al., 2014). Since the support for different languages is a key feature of the Language Quiz, this feature is very important for the overall success of the game. Players also remain motivated because they know that the task would contain test questions and that their payment (and indirectly their Crowdfunder level) depends on the accuracy of their answers.

## Summary

The human computation framework of uComp is flexible and modular in terms of progress monitoring, quality control, and incentive mechanisms for different types of tasks. This flexibility is reflected in the development and launch of two rather different applications, the Climate Challenge and the Language Quiz:

- For *gameplay-driven* games like the Climate Challenge, incentive mechanisms include leaderboards, real-time level indicators and progress graphs. Cross-player validation prevents cheating, while social sharing and other social display mechanics increase the applications' visibility.
- *Task-centric* games like the Language Quiz can benefit from promotional campaigns on marketplaces such as Crowdfunder to recruit players and kickstart the game. This approach includes test questions and leverages internal Crowdfunder mechanism (e.g., selecting only Level 3 users).

The human computation framework of uComp is well suited to handle both type of games, and includes the necessary tools to keep heterogeneous user groups engaged while ensuring ensure a consistent high quality of answers.

## References

- Huberman, B.A. (2008). "Crowdsourcing and Attention", *IEEE Computer*, 41(11): 103-105.
- Rafelsberger, W. and Scharl, A. (2009). Games with a Purpose for Social Networking Platforms. 20th ACM Conference on Hypertext and Hypermedia. C. Cattuto et al. Torino, Italy: Association for Computing Machinery: 193-197.
- Sabou, M., Bontcheva, K., Derczynski, L. and Scharl, A. (2014). Corpus Annotation through Crowdsourcing: Towards Best Practice Guidelines. 9th Language Resources and Evaluation Conference (LREC-2014). Reykjavik, Iceland: European Language Resources Association (ELRA).
- Sabou, M., Bontcheva, K., Scharl, A. and Föls, M. (2013). Games with a Purpose or Mechanised Labour? A Comparative Study. 13th International Conference on Knowledge Management and Knowledge Technologies (i-KNOW 2013). Graz, Austria: Forthcoming.
- Sabou, M., Scharl, A., and Föls, M. (2013). Crowdsourced Knowledge Acquisition: Towards Hybrid-Genre Workflows. *International Journal on Semantic Web and Information Systems*, 9(3), 14-41.
- Scharl, A., Föls, M. and Herring, D. (2015). "Climate Challenge - Raising Collective Awareness in the Tradition of Games with a Purpose", 14th Brazilian Symposium on Human Factors in Computer Systems (IHC-2015). Salvador, Brazil. 506-509.
- Scharl, A., Sabou, M., Gindl, S., Rafelsberger, W., & Weichselbraun, A. (2012). Leveraging the wisdom of the crowds for the acquisition of multilingual language resources.

## Extended Bibliography

- Ahn, L.v. (2006). "Games with a Purpose", *Computer*, 39(6): 92-94.
- Ahn, L.v. and Dabbish, L. (2008). "Designing Games with a Purpose", *Communications of the ACM*, 51(8): 58-67.
- Ben Kirman, S.L., Conor Linehan (2009). *Gaming On and Off the Social Graph: The Social Structure of Facebook Games*. 4th IEEE International Conference on Computational Science and Engineering. Vancouver, Canada: IEEE Computer Society: 627-632.
- Chernova, S., Orkin, J. and Breazeal, C. (2010). *Crowdsourcing HRI through Online Multiplayer Games*. AAAI Fall Symposium. Arlington, USA: 14-19.
- Dahab, K. and Belz, A. (2010). *A Game-based Approach to Transcribing Images of Text*. International Conference on Language Resources & Evaluation (LREC-2010). N. Calzolari et al. Valletta, Malta: European Language Resources Association.
- Doan, A., Ramakrishnan, R. and Halevy, A.Y. (2011). "Crowdsourcing Systems on the World Wide Web", *Communications of the ACM*, 54(4): 86-96.
- Heer, J. and Bostock, M. (2010). *Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design*. 28th ACM Conference on Human Factors in Computing Systems (CHI-2010). Atlanta, USA: 203-212.
- Hees, J., et al. (2011). "BetterRelations: Using a Game to Rate Linked Data Triples", 34th Annual German Conference on Artificial Intelligence (= LNCS, Vol. 7006). Eds. J. Bach and S. Edelkamp. Berlin: Springer. 134-138.
- Johnson, P.M., Xu, Y., Brewer, R.S., Lee, G.E., Katchuck, M. and Moore, C.A. (2012). *Beyond kWh: Myths and Fixes for Energy Competition Game Design*. Meaningful Play Conference. East Lansing, USA.
- Kirman, B., Lawson, S., Linehan, C., Martino, F., Gamberini, L. and Gaggioli, A. (2010). *Improving Social Game Engagement on Facebook through Enhanced Socio-Contextual Information*. 28th ACM Conference on Human Factors in Computing Systems. Atlanta, USA: ACM: 1753-1756.
- Kittur, A., Smus, B. and Kraut, R. (2011). *CrowdForge: Crowdsourcing Complex Work*. Conference on Human Factors in Computing Systems (CHI-2011). Vancouver, Canada: Association for Computing Machinery: 1801-1806.
- Kuo, Y.-I., Lee, J.-C., Chiang, K.-y., Wang, R., Shen, E., Chan, C.-W. and Hsu, J.Y.-J. (2009). *Community-Based Game Design: Experiments on Social Games for Commonsense Data Collection*. ACM SIGKDD Workshop on Human Computation (HCOMP-2009). New York, USA: 15-22.

Lieberman, H., Smith, D.A. and Teeters, A. (2007). Common Consensus: A Web-based Game for Collecting Commonsense Goals. International Conference on Intelligent User Interfaces (IUI-2007). Honolulu, USA.

Munro, R., Bethard, S., Kuperman, V., Lai, V.T., Melnick, R., Potts, C., Schnoebelen, T. and Tily, H. (2010). Crowdsourcing and Language Studies: The New Generation of Linguistic Data. NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk. Los Angeles, USA: 122-130.

Parent, G. and Eskenazi, M. (2011). Speaking to the Crowd: Looking at Past Achievements in Using Crowdsourcing for Speech and Predicting Future Challenges. 12th Annual Conference of the International Speech Communication Association (Interspeech-2011). Florence, Italy: 3037-3040.

Sabou, M., Bontcheva, K. and Scharl, A. (2012). Crowdsourcing Research Opportunities: Lessons from Natural Language Processing. 12th International Conference on Knowledge Management and Knowledge Technologies (iKnow-2012). S. Lindstaedt and M. Granitzer. Graz, Austria: Forthcoming.

Siorpaes, K. and Hepp, M. (2008). "Games with a Purpose for the Semantic Web", IEEE Intelligent Systems, 23(3): 50-60.

Vickrey, D., Bronzan, A., Choi, W., Kumar, A., Turner-Maier, J., Wang, A. and Koller, D. (2008). Online Word Games for Semantic Data Collection. Conference on Empirical Methods in Natural Language Processing (EMNLP-2008). Stroudsburg, USA: Association for Computational Linguistics: 533-542.

Wang, A., Hoang, C.D.V. and Kan, M.-Y. (2011). "Perspectives on Crowdsourcing Annotations for Natural Language Processing", Language Resources and Evaluation: Forthcoming.

Weichselbraun, A., Gindl, S. and Scharl, A. (2011). Using Games with a Purpose and Bootstrapping to Create Domain-Specific Sentiment Lexicons. 20th ACM Conference on Information and Knowledge Management (CIKM-2011). Glasgow, UK: Association for Computing Machinery: 1053-1060.

Wiggins, A. and Crowston, K. (2011). From Conservation to Crowdsourcing: A Typology of Citizen Science. 44th Hawaii International Conference on System Science (HICSS-44). Koloa, USA: IEEE Computer Society.

Wolf, L., Knuth, M., Osterhoff, J. and Sack, H. (2011). RISQ! Renowned Individuals Semantic Quiz: A Jeopardy Like Quiz Game for Ranking Facts. 7th International Conference on Semantic Systems (I-Semantics 2011). C. Ghidini et al. Graz, Austria: ACM.