

Evaluating a User Story Based Recommendation System for Supporting Development Processes in Large Enterprises

Maria Lusky, Matthias Jurisch, Stephan Böhm

Faculty of Design – Computer Science – Media
RheinMain University of Applied Sciences
Wiesbaden, Germany

Email: {maria.lusky, matthias.jurisch,
stephan.boehm}@hs-rm.de

Katharina Kahlcke

User Experience Consulting
DB Systel GmbH
Frankfurt, Germany

Email: katharina.kahlcke@deutschebahn.com

Abstract—In mobile application development projects, large enterprises have to face special challenges. To meet these challenges and to meet today’s high expectations on user centered design, inter-project knowledge transfer of software artifacts becomes an important aspect for large software development companies. For supporting this kind of knowledge transfer, we propose an approach for a recommendation system based on textual similarity of user stories that are computed via standard information retrieval techniques. We also present a three-step evaluation of this approach, comprising data analysis, a survey and a user study. The results tend to support our approach and not only show that user story similarity rated by users and rated by an algorithm correlates, but also demonstrate a strong relation between user story similarity and their usefulness for inter-project knowledge transfer.

Keywords—Mobile Enterprise Applications; User Stories; Recommendation Systems; User Centered Design.

I. INTRODUCTION

In recent years, the user centered design approach has become an integral part of software development, and also for mobile application (app) development. Often, at the beginning of an agile development process, requirements for an app are analyzed and are written down in the form of user stories. They are short requirement descriptions from the user’s point of view. Based on these user stories, during the further development process other software artifacts, such as documentation, screen designs, or source code are created to support the development process. Especially in large enterprises, the reuse of these software artifacts can save time and resources, since large software development companies are facing several challenges: There are multiple development projects at the same time, resulting in a large number of software artifacts. Due to a lack of time, these artifacts are often not properly documented in order to support a reuse of these materials. Furthermore, team members often do not know if there is a project with similar requirements and which coworker they can contact about a reuse of software artifacts. In general, large software development companies deal with lack of transparency in development projects, contact persons, and software artifacts.

Saving time and resources through reuse is especially desirable for enterprises in the Mobile Enterprise Application (MEA) market: due to digitalization trends, these enterprises

are developing many apps for various customers at the same time [1]. Nevertheless, quick time to market is important because of a rapidly changing mobile ecosystem. Additionally, enterprises can only access a limited number of specialists that should focus on demanding tasks and work on difficult problems that have not been solved in the company already. This is the context of the project PROFRAME [2], that tries to support the reuse in software project using several approaches. Given this background, in this paper we propose an approach that supports the reuse of software artifacts in mobile app development projects based on textual similarity of user stories. In a previous paper on this problem, we showed that similar user stories can be identified via classical methods of information retrieval [3]. In the following evaluation, we investigate how well these methods work in a real world dataset and how useful our approach is for employees of a large software development company. Therefore, Section II provides an overview on related work. Section III introduces our approach. An evaluation is described in Section IV and its results are presented in Section V. These results are discussed afterwards in Section VI. Section VII concludes this paper and gives an outlook on further research.

II. RELATED WORK

Supporting reuse in the context of software development can be facilitated in many ways. One of these ways is best practice sharing, where good solutions for common problems are shared. However, this requires a lot of work and time to find common problems and respective solutions. Especially in the context of MEA development, time is an important factor. Therefore, supporting this process with automated approaches seems to be promising. An automated approach in this area is the use of recommendation systems [4]. Recommendation systems recommend items to users based on item similarities or preferences of similar users. This idea can be applied to software engineering, where a recommendation system can recommend software engineering artifacts to developers [5]. The goal of these recommendation systems is mainly to support the software development process, especially focused on the implementation phase and fixing bugs.

An important field in this area is issue triage. This field deals with supporting the management of bug reporting systems. This includes both recommending specific developers for

a given bug report, as well as detecting duplicate bugs. Usually, standard information retrieval techniques or shallow machine learning approaches are used [6]: An approach by Runeson et al. [7] is based on information retrieval techniques and tries to detect duplicates. Other approaches build on including other information besides text, e.g., execution information [8]. In [9], a framework for building recommendation systems for issue triage is presented. This is done by linking both developers and bug reports to project components. Besides using classical information retrieval methods, more recent approaches use deep learning-techniques like word embeddings [10]. While this area also includes computing similarities between textual descriptions in the area of software development, there are some important differences: (1) bug reports are often written from a very developer-centric perspective. (2) They usually contain a lot of technical information like log output. (3) The main goal of issue triage is not to support reuse, but to support bug management tasks.

Another approach to improve the knowledge management in software development projects is to document and store project information in an accessible way, e.g., in architectural knowledge management tools [11]. These approaches have also been applied in industrial case studies and were deemed fit for usage in an industry context: [12] evaluated a semantic architectural knowledge management tool that is based on existing data on software design patterns and their application in software projects. However, if this kind of data is not already available the overhead for documenting usage of design patterns can be too high for an application in practice. This is especially an issue for the fast-paced market of mobile enterprise applications.

In the last decade, user stories as a user-centric representation of requirements were introduced [13]. A typical user story is at most two sentences long and consists of a reference to a certain type of user, a description of an activity this user wants to do with the software and a reason why this will help the user. As an attachment to the user story, acceptance criteria add more detailed information to the user story. Only few approaches exist to support software reuse in the context of user stories: [14] proposes a recommendation system based on user stories and evaluates this system on a project history. However, it is not clear how helpful these recommendations would be when actually working on a new project. In our previous work [3], we evaluated how well information-retrieval-based approaches can distinguish between two types of user stories and which aspects of the user story are important to it.

In this paper, we expand this previous work to an evaluation of how useful recommendations on a real-world software engineering dataset are and what information needs to be contained in these recommendations to make them actually helpful. This issue has not been addressed by the approaches we mentioned in this section and is required to make recommendations in the context of user stories usable in practice. The only way to work on this issue is to conduct a survey with practitioners from the industry.

III. RECOMMENDATION APPROACH

Our general approach to supporting reuse is to use established techniques from information retrieval in order to recommend textually similar user stories to the story a participant in an app development-project is currently working on. The

information that is attached to the recommended user stories (e.g., screen designs, textual documents or source code) can then be used to support current efforts. In this way, team members could reuse results from different projects without previously knowing about these projects.

To find textually similar user stories, a search based on the well-known information retrieval approach Term Frequency-Inverse Document Frequency (TF-IDF) and stop word removal is used. Stop words are words that occur frequently in texts so that they do not contain useful information. Examples for stop words are "I", "the", "a", etc. These words are removed from the user stories before processing the user stories with TF-IDF, which represents texts as follows: Each document d (i.e., a user story) is represented by a vector w_d , that contains an entry for each term used in the dataset. Each vector component $w_{d,t}$ represents the importance of a term t for the document d . This representation is computed by the frequency of a given term in the document $tf_{d,t}$ multiplied by the inverse document frequency $\log \frac{N}{df_t}$, where N is the number of all documents and df_t is the number of occurrences for a given term in all documents. This yields the following formula for a document's vector representation:

$$w_{d,t} = tf_{d,t} * \log \frac{N}{df_t}$$

To compute the similarity between documents, the cosine of the angle of two vectors is used. The naive approach for similarities would be to compute the euclidian distance between vectors, however this would favour documents with similar lengths.

Cosine similarities are then used to order texts regarding their relative similarities. Thus, we do not use similarity scores as an absolute value, but only to distinguish between more and less similar documents. To find similar user stories to one given user story, the similarity is computed according to the described procedure. User stories are then ordered by their similarity and the user story with the highest score is considered the most similar.

IV. EVALUATION

We evaluated the recommendation approach described in Section III and the importance of using recommendation systems for reuse in development processes in large enterprises by investigating the following research questions:

- 1) Which kind of knowledge transfer is already being practiced?
- 2) Can an automated recommendation system be useful for supporting knowledge transfer?
- 3) Is there a relation between user story similarity and their usefulness?

A. Methodology

To answer these research questions, we conducted an evaluation comprising three steps: In the first step, we analyzed a dataset of user stories from a large German software development company. In the second step, we distributed a questionnaire covering questions about practices in inter-project knowledge transfer in general. In the third step, we invited participants to single sessions where they were asked to

solve tasks focusing on user stories in inter-project knowledge transfer.

The number of participants in this study was limited to a small number due to the testing requirements: all participants had to come from the same company with a specific expertise on the implementation of the user stories and as references for the similarity comparison. Thus, the results are far from representative and cannot be considered as an empirical validation. However, the results of this prestudy can give some critical and usable expert feedback on the potential usefulness and applicability of our approach.

B. Dataset

To evaluate the usefulness of recommendations based on user stories in the area of Mobile Enterprise Application Development, we used a dataset of real-world user stories out of nine app development projects provided by an industry partner. The dataset contains 591 user stories, of which 355 are long enough to contain meaningful information. User stories were considered long enough when they were at least 80 characters long, which is roughly two times the length of only the formal aspects of a user story description. This boundary was set by investigating example user stories. A histogram of story length (in characters) is shown in Figure 1. From the distribution of story length and the standard deviation, we can already conclude that the dataset is very heterogeneous, as could be expected in a real-world dataset. The data is not only heterogeneous regarding the textual length, but also regarding their specificity and their degree of abstraction. For example, some user stories describe fixing typos in data protection regulation informations, while others describe a high level view of a location-based service.




Figure 1. Distribution of User Story Length

For each user story in the dataset, we computed the top five most similar user stories according to TF-IDF, with cosine similarity both for stories from different projects as well as stories from the same project. A histogram of cosine similarity values between all user stories is shown in Figure 2. Stories in the same project are given higher similarity values than stories from different projects, which indicates that it is possible to differentiate between projects using cosine similarities.

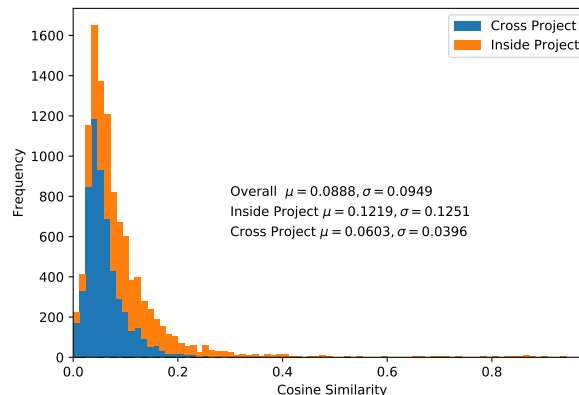


Figure 2. Distribution of User Story Similarities

C. Survey

To get an overall overview of the requirements in user story-centered reuse, we designed a questionnaire that comprised ten questions about inter-project knowledge transfer. The questionnaire was online for 17 days and was distributed among the employees of a large German software development company. It was also used for recruitment of participants for the user study. First, the participants had to specify their field of activity. We then explained to them our approach for inter-project knowledge transfer that underlay the questions, that is the re-use of software artifacts and knowledge, such as user stories, screen designs, documentation or source code, during development projects of mobile applications.

We asked the participants, how useful such a knowledge transfer is and in which way and how regularly it is already being practiced in their department. Further, they had to name obstacles that occur with inter-project knowledge transfer. Then, we asked them to rate the usefulness of particular software artifacts in this context and they had to assess the viability of such a knowledge transfer in their department. They were asked to rate the usefulness of a software that would support knowledge transfer. Lastly, the participants were asked to rate the importance of additional information to specific software artifacts on a five-point scale. Importance may differ from usefulness in certain situations, since it is used to prioritize between different potentially useful artifacts.

D. User Study

The user study was carried out in one-on-one sessions with employees of a large German software development company. Each session lasted 20 to 30 minutes. We selected three user stories for which related user stories were known to be in the dataset. For each of these *reference* user stories, we computed the most similar user stories from the different projects with varying levels of similarity: one user story that the algorithm listed most similar, one that it listed as medium similar, and one that it listed as less similar, which lead us to three user story groups, one for each reference story.

Based on the reference user stories, the participants were asked to solve three tasks. First, they had to rank the user stories obtained by the algorithm regarding their similarity to

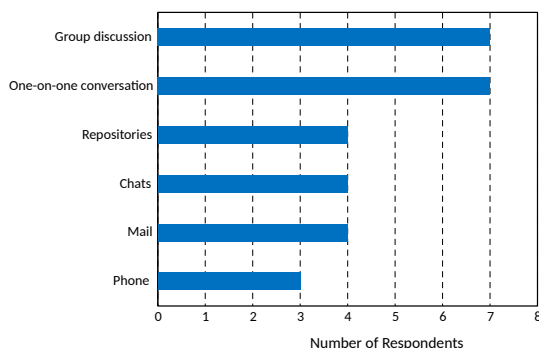


Figure 3. Currently Used Types of Knowledge Transfer

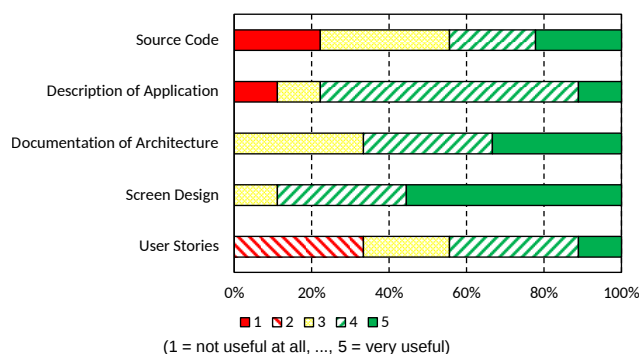


Figure 5. Perceived Usefulness of Artifacts

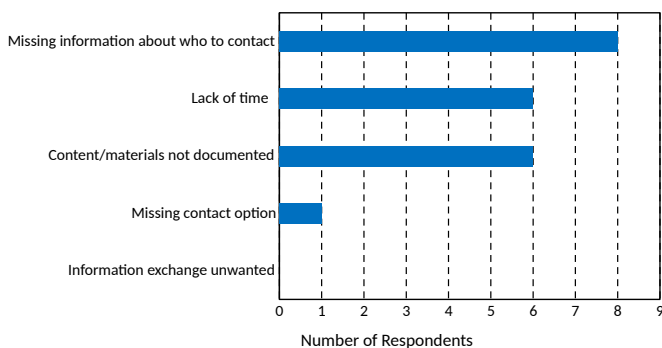


Figure 4. Existing Barriers for Knowledge Transfer




Figure 6. Responses on Potential Recommendation System Artifacts

the reference user story from the most similar to the least similar one. Then, they should rate the usefulness of each of the similar user stories. To determine the usefulness of the user stories, participants were told to estimate, how much artifacts (e.g., source code, design documents or documentation) produced during an implementation of a ranked user story could contribute to the implementation of the corresponding reference user story. Note that this is not included in similarity aspects: user stories can cover a roughly similar topic, however different levels of abstraction, different user types or platforms or technical aspects could make it impossible to actually reuse the results of the implementation of a user story in a different context. Such user stories would be considered similar by users, but finding these stories would not actually support reuse of artifacts related to one user story to another. Concluding the session, they were asked to name additional information that should be provided by the recommendation system in order to support the implementation of the reference user story.


V. RESULTS

The questionnaire was answered by nine employees of a large German Software development company. While nine participants are not enough to allow a detailed statistical analysis, this number is in general considered enough for usability testing [15]. Eight participants specified their field of expertise as conception, one as implementation.

All of the participants rated the knowledge transfer described by us (that is, the re-use of software artifacts and

knowledge, such as user stories, screen designs, documentation or source code, during development projects of mobile applications), using a five-point scale, as useful or rather useful (median=5; maximum=5; minimum=4). The currently used types of knowledge transfer selected from a list of pre-made options are shown in Figure 3. All of the participants stated that they already practiced this kind of knowledge transfer, seven via one-on-one conversations or group conversations, four via e-mail, chats or by using knowledge bases, and three via phone calls. On average, each person practices three methods of knowledge transfer, and eight practice it as needed. Obstacles for knowledge transfer selected by participants from a list of possible obstacles are shown in Figure 4. The most often named obstacle was missing information about a contact person (eight), followed by missing documentation of content and materials and lack of time (six respondents each). The participants described further obstacles as being unaware of the existence of reusable materials, as well as not knowing where to look for information regarding reusable artifacts.

User ratings for usefulness of artifacts for Knowledge transfer on a five-point scale are shown in Figure 5. Screen designs were rated as most useful (median=5, maximum=5, minimum=3), followed by documentation of the software architecture (median=4, maximum=5, minimum=3). Ratings for potential usefulness, implementability and importance are shown in Figure 6. Regarding the viability of such a knowledge transfer in their department and in relation to specific



(a) Ranking of the user stories by the algorithm and mean ranking by the participants. (b) Mean estimated usefulness and mean ranking by participants per user story.

Figure 7. User study results

software artifacts, the highest implementability was considered for screen designs, followed by documentation of the software architecture and use case descriptions. Furthermore, the answers revealed that most knowledge transfer that is already practiced concerns screen designs (done by 4 participants), documentation of the software architecture (2 participants) and user stories (1 participant).

The participants rated a software solution for supporting knowledge transfer on a five-point scale as rather useful (median=4; maximum=5; minimum=1), with 6 participants considering it rather useful or useful. Regarding the importance of additional information, information on use case descriptions were rated as most useful (median=4; maximum=5; minimum=3). In general, any kind of additional information (e.g., source code, screen designs, architecture documentation) was rated as "rather important" for all kinds of software artifacts.

Of the eight participants of the user study, all were working in conception respectively design. Results of the user study are shown in Figure 7. Results of the first task show that the participants ranked the user stories similar to the ranking of the algorithm. Figure 7a shows the ranking of the result user story similarity to the reference story by the algorithm and the mean ranking by the participants. The data shows that user story similarity of the algorithm seems to resemble the perceived user story similarity by humans: The three user stories ranked as most similar by the algorithm also got the highest similarity rankings by the participants. The user stories ranked as second by the algorithm were partially ranked as more and partially as less similar, but in general reflect the algorithmic ranking. For stories that are not obviously the least or most similar, this result was to be expected. Accordingly, the three least similar user stories of the participants match those ranked by the algorithm.

Further on, for each user story in the three user story groups, we calculated the mean similarity ranking to the reference story given by the participants and the mean useful-

ness rating, according to the rated usefulness of a computed user story for the implementation of the reference story. As Figure 7b shows, there are two groups of user stories that are delineated from each other. The first group has higher usefulness values (3,50 to 4,88) and higher similarity rankings (1,00 to 1,75), while the second group has lower usefulness values (1,63 to 2,25) and lower similarity rankings (2,38 to 3,00). However, the user story with the highest usefulness (4,88) is ranked with medium similarity (1,75), while the two user stories with the lowest usefulness (1,63 and 1,75) are ranked as least similar.

VI. DISCUSSION

The results of the questionnaire show that, in general, people appreciate knowledge transfer and that our approach for it meets the users' needs. This is also reflected in the fact that 7 of our participants already practice this kind of knowledge transfer. Although, direct contact and personal conversations are the preferred ways of doing so. Electronic ways for contacting each other are rarely used. One reason for that might be that electronic methods, such as emails, chats or knowledge bases, do not meet the user needs for knowledge transfer. Nevertheless, for all these methods the user needs to know, which person can be contacted for further information – our approach takes this important feature into account, that was also the most often named obstacle for knowledge transfer. The second obstacle is insufficient documentation – this problem is also taken account of in our approach, since it simplifies documentation by searching existing software artifacts based on similarities. In general, the answers confirm that our approach addresses the right issues and helps to eliminate the obstacles for knowledge transfer that have been named by the participants. The software tool proposed by us was said to be most useful for screen design. This answer is not unusual, since almost all participants of our questionnaire work in design and conception. However, screen design is their main field of activity and thus, we consider it positive that our

approach is rated as useful for this field. Further, the viability was rated highest for screen designs. These results give some evidence that our approach can create additional value in one of the most important fields for knowledge transfer. All in all, two thirds of the participants consider our approach useful or very useful.

The results of the user study show that the user story similarity of the algorithm seems to be connected to the perceived human user story similarity. Furthermore, user story similarity mostly coincides with their usefulness. The data indicates that a low similarity implies a low usefulness. However, the most similar stories are not always the most useful ones, but to some extent a high similarity seems to be connected to higher usefulness. This confirms our assumption that similarity between user stories and usefulness for the implementation of another user story are not necessarily the same.

VII. CONCLUSIONS AND FURTHER RESEARCH

In conclusion, the evaluation provided valuable findings, so that our research questions can be answered as follows:

- 1) Which kind of knowledge transfer is already being practiced?
Mainly, knowledge transfer takes place in personal conversations between two people and groups. It is not carried out on a regular basis, but as required. Our results suggest that everyone does practice knowledge transfer in one way or the other.
- 2) Can an automated recommendation system be useful for supporting knowledge transfer?
Our results show that an automated recommendation system is a useful tool for supporting knowledge transfer, especially for screen designs.
- 3) Is there a relation between user story similarity and their usefulness?
The results of our user study indicate a relation between user story similarity and their usefulness. Further, there also is a connection between user story similarity rated by an algorithm on the one hand and humans on the other hand.

As a next step, another iteration of the evaluation could be made in different companies, in order to receive results that are applicable in several contexts of work. Also, more approaches for computing similar user stories could be evaluated: A comparative study of textual similarity approaches such as word movers distance or taking metadata into account, would provide valuable insights.

ACKNOWLEDGMENT

This work was funded by the German Federal Ministry of Education and Research, grant no. 03FH032PX5; the PROFFRAME project at RheinMain University of Applied Sciences. All responsibility for the content of this paper lies with the authors.

REFERENCES

- [1] A. Giessmann, K. Stanoevska-Slabeva, and B. de Visser, "Mobile enterprise applications—current state and future directions," in *System Science (HICSS)*, 2012 45th Hawaii International Conference on, Jan 2012, pp. 1363–1372.
- [2] M. Jurisch, B. Iglar, and S. Böhm, "PROFRAME: A Prototyping Framework for Mobile Enterprise Applications," in *CENTRIC 2016, The Ninth International Conference on Advances in Human-oriented and Personalized Mechanisms, Technologies, and Services*, 2016, pp. 7–10.
- [3] M. Jurisch, M. Lusky, B. Iglar, and S. Boehm, "Evaluating a recommendation system for user stories in mobile enterprise application development," *International Journal On Advances in Intelligent Systems*, vol. 10, no. 1 and 2, 2017, pp. 40–47.
- [4] M. Robillard, R. Walker, and T. Zimmermann, "Recommendation systems for software engineering," *IEEE Software*, vol. 27, no. 4, 2010, pp. 80–86.
- [5] D. Cubranic, G. C. Murphy, J. Singer, and K. S. Booth, "Hipikat: A project memory for software development," *IEEE Trans. Softw. Eng.*, vol. 31, no. 6, Jun. 2005, pp. 446–465. [Online]. Available: <https://doi.org/10.1109/TSE.2005.71>
- [6] A. Goyal and N. Sardana, "Machine learning or information retrieval techniques for bug triaging: Which is better?" *e-Informatica Software Engineering Journal*, vol. 11, no. 1, 2017, pp. 117–141.
- [7] P. Runeson, M. Alexandersson, and O. Nyholm, "Detection of duplicate defect reports using natural language processing," *Proceedings - International Conference on Software Engineering*, 2007, pp. 499–508.
- [8] X. Wang, L. Zhang, T. Xie, J. Anvik, and J. Sun, "An Approach to Detecting Duplicate Bug Reports using Natural Language and Execution Information," *Proceedings of the 30th international conference on Software engineering*, 2008, pp. 461–470.
- [9] J. Anvik and G. C. Murphy, "Reducing the Effort of Bug Report Triage: Recommenders for Development-Oriented Decisions," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 20, no. 3, 2011, pp. 10:1–10:35.
- [10] S. Mani, A. Sankaran, and R. Aralikkatte, "Deeptriage: Exploring the effectiveness of deep learning for bug triaging," *arXiv preprint arXiv:1801.01275*, 2018.
- [11] R. Capilla, A. Jansen, A. Tang, P. Avgeriou, and M. A. Babar, "10 years of software architecture knowledge management: Practice and future," *Journal of Systems and Software*, vol. 116, 2016, pp. 191–205.
- [12] M. Sabou et al., "Exploring enterprise knowledge graphs: A use case in software engineering," in *European Semantic Web Conference*. Springer, 2018, pp. 560–575.
- [13] M. Cohn, *User Stories Applied: For Agile Software Development*. Redwood City, CA, USA: Addison Wesley Longman Publishing Co., Inc., 2004.
- [14] H. Pirzadeh, A. D. S. Oliveira, and S. Shanian, "ReUse : A Recommendation System for Implementing User Stories," in *International Conference on Software Engineering Advances*, 2016, pp. 149–153.
- [15] J. Nielsen, "How many test users in a usability study?" <https://www.nngroup.com/articles/how-many-test-users/>, website. [retrieved: August, 2018], 2012.