Quantitative Requirements Prioritization from a Pre-development Perspective

Enrico Johansson[1](✉), Daniel Bergdahl[1], Jan Bosch[2], and Helena Holmström Olsson[3]

[1] Sony Mobile Communications Inc, Lund, Sweden
{Enrico.Johansson,Daniel.Bergdahl}@sonymobile.com
[2] Department of Computer Science and Engineering,
Chalmers University of Technology, Gothenburg, Sweden
Jan.Bosch@chalmers.se
[3] Department of Computer Science, Malmö University, Malmö, Sweden
Helena.Holmstrom.Olsson@mah.se

Abstract. Feature content in system releases tends to be prioritized using limited amounts of qualitative user input and based on the opinions of those in product management. This leads to several problems due to the wasteful allocation of R&D resources. In this paper, we present the results of our efforts to collect quantitative customer input before the start of development using mock-ups and surveys for a mobile application developed by Sony Mobile. Our research shows that (1) collecting quantitative feedback before development is feasible, (2) the data collected deviates from the original feature prioritization, i.e. it is beneficial and (3) the data gives further insight in requirement prioritization than a qualitative method could have provided.

Keywords: Requirements engineering · Customer data · Survey · Mock-up · Data-driven development · Case study

1 Introduction

In most software development companies, the pre-development is the phase in which decisions are taken on whether to develop a feature or not. In this phase, during requirement prioritization, the expected value of a feature is estimated and if the outcome is positive the feature is developed. There are, however, a number of problems associated with this.

First, the estimated value of a feature is typically based on very limited data that can prove whether the estimation is correct. As a result, feature prioritization becomes an opinions-based process [1].

Second, the estimations that are done in the pre-study phase are typically based on limited amounts of qualitative feedback from customers. Data is collected by asking customers what they want and by observing what they do and the output is a limited set of individual customer opinions and experiences regarding product use. While this feedback is valuable, it does not represent a large customer base and it does not reveal actual product usage. Ideally, qualitative customer feedback in the pre-development

T. Rout et al. (Eds.): SPICE 2015, CCIS 526, pp. 58–71, 2015.
DOI: 10.1007/978-3-319-19860-6_6

phase should be complemented with quantitative data that confirm individual perceptions but this has proven difficult to accomplish [2]

Third, due to lack of mechanisms to measure feature usage, companies invest in developing features that have no proven customer value. Often, and as recognized in previous research [3], the majority of features in a system are never, or only very seldom, used. This situation could be avoided if accurate data collection mechanisms were in place, allowing companies to allocate resources to development with proven customer value.

There is a need to overcome these problems to stay competitive. The one who makes the best prediction (by prioritizing the features with highest value) not only wins the market shares but also reduces waste in the development cycle. However, prioritization is a challenging part of the requirement engineering processes since it is trying to predict the future. This is especially true in a market driven context addressing end-users as customers.

A number of qualitative prioritization methods are defined in requirement engineering processes. Qualitative prioritization methods are often by nature subjective and involve for example guessing or weighing requirements against each other. An alternative approach to find the requirement priority is to quantitatively measure usage by introducing mock-ups to collect what users find interesting.

The paper does an assessment if it is valuable to include quantitative prioritization methods in the overall requirement engineering process during pre-development. The assessment is based on a case study conducted at Sony Mobile Communications Inc. We explore data collection techniques that allows for collection of quantitative data in the pre-development phase, i.e. before development of a feature starts. The exploration is done by assessing the possibility of pre-development quantitative data collection as outlined by a specific model for quantitative and qualitative data collection. Our research shows that (1) collecting quantitative feedback before development is feasible, (2) the data collected deviates from the original feature prioritization, i.e. it is beneficial, and (3) the data gives further insight in requirement prioritization than a qualitative method could have provided.

The remainder of this paper is organized as follows. The next section presents the background research for this paper. Section 3 describes the problem statements that can be addressed via quantitative data collection methods before the start of development. Section 4 describes the method used to assess the pre-development phase. Section 5 describes the case study and data collected. Section 6 contains final analysis and discussion of the result and possibilities for further work.

2 Background

Companies use a range of techniques to collect customer feedback in the early stages of product development. In the pre-development phase, techniques such as customer interviews, customer observation and customer surveys are typically used to get an understanding for customer perceptions of new product functionality [2][4][5][6][7][8]. Furthermore, mock-ups and different prototyping techniques are common to have customers try early versions of the product and for evaluating e.g. user interfaces.

Typically, these techniques provide a limited set of qualitative data reflecting individual customer needs [2]. In addition, there exist a number of techniques that can be used to validate customer needs during the development process, e.g. the HYPEX model [1]. Inspired by the 'Build-Measure-Learn' loop as outlined in the Lean Startup literature [9], they emphasize the need to build smaller increments of features that can be frequently validated by customers to avoid developing features that are not appreciated by customers. In a number of recent studies [1][10][11], the notion of frequent customer validation is described as 'innovation experiment systems' in which the R&D organization responds based on instant feedback from customers. This requires features and products with instrumentation for data collection so that product use can be continuously monitored. In this way, companies can learn about customer behaviors on a continuous basis and improve their products accordingly. In addition, and similar to common practice in the Web 2.0 and the SaaS domain [12], companies can run feature experiments, e.g. A/B or split testing, in which different customer groups receive different versions of the same feature and where data is collected to determine which version is the most appreciated one. In [11], this is referred to as the most efficient way to learn about customers. According to this author, the faster an organization learns about its customer and the ways in which they use the products, the more value it will provide. In [1] the process for feature experiments is further elaborated upon, and the authors present the HYPEX model, i.e. a process model for initiating, running and evaluating feature experiments. Similar to [1], [13] describe this as continuous experimentation and emphasize that customer experiment results need to be closely linked to feature prioritization and road mapping in order to support a more flexible business strategy.

As one attempt to capture the wide range of available customer feedback techniques, [2] present a model in which they identify different techniques, the type of data that is collected and the development phases in which the techniques are typically used. They picture the early development stages as characterized by direct customer feedback, and with small amounts of qualitative data being collected. In later stages, and after commercial deployment of the product, companies observe customers and use indirect feedback techniques to collect large sets of quantitative data. In ongoing work, and as a way to further detail qualitative and quantitative customer feedback techniques, [14] present the 'Qualitative/quantitative Customer-driven Development' (QCD) model. In this model, qualitative and quantitative customer feedback techniques are used to validate hypotheses derived from a backlog representing product concepts and ideas. The model suggests an approach in which requirements are treated as hypotheses that are continuously validated with customers, and only those that prove customer value are fully developed and deployed.

However, despite a wide range of available techniques, there are few techniques that help companies collect quantitative customer data already before investing in development of a feature. In what follows, we detail the problems that we encountered in previous research and in the interactions with companies in the software development domain, and we explore techniques that allow for the collection of qualitative data also before investing in the idea.

3 Problem Statement

Based on the research presented above, we have identified three problems often occurring in companies developing software-intensive systems. Below we describe each problem in more detail.

3.1 Release Content Cast in Stone

Most companies use a release model where the feature content for each release is decided upon before the start of development. Companies lack mechanisms to continuously validate feature content with customers, and find it difficult to re-prioritize pre-study outcomes. This causes companies to complete the building of features even if during development it becomes obvious that the feature clearly doesn't provide value to customers. This causes a sizeable part of the R&D resources to be allocated to wasteful activities and deteriorates the competitive position of companies over time.

3.2 Featuritis

There is evidence that a majority of features are seldom or never used and that customers seldom use the full potential of the functionality they receive [3]. Often referred to as "featuritis" [15], this means half or more of the R&D effort of a company is wasted on non-value adding activities. Similar to the previous problem, if competitors manage to have less waste in their R&D activities, over time the market position of the company is affected negatively.

3.3 Everything and the Kitchen Sink

Although often treated as atomic in research, features can be implemented iteratively and to a lesser or greater extent. Engineers often have a tendency to build features such that all use cases, exceptions and special situations are taken into account. Often, however, the value of a feature to customers is already accomplished after building a small part of the feature that provides the greatest value. Further development does not lead to (much) more value for customers. However, companies find it difficult to decide on when and how to stop building a feature when further iterations fail to add value to customers due to a lack of mechanisms for collecting feedback before, during development and after deployment of functionality [2].

In the remainder of the paper, we present a case study in which we evaluate two techniques that help companies validate customer value already in the pre-development phase before R&D investment has been made.

4 Research Method and Process

The research reported in this paper is based on a case study conducted at Sony Mobile. In our study, we focus on the data collection practices, and especially how to collect quantitative data already in the pre-development phase. To study this, and to evaluate two techniques that allow for this, we conducted case study research [17][18] based on interviews, workshop sessions, participant observations and active interventions in the case company where two of the authors are also employed. As a research method, case study research is typically used to contribute to our knowledge of individual, group, organizational and social phenomena, and is typically used to explain 'how' and 'why' and questions that require an extensive and in-depth understanding of the phenomenon [17][18]. Below, we describe how we selected an appropriate product and application for our study, and how to assess the data collection practices in the pre-development phase. As can be seen below, the first step is to select a product and application that meets the requirements to make the assessment possible. The second step is to identify what aspects of the pre-development phase the assessment should target. The final step involves how the assessments are performed and evaluated.

4.1 Choice of Product and Application

The choice of product landed on Sony Xperia™ phone and a specific Android application. This choice of product and application has been governed by compliance with three main requirements.

1. Large number of interactions with end consumers
2. Main assumption and statistics is that people using the app are first time users.
3. Possibility to change the feature set for selected users by showing a mock-up of a new feature set.

Requirement 1 targets the external validity [16][17] of the study. External validity refers to how well data and theories from one single case study can be generalized. Requirement 2 targets the internal [16][17] validity of the study. Internal validity refers to how well the case study is designed to avoid confounding factors. A confounding factor can be described as possible independent variable causing the effect, rather the variable concerned in the case study.

Requirement 3 targets the technical aspects of the android applications. The challenge was to create a mock-up of an existing application and replace the application in the already deployed product.

4.2 Assessment of Pre-development Phase

Two aspects of the pre-development phases are within the assessment scope of the case study. These two aspects are:

1. How well can we rely on pre-development quantitative data?
2. Is it possible to use pre-development quantitative data to give input to the hypothesis used for further experimentation?

By exploring the result from three data sets an assessment can be made of how valuable quantitative data collection in pre-development phases. One data set (FS-I) collects data from the real customer usage, the second data set (FS-II) is collected through a on-line survey on the web and the third data set (FS-III) is collected by mockup of the application presenting possible features considering the branding and name of the application. The criteria for giving a positive assessment result for aspect 1 is if the mockup (FS-III) and real application (FS-II) show similar pattern in customer usage for similar features. The criteria for giving a positive assessment result for aspect 2 is if new, edited, deleted hypothesis are elicited with use of the data collected (FS-II and FS-III).

5 The Sony Mobile Case

The case study has been performed at Sony Mobile Communications Inc. (Sony Mobile). Sony Mobile is a wholly owned subsidiary of Tokyo-based Sony Corporation, a leading global innovator of audio, video, game, communications, key device and information technology products for both the consumer and professional markets. Due to confidentiality reasons we cannot reveal the actual application. However, we can provide the following base data: The application was first deployed to live users in January 2013 and at present it is available in approximately 40M devices globally. With a few exceptions, the application is shipped with every mobile phone and tablet that Sony Mobile ships. Every month the application is in use approximately 3.5M times where every use on average involves three main use cases. The original application provides 4 + 1 features (see description of feature set one (FS-I) below) where each feature requires two to five interactions in the normal case and based on collected data we know that the application has approximately 9,5M interactions per month. The use of the application is consistent over the year and does not show e.g. seasonal variations or variations due to product releases. The application is considered to be one of the base line applications delivered with Sony Mobile products.

5.1 Application Features

In order to not break confidentiality we have coded the available and possible features into feature types (FT) and enumerated these from FT01 to FT12. Of particular interest for this study are the “no feature” id est users that are not using the application for any particular reason, rather just exploring the application. This feature was added to Feature Set II and Feature Set III. The “nothing feature” is enumerated as FT13.

Throughout this study we have used three methods to capture what the user really wants to use; actual usage, survey and mock-up. Due to limitations, technical possibilities and semantic limitation in the three methods all features could not be made available in all three methods. Instead a selection had to be made. An explanation of what features were selected in respective feature set follows.

5.1.1 Feature Set I – (FS-I)

The first feature set to be evaluated is what was actually deployed in the first version of the application. The application itself contains a tracking mechanism that collects usage data into Google analytics. Hence the data was readily available for our study.

During the measured time frame the application was used 3.652.796 times. The application did at that time present four main features (FT01 - FT04) as well as a less prominently displayed feature (FT05). During that period the features were used a total of 2.618.513 times. All other usage of the application a total of 1.034.283 are considered to be in FT13 group (no usage/nothing group)

Expressed as percentage of usage the distribution of the six identified feature types looks like in Table 1 below:

Table 1. Response distribution FS-I

Feature	Relative Usage
FT01	28,56%
FT02	21,53%
FT03	11,49%
FT04	6,79%
FT05	3,31%
FT13	28,31%
Total	100,00%

5.1.2 Feature Set II – (FS-II)

As we were looking to extend the feature set for the product, but wanted to collect quantitative data on the customer needs we employed two techniques. The first, a survey, is discussed in this section. The second, a mockup, is discussed in the next section.

The second feature set to evaluate is what users answers when prompted in a survey. Users visiting a specific site were randomly selected to answer a survey about their reason for using the application. A total of seven (7) different features where selectable and only one answer could be given. The question was formulated as "What is the primary feature you are looking for", thereby forcing the user to give a distinct answer even if the user had several reasons for using the application.

The selectable features were basically the same as in the original application however the least prominent feature (FT05) was not included and two more features (FT06 and FT07) were possible to select. However FT01 through FT04 were selected to be same as well as the nothing feature (FT13). This slight separation between features was necessary in order to be able to present the various features in a sensible manner and not making it obvious to the user that additional data was collected so that we could get a clear user aware set-up. A total of 119.370 survey answers were collected during one month. The answer relative distribution can be seen in Table 2 below.

Table 2. Response distribution FS-II

Feature	Relative Usage
FT01	24,39%
FT02	33,55%
FT03	11,62%
FT04	13,09%
FT06	1,77%
FT07	2,50%
FT13	13,07%
Total	100,00%

5.1.3 Feature Set III – (FS-III)

Based on the two first feature sets, basically collected from existing data and/or without modifying the actual application it was seen that additional features could be possible candidates for further development. In line with the HYPEX model [1], we designed a new version of the application as a mockup that would allow for serving a very large number of features. Imperative was that the new design should be the minimal (least expensive/requiring the least effort) viable feature (MVF). In order to find possible features, in addition to the existing (FS-I) and the ones found when doing the survey (FS-II) a third feature set FS-III was constructed. To find additional possible features, similar applications from other manufacturers were surveyed. The features FT08 through FT12 were added. FT06 had to be excluded from the application due to its nature and in the context of the application it was too similar to FT07. In the choice between FT06 and FT07 it was judged that FT07 was a better match.

The application was modified to present the newly designed feature set and in order to keep the development and implementation as minimal as possible but still viable for use the underlying functionality was strictly limited, in fact so much that it just barely resembled the promised feature. To mitigate the risk of causing damage to the product extra work was put into making sure that any given user would be exposed to the new design only once.

As we presented a much larger feature set in the mockup and rather than implementing each feature, by and large linked to existing features or served rather simple screens, we view this as a pre-development activity as the features are not actually developed. Rather the mockup was “inserted” in front of the normal application and linked to existing functionality where possible.

The new design of the application forced the user to make a selection upfront leaving out all other elements of the application at this stage. The order in which the features were presented was randomized in order to avoid the risk of the user selecting e.g. the first or last feature.

The prototype was launched into production for a period of 10 days, during this time we collected 34.393 interactions. The relative distribution of answers is shown in table 3 below.

Table 3. Usage distribution FS-III

Feature	Relative Usage
FT01	27,08%
FT02	9,68%
FT03	9,79%
FT04	4,20%
FT05	4,57%
FT07	6,57%
FT08	8,90%
FT09	8,42%
FT10	3,63%
FT11	3,43%
FT12	2,26%
FT13	11,46%
Total	100,00%

5.2 Feature Set Usage Description

In the previous section, we introduced the features that were used for the baseline activity, the survey approach and the mockup approach. In this section, we present in more detail the usage of features for each approach and indicate relevant aspects of each step in our research process.

5.2.1 Feature Set – I (FS-I)

Looking at FS-I, with its limited number of features FT01, FT02 and FT13 together amounts for almost 80%, noting as well that the nothing feature (FT13) is almost as large as the real feature FT01, 28,31% for FT13 as compared to 28,56% for FT01. FT05 in FS-I is the feature that has a less prominent display in the application that could explain its very low usage (3,31%).

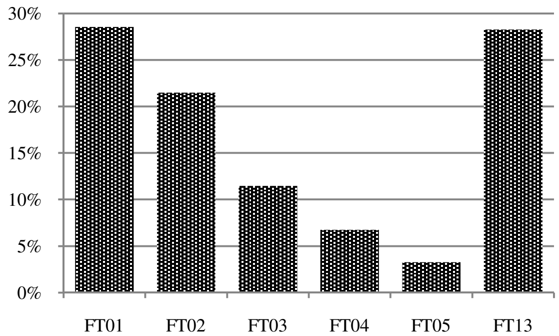

Fig. 1. FS-I Relative distribution of feature usage

5.2.2 Feature Set – II (FS-II)

In FS-II where the user was asked to actively participate in a survey the distribution of answers can be illustrated in figure 2.When asked the users appear to favor FT02 though the difference to FT01 is not that large. Looking at the graph it may appear such that the features group (FT01 and FT02, FT03 and FT04, FT06 and FT07), however looking at the underlying features no such grouping makes sense and this grouping is random rather than a true correlation. To our surprise FT013 stands out in this feature set as well: even in a survey, users of the app were willing to indicate that they were not looking for anything in particular.

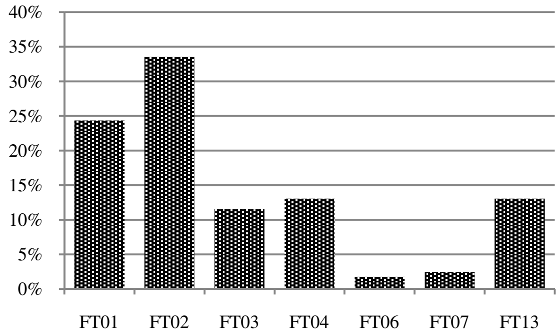

Fig. 2. FS-II relative distribution of feature type usage

5.2.3 Feature Set – III (FS-III)

Using the mock-up application, the large number of available features seems to have leveled out the choice users make, though FT01 as well as FT13 stands out. The relative usage is depicted in Fig 3 below.

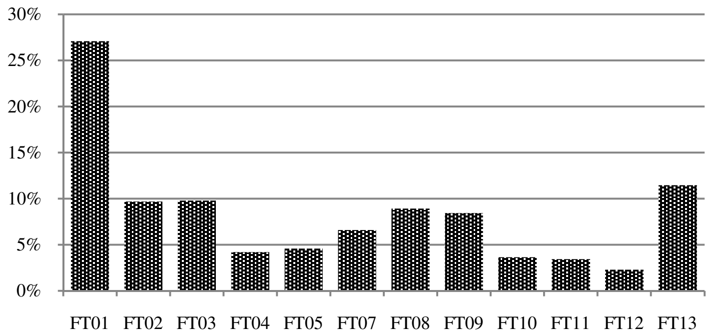

Fig. 3. FS-III relative distribution of feature usage

5.3 Feature Comparison Between Collection Methods

Looking at the differences between the three methods it can be clearly seen that the results match well between the various methods. The overall distribution is depicted in Fig. 4. below. In all three collection methods, FT01 and FT13 have high usage. An obvious conclusion from this is that carefully selecting the feature to put first as well as designing a good user experience for users that arrive at the app without a clear goal in mind is particularly important. When a wider selection is available as in FS-III the user preference for the other features is more nuanced and requires more careful interpretation of the data. For instance, the selection frequency for FT02 and FT04 is very different for the survey and the mockup. As the survey asks customers explicitly for a choice, the technique measures what users say they want. The mockup, on the other hand, is concerned with measuring what users will actually do in practice. The latter is obviously more relevant but understanding why the gap between espoused and enacted behavior in users exists is important for product management and R&D to understand.

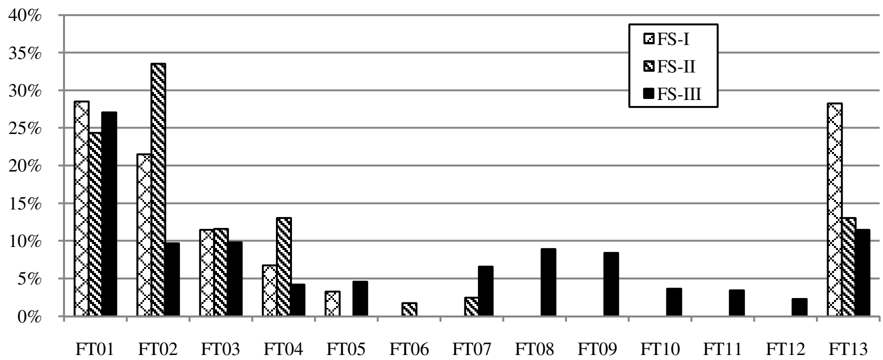

Fig. 4. Comparison of the feature types between the feature sets

6 Analysis and Discussion

6.1 Assessment of Aspect 1

The data collected from all set can be visualized by ordering the feature in descending order concerning the usage percentage. Considering the similarity in patterns especially with features with high usage is that the conclusion is that pre-development collection methods are valid for especially for the features with high usage.

For the different sets (FS-I, FS-II, FS-III) the lists are the following:

i. :{FT01, FT13, FT02, FT03, FT05, FT04}

ii. :{FT02, FT01, FT04, FT13, FT03, FT07, FT06}

iii. :{FT01, FT13, FT03, FT02, FT08, FT09, FT07, FT06, FT05, FT04, FT10, FT11, FT12}

Table 4. Feature usage in usage order

FS-I	FS-II	FS-III
FT01	FT02	FT01
FT13	FT01	FT13
FT02	FT04	FT03
FT03	FT13	FT02
FT04	FT03	FT08
FT05	FT07	FT09
	FT06	FT07
		FT05
		FT04
		FT10
		FT11
		FT12

6.2 Assessment of Aspect 2

Interviews have been held with senior staff members at Sony Mobile. The scope of the interviews has been to study and compare the different data set to evaluate if hypothesis can been elicited. These hypotheses would be candidates to populate new, edit or delete hypotheses to be used as experiments in the product development. The result of the interviews is that it was possible to elicit hypothesis. The following hypotheses were elicited. FT07, FT08, FT09 are not available in existing application (FS-III) and have a relative high usage percentage in the mockup. Therefore they are considered valid entries to the hypothesis backlog as additional features. FT10 has a low usage rate in the mockup and the hypothesis is that there are a number of better features to invest in. FT13 is a special case since it defines a usage where the customer opens the application but does not take any actions. A hypothesis is that there is an opportunity to capture this kind of users and turn them to active users within the application.

6.3 Discussion

In this paper, we present a case study conducted at Sony Mobile in which we explore the feasibility to collect quantitative customer data also before development starts, i.e. in the pre-development phase. While this has proven difficult, we evaluate two techniques that allow for assessment of whether a feature is worthwhile developing or whether R&D resources should be allocated somewhere else. Our case study shows that collection of quantitative data in the pre-development phase is both feasible and useful in the feature prioritization process, and that this data is also reliable. In addition, the combination of techniques where users are aware that they are asked for input (espoused behavior) [19], and techniques where enacted user behavior is captured provides for valuable insight in to the difference between the two. This allows for further hypothesis development to explain the gap between espoused and enacted behavior.

In our study, the frequent use of the "do nothing" feature indicates that there are additional opportunities to engage users than the basic "select a feature" functionality. As the user opened the app without a clear goal in mind, the app can propose and nudge users to engage with the app or other Sony Mobile applications, as the user almost seems to expect the app to take the initiative. These types of insights, i.e. those user behaviors that are not necessarily captured by qualitative data collection techniques, are extremely valuable for product managers and R&D teams. Also, if these can be captured already before development starts, as shown in our study, product managers and R&D teams can act pro-actively and, if needed, re-prioritize feature content.

Finally, there are a number of opportunities for further study that we are considering, especially concerning the empirical design. For example, expanding the analysis to more applications, and other types of applications, making the feature sets in pre-development more comparable by using same number of features. Also, we plan to assess other pre-development quantitative methods.

Acknowledgments. The development work needed for of the case study has been funded and supported by Sony Mobile. This research was in part funded by Software Center (www.software-center.se).

References

1. Olsson, H.H., and Bosch, J.: From opinions to data-driven software R&D: a multi-case study on how to close the `Open Loop' problem. In: Proceedings of EUROMICRO, Software Engineering and Advanced Applications (SEAA), Verona, Italy, 27-29 August 2014
2. Bosch-Sijtsema, P., Bosch, J.: User Involvement throughout the Innovation Process in High-Tech Industries. Journal of Product Innovation Management. (2014). doi:10.1111/jpim.12233
3. Backlund, E., Bolle, M., Tichy, M., Olsson, H.H., Bosch, J.: Automated user interaction analysis for workflow-based web portals. In: Lassenius, C., Smolander, K. (eds.) ICSOB 2014. LNBIP, vol. 182, pp. 148–162. Springer, Heidelberg (2014)
4. Hofman, H.F., Lehner, F.: Requirements engineering as a success factor in software projects. IEEE Software **18**, 58–66 (2001)
5. Kabbedijk, J.; Brinkkemper, S.; Jansen, S.; van der Veldt, B.: Customer Involvement in Requirements Management: Lessons from Mass Market Software Development, Requirements Engineering Conference (2009)
6. Yiyi, Y., Rongqiu, C.: Customer Participation: Co-Creating Knowledge with Customers-Wireless Communications, Networking and Mobile Computing (2008)
7. Olsson H.H., Bosch, J.: 2013a. Towards data-driven product development: a multiplecase study on post-deployment data usage in software-intensive embedded systems. In: Proceedings of the Lean Enterprise Software and Systems Conference (LESS), Galway, Ireland, 1-4 December 2013

8. Holmström Olsson, H., Bosch, J.: Post-deployment data collection in software-intensive embedded products. In: Herzwurm, G., Margaria, T. (eds.) ICSOB 2013. LNBIP, vol. 150, pp. 79–89. Springer, Heidelberg (2013)
9. Ries, E.: The Lean Startup: How Constant Innovation Creates Radically Successful Businesses. Penguin Group, London (2011)
10. Olsson, H.H., Alahyari, H., Bosch, J.: Climbing the "Stairway to Heaven": a multiple-case study exploring barriers in the transition from agile development towards continuous deployment of software. Proceedings of the 38th Euromicro Conference on Software Engineering and Advanced Applications, September, pp. 5–7. Cesme, Izmir (2012)
11. Bosch, J.: Building products as innovations experiment systems. Proceedings of 3rd International Conference on Software Business, pp. 18–20. Massachusetts, Cambridge (2012)
12. Kohavi, R., Crook, T., Longbotham, R.: Online experimentation at microsoft. In: Third Workshop on Data Mining Case Studies and Practice Prize (2009)
13. Fagerholm, F., Sanchez, G., Mäenpää, H., Münch, J.: Building blocks for continuous experimentation. In: The Proceedings of the RCoSE 2014 Workshop, Hyderabad, India, 3 June 2014
14. Olsson, H.H., Bosch, J.: The QCD Model: From Requirements To Continuous Reprioritization Of Hypotheses. Submitted to an international software engineering conference (Submitted)
15. Wikipedia. http://en.wikipedia.org/wiki/Feature_creep
16. Campbell, D.T., Stanley, J.C.: Experimental and quasi-experimental designs for research, p. 1983. Houghton Mifflin Company. Clark, R. E, Boston (1966)
17. Robson, C.: Real World Research, 2nd edn. Blackwell, Oxford (2002)
18. Runesson, P., Höst, M.: Guidelines for conducting and reporting case study research in software engineering, Empirical Software Engineering, vol, 14 (2009)
19. Argyris, C., Schön, Donald, A.: Organizational learning II. Reading, MA: Addison-Wesley (1996)