



The unfulfilled potential of data-driven decision making in agile software development

Downloaded from: <https://research.chalmers.se>, 2024-04-19 07:42 UTC

Citation for the original published paper (version of record):

Berntsson Svensson, R., Feldt, R., Torkar, R. (2019). The unfulfilled potential of data-driven decision making in agile software development. *Lecture Notes in Business Information Processing*, 355: 69-85. http://dx.doi.org/10.1007/978-3-030-19034-7_5

N.B. When citing this work, cite the original published paper.



The Unfulfilled Potential of Data-Driven Decision Making in Agile Software Development

Richard Berntsson Svensson^(✉), Robert Feldt, and Richard Torkar

Department of Computer Science and Engineering,
Chalmers|University of Gothenburg, Gothenburg, Sweden
`richard.berntsson.svensson@gu.se`, `robert.feldt@chalmers.se`,
`richard.torkar@cse.gu.se`

Abstract. With the general trend towards data-driven decision making (DDDM), organizations are looking for ways to use DDDM to improve their decisions. However, few studies have looked into the practitioners view of DDDM, in particular for agile organizations. In this paper we investigated the experiences of using DDDM, and how data can improve decision making. An emailed questionnaire was sent out to 124 industry practitioners in agile software developing companies, of which 84 answered. The results show that few practitioners indicated a widespread use of DDDM in their current decision making practices. The practitioners were more positive to its future use for higher-level and more general decision making, fairly positive to its use for requirements elicitation and prioritization decisions, while being less positive to its future use at the team level. The practitioners do see a lot of potential for DDDM in an agile context; however, currently unfulfilled.

Keywords: Data-Driven Decision Making · Agile · Survey

1 Introduction

When developing software-intensive products, agile methods have become the *de facto* way to develop software across almost every industry. The introduction of agile methodologies has changed the way software is developed [1], how Requirements Engineering (RE) is conducted [2], and how decisions are made [3]. In transitioning to Agile Software Development (ASD), learning about the customers, collecting customer/user feedback, and involving a customer representative in development, requirements engineering, and decision making, are important [4]. In addition, ASD teams, due to delivering working software in short iterations, are frequently involved in short-term decisions and need to adopt to a fast decision making process [5].

With digital networks connecting an increasing number of people, devices, and products, a vast amount of diverse data is available. Industries gather data

and knowledge from their customers, suppliers, alliance partners, and competitors. For example, mobile phones, cars, transportation vehicles, and automation systems, are developed to generate data about their customers and usage of their activities. This diverse data is not only generated internally within software-intensive companies, but also from public, proprietary, and purchased sources [6]. Software developing companies need to focus on exploiting the available data to gain competitive advantages [6], which will transform how business are generated, how RE is performed, and how decisions are made [7]. In particular, the recent resurgence of interest in artificial intelligence (AI) and machine learning (ML) accelerates these trends due to their promise of more automated and powerful data analysis.

However, despite the vast amount of data that is available for decision making, the decisions and selection of what to include in the next product release cycle, are commonly based on the product managements and/or stakeholders' previous experiences, opinions, intuitions, various criteria, arguments, or a combination of one or several of these information sources [4, 7]. These decisions are typically subjective, frequently inconsistent, and often lack explanations as well as links to which data and evidence they were based on. Moreover, when stakeholders make decisions based on, e.g., opinions, intuitions, and arguments, the decisions are more likely to be influenced by politics and individual agendas [8–10] rather than, e.g., business opportunities or customer value. In addition, even when data is more clearly being taken into account in decisions, too much data and information may distract the decision maker rather than inform them. According to Wnuk et al. [12], irrelevant information is visible in practitioner backlogs to a large extent today, and recent research shows that it can negatively impact decisions [13].

In order to benefit from data-driven decision making (DDDM), not only is the quality of the processing techniques and tools directly related to the quality of the decisions [17], but also the quality of the visualizations used to support decision makers [17]. While visualization of software engineering data has shown promise in supporting practitioners' decisions, the focus has often been on specific phases or problems, e.g., testing and quality assurance [11], rather than throughout development processes and in agile settings. In the literature, most of the attention in DDDM has focused on the development of new techniques, technologies, and tools for data processing [14], while few (if any) have investigated DDDM from the practitioners' perspectives and the specific and important context of agile development has not been in focus.

This paper presents the results of an empirical study that includes data collected through an emailed questionnaire with 84 respondents from 28 agile software developing companies from 9 domains. The study investigate how common the use of data for decision making is in industry today, how often data is used, the respondents opinions about the usage of data in the future, and how data can improve decision making.

The remainder of this paper is organized as follows. In Sect. 2, we outline the background to data-driven decision making. Section 3 describes the research

methodology, while Sect. 4 presents an overall statistical analysis of the data. Section 5 presents and discuss the results, and finally Sect. 6 presents the conclusions.

2 Background

Data-driven decision making (DDDM) has become a critical ability for organizational success. Several studies have demonstrated the benefits of DDDM, e.g., Brynjolfsson et al. [16] showed that DDDM is strongly related to higher productivity, higher return on assets, return on equity, and market value.

In the literature, there are several defined steps in DDDM, starting with data capturing and resulting in decision making. For example, Chen and Zhang [14] identify five steps; data recording, data cleaning/integration/representation, data analysis, data visualization/interpretation, and decision making. Although steps are identified, most of the attention in the literature has focused on the development of new techniques, technologies, and tools. Techniques for DDDM involve a number of disciplines with a number of specific techniques and tools in each discipline. For example, fundamental mathematics, statistics, and optimization tools are used as input to data analysis techniques such as data mining, machine learning, neural networks, signal processing, and visualization methods [14]. Current DDDM tools can be divided into three categories: batch processing tools, stream processing tools, and interactive analysis tools [14]. For more details about different techniques, technologies, and tools, we refer to [14]. We also see an increased interest in applying AI and machine learning in a software engineering context [15] and supporting decisions during development is one of the key application types.

The quality of the decisions when using DDDM may improve or degrade based on the quality of the data and the processing techniques and tools [17]. However, the quality of the decisions are not only based on pre-processing techniques, processing techniques and tools, it is also related to the quality of the visualizations of the data to the decision makers, the decision makers' understanding and knowledge about the data sources, the decision makers' ability to interpret data processed data, and the decision makers' knowledge about the relationships of the data [17]. As one example, Feldt et al. [11] showed how visualisation of testing-related data, without any advanced modeling, could foster understanding and support decisions around software quality in an iterative development context. Thus, in order to benefit from DDDM, it is important to focus also on other aspects than just the pre-processing and processing techniques, technologies, and tools.

3 Research Method

The objective of this study was to investigate how common the use of data for decision making is in industry today, how often data is used, and the respondents' opinions about the usage of data in the future, with a special focus on the agile

context in which modern-day software is developed. Given the objective, and that the research questions are geared towards the opinions of the respondents, we chose to use a survey as the research method and emailed a questionnaire for data collection. Surveys are an appropriate strategy for getting empirical descriptions about trends, attitude and/or opinions of the studied population [18, 19]. In addition, surveys are useful for analyzing large populations, given an adequate response rate [20, 21]. The motivation for using an emailed questionnaire was to maximize coverage and participation. The following research questions provided the focus for the empirical investigation:

- **RQ1:** How do software practitioners view data as part of decision making in agile software developing companies?
- **RQ2:** To what extent is data used for decision making and requirements engineering in agile software developing companies?
- **RQ3:** How can data be used to improve future decisions in agile software developing companies?

3.1 Survey Study

The survey was executed through the creation of an emailed questionnaire that was designed based on the research questions using a mix of open-ended and closed questions [19]. In order to test the reliability and validity of the survey instrument, a pilot study was conducted with one industry practitioner. Based on the feedback from the pilot study, the survey instrument was (lightly) revised. The instrument (see Table 1) had three parts. The first part gathered demographic information about the respondents. The second part mainly addressed how, and how often data is used in decision making today, while the third part focused mainly on how data can be used for decision making in the future. Part 1 only contained free-text questions. All of the questions in Parts 2 and 3 contained Likert-type scale and free-text questions. The free-text area was added to allow the respondents to expand and/or explain their answer.

Data Collection. Subjects were sampled primarily through personal contacts and previous collaborators in industry and we encouraged them to also spread the survey within their organisations. Hence, the sample can be described as convenience sampling [19]. We provided the contacts with the questionnaire (emailed questionnaires) and information about the goals of the survey, and asked them to answer the questions and to spread the questionnaire to their colleagues. Each contact person reported back how many people they had forwarded the questionnaire to. A total of 124 subjects received the questionnaire, and 84 completed the mandatory questions and returned the questionnaire to the researchers. That is, we obtained a response rate of 67.7%. Without going through personal contacts in industry we likely would not have been able to get this high a response rate.

Data Analysis. The data was analyzed using descriptive statistics with diverging stacked bar charts for the graphical visualization. In addition, we built a linear

Table 1. Survey instrument

ID	Question
Q0	What company do you work for?, How many employees does your company have?, What role do you generally have in your work?, What software development process do you use?
Q1	Data is important for decision-making
Q2	Data is highly valued for decision-making
Q3	Data is treated as an asset
Q4	Data is used to identify new business opportunities
Q5	Data is used to predict future trends and behavior
Q6	Decision makers use data for decision-making
Q7	Teams use data for decision-making
Q8	Data is used as part of requirements elicitation/identification
Q9	Data is used for prioritization of requirements/features
Q10	Data should be important for decision-making
Q11	Data should be highly valued for decision-making
Q12	Data should be treated as an asset
Q13	Data should be used to identify new business opportunities
Q14	Data should be used to predict future trends and behavior
Q15	Decision makers should use data for decision-making
Q16	Teams should use data for decision-making
Q17	Data should be used as part of requirements elicitation/identification
Q18	Data should be used for prioritization of requirements/features

model (ordered logit) using a Bayesian approach [22,23] to statistically analyse the data. The analysis is described in more detail in Sect. 4.

3.2 Validity Threats

To avoid evaluation apprehension (*construct validity*) [24], we guaranteed the respondents complete anonymity. Another threat is ‘hypothesis guessing’ [24], which was minimized by clearly expressing the need for honesty in the instructions to the respondents; however, it is not possible to completely dismiss this threat. In addition, the background of the subjects, e.g., experience, may influence the results; however, since the respondents have different competences and roles we believe that this risk is limited. It is not possible to exclude the possibility that the respondents misunderstood the questions (*conclusion validity*) [24]. To minimize this threat, we conducted a pilot study with an industry practitioner, which also minimized the threat of instrumentation (*internal validity*) [24]. One threat that cannot be ignored is the interest of the respondents in the topic, which may influence the representativeness. This is difficult to counter

since the willingness to participate and the interest in the topic may be linked. There are also threats to validity based on selection bias and the convenience sampling; even though we sent to most of our contacts in agile software organisations and approached them in a standardised way, the final sample might not be representative for a global population of developers. For example, they were all from organisations in Sweden.

4 Analysis

To plot and assess visually the difference between distributions of responses in Likert scale data is hard. As an example, if we examine Fig. 1, we see that there is a difference between the distribution of answers on two questions (Q16, on top in the figure, and Q17, on bottom) but it is not clear how to judge how large the difference is. Also, if we only use descriptive statistics, which is the default analysis technique for survey data in software engineering, it is difficult to assess the uncertainty of our conclusions. In contrast, a Bayesian statistical analysis does not have the same problem. Thus, in line with recent arguments for use of Bayesian methods in empirical software engineering we thus, first, start with such an analysis [25, 26].




Fig. 1. Plots of Q16 (top) and Q17 (bottom). Left: Histogram of discrete response in the sample. Middle: Cumulative proportion of each response. Right: Logarithm of cumulative odds of each response. Note that the log-cumulative-odds of Level 5 is infinity if there are responses among all five levels $[1, \dots, 5]$, as in Q17 (for Q16 there were no responses on Level 5).

In order to assess differences in Likert scale data one could assume normality and use a t -test, or make use of some of the non-parametric tests such as Mann-Whitney U or χ^2 .

However, Likert scale data is not only categorical, it is also of an ordered nature but where we cannot assume that the ‘distance’ between consecutive pairs of answers is the same. Thus it is not clear that we can assume the data is normally distributed or that the distribution of scores for different answers has the same shape (distribution family) [27]. Given these problems, in our view, the most conservative approach to analyze Likert scale data is to build a simple linear model using a Bayesian approach but keeping data categorical [22, 23]. This way we will get a posterior distribution with which we can assess uncertainty. To this end we build two overall models to study the general trends in our data:

$$\begin{aligned} R_i &\sim \text{Ordered}(p_i) \\ \text{logit}(p_i) &= \beta_T * \text{temporal}_i \\ \beta_T &\sim \mathcal{N}(0, 10) \end{aligned} \tag{1}$$

$$\begin{aligned} R_i &\sim \text{Ordered}(p_i) \\ \text{logit}(p_i) &= \beta_Q * \text{req}_i \\ \beta_Q &\sim \mathcal{N}(0, 10) \end{aligned} \tag{2}$$

where R_i is the i th response with an ordered categorical outcome, and Model 1 (Eq. 1) compares the answers for questions about the present (Questions 1–9, see Fig. 4) versus future (Questions 10–18, see Fig. 4) use while Model 2 (Eq. 2) compares the non-RE (Questions 13–16, see Fig. 5) versus the RE-specific (Questions 17–18, see Fig. 5) questions. We use the logit link function to translate the linear model’s real numbers to probability mass (and hence constrain it to lie between zero and one). The linear model (in Eq. 1) then is simply a parameter β_T that we will estimate given the data at hand (temporal). The data is coded as 0/1, representing ‘present’ (today) and ‘future’, respectively. Finally, we assign a prior to β_T , $\mathcal{N}(0, 10)$, with a mean of 0 and a large variance of 10. This is a (very) weakly informative prior that only gives a pressure towards realistic parameter values. We also verified that the analysis was not sensitive to the prior selection (i.e., a sensitivity analysis was conducted).

For the other model (Eq. 2) we simply change the parameter. Instead of estimating β_T using ‘temporal’ data, we estimate β_Q for our variable ‘question’, which is coded 0/1, representing question with a ‘non-RE’ (Q13–16) and ‘RE’ focus (Q17–18), respectively.¹

Figure 2 visualizes the results from running the first model and drawing 250 samples from the posterior distribution. It is obvious that low Likert scale values are much more common for the ‘present’ compared to the ‘future’ category. For example, we see that the number of answers of option 1 (‘Strongly disagree’) is roughly around 70% for questions about the present (today) state but decreases down to only 5% for the future state. We can also see that the uncertainty is not large with variations only in the range of 1–7% for all the answer alternatives.

¹ The models overall sampled well with mixed chains, $\hat{R} \ll 1.1$, and an effective sample size of $n_{\text{eff}} \gg 0.2$.

When comparing non-RE and RE questions using Model 2 in Fig. 3, we can also see some trends even if they are less clear and the uncertainty is higher as visualized by the, relatively speaking, broader bands of posterior predictions. However, the model clearly shows that we see a difference between non-RE and RE related questions with the average of the β_Q , being $\mu = -0.53$ HPDI_{95%} $[-0.87, -0.19]$, i.e., the 95% highest posterior density estimate (HPDI) does not cross 0. This indicates that answers to the RE questions are generally lower (i.e. towards more disagreement with the statement in the question) than for the non-RE ones and that this difference is clear.

After this detailed, statistical analysis of the general trends in the responses the following Section will discuss the results in more detail.




Fig. 2. Posterior predictions (250 draws) of the ordered categorical model (present vs. future perspective). As is clearly evident, the probability for lower Likert scale values, e.g., 1 or 2, is much higher when the perspective is ‘present’, compared to ‘future’, i.e., everything is shifted upwards. This indicates less agreement at present and more agreement for the future, i.e. there is unfulfilled potential since the present state has a higher percentage of low disagreement answers.

5 Results and Discussion

This section presents the results of the survey, organized according to the research questions in Sect. 3.

5.1 Survey Respondent Demographics

A total of 84 industry practitioners completed the questions of the survey. The respondents come from 28 agile software developing companies varying in size and domain. In total, the respondents came from nine different domains, with the top three being Telecommunication (27%), Consulting (18%), and Transportation (13%), see Table 2. The size of the companies where the respondents




Fig. 3. Posterior predictions (250 draws) of the ordered categorical model (non-RE vs. RE perspective). We see some trends, i.e., respondents are more positive in non-RE focused questions, but there is quite much uncertainty visualized here by the broader clusters of lines. However, the parameter β_Q , which represents the domain focus, indicates that the trend is non-negligible ($\mu = -0.53$ HPDI_{95%} $[-0.87, -0.19]$).

work, in terms of number of employees, ranges from 25 up to 5,000. With respect to the respondents' roles, see Table 3, the top three are developers (17%), scrum masters (15%), and product owners (14%) with a fairly even distribution of other, common roles also represented. For the development processes used at the companies see Table 4 where Scrum (43%) is the most used, followed by (the general option) Agile (29%), Kanban (15%), and then DevOps (12%). Note that the Agile category means that a respondent did not specify which agile methodology they used. Overall, we consider these respondents representative for a broad set of domains, roles and sizes of companies, even if they are all active in a Swedish context. The one role that is less clearly represented is Requirements Engineer although several of the respondents also partly do work with requirements in one form or another, as is common in agile development.

5.2 View of Data in Decision Making (RQ1)

In analyzing Research Question 1 (RQ1), this section examines the respondents' view of data as part of decision making in ASD companies. In Fig. 4, we can see the respondents' answers to each question. Each row shows the distribution of answers for that question aligned horizontally so that positive responses are to the right of the mid (zero) line while negative responses are to the left.² This makes it possible to compare the answers between different questions.

In general, looking at Fig. 4, we can see that it follows the general trend identified in the statistical analysis above, i.e., respondents disagreed with the

² Note that the neutral, mid answer option (on the 5-category Likert scale) is split in half, with half of them shown in a lighter (gray) color to the left and the other half in darker (gray) color to the right of the mid (zero) line.

Table 2. Distribution of respondents based on domains

Domain	Respondents
Telecommunication	27%
Consulting	18%
Transportation	13%
Consumer electronics	11%
Surveillance	10%
Control systems	8%
Retail	5%
Camera	5%
Banking	4%

Table 3. Distribution of respondents based on roles

Roles	Respondents
Developer	17%
Scrum master	15%
Product owner	14%
Project manager	11%
Tester	11%
Senior software engineer	11%
Product manager	10%
Architect	6%
Requirements engineer	6%

Table 4. Distribution of respondents based on development process

Development process	Respondents
Scrum	43%
Agile	29%
Kanban	15%
DevOps	12%
XP	1%




Fig. 4. Respondents' view of data as part of decision making. Present (Q1–Q3) and future (Q10–Q12).

statements more in questions about the current state while agreeing more in questions about the future. For example, we see that a majority of the respondents disagreed or strongly disagreed that data is important (66% for Q1) and highly valued (79% for Q2) in today's decision making. However, a majority of

the respondents agreed or strongly agreed that data should play an important role (71% for Q10) and be highly valued (87% for Q11), when making decisions in the future. Examining if data is treated as an asset today (Q3), 93% of the respondents disagreed or strongly disagreed, while 63% of the respondents agreed or strongly agreed that data should be treated as an asset in the future (Q12). Although the respondents have a positive view of how data should ideally be viewed for decision making, their answers indicate this is not how it is being viewed at present in their organisations.

5.3 Use of Data in Decision Making (RQ2)

In analyzing Research Question 2 (RQ2), this section examines to what extent data is used (present) and should be used (future) in decision making and requirements engineering in ASD companies, as illustrated in Fig. 5. Figure 5 is constructed in the same way as Fig. 4, with the exception that the zero line, i.e., the neutral answer, is set to the answer ‘About half of the time’. In general, Fig. 5 shows that data is seldom (never or sometimes) used in today’s decision making or in Requirements Engineering (RE) (Q4–Q9 in Fig. 5). However, a vast majority of the respondents believe that data should be used most of the time or always in future decision making and RE (Q13–Q18 in Fig. 5).




Fig. 5. Use of data as part of decision making. Present (Q4–Q9) and future (Q13–Q18)

Looking closely into what extent data is used in today’s decision making, for all questions (Q4–Q9), more than 90% of the respondents stated that they never or only sometimes use data in decision making and RE, where more than 73% of the respondents stated that they never use data today. No respondent stated that they always use data. Only 1% of the respondents stated that they use data most of the times for requirements elicitation/identification (Q8) and requirements prioritization (Q9). Instead of using data, the respondents explained in the free-text answer that decisions are mainly based on ‘gut-feeling’, the decision-makers’ experiences, or the value for customers.

That is, the decisions may be subjective [7], politically influenced [8], and/or biases could be involved [13]. Instead of using data when prioritizing requirements, respondents detailed that requirements are prioritized using various criteria (e.g., cost, cost/benefit, customer value, business value), numerical assignment, experiences, ‘gut-feeling’, or a combination of these. This is inline with other studies on how requirements are prioritized in ASD companies today [28].

When asking the respondents to what extent data should be used in decision making in the future, 93% of the respondents believe that decision makers should always, or most of the time use data for decision making (Q15), 85% believe that data should always, or most of the time be used to identify new business opportunities (Q13), and almost 75% believe that data should always, or most of the time be used to predict future trends and behaviours (Q14). Only 8% of the respondents believe that (agile) teams should always, or most of the time use data for decision making (Q16), while almost half of the respondents (43%) believe the (agile) teams should never, or only sometimes use data when making decisions. No explanation was provided by the respondents in the free-text answers for these questions.

One possible explanation may be that the respondents believe that DDDM is only useful and beneficial for high-level decisions. This is supported by the high confidence in using DDDM for identifying business opportunities (Q13) and to predict future trends and behaviours (Q14). When such high-level decisions are made, including creating product strategies, road-maps, and release plans, the respondents may believe that teams do not need DDDM when, e.g., breaking down high-level requirements to low-level ones. Another explanation may be related to today’s development processes and short sprints, which may not be well suited for DDDM at the team level.

To create and rapidly release software-intensive products in the future, it is crucial that the products are based on data and real-time feedback from the customers [7]. Thus, when moving from a subjective decision-making process, mainly based on experiences, to a DDDM process, changes in infrastructure and methodologies are needed in the development processes [7].

For RE, 60% of the respondents believe data should always, or most of the times be used when eliciting/identifying requirements in the future (Q17), while 15% believe data should never, or only sometimes be used for requirements elicitation/identification. Only 35% of the respondents believe data should always, or most of the time be used when prioritizing requirements, 25% believe it should never, or only sometimes be used, while as many as 40% answered that data should be used about half of the times when prioritizing requirements (Q18).

When we analyzed the data by building a simple linear model (Eq. 1) using a Bayesian approach, the results show a difference between today (‘present’ in Fig. 2) and the future. In Fig. 2, we see that the lower Likert scale values (e.g., answers ‘never’ and ‘sometimes’) are more common for Present, while the higher Likert scale values (e.g., answers ‘always’ and ‘most of the time’) are more common for the Future. That is, the respondents, with a high certainty, are positive to use DDDM in the future. When comparing RE related questions (Q17 and

Q18) with non-RE related questions (Q13–Q16), the Bayesian model (Eq. 2) shows a difference, as shown in Fig. 3. That is, although the respondents are positive to use DDDM in the future in general (as shown in Fig. 2), the respondents are more positive to use DDDM in non-RE related decisions compared to RE-related decisions.

Reasons for Using (not Using) Data. We asked the respondents what the reasons for using data in today’s decision making is. According to the respondents, the main reason is that DDDM improves the decisions. One respondent explained that when data has been used as input to decision makers, the decisions have been more informed and more transparent. Another reason mentioned by the respondents was, if data is available, then we use it.

A few respondents also gave reasons for partial data use: although the data is there and can improve decisions, it requires a lot of work to filter the data and to present the data in a way that is useful for the decision makers; thus it is only used sometimes for critical/important products/strategies.

Looking at Table 5, we see that *data is not available to us at the company* is the most common reason (82% of the respondents). Most of the respondents who stated that data is not available, also mentioned several other reasons for not using DDDM, including *too much data is available out there* (79% of the respondents), *do not know how to use the data* (73% of the respondents), and *do not know how to make the data relevant to us* (70% of the respondents). Several of the most mentioned reasons for not using DDDM are related to the decision makers’ understanding of the data (including the visualization), and how to make use of it. This confirms the findings in [17]. In order to fully benefit from DDDM, the quality of the data is important as it is directly related to the quality of the decisions [17]. Therefore, it is surprising that only 6% of the respondents mentioned that data is not used in today’s decision making due to the quality of the data. Either, decision making in agile is different or respondents are less aware of these important considerations.

5.4 How Can Data Improve Decision Making (RQ3)

We asked the respondents if they believe data could help them in making better decisions (Q19 in Table 6). Eleven percent of the respondents believe data will improve their decisions (answered ‘yes’), while a majority (58%) believe that data, in combination with other aspects (described below), will lead to better decisions. Close to a third (29%) of the respondents believe data may help in making better decisions but they weren’t sure (i.e., they answered ‘maybe’). Their stated reasons were: (1) have not used data hence do not know if it will lead to better decisions, (2) it depends on which data, the quality of the data, and who makes decisions, (3) and what kind of decisions and when the decisions are made. Only 2% of the respondents do not believe data will help in making better decision. One respondent explained this by stating “*data can never replace my own experiences and gut-feeling*”.

Table 5. Reasons for not using data in decision making

Reason	Respondents
Data is not available to us at the company	82%
Too much data is available out there	79%
Do not know how to use the data	73%
Do not know how to make the data relevant for us	70%
Do not know how to link/use data in relation to decisions	52%
Do not have appropriate tools	31%
Which data should be used?	23%
Cannot trust the data	11%
Do not know how to access the data	7%
Not sure about the quality of the data	6%
Too many systems/tools that store the data	4%

Table 6. Respondents' views if data improves decision making

Q19: Do you believe that you could have made better decisions if data was used as input to decision making?	Respondents
Yes	11%
Yes, if combined with...	58%
Maybe	29%
No	2%

The respondents identified five aspects that needs to be combined with DDDM in order to make better decisions. The five aspects are: (1) own experience, (2) business value, (3) customer value, (4) input from key stakeholders, and (5) experiences from others.

In order to be able to use the full potential of DDDM and thus truly change how decisions are made in ASD, new approaches to provide and visualise constructive and understandable data (information) to the decision makers are needed. By combining understandable visualizations of data and human expertise, the future of DDDM in ASD looks promising.

6 Conclusions

There is a general trend towards data-driven decision making (DDDM), i.e., basing and driving decision making on and with data. However, there has been a lack of studies on how software practitioners view and use this and, in particular, in an agile context. In this study we thus performed a survey and collected questionnaire responses from 84 software practitioners working with agile software development.

Our main result is that the practitioners see a lot of potential for DDDM but that this potential is currently unfulfilled. While very few respondents indicated more wide-spread data-driven decision making in their current practice, a clear majority saw it as important and highly valued in the future. They were more positive to its future use for higher-level and more general decision making, fairly positive to its use for requirements elicitation and prioritization decisions, while being less positive to its future use at the team level. Multiple reasons were given for data not being used today, in particular it may not be available, be available in too large quantities, or it may not be clear how to use it, make it relevant and link it to decisions. Notably, respondents seemed less concerned about quality and trust issues around data.

Our results show that there is an unfulfilled potential for data-driven decision making in agile software development contexts. Future research should investigate this in more detail and also develop new automated data collection, analysis and visualisations techniques and methodologies that augments existing, agile decision processes by linking relevant data to specific decision contexts.

References

1. Petersen, K., Wohlin, C.: A comparison of issues and advantages in agile and incremental development between state of the art and an industrial case. *J. Syst. Softw.* **82**, 1479–1490 (2009)
2. Schön, E-M., Winter, D., Escalona, M.J., Thomaschewski, J.: Key challenges in agile requirements engineering. In: *Agile Processes in Software Engineering and Extreme Programming*, XP, pp. 37–51 (2017)
3. Moe, N.B., Aurum, A., Dybå, T.: Challenges of shared decision-making: a multiple case study of agile software development. *Inf. Softw. Technol.* **54**, 853–865 (2012)
4. Olsson, H.H., Bosch, J.: From opinions to data-driven software R&D: a multi-case study on how to close the ‘Open Loop’ problem. In: *40th Euromicro Conference on Software Engineering and Advanced Applications*, pp. 9–16 (2014)
5. Cockburn, A., Highsmith, J.: Agile software development: the people factor. *Computer* **34**(11), 131–133 (2001)
6. Provost, F., Fawcett, T.: Data science and its relationship to big data and data-driven decision making. *Big Data* **1**(1), 51–59 (2013)
7. Maalej, W., Nayeibi, M., Johann, T., Ruhe, G.: Toward data-driven requirements engineering. *IEEE Softw.* **33**(1), 48–54 (2016)
8. Milne, A., Maiden, N.: Power and politics in requirements engineering: embracing the dark side? *Requir. Eng. J.* **17**(2), 83–98 (2012)
9. Magazinius, A., Börjesson, S., Feldt, R.: Investigating intentional distortions in software cost estimation-an exploratory study. *J. Syst. Softw.* **85**(8), 1770–1781 (2012)
10. Magazinius, A., Feldt, R.: Confirming distortional behaviors in software cost estimation practice. In: *37th EUROMICRO Conference on Software Engineering and Advanced Applications*, pp. 411–418 (2011)
11. Feldt, R., Staron, M., Hult, E., Liljegren, T.: Supporting software decision meetings: heatmaps for visualising test and code measurements. In: *39th EUROMICRO Conference on Software Engineering and Advanced Applications*, pp. 62–69 (2013)

12. Wnuk, K., Gorschek, T., Zahda, S.: Obsolete software requirements. *Inf. Softw. Technol.* **55**(6), 921–940 (2013)
13. Gren, L., Svensson, R.B., Unterkalmsteiner, M.: Is it possible to disregard obsolete requirements? - an initial experiment on a potentially new bias in software effort estimation. In: 10th International Workshop on Cooperative and Human Aspects of Software Engineering, pp. 56–61 (2017)
14. Chen, C.L.P., Zhang, C.-Y.: Data-intensive applications, challenges, techniques and technologies: a survey on big data. *Inf. Sci.* **275**, 314–347 (2014)
15. Feldt, R., Neto, F.G., Torkar, R.: Ways of applying artificial intelligence in software engineering. In: 6th International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering, pp. 35–41 (2018)
16. Brynjolfsson, E.: Strength in numbers: how does data-driven decision making affect firm performance? SSRN (2011). <https://ssrn.com/abstract=1819486>
17. Janssen, M., van der Voort, H., Wahyudi, A.: Factors influencing big data decision-making quality. *J. Bus. Res.* **70**, 338–345 (2017)
18. Punter, T., Ciolkowski, M., Freimut, B., John, I.: Conducting on-line surveys in software engineering. In: International Symposium on Empirical Software Engineering, pp. 80–88 (2003)
19. Robson, C.: *Real World Research*. Blackwell, Oxford (2002)
20. Creswell, J.W.: *Educational Research: Planning, Conducting, and Evaluating Quantitative and Qualitative Research*. Pearson, London (2011)
21. Gliner, J.A., Morgan, G.A.: *Methods in Applied Settings: An Integrated Approach to Design and Analysis*. Lawrence Erlbaum Associates, Mahwah (2000)
22. Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B.: *Bayesian Data Analysis*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis, Abingdon (2013)
23. Carpenter, B., et al.: Stan: a probabilistic programming language. *J. Stat. Softw.* **76**(1), 1–32 (2017)
24. Wohlin, C., Runeson, P., Höst, M., Ohlsson, M.C., Regnell, B., Wessén, A.: *Experimental Software Engineering - An Introduction*. Kluwer Academic Publisher, Dordrecht (2000)
25. Furia, C.A., Feldt, R., Torkar, R.: Bayesian data analysis in empirical software engineering research. arXiv preprint [arXiv:1811.05422](https://arxiv.org/abs/1811.05422) (2018)
26. Torkar, R., Feldt, R., Furia, C.A.: Arguing practical significance in software engineering using Bayesian data analysis. arXiv preprint [arXiv:1809.09849](https://arxiv.org/abs/1809.09849) (2018)
27. Bürkner, P.-C., Charpentier, E.: monotonic effects: a principled approach for including ordinal predictors in regression models. *PsyArXiv*: psyarxiv.com/9qkhj (2018)
28. Daneva, M., et al.: Agile requirements prioritization in large-scale outsourced system projects: an empirical study. *J. Syst. Softw.* **86**(5), 1333–1353 (2013)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

