

# Determining Relevancy: How Software Developers Determine Relevant Information in Feeds

Thomas Fritz and Gail C. Murphy  
University of British Columbia  
{fritz,murphy}@cs.ubc.ca

## ABSTRACT

Finding relevant information within the vast amount of information exchanged via streams, such as provided by Twitter or Facebook, is difficult. Previous research into this problem has largely focused on recommending relevant information based on topicality. By not considering individual and situational factors that can affect whether information is relevant, these approaches fall short. Through a formative, interview-based study, we explored how five software developers from a team determined relevancy of items in two kinds of project news feeds. We identified four factors that the developers used to help determine relevancy and found that placement of items in source code and team contexts can ease the determination of relevancy.

## Author Keywords

News feeds, relevancy, context, interview study.

## ACM Classification Keywords

H.5.3 Group and Organization Interfaces.

## INTRODUCTION

Information streams—tweets, status updates from services such as Facebook<sup>1</sup> and other RSS feeds—enable users to stay aware of information that might be relevant to them. Users of such information streams face two problems. First, they must determine which streams contain information that may hold relevance for them. Second, in many cases, not all of the information in a stream is relevant, forcing the user to separate the wheat from the chaff. Malone identified this second problem in the context of distribution lists more than twenty years ago [5].

Surprising to us, there appears to be very little formative research about the question of how to determine relevant information in a stream. Instead, the research has focused on how to recommend items within a stream, largely based on topicality; for example, the approach of Chen *et al.* [2] recom-

<sup>1</sup>[www.facebook.com](http://www.facebook.com), verified 24/09/10

mends tweets by matching terms recorded in a profile of the user and terms in the tweets. Schamber *et al.* [7] argue that a determination of relevance should go further and should involve understanding how users perceive information relative to their current situation. They suggest the use of more qualitative research methods to better understand how users determine which parts of information are relevant.

In this paper, we take this direction and present the results of a small formative study that investigated two questions: how does a user determine relevance and can the context of the situation be used to help ease the determination of relevance. Our aim was to allow the formation of data-driven hypotheses that can be subjected to subsequent, more detailed testing. We chose to study a team of software developers because this situation allowed us to study multiple individuals working with similar information items in a similar work environment.

Our exploratory study involved a team of five developers who use IBM's Rational Team Concert (RTC) environment in their daily work. RTC includes a default feed reader that provides news item updates about several information streams, including changes to shared work items (e.g., tasks or defects) and changes to shared code (e.g., change sets). Over the period of a week, we interviewed each developer two times, presenting each with a subset of the information items in their feed reader and asked them questions related to the relevance of those items. In all, we studied 291 news items.

We found that each user had their own definition of relevance, ranging from a news item being relevant because it impacted the individual's work in the present to more generic determinations that the item may help to do one's job in the future. We also found that users use the kind of information stream, the content of an item in the stream, the target of the content, their relationship with who created the item and whether or not the news had been discussed elsewhere to determine relevancy of individual items. Finally, we found that presenting news items in the context of a developer's work does appear to hold promise for easing relevancy determination. These findings lead to hypotheses that can inform future studies and that can suggest new approaches for automated support for managing information streams.

## RELATED WORK

While there have been studies of the workload imposed on machines by feeds [4], we were able to find only one study that looked at the usage of feeds by humans. Treude *et al.* [9]

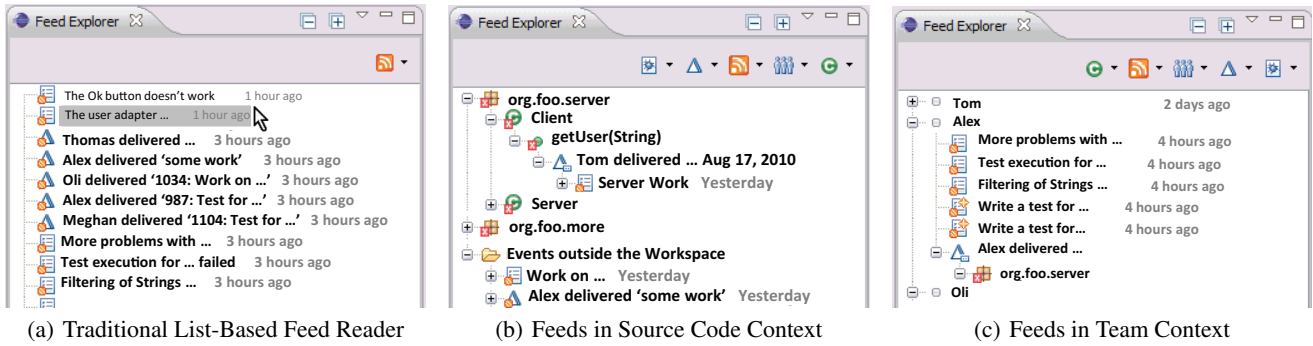


Figure 1. Presentation of Feeds.

studied how software developers use feeds in RTC for accomplishing awareness. They did not consider how developers assess the relevancy of the feeds or the items within, which is our focus in this paper.

A number of researchers have attempted to determine relevancy automatically using recommender-based approaches. Chen *et al.* [2] explore a variety of recommender algorithms, largely based on topic relevance, to find interesting URLs in Twitter feeds. Dabbish *et al.* [3] predict the importance of an e-mail message to a user based on predefined categories. These approaches do not try to understand how users define an item of information as relevant to their current situation.

Other research has looked at the problem of staying aware of relevant information. Palantir [6], for example, provides awareness of changes to the source code in context of a developer’s workspace. FASTDash [1] extends the approach by also visualizing which files a developer is currently working with. This research mainly assumes an information item as relevant if it affects the source code in the developer’s workspace or team.

## INTERVIEW METHODOLOGY

For our study we interviewed a sample of five developers (4 male, 1 female) from a development team within a large corporation. The team was chosen based on their use of an environment that contains project feeds (IBM’s RTC) and the access we were granted to their project specific changes and information. The participants had between 4 and 24 years of professional development experience ( $M=10$ ,  $SD=9$ )<sup>2</sup> and at least 20 months of experience using RTC ( $M=23$ ,  $SD=7$ ).

We interviewed each participant twice, with four days in between the first and the second interview. In each interview, we presented the participant with a list of 30 items<sup>3</sup> from two default feeds that exist in RTC. One feed, *change set news*, captures changes to the code and the other feed, *work item news*, captures changes to work items, which are similar to bugs, issues or tasks. For each participant, we randomly selected the presented items from the set of all items that had occurred for the participant over the previous three days, en-

suring that the presented items represented both feeds and that the participant was not the creator of the news item. In each of the two sessions, we asked each participant to review the items, one at a time, and to tell us whether or not the item was relevant, how relevant the item was and why the item was relevant (or not). Participants had the option to read and explore the item and related information as much as desired. We deliberately did not define relevance for the participants to allow investigation of differing ideas of relevance between the participants.

In the first interview session, the news items were presented in a list, based on the interface of traditional feed readers (Figure 1(a)). In the second interview session, we asked the participants to review an additional 30 news items. This time, we showed the participant the news items in the context of the source code on which the team worked. If a news item could be related to the source code, the news item was shown hierarchically beneath the related source code (Figure 1(b)). If one of the news items did not relate to the source code, it was shown in a flat list below the code. We asked the participant whether the context changed their determination of the relevance of an item and whether or not the context was helpful. Finally, we showed the news items in terms of a team context (Figure 1(c)), in which a news item was shown underneath the team member who caused the creation of the news item. We asked the participant the same questions for the team context as the source code context.

Each interview session took between 20 and 60 minutes. All interviews were recorded and the interviewer (first author of this paper) took handwritten notes. Due to the exploratory nature of our study, we parsed our data using an open coding technique to develop and identify categories of data [8].

## FEEDS

During the study period, a mean of 50 ( $SD=31$ ) items per day were available to the participant in their feed reader, created by a mean of 5 ( $SD=2$ ) authors, excluding the participant him or herself. Most, 46 of 50 ( $SD=32$ ), items were about work items; the remaining 4 of 50 ( $SD=3$ ) were about change sets.

In total, our study gathered feedback on 291 news items. Overall participants perceived 79 of the 291 (27%) news items as relevant, 24 (8%) as somewhat relevant and 188

<sup>2</sup> $M$ ,  $SD$  stand for mean and standard deviation respectively.

<sup>3</sup>In one case, only 21 news items were available.

(65%) as not relevant. Participants perceived more of the change set news items as relevant compared to work item news; of the 81 change set news items, participants perceived 54 of the 81 (67%) as relevant or somewhat relevant, whereas only 49 of the 210 (23%) items related to work items were considered relevant or somewhat relevant. Since all participants worked on the same team, many of the news items overlapped.

In the second interviews, we placed the items into the source code context. Of the 150 news items considered in these interviews, 37 (25%) were related to the source code of the project the team was working on and thus could be placed into context. In the team context, all the news items could be related to team members and were therefore shown in context. However, the team was broadly defined and thus comprised developers that were not seen as direct members of the team of the 5 participants.

Although all participants knew of the feed reader in RTC, only one participant described using the view on a regular basis. Others just used it occasionally or very rarely. Instead, they often used other, more specialized views available within RTC to monitor changes. In particular, all participants stated that they regularly monitor changes to the source code, the information summarized in the change set news, using a view called the Pending Changes View.

### USER-ORIENTED VIEW OF RELEVANCE

Each participant defined relevancy of a news item differently: a news item is relevant when it impacts the code that I am familiar with (P1), when the information is useful to me now or in the foreseeable future (P2), when it's in an area in which I am an expert of, interested in and almost are lacking expertise (P3), when it helps me to do my job at some point (P4) and if it is in an area that I see myself as a keeper or knowledgeable in or critical to the project (P5).

Classification of the detailed reasons given by participants when determining relevancy suggests nine broad categories. Figure 2 presents these categories and shows how many items fit into each category. Some of the categorized reasons were used mainly for identifying an item as relevant. For example, participants mentioned an item to be about a change set in 16 cases and in all these cases the reason was used to explain that the item was relevant. Others, such as an item being not from someone on the participant's team, were used to explain items as being irrelevant. Finally, some of the categorized reasons were used to identify items as relevant as well as irrelevant.

Participants did not apply consistent criteria for determining relevancy. For instance, for only 16 of the 81 (20%) change set news items did participants indicate the fact that the news item was a change set that made the item relevant. Although the perception of relevance is individual and situation-based, we have identified four dimensions along which participants determined the relevancy of an item: the content of the item, the target of the content, the relation with the creator and the previous interaction with the topic of the news item.

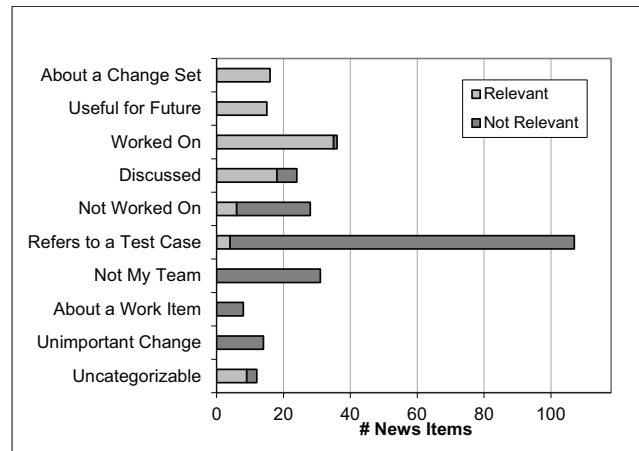


Figure 2. Categorized Reasons Developers Mentioned for Identifying a News Item as Relevant or Irrelevant.

### Content

All participants generally rated change set news as more relevant than work item news. One participant even went as far as saying that all change set news for the project are relevant but that work item news are not relevant as too many items appear and much of the information was available for the team on a physical white board (P4). This content-based approach is well-supported by today's tools that support feed subscription based on content. However, as only 67% of all change set news and 23% of work item news were considered as relevant, this statement is not well supported by the data from all participants. Furthermore, this same participant (P4) admitted that because he generally ignores work item news, he also misses questions his team lead asks of him through work items, suggesting a content-based approach has significant drawbacks. Interestingly, in the interviews, the same participant considered several work item news as relevant because of the target of the content or the relation with the creator.

Other participants considered the type of the content, as opposed to the type of the feed, to determine relevance. For example, most participants stated that work item news concerning test cases were not relevant. Other factors did override these general trends. For instance, the participants described that the state of the project affects the relevancy of work item news referring to test cases. In some project states, such as right before releases or when the system fails to run, work item news related to testing is more likely to be deemed relevant.

This trend of using the type of an item was also seen in other work item news deemed generally unimportant and irrelevant, such as items on the addition of a work item to a feature (P1, P2) or on the change of a subscriber to a work item (P5).

### Target of Content

The target of the content of a news item was often mentioned as a reason for an item being relevant or not. One participant stated, "This is not relevant since I have not worked on that

[part of the code] for a while” (P1). Another stated, “[this is relevant as] the area is related to what I worked on” (P5). As with the content, the relevance based on the target of the content is also very dependent on the situation and the individual. For instance, (P2) stated that even though an item is not in an area he worked on, it could be useful for the future and thus perceived as relevant. Tools could be built to help automate this factor of relevance determination; tracking data about which code a developer touches and changes could be used to filter information streams to only show items on which a developer does (or does not) work.

#### *Relation with Creator*

In most cases, an item created by a developer not on the direct team of the participant was perceived as less relevant (P1, P2, P4, P5). However, in case the participant worked together with the developer from the different team on a regular basis, the item was perceived as highly relevant (P3). Some participants stated that a combination of factors may make an item relevant; for instance, if an item was created by a certain teammate and touched an area of related work it was relevant (P1). Tools could be built to support the use of combinations of factors in automating determining relevance of items.

#### *Previous Interaction*

Whether or not the news had been discussed in another setting had an impact on the perceived relevance of the corresponding news item when it appeared in the feed. Some stated that work item news were not interesting because the news had already been discussed in the planning sessions (P4). Others stated that seeing a news item related to a previously discussed issue was relevant as it confirms what is known (P3). One participant (P5) had mixed views and identified one item as relevant and another one as irrelevant based on having previously discussed these items. Providing support to automatically determine if news had been seen in another setting appears much more challenging. A better characterization of the impact of this factor would be helpful to ascertain in future studies.

#### **CONTEXT TO REDUCE INFORMATION OVERLOAD**

Although participants stated that their knowledge about the part of the code to which the news item related was used in determining relevance, only 37 of 150 (25%) were easily determined to relate to the code on which the participants worked. In 20 of the 37 (54%) cases, participants considered the source code context to be helpful. In 9 of 37 (24%) cases the participants referred to it as being somewhat helpful and in 8 (22%) cases as being not helpful. In particular, when the short form of the summary used to identify the news item was too generic, such as “delivered 2 change sets”, the context provided a significant benefit to the participant in explaining the item.

When asked about source code context in general, three out of the five participants considered the placement of news items in context helpful (P1,P2,P5). One of the three stated that he “had problems with the [news items] out of context, but [for] the ones within, [he could determine relevancy]

pretty quick” (P1). Another participant preferred the team context, stating that the source code context provided too much information but that the team context allows skimming over people quickly (P3). Three more participants (P1,P2,P5) conceived the team context as helpful, in particular in situations in which you do not have to work with a lot of source code, such as quality control.

#### **CONCLUSION**

A better understanding of how users determine relevance of items from information streams, or feeds, is needed to build better tool support to help users deal with the vast amount of information flowing to them in this form. Our exploratory study of how developers determine the relevancy of items in project-related feeds resulted in a description of four factors that impact relevancy. Only one of these factors, the type of information stream, is well-supported by today’s tooling. Other factors such as the target of the content and the relationship with the creator appear plausible to develop. The fourth factor, previous interaction, requires additional study to better understand its impact in the determination of relevancy. Our study also shows that the placement of information stream items into user-oriented context can be helpful in determining relevance. Further study is needed to understand how these factors generalize into other settings.

#### **REFERENCES**

1. J. T. Biehl, M. Czerwinski, G. Smith, and G. G. Robertson. Fastdash: a visual dashboard for fostering awareness in software teams. In *Proc. of CHI’07*, 2007.
2. J. Chen, R. Nairn, L. Nelson, M. Bernstein, and E. Chi. Short and tweet: experiments on recommending content from information streams. In *Proc. of CHI’10*, 2010.
3. L. A. Dabbish, R. E. Kraut, S. Fussell, and S. Kiesler. Understanding email use: predicting action on a message. In *Proc. of CHI’05*, 2005.
4. H. Liu, V. Ramasubramanian, and E. G. Sirer. Client behavior and feed characteristics of rss, a publish-subscribe system for web micronews. In *Proc of IMC ’05*, 2005.
5. T. W. Malone, K. R. Grant, F. A. Turbak, S. A. Brobst, and M. D. Cohen. Intelligent information-sharing systems. *Commun. ACM*, 30(5), 1987.
6. A. Sarma, Z. Noroozi, and A. van der Hoek. Palantír: raising awareness among configuration management workspaces. In *Proc. of ICSE’03*, 2003.
7. L. Schamber, M. B. Eisenberg, and M. S. Nilan. A re-examination of relevance: toward a dynamic, situational definition. *Information Processing & Management*, 26, 1990.
8. A. C. Strauss and J. Corbin. *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*. Sage Publications, Inc, 2nd edition, 1990.
9. C. Treude and M.-A. Storey. Awareness 2.0: staying aware of projects, developers and tasks using dashboards and feeds. In *Proc. of ICSE ’10*, pages 365–374, 2010.