

TaggedComments: Promoting and Integrating User Comments in Online Application Tutorials

Andrea Bunt¹, Patrick Dubois¹, Ben Lafreniere², Michael Terry², David Cormack¹

Department of Computer Science¹

University of Manitoba,

Winnipeg, Manitoba, Canada

bunt@cs.umanitoba.ca, umdubo26@cc.umanitoba.ca

Cheriton School of Computer Science²

University of Waterloo

Waterloo, Ontario, Canada

{bjlafren, mterry} @cs.uwaterloo.ca

ABSTRACT

User comments posted to popular online tutorials constitute a rich additional source of information for readers, yet current designs for displaying user comments on tutorial webpages do little to support their use. Instead, comments are separated from the tutorial content they reference and tend to be ordered according to post date. We propose and evaluate the TaggedComments system, a new approach to displaying comments that users post to online tutorials. Using tags supplied by commenters, TaggedComments seeks to enhance the role of user comments by 1) improving their visibility, 2) allowing users to personalize their use of the comments according to their particular information needs, and 3) providing direct access to potentially helpful comments from the tutorial content. A laboratory evaluation with 16 participants shows that, in comparison to the standard comment layout, TaggedComments significantly improves users' subjective impressions of comment utility when interacting with Photoshop tutorials.

Author Keywords

Web-based Tutorials; Feature-Rich Software; Tagging

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Feature-rich software can be difficult for users to learn and use [9]. While most applications include in-application help pages, users' help-seeking practices typically involve turning to their favourite search engines [5], where they are directed to a large array of web-based tutorials. Given the now important role of online tutorials, there has been strong interest within the research community in improving the experience of interacting with tutorials (e.g., [2,6,16]).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2014, April 26 - May 01 2014, Toronto, ON, Canada
Copyright is held by the authors. Publication rights licensed to ACM.
ACM 978-1-4503-2473-1/14/04...\$15.00.
<http://dx.doi.org/10.1145/2556288.2557118>

An element of online tutorial design that has received little attention is the comments section. A recent analysis found that user comments on popular text- and image-based tutorials constitute a rich and diverse source of information [15]. For example, readers thank the author, write content-specific questions, provide answers to other users' questions, offer tips, and make requests for help on tasks that are extensions of the given tutorial, or are only tangentially related. Despite the utility of these comments, current comment display interfaces place the comments at the bottom of the tutorial and order them chronologically, doing little to leverage their value. This design is problematic when the reader seeks information likely found in the comments, as the next example illustrates.

Consider a user attempting to solve a problem that a previous user has faced with the tutorial. To find relevant comments, she would need to scroll to the bottom of the tutorial, and read through the chronologically-ordered list of comments, sifting through dozens of unrelated comments (e.g., social comments, requests for help on other tasks). If she does manage to find a relevant comment, she would then need to re-locate her place in the tutorial content before continuing the tutorial. This added cognitive load would be on top of any created by trying to follow and learn from the original tutorial content.

We present an alternative design for tutorial comment sections of online text- and image-based tutorials that allows users to tag their comments, and optionally pin them to particular locations in the tutorial content. Based on this design, we illustrate how TaggedComments 1) allows for *personalized comment use* by allowing tutorial authors and users to filter the comments according to their information needs; 2) provides *tutorial help in context*, by integrating comments with the tutorial content they reference; and 3) *highlights community feedback* to aid in the tutorial selection process and to assist with community building.

We also describe the results of a laboratory evaluation examining TaggedComments' impact on users' experience completing a Photoshop tutorial. Our results show that users' perceptions of comment utility were significantly higher with TaggedComments than with the standard comment organization. Our results also suggest that TaggedComments helped create a sense of community surrounding the tutorial.

RELATED WORK

Our coverage of related work focuses on approaches to improving application tutorials, alternative forms of crowdsourced help for feature-rich software, and comment tagging in support of other types of complex tasks.

Enhancing Application Tutorials

Prior work on enhancing application tutorials falls into three primary areas: automatically generating tutorials, supporting tutorial selection, and supporting use of tutorials. We discuss each area in turn.

Automatic Generation

One line of work has looked to improve application tutorials by making it easier for authors to create them. Much of this work has focused on automatically generating tutorials based on workflow demonstrations. For example, Grabler *et al.*'s system automatically generates text- and image-based tutorials through a combination of image-recognition techniques and author demonstration [8], whereas MixT adds short video demonstrations to tutorial instructions [2]. Taking a different approach, Chronicle provides interactive, annotated document histories [10].

Supporting Tutorial Selection

Work on helping users select between tutorials has primarily focused on highlighting the application commands used in the tutorial. Delta supports tutorial selection by displaying the commands used for each workflow (and their order) [14]. Ekstrand *et al.* use knowledge of the user's command history to augment Google's search results with the icons for the commands referenced within each result page [5]. In contrast to these approaches, TaggedComments uses a summary of the community's response to a tutorial to help guide selection.

Supporting Tutorial Completion

Lastly, there have been a number of approaches to designing more effective and engaging tutorials. Stencil-based tutorials scaffold the user's progress through the tutorial by highlighting the commands required at each step [13]. Sketch-Sketch Revolution guides the user through a drawing tutorial, providing feedback and stroke guidance [6]. TApps use automation to reduce the effort involved in completing a tutorial, while still allowing the user to experiment with parameter settings [17]. Recent work has also proposed gamifying tutorials as a mean of increasing tutorial engagement [19].

There has been little work on enhancing tutorials using feedback from its user community. One exception is FollowUs, which captures video demonstrations from users and makes them available on a step-by-step basis [16]. An evaluation showed that users experienced less frustration with FollowUs than they did in a control condition where only the author's demonstration was available.

Crowdsourced Help for Feature-Rich Software

By integrating user comments with the tutorial content they reference, TaggedComments enables a form of

crowdsourced help. In this sense, our work shares similarities with systems like IP-QAT [22] and LemonAid [3], which provide in-application assistance for use of software through integrated, crowdsourced Q&A. In contrast to this prior work, tutorials are typically multi-step documents describing complete and often complex workflows. As a result, most tutorial comments are targeted toward the *tutorial*, the workflow it describes, and the use of the software to accomplish that workflow, as opposed to use of the software in general. Tutorials also include multiple stakeholders and interested parties that must be addressed in any design: the readers and the tutorial author.

Tagged Contributions to Support Complex Tasks

While we are not aware of other systems that support tutorial comment tagging, our approach is motivated by prior work demonstrating the value of tagged contributions in support of other complex tasks. CommentSpace, a collaborative visual analytics tool, encourages users to tag their comments as hypotheses, questions or to-do items [23]. An evaluation showed that when using CommentSpace, users generated higher-quality responses for a data synthesis task than when they performed the task using untagged user comments. Similarly, PathFinder supports citizen science by allowing individuals to annotate findings and comments [20]. An evaluation of Pathfinder showed that the increased organization afforded by tagged comments had a number of positive subjective impacts on users' perceived utility of the system.

COMMENT-ENHANCED TUTORIALS

In this section we describe our design for the TaggedComments system (see Figure 1), which aims to enhance the role of user comments within a tutorial. In designing the system, our goals were as follows:

- *Comment visibility:* User comments contain a wealth of information, but are typically deemphasized at the end of the tutorial. Our goal was to promote comments, increasing their visibility to other tutorial followers.
- *Comment personalization:* Comments contain diverse types of information, from bug fixes to words of praise for the author. Our goal was to allow both tutorial authors and consumers to personalize their use of comments according to their information needs.
- *Comment/content integration:* Many posted comments pertain directly to content of the tutorial, for example by pointing out problems with the instructions, providing quality assurance (e.g., testing with multiple versions) or suggesting a solution to an error-prone instruction. We wanted to reduce the separation between comments and their target tutorial content.

To accomplish these goals, the system asks users to both tag their comments and, when appropriate, to "pin" their comments next to the tutorial content to which their comments refer. The next section describes how users tag and pin their comments, followed by an illustration of how TaggedComments leverages this information.

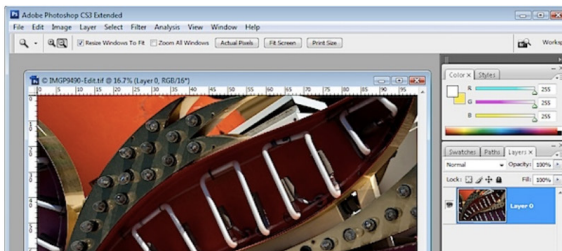
How to Create Kaleidoscope Patterns from Your Photos

If you are like me you had a kaleidoscope as a kid. You would look through one end and turn a dial and the world would be displayed as a mirrored fractured shape. Thanks to Photoshop you can create kaleidoscopes from your photos. Any image with interesting color and shapes will work just fine.



Step 1

Convert the background layer to a regular layer by double clicking it and click OK.



Filter 34 comments by category

Category	Count
All comments	34
Replies	5
Appreciation	9
Suggestions	4
Help with related task	5
Help completing tutorial	5
Images	5
Links	1
Other	0

User-submitted images

Filter comments

All comments

Comments

On 30/07/2013 11:32:13 AM, Anonymous said:
I'm having trouble completing step 3. I can't find the Contract command in my version of the application.
Program Version: 7.0

Reply Go to pin

On 18/04/2013 2:13:55 PM, Anonymous said:
I need help completing step 6.
Program Version: Photoshop CS6

Reply Go to pin

Submit a comment anonymously

Please select a section that best describes your comment.

Asking for help completing the tutorial

Enter the version of the program you are using, if known.

7.0

You have 3896 characters left.

I'm having trouble completing step 3. I can't find the Contract command in my version of the application.

☐ Pin comment into tutorial? Pins?

Submit Comment

Figure 1: Left: The TaggedComments interface, consisting of: the comments section (A); a graph representation of community response (B); an area for users to “pin” comments to the tutorial content (C); an image carousel of user-submitted results (D). Tutorial source: <http://digital-photography-school.com>. Right: Submitting a comment.

Allowing Users to Tagging and Pinning Their Comments

TaggedComments augments the traditional tutorial comment posting interface with a request for information to categorize the comment. As shown in Figure 1(Right), when the user enters a comment, s/he is asked to categorize the comment using one of six pre-defined tags. These categories, based on Lafreniere *et al.*'s analysis of tutorial comments [15] and refined through pilot testing, are:

- *Appreciation*: praise or encouragement for the tutorial author.
- *Suggestion*: a suggestion or tip that may help with completing the tutorial.
- *Help completing the tutorial*: a request for help with some aspect of completing the tutorial.
- *Help with a related task*: a request for help with a task that goes beyond the focus of the tutorial.
- *Images*: a link to the results achieved by following the tutorial.
- *Link to related resource*: a link to resources that are helpful for completing the tutorial (e.g., brush sets).

We chose to focus on pre-defined tags for two reasons. First, prior work suggests that pre-defined tags offer a number of advantages over free-form tags (e.g., [1,7]). Second, using pre-defined tags allowed us to build dialogs into the system that prompt users for additional category-specific information to make the comments more valuable. For ex-

ample, as shown in Figure 1(Right), when requesting help, the user is asked to specify which version of the application s/he is using. To accommodate emerging comment categories, users also have the option to enter their own free-form tags.

In addition to tagging their comments, users can choose to associate a comment with a position in the tutorial content by placing a “pin” on the vertical strip separating the tutorial content and the commenting interface (see Figure 1(C) and Figure 2(Right) for a closer view). Each pin displays the number of user comments associated with that location and a callout to help reinforce the association between the pin and the comments. We initially experimented with allowing users to pin the comments within the text itself (as opposed to alongside the content), but pilot testing revealed that readers had difficulty locating the pins with such tight embedding.

The addition of tags and pins to comments enables the system to highlight community feedback, support personalized comment use, and provide in-context community help. We describe each feature below.

Highlighting Community Feedback

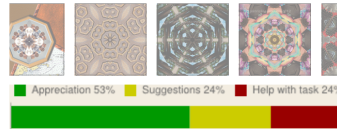
To help users decide whether a tutorial is worth their time and effort, the system provides a succinct summary of community feedback based on the tags assigned by users. At the top of the TaggedComments sidebar on the tutorial page, feedback is displayed using a color-coded bar graph

Hide [How to Create Kaleidoscope Patterns from Your Photos](#)

If you are like me you had a kaleidoscope as a kid. You would look through one end and turn a dial and the world would be displayed as a mirrored fractured shape. Thanks to Photoshop you can create kaleidoscopes from your photos. Any image with interesting color and shapes will work just fine.

Author: Helen
Program: Photoshop
comments: 33 comments

Date: October 31st 2009
Time: 20 mins



Hide [3D Pixel Stretch Effects in Photoshop](#)

Stretching a single line of pixels is an easy way to create special digital effects in graphic design. By taking this technique a step further, we can create wonderful, vibrant 3D effects right in Photoshop!

Author: David
Program: Photoshop
comments: 109 comments

Date: July 23rd 2008
Time: 30 mins

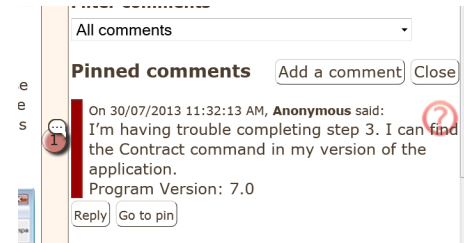
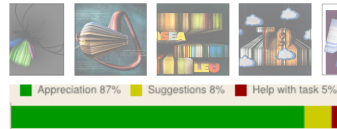


Figure 2: Left: Displaying comment prevalence for a set of tutorials. Right: Accessing a pinned comment.

depicting the prevalence of each comment type (Figure 1(B)). Below the graph, the system displays an image carousel of example images (or other content) that users have produced by following the tutorial (Figure 1(D)).

This information could also be embedded in search results, or a list of tutorials in a repository (Figure 2(Left)), to enable users to quickly evaluate a number of potential tutorials based on the percentage of appreciation comments, suggestions, and requests for help associated with each. This kind of summary mechanism captures community response to the tutorial and could complement command-oriented approaches (e.g., [14]) for helping users to select between tutorials.

In addition to guiding tutorial selection, highlighting community feedback reinforces the notion of the tutorial as a “community meeting place” [15] and may encourage users to post. There is also evidence that the act of tagging can encourage participation and create a greater sense of community [18].

Personalized Comment Use

To enable users to personalize the comment display to their particular information needs, TaggedComments includes features to filter and organize comments. For example, a tutorial author might want to look only at requests for help when attempting to improve a tutorial, or requests for help on related tasks when planning to author a new tutorial. Tutorial readers, on the other hand, might wish to focus only on suggestions and tips. With TaggedComments, users can filter comments by clicking on the corresponding label in the bar graph or by selecting a category from the drop down menu. Figure 3 shows an example of both an original set of comments and those filtered according to the “Suggestions” tag. For this particular tutorial, tagged by us for illustration, there are 6 suggestion comments buried within 68 social comments. This is a typical ratio for tutorials with comments on the web [15].

In-Context Community Help

Finally, pinned comments enable users to view content-related comments in context, alongside the content they are referencing. For example, when in the process of perform-

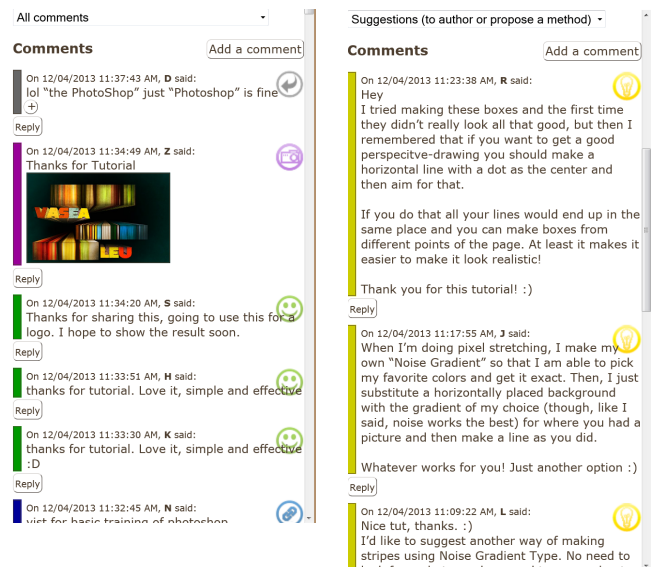


Figure 3: All comments colour-coded according to category (Right) and filtered to display only “Suggestions” (Left).

ing a tutorial, users can view any difficulties other tutorial users have experienced at the current step, and any suggestions for overcoming difficulties. Clicking on a pin in the vertical strip next to the tutorial brings the user to the associated comment(s) as shown in Figure 2(Right). Users can also navigate to pins from the comments section using “Go to Pin” buttons (see the comment in Figure 2, Right).

STUDY

We conducted a laboratory study comparing TaggedComments to a control condition designed to emulate existing tutorial comment systems. The goal of our study was to see how users would react to TaggedComments’ interface, and to understand how its features would change how users referred to comments while completing tutorials. We chose Adobe Photoshop, a feature-rich application with one of the widest showings of web-based tutorials, as our target application.

In the TaggedComments condition, participants completed a Photoshop tutorial with comments marked up and dis-

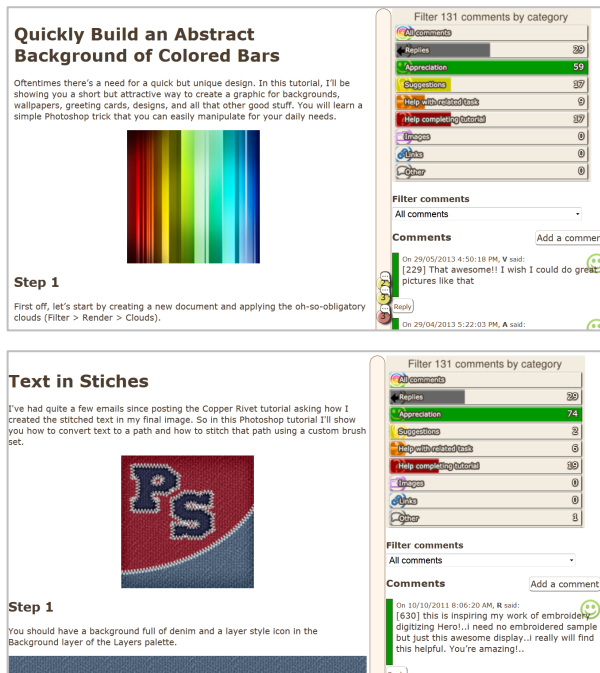


Figure 4: The two study tutorials displayed using the TaggedComments interface: the Background tutorial (Top) and the Stitches tutorial (Bottom).

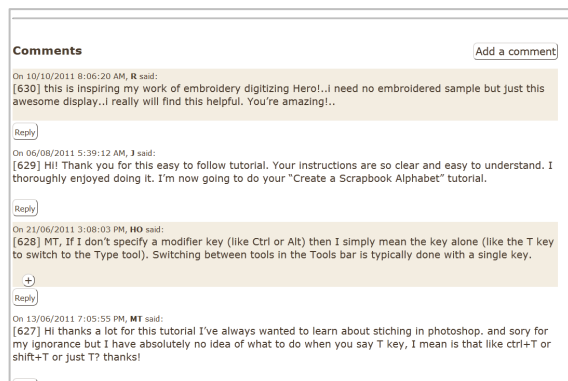


Figure 5: Default condition: Comments are at the bottom of the tutorial content and ordered reverse chronologically.

played using the system as described above (and shown in Figure 4). In the Default condition, participants completed a Photoshop tutorial with comments displayed below the tutorial content in reverse chronological order by posting date (Figure 5). This is the standard design for comment systems on web-based tutorials.

Participants

We recruited 16 volunteers (12 male, 4 female), ages 18-40 via signs posted around a university campus. Participants were pre-screened to ensure that they had some experience using Photoshop. Participants had a range of Photoshop experience levels; four participants indicating using the software multiple times per week, six used it weekly, five

	<i>Background</i>	<i>Stitches</i>
Tutorial Source	psd.tutsplus.com	pshero.com
Original Steps	11	16
Original Comments	190	128
Experiment Steps	8	8
Experiment Comments	131 (41 content-specific)	131 (41 content-specific)

Table 1: Properties of the two study tutorials.

used it monthly, and one yearly. Participants were compensated with a \$15 honorarium.

Tutorials

Our study followed a within-subjects design, with each participant attempting two different text- and image-based tutorials, one per condition. We chose a within-subjects design to elicit comparative statements and to account for individual variability (e.g., owing to Photoshop expertise). To increase external validity, we used tutorials from the web and the user comments already posted on them as a starting point for the tutorials used in the study.

In selecting tutorials from the web, we had three primary criteria. First, we wanted each tutorial to include a reasonable number of comments, some of which could be helpful in completing the tutorial. Second, we wanted two tutorials of similar length and complexity. Finally, we wanted tutorials for tasks likely to be achievable by participants with moderate levels of Photoshop expertise, to ensure access to a sufficiently wide participant pool.

The two tutorials selected are shown in Figure 4 and summaries of their characteristics are listed in Table 1. The *Background* task involved creating an abstract background and the *Stitches* task involved creating a stitches effect around the borders of text. As can be seen in Table 1, we modified the tutorials to balance their length, and to keep them feasible to complete within a 90-minute study session. These modifications consisted of removing steps from the beginning (i.e., providing the participant with a more advanced starting point) and end of the tutorials.

Both tutorials included two steps that were difficult to complete by following the instructions as written, but for which there was helpful advice contained in the user comments. For example, in Step 3 of the Background tutorial, the author suggests using Motion Blur to turn a noisy image into a series of crisp vertical bars. This effect, however, does not work properly with all canvas sizes (including the size that our participants started with). Below are two comments pertaining to this issue, a request for help and a reply with a suggestion:

Photoshop CS4 creates that stripe pattern, but there is always a strange cloud in the middle.

If you get a blurry centre of the image apply the motion blur 4-6 more times. Then increase the contrast to 100% and you should get a decent result.

Following the tip suggested in the comment above would allow the user to complete Step 3 properly. Each tutorial included an additional step where comments could potentially be helpful (e.g., if users deviated from the author's recommend path) but were not as critical.

While we initially intended to use only original user comments, we found that there were more helpful comments on the Background tutorial than there were on the Stitches tutorial. To compensate, we created and added 15 content-specific comments (a mix of questions, comments and tips) to the Stitches tutorial, using similar terminology and commenting styles as the original comments. We then selectively removed non-content related comments from both tutorials to balance the number of comments on each tutorial and the percentage of comments pertaining to tutorial content. Finally, we manually added a tag to each comment, using the tagging scheme described earlier in the paper.

Procedure

After completing a demographics questionnaire, participants completed the two tutorials described Table 1, one with TaggedComments, and one with the Default interface. Interface and task order were fully counterbalanced across participants.

Prior to beginning each tutorial, participants were given a brief demonstration of the interface to be used in that condition. While completing the tutorial, we asked participants to write down the three-digit code that we assigned to each tutorial comment for any comment that they tried to use to help complete the tutorial. Participants were given 30 minutes to complete each tutorial.

After completing each tutorial, participants completed the NASA Task-Load Index (TLX), which measures perceived mental workload [11]. In addition to the six NASA TLX questions, we asked two questions to gauge participants' perceptions of the comments in that condition (using the same 20-point scale as the TLX): 1) "Comment Helpfulness: How much did the comments help you complete the task?" and 2) "Ease of Finding Comments: How easy was it to find helpful comments?"

At the end of the session, participants completed a post-session questionnaire and participated in a short semi-structured interview. Each session lasted approximately 90 minutes.

Data were captured via log files, videotape and audiotape. The videotape and audiotape recordings were later coded and transcribed for analysis purposes.

The experiment was run on a desktop computer with 12 GB of RAM and a 21-inch monitor with 1920×1080 resolution. The TaggedComments interface was implemented using JavaScript.

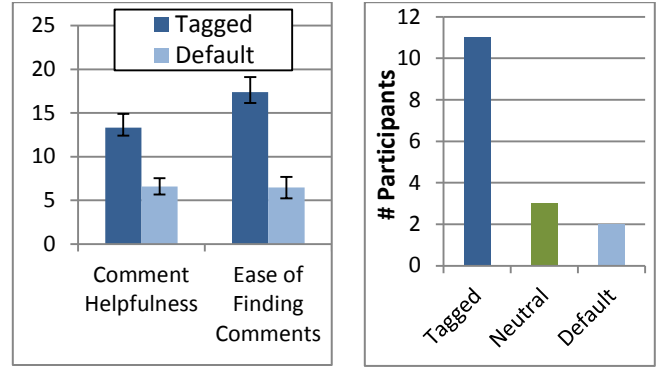


Figure 6: Left: Subjective impressions of comment utility measured after each condition (measured on a 20-point scale). Right: Participants' stated preference for the two interfaces.

Results

We analyzed quantitative data using a repeated-measures ANOVA with condition (TaggedComments vs. Default) as the within-subjects factor. We did not find significant effects or interactions for the order of the conditions or tasks, so we do not discuss these in our analysis below. In all figures, error bars show standard error.

We begin by describing the impact of the two interfaces on participants' subjective impressions of comment utility and their ability to find comments, followed by their impact on task performance and comment use.

Subjective Response

Responses to our questions regarding comment helpfulness, and participants' ability to find comments both clearly favoured TaggedComments (Figure 6, Left). On both measures, the mean values for TaggedComments were more than double those for Default, and both differences were found to be significant (Comment Helpfulness: TaggedComments 13.3, s.e. 1.6 vs. Default 6.6, s.e. 1.7, $F_{1,14} = 12.350$, $p = .003$, $\eta^2 = .469$; Ease of Finding Comments: TaggedComments 17.4, s.e. 0.9 vs. Default 6.5, s.e. 1.2, $F_{1,12} = 35.594$, $p < .001$, $\eta^2 = .656$).¹

The majority of participants also expressed a preference for the TaggedComments system in a post-study questionnaire, with 11 participants preferring TaggedComments, 2 preferring the Default interface and 3 expressing no preference. This difference was found to be significant according a chi-squared test ($\chi^2 = 9.125.0$, $p = .01$).

Despite this clear difference in subjective preference, we did not find a difference in perceived mental workload as measured by the NASA TLX, either overall ($F_{1,15} = 0.224$, $p = .634$, $\eta^2 = .015$) or on any of the individual subscales.

¹ After not accessing comments, some participants responded "N/A" (1 for "Comment Helpfulness" and 3 for "Ease of Finding Comments").

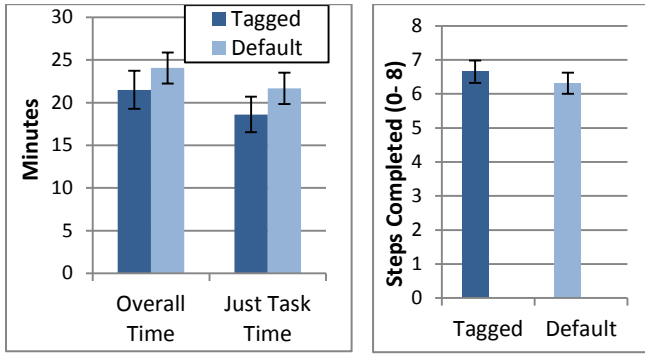


Figure 7: Left: Task time per condition overall and removing the time spent in the comments section. Right: The number of tutorial steps completed per condition.

	Tagged mean (s.e.)	Default mean (s.e.)	$F_{1,15}$	p	η^2
Visits	7.7 (1.5)	4.8 (1.7)	1.493	0.241	0.091
Time (min)	2:52 (0:32)	2:24 (0:34)	0.434	0.520	0.34
Read	10.4 (2.0)	5.8 (1.8)	2.586	0.129	0.147
Used	2.3 (0.4)	1.3 (0.4)	2.00	0.178	0.118

Table 2: Comment use by interface.

Task Performance

We analyzed three measures of task performance: overall time spent completing the tutorial, time spent on the tutorial not including time spent in the comments section, and task completeness. We determined time spent in the comment section from the video by using mouse position as a proxy for gaze, a technique validated in prior work [12].

Figure 7 (Left) displays the impact of condition on task time. On average, participants were faster overall with TaggedComments (21min 30s, s.e. 2min 14s) than with Default (24min 30s, s.e. 2min 5s), however, this difference was not significant ($F_{1,15} = 1.002$, $p = .333$, $\eta^2 = .063$). Considering only time spent working in the tutorial, participants spent less time on average with TaggedComments (18min 40s, s.e. 1min 50s) than with Default (21min 40s, s.e. 1min 50s), but this difference also did not reach significance ($F_{1,15} = 1.957$, $p = .18$, $\eta^2 = .16$). Participants completed roughly the same numbers of steps in each condition (Tagged: 6.7 steps, s.e. 0.3; Default: 6.3 steps, s.e. 0.3; $F_{1,15} = 1.634$, $p = .221$, $\eta^2 = .098$).

The lack of significant performance differences between the two conditions indicates that, at a minimum, promoting user comments did not interfere with participants' abilities to complete the tutorials. Given the large variability in completion times, however, it is also possible that we might start to see impacts of TaggedComments on performance with a larger number of participants.

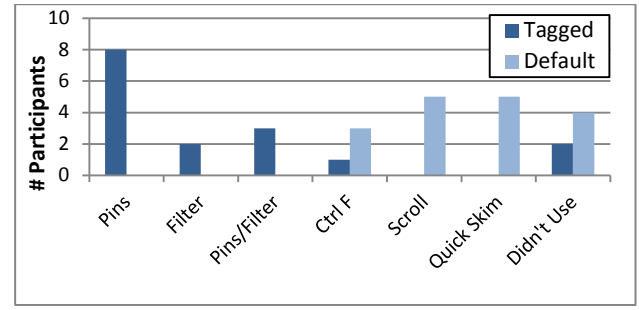


Figure 8: Strategies for finding relevant comments.

Comment Use

Using the video recordings and the log files, we analyzed the extent to which participants were relying on the comments in the two conditions and the strategies that they used to search through the comments.

Table 2 displays the following: the number of visits to the comment section (Visits), the total time spent in the comment section (Time), the number of comments participants appeared to read (Read), and the number of comments that participants attempted to apply to their task (Used). While there are no significant differences between the two conditions, the data suggest two possible trends. First, TaggedComments might have encouraged people to read through more comments than the Default interface. Second, participants might have tried to apply a larger number of comments to their task with TaggedComments than with Default. As with the impact of TaggedComments on task performance, further data would be needed to verify these trends.

Figure 8 summarizes the different strategies that participants used to find comments in the two interfaces. With TaggedComments, 14/16 participants visited the comments section. Participants' predominant strategy was to use the pins (8/14); however, filtering was also used (5/14). With the Default interface, only 7/16 participants made meaningful use of the comment section, with the remaining nine participants either not visiting the comments at all or scrolling through the comments quickly. The 7 participants who did look at the comments did so by either systematically scrolling through them or attempting to locate relevant comments using keyword search via Ctrl-F.

The above results suggest that a larger number of participants made a concerted effort to use comments as an information source when using TaggedComments than with the Default interface. That scrolling and quickly skimming comments were not used extensively also indicates that the features offered by TaggedComments provide participants with a viable alternative strategy for making use of comment data.

Participant Attitudes Towards TaggedComments

In the interviews, most participants expressed enthusiasm for the TaggedComments interface. For example, two par-

ticipants suggested that they would like to see the TaggedComments interface more widely adopted:

I'd like to see that more often -- I like that. I've never really thought about the old way until seeing that and it's like yeah -- the old way could use some work. [P10]

It seems really good. I would actually like to see that adopted more in tutorials. [P16]

Conversely, with the default design, a number of participants felt that the comments were close to unusable, as the following quote illustrates:

In the [Default interface] the comments were almost useless. [...] I thought that going through the comments [with the Default interface] would consume too much time. I [would] try to do the instructions two or three times first. [P5]

Comment/Integration: The above quotes were part of a larger theme of appreciation for the greater integration of comments and content that TaggedComments provides, which was seen to provide two complementary benefits. The first was the ability to retrieve relevant help in context:

With the pinned comments, it was not like I had to search through all of the comments, because I could just pick the comment [...] It was very useful for me. [P6]

Because usually if you are looking through the comments for help you are going to be looking for a certain step, not generic -- you're not looking for the generic feedback, you are looking for one thing that someone is commenting on. [P10]

Participants also liked being able to interact with the tutorial content and the comments simultaneously:

[I] liked that you could scroll in the comments without scrolling in the tutorial itself, so that you could see both at the same time, which was nice. [P11]

I like how it is side-by-side with the tutorial so I don't have to be scrolling down, scrolling back up -- it gets annoying. [P15]

Personalized Comment Use: Participants valued the ability to filter the information according to their needs either explicitly or using the colours (or icons) as visual guides. As expected given the experimental tasks, participants wanted to focus on the content-specific comments and to ignore social comments intended for the author of the tutorial:

I think that it's good that you differentiate 'appreciations' [from] 'how to do it'. In the [Default interface] there were too many comments about just appreciations, just thanking the author [...] You go through many comments thanking the author of tutorial and you cannot actually find the answer to what you are looking for [P2]

The colour coding makes it easy to tell which ones are actually going to be useful. Like the little 'yay thumbs up' stuff -- much easier when you can just skip over that. [P16]

Given that the percentage of appreciation comments in the tutorials we selected is representative of that in tutorials with comments on the web, this desire to isolate content-pertinent comments is likely to extend beyond these particular tutorials.

Visual Complexity: While most interview comments were in favour of the TaggedComments' design, two participants spoke to the additional visual complexity to the tutorial, though in both cases this was not cited as a major issue:

It is a little bit busier, so there's a lot more on your screen. [P10]

It was a little distracting, but maybe not that much. [P14]

It may be possible to adjust the design of the system to present relevant information more subtly; however, any design that promotes and integrates user comments will inevitably add some visual complexity.

Promoting a Sense of Community

To explore TaggedComments' potential role in promoting a sense of community surrounding the tutorial, we asked participants in the interviews to reflect on if and how the design might influence their willingness to post comments.

Participants spoke to two reasons for an increased willingness to contribute. The first was a desire to help others:

I think [I would post comments] because it would help others the way it helped me -- they would be able to learn the things they don't know and they won't have to use a lot of effort to learn that, it will be categorized and organized [P5]

One participant also felt his posts would be more likely to be seen and responded to:

I would probably feel that the odds of my question getting answered relative to what I'm asking would be higher. With the other format, it lends itself more to people saying 'oh well you should read the question better'. Whereas this one, I feel like -- and it might just be my perception -- because it is linking much more closely and much more organized, if people were actually trying to help, they would be more directed to do so. [P16]

In both cases, the manner in which TaggedComments' organizes contributions appears to be helping participants view user comments as a legitimate help resource, one that they are interested in contributing to.

Not all participants, however, felt that TaggedComments would encourage them to post comments. Two participants indicated that they often work with old tutorials and feel that there is little chance that their posts will be seen or responded to:

One of the big factors -- it is so silly -- but the comments are just so old. So of course there is nobody who is contributing anymore to the tutorial that's been there for 5 years on-line and nobody ever looks there or needs help anymore. [P11]

Generally the tutorials that I'm looking at are years old so people don't really respond back. [P15]

These comments show that users are picking up a sense of how active the user community around a tutorial is, and that this is influencing their choice of whether or not to contribute. This is encouraging for the general idea of improved systems for commenting and creating a sense of community around tutorials, and suggests that an indicator of the level

of community activity around a tutorial might be a useful addition to TaggedComments' interface.

DISCUSSION

In our study, TaggedComments had a positive impact on users' subjective experience of interacting with tutorial comments. Participants found comments significantly more helpful and easier to find with the TaggedComments interface, and they expressed a strong preference for its design over the status quo. These positive subjective impacts are encouraging on a number of fronts. Our interview data suggests that improved perception of comment utility might increase community participation, which could improve the quality of community help over time. Designs that display community feedback in a more attractive, accessible manner could also make completing tutorials more compelling, motivating more users to expand their skill sets.

Impact on Performance

While our study results suggest that TaggedComments improved the subjective experience of completing tutorials, we did not find a significant improvement in task performance. We believe that there are several reasons why this may have been the case.

First, conducting our study on the granularity of entire tutorials allowed us to investigate participants' experience of using TaggedComments in a realistic setting, but it also introduced noise into our performance measures. For example, some participants had trouble completing tutorial steps leading up to those that were intended to be challenging, and for which useful comments existed. When participants did reach these steps, some proceeded with incorrect results or skipped over the step entirely. A more tightly controlled study could address these issues by asking participants to complete a series of individual tutorial steps presented in isolation. This would eliminate the impact of difficulties leading up to a step.

Second, not all participants had sufficient Photoshop skills to complete the tutorials. As a result, some participants were able to find a tip in the comments, but did not have the necessary skills to apply it successfully. While it is difficult to accurately gauge knowledge of complex software such as Photoshop, we could try using a pre-screening sample task in place of self-reports of expertise.

Finally, our study did not hide the fact that tutorial comments were potentially useful, and that comment use was a feature being studied. This allowed us to elicit comparisons of the two systems from participants, but it may have also diminished the effect of a feature of TaggedComments -- that it reveals to users that comments contain valuable information. We suspect that more participants would have disregarded comments entirely in the control condition were they not aware that we were interested in measuring their use. This could be tested using a between-subjects study design, where the role of comments is not revealed to participants.

Generalizability

While we evaluated TaggedComments as applied to Photoshop tutorials, there are a number of aspects of the design of the system and findings from the evaluation that are likely to generalize beyond this particular application. First, Photoshop tutorials follow a format that is typical of many text/image-based tutorials for software use (i.e., step-by-step instructions followed by comments). The one caveat is that TaggedComments' image-carousel feature would need to be modified to provide previews for other media, such as audio. If one considers *video-based* tutorials, the notions of promoting, personalizing and integrating comments are also applicable, and there are existing analogues to pinning comments for time-based media. For example, SoundCloud provides facilities for users to post comments at particular points of time for streaming audio files, with these comments visible within the timeline. Beyond tutorials, we also envision applications of TaggedComments to other forms of long-form authored content, such as blog posts and news articles.

Community Participation

An open question with respect to the TaggedComments' design is whether users would be willing to tag their comments. While only a deployment could provide a definitive answer, we believe there are reasons to be optimistic. Our interviews suggest that users see value in tagging both for their own benefit and that of the community, and there are precedents in online communities that people are willing to annotate their content if they see value in doing so (e.g. YouTube video posters tagging with keywords). Wikipedia and StackOverflow [21] provide two further examples of communities where users align with cultural norms surrounding contribution. One could also imagine designs that explicitly reward tagging, such as promoting only tagged comments (increasing the likelihood of them receiving a response), or delaying the appearance of non-tagged comments for a certain period of time. Finally, one could explore supplementing user-supplied tags with automatically generated ones, following work on automatically identifying recipe refinements within user comments [4].

Alternative Models of Community Refinement

In a sense, TaggedComments allows people to become co-authors of a tutorial after the fact, while leaving the original tutorial content unchanged. This represents an interesting alternative to community refinement via wiki-style interfaces, where content continually evolves. A system like TaggedComments preserves author control and recognition, while still enabling the community to tweak and improve the content. This is particularly important in tutorials for content-generation applications such as Photoshop, where authors often use tutorials as a platform to demonstrate their skills and abilities. TaggedComments allows individuals to contribute to improving a tutorial's content in an organized manner, while allowing the tutorial author to maintain control over his/her work.

SUMMARY AND FUTURE WORK

TaggedComments introduces a novel way of incorporating community feedback within a tutorial. In TaggedComments, user comments are highly visible, tightly integrated with the tutorial content, and personalizable to the individual needs of both tutorial followers and authors. Our study revealed that this design led to greater perceived comment utility and was preferred to the status quo. Many users also responded positively to the increased sense of community surrounding the tutorial.

As future work, we plan to evaluate TaggedComments' role in aiding tutorial selection, as well its impact on user commenting quality and quantity. We are also interested in understanding the utility of the design from the tutorial author's perspective, including whether or not it helps focus their efforts, and their attitudes towards this model of community refinement. Finally, we intend to investigate the generalizability of the design and findings to other types of tutorials and long-form content.

ACKNOWLEDGEMENTS

This work was supported by the GRAND Network Centre of Excellence and the Natural Sciences and Engineering Research Council (NSERC).

REFERENCES

1. Chi, E.H. and Mytkowicz, T. Understanding the efficiency of social tagging systems using information theory categories and subject descriptors. In *Proc. HT*, (2008), 81–88.
2. Chi, P.P., Ahn, S., Ren, A., Dontcheva, M., Li, W., and Hartmann, B. MixT: Automatic generation of step-by-step mixed media tutorials. In *Proc. UIST*, (2012), 93–102.
3. Chilana, P.K., Ko, A.J., and Wobbrock, J.O. LemonAid: Selection-based crowdsourced contextual help for Web applications. In *Proc. CHI*, (2012), 1549–1558.
4. Druck, G. and Pang, B. Spice it up? Mining refinements to online instructions from user generated content. In *Proc. ACL*, (2012), 545–553.
5. Ekstrand, M., Li, W., Grossman, T., Matejka, J., and Fitzmaurice, G. Searching for software learning resources using application context. In *Proc. UIST*, (2011), 195–204.
6. Fernquist, J., Grossman, T., and Fitzmaurice, G. Sketch-Sketch Revolution: An engaging tutorial system for guided sketching and application learning. In *Proc. UIST*, (2011), 373–382.
7. Golder, S.A. and Huberman, B.A. Usage patterns of collaborative tagging systems. *Journal of Information Science* 32, 2 (2006), 198–208.
8. Grabler, F., Agrawala, M., and Erato, J.S.T. Generating photo manipulation tutorials by demonstration. In *Proc. SIGGRAPH*, (2009), 1–9.
9. Grossman, T., Fitzmaurice, G., and Attar, R. A survey of software learnability: Metrics, methodologies and guidelines. In *Proc. CHI*, (2009), 649–658.
10. Grossman, T., Matejka, J., and Fitzmaurice, G. Chronicle: Capture, exploration, and playback of document workflow histories. In *Proc. CHI*, (2010), 143–152.
11. Hart, S.G., California, M.F., and Staveland, L.E. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In P. Hancock and N. Meshkati, eds., *Human Mental Workload*. 1988, 139–183.
12. Huang, J., White, R.W., and Dumais, S. No clicks, no problem: Using cursor movements to understand and improve search. In *Proc. CHI*, (2011), 1225–1234.
13. Kelleher, C. and Pausch, R. Stencils-based tutorials: Design and evaluation. In *Proc. CHI*, (2005), 541–550.
14. Kong, N., Grossman, T., Hartmann, B., Fitzmaurice, G., and Agrawala, M. Delta: A tool for representing and comparing workflows. In *Proc. CHI*, (2012), 1027–1036.
15. Lafreniere, B., Bunt, A., Lount, M., and Terry, M. Understanding the roles and uses of Web tutorials. In *Proc. ICWSM*, (2013), 303–310.
16. Lafreniere, B., Grossman, T., and Fitzmaurice, G. Community enhanced tutorials: Improving tutorials with multiple demonstrations. In *Proc. CHI*, (2013), 1779–1788.
17. Laput, G., Adar, E., Arbor, A., Dontcheva, M., and Li, W. Tutorial-based interfaces for Cloud-enabled applications. In *Proc. UIST*, (2012), 113–122.
18. Lee, K.J. What goes around comes around: An analysis of del.icio.us as social space. In *Proc. CSCW*, (2006), 191–194.
19. Li, W., Grossman, T., and Fitzmaurice, G. GamiCAD: A gamified tutorial system for first time AutoCAD users. In *Proc. UIST*, (2012), 103–112.
20. Luther, K., Counts, S., Stecher, K.B., Hoff, A., and Johns, P. Pathfinder: An online collaboration environment for citizen scientists. In *Proc. CHI*, (2009), 239–248.
21. Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., and Hartmann, B. Design lessons from the fastest Q&A site in the west. In *Proc. CHI*, (2011), 2857 – 2866.
22. Matejka, J., Grossman, T., and Fitzmaurice, G. IP-QAT: In-product questions, answers & tips. In *Proc. UIST*, (2011), 175–184.
23. Willett, W., Heer, J., Hellerstein, J.M., and Agrawala, M. CommentSpace: Structured support for collaborative visual analysis. In *Proc. CHI*, (2011), 3131–3140.