

Improving Warning Messages Adherence: Can Maya Security Bot Advisor Help?

ABSTRACT

Despite several different methods that have been suggested to improve users' adherence to warning messages, users are still left alone in their decision-making process and have to deal with possible negative consequences of their decisions. To improve warning message adherence and to help users make a more informed decision, we implemented a real-time security bot advisor Maya.

The results suggest significant improvement in the warning adherence (from 38% to 61%) with the highest effect on older users. We demonstrated that the security bot advisor can be of a great help in influencing a user's decision-making process. We suggest several ways that our findings can be used by warning designers and provide new directions for future research.

Author Keywords

computer warning; security bot; information security

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; See <http://acm.org/about/class/1998> for the full list of ACM classifiers. This section is required.

INTRODUCTION

Similar to real-world situations (e.g., traffic lights), in the digital realm, individuals are confronted with computer warnings where they are asked to make a security decision. Computer warnings represent security communication informing users about potential hazards that may occur, if the user continues with his/her action [1]. The user's action usually corresponds to a binary decision, continue or exit, where the end result can have important security implications. It can eventually lead to a loss of information and data or a breach of user's privacy information. In this human-computer interaction process, users are considered to be the weakest link, since they often ignore the warning messages and continue with their actions.

Previous research has explained that this user's behavior is due to habituation [2, 3], where users simply do not read

warning messages and, therefore, ignore and skip them [e.g. 2, 4]. Researchers have recently used polymorphic warnings [5] which could lead to an habituation decrease and multi-modal neuro-physiological measures [6] to suggest that users are reading, comprehending, and eventually heeding the warning message content. Although research has improved the current understanding of what drives habituation [7] and what the mental model process looks like during the decision-making process [8], very few studies have tried to propose solutions to increase user's adherence and higher compliance to security warning messages. The current level of adherence is relatively low: only one third of Google Chrome users showed the compliant behavior (adhered to the warning message) [9].

A study done by [10] suggested that SSL warning message comprehension and adherence can be improved through opinionated design, where security is leveraged with visual cues. Several other studies [e.g., 2, 11-16] suggested different ways to improve adherence by mostly improving warning message design elements.

However, the problem with all of these improvements is the fact that the user is still left alone in his/her decision-making process. Regardless of the visual cue improvements, the decision is still purely static and based on the limited information that the user receives through warning content or visual cues.

We argue that in order to improve the user's adherence, additional help is needed so that the user can be better advised, informed or given time to think about the consequences of his/her actions.

We set out to design a Security Bot (Chat Bot) as part of the warning message. The objective was to engage the user in conversation through textual means, in order to provide a real-time conversational support to the user. We used SIML (Synthetic Intelligence Markup Language) for this purpose in order to build the security knowledge base for our Security Bot (named Maya).

Our newly designed warning message with security bot was tested with 312 users through a field experiment. The results showed a significant improvement in adherence for the warning with security bot (61%) compared to the warning without security bot (38%) - control group.

BACKGROUND

Chat Bot

In order to build a Security Bot, we used SIML (Synthetic Intelligence Markup Language) through the Syn Bot Framework which provides an interpreter (SimlBot) for

Paste the appropriate copyright/license statement here. ACM now supports three different publication options:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single-spaced in Times New Roman 8-point font. Please do not change or modify the size of this text box.

Each submission will be assigned a DOI string to be included here.

SIML. The Syn Bot is running on .NET Framework 4.5 that we used to develop the Security Bot. It is a simple chat bot software that provides a conversational feature during the interaction with the user.

In order for Syn Bot to work, a knowledge base has to be constructed. The knowledge base can provide both simple and complex interactions, such as complex pattern recognitions or conditional patterns.

The overall architecture consists of two principles: concepts and models. A concept represents a topic under which models are declared. A model is the basic unit of knowledge composed of a pattern and a response. In Figure 1, a simple example is provided where to the pattern “Hello bot”, bot would reply to the user “Hello User!”

```
<Siml>
  <Concept Type="public" Name="Hello Bot">
    <Model>
      <Pattern>HELLO BOT</Pattern>
      <Response>Hello User!</Response>
    </Model>
  </Concept>
</Siml>
```

Figure 1. Example of pattern/response interaction

In this research, our goal was to implement an as simple as possible interaction between the bot and the user, so that the user does not spend too much time in the communication process with the bot. We wanted to avoid the security fatigue which was found to be “one cost that users experience when bombarded with security messages, advice, and demands for compliance—and that this cost often results in what security experts might consider less secure online behavior” [17, p.27].

We, therefore, kept the interaction simple and started by explaining Maya’s (our Security Bot name) purpose “Hi I am Maya, I am here to help you with your security decision “, and then asking “Can you tell me if your intention was to click on 'Continue' or on 'Exit' button?”.

If the reply was ‘Exit’, then we would simply confirm that it is a good choice. If the reply was ‘Continue’, we would then engage in the conversation trying to advise the user on the best decision to take. For example, we asked user “OK, let me ask you a few more questions. Does the web link look authentic?”

We did not want to influence the user by providing any recommendations, but rather tried to influence the ‘habituation’ and ‘security fatigue’ effects by suggesting where and what to look for, so the user could think again before making the decision.

Our final implemented Security Bot is depicted in Figure 2.



Figure 2. Maya – Security Bot implementation

Security Warnings

Warning adherence was investigated from a number of different perspectives. It was demonstrated that the color of the warning message is an important factor that is mostly neglected, since all warnings are the same in existing web browsers across all cultures [18, 19]. In this context, users prefer to simply skip the warning message and not read it at all. The Firefox SSL warning was ignored by 68% of users [20]. Similar behavior was observed for the Internet Explorer 7’s SSL warning message, where more than one third of users ignored the warning and revealed their confidential data [21].

Other studies received mixed results. For example, users of Mozilla Firefox showed 70% adherence rates for SSL warnings, while the same was not confirmed for Google Chrome warnings which only had 30% rates [9]. Another suggestion to improve adherence was found to be related to the thermal feedback [22]. Results suggest that users may associate certain temperatures (cold or warm) with a secure or insecure page. In order to improve security warning adherence, opinionated design [10] and attractors [23] were recently introduced. Although both methods did show some encouraging results in improving user’s adherence, we believe that the user’s challenge of ‘being alone’ in front of the warning message still holds and could back-fire in the long run, since the user will most likely switch back to habituation and ignorance.

ADDING SECURITY BOT TO WARNING

Design

To build the new warning message, we relied on the existing warning look and feel from Google Chrome (version 42)




Figure 3. Warning Message (Google Chrome v42)

To keep the whole design simple, we simplified the decision buttons (continue and exit) so that they have the same design. For the warning content, we also removed any non-technical words (e.g., malware) that may confuse novice users and may not be understandable. We focused instead on the malicious web site context, informing the user that the upcoming site is not safe (Figure 3).

This is consistent with recommendations from past studies, which suggest that warnings should use simple and non-technical jargon [24]. Therefore, we removed all technical jargon and used simple language. This reduced the need to deal with comprehension challenges, since our overall objective was to influence adherence.

Security Bot Integration with Warning Message

As we built the knowledge base, the next step was to implement the Security Bot (Maya) into the warning message.

To accomplish this step, we created an experiment where we combined the Security Bot with the warning. We used the existing Syn.Bot library that we adapted to the ASP.Net environment. The Security Bot was a modal dialog pop-up window that would appear when the warning message page was presented to the user (Figure 4).




Figure 4. Warning message with Maya Security Bot

The user, therefore, had the option whether or not to interact with the Security Bot. Every user's interaction as well as all pattern/response communication was registered in our backend database.

METHODOLOGY

We used field experiments as our research method. After receiving IRB approval, we proceeded with the experiment.

We used participants from Amazon mturk, who were compensated for the task they had to accomplish. Although some studies criticized mturk uses for surveys, recent research argues that not only can high quality data collections can be achieved with mturk but also that there are several advantages in doing so [25].

We used the deception method to increase the study's realism. We proceeded as follows. Participants, after choosing the task, were told that the purpose of the study was to do a search task. The search task consisted of finding the answer to a question that was located in the text (Figure 5). We purposely made the task more difficult, since it demanded that each participant carefully read the text (22 answers were rejected since the participants provided the wrong answers). We did this because we wanted to minimize suspicion about the study's real objective and, therefore, increase the study's realism.

CIOs vastly underestimate extent of shadow IT

New study from Cisco reveals CIOs are operating in the dark, failing to meet end users' needs for cloud applic services.



Credit: Thinkstock


By Kenneth Corbin | Follow
CIO | Aug 10, 2015 5:47 AM PT

RELATED TOPICS

CIO Role
IT Strategy

Most CIOs have an inkling that employees in their enterprise have snuck a few applications past the IT department, but a new study by Cisco indicates that they are vastly underestimating the extent that unauthorized apps and services have infiltrated the enterprise.

MORE LIKE



on 100 Answers:
How is data stc
cloud??

Figure 5. Search task

Participants were instructed to click on the submit button in order to proceed to the survey site. At this point, the warning message (Figure 4) was displayed either with the Maya Security Bot or without it (control box).

In the background, our application tracked all of the user's actions through following variables: 1) time - we measured the time needed to make a decision (before continue or exit is clicked); 2) decision: 0 or 1 which corresponded to the 'Continue' or 'Exit' action. We have also tracked the users that clicked on the 'back' button in the browser. In this case, users were directly taken to the survey site. However, we did

not track users that closed the internet browser, due to the technical limitations of capturing such events.

In both cases (clicked on continue/exit or back button) users were taken to the survey site, where we asked them several questions to better understand their decisions and their backgrounds.

Finally, our application was hosted on the Amazon AWS site and had the following link extension xxx-dev.us-west-2.elasticbeanstalk.com. We believe that the link is appropriate for the study, since it can easily raise suspicions and authenticity questions, concerning its validity. However, for advanced users, it would be sufficient to test the domain name (elasticbeanstalk.com) in the web browser to quickly see that the link is authentic and valid and takes user to the authentic Amazon web site.

RESULTS

We received answers from a total of 349 participants. We removed 10 answers for implausible answer times (less than two seconds to make the decision or less than one minute to fill in the survey). We also removed 22 participants, since they did not provide the correct answers to the search task questions and we considered them to be invalid answers. Our final sample accounted for 312 answers.

Our participants were mainly from India (30.4%), the U.S. (23.7%) and other countries (45.9%). The average age was 30.8. We had 184 men and 128 women that completed the experiment and the survey. The average computer experience was 10.5 years and the education level was 13.5 years.

The results can be found in Table 1 where ‘without’ and ‘with Security Bot’ use results are displayed.

We can see that 38% of the users clicked on the Exit button, when they did not use the Security Bot. The corresponding result for the ‘with bot’ condition increased to 61%. Similar results but in the opposite direction can be observed for the Continue action. From 62% for without Security Bot use, users decided to continue only in 39% of cases when the Security Bot was used as the security advisor.

Decision	Without BOT		With BOT	
	in %	N	in %	N
Exit	38%	56	61%	100
Continue	62%	90	39%	65
Total		147		165

Table 1. Results without and with Security Bot use

This first results indicate a positive influence of the Security bot on a user’s decision, since we can see an improvement in a user’s adherence to the security warning message by 23 points (from 38% to 61%) in the presence of the security advisor Maya ($\chi^2 = 15.3$, $p = 0.00009$, the result is significant at $p < .05$).

We also wanted to understand if any differences were observed between the male and female population. In Table 2, the gender results are summarized.

We can see that improvement to warning adherence is much more pronounced for males (from 32% to 61%) compared to females (from 47% to 60%). However, the Z-Score calculations did not show any statistical difference between the male and female results.

Decision	Without BOT		With BOT	
	in %	N	in %	N
Exit	32%	30	61%	54
Continue	68%	65	39%	35
Total		95		89

Decision	Without BOT		With BOT	
	in %	N	in %	N
Exit	47%	24	60%	46
Continue	53%	28	40%	30
Total		52		76

Table 2. Gender results without and with Security Bot use

We also wanted also to understand if experience had any effect on the adherence rates for both groups: with or without Security Bot use.

In Table 3, we split users experience in two categories: less than five years and more than or equal to five years. We can see that the positive effect of the security bot is more pronounced on the more experienced users, where exit action was chosen by 51% of participants compared to 41% for less experienced users.

This is somewhat unexpected behavior, since we were expecting less experienced users to follow more security bot advisor suggestions than more experienced users. However, we must caution against the relatively low number of participants with <5 years’ experience. However, Z-score calculations did not show any statistical difference between the two experience results.

Decision	Without BOT		With BOT	
	in %	N	in %	N
Exit	33%	12	41%	20
Continue	67%	24	59%	28
Total		36		48

Decision	Without BOT		With BOT	
	in %	N	in %	N
Exit	39%	43	51%	60
Continue	61%	68	49%	57
Total		111		117

Table 3. Experience results without and with Security Bot use

Finally, we also report the results based on the user’s age (Table 4). It shows that age has an important effect for

population older than 30 years, where adherence reaches 76% when security bot is used compared to 39% when it is not. For population younger than 30 years, there is still some effect (from 36% to 44%), but it is much weaker than for the population over 30 years old.

This is a very interesting finding, since it suggests that the older population (over 30 years old) is more sensitive to receiving security advice from a security bot advisor (Z-Score is 4.1462, significant at $p < 0.05$).

<u>< 30 years</u>	Without BOT		With BOT	
Decision	in %	N	in %	N
Exit	36%	26	44%	36
Continue	64%	46	56%	44
Total		72		80

<u>>= 30 years</u>	Without BOT		With BOT	
Decision	in %	N	in %	N
Exit	39%	29	76%	65
Continue	61%	46	24%	20
Total		75		85

Table 4. Age results with and without Security Bot use

In the post-experiment survey, we asked participants to provide some answers about their perceptions on using the security bot.

Can the security bot (Maya) help you to make a better security decision in the future?

When asked this question, 76% of participants confirmed that the security bot advisor Maya could be helpful in the future. One user commented in the open-ended area of the survey “*Maya will be helpful, if the site does not look genuine*”. Another one added “*Maya helped me to think about it*”.

When asked if the security advisor made them spend more time in making their security decision, 59% replied yes. Interestingly, 43% of the participants confirmed that they are able to recognize a malicious web site or link that is presented to them. We think that this is a high number of users who are not capable of identifying which web site could be potentially harmful. In this context, using Maya as the security advisor could be even more important and interesting for this population of users.

Finally, when asked if they care about possible negative consequences (e.g. loss of disk drive) when clicking on potentially malicious links or visiting malicious web sites, 86% confirmed that they do not really care about the negative outcomes. However, 75% of the participants confirmed that by continuing their actions (e.g., surfing potentially malicious web sites) can be dangerous.

This suggests that users, despite the fact that they are aware of the danger ahead, will most likely continue with their action, because they do not really care about the possible negative

outcomes. One possible interpretation could be that they are not fully aware, educated or properly informed about the real consequences and how serious they can be.

DISCUSSION

To the best of our knowledge, this is the first study that proposed and implemented a real-time security advisor able to advise users on their security decision in the warning message context. We not only provide an informed decision for the user, but we also guide the user in his/her decision-making process by providing a real-time risk assessment done by the users themselves. The habituation and ignorance effects are mitigated through the user’s communication with the security bot.

We found that security bot advisor Maya considerably improved warning message adherence (from 38% to 61%), when compared to those that used the bot against the control group. This is an important finding, since it suggests that users are in need of real-time security advice that would help them to better identify the security risks. Security advisor clearly seems to have influenced the user’s decision-making process, by asking them the right questions. It is possible that by being engaged in this human-bot communication process, users start to think about their previously formed decision, since they have now received some new informational clues to think about.

This is also confirmed through the survey, since 76% of the participants found security advisor Maya to be helpful in making a security decision.

This suggests that a possible direction to improve warnings adherence would be to not leave users alone in their decision-making process. They should instead be provided with an alternative course of action that engages them in the real-time human-bot communication. A similar process was implemented in other contexts such as the social media, where it was found that chatbot [26] helps users to cope with emotional situations as well as human agents.

We believe that one possible implication for practitioners could be the implementation of a security advisor into the existing web browsers, where users could seek additional help if they do not feel comfortable with making the decision. One possible drawback of this optional security advisor, could be that users may not want to spend additional time interacting with the bot. However, this could be tackled by creating short and simple interactions where, for example, in a maximum of four interactions, the user would get the right security advice or where the security advisor would provide guidance so the user would be able to better understand the challenge and the risk ahead.

From the restudy results, we saw that older users (>5 years) are showing higher adherence to warnings when compared to less advanced users. It is possible that their security fatigue is already at high levels. When they see the security advisor,

it creates a sort of mental interruption. This finding is in line with their computer experience, where more experienced users (>5 years) show similar trends. We don't have a complete understanding of why this is happening, since we were expecting to see younger and less advanced users be more sensitive and influenced by the security advisor.

Interestingly, we saw that females are much less affected by the security advisor compared to males. This is also a surprising result, since we were expecting to see at least equal results. More research would be needed in this area to better understand the motives behind this result.

With this research, we advance our current understanding of how to better cope with warning message adherence by implementing a new way of providing real-time security support to users through a security bot implementation.

There are several ways that future research can further extend our findings. First, we built the knowledge base specific to the warning message context. It would be interesting to further extend this context to other security situations to understand how and if security advisor can help in other areas of information security. Second, more qualitative studies would be welcome to understand, why, for example, females are less sensitive to security advisor help than males. It would also be helpful to understand why more experienced and older users were more influenced by the security bot advisor than less experienced and younger users. We believe that to fully explore the security bot advisor potential to influence the user's decision-making processes, more research is needed that incorporates various psychometric variables that would further clarify the role and importance of the security bot advisor. Finally, our study was a one-time investigation of the security advisor influence on user's decision-making processes. It would be interesting to understand how the security advisor influences the user's decision over a longer period of time. For example, would the security advisor have the same effect when the user sees the security advisor a second time?

Our study is also limited by several factors. We did not use the real malicious links, but rather the AWS website link, which some users could easily recognize. We attempted to address this by removing several users for implausible response times. We used mturk participants who could have been already engaged in similar studies. However, we do not believe this was the case. We also explained why we believe that mturk participants are a good population to use in a research study. In some cases, they are even better to use than some other general internet user samples.

CONCLUSION

We proposed and evaluated a new way of improving warning adherence through the implementation of a real-time security bot advisor. Our solution showed important improvement in the warning adherence (from 38% to 61%), demonstrating that a security bot advisor can be a great help in influencing the users' decision-making process.

REFERENCES

1. Silic, M., Barlow, J., Ormond, D.: Warning! A Comprehensive Model of the Effects of Digital Information Security Warning Messages. The 2015 Dewald Roode Workshop on Information Systems Security Research, IFIP, pp. 1-32. IFIP, Dewald (2015)
2. Sunshine, J., Egelman, S., Almuhiemedi, H., Atri, N., Cranor, L.F.: Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In: USENIX Security Symposium, pp. 399-416. (Year)
3. Anderson, B., Vance, T., Kirwan, B., Eargle, D., Howard, S.: Users Aren't (Necessarily) Lazy: Using NeuroIS to Explain Habituation to Security Warnings. (2014)
4. Akhawe, D., Felt, A.P.: Alice in Warningland: A large-scale field study of browser security warning effectiveness. Proceedings of the 22nd USENIX Conference on Security, pp. 257-272. USENIX Association (2013)
5. Anderson, B.B., Kirwan, C.B., Jenkins, J.L., Eargle, D., Howard, S., Vance, A.: How polymorphic warnings reduce habituation in the brain: Insights from an fMRI study. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 2883-2892. ACM, (Year)
6. Neupane, A., Rahman, M.L., Saxena, N., Hirshfield, L.: A multi-modal neuro-physiological study of phishing detection and malware warnings. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 479-491. ACM, (Year)
7. Anderson, B.B., Vance, A., Kirwan, C.B., Eargle, D., Jenkins, J.L.: How users perceive and respond to security messages: a NeuroIS research agenda and empirical study. European Journal of Information Systems 25, 364-390 (2016)
8. Bravo-Lillo, C., Cranor, L.F., Downs, J., Komanduri, S.: Bridging the gap in computer security warnings: A mental model approach. IEEE Security & Privacy 9, 18-26 (2011)
9. Akhawe, D., Felt, A.P.: Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In: Usenix Security, pp. 257-272. (Year)
10. Felt, A.P., Ainslie, A., Reeder, R.W., Consolvo, S., Thyagaraja, S., Bettess, A., Harris, H., Grimes, J.: Improving SSL warnings: Comprehension and

- adherence. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 2893-2902. ACM, (Year)
11. Trommelen, M.: Effectiveness of explicit warnings. *Safety Science* 25, 79-88 (1997)
 12. Silic, M.: Understanding colour impact on warning messages: evidence from us and India. In: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 2954-2960. ACM, (Year)
 13. Silic, M., Silic, D., Oblakovic, G.: Restrictive Deterrence: Impact of Warning Banner Messages on Repeated Low-trust Software Use. 18th International Conference on Enterprise Information Systems (ICEIS 2016), vol. 2, pp. 435-442. SCITEPRESS (2016)
 14. Silic, M., Silic, D., Oblakovic, G.: The Effects of Colour on Users' Compliance with Warning Banner Messages Across Cultures. ECIS 2016, Istanbul (2016)
 15. Silic, M., Cyr, D.: Colour Arousal Effect on Users' Decision-Making Processes in the Warning Message Context. In: International Conference on HCI in Business, Government and Organizations, pp. 99-109. Springer, (Year)
 16. Felt, A.P., Reeder, R.W., Almuhiemedi, H., Consolvo, S.: Experimenting at scale with google chrome's SSL warning. In: Proceedings of the 32nd annual ACM conference on Human factors in computing systems, pp. 2667-2670. ACM, (Year)
 17. Stanton, B., Theofanos, M.F., Prettyman, S.S., Furman, S.: Security Fatigue. *IT Professional* 18, 26-32 (2016)
 18. Silic, M., Njavro, M., Oblakovic, G.: Understanding Color Risk Appropriateness: Influence of Color on a User's Decision to Comply with the IT Security Policy—Evidence from the US and India. In: International Conference on HCI in Business, Government, and Organizations, pp. 412-423. Springer, (Year)
 19. Silic, M., Cyr, D., Back, A., Holzer, A.: Effects of Color Appeal, Perceived Risk and Culture on User's Decision in Presence of Warning Banner Message. (2017)
 20. Dhamija, R., Tygar, J.D., Hearst, M.: Why phishing works. In: Proceedings of the SIGCHI conference on Human Factors in computing systems, pp. 581-590. ACM, (Year)
 21. Schechter, S.E., Dhamija, R., Ozment, A., Fischer, I.: The emperor's new security indicators. In: Security and Privacy, 2007. SP'07. IEEE Symposium on, pp. 51-65. IEEE, (Year)
 22. Wilson, G., Maxwell, H., Just, M.: Everything's Cool: Extending Security Warnings with Thermal Feedback. In: Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 2232-2239. ACM, (Year)
 23. Bravo-Lillo, C., Komanduri, S., Cranor, L.F., Reeder, R.W., Sleeper, M., Downs, J., Schechter, S.: Your attention please: designing security-decision UIs to make genuine risks harder to ignore. In: Proceedings of the Ninth Symposium on Usable Privacy and Security, pp. 6. ACM, (Year)
 24. Wogalter, M.S., Conzola, V.C., Smith-Jackson, T.L.: based guidelines for warning design and evaluation. *Applied ergonomics* 33, 219-230 (2002)
 25. Lowry, P.B., D'Arcy, J., Hammer, B., Moody, G.D.: "Cargo Cult" science in traditional organization and information systems survey research: A case for using nontraditional methods of data collection, including Mechanical Turk and online panels. *The Journal of Strategic Information Systems* 25, 232-240 (2016)
 26. Xu, A., Liu, Z., Guo, Y., Sinha, V., Akkiraju, R.: A New Chatbot for Customer Service on Social Media. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, pp. 3506-3510. ACM, (Year)