

A Design Science Artefact for Cyber Threat Detection and Actor Specific Communication

Research-in-progress

Kaan Eyilmez

Virtimo AG
Berlin, Germany
Email: kaan.eyilmez@virtimo.de

Ali Sercan Basyurt

Department of Computer Science and Applied Cognitive Science
University of Duisburg-Essen
Duisburg, Germany
Email: ali-sercan.basyurt@uni-due.de

Stefan Stieglitz

Department of Computer Science and Applied Cognitive Science
University of Duisburg-Essen
Duisburg, Germany
Email: stefan.stieglitz@uni-due.de

Christoph Fuchss

Virtimo AG
Berlin, Germany
Email: fuchss@virtimo.de

Marc-André Kaufhold

Department of Computer Science
Technical University of Darmstadt
Darmstadt, Germany
Email: kaufhold@peasec.tu-darmstadt.de

Christian Reuter

Department of Computer Science
Technical University of Darmstadt
Darmstadt, Germany
Email: reuter@peasec.tu-darmstadt.de

Milad Mirbabaie

Department of Information Systems
Paderborn University
Paderborn, Germany
Email: milad.mirbabaie@upb.de

Abstract

Over the past few decades, the number and variety of cyberattacks and malware patterns have increased immensely. As a countermeasure, computer emergency response teams were established with the responsibility of securing the cyber environment. However, recent studies revealed that currently performed manual processes and the unavailability of adequate tools impede the achievement of cybersecurity. To address these challenges, we followed the Design Science paradigm to develop an artefact that improves the evaluation of open-source intelligence obtained from Twitter as well as the actor-specific communication of cyber threat information. Subsequently, the implemented artefact will be evaluated through semi-structured interviews with subject matter experts. This research in progress article presents the identified research gap and describes the development process and the endeavor to contribute to the cybersecurity domain theoretically with design principles for the development of an instrument and practically by implementing an artefact that supports domain experts in their work.

Keywords Computer Emergency Response Teams, Cybersecurity, Event Detection, Communication, Design Science Research

1 Introduction

Our dependency on information and communication technologies as well as on interconnected networks has increased significantly (Kaufhold et al. 2021). This interconnectedness offers many benefits, however, it also leads to a significant increase in cyber-attacks (Sillaber et al. 2016, Kaufhold et al. 2021). Consequently, ensuring the security and privacy of the data of individuals, organizations and companies became one of the biggest challenges in this current era of digital information technologies (Conti et al. 2018). As a countermeasure, Computer Emergency Response Teams (CERT) were established to adequately protect IT infrastructures against cyberattacks. The responsibilities of these CERTs include monitoring and analyzing cybersecurity events and raising awareness by distributing warning messages to distinct stakeholders (Basyurt et al. 2022). In their current cyber threat detection workflow, CERTs consider various closed and open sources, which include social media platforms such as Twitter. Although Twitter has proven to be useful for achieving cybersecurity, its potential is currently not being fully utilized. Recent interviews with CERTs revealed that only the information provided by subject matter experts is considered, as the inclusion of more data would necessitate more resources. Due to this omission of potentially relevant information, the study concluded that a novel system is required to address this issue (Basyurt et al. 2022). Additionally, another frequently mentioned area for improvement is the communication of Cyber Threat Intelligence (CTI) through warning messages. Research and practice have identified a great demand for the exchange of information, data and knowledge between cybersecurity stakeholders to appropriately manage vulnerabilities and threats (Sillaber et al. 2016, Basyurt et al. 2022). However, CTI communication is not trivial, and reaching the various affected stakeholders has been described as a challenge by CERTs (Basyurt et al. 2022). The process is complicated by the varying IT knowledge of different target groups, and stakeholders from a non-technical background often do not have the technical skills to fully comprehend the distributed information (Ask et al. 2021). In addition, simply providing information does not change the behavior of non-experts for the benefit of their cybersecurity (de Bruijn and Janssen 2017) and consequently, not only the terms used but also the wording must be adjusted according to the IT knowledge of the target groups. Furthermore, the current warning message generation process is mostly done manually and is very labour-intensive, emphasizing the importance of a certain degree of automation (Wagner et al. 2019) that is not sufficiently provided by the available information sharing software (Sauerwein et al. 2017, Basyurt et al. 2022). The unavailability of such tools and the lack of research on the actor-specific communication of cyber warning messages reveals a research gap that we aim to address with the purpose of improving the handling of cyber threats.

For this purpose, we utilized the Design Science Research (DSR) process described by Peffers et al. (2007) to develop a novel tool that supports CERTs in both the assessment of Open-Source Intelligence (OSINT) obtained from Twitter and in the actor-specific communication of CTI. To identify relevant cybersecurity events on Twitter in real-time, similar to the method of Sceller et al. (2017), we applied the Locality Sensitive Hashing (LSH) approach proposed by Petrović et al. (2010) to cluster related posts and thereby detect cybersecurity-related events. To improve the present CTI communication, the proposed approach of Basyurt et al. (2022) was utilized who suggested a system of prewritten text fragments that are selectable with corresponding keywords, with the additional option to manually customize the generated warning messages. Additionally, different communication channels such as social media and email are provided to ensure that the generated warning messages can be disseminated appropriately. To be able to determine whether the developed technical solution would facilitate the current workflows, semi-structured interviews with cybersecurity experts will be conducted. The results will subsequently be used to derive design principles for future technical solutions for social media event detection and CTI communication in the context of cybersecurity. Hence, the following two research questions will be answered with this study:

RQ1: How can a holistic tool be designed to support cybersecurity professionals to automatically detect relevant events on Twitter and disseminate actor-specific warning messages?

RQ2: How should cybersecurity warning messages be designed to appropriately reach stakeholders?

By addressing these research questions, a theoretical contribution is made in the form of design principles to develop a holistic tool to support cybersecurity professionals and their needs. Furthermore, by addressing their needs with an artefact that supports experts in their work a practical contribution is made. The remainder of this paper is structured as follows: In the next section, the related work on cybersecurity management, CERTs and the current state of event detection on social media and CTI communication is explained. Subsequently, information on the applied DSR method and the technical approaches for the development of the holistic tool are explained. Afterward, we will conclude by specifying the expected results as well as the theoretical and practical contributions.

2 Related Work

2.1 Computer Emergency Response Teams

Cybersecurity encapsulates the ambition of governments and organizations to achieve a secure cyberspace for everyone. However, ensuring cybersecurity is becoming increasingly difficult due to the growing number and variety of cyberattacks (Conti et al. 2018). Based on an in-depth analysis of real-world data breaches at 500 companies, the cost of a cybersecurity breach averages \$4.24 million per incident, the highest cost in 17 years (IBM 2021). As a countermeasure against cyber threats, CERTs were established to appropriately protect IT infrastructures. CERTs can take many forms but often consist of cybersecurity experts who determine the response to cybersecurity breaches, forensic investigators who gather evidence for potential legal actions, and engineers who maintain the specialized technologies of the operation centers (Horne 2014, Agyepong et al. 2020). Their main objective is to contain and minimize the damage caused by cybersecurity incidents, to ensure effective recovery and response, and prevent future incidents (Basyurt et al. 2022). The literature revealed that challenges CERTs face are often results of the intensity and complexity of their ambition (Zhong et al. 2018). As a result, several efforts have been made to understand and support the daily work of CERTs, however, it remains unclear whether researchers obtained a comprehensive insight into the domain (Agyepong et al. 2020). Currently, the workflow generally can be divided into three major activities. Starting with the collection of data from various sources, followed by the analysis of the cyber situation, and ending with the communication of the identified cyber threats. Recent results revealed several challenges in the current workflow (Basyurt et al. 2022). For instance, the amount of data to be considered is a major challenge, especially when analyzing social media data that contains a lot of redundant information and consequently is examined to a very limited extent. However, the inclusion of more data would require more resources that may not be available to CERTs. Consequently, the detection and analysis phases were perceived as areas for potential research (Saudi et al. 2016). Additionally, the communication of CTI through warning messages was seen as worthy of research and improvement. Cybersecurity experts described reaching specific audiences with effective warning messages as a major problem since the IT knowledge of the various stakeholders, such as governments, organizations, individuals, municipalities or authorities, are different and must be considered in the formulation (Basyurt et al. 2022). Furthermore, the generation is also labour-intensive as the process is mainly done manually (Wagner et al. 2019). Consequently, a technical improvement was perceived as necessary in both areas (Basyurt et al. 2022). To address these challenges, the implemented artefact includes an event detection module for Twitter and a communication module with customizable, prewritten text fragments.

2.2 Social Media Event Detection

Social media has become an important driver for the acquisition and dissemination of information in a variety of fields (Stieglitz et al. 2018) including the domain of cybersecurity (Sceller et al. 2017). It was revealed in previous studies that social media platforms can provide valuable insight into evolving cyber threats and signal imminent attacks before the malicious activity can be detected by the target system. For instance, Sabottke et al. (2015) found that computer exploits are often discussed on Twitter before they are publicly disclosed. This was also the case with the "Petya/NotPetya" global ransomware outbreak, which was discussed extensively on Twitter before it was reported in the mainstream media (Le et al. 2019). In addition, denial-of-service attacks are often first reported by the users of the website or the services that are attacked (Sceller et al. 2017). However, due to the high number and dynamics of users combined with the technical requirements the extraction of relevant information from social media during critical events is very complex and multifaceted. One feasible approach for this purpose is to identify the first story discussing an emerging event, which is also referred to as event detection (Petrović et al. 2010). The benefits of event detection, such as the advantage of an immediate response to a possible threat, was discussed in previous studies (Sceller et al. 2017) making this method suitable to support filtering and aggregation of shared content and detect relevant events. While there are multiple approaches for event detection, the selection of the appropriate method depends on the objectives. To address present challenges, we selected the LSH as an adequate approach following the example of Sceller et al. (2017), who successfully utilized the method to identify cybersecurity events on Twitter. Subsequently, the developed event detection was connected to the communication module.

2.3 Cyber Threat Communication

The increasing numbers of cyber threats require new forms of cross-organizational information sharing to be able to counter attacks adequately in an early stage (Skopik et al. 2016). The process of sharing knowledge and analyzing the cyber threat landscape is currently mostly done on CTI platforms, while other affected stakeholders who do not have access to the sharing platforms are contacted via mail or

specific apps provided by the organization (Basyurt et al. 2022). In general, not much attention has been paid to the communication of cyber threats, and the nature and quality of the conducted studies vary. In particular, human-to-human communication has received insufficient consideration within the existing literature (Ask et al. 2021). Currently, the communication of cyber threats is one of the major challenges for CERTs (Agyepong et al. 2020). One possible reason is that the present sharing process heavily relies on manual input and is therefore very labour-intensive and the information sharing platforms currently in use offer limited automation capabilities and mechanisms (Sauerwein et al. 2017). Similarly, three out of seven interviewed cybersecurity experts described the manual effort required to post a warning message as demanding (Basyurt et al. 2022), particularly when the target audience of the warning message belongs to a non-technical sector and lacks the technical skills to fully comprehend the shared information (Ask et al. 2021). In addition, interviewed cybersecurity experts described reaching specific target groups as difficult since not all CERTs use multiple communication channels. For instance, it was identified that the dissemination of warning messages on social media is often unexploited, resulting in the underutilization of the potential of the platforms (Basyurt et al. 2022).

Despite the challenges and need for CTI exchange, insufficient work has been conducted to develop standardization of exchanges. One reason for the limited tools is that CTI sharing and automating the process are still relatively new as research disciplines (Wagner et al. 2019). Despite the absence of appropriate systems, there are many studies that have investigated the design of warning messages to motivate individuals to improve their online security behaviour. One popular approach is message framing, a strategy that aims to convey the main arguments of a complex problem in an understandable and unchallengeable manner (de Bruijn and Janssen 2017). However, one approach is particularly popular in research on understanding the effects of framing and is known as the prospect theory (Kahneman and Tversky 1979). This theory states that people assess obtained information either in terms of potential gains (positive framing) or potential losses (negative framing). The need for message framing to motivate desirable behaviour from non-experts was discussed in previous literature (de Bruijn and Janssen 2017) and was consequently considered for the message generation in the developed artefact. In general, a lack of research is apparent on the effective dissemination of cyber warning messages (Kaufhold et al. 2021) that we want to address with our study.

3 Method

3.1 Design Science Research

DSR aims to generate knowledge on how things should be designed to achieve a desired set of goals. Two dominant types of contributions can be derived by utilizing this approach. On the one hand, an artefact is generated to address real-world challenges. On the other hand, a theoretical contribution can be made through the formulation of design principles which can be beneficial for future development of similar solutions (Baskerville et al. 2018). Within this study, the DSR methodology proposed by Peffers et al. (2007) is utilized which distinguishes between six distinct activities. This starts with (1) the identification of a problem and the motivation of the research, followed by (2) the definition of objectives that are used for the subsequent activity of (3) designing and developing an artefact. Afterward, a comprehensive assessment is conducted in the (4) demonstration and (5) evaluation activities to assure that the achieved functionality adequately addresses the defined objectives. Lastly, (6) obtained knowledge is communicated in publications. Currently, we completed the first three activities of the DSR methodology. We identified the current problems in the cybersecurity domain with a comprehensive literature review and existing interviews with CERTs (Basyurt et al. 2022). The identified challenges were subsequently used to derive the following performance objectives as the second DSR activity:

- The event detection considers all users without a required preselection of expert accounts.
- The identified clusters can be displayed or hidden individually by the user depending on the redundancy of the information they contain.
- The event detection based on Twitter data executes automatically without additional user input.
- The communication module supports the information exchange by automatically generating text fragments, thus accelerating the process of creating warning messages.
- Different keywords are provided to generate actor-specific warning messages. The selectable keywords depend on the selected warning message type and the target group.
- The selectable keywords are based on the required information for the specific target group and the wording of the prewritten text fragments depend on the presumed IT knowledge

- The communication module provides multiple communication channels for warning message dissemination such as the communication via mail or social media like Twitter and Reddit.
- The communication module allows the export of the generated warning message as PDF files.

3.2 Artefact Design and Development

Event detection is a clustering task in which the system receives a stream of new documents, in this case tweets obtained through cyber incident related keyword inputs by users monitored through the Twitter API, and has to organize them according to the most appropriate event (Sceller et al. 2017). One such method is the nearest neighbour search which finds the closest point to a query point Q within a set of points $P = \{P_1, P_2, P_3, \dots, P_N\}$ represented as a matrix M . This is achieved by comparing the computed distance between every point within P with the query point. However, this approach is very computationally intensive, especially at high dimensions, for instance, in the case of a vector space model of a corpus with a large vocabulary (Kannan et al. 2017). To counteract this problem, the approximate nearest neighbour approach was proposed, in which the goal is to find the closest point P' to the query point Q within a certain radius. One of the initial approaches to solve the nearest neighbour approximation problem in sublinear time was the LSH method (Petrović et al. 2010). In the LSH method, each new query point is hashed into buckets in a specific way so that the probability of a collision of two similar points is increased in cases where they are close to each other. Similar to the LSH approach of Sceller et al. (2017) for the detection of cybersecurity events, we used the method and parameters proposed by Petrović et al. (2010), which are considered as the state of the art in LSH event detection. Although there are some drawbacks such as the susceptibility to fragmentation and difficulties in separating similar events that occur simultaneously, we perceived the benefits of avoiding the high computational costs of analyzing large volumes of data in real time as essential to address the identified performance objectives in the previous survey (Basyurt et al. 2022). Since the algorithm compares present words within the collected dataset, it is important to ensure that the compared tweets use similar languages, which was considered in the collection with the Twitter API and in preprocessing, which includes the conversion of tweets into a sequence of word tokens, removal of stopwords and non-English or non-German Tweets. Figure 1 illustrates the developed user interface for identified events.

Thread Number	Thread ID	Number of Posts	Shannon Entropy
1	3	9	1.83
2	4	14	0.00
3	6	8	2.08
4	8	6	1.79
5	9	6	1.79

Figure 1: User interface of the LSH event detection module

The goal of automated CTI exchange is to achieve shared situational awareness through a simplified and accelerated exchange process (Wagner et al. 2019). There are several challenges that are not addressed by the currently available sharing platforms and tools. These include the lack of offered automation mechanisms that impede the process and allow for human error (Sauerwein et al. 2017; Wagner et al. 2019) and reaching different target audiences while taking their varying levels of IT expertise into account (Basyurt et al. 2022). To address these challenges, the artefact includes a communication module to assist CERTs in the process of sharing actor-specific cyber threat information. For this purpose, customizable warning messages with selectable keywords were desired by recently interviewed cybersecurity experts (Basyurt et al. 2022), which was implemented in our artefact. The communication module provides pre-formulated text fragments that can be selected through corresponding keywords and subsequently altered. Depending on the selected keywords, the contained information and the

wording are adapted to the warning message type and the IT knowledge of the selected target group. Selectable keywords like cyber threat types and target groups as well as appropriate warning message formulation were provided by CERTs upon request. Figure 2 displays the user interface of the second module with selectable keywords on the left and the generated customizable message on the right.

Figure 2: User interface of the communication module for actor-specific warning messages

The artefact itself, incorporating both modules, was implemented using Python, Node.js, MongoDB, Vue.js and the Express.js framework, thereby completing the third DSR activity.

4 Conclusion and Next Steps

The goal of the developed artefact is to improve the capabilities of CERTs to manage cyber threats effectively by automatically detecting relevant events and by providing them with support in creating and communicating target-specific warning messages. We aim to further develop the artefact to better address the needs of CERTs and incorporate more social media sources, in future iterations. Through the implementation of the artefact and its future evaluation a contribution to research is made by presenting appropriate approaches to support experts and areas for improvement. Moreover, by providing CERT members with an instrument that addresses their needs and supports their work a practical contribution is made. Our next steps include the demonstration of the artefact's ability to successfully address the identified objectives, which is the fourth activity of the DSR methodology by Peffers et al. (2007). This is followed by the fifth activity dedicated to the evaluation of the artefact and its functionality, which we will perform by conducting semi-structured interviews with CERT members to acquire their opinions on the developed artefact. Lastly, as the sixth activity we will publish our results in a scientific article and describe the technical features of the artefact together with design principles in detail enabling the development of similar tools and their improvement through further research.

5 References

- Agyepong, E., Cherdantseva, Y., Reinecke, P., and Burnap, P. 2020. "Challenges and Performance Metrics for Security Operations Center Analysts: A Systematic Review," *Journal of Cyber Security Technology*. (<https://doi.org/10.1080/23742917.2019.1698178>).
- Ask, T. F., Lugo, R. G., Knox, B. J., and Sütterlin, S. 2021. "Human-Human Communication in Cyber Threat Situations: A Systematic Review," *HCI International 2021*, Springer, Cham, pp. 21–43.
- Baskerville, R., Baiyere, A., Gregor, S., Hevner, A., and Rossi, M. 2018. "Design Science Research Contributions: Finding a Balance between Artifact and Theory," *Journal of the Association for Information Systems* (19:5), pp. 358–376. (<https://doi.org/10.17705/1jais.00495>).
- Basyurt, A. S., Fromm, J., Kuehn, P., Kaufhold, M.-A., and Mirbabaie, M. 2022. "Help Wanted - Challenges in Data Collection, Analysis and Communication of Cyber Threats in Security Operation Centers," in *Wirtschaftsinformatik 2022 Proceedings*.
- de Bruijn, H., and Janssen, M. 2017. "Building Cybersecurity Awareness: The Need for Evidence-Based Framing Strategies," *Government Information Quarterly* (34:1), The Author(s), pp. 1–7.

- Conti, M., Dargahi, T., and Dehghantanha, A. 2018. "Cyber Threat Intelligence: Challenges and Opportunities," *Advances in Information Security* (70), pp. 1–6.
- Horne, B. 2014. "On Computer Security Incident Response Teams," *IEEE Security and Privacy* (12:5), IEEE, pp. 13–15. (<https://doi.org/10.1109/MSP.2014.96>).
- IBM. 2021. "IBM Report: Cost of a Data Breach Hits Record High During Pandemic." (<https://newsroom.ibm.com/2021-07-28-IBM-Report-Cost-of-a-Data-Breach-Hits-Record-High-During-Pandemic>, accessed November 3, 2021).
- Kahneman, D., and Tversky, A. 1979. "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* (47:2), pp. 263–291.
- Kannan, J., Shanavas, A. M., and Swaminathan, S. 2017. "SportsBuzzer: Detecting Events at Real Time in Twitter Using Incremental Clustering," *Transactions on Machine Learning and Artificial Intelligence* (6:1), pp. 1–23. (<https://doi.org/10.14738/tmlai.61.3861>).
- Kaufhold, M.-A., Fromm, J., Riebe, T., Mirbabaie, M., Kuehn, P., Basyurt, A. S., Bayer, M., Stöttinger, M., Eyilmez, K., Möller, R., Fuchß, C., Stieglitz, S., and Reuter, C. 2021. "CYWARN: Strategy and Technology Development for Cross-Platform Cyber Situational Awareness and Actor-Specific Cyber Threat Communication," *Workshop-Proceedings Mensch Und Computer 2021*, pp. 1–9.
- Le, B. D., Wang, G., Nasim, M., and Babar, M. A. 2019. "Gathering Cyber Threat Intelligence from Twitter Using Novelty Classification," *Proceedings - 2019 International Conference on Cyberworlds, CW 2019* (June), pp. 316–323. (<https://doi.org/10.1109/CW.2019.00058>).
- Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. 2007. "A Design Science Research Methodology for Information Systems Research," *Journal of Management Information Systems* (24:3), pp. 45–77. (<https://doi.org/10.2753/MISO742-1222240302>).
- Petrović, S., Osborne, M., and Lavrenko, V. 2010. "Streaming First Story Detection with Application to Twitter," *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 181–189.
- Sabottke, C., Suci, O., and Dumitras, T. 2015. "Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits," *Proceedings of the 24th USENIX Security Symposium*, pp. 1041–1056.
- Saudi, M. M., Basir, N., Nabila, N. F., Ridzuan, F., and Pitchay, S. A. 2016. "An Efficient Easy Computer Emergency Response Team Malware Reservoir System(EZCERT)," *Proceedings - UKSim-AMSS 17th International Conference on Computer Modelling and Simulation*, pp. 142–146.
- Sauerwein, C., Sillaber, C., Mussmann, A., and Breu, R. 2017. "Threat Intelligence Sharing Platforms: An Exploratory Study of Software Vendors and Research Perspectives," *The 13th International Conference on Wirtschaftsinformatik*, pp. 837–851.
- Sceller, Q. Le, Karbab, E. M. B., Debbabi, M., and Iqbal, F. 2017. "SONAR: Automatic Detection of Cyber Security Events over the Twitter Stream," *ACM International Conference Proceeding Series*.
- Sillaber, C., Sauerwein, C., Mussmann, A., and Breu, R. 2016. "Data Quality Challenges and Future Research Directions in Threat Intelligence Sharing Practice," in *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, pp. 65–70.
- Skopik, F., Settanni, G., and Fiedler, R. 2016. "A Problem Shared Is a Problem Halved: A Survey on the Dimensions of Collective Cyber Defense through Security Information Sharing," *Computers and Security* (60), Elsevier Ltd, pp. 154–176. (<https://doi.org/10.1016/j.cose.2016.04.003>).
- Stieglitz, S., Mirbabaie, M., Ross, B., and Neuberger, C. 2018. "Social Media Analytics – Challenges in Topic Discovery, Data Collection, and Data Preparation," *International Journal of Information Management* (39:October 2017), Elsevier, pp. 156–168.
- Wagner, T. D., Mahbub, K., Palomar, E., and Abdallah, A. E. 2019. "Cyber Threat Intelligence Sharing: Survey and Research Directions," *Computers and Security* (87).
- Zhong, C., Lin, T., Liu, P., Yen, J., and Chen, K. 2018. "A Cyber Security Data Triage Operation Retrieval System," *Computers and Security* (76), Elsevier Ltd, pp. 12–31.

Acknowledgements

This work has been funded by the German Federal Ministry of Education and Research (BMBF) in the project CYWARN (13N15407-13N15410).

Copyright © 2022 Kaan Eyilmez, Ali Sercan Basyurt, Stefan Stieglitz, Christoph Fuchss, Marc-André Kaufhold, Christian Reuter, Milad Mirbabaie. This is an open-access article licensed under a [Creative Commons Attribution-Non-Commercial 3.0 Australia License](https://creativecommons.org/licenses/by-nc/3.0/au/), which permits non-commercial use, distribution, and reproduction in any medium, provided the original author and ACIS are credited.