How we Guide, Write, and Cite at CHI

Henning Pohl

University of Copenhagen
Copenhagen, Denmark
henning@di.ku.dk

Aske Mottelson

University of Copenhagen
Copenhagen, Denmark
amot@di.ku.dk

Figure 1: We analyze all papers from CHI 1982–2018. The number of CHI papers per year has recently been growing. Thus, our dataset is biased towards newer work.

ABSTRACT

There are many opinions on how to write an influential CHI paper, ranging from writing in an active voice to including colons in the title. However, little is known about how we actually write, and how writing influences impact. We conducted quantitative analyses of the full text of all 6578 CHI papers published since 1982 to investigate. We looked at readability, titles, novelty, and name-dropping and related these measures to the papers’ citation count; overall and for different subcommittees. We found that CHI papers are more readable than papers from other fields. Furthermore, readability, title length, and novelty markers all influence citation counts.

CCS CONCEPTS

• **Human-centered computing** → HCI theory, concepts and models; • **General and reference** → *Surveys and overviews*;

KEYWORDS

Scientometrics; writing style; readability; citations; novelty; CHI; navel-gazing

ACM Reference Format:

Henning Pohl and Aske Mottelson. 2019. How we Guide, Write, and Cite at CHI. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI’19 Extended Abstracts)*, May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3290607.3310429>

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’19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5971-9/19/05...\$15.00

<https://doi.org/10.1145/3290607.3310429>

Note: Error bars in all figures below show bootstrapped 95 % confidence intervals. Regression plots also show the 95 % confidence interval of the regression estimate.

INTRODUCTION

The CHI conference is the focal point of human-computer interaction research. While the individual papers tackle diverse topics and vary in style, in aggregate, they represent the current state of the field. Thus, over time, larger trends and shifts within the field also manifest themselves in the papers appearing each year.

An example of the trends revealed by analysis of CHI paper data is Bartneck and Hu's work [1] on *who* contributes to CHI. Similarly, data on authorship and gender was also reported by Kaye [3]. Liu et al. investigated [7] *what* kind of topics are written about at CHI, in which ways they relate to each other, and how they change over time. Yet, little is known about *how* papers are written for CHI.

We address this gap and investigate how measures of writing style vary over the history of CHI. Where the analyses above were run on meta-data and author-picked keywords, our investigation of paper writing is done on the actual *paper content*. This allows for analysis of the *readability* of papers, as well as into more idiosyncratic aspects of them. For example, we looked into how papers write about novelty and how papers from individual subcommittees differ.

Yet, investigating how papers *are* written has limited guiding power for how papers *should* be written. To provide an estimate for the impact of different writing styles, we correlated them with citation data from Google Scholar.

CHI PAPER WRITING

Opinions on what entails good CHI writing differ and several guidelines exist for this. General advice is also given in official documents like the CHI 2019 “*guide to a successful submission*”¹. Here, both *clarity* and *conciseness* are stated as the main goals to strive for in writing. Added emphasis is placed on language, such that papers are written in a way “*that effectively communicates across national and cultural boundaries*.” Examples of things to avoid are “*long, complex sentences, and [...] regional colloquialisms, jokes, or puns that could be difficult for someone outside your culture to understand*.”

In addition to writing advice on the CHI website, several members of the community have created their own guides to CHI writing. An example is Lennart Nacke's course [8] on *How to Write and Review CHI Papers*². He, for example, advises writers to avoid passive voice, reduce jargon, and remove words that do not contribute to a sentence's meaning. Other examples of writing guides include blog posts by Philip Guo³, Andrew Ko⁴, Lana Yarosh⁵, and Jacob Wobbrock⁶. Apart from general writing advice, these guides commonly also include advice on how to structure a research paper and what kind of content should go into each section. For example, several of the guides advise structuring the introduction along 5–6 key paragraphs. In our investigation of writing, we focused on writing style and did not look into the structure of papers.

¹<http://chi2019.acm.org/authors/papers/guide-to-a-successful-submission/>

²<http://chicourse.academic.com/>

³<http://pgbovine.net/how-to-write-hci-research-paper.htm>

⁴<https://faculty.washington.edu/ajko/-advice#goodpaper>

⁵<http://lanayarosh.com/2012/10/how-to-get-me-to-positively-review-your-chi-paper/>

⁶<http://faculty.washington.edu/wobbrock/pubs/-Wobbrock-2015.pdf>

Figure 2: We collected citation data for all CHI papers from Google Scholar. The blue line shows the average number of citations for papers from each CHI; the orange line shows the citation count for the 10 % most cited papers.

Figure 3: The cumulative probability of citations per year varies by decade, with 2000–2010 having the highest impact.

Figure 4: Shown here are the normalized citations counts for all CHI papers, suggesting declining impact in recent years.

SOURCE DATA COLLECTION

For our analyses we used the CHI papers' content, and citation data for them from Google Scholar.

CHI Papers

We gathered PDFs of all papers published at CHI, growing from around 100 per year through the 90's and 00's to more than 600 per year in recent years (see Figure 1). The data set contains the full texts of all CHI papers (including full and short papers, excluding extended abstracts). The complete data set holds 6578 papers over 36 years (1982–2018), excluding 1984, where CHI was not held.

From each CHI paper, we extracted the full text from the PDF using pdftotext. Where this failed (e.g., because some older papers were scanned PDFs) we performed OCR using Adobe Acrobat. We analyzed the quality of the extracted text by looking at the amount of unknown words in each paper, and based on that analysis manually transcribed 27 papers.

With a Python script we combined string manipulations and approximately 100 regular expressions, to remove papers' meta information, such as title, author information, ACM classification, keywords, copyright information, session title, references, and acknowledgements. All analyses are conducted on the body text of the papers.

Citation Data

While the ACM Digital Library provides citation counts for papers, these counts underestimate the actual number of citations of a paper, because the ACM Digital Library is not well-equipped for tracking citations from academic work published outside ACM itself. To gather more accurate citation data, we hence acquired citation metrics from Google Scholar. We scraped Google Scholar using a Python script and stored the citation count for every CHI paper during October 2018. We were unable to automatically match about 70 papers with Google Scholar; for these we manually searched Google Scholar and found the citation count.

Figure 2 shows how the average number of citations for papers at CHI varies over the years. It also highlights how the 10 % most cited papers each year are cited more than four times as frequently compared to the overall average. However, because our dataset spans work from 36 years, the total number of citations of papers are not directly comparable. Instead, we derive the measure of *citations per year* for normalization. We take into account fractional years since publication. For example, between CHI 2018 and the scraping of citation data in October, only about half a year had passed.

Figures 3 and 4 show how citations per year for CHI papers varies between years. Since about 2010, the average number of citations a CHI paper receives per year has been declining, which, however, may be a product of their infancy. In contrast, between 1990 and 2010, many papers were published that have seen a continuously strong reception.

Figure 5: NDC scores translate to grade levels where a text is easily understood.

Figure 6: Readability of CHI papers over time; the lower the score, the more readable a paper is. CHI papers have become easier to read over time.

Figure 7: Readability has an impact on how much papers are cited. Too hard and too easy to read correlate with lower citation rates.

READABILITY

Readability concerns the ease with which a reader can understand text. As previously noted, CHI submissions are asked to be clear and concise, complex sentences and jargon should both be avoided. Adhering to these guidelines should increase the readability of a text. Several metrics that quantify the readability of text exist; in this work we employ the New Dale-Chall Readability Formula (NDC) [2]:

$$\text{NDC} = 0.1579 \left(\frac{DW}{W} \times 100 \right) + 0.0496 \frac{W}{S} + 3.6365,$$

where W , DW , and S are the number of words, difficult words, and sentences, respectively. For NDC, lower scores indicate higher readability and scores allow for interpretation as required grade levels (see Figure 5). The formula uses a list of 3000 words that most fourth-grade US students understand (all other words are treated as difficult). We used the Python package `textstat` to compute the values used for this analysis.

NDC scores across the history of CHI are decreasing slowly; $r = -.33$, 95% CI $[-.36, -.29]$, $p < 0.001$. This suggests that readability over the history of CHI has slightly changed towards more readable (see Figure 6). This is in contrast to other academic disciplines, where the readability has been decreasing over time [9]. The relatively short existence of CHI compared to other academic disciplines, of course, makes this comparison somewhat difficult. However, the average writing style of CHI is much simpler than that of other select academic disciplines, leading to higher readability. In 2015, the average CHI paper had an NDC of 7.5, while a large PubMed analysis yielded an average of 13 NDC for the year 2015. Where 7.5 NDC is considered comprehensible to grades 9–10, an NDC score of 10 matches the comprehension level of college graduates. This difference in readability could increase over time, as the readability of scientific writing in general is decreasing, while the readability at CHI seems to follow the opposite trend.

Impact of Readability

We found that readability is slightly correlated with how many citations a paper receives per year; Kendall's $\tau(6576) = 0.07$, $p < 0.001$. However, most CHI papers fall into a small NDC range: 50 % have an NDC between 7.0–7.7 and 99 % are within an NDC of 6.1–8.6. Figure 7 shows how a paper's citations per year varies within this range. How many citations a paper receives seems to peak around an NDC of 6.5–7.5.

Multiple factors could contribute to the observed patterns. Readability has been increasing over the years and there are many more recent papers than older ones. Whether the higher readability of recent work is responsible for that remains to be explored. Writing too complex or too simple papers does seem to harm citations counts.

Figure 8: Distribution of paper title lengths in CHI papers.

Figure 9: The length of paper titles negatively correlates with the amount of citations a paper receives.

Figure 10: Most marks used in titles make no difference with respect to a paper's citations. However, papers using a comma in the title were cited significantly more than those without.

⁷<http://sigchi.tumblr.com/post/104956615720/-what-should-you-title-your-paper>

CHI PAPER TITLES

Paper titles are the first thing readers see and thus the question of how to choose a title has garnered much interest. The SIGCHI Tumblr, for example, has a post on paper titles and how they impact citations⁷. For that analysis 90668 articles from the ACM Digital Library were used. The analysis is then only in the correlation of the use of a colon or question mark in the title with the total number of citations (presumably per the Digital Library) a paper has received. It found that conference papers with a colon in the title received about two citations more than conference papers without said colon. Respectively, conference papers with a question mark in the title received about three citations less than those without. However, this data is not specific to CHI or even just HCI conferences.

Analyzing the 6578 titles of CHI papers, we find that colons are used in 58.6 % (3871) of them. Other marks, like semicolons (0.3 %, 18), commas (7.6 %, 505), periods (0.9 %, 57), question marks (4.9 %, 327), or exclamation points (1.1 %, 70) were used much less. Acronyms (defined for our purposes as sequences of at least three capital letters) were used in 7.9 % (519) of CHI paper titles.

CHI paper titles are between 5–187 characters long ($M=74$, $SD=23$). Figure 8 shows how the distribution of paper title lengths is multimodal. There is a noticeable dip around a title length of 62 characters. While the paper format changed over the years, this roughly coincides with the length of one line. Authors appear to avoid titles that result in single words on a separate line. This also explains why there are few papers with more than ~120 characters.

Impact of Paper Titles

As shown in Figure 9, papers with longer titles tend to be cited less than papers with shorter titles. The length of paper titles significantly predicts the number of citations per year; $r^2 = 0.16$, $p < 0.001$. A similar relationship has previously been observed for a sample of journal articles [5]. Why shorter titles are better is unclear, though. Maybe readers are less likely to even open a paper with a long title.

As shown in Figure 10, most marks used in titles do not make an impact on how much a paper is cited. We used Welch's t-tests with Bonferroni corrections for multiple comparisons to analyze this impact. The only significant difference was between papers with a comma in the title and those without; $p < 0.05$. As in a previous analysis⁷, we see a small increase of citations if a paper has a colon in the title. However, this is only a slight difference and was not significant after correction for multiple comparisons; $p = 0.3$.

While the mean NDC of the body of CHI papers is 7.6 ($SD=.5$), interpreted as “easily understood by an average 9th or 10th-grade student,” paper titles are, on average, more sophisticated. The mean NDC for CHI paper titles is 13.1 ($SD=2.6$)—about 70 % harder to read than the body of the papers. A linear regression showed no significant correlation between title NDC and citation count: $r^2 = .006$, $p = .29$. This suggests it might be worthwhile to use more complex language in order to keep a title shorter.

Figure 11: Development of the use of the words ‘new’ and ‘novel’ throughout CHI.

Table 1: Top-5 bi-grams used with ‘new’ and ‘novel’ in CHI papers.

Rank	Bigram ‘new’	Bigram ‘novel’
1.	new ways	novel interaction
2.	new technologies	novel approach
3.	create new	novel design
4.	new design	novel ways
5.	new forms	novel interfaces

Figure 12: The analysis shows a significant difference in citation count comparing manuscripts that include ‘new’ and manuscripts that do not; not using the word ‘new’ seems to yield lower impact.

⁸<http://chi2019.acm.org/authors/papers/guide-to-a-successful-submission/>

NOVELTY AT CHI

All CHI reviewers are asked to judge submissions based on the same criteria, one of which is “*Originality of the work: what new ideas or approaches are introduced?*”⁸. Other research disciplines are increasingly using ‘new’ and ‘novel’ in publications [11]; we were similarly interested to see if that is also the case for CHI.

Across all CHI papers we checked if the words ‘new’ or ‘novel’ were present in the body of the paper. We then found all bigrams including ‘novel’ and ‘new’ to see what types of things are described as new at CHI. The bi-grams refer to two adjacent words, such as the bigram ‘quite new’ or ‘novel interface’. Figure 11 shows that CHI papers often write about novelty; at CHI’18, for instance, 93 % of the papers included the word ‘new’ and 46 % contained ‘novel’. While the word ‘new’ has been used in around 90 % of the papers at each CHI, the use of ‘novel’ is on the rise, with around 10 % of papers using ‘novel’ in the 80’s, 20% around year 2000, and 40 % in the most recent years. This trend is similar to that of other scientific fields, as the use of ‘novel’ in PubMed records has increased almost 4000% since 1970 [11].

A linear extrapolation suggests that every scientific (PubMed) publication will use the word ‘novel’ in the year 2123. Similarly, this would occur already in 2060 for CHI; $r^2 = .84$, $p < .001$.

What’s ‘new’?

To shed light on how the CHI community writes about novelty, we looked at the sentences containing ‘novel’ and ‘new’. Table 1 shows the most common words associated with ‘new’ and ‘novel’, respectively. The most commonly used words to claim novelty at CHI are: ‘ways’, ‘forms’, ‘interactions’, ‘approaches’, ‘technologies’, ‘designs’, ‘interfaces’, and ‘systems’.

Additionally, novelty is often claimed as an introduction: ‘create new’, ‘present a novel’, or ‘propose a novel’ are among the most common bi- and trigrams of ‘new’ and ‘novel’. The word ‘novel’ tends to come a bit before ‘new’ within a paper. ‘Novel’, on average, is placed a bit before the middle of the paper with (40 %, $SD = 36$), while ‘new’ is, on average, used at the middle of a paper (48 %, $SD = 32$).

The Impact of Novelty

We checked whether the presence of ‘new’ and novelty markers (‘novel’ and ‘novelty’) influences citation counts (see Figure 12). On average, the inclusion of novelty markers does not have a notable influence on citation count, with papers including novelty markers garnering 8.7 citations per year, and papers without 8.1 citations per year; $t(5229.5) = 1.9$, $p = .06$. Including ‘new’, however, seems to increase citation count (or rather, papers not including ‘new’ have decreased citation counts). The average citation count per year for papers including ‘new’ is 8.6 while the average for papers without is only 6.5 citations per year. The difference is also significant: $t(867.2) = 4.4$, $p < .001$.

Figure 13: Over the years, the percentage of CHI papers that refer to famous philosophers and thinkers has been steadily increasing. In addition to the overall ratio of name dropping, we show the development for the four most used ones. Currently, about every 30th paper refers to famous philosophers.

Figure 14: We compared how papers name-dropping and papers not doing so fare with respect to citations per year. We found a significant difference between the two, where name-dropping leads to more citations per year.

NAME-DROPPING IN CHI PAPERS

Name-dropping refers to the practice of sneaking names of authoritative people into a conversation to impress others. Here, we define it as referring to famous people within a paper, which is different from referencing. For example, we can sensibly point to Liu et al.’s work here [6]. On the other hand, if we were to mention that Francis Bacon pointed out in *Of Studies*, that “*Reading maketh a full man; conference a ready man; writing an exact man*”, it would be in loose connection to this paper’s topic, yet quite unnecessary and a bit pretentious.

Name-dropping is not per se bad writing. In fact, a CHI paper might have good reasons to refer to Wittgenstein or Popper. However, we regard them as a marker for writing that attempts to be more theory-heavy or to connect to larger streams of thought outside of HCI. To find out how much CHI papers engage in this practice, we searched the full texts (sans references) for names of famous thinkers. We found these on the Wikipedia portal on Philosophy, and included names with at least five references throughout the CHI corpus: Adorno, Aristotle, Barthes, Baudelaire, de Beauvoir, Bourdieu, Camus, Confucius, Deleuze, Derrida, Descartes, Foucault, Freud, Habermas, Hegel, Heidegger, Horkheimer, Husserl, Jesus, Kant, Kierkegaard, Kuhn, Laclau, Locke, Marx, Nietzsche, Plato, Popper, Sartre, Schopenhauer, Socrates, Spinoza, Wittgenstein, and Zizek (including spelling variations). As shown in Figure 13, the share of CHI papers that name-drop has been steadily increasing over the last decades. In the ongoing decade, 3.4 % of papers mention at least one of the above people, compared to the 80s where this was only true for 1.5 % of papers.

Impact of Name-Dropping

The average CHI paper receives 7.7 citations per year. As shown in Figure 14, there is a noticeable difference for papers not name-dropping (7.7 citations per year) and name-dropping (10.1). A Welch’s t-test also shows a significant difference between the two; $t(193.0) = 2.5$, $p < 0.05$. So overall, name-dropping does correlate with how much a paper is cited. However, there might be differences depending on *whose* name is dropped. We take a closer look at the effects of naming people that were referred to in at least ten CHI papers. Only two of them relate to papers being cited less than average: Jesus (6.5 citations per year, 13 instances), and Habermas (7.0, 12). However, nine people had a more positive impact: Kant (7.9, 13), and Aristotle (9.6, 17), Bourdieu (10.1, 11), Wittgenstein (10.9, 11). Plato (10.9, 11), Kuhn (11.0, 20), Heidegger (12.3, 21), Marx (12.3, 21), and Foucault (13.3, 32).

Likely, just mentioning Foucault will not increase a paper’s impact. Name-dropping Locke might help (those five paper have an average of 44.3 citations per year), but could also be a spurious result. However, there could be a link where papers on certain topics are more likely to, for example, name-drop Foucault. A researcher looking for guidance, might thus want to inspect *why* papers pointing to Marx tend to have a higher impact.

Figure 15: We classified CHI papers into subcommittees, based on the example papers for each of CHI 2019's subcommittees (shown as crosses). Shown here is a 2D embedding generated via t-SNE.

Figure 16: Readability and citations per year for the subcommittees as above.

DIFFERENCES BETWEEN SUBCOMMITTEES

CHI as a conference spans multiple subareas of HCI. Because of this diversity and because of the size of the conference, paper submissions are handled by one of several subcommittees. For CHI 2019, for example, there are 12 committees on areas such as *Health* or *User Experience and Usability*.

The subcommittees have changed over the years, both in numbers and themes, and early CHIs did not have subcommittees at all. We used the example papers provided for each subcommittee on the CHI 2019 website as seed data and then classified other papers based on their similarity to these. To determine this similarity, we represented papers with their document embedding using Doc2Vec [4, 10], resulting in 100-dimensional document vectors. We used a k -nearest neighbor classifier to assign all other papers published at CHI to subcommittees. A good classification should keep the relative size of the subcommittees constant; we found that a value of $k = 1$ yielded the lowest deviation.

Figure 15 shows a t-SNE embedding of the classification. To compare how closely related subcommittees are, we bootstrapped the average distance between their papers. The most closely related pair of subcommittees are *Interaction Techniques, Devices, and Modalities* and *Engineering Interactive Systems and Technologies*. This is followed by the *Design* subcommittee, which is closely related to *Accessibility and Aging*, *Specific Applications Areas*, and *Interaction Beyond the Individual*. On the other hand, the *Games and Play* subcommittee shows the lowest relationship to other subcommittees.

Subcommittees also differ in coherence: how closely related papers are *within* it. We find that the *Accessibility and Aging* and *Interaction Techniques, Devices, and Modalities* subcommittees exhibit comparably high coherence. The largest subcommittee, *Understanding People: Theory, Concepts, Methods*, however, is also the least coherent. It is followed by *Privacy, Security and Visualization*, where papers on visualization possibly have not much in common with those on privacy and security.

The upper plot in Figure 16 shows how readability varies between the subcommittees. While the differences overall are comparably small, we found that papers from the *Games and Play* subcommittee had the highest readability. On the other hand, papers from the *Design* subcommittee had a slightly lower readability. This does not directly translate to the number of citations papers from these subcommittees receive (see lower plot in Figure 16). Here, *Learning, Education and Families* fares worst and *Health* best. However, this analysis based on CHI 2019 subcommittees is likely confounded in several ways. For example, the average paper in *Games and Play* was published in 2014, while the average *Design* paper was published in 2008. Furthermore, *Design* papers have shorter titles on average than *Health* papers (71 vs. 124 characters). There are also limitations to our subcommittee classification, as it is only based on the 235 example papers given for the subcommittees.

Overall, this analysis shows that there are slight differences in writing between the subcommittees. This can be used to derive relationships between subcommittees, but also seems to slightly impact how many citations papers from each subcommittee receive.

DISCUSSION

We have presented several measures of writing style, evaluated on the dataset of all CHI papers. With a conference the size and breadth of CHI, a large amount of variability is to be expected which in fact shows in our analysis. However, we believe there are a few overall guidelines and questions that emerge from this exploration.

The biggest trend we have seen is a decline in impact over recent years. This goes along with more readability, more ‘novelty’, and more name-dropping. However, we cannot say whether these measures are linked to the decrease in citations per year. We also looked at averages and every year there are, of course, still impactful papers published at the conference.

We have also seen how the split of CHI papers into subcommittees is fuzzy and there is a strong overlap between most of them. Interestingly, though, some subcommittees are more connected to others. For example, if writing of papers in the *Interaction Technologies, Devices, and Modalities* and *Engineering Interactive Systems and Technologies* is so similar, why are these individual committees? Similarly, the *Privacy, Security and Visualization* committee seems to span weakly-connected fields, and it also seems like *Games and Play* might be a bit of an outsider committee as well. This by no means is meant to suggest that these works have no place at CHI, but raises the question whether CHI adequately handles the contained diversity of research.

In closing remarks, it should be noted that the correlations found do not imply causation; for instance, because papers with longer titles are not cited as much as those with shorter titles, it is not necessarily the underlying reason. Furthermore, because of the exponential growth in the number of CHI papers, analyses of characteristics (where not done by year) are biased towards recent work.

CONCLUSION

How to write for CHI is a question several members of the community have tried to answer in blog posts or lectures. However, each individual’s perspective is inherently limited to the subset of papers they have seen. With 6578 paper currently published at CHI, there likely is nobody who read all of them. By applying quantitative analysis methods to this dataset, we were able to provide some information on the ways papers are written and how that impacts citation counts.

Yet, much work remains and the purpose of this paper is more to provide data for an ongoing conversation than to provide the ultimate guide to CHI paper writing. We have tried to consider several of the measures presented above in the writing of this paper. For example, the NDC of this paper itself is 6.6 (7th or 8th-grade), which our analysis indicates a decent probability of impact for. We also name-drop Locke and put a comma in the title (short, of course). Only time will tell whether this caused others to cite this paper as much as our analysis suggests.

“The better to understand the nature, manner, and extent of our knowledge, one thing is carefully to be observed concerning the ideas we have; and that is, that some of them are simple and some complex.”

— John Locke, *An Essay Concerning Human Understanding*, 1689

ACKNOWLEDGMENTS

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement 648785).

REFERENCES

- [1] Christoph Bartneck and Jun Hu. 2009. Scientometric Analysis of the CHI Proceedings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 699–708. DOI : <http://dx.doi.org/10.1145/1518701.1518810>
- [2] J.S. Chall and E. Dale. 1995. *Readability revisited: the new Dale-Chall readability formula*. Brookline Books.
- [3] Joseph 'Jofish' Kaye. 2009. Some Statistical Analyses of CHI. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems (CHI EA '09)*. ACM, New York, NY, USA, 2585–2594. DOI : <http://dx.doi.org/10.1145/1520340.1520364>
- [4] Quoc Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *Proceedings of the 31st International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Eric P. Xing and Tony Jebara (Eds.), Vol. 32. PMLR, Beijing, China, 1188–1196. <http://proceedings.mlr.press/v32/le14.html>
- [5] Adrian Letchford, Helen Susannah Moat, and Tobias Preis. 2015. The advantage of short paper titles. *Royal Society Open Science* 2, 8 (2015). DOI : <http://dx.doi.org/10.1098/rsos.150266>
- [6] Yong Liu, Jorge Goncalves, Denzil Ferreira, Simo Hosio, and Vassilis Kostakos. 2014a. Identity Crisis of Ubicomp?: Mapping 15 Years of the Field's Development and Paradigm Change. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 75–86. DOI : <http://dx.doi.org/10.1145/2632048.2632086>
- [7] Yong Liu, Jorge Goncalves, Denzil Ferreira, Bei Xiao, Simo Hosio, and Vassilis Kostakos. 2014b. CHI 1994-2013: Mapping Two Decades of Intellectual Progress Through Co-word Analysis. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3553–3562. DOI : <http://dx.doi.org/10.1145/2556288.2556969>
- [8] Lennart E. Nacke. 2017. How to Write and Review CHI Papers. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 1228–1231. DOI : <http://dx.doi.org/10.1145/3027063.3027097>
- [9] Pontus Plavén-Sigraý, Granville James Matheson, Björn Christian Schiffler, and William Hedley Thompson. 2017. Research: The readability of scientific texts is decreasing over time. *eLife* 6 (sep 2017), e27725. DOI : <http://dx.doi.org/10.7554/eLife.27725>
- [10] Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. ELRA, Valletta, Malta, 45–50. <http://is.muni.cz/publication/884893/en>.
- [11] Christiaan H. Vinkers, Joeri K. Tjink, and Willem M. Otte. 2015. Use of positive and negative words in scientific PubMed abstracts between 1974 and 2014: retrospective analysis. *BMJ* 351 (2015). DOI : <http://dx.doi.org/10.1136/bmj.h6467>