

29410 unique referenced or cited documents have DOIs.⁴ We use DOIs to link 11035 (51.4%) references and 17203 (56.7%) citations to DBLP, from which we derive normalized venue metadata. An additional 4464 references and 3783 citations are linked to Semantic Scholar; the venue data in Semantic Scholar is not normalized—i.e. the venue is a string value that must be mapped to a normalized venue, e.g., “CHI ’19” and “The 2019 ACM CHI Conference on Human Factors in Computing Systems” must both be mapped to *CHI*. We heuristically and manually map these venue strings to normalized venues. We are unable to find venue information for 2658 unique references and 4760 unique citations. Most (71.5%) of these venue-less documents lack DOIs, making them challenging to identify or link. Of those with DOIs, we investigate a sample to better understand what they are and how their lack of venue information could impact our analysis. An assessment of the 100 most commonly occurring DOIs within this set reveals that most of these (73 of 100) resolve to books, book chapters, reports, or other document types without associated venues. Of the documents that have a publication venue unknown to DBLP or Semantic Scholar (23 of 100), all are from less well-known venues, and none are associated with the venues selected for our analysis. Therefore, we anticipate minimal bias to our results due to missing venue data. Details on this error analysis and additional commentary are available in supplementary files.

Finally, we map all papers to the Microsoft Academic Graph (MAG) to derive their fields of study [24, 25]. The MAG fields are organized into a six-level hierarchy, and we retain and analyze all fields in the upper two levels (L0-L1). L0 is the highest level, and includes 19 fields such as Medicine, Psychology, and Computer Science. L1 fields are more granular, including things like human computer interaction, computer vision, developmental psychology, physical therapy, etc. Though the hierarchy continues into L2 and beyond, the fields quickly become too specific, which is why we elect to perform analysis over only the top two levels. Each paper can be associated with multiple fields of study at each level. Each field of study may have multiple parents, though we default to selecting a primary parent when displaying the L0 information associated with any particular L1 field. Of the 21464 references and 30355 citations of the core set, we identify field of study information for 19252 (89.7%) references and 26997 (88.9%) citations.

To provide context for interpretation, we select a set of 13 conferences that publish accessibility and HCI research (including ASSETS and CHI) for comparison. These publication venues (Table 1) are selected based on proximity and prestige to accessibility and HCI. They include general HCI venues like CHI and HCI, sub-discipline specific venues like TEI and UIST, as well as regional conferences like OZCHI and NordiCHI that are similar in size to ASSETS. Note that the number of entries for ASSETS in Table 1 is much higher than the paper count in the core dataset, which includes all *full-length* accessibility papers at both ASSETS and CHI and no extended abstracts. In contrast, Table 1 is derived from DBLP and includes full length papers along with other types of accepted submissions such as posters, late-breaking work, and/or

⁴DOIs are provided by most large academic publishers, and are the most widely used identifiers for scholarly publications. However, not all publications receive DOIs, e.g. some conferences and workshops do not acquire them, some books may only have ISBNs, etc. A coverage of 78% is fairly standard.

demographics. Given this distinction, none of these venues are directly comparable to the core set, and are rather used to provide context for the expected reference and citation diversities in similar venues. Similar to the core set, references and citations for papers in these comparative venues are derived from Semantic Scholar, and venue and field of study information from DBLP, Semantic Scholar, and MAG as previously described.

3.2 Analyses

We examine citation patterns in accessibility research, references and citations to/from other fields, and temporal trends.

Analysis of venues: We aggregate all references and citations for the core dataset. Identifying the venues of these papers, we then determine the top venues referenced by and citing these accessibility papers. Venue is one of two proxy measures we use to distinguish between research communities.

Analysis of fields of study: We analyze the MAG L1 fields of study of each paper referenced by or citing a paper in the core set. We also analyze temporal trends to determine whether the proportional representation of certain fields is increasing, decreasing, or remaining stable over time. Though we refer to these as fields based on MAG’s nomenclature, each field more closely resembles a topic. Therefore, though a paper may be published in a computer science venue like CHI, it may be about a combination of topics, including ones in computer science like human-computer interaction or computer security, but also outside of computer science like epistemology or ethics.

Comparative analysis: Finally, we perform field of study analysis across all 13 comparative venues. To compare the diversity among referenced and citing papers of these venues, we compute a diversity index over their MAG fields of study. The LCDI is computed over the L1 fields for the references and citations of each paper, and is defined as:

$$LCDI = \left(\sum_{i,j=1}^N s_{ij} p_i p_j \right)^{-1} \quad (1)$$

where s_{ij} gives the similarity between two fields of study, p_i is the proportion of references in field i out of N total fields, and j the field of the paper of interest. We derive the similarity s_{ij} between fields using the hierarchy defined by MAG as $\frac{1}{2^n}$, where n is the number of levels of hierarchy that must be traversed to find a common parent. The larger the LCDI, the more diverse the fields of study are among the reference or citation pool for that paper. In the case where all referenced papers are in the same field of study as the paper of interest, the LCDI equals 1. For each comparative venue, we compute the LCDI of all references and citations, and compare the distributions of these diversity scores.

4 RESULTS

The 836 papers in the core accessibility set reference a median of 23 papers (mean = 25.7; SD = 17.0) and are cited by a median of 24 papers (mean = 36.3; SD = 43.6). Figure 1 shows the distribution of reference and citation counts per paper in the core dataset split by venue. The average number of references is much higher for CHI (mean = 34.5; SD = 18.3) than ASSETS (mean = 21.7; SD = 14.7). Though the average citation count is also higher for CHI (mean =

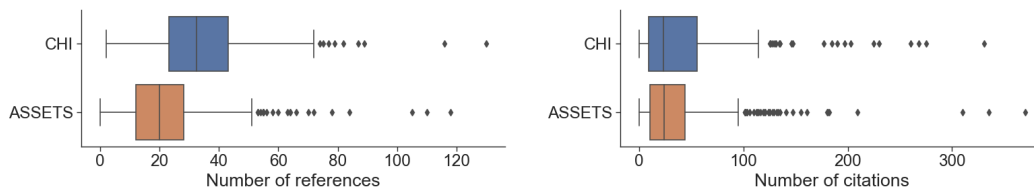


Figure 1: Distribution of reference (outbound citations) and citation (incoming citations) counts for accessibility papers in the core set split by ASSETS (N=576) and CHI (N=260). CHI accessibility papers have a median of 32.5 references (mean=34.5; SD=18.3) and ASSETS 20.0 references (mean=21.7; SD=14.7). For citations, CHI accessibility papers have a median of 23.0 citations (mean=41.9; SD=52.1) and ASSETS 24.0 citations (mean=33.8; SD=38.9).

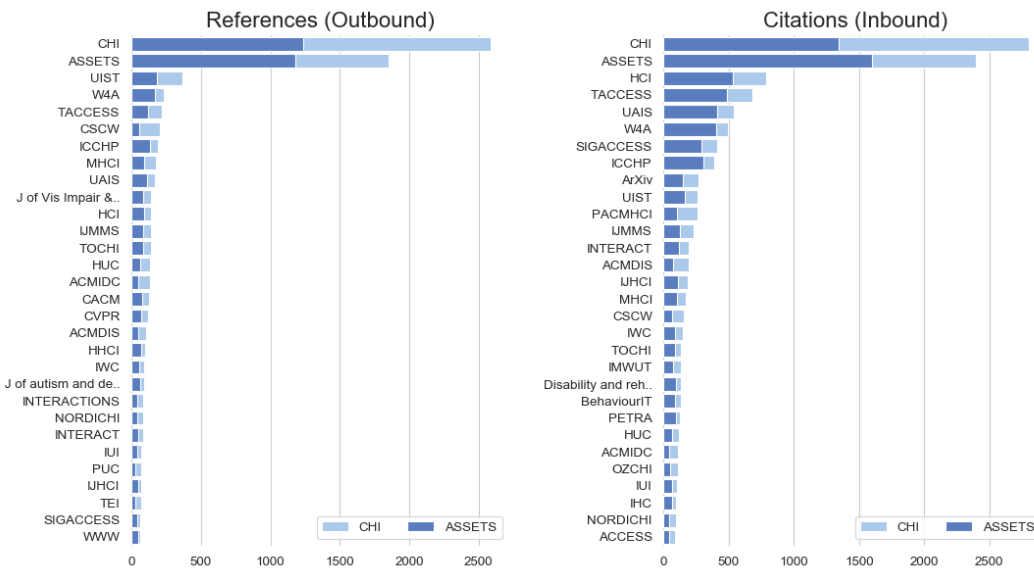


Figure 2: Top venues of papers referenced by (left) and citing (right) accessibility papers in the core set. References and citations are both dominated by papers from CHI and ASSETS, though a relatively larger proportion of citations arrive from other publication venues.

41.9; SD = 52.1) than ASSETS (mean = 33.8; SD = 38.9), the median is similar for both venues, 23 for CHI and 24 for ASSETS.

Figure 2 shows the top venues represented among referenced and citing papers. CHI and ASSETS papers make up a substantial portion of references and citations, and are especially well-represented among references. For references, ASSETS papers cite a similar number of papers in CHI and ASSETS, though CHI papers in our core set are around twice as likely to cite CHI papers as ASSETS papers ($\chi^2 = 108.1, p < 0.001$). Other reference behaviors are similar between the two subsets, though CHI papers are more likely to cite papers published in CSCW and IDC. Citations, on the other hand, are more likely to come from papers in the same venue, i.e., citations to ASSETS papers are more likely to come from ASSETS papers, and citations to CHI papers from CHI papers ($\chi^2 = 182.4, p < 0.001$) (see contingency tables in supplementary files). The most commonly occurring non-computer science venues among references are disability-related journals like the *Journal of Visual Impairment & Blindness* (141 references) and *Journal of Autism and Developmental Disorders* (89 references), and among citations, *Disability and Rehabilitation: Assistive Technology* (133

citations) and *IEEE Trans. Neural Systems & Rehab. Engineering* (60 citations).

Figure 3 shows the distribution of fields among referenced and citing papers for the core set, split by CHI and ASSETS. Papers referenced by accessibility papers tend to be about a more diverse set of topics (relative counts of non-HCI papers are higher among references than citations). Though these papers are dominated by computer science subfields, many subfields of psychology and medicine are represented. Citations are more niche, originating predominantly from topics in computer science. We aggregate the non-CS subfields among the top 30 fields shown in Figure 3 into their L0 parent fields, and present temporal trends for these aggregate L0 fields in Figure 4. Over time, papers on the topic of psychology appear more frequently among both references and citations, and to a lesser degree, we see an increasing proportion of citations from papers on the topic of medicine. These increases can be partially attributed to the growing diversity of fields represented by papers in the core set, e.g., in recent years, more papers in the core set are being classified into subfields of psychology. We show this change in Figure 4 (left), which plots the distribution of

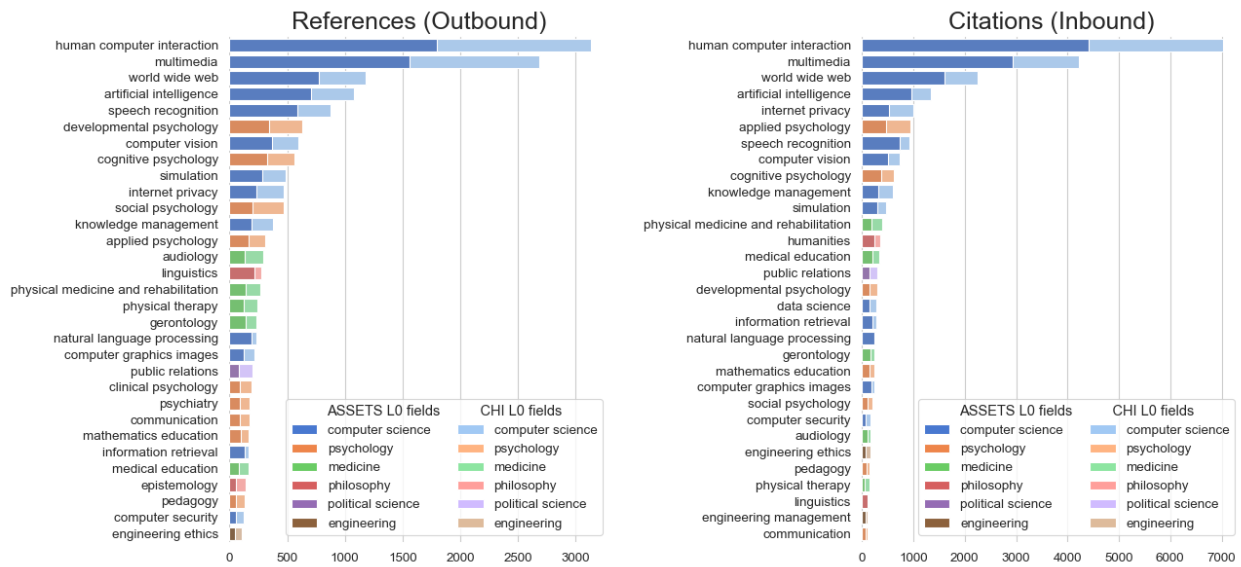


Figure 3: Top L1 fields of study of papers referenced by (left) and citing (right) accessibility papers in the core set. The primary color of each bar is determined by the L0 parent of that field, e.g., subfields of computer science are blue and subfields of psychology are orange. References show a greater diversity of fields of study (especially those outside of computer science) compared to citations.

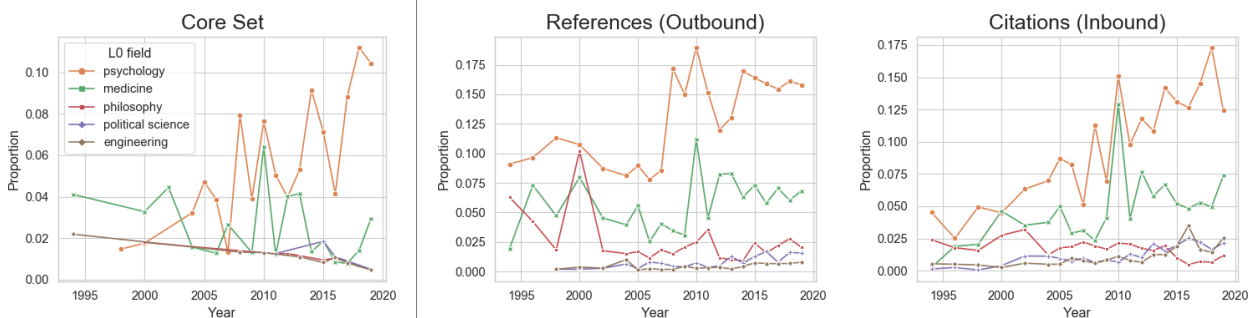


Figure 4: Proportions of non-CS fields represented among papers in the core set over time (left), and among the papers referenced by (center) and citing (right) papers in the core set. The proportions of papers on the topic of psychology, and medicine to a lesser degree, have increased over time, especially among citations. Only years with greater than 5 papers are shown.

non-CS L0 fields associated with L1 fields seen among the papers of the core set. A reference or citation to an accessibility paper that is classified into a subfield of psychology will artificially inflate the representation of psychology among the referenced and citing papers. However, this increase is not seen for medicine, and the increase in incoming citations to accessibility papers from papers on subtopics of medicine may derive from other sources.

Figure 5 shows the distribution of the LCDI for each venue. As expected, more general HCI venues like CHI and UbiComp have higher diversity among references and citations than subdomain-focused venues like ASSETS, ICCHP, or TEI. Among these venues, papers in CHI, UbiComp, CSCW, and IDC influence the most diverse set of papers. These LCDI values show that the citation diversity for these HCI venues is generally lower than their reference diversity (note different scales). For venues like CSCW or UIST, there is

minimal difference between reference and citation diversity; yet the difference is pronounced for venues like the HCI conference. ICCHP has relatively lower reference and citation diversity, suggesting that it is a more niche conference in general. ASSETS and CHI have fairly standard levels of reference diversity and enjoy comparable or higher levels of citation diversity compared to other HCI venues.

5 DISCUSSION

Citation diversity and measures of interdisciplinarity allow us to comment on the strength of relationships between fields—how often one field cites another or builds upon their work. In this work, we focus on the citational diversity of accessibility research, a subfield of HCI that is cross-disciplinary by nature, as it draws influence from not only the broader HCI community, but also from innovations in rehabilitation medicine, gerontology, psychology,

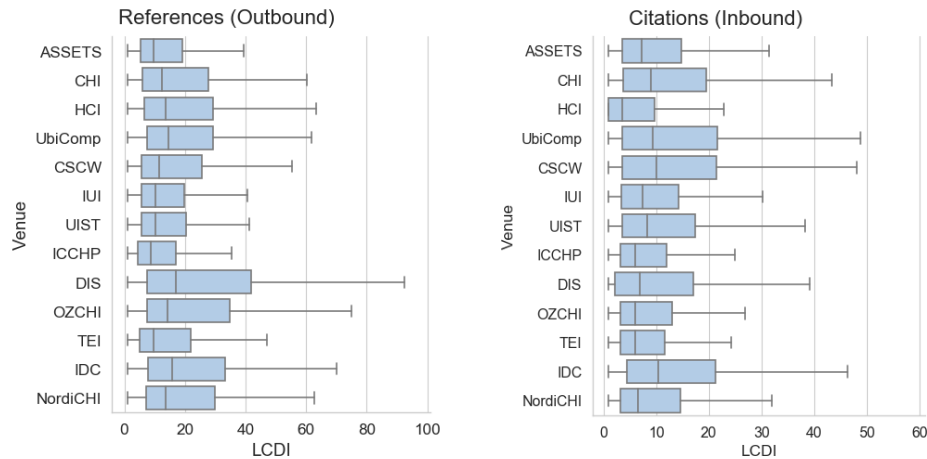


Figure 5: Distribution of the LCDI [15] computed over the L1 fields of study of referenced papers (left) and citing papers (right) for each of 13 venues in the comparative dataset. Higher LCDI indicates higher diversity among the fields of the references or citations. LCDI is similar across venues but is lower for those focused on a particular subdomain (e.g. ASSETS, IUI, ICCHP, TEI) or regional conferences (OZCHI, NordiCHI) and higher for more general HCI conferences (CHI, UbiComp).

education, and more. For the most part, the venue in which an accessibility paper is published, CHI versus ASSETS, does not affect major differences in a paper’s eventual citational impact; the median citation count is similar between the two conferences. Reference and citation patterns between accessibility papers published in the two venues are also similar, perhaps due to an overlap in the authorial community. When applying LCDI over all publications (not just those on accessibility) in ASSETS and CHI, we observe that the relative diversity of ASSETS references and citations are lower than those of CHI. This is unsurprising, since ASSETS is focused on the sub-discipline of accessibility, while CHI represents the broader HCI community. One could conjecture the benefits of both venues: ASSETS focuses on accessibility and papers published there reach a targeted community, while CHI is less thematic but grants exposure to a potentially more diverse research audience among its attendees.

The primary limitation of this study stems from imperfect paper and citation metadata. No database of paper metadata is complete, and we attempt to offset the brunt of this issue by sourcing metadata from two databases. We quantify the bias introduced by data missingness through error analysis, the results of which suggest that there should be minimal impacts to our results. Additionally, citations are only one way in which researchers from different disciplines interact, and they do not fully capture interdisciplinary relationships. Explicit collaborations between authors from different departments, schools, and institutions can also be used as a measure of interdisciplinarity, perhaps in future work.

Another direction for future work is to explore the nature of venue differences and how they impact individual papers. Though we do not analyze authors in this work, the authorial composition of a paper likely contributes to a paper’s reference choices and citation outcomes. Future work could also explore whether a paper’s reference diversity is correlated with its citation diversity, i.e., whether a paper that positions itself as more interdisciplinary actually contributes to the furthering of knowledge across broader fields of study.

6 CONCLUSION

Periodic top-down examination of a field’s relation to other fields can help the community reflect upon the broader impacts of their work. Our analysis of citation diversity for accessibility papers reveals that though these papers are predominantly influenced by other works in accessibility and HCI, they also draw influence from disability studies, psychology, and other fields. Whether an accessibility paper is published in CHI or ASSETS produces little difference in its citational outcome, though the venues as a whole are rather different. ASSETS exemplifies a more targeted venue, focused on research in accessible computing, while CHI, as a general HCI venue, demonstrates higher reference and citational diversity among its publications. There are complementary benefits to publishing in both venues. We also encourage our fellow researchers to continue drawing inspiration broadly, and look to increasing interdisciplinarity as a way of seeking new avenues for innovation in accessibility.

ACKNOWLEDGMENTS

This work was funded in part by the National Science Foundation under grants IIS-1818594 and IIS-1652339 as well as the UW Center for Research and Education on Accessible Technology and Experiences (CREATE).

REFERENCES

- [1] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu A. Ha, Rodney Michael Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler C. Murray, Hsu-Han Ooi, Matthew E. Peters, Joanna L. Power, Sam Skjonsberg, Lucy Lu Wang, Christopher Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. Construction of the Literature Graph in Semantic Scholar. In *NAACL-HLT*.
- [2] A. Cavacini. 2014. What is the best database for computer science journal articles? *Scientometrics* 102 (2014), 2059–2071.
- [3] Constantinos K. Coursaris and N. Bontis. 2012. A Meta Review of HCI Literature: Citation Impact and Research Productivity Rankings.
- [4] T. Dillon. 1995. Mapping the discourse of HCI researchers with citation analysis. *ACM Sigchi Bulletin* 27 (1995), 56–62.

- [5] Sergey Feldman, Kyle Lo, and Waleed Ammar. 2018. Citation Count Analysis for Papers with Preprints. *ArXiv abs/1805.05238* (2018).
- [6] Yassine Gargouri, Chawki Hajjem, V. Larivière, Y. Gingras, L. Carr, T. Brody, and S. Harnad. 2010. Self-Selected or Mandated, Open Access Increases Citation Impact for Higher Quality Research. *PLoS ONE* 5 (2010).
- [7] Fu Guo, Fengxiang Li, Wei Lv, Li Liu, and Vincent G. Duffy. 2020. Bibliometric Analysis of Affective Computing Researches during 1999–2018. *International Journal of Human Computer Interaction* 36, 9 (2020), 801–814. <https://doi.org/10.1080/10447318.2019.1688985>
- [8] S. Haustein, R. Costas, and V. Larivière. 2015. Characterizing Social Media Metrics of Scholarly Papers: The Effect of Document Properties and Collaboration Patterns. *PLoS ONE* 10 (2015).
- [9] L. Holman, D. Stuart-Fox, and Cindy E. Hauser. 2018. The gender gap in science: How long until women are equally represented? *PLoS Biology* 16 (2018).
- [10] Mario Karlovac and D. Mladenic. 2014. Interdisciplinarity of scientific fields and its evolution based on graph of project collaboration and co-authoring. *Scientometrics* 102 (2014), 433–454.
- [11] M. Kas, Kathleen M. Carley, and L. Carley. 2011. Trends in science networks: understanding structures and statistics of scientific networks. *Social Network Analysis and Mining* 2 (2011), 169–187.
- [12] S. Klar, Yanna Krupnikov, J. Ryan, Kathleen Searles, and Yotam Shmargad. 2020. Using social media to promote academic research: Identifying the benefits of twitter for sharing academic work. *PLoS ONE* 15 (2020).
- [13] Konstantinos Koumaditis and Tajammal Hussain. 2017. Human Computer Interaction Research Through the Lens of a Bibliometric Analysis. In *Human-Computer Interaction. User Interface Design, Development and Multimodality*, Masaaki Kurosu (Ed.). Springer International Publishing, Cham, 23–37.
- [14] J. Lee, Andrea Cassano-Pinché, and K. Vicente. 2005. Bibliometric Analysis of Human Factors (1970–2000): A Quantitative Description of Scientific Impact. *Human Factors: The Journal of Human Factors and Ergonomics Society* 47 (2005), 753–766.
- [15] T. Leinster and C. Cobbold. 2012. Measuring diversity: the importance of species similarity. *Ecology* 93 (2012), 477–89. Issue 3.
- [16] M. Ley. 2009. DBLP - Some Lessons Learned. *Proc. VLDB Endow.* 2 (2009), 1493–1500.
- [17] Loet Leydesdorff and Ismael Rafols. 2011. Indicators of the interdisciplinarity of journals: Diversity, centrality, and citations. *Journal of Informetrics* 5, 1 (2011), 87–100. <https://doi.org/10.1016/j.joi.2010.09.002>
- [18] Kelly Mack, Emma McDonnell, Dhruv Jain, Lucy Lu Wang, Jon Froehlich, and Leah Findlater. 2021. What Do We Mean by “Accessibility Research”? A Systematic Review of Accessibility Papers in CHI and ASSETS from 1994 to 2019. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems 2021*.
- [19] Andrea Mannocci, Francesco Osborne, and Enrico Motta. 2019. The evolution of IJHCS and CHI: A quantitative analysis. *International Journal of Human-Computer Studies* 131 (2019), 23–40. <https://doi.org/10.1016/j.ijhcs.2019.05.009> 50 years of the International Journal of Human-Computer Studies. Reflections on the past, present and future of human-centred technologies.
- [20] J. Marshall, Conor Linehan, J. Spence, and Stefan Rennick Egglestone. 2017. Throwaway Citation of Prior Work Creates Risk of Bad HCI Research. *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (2017).
- [21] Alexis-Michel Mugabushaka, Anthi Kyriakou, and T. Papazoglou. 2016. Bibliometric indicators of interdisciplinarity: the potential of the Leinster–Cobbold diversity indices to study disciplinary diversity. *Scientometrics* 107 (2016), 593–607.
- [22] M. Newman. 2001. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences of the United States of America* 98 (2001), 404–9. Issue 2.
- [23] L. Ribeiro, M. Rapini, Leandro Alves Silva, and E. Albuquerque. 2017. Growth patterns of the network of international collaboration in science. *Scientometrics* 114 (2017), 159–179.
- [24] Zhihong Shen, Hao Ma, and Kuansan Wang. 2018. A Web-scale system for scientific knowledge exploration. In *Proceedings of ACL 2018, System Demonstrations*. Association for Computational Linguistics, Melbourne, Australia, 87–92. <https://doi.org/10.18653/v1/P18-4015>
- [25] Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Darrin Eide, Yuxiao Dong, Junjie Qian, Anshul Kanakia, Alvin Chen, and Richard Rogahn. 2019. A Review of Microsoft Academic Services for Science of Science Studies. *Frontiers in Big Data* 2 (2019).
- [26] L. Zhang, R. Rousseau, and W. Glänzel. 2016. Diversity of references as an indicator of the interdisciplinarity of journals: Taking similarity between subject fields into account. *Journal of the Association for Information Science and Technology* 67 (2016).