

From Users to (Sense)Makers: On the Pivotal Role of Stigmergic Social Annotation in the Quest for Collective Sensemaking

RONEN TAMARI, DAOStack, Hebrew University of Jerusalem, Israel

DANIEL A FRIEDMAN, University of California, Davis, USA

WILLIAM FISCHER and LAUREN HEBERT, Veeo, USA

DAFNA SHAHAF, Hebrew University of Jerusalem, Israel

The web has become a dominant epistemic environment, influencing people’s beliefs at a global scale. However, online epistemic environments are increasingly polluted, impairing societies’ ability to coordinate effectively in the face of global crises. We argue that centralized platforms are a main source of epistemic pollution, and that healthier environments require redesigning how we collectively govern attention. Inspired by decentralization and open source software movements, we propose Open Source Attention, a socio-technical framework for “freeing” human attention from control by platforms, through a decentralized eco-system for creating, storing and querying stigmergic markers; the digital traces of human attention.

CCS Concepts: • **Human-centered computing** → **Social content sharing**; *Social tagging systems*.

ACM Reference Format:

Ronen Tamari, Daniel A Friedman, William Fischer, Lauren Hebert, and Dafna Shahaf. 2022. From Users to (Sense)Makers: On the Pivotal Role of Stigmergic Social Annotation in the Quest for Collective Sensemaking. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media (HT '22)*, June 28–July 1, 2022, Barcelona, Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3511095.3536361>

1 INTRODUCTION

The web has become a dominant epistemic environment, shaping peoples’ beliefs and knowledge on a global scale. The web, however, is also currently a severely polluted epistemic environment [12], due to highly centralized and opaque information ecologies, coupled with incentive misalignment and unprecedented information overload. A small number of major web platforms such as Google and Facebook have gained immense control over the means to search, create, and distribute information [13]. Centralization leads to opacity, in which network data as well as algorithms for content creation, search, and distribution are effectively hidden away from public, scientific, and ethical oversight [1].

Platform incentives are fundamentally misaligned with those necessary for healthier epistemic environments [24]. For example, centralization and control of data are necessary for running lucrative “attention markets”, but ultimately hinder attempts to address information overload, and undermine both user autonomy [19] as well as the open information networks necessary for healthy democracies [11, 25]. Platforms are implicated in a host of problematic social phenomena, including the spread of false information, behavioral changes, and societal polarization, epistemic distraction and degradation of individual and collective sense-making capacities [13, 23]. Impending global ecological and societal crises lend increased urgency to addressing these problems: astute collective sense- and decision-making have perhaps never been more needed [1, 22, 23].

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

Manuscript submitted to ACM

Laudable recent efforts have called attention to this precarious state of affairs of collective online sensemaking [1, 11, 13]. However, while providing invaluable insights, they have largely focused on improving platforms through regulatory action, whether internally or externally imposed. Such platform-centric approaches are an important step towards healthier epistemic environments, but face inherent limitations (§3), and are fraught with many impediments as they often run counter to powerful platforms’ core business models. Perhaps more crucially, platform-centric initiatives cannot adequately account for the fundamentally distributed [8], self-organizing [7], and stigmergic [15] nature of collective intelligence.

We argue that reflecting these considerations in practice would benefit from a more radical redesign of our epistemic environments, centered around guiding principles of agency, transparency, interoperability, decentralization, and a collective conceptual transition to a “maker” mindset [10]; from passive users, to more active (sense)makers.

Inspired by both theoretical and practical breakthroughs of decentralization and open source software movements contra entrenched centralized systems, we propose Open Source Attention (OSA), a conceptual framework and “call to movement” towards decentralized, open-source, stigmergic annotation. We envision our framework as a step towards systems for distributed governance, education, and control of collective sense-making and attention.

Hypertext and social annotation play a pivotal role in our proposed transition: in the current *platform-centric* ecology, user annotations (such as likes, retweets, etc) are locked across platforms’ data siloes where they serve to optimize *platform* growth. OSA aims to empower *maker-centric* ecologies by employing distributed content creation and storage technology (e.g., Solid [18]). In this way, makers will control creation and dissemination of their annotations, which can then be leveraged to optimize personalized *human* growth and learning for individuals and collectives.

2 DISTRIBUTED, STIGMERGIC FOUNDATIONS OF COLLECTIVE SENSE-MAKING

Sense-making refers to processes by which agents make sense of their environment, achieved by organizing sense data until the environment is understood well enough to enable reasonable decisions [21]. Theories of extended [5] and stigmergic [15] cognition highlight the integral role of *environment modification* in sense making; agents actively change their environment to assist internal cognitive processes (e.g., writing to-do notes) as well as indirect *stigmergic* communication with others (e.g., ant pheromone trails). Stigmergy is particularly relevant for the setting of collaboration of large-scale groups [6]. In stigmergic communication, the environment acts as a kind of distributed memory; modifications left by others provide cybernetic feedback, driving both *emergence* of novel system-level behavior from local

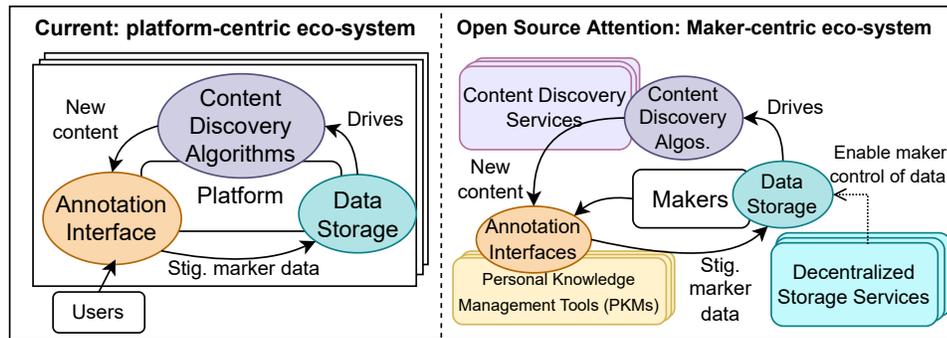


Fig. 1. Platforms leverage control of both data and content discovery algorithms to drive growth at the expense of users (left); Decoupling data and algorithms incentivizes content discovery services oriented towards human-centered growth (right).

interactions of agents, and *immergence* (individual interactions informed by a global state of affairs) [15]. Sense-making is thus inherently co-created, through agents modifying their environment and reacting to changes made by others.

What kind of environment modifications are relevant to consider for sense-making in vast digital spaces? The literature broadly distinguishes between two types of modifications: *sematectonic stigmergy*, which directly alters the environment state (in the digital case: creating new content, such as publishing a blog post), and *stigmergic markers*, which do not directly modify content, but rather serve as signalling cues (in the digital case: likes, annotations, hyperlinking of text). Importantly, stigmergic markers play a central role in assessing epistemic quality of content, both for humans [12, 15] and machines [13], due to the sheer volume of information as well as challenges in endogenous content interpretation. Stigmergic markers may be explicitly left by users (e.g., likes) or implicitly recorded through their behavior (e.g., link click-through data, reading time).

3 OPEN-SOURCING STIGMERGIC MARKERS FOR HEALTHIER EPISTEMIC ENVIRONMENTS

Polluted epistemic environments are often framed as casualties of the “attention economy”; platforms selling user data to advertisers and putting up ads in social media feeds, with the aim of “capturing” users’ attention and seducing them to make yet another purchase. While “data” and “attention” are popular abstractions, the stigmergic perspective is valuable in guiding practical redesign of epistemic environments. Stigmergic markers can be thought of as digital traces of human attention, whose primacy as indicators of epistemic value makes them precious resources, whether for extractive (e.g., ad-tech) or constructive (e.g., collective sense-making) purposes. In the following sections, we illuminate the role of stigmergic markers in nourishing healthier epistemic environments.

3.1 From attention to intention

Healthier epistemic environments involve moving from exploitation of attention to supporting our intentions [23]. This transition requires two paradigm shifts. First, a mindset shift on the human side, from passive, unwitting users consuming “unhealthy information diets” [9], to active makers, who cultivate growth-oriented intentions and are mindful of the (stigmergic) traces they leave, as well as their role as co-creators in the larger digital and physical ecology. Realistically, humans stand no chance of making the transition in isolation; content discovery algorithms are indispensable for navigating vast digital landscapes, but to a large degree are controlled by platforms [13]. Accordingly, the second shift involves re-designing our epistemic environments to support this transition by empowering makers through human-centric content discovery. As shown in Fig. 1, current content discovery is platform-centric: platforms enjoy a closed feedback loop consisting of both the content discovery algorithms as well as the stigmergic marker data needed to drive algorithmic optimization towards platform growth [20]. Human-centric content discovery requires supplanting this degenerative cycle with a more symbiotic information ecology, in which makers create and control their stigmergic markers, and thus are empowered to share their data to content discovery services oriented towards personalized individual or collective growth. Content moderation is an important representative example [16]: moderation is intractable in centralized systems, due to inherent limitations of AI capabilities as well as the scale of complex human adjudications needed. In contrast, decentralized eco-systems enable a “marketplace of filters”, where different individuals and organizations can create and tune content moderation systems for their own needs.

3.2 Open Source Attention: maker-centered information ecology

Analogously to open-source code and common domain knowledge [10], stigmergic markers can be thought of as a public good. However, despite their unique importance, surprisingly little work has specifically targeted their decentralization

(§4). Stated simply; where open source software is a movement to “free” software, similarly OSA is a movement to “free” stigmergic markers, starting from basic hypertext primitives: emotional valence (e.g., likes), bi-directional links, span highlighting, semantic categorization (tags, bookmarks), and textual annotation. We envision decentralized, maker-centered ecologies, comprised of three main architectural elements (see also Fig. 1):

Annotation tools. Enable makers to easily create markers attached to any URL or content element included therein, not just where platforms provide like buttons [3]. Some types of markers should themselves be mark-able, allowing for example the option to “like” a particular annotation, or link between two annotations. Future extensions can address implicit stigmergic markers such as read-time or click-through counts [14]. Apps recording these function as automatic annotation tools, though their implicit nature requires extra caution with regard to consent and data privacy issues.

Self-sovereign storage. Makers own their markers and control their visibility (private, public, etc) to other people or services. Identity provision is a key related service that can (but does not have to be) provided along with storage [20].

Content discovery services. Rather than platforms’ monolithic and opaque feeds, a decentralized ecology encourages a market of diverse and human-centered content discovery services. For example, competing interfaces for social media that better moderate trolls, promote thought-provoking stories, or provide customizable feed controls [16].

4 DISCUSSION

While the idea of leveraging stigmergic markers for collective sense-making has a long history [2], most contemporary open-source and decentralization efforts have focused on sematectonic (content-creating) stimergy, such as code, social media, financial ledgers and executable contracts [4]. Closest to our proposal is the Solid eco-system [18] that, similarly, targets “re-decentralizing the web” [20], and empowers individuals to control their data. Solid also features a marketplace of services, including the dokieli decentralized annotation client for scientific research [3]. While Solid and dokieli are inspiring initial steps, they are limited with regard to content discovery services or social incentives. More broadly, where Solid is primarily a technology, OSA proposes an ecological perspective accounting for the embeddedness of such technologies in wider social, educational and economic contexts. For example, a key extra-technological challenge concerns changing norms around knowledge work. Similarly to how platforms changed the culture around certain kinds of content creation, effectively turning us all into performers, well designed social networks could help shape the norms and prestige associated with sense-making activities. Academic Twitter demonstrates that even without direct economic incentives, social incentives lead experts to freely share high-quality information publicly [17].

Another key question concerns scale: for global-scale sense-making, any proposal must necessarily compete with massive, well established platforms. While recent years are seeing a resurgence of *personal* knowledge management apps (PKMs) enabling content creation and annotation, knowledge tends to remain siloed at the individual level; adaptation to *collective* knowledge management (CKM) has been limited.¹ OSA is naturally congruent with the promising “protocols, not platforms” approach [16]; rather than head-to-head competition between PKMs, existing PKM growth can be bootstrapped for CKM by introducing interoperable protocols and storage for stigmergic primitives (e.g., links, tags). In this way, data from across diverse PKM apps could be shared to contribute to collective sense-making efforts.

Many open questions remain out of scope of this short piece, which is best seen as a call to attention; success in surmounting the formidable challenges faced today by humanity requires that “we give the right sort of attention to the right sort of things” [23]. We have claimed that attending “to the right things” will require re-imagining our

¹<https://athensresearch.ghost.io/season-2/>

socio-technological systems for governing collective attention; we hope our proposal will help galvanize action towards this vital cause.

ACKNOWLEDGMENTS

We thank Nimrod Talmon and the DAOStack team for thoughtful feedback and support, and we thank Metagov and RadicalXChange for cultivating the wonderful real-world and online spaces that seeded this collaboration.

REFERENCES

- [1] Joseph B. Bak-Coleman, Mark Alfano, Wolfram Barfuss, Carl T. Bergstrom, Miguel A. Centeno, Iain D. Couzin, Jonathan F. Donges, Mirta Galesic, Andrew S. Gersick, Jennifer Jacquet, Albert B. Kao, Rachel E. Moran, Pawel Romanczuk, Daniel I. Rubenstein, Kaia J. Tombak, Jay J. Van Bavel, and Elke U. Weber. 2021. Stewardship of global collective behavior. *Proceedings of the National Academy of Sciences* 118, 27 (2021), e2025764118. <https://doi.org/10.1073/pnas.2025764118>
- [2] Vannevar Bush. 1945. As we may think. *The atlantic monthly* 176, 1 (1945), 101–108.
- [3] Sarven Capadisli. 2020. Linked research on the decentralized Web. (2020).
- [4] Fran Casino, Thomas K Dasaklis, and Constantinos Patsakis. 2019. A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics* 36 (2019), 55–81. <https://doi.org/10.1016/j.tele.2018.11.006>
- [5] Andy Clark and David Chalmers. 1998. The extended mind. *analysis* 58, 1 (1998), 7–19.
- [6] Mark Elliott. 2006. Stigmergic Collaboration: The Evolution of Group Work: Introduction. *M/C Journal* 9, 2 (may 2006). <https://doi.org/10.5204/mcj.2599>
- [7] Nigel R. Franks and J. L. Deneubourg. 1997. Self-organizing nest construction in ants: individual worker behaviour and the nest's dynamics. *Animal Behaviour* 54 (1997), 779–796.
- [8] Todd M. Gureckis and Robert L. Goldstone. 2006. Thinking in groups. *Pragmatics & Cognition* 14 (2006), 293–311.
- [9] C Johnson. 2011. *The Information Diet: A Case for Conscious Consumption*. O'Reilly Media. <https://books.google.com/books?id=QrW62y9l3lYC>
- [10] Vasilis Kostakis, Vasilis Niaros, George Dafermos, and Michel Bauwens. 2015. Design global, manufacture local: Exploring the contours of an emerging productive model. *Futures* 73 (2015), 126–135. <https://doi.org/10.1016/j.futures.2015.09.001>
- [11] Anastasia Kozyreva, Stephan Lewandowsky, and Ralph Hertwig. 2020. Citizens Versus the Internet: Confronting Digital Challenges With Cognitive Tools. *Psychological Science in the Public Interest* 21, 3 (2020), 103–156. <https://doi.org/10.1177/1529100620946707>
- [12] N Levy. 2021. *Bad Beliefs: Why They Happen to Good People*. OUP Oxford. https://books.google.com/books?id=C%5C_ZQEAAAQBAJ
- [13] Philipp Lorenz-Spreen, Stephan Lewandowsky, Cass R Sunstein, and Ralph Hertwig. 2020. How behavioural sciences can promote truth, autonomy and democratic discourse online. *Nature human behaviour* 4, 11 (2020), 1102–1109.
- [14] Artur Sancho Marques and José Figueiredo. 2013. Stigmergic hyperlink: A new social web object. *Information Systems and Modern Society: Social Change and Global Development* (2013), 260–272. <https://doi.org/10.4018/978-1-4666-2922-6.ch016>
- [15] Leslie Marsh and Christian Onof. 2008. Stigmergic epistemology, stigmergic cognition. *Cognitive Systems Research* 9, 1-2 (2008), 136–149. <https://doi.org/10.1016/j.cogsys.2007.06.009>
- [16] Mike Masnick. 2019. Protocols, Not Platforms. *Knight First Amendment Institute* (2019).
- [17] Daniel S. Quintana. 2020. Twitter for Scientists [eBook edition]. <https://doi.org/10.5281/ZENODO.3707741>
- [18] Andrei Sambra, Amy Guy, Sarven Capadisli, and Nicola Greco. 2016. Building Decentralized Applications for the Social Web. In *Proceedings of the 25th International Conference Companion on World Wide Web* (Montréal, Québec, Canada) (*WWW '16 Companion*). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1033–1034. <https://doi.org/10.1145/2872518.2891060>
- [19] Chirag Shah and Emily M. Bender. 2022. Situating Search. In *ACM SIGIR Conference on Human Information Interaction and Retrieval* (Regensburg, Germany) (*CHIIR '22*). Association for Computing Machinery, New York, NY, USA, 221–232. <https://doi.org/10.1145/3498366.3505816>
- [20] Ruben Verborgh. 2022. Re-decentralizing the Web, for good this time. In *Linking the World's Information: A Collection of Essays on the Work of Sir Tim Berners-Lee*, Oshani Seneviratne and James Hendler (Eds.). ACM. <https://ruben.verborgh.org/articles/redecentralizing-the-web/>
- [21] K E Weick and K E W Weick. 1995. *Sensemaking in Organizations*. SAGE Publications. <https://books.google.com/books?id=nz1RT-xskeoC>
- [22] Jevin D West and Carl T. Bergstrom. 2021. Misinformation in and about science. *Proceedings of the National Academy of Sciences* 118, 15 (apr 2021), e1912444117. <https://doi.org/10.1073/pnas.1912444117>
- [23] James Williams. 2018. *Stand out of our light: freedom and resistance in the attention economy*. Cambridge University Press.
- [24] S Zuboff. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs. <https://books.google.com/books?id=IRqrDQAAQBAJ>
- [25] Ethan Zuckerman. 2020. The Case for Digital Public Infrastructure. *The Tech Giants, Monopoly Power, and Public Discourse: An Essay Series by the Knight Institute, Columbia University* (2020). <https://knightcolumbia.org/content/the-case-for-digital-public-infrastructure>