

This version of the article has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at:

<https://link.springer.com/article/10.1007/s11245-024-10139-3>

Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use <https://www.springernature.com/gp/open-research/policies/acceptedmanuscript-terms>

“A place of very arduous interfaces”¹. Social media platforms as epistemic environments with faulty interfaces

By Lavinia Marin (TU Delft)

Email: l.marin@tudelft.nl

Abstract

I argue that the concept of an epistemic interface is a useful one to add to the epistemic ecology toolkit in order to enrich our investigations concerning the complex epistemic phenomena arising on social media. An epistemic interface is defined as any informational interface (be it technical, human or institutional) that facilitates the transfer of epistemic goods from one epistemic environment to its outside, be that another epistemic environment or a person. When assessing the kinds of epistemic environments emerging on social media platforms, we should not only look at their epistemic health or hostility as environments, but also look at how these environments interact with each other, at their interfacing. I show that several epistemic problems specific to social media such as epistemic bubbles, echo chambers and epistemic flooding can be understood also in terms of faulty epistemic interfacing, which adds a layer of complexity to the discussion. The paper proposes two criteria for healthy epistemic interfacing: preserving the flourishing of epistemic agents, and preserving in some form the value of epistemic outcomes as these cross the environment to the outside world. These criteria should give rise to specific norms and design requirements for healthy interfacing, but such norms will need more empirical work before these get formulated.

1. Introduction. Social media platforms as hostile epistemic environments or a hostile ecosystem?

Mainstream social media platforms such as Facebook, YouTube, X, TikTok, Instagram, Reddit, etc., are increasingly recognised as epistemic environments (Battaly, 2018), namely spaces of epistemic interaction between various agents, be they human or artificial (Floridi, 2013). These social media platforms (SMPs from now on) are increasingly seen as hostile or detrimental epistemic environments for their users (Nguyen, 2023) whereby the epistemic hostility of SMPs consists in them being environments with “features [that] exploit our

¹ This is from a quote from Marshall McLuhan. The full quote goes: “The global village is a place of very arduous interfaces and very abrasive situations.”(cited from McLuhan & Powers, 1989)

cognitive vulnerabilities.” (Nguyen, 2023, p. 2) The challenge at hand for epistemic ecology (Carmona, 2022) and for social epistemology at large is to comprehensively describe and explain these hostile features of the online epistemic environment.

What is epistemically hostile about SMPs has received several explanations targeting a variety of hostile design features or phenomena emerging from the interaction between agents in the online world: fake news and misinformation sharing at a large scale (Frost-Arnold, 2023; Harris, 2024; Rini, 2017), epistemic flooding (Anderau, 2023), echo chambers and epistemic bubbles forming online (Nguyen, 2020), ignorance propagation (Arfini et al., 2019), decontextualised sharing and lack of epistemic norms for sharing (Record & Miller, 2022), context collapse (boyd 2008, Frost-Arnold 2021), unusual networked topologies of influence (Alfano, 2016), gamification of values (Nguyen, 2021), confusing social reputation for ranking (Origgi, 2018), mistaking others as their beliefs (Figà Talamanca & Arfini, 2022), artificial users mimicking human actors while acting in bad faith (Ienca, 2023), and several others. These explanations have their merits in unpacking the complex phenomena arising on social media platforms (SMPs), and these have been used to build an emergent field of inquiry into the social epistemology of social media. Still, there seems to be something not yet covered by these explanations and which the current paper seeks to address: the matter of how online epistemic environments interface.

When discussing SMPs as epistemic environments, we need to be careful about the terms we use and whether we use a plural or a singular. There is no single social media environment, nor can we imply that a particular social media platform constitutes a single epistemic environment. First, while there is a multitude of social media platforms out there working as social environments, we cannot say that one platform maps unto one social environment. As social media scholars have argued, we should instead say that these are personal “social media ecosystems” (Carter et al., 2023) or social media ecologies (Zhao et al., 2016) to do justice to the fact that every online user uses several platforms at the same time, connecting them in a variety of ways, and also noticing that the same platform is experienced by a user as a multitude of social environments, depending on how many groups, pages or influencers they follow. SMP users do not experience one platform as one clearly delimited social environment, rather they mix and match the platforms, seamlessly transitioning from one platform to another, and experiencing this flow as their own social ecosystem. We can think of how people carry a conversation across multiple media, from an instant messaging app to a post, to a comment thread, and then in a face-to-face dialogue in real life. Secondly, the distinction online/ offline is analytically correct for research purposes, yet human social life is a mix of online and offline experiences, with often the “onlife” being undistinguishable from the offline life (Floridi, 2015). Because there is a genuine continuum of online and offline social experiences, it becomes difficult to state that SMPs as a whole world or one platform in particular constitute a social environment, at least from the perspective of what the social media user experiences. If SMPs are a multiplicity of social environments mixing and flowing into each other, then we should be able to say the same about SMPs *as epistemic environments*: not one, but many; not static, but in a flux, interacting. What does this state of multitude mean then for ascribing epistemic hostility to these environments? We need to look at how these environments interact with each other and with other epistemic environments. For this purpose, we need to flesh out the concept of an *epistemic interface*.

2. The concept of an epistemic interface: technical and human interfaces

Epistemic interface is a concept derived from the related concept of an informational interface, widely used in human-computer interaction (HCI) studies. An informational interface is a "space[s] of interactions between agents (human, artificial, or hybrid) who want something from us, and our resources, something that we possess and that these other agents want or desire (at the very least in theory) for themselves." (Floridi, 2023, p. 95) Humans, as well as designed artefacts can act as informational interfaces. If the informational system is technical, the informational interface is found at the outermost layer of an informational system with which users interact, a layer usually called a user interface (see Figure 1 below). Every informational system has at least one user interface, if it is intended to be used by humans². In these cases, we call it an "user interface" or simply interface to describe this space of interactions between the human user and the system.

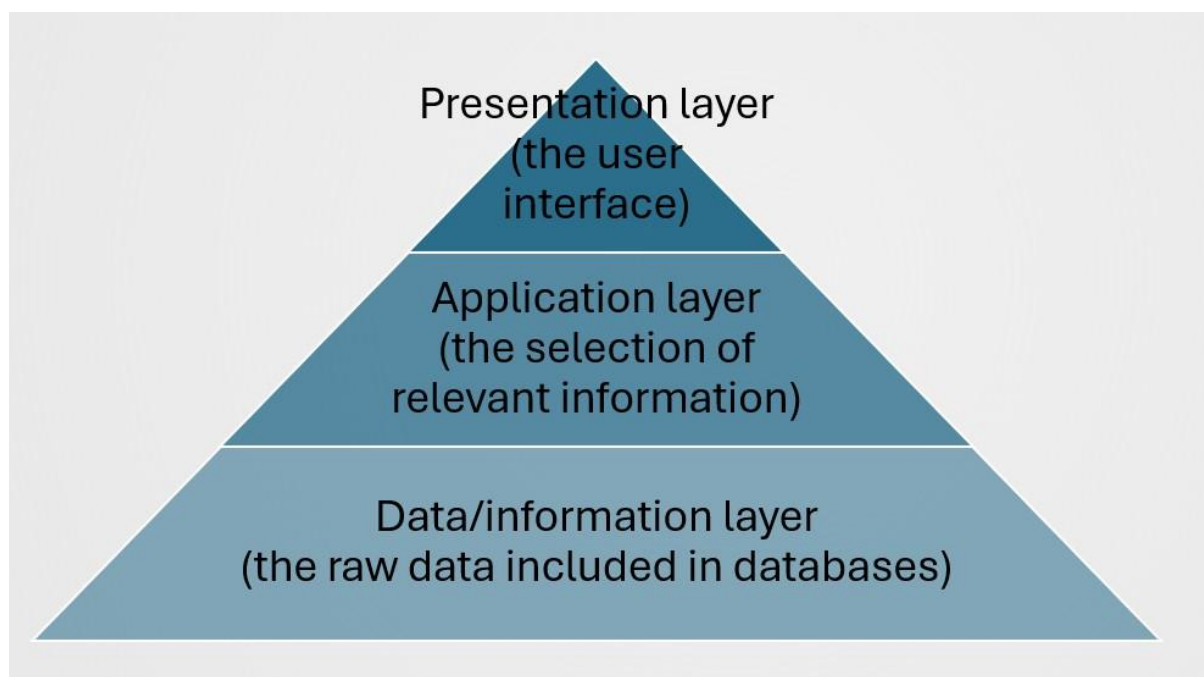
Technical informational interfaces

Interface studies in HCI have mostly focused on the concept of a user interface as a site of affordances for the user – what the user is allowed to do, nudged or hindered in doing (Tollon, 2022). The user interface makes some information become user accessible. Behind every website are thousands of lines of code that encapsulate the information as well as the possible actions that a user could do with that information. The user interface mediates the user's access to an informational resource with various ends in sight: control of information, access to information, doing things with said information, interacting through the information with other users, entertainment, knowledge or understanding, etc. Simply put, without the user interface, the user would not be able to access said information without coding skills. A similar idea is expressed in media studies, where interfaces are understood as "the ways to represent ('format') and control the signal" (Manovich, 2013, p. 12). Regardless of whether we approach interfaces informed by philosophy of information, HCI, or media-studies, all three fields see informational interfaces in roughly the same way: as transitional spaces between the agent and some informational resource, where the space is never fully transparent and introduces a distortion in how the agent perceives the information. Any user interface mediates and transforms as it transfers information. Thus, interfaces manipulate, curate, select and represent information in ways consistent with specific models of the world and certain values and norms (which are not ethical in themselves, we can think of norms of translation of what counts as good or bad access to information). For something to count as an epistemic interface, in addition to being an informational interface, it also needs to mediate some epistemic process for its users, be that upgrading of information to knowledge, information gathering, justification practices, or to shape their experience of epistemic agency.

A few examples of technical informational interfaces will clarify the concept further such

² An informational system without human users could be a huge dataset that is trawled and analysed by machine learning algorithms in view of making predictions. The algorithms do not need access to the dataset in any human-readable form, hence that system has no user interface. However, once there is an output of the machine learning algorithm, this output has to go through an user interface to be translated for human eyes, and then a user interface can be added on top of the system.

that I can build upon it next to define epistemic interfaces more analytically. A search engine (Google, or Bing, Yahoo, etc.) is an informational interface that most of us encounter daily. When we do a search query in this informational system, the data and the algorithms behind the interface compile and collapse unimaginable quantities of information into a ranked list of websites that will contain what we are looking for. To understand what is retrieved and shown to the user, we can look at Figure 1 for understanding the general architecture of an informational system with a tripartite structure, a model used widely in informational systems (see Cheong, 2023): The data layer (with all the raw data available for an informational system), the application layer (which does some selection of the data in view of some criteria, also translating it into information), and the presentation layer which is the user interface (this is where the information is accessed and understood).



In the case of the search engine, the data is stored in data tables (SQL or other formats), on some servers. The search engine uses an algorithm to retrieve that data when a search query is initiated, and then to also rank the pages into an order that is then presented to the user in the window with the results. Thus, the informational interface of the search engine is both the query field and the page listing the results. The search engine's interface shows to the user some results translated from a computer-readable format to a human-readable format. The user interface translates when it renders information, and, with the ranking of relevant information being done behind the interface by the application layer, the search engine algorithms that decide what is relevant or not and the order of importance.

Another example of an informational interfaces is ChatGPT or any other LLMs that users can query through a chat window or search field. ChatGPT answers questions in a long answer format, and some users will use it just like a search engine. Except that ChatGPT does not tell the users where it got this information and how it was compiled, nor does it give any indication of when it is hallucinating its answers or when it is compiling a variety of reliable sources. Because of this opacity of sources used and of how these are combined, ChatGPT is

an unreliable source of information and, furthermore, when used like a search engine to look for actual answers, is a problematic informational interface. All technical user interfaces mediate and interpret the information that is hidden from the user by selecting and prioritising certain norms in their handling of information, which often has epistemic consequences, depending on the values and norms used to shape these processes of translation.

Epistemic interfaces of the technical kind

Epistemic interfaces are spaces of interaction between two epistemic environments or an epistemic environment and an outsider to that environment. The interaction can be one of exchange - something is exchanged through the two sides, usually information, but also norms, epistemic practices, data, testimony, etc. - but it also can be one of connection or putting resources together. To distinguish between mere informational interfaces and epistemic interfaces, I add the condition that the interface has to have some effect on the epistemic agency³ of one of the parties involved, be it a user or the environment. Going back to the previous two examples, if a someone constantly uses the search engine as a source of reliable information that then they incorporate and build into knowledge⁴, this informational interface is, for this user, an epistemic interface. The same holds for ChatGPT: if some people use it as a source of information, or as a source of testimony (Freiman, 2024), relying on it as if it were an expert, then ChatGPT is an epistemic interface for them. We can of course argue that this is not genuine access to expertise, and this epistemic interface is faulty in some way, but for these users, ChatGPT is an integral part of how they build knowledge and justify it, so for their epistemic agency, the interaction with ChatGPT is essential. For my concept of an epistemic interface, it matters if the user(s) experience the interface as bearing on their epistemic agency, not if the interface was designed with such purposes in mind. Something can be an epistemic interface even when systematically mis-used or misunderstood. But why not simply say that the search engines and ChatGPT are simply (unreliable) sources of testimony, instead of saying that these are epistemic interfaces? Because, and this is a crucial difference, the search engine and ChatGPT mediate the transfer of information from a different epistemic environment than the user is currently inhabiting. If I am a student at the university and the instructor says something I do not understand, but then I rely on it, this is only about testimony and information picked up in the same environment because, as a student, I am supposed to know the epistemic norms of testimony governing my local environment; the instructor's testimony is not that alien to me. With technical epistemic interfaces, most users find themselves as outsiders to a system that they have no idea about: no knowledge of where its information comes from, not inking of its epistemic norms, no clear contribution can be made by the user to this epistemic environment that is merely presented for them to use, not to build or act upon.

The concept of an epistemic interface should help when evaluating how various epistemic environments interact among themselves or with their outsiders. I will argue further that it

³ Epistemic agency is defined here as the human capacity to decide and be responsible for the beliefs one holds, assuming that people build their own knowledge through autonomous choices, and hence that people are accountable for their beliefs and how they arrived at them (Reider 2016, p. ix).

⁴ For how information upgrades to knowledge, see Luciano Floridi's (2012) account: information is upgraded to knowledge is the agent can provide a network of account for why that information is true.

is not enough to say that an epistemic environment is healthy, it matters also how this environment interfaces with its outside in order to evaluate its epistemic hostility. We can thus distinguish then between healthy interfacing and hostile or toxic interfacing. Hostile epistemic interfacing can happen in a variety of ways, not limited to but, for example, by silencing or nullifying the epistemic outcomes of one environment when these should carry over to another environment, or by hindering the epistemic agency of agents crossing two environments, and many other ways.

Humans and institutions as epistemic interfaces

An epistemic interface can be any device or agent (be it human or non-human) that enables the connection between the two epistemic environments by allowing the communication/ transfer of the epistemic outcomes from one environment to another. The interface need not be technical, it can be enacted by a single human (Floridi, 2023), groups or institutions. Institutions are sets of rules intended to solve certain collective problems (Werner, 2022) and which are visible in how they constrain or enable the local access to the cognitive resources needed to solve that problem (Werner, 2022, pp. 168-169). In this definition, institutions act as epistemic interfaces inside the same epistemic environment, without there being any visible user interface that one needs to access. Rather, the user/ citizen gets access to the cognitive resources needed by following the rules prescribed by that institution. The institution is experienced as an epistemic interface between the solitary user and the epistemic environment which is experienced as alien to the user, hence in need of some sort of translation. We can think here of Kafka's character K. trying to understand how to get to talk to someone in the Castle, or of an applicant trying to understand how to apply to a funding call, trying to elucidate what all these forms mean for those who wrote them and for the jury.

Standard examples of institutional epistemic interfaces are libraries; in there, users get access to a scarce epistemic resource (the books and journals) and sometimes they get training in how to access and use these resources, as many libraries offer workshops on how to do a literature search, how to differentiate the predatory journals from legitimate scientific resources, how to use referencing software, how to do a literature review, how to collect notes, etc. The aims of institutional epistemic interfaces can be educational when they empower outsiders to gain the skills to be self-sustaining in performing certain epistemic tasks, or to facilitate understanding of a specific set of cognitive resources they hold monopoly over (e.g. what is a grant-worthy application, what counts as innovation in a field). A person can act as an epistemic interface for others (Floridi, 2023) if they act as representative of a different epistemic environment. A teacher at a parent-teacher conference is the interface between the school's epistemic environment and the outside world. An accreditation expert is the epistemic interface between an institution seeking educational accreditation and the Ministry of Education, which represents the epistemic priorities of a society - what do we want from educational institutions, and who should get to call themselves one? Yet most non-technical epistemic interfaces are attached to an institution or organisation and this give them authority, since the institution has the power to enforce epistemic norms by granting or withholding access to its resources.

Epistemic interfaces have a variety of functions: to preserve the epistemic norms native to

the epistemic environment, to defend these norms to outsiders, translate their outcomes to the outside world, and to make these epistemic outcomes usable by outsiders. The teacher helps students gain knowledge and skills they can use in other epistemic environments which pupils inhabit in their life after school. Yet the school's epistemic norms – norms about what counts as knowledge, for example – do not carry so well outside, this is why teachers, schools, and educational systems get a constant push-back from society. The norms of an epistemic environment do not carry easily over to another environment, but the epistemic outcomes do, insofar as these are seen as useful by society or other actors.

Epistemic interfaces also work as heuristics for environments: only when we are in the presence of an interface do we realise that two epistemic environments are interacting there. For example, the teacher is an epistemic interface whenever they interact with the parents and society in their role as a teacher. It is during the discussions between teachers and parents that one may find out that some parents are flat-earthers or anti-evolutionists because during those meetings, the school teachings are discussed as epistemic objects, something to be contested and debated. The theories taught are taken for granted during school hours by the pupils, but only when there is interaction with outsider which contest those theories –parents, NGOs – then an epistemic interface needs to be enacted by the teacher, sometimes by the school principal.

Epistemic interfaces are the connecting tissue between two epistemic environments and what makes the interaction possible such that epistemic outcomes can get outside the environment. No healthy epistemic environment is an island, rather, each epistemic environment interacts with other pre-existing epistemic environments, and if this interaction is constructive, environments working together become ecosystems that transfer epistemic outcomes reliably and predictably from one environment to another. The interesting cases emerge when the epistemic environments are not functioning together in an ecosystem, seamlessly interfacing with each other, but instead are at odds, trying to nullify each other's epistemic outcomes. We can think of a scientific research group on evolution and a religious community endorsing creationism: they cannot interface effectively since they have no use for the other's epistemic outcomes; if these communities are enclosed epistemic environments, there is no actual interfacing in between them.

My aim here is to show that the concept of an epistemic interface helps with normatively assessing the interactions between epistemic environments without borrowing epistemic norms from one epistemic environment and carrying them illegitimately into another epistemic environment that would find them as alien. The next section briefly summarises the ways in which social media platforms are understood as hostile epistemic environments in the ongoing social epistemological discussions, and then I will build on top of these my account of interfacing.

3. Hostile epistemic environments

Mainstream online social platforms are considered to be epistemic environments on their own by social epistemologists (Ryan, 2021; Anderau, 2023; Levy, 2023; Amico-Korby et al., 2024), and often cast in a negative light. The concept of epistemic hostility, made popular by Thi Nguyen (2023), was inspired by Heather Battaly's coining of the term "epistemically

hostile environments” (Battaly, 2018, p. 23). Nguyen defines an epistemically hostile environment one in which the agent’s “good-faith efforts [towards truth] are being subverted by features of the world.” (Nguyen, 2023, p. 2) This subversion of our efforts can be intentional or not, it does not matter to qualify for the status of epistemically hostile environment. Such environments are those whose “features exploit our cognitive vulnerabilities” (Nguyen, 2023, p. 2). According to Nguyen’s definition and examples, an epistemically hostile environment is cognitively hostile if it takes advantage of the users’ tendencies to rely on cognitive heuristics.

To identify what makes SMPs epistemically hostile, some cognitive tactics that target the users have been already discussed: dark patterns in the user interface (Gray et al., 2018), click-bait, rage baiting, algorithmic selection of information, short form content, infinite scrolling, attention capturing and manipulation – all examples of features of the informational environment that make being on social media an overall cognitively exhausting experience (Anderau, 2023; Srinivasan, 2023). But do these features make social media be an epistemically hostile environment overall (or ecosystem, for that matter)? The problem is here of distinguishing between cognitively hostile interactions (such as those facilitated by dark patterns⁵) between the user and the platform, and then generalising this to a pattern of hostile interactions and then inferring that this is what makes the platform epistemically hostile. At first sight, this seems a legitimate move. But upon closer examination, we should notice that the social media ecosystem is only hostile for some users, while others thrive as epistemic agents and are not affected by it. I do not think we can infer from an environment being designed to be cognitively hostile for its users to it being automatically also epistemically hostile for them. Experiencing an environment as epistemically hostile depends on the user and their intent when being on social media. If the user wants to get information that matters for a decision or for building their knowledge, and they are constantly distracted by irrelevant content, then this should count as an epistemically hostile environment. *For them*. But if they merely want to have fun, then SMPs are not epistemically hostile for those users in that context of use even though it can be claimed that these environments are systematically cognitively hostile, occasionally becoming epistemically hostile⁶.

For my current purpose of defining epistemic hostility of an environment, any environment that systematically provides false information, misleadingly de-contextualised information, or manipulates its users to arrive at false beliefs should be seen as epistemically hostile. However, we need to notice that SMPs are not that systematic in providing their users with false or manipulative information. Online platforms expose their users to both false and

⁵ Dark patterns in user design are those “instances where designers use their knowledge of human behavior (e.g., psychology) and the desires of end users to implement deceptive functionality that is not in the user’s best interest” (Gray et al., 2018, p. 1)

⁶ One could argue that passively encountering false information while scrolling online is epistemically harmful in the long run because this information will be recalled at some point, inadvertently or not [Author, 2021]. In this vein, Miller and Record (2013) have argued that users are and should be held responsible for the mistaken beliefs they passively gather online (Miller & Record, 2013, p. 118). They argue that at least the user should be held responsible for entering such a hostile environment, even if they do not know how it works or why it misleads them (Miller & Record, 2013, p. 131). Encouraging users to be irresponsible about the information they acquire online is a way of undermining their epistemic agency. This is a relevant objection, but I do not have the space to tackle it here.

true information, both to manipulative claims and true ones, sometimes even useful. It is possible to use SMPs in a way that helps one become a better epistemic agent, learn something and refine one's epistemic skills, ultimately improving one's epistemic agency. SMPs should not be considered as epistemically hostile just because the information found in them is sometimes misleading or manipulative for some users. If SMPs were the only way to reach a certain kind of information, then we these would need to be more reliable but, as it stands, SMPs are one source of information among many others, with most others being more reliable than SMPs.

One major difficulty in giving an account of the epistemic hostility of SMPs as environments comes from the difficulty of defining how a healthy epistemic environment should look like. We need to look at what counts as epistemic environment to try to understand what counts as environmental health there. Recently, Ryan has defined an epistemic environment as something "constituted by facts, whether social or physical, that bear on epistemic attainment (gaining knowledge, understanding, and so on) in that environment." (Ryan, 2021, p. 607) Healthy epistemic environments facilitate the "epistemic attainment" of their inhabitants⁷. This can be spelled out further as "what counts as an epistemically healthy environment will also depend on the epistemic capabilities, tools, and resources of the agents that occupy them" (Amico-Korby et al., 2024, p. 3), in other words the environment can maximise the epistemic attainment of its inhabitants only insofar as these have certain skills and access to resources. We can assume that if an environment hinders the epistemic attainment of its members, it should be seen as hostile, while if it facilitates it, probably it's a healthy epistemic environment. This definition of health works well for describing the epistemic environment as it interacts with its native inhabitants, the agents that rely on this environment and use this for their epistemic attainment. To assess the epistemic attainment of the inhabitants and how well this is facilitated, we should know the internal epistemic norms of that environment. For example, we can assume that what counts as epistemic attainment differs from a scientific research group to a theology seminar. Critical thinking probably not so much an epistemic attainment in the theology seminar as much as agreeing to the dogma and skilfully interpreting it to overcome the cognitive dissonances. While, from the outside, it may seem that a theology-infused environment is hostile to epistemic attainment, we should be careful not to impose our own external norms, whereby critical thinking is always seen as an epistemic good. There is epistemic attainment in the religious seminars, it just looks different from what one would expect in a more scientific-oriented environment. We need to be immersed in an epistemic environment to understand what counts as attainment there because epistemic attainment is grounded in the local language games and forms of life. What holds together one epistemic environment seems to be its epistemic norms and the ability of certain actors to enforce them. An epistemic environment with institutions attached to it is much easier to delimit since we can point to the agreed norms, epistemic goods and the enforcers of those norms, hence also the interface becomes obvious. Yet with SMPs, we are in the presence of a wild epistemic ecosystem, with no clear institutions to mediate and enforce norms [Author, 2024].

⁷ The notion of epistemic attainment is also favored by other epistemologists because it seems to remain neutral to what is being attained, the epistemic goods - see (Anderau, 2023, p. 3).

Healthy epistemic interfacing

Now that we know what counts as epistemic healthy environment, we also need a definition of healthy interfacing between an environment and the outside world (be that an individual or another environment). Even if we have clarity about the internal epistemic norms, these do not give an external reference point to evaluate the healthy from a hostile interfacing between environments. Amico-Korby and colleagues distinguish between two types of norms in epistemic environments: the environment-specific ones and those that apply to all environments (Amico-Korby et al., 2024, p. 3), hence external or universal. Such norms are not clearly defined but these should be “at least consistent with the epistemic flourishing of individuals” (Amico-Korby et al., 2024, p. 16). It is indeed difficult to give a universal account of epistemic norms across all environments, since flourishing will look differently across environments. But we can use this minimal account oriented towards flourishing of epistemic agents to infer what healthy epistemic interfacing should look like.

I offer two conditions for healthy epistemic interfacing: a) the environment-to-environment interfacing and b) individual-to-environment interfacing.

- a) Healthy epistemic interfacing between two environments happens when local epistemic outcomes (or goods) from an environment are carried over to another epistemic environment. Interfacing comes with translation and mediation, so we should not expect these goods to be carried over in their original form, some transformation is required. But what matters here is that not all information lost in the interaction between the two distinct environments, and that something valuable for one environment gets picked up by the other one.
An example will clarify this: in a healthcare context we have the environment of medical professionals, ruled by medical epistemology, and that of the patients and their caretakers at home. If the untrained care-taker notices something off with the patient and reports this to the medical staff, there is an encounter of two epistemic environments. If the medical professionals ignore the observations of the caretaker, because supposedly these sound naïve and are not well-articulated, the patient will suffer, but it can also said that the two environments did not interface, since no there was no use for the observations of the care-taker. However, if the medical professionals translate these observations into clinically relevant terms and then acts upon them, we have a successful epistemic interfacing. In this successful interfacing, the norms of one environment will not be carried over, just the outcomes (in this case, the observation that there is something off with the patient). The example was about one individual interfacing with other professionals; to qualify this as a successful interfacing of epistemic environments we would need to have some rules in place that allow the medical environment to systematically uptake and transform the patient-world outputs, not leaving it to the goodwill of a few professionals.
- b) An environment successfully interfaces with an outside individual when their individual’s epistemic flourishing is preserved to some extent in the new environment. Taking the previous example again, the medical environment would need, as a minimal condition for healthy interfacing, to take seriously the caretakers and patient’s input in a systematic way, and then, as a maximal condition, it could also explain to them why their observations matter and what these signify. This kind

of explanation would enhance the outsider's epistemic flourishing. Any environment that seeks to educate and upskill those who are not its regular inhabitants is also contributing to their epistemic flourishing and thus, interfacing in a healthy way.

4. Social media platforms and epistemic hostility as faulty interfacing

SMPs constitute a novel kind of epistemic ecosystem: never before in the history of humanity was everyone's epistemic attainment visible to others, and never before was it so easy to influence other agents' epistemic dispositions without being an epistemic authority. Before the advent of the SMPs, we never knew what people in our community actually believed, only sociologists could claim to have an overview of the distribution of people's beliefs over an area. After the emergence of SMPs, with the localised version of instant messaging groups for the neighbourhoods, we gained access to what our neighbours believed and liked, and to how they justified those beliefs. What holds for the neighbourhood applies to our communities, both distant and close, family and friends, even influencers that we follow online. After SMPs became widely used, we were suddenly more aware of the epistemic environment we inhabit because we gained the technical means to see what hundreds - perhaps thousands - of social media users claim to believe and how they justify those beliefs. And their beliefs influenced us, whether we wanted or not. This is not to say that any SMP user has a reliable overview of what others believe, their perspective is still distorted and fragmentary, ordered by what personalisation algorithms want them to see. But, through SMPs, we are certainly exposed to more beliefs than we can handle, and we are more aware of how differently people think from us. This is happening primarily because of two features that SMPs uniquely have: epistemic transparency and emotional contagion.

SMPs seem to be transparent environments: we have a heightened awareness of what others (claim to) believe and others have access to our own beliefs. Being aware about the epistemic dispositions of those in our close social circle is bound to have some effect on our own epistemic attainment. If someone from our family posts about a conspiracy theory on social media and we see it in our feed, this should affect us in some way: either we diminish our propensity to trust that family member about political or social matters, or we may start looking at that theory deeper, giving it more credence since someone from our close circle endorses it, so maybe there is something there worthwhile after all. No matter how autonomous we may consider ourselves to be, knowing that most people around us hold a belief, it is hard to escape its influence on us due to the emotional contagion going on in online environments (del Vicario et al., 2016). The emotional contagion surrounding information propagation and the experienced exposure to other's epistemic attainment are only two mechanisms that differentiate SMPs from other socio-technical environments and that have a bearing on the epistemic attainment of the users, and often in a negative way. The issue at stake is that these two mechanisms also influence how epistemic environments interface on one platform or across multiple social media platforms.

As previously mentioned, SMPs are ecosystems, a multitude of environments which interact with each other in a variety of ways, depending on the user's perspective. SMPs gather in one place several epistemic environments and this collection generates new interactions by

itself. Thus, the discussion on SMP as hostile epistemic environments should be nuanced to account for the multiplicity of environments and how these interface. Many accounts of SMPs as hostile epistemic environments assume an external point of view, one notion of epistemic health and one set of epistemic norms. However, if we assume that multiple epistemic environments are in the online ecosystem, how do these interface?

Specific examples of epistemic interfaces on SMPs are human users and algorithms. Any member of two or more distinct epistemic environments will enact an epistemic interface whenever they post content that is reachable by the members of the other epistemic environment. Influencers may find them interfacing more often than they wish to, but also regular users who happen to be visible to more than one epistemic environment. In addition, any kind of algorithmic filtering mechanism - be it for entertainment, engagement or information - act *de facto* as epistemic interfaces. They expose users to information and worldviews that one would not normally see. Looking at two standard examples of online hostile epistemic environments, such as echo chambers (Nguyen, 2020) or epistemic flooding (Anderau, 2023), we should notice that what makes these phenomena hostile to the epistemic attainment of online users is also related to how the environment interacts and interfaces with another one in ways that are detrimental to the environment itself, to its inhabitants, or to outsiders.

In a 2020 paper, Thi Nguyen argued that echo chambers and epistemic bubbles are "structures of exclusion ... [that] reinforce ideological separation" (p. 141-142), with a significant difference in the mechanisms achieving this separation. Exclusion is the mark of an epistemic environment trying to be separate from other environments and we can read this as refusing to interface. Both epistemic bubbles and echo chambers are epistemic environments that effectively separate their inhabitants from other environments, with bubbles occurring inadvertently by excluding "some relevant voices" (Nguyen, 2020, p. 142), while echo chambers are actively trying to exclude or discredit other perspectives, usually in bad faith. Both echo chambers and epistemic bubbles have faulty epistemic interfacing: the echo chamber ideally has no interfacing, which means that the epistemic flourishing of its members is heavily reliant on the echo chamber community, and, when an echo chamber member escapes to the outside world and tries to bring with them the epistemic outcomes from the echo chamber (certain insights about, let's say, conspiracies or vaccines), these outcomes are not valuable in any way to the outside world. It could be said that the outside world fails to interface with the echo chamber since it does not take seriously its epistemic outputs, but, given the intended isolation of the echo chamber from the beginning, these outcomes are not seriously usable in the outside, since these were made to work only in the closed world of the echo chamber, with no intention for dialogue. In the case of echo chambers, their outcomes can survive the transfer to other environments, depending on the situation. Still, there is faulty interfacing in the epistemic bubbles because the very lack of interacting with other epistemic environments is what allows the epistemic bubble to survive intact for so long. Inferring from these two cases an ad-hoc norm for interfacing, no epistemic environment should become an isolated world, unconnected with other environments. There should always be some interfacing with other epistemic environments. And because SMPs allow for such easy isolation in groups with like-minded views, remaining unchallenged and self-sufficient in their lack of interaction with the outside world, SMPs enact epistemic environments with faulty interfacing.

Epistemic flooding has been described “when epistemic agents find themselves in epistemic environments in which they are routinely confronted with more information and evidence than they can diligently process” (Anderau, 2023, p. 3). Flooding is a particular phenomenon characteristic of SMPs, as it is triggered by design choices which favour the sharing and multiplication of information online, due to the algorithmic nature of social media platforms whose “feeds thrive off of engagement and the prolific sharing of content” (Anderau, 2023, p. 9). We can think of YouTube influencers who pick up a theme from other influencers and then make their own spin on it, repeating the same kind of information but with their own style. The flooding occurs sometimes by repetition: the same information repeated to users ad nauseam may end up giving a false idea of consensus. The second problem is that this information answers questions we did not ask, but it pre-emptively answers them and may be seen as a weak form of inquiry (Anderau, 2023, p. 8). To break from the flooding’s overwhelming effects, we would need access to a diverse range of information, different opinions and interpretations of said information, but, since we are under a constant overwhelm of information on social media, we are not inclined to look for more. The problematic aspect of epistemic flooding occurs both because of the sheer quantity of information to which users are exposed (which is cognitively hostile in itself), and because of the repetition of the same kind of information (which is epistemically misleading). Here we have an opposite feature of interfacing from epistemic bubbles: too much information is flooding the users. Granted, the traffic of information can happen also inside the same epistemic environment, and this is not a problem of interfacing, rather a problem of the epistemic environment itself that is simply too informationally rich. But when the information comes from an outside epistemic environment, this has detrimental effects for the environment it is flooding: the norms and epistemic outcomes of the environment being flooded cannot be safely preserved or used, since there is a mix of what counts as a valuable epistemic outcome, and ultimately a diminishment of the capacity of the inhabitants of discerning epistemic value from junk. If the internal norms for what counts as true and justified information are consistently over-ridden by epistemic outcomes from another environment (which does not challenge the norms as such, just consistently provides epistemic outputs based on different norms and values), then we can only conclude that there is something wrong with our norms, or, nihilistically, that anything goes, and all norms are entitled. This can make the epistemic environment lose its unique contribution to the world and nullify its own epistemic outcomes which are constantly under siege. In cases of faulty epistemic interfacing, we can say that flooding results in an effective diminishing of the epistemic attainment of the inhabitants of one environment. Of course, not all cases of flooding result in that effect.

One could argue that we do not need the new concept of an epistemic interface, since there are previous concepts that are similar enough in social epistemology, such as the concept of an epistemic conduit (Alfano, 2021; Alfano & Sullivan, 2023). An epistemic conduit is any actor (human or non-human) that is placed in a communication network and that “pass[es] on information from sources (or from other conduits)” (Sullivan & Alfano, 2023, p. 8). It seems that epistemic interfaces do simply what a conduit does: moving information around and making it available to others. However, there is a crucial difference here between the two concepts: epistemic interfaces move information around from at least two distinct epistemic environments, whereas the network topology described by Alfano and Sullivan is

neutral towards how many epistemic environments are crossed by a conduit. A conduit functions inside an epistemic environment and may sometimes cross two distinct epistemic environments. However, by its very nature, the interface is placed between two different environments, translating between them. The second distinction is that epistemic interfaces aim to foster or at least not interfere with the epistemic attainment of the agents in an environment, whereas conduits simply transfer information from one network point to another.

Another similar concept is that of context collapse, defined as “the flattening out of multiple distinct audiences in one’s social network, such that people from different contexts become part of a singular group of message recipients” (Vitak, 2012, p. 451). Context collapse is often found in SMPs such as X/Twitter (boyd, 2008), resulting in significant epistemic consequences (Frost-Arnold, 2021) such as facilitating misunderstandings. Could we not then say that epistemic interfacing is simply a mere context switch that is well marked? As mentioned above, I am interested in epistemic environments insofar as they allow for epistemic outcomes to circulate outside their native environment. Context collapse occurs when we misunderstand the communicative intent of the speaker because we have no idea in which context this was uttered, for example when someone says something jokingly and we, several tweets removed from it, perceive it as a serious political statement. The context we misunderstand can be also an altogether different epistemic environment with its norms and language games. If two environments interface healthily, the epistemic outcomes of a user get picked up in the other environment as epistemic outcomes, not as jokes or moral statements. Context collapse can happen in a variety of situations, inter- and intra-environments, so it is not specific to interfacing. However, there are cases when context collapse contributes to faulty epistemic interfacing, because the original context of the utterance is not clear and this hides the nature of the communicative intent, epistemic or otherwise. This is to say that not all cases of context collapse are faulty epistemic interfacing, but that context collapse can effectively contribute to the faulty interfacing, especially when the communicative act crosses the barrier of an epistemic environment. Faulty epistemic interfacing is also a systematic phenomenon pertaining to epistemic environments which simply cannot make use of each other’s epistemic outputs, because of the different norms and values they hold, whereas context collapse is limited to individual speech acts.

6. Conclusions. Why faulty epistemic interfacing matters

The aim of this paper was to introduce the concept of epistemic interfacing as the dynamic space of interaction between an epistemic environment and its outside world. Concerning SMPs, this work pleads for coming up with norms of healthy interfacing besides the work already done in delimiting what makes an epistemic environment healthy or not. While I agree with most scholars that analyse SMPs as epistemic environments, we should also be aware of SMPs as ecosystems, collections of environments that interface in a variety of ways, often unpredictably or hostile to their inhabitants. SMPs are ecosystems of epistemic environments with no explicit norms about what counts as a legitimate interaction or successful interfacing. This is in line with previous work highlighting that we need more explicit epistemic norms to account for the epistemic operations going on online, for

example norms for sharing content online (Record & Miller, 2022; Author, 2021), but this work also qualifies the difference between inter and intra-environmental epistemic norms. However, with norms for interfacing, we find ourselves in the strange space between two epistemic environments, each with their own norms and values, theoretically very difficult to bridge. Norms for epistemic interfacing need to be a kind of meta-norms, ensuring the epistemic attainment of both sides, without committing to a particular form of epistemic attainment, agency or value.

We need explicit norms that are not bound to a specific epistemic environment, and we need online institutions to enforce these without crushing this newly found epistemic agency for regular users. So far, enforcement has meant algorithmic filtering, blocking and reporting, and misinformation tagging. However, this kind of enforcement stays at the two basic levels of the informational system (information and application level) because we still do not have interface norms.

SMPs are the hosts of a variety of epistemic environments that co-exist interact with each other in the same digital space, also spilling in the offline world. Some of these epistemic environments are hostile because of how they are run (echo chambers) or the inadvertent isolation from different perspectives (epistemic bubbles enforced by algorithmic filters), or because they overwhelm their users with information that cannot be meaningfully engaged with. But, in addition to the internal hostility of these environments, there are often problems with how the environments interface with their outside realm: either not letting anything pass through from the outside (accidentally or intended), or by being too permeable to the outside world, by allowing for context collapse and anything to count as an epistemic outcome. It is still unclear how healthy epistemic interfacing should look like, but a minimal set of requirements could be encapsulated in these tentative norms: that interfacing should aim to preserve the epistemic agency of the inhabitants, the interfacing should preserve the epistemic outcomes from an environment to another, and there should always be some kind of interfacing between epistemic environments, such that there is no isolated environment. These norms are minimal and will need more backing in future work, probably with some empirical measures of kinds of interfacing and healthy interactions.

References

[Author, 2021]

[Author, 2024]

Alfano, M. (2016). The Topology of Communities of Trust. *Russian Sociological Review*, 15(4), 30–56. <https://doi.org/10.17323/1728-192X-2016-4-30-56>

Amico-Korby, D., Harrell, M., & Danks, D. (2024). Environmental epistemology. *Synthese*, 203(3), 81. <https://doi.org/10.1007/s11229-024-04504-3>

Anderau, G. (2023). Fake news and epistemic flooding. *Synthese*, 202(4), 106.

<https://doi.org/10.1007/s11229-023-04336-7>

Arfini, S., Bertolotti, T., & Magnani, L. (2019). Online communities as virtual cognitive niches. *Synthese*, 196(1), 377–397. <https://doi.org/10.1007/s11229-017-1482-0>

Battaly, H. (2021). Engaging closed-mindedly with your polluted media feed. In M. Hannon & J. de Ridder (Eds.), *The Routledge Handbook of Political Epistemology* (1st ed., pp. 312–324). Routledge. <https://doi.org/10.4324/9780429326769-38>

Battaly, H. (2018). Can Closed-mindedness be an Intellectual Virtue? *Royal Institute of Philosophy Supplement*, 84, 23–45. <https://doi.org/10.1017/S135824611800053X>

Blake-Turner, C. (2020). Fake news, relevant alternatives, and the degradation of our epistemic environment. *Inquiry*, 1–21. <https://doi.org/10.1080/0020174X.2020.1725623>

Boncompagni, A. (2022). *Wittgenstein on Forms of Life* (1st ed.). Cambridge University Press. <https://doi.org/10.1017/9781108946513>

Carter, M. C., Cingel, D. P., Ruiz, J. B., & Wartella, E. (2023). Social media use in the context of the personal social media ecosystem framework. *Journal of Communication*, 73(1), Article jqac038, 25–37. <https://doi.org/10.1093/joc/jqac038>

Cheong, M. (2023). *Social Media Harms as a Trilemma: Asymmetry, Algorithms, and Audacious Design Choices* (arXiv:2304.14679). arXiv. <http://arxiv.org/abs/2304.14679>

Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2016). Echo Chambers: Emotional Contagion and Group Polarization on Facebook. *Scientific Reports*, 6(1), 37825. <https://doi.org/10.1038/srep37825>

Figà Talamanca, G., & Arfini, S. (2022). Through the Newsfeed Glass: Rethinking Filter Bubbles and Echo Chambers. *Philosophy & Technology*, 35(1), 20. <https://doi.org/10.1007/s13347-021-00494-z>

Floridi, L. (2013). *The ethics of information* (First edition). Oxford University Press.

Floridi, L. (Ed.). (2015). *The Onlife Manifesto*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-04093-6>

Floridi, L. (2023). *The green and the blue: Naive ideas to improve politics in an information society*. John Wiley & Sons.

Frost-Arnold, K. (2021). Epistemic Dangers of Context Collapse Online. In Jennifer Lackey (Ed.), *Applied Epistemology* (pp. 437–456). Oxford University Press.

Frost-Arnold, K. (2023). *Who should we be online? A social epistemology for the internet*. Oxford University Press.

Freiman, O. (2024). AI-Testimony, Conversational AIs and Our Anthropocentric Theory of Testimony. *Social Epistemology*, 1-15.

Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018, April). The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI conference on human factors in computing systems* (pp. 1-14).

Harris, K. R. (2024). *Misinformation, Content Moderation, and Epistemology: Protecting Knowledge* (1st ed.). Routledge. <https://doi.org/10.4324/9781032636900>

Huemer, Michael. 2005. "Is Critical Thinking Epistemically Responsible?" *Metaphilosophy* 36:522–31.

Ienca, M. (2023). On Artificial Intelligence and Manipulation. *Topoi*. <https://doi.org/10.1007/s11245-023-09940-3>

Kerr, J. (2014). Western epistemic dominance and colonial structures: Considerations for thought and practice in programs of teacher education. In *Decolonization: Indigeneity, Education & Society*.

Kurtulmus, F. (2020). The Epistemic Basic Structure. *Journal of Applied Philosophy*, 37(5), 818–835. <https://doi.org/10.1111/japp.12451>

Levy, N. (2023). It's Our Epistemic Environment, Not Our Attitude Toward Truth, That Matters. *Critical Review*, 35(1–2), 94–111. <https://doi.org/10.1080/08913811.2022.2149108>

Manovich, L. (2013). Media After Software. *Journal of Visual Culture*, 12(1), 30–37. <https://doi.org/10.1177/1470412912470237>

McLuhan, M., & Powers, B. R. (1989). *The global village: Transformations in world life and media in the 21st century*. Communication and society.

Miller, B., & Record, I. (2013). Justified Belief in a Digital Age: On the Epistemic Implications of Secret Internet Technologies. *Episteme*, 10(2), 117–134. <https://doi.org/10.1017/epi.2013.11>

Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17(2), 141-161.

Nguyen, C. T. (2021). How Twitter Gamifies Communication. In J. Lackey (Ed.), *Applied Epistemology* (1st ed., pp. 410–436). Oxford University Press/Oxford. <https://doi.org/10.1093/oso/9780198833659.003.0017>

Nguyen, C. T. (2023). Hostile epistemology. *Social Philosophy Now*, 39, 9–32. <https://doi.org/10.5840/socphiltoday2023391>

Origg, G. (2018). *Reputation: What it is and why it matters* (S. Holmes & N. Arikha, Trans.; English edition). Princeton University Press.

Record, I., & Miller, B. (2022). People, posts, and platforms: reducing the spread of online toxicity by contextualizing content and setting norms. *Asian Journal of Philosophy*, 1(2), 41.

Reider, P. J. (Ed.). (2016). *Social Epistemology and Epistemic Agency*. Rowman & Littlefield International.

Rini, R. (2017). Fake news and partisan epistemology. *Kennedy Institute of Ethics Journal*, 27(2), E-43.

Ryan, S. (2018). Epistemic Environmentalism: *Journal of Philosophical Research*, 43, 97–112. <https://doi.org/10.5840/jpr201872121>

Ryan, S., Palermos, S. O., & Farina, M. (2023). Smart Environments. *Social Epistemology*, 1–20. <https://doi.org/10.1080/02691728.2023.2277230>

Srinivasan, P. (2023, July). Information Everywhere! Cognitive Load, Elaboration, and Knowledge Gain from Consuming Social Media News Posts. In *International Conference on Human-Computer Interaction* (pp. 119-125). Cham: Springer Nature Switzerland.

Tollon, F. (2022). Artifacts and affordances: From designed properties to possibilities for action. *AI & SOCIETY*, 37(1), 239–248. <https://doi.org/10.1007/s00146-021-01155-7>

Trompette, P., & Vinck, D. (2009). Revisiting the notion of boundary object. *Revue d'anthropologie des connaissances*, 3(3-1).

Vitak, J. (2012). The Impact of Context Collapse and Privacy on Social Network Site Disclosures. *Journal of Broadcasting & Electronic Media*, 56(4), 451–470. <https://doi.org/10.1080/08838151.2012.732140>

Werner, K. (2022). Institutions as Cognitive Niches: A Dynamic of Knowledge and Ignorance. In S. Arfini & L. Magnani (Eds.), *Embodied, Extended, Ignorant Minds* (Vol. 463, pp. 161–189). Springer International Publishing. https://doi.org/10.1007/978-3-031-01922-7_8

Zhao, X., Lampe, C., & Ellison, N. B. (2016). The social media ecology: User perceptions, strategies and challenges. In J. Kaye, A. Druin, C. Lampe, D. Morris, & J. P. Hourcade (Eds.), *Proceedings of the 2016 CHI conference on human factors in computing systems* (pp. 89–100). ACM Press. <https://doi.org/10.1145/2858036.2858333>