A Knowledge Producer's View on the Knowledge Commons

M. Noual[1]

[1]Université Paris-Saclay, CEA, List, Palaiseau, France

2023

Hardin [1] introduced the notorious concept of "*tragedy of the commons*" (TOTC). Worrying about the consequences of human overpopulation on the planet, he discussed what I will refer here to as "*hard problems*": problems with no technical solutions, that can only be addressed by way of an evolving morality. Hardin's TOTC predicts that *the hard problem of human population growth directly implies a hard problem of overuse or pollution of the commons*. Here, I focus on the *knowledge* commons. I argue that it is not clear that the TOTC applies to the knowledge commons, for reasons similar to those that have protected agrarian commons for centuries. Even if the knowledge commons satisfies the TOTC's necessary conditions, it is not clear that the ensuing problems are *hard*.

In section 1, I present the point of view that I propose to take on the knowledge commons. It emphasises the notion of *human attention*. I also discuss conditions under which the TOTC applies to the knowledge commons. Section 2 introduces the main principles of a technical solution called the "MMM" [2] for structuring a digital knowledge record. Relying on these technically implementable principles, section 3 proposes solutions to push back and narrow down the hardness of familiar problems related with the knowledge commons [3, 4].

As will be justified below, in this paper, I don't distinguish between knowledge and information.

1 Tragedy or not?

The TOTC is usually considered not to apply to the knowledge commons. In this section I argue that there is a perspective on the knowledge commons from which the TOTC does apply. To predict the ineluctable overuse of a commons, Hardin made a number of assumptions. Primarily, he assumed consumption of the commons was *rivalrous* and *non-excludable*.

Rivalry is due to the consumed resource being *substractive*. This is usually not considered to be the case with knowledge [5, 6]: if I use a piece of knowledge, you can also still use that same piece of knowledge. Knowledge consumption isn't rivalrous and the conditions for Hardin's TOTC aren't satisfied. There is no risk of overusing knowledge. Arguably it may even be the case that the more there are other individuals consuming knowledge the more educated society is and the more probable it is for each individual to consume knowledge. This is a **knowledge consumer view** on the knowledge commons. It focuses on the consumption of a non-material non-substractive resource, namely knowledge, assumed to have already been produced. With this view, production of the resource (e.g. by Fox news or by academic research) and also maintenance (e.g. fact-checking) are dissociated from consumption. They can be centralised. The skills and means for emission and maintenance can be entrusted to different entities[1] than those who consume the information. The reality of the so-called information age challenges this view. Informational resources are mass produced. Centralised maintenance is not realistic. Institutional fact-checking is conspicuously failing at making the Web safe while Wikipedia continues to demonstrate the comparative success of distributed informational curation. An alternative **knowledge producer view** is necessary to account for the production of the knowledge good, upstream from knowledge consumption.

I assume that all **pieces of knowledge** (unlike thoughts) are produced with an intention for them to be consumed. With this assumption and the knowledge producer view, production and maintenance are indissociable from consumption.

[1]e.g. Google – Google's guidelines [7] define at length what good information is under the hypothesis that good information is measured in terms of a fit between a document and a user's expectancies for this document. The webpage of a fascist conspiracist comic can be deemed of excellent quality if it is labelled as containing fascist, comic content. Serving this page to users asking for fascist comic content will bring about customer satisfaction, which is the purpose of "eyeball selling" digital businesses like Google [8, 9]. However, arguably, good informing is a different activity from bringing about customer satisfaction. The Google founders themselves originally noted the incompatibility between good informing and eyeball selling [10].

Knowledge itself isn't subtractive, but the means of sharing it can be. If you borrow the library's only book on Category Theory, then I can't borrow that book simultaneously. Digital technology mitigates the substractiveness of classical means of sharing knowledge [11]. The replenishabilty of network bandwidth facilitates the replenishing of the potential for knowledge consumption. The replicability of digital files facilitates the repetition of knowledge consumption experiences by multiple knowledge consumers. It would be a mistake however to confuse knowledge with knowledge media. Just like the book isn't the knowledge it conveys, radio frequency, network bandwidth, digital files are not the knowledge they convey. The library book can serve as a door wedge. Its digital version can be used as part of a digital art piece. I propose to relate knowledge consumption with *understanding and assimilating knowledge* rather than with *accessing the (digital) resources that convey the knowledge*. The relevant bandwidth to consider now is that of the human mind: **attention**, or readiness to knowledge consumption [12]. Attention is also replenishable, and arguably, non capitalisable (present mental bandwidth can't be saved in order to have more bandwidth in the future). Considering attention reintroduces the question of substractiveness [12, 13]. If all my present attention is consumed in the intellectual effort to understand Category Theory, my mind is presently unavailable to learn about AI transformers. In 2004, P. Le Lay, then president of a French TV channel, notoriously said that what the TV channel was selling to Coca-Cola was "available brain time" [14]. He specified that for an advertisement to be perceived, the viewer's brain must be available, and that the purpose of the broadcasted TV content is to prepare the brain for the reception of advertisement messages. Devenport and Beck introduce the notion of "attention economy" in which attention is a commodity, sometimes even a currency [12]. Meanwhile, the digital marketing industry supports the "eyeball economy" where the good being sold by media businesses is not digital information (that is provided for free to consumers) but consumers' "eyballs", i.e., their attention [8, 9]. H. Simons said: "*In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.*" [3]

While there is no rivalry in knowledge consumption, there is rivalry to be expected in the consumption of the collective attention resource [15, 16, 17]: (1) epistemic **rivalry between pieces of knowledge** and (2) (economic) **rivalry between knowledge producers**. First, prior consumption of one specific piece or type of information may affect the future consumption of another [18, 16]. An individual's assimilation of the theological argument may affect their capacity to understand processes involved in biological evolution, and vice versa. The prevalence of an understanding of the theological argument in a community may affect that community's capacity to process biological knowledge, and vice versa. Of course not all pieces of knowledge are in competition with one another. Some synergise. One first piece of knowledge may make a second piece of knowledge easier to consume. A step further would be to consider a notion of "**epistemic pollution**" of information spaces. Just like air pollution degrades individual humans' experience of the air, epistemic pollution degrades their experience of information. It consumes consumers' attention and their ability to be well-informed. A notion of epistemic pollution would require a notion of *epistemic purity* which could be defined in different ways reflecting different biases on what pure/polluted knowledge is considered to be. From lack of a need to compare the qualities of different pieces of knowledge in terms of a single informational quality, the present article doesn't need to formalise a notion of epistemic pollution. For the same reason it needn't distinguish between *information* and *knowledge*. I use the terms interchangeably here[2]. Rivaly between pieces of knowledge competing for attention, makes knowledge production rivalrous for knowledge producers. If I produce a piece of knowledge that catches your attention, I may have affected your ability to pay attention to someone else's work. Knowledge producers, like/including the media industry's advertisers, compete for knowledge consumers' attention. If the milk industry has already convinced the population that milk is scientifically proven to be good for your health, the scientific community may struggle to get the nuancing pieces of knowledge it produces across.

To meet the conditions of the TOTC, knowledge production, or more precisely attention consumption, needs to be "**non-excludable**" in addition to rivalrous. *Physical* information media (e.g. printed newspapers) preserve a form of informational freedom or "intellectual privacy" [19]. They tend to be slower and fewer at capturing our attention than digital media (cf smartphone notifications). The ubiquity of information and communication technology makes it difficult to exclude knowledge producers from tapping into the collective attention resource [20]. The non-excludability condition is

[2] Mocanu et al. [16] provide an additional reason not to make a formal distinction between different qualities of information: "*attention patterns when faced with various contents are similar despite the different qualitative nature of the information*".

satisfied and the knowledge commons is subjected to the TOTC after all.

2 The Technical MMM Proposal

The conditions of the TOTC being satisfied, Hardin predicts a problem of scale with no technical solution. Modern information and communication technology opens consumption of the substractive collective attention resource to all. The size of the human population is bound to bring the knowledge commons to its ruin as the numerous knowledge producers acting in their own interest are bound to deplete the resource. Humans will then be deprived of the attention necessary to process information [12] and the hard problem of overpopulation will become a hard problem of overly stupid population. To check Hardin's prediction, it remains to check the *hardness* of problems involved. In this section I present a technical solution that invalidates Hardin's prediction by pushing back and narrowing down the need for morality to solve the issues of the knowledge commons (especially the need for a fresh new morality that wouldn't already exist).

This section lists the main ingredients of the MMM proposal which is detailed in [2]. The idea of the proposal is to support the knowledge commons by organically materialising a "backbone" for it using the "MMM format" (cf §2.2).

2.1 Pieces of Knowledge

In the MMM proposal [2], the atomic informational unit considered is a piece of knowledge. Traditional documents (e.g. articles) usually contain multiple pieces of knowledge and sometimes overlap (one piece of knowledge appears in multiple documents). Like the Semantic Web proposal, the MMM proposal departs from mainstream information technology approaches: it supports a finer granularity of information than what is supported by traditional documents. The MMM format's (see next paragraph) definition of pieces of knowledge however, is significantly looser than the SW's. Individual questions, in particular, are typical MMM pieces of knowledge – cf Fig. 1.

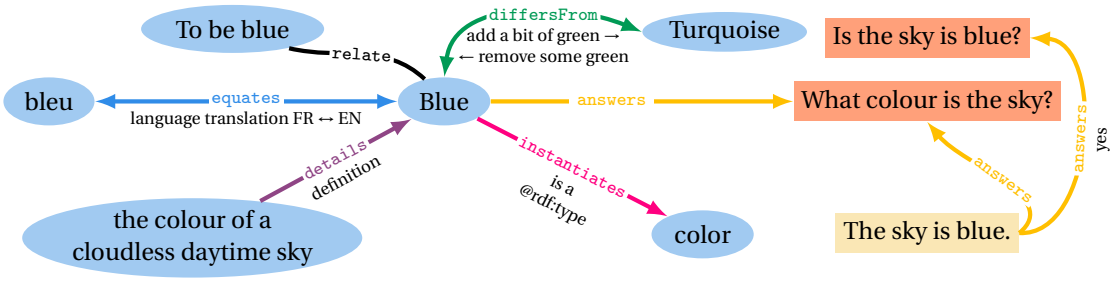

Figure 1: The MMM format defines different **types** of pieces of knowledge. Some are represented in the figure above as **nodes** of a graph, others as **edges** connecting two nodes. Any piece of knowledge documented in MMM format must be assigned a type. In this figure, yellow rectangle nodes represent pieces of knowledge of type "`narrative`" which is the **default type**. Orange rectangles and blue ovals respectively represent pieces of knowledge of type "`question`" and"`existence`". Yellow arrows represent pieces of knowledge of type "`answers`". *etc.* The visual choices made in this figure are arbitrary. MMM formatted information doesn't even need to be graphically represented [21]. MMM edge types (e.g. `equates`, `differsFrom`, `instantiates`, `details`, `nuances`, `questions`), not all of which are illustrated above, have loose semantics which can be specified using an edge **label** – e.g. the epistemic relation conveyed by the `details` edge is specified by the label "definition".

2.2 A Common Documentation Format

A comparison with the academic journal article format (AJAF) can help understand the MMM format. The AJAF is a common documentation format for academic researchers' work. It imposes a loose structure: documented knowledge is organised into standard types of "information containers" such as: *title, abstract, introduction, numbered section 2.3, etc.* The *introduction* of an AJAF document for instance can contain multiple pieces of knowledge. The MMM format is also

a generic documentation format that imposes a loose structure. Unlike information containers of the AJAF, those of the MMM format are meant to contain a single piece of knowledge. The MMM format defines a small set of generic **types** of information containers which are presented in Figure 1. Documenting notes in MMM format means decomposing them into a network of typed pieces of knowledge. An especially important way of documenting MMM content is to **annotate** pre-existing MMM content. Annotating a piece of knowledge k (e.g. nuancing, questioning, detailing it) means documenting other pieces of knowledge and linking them to k using the appropriate, predefined types of MMM edges (e.g. `nuances`, `questions`, `details`).

MMM formatted content can be saved in a human friendly MMM-JSON format [21]. Existing software (e.g. documentation, publication, communication tools) may be enhanced with MMM formatting features, and various new tools may be developed to support and to customise MMM editing and navigating experiences. An individual's access to the MMM formatted knowledge that they have produced and gathered is however not dependent on any such tool.

2.3 Local Epistemic Territories

A set of MMM pieces of knowledge is called an **epistemic territory** or **land** [2]. An individual's **local epistemic territory** is the set of pieces of knowledge that she is acquainted with – the area covering the extent of what she knows, cf Figure 2. The local epistemic territory of a individual is stored on local machines that she owns and/or on remote servers that she controls. An individual has full agency of her territory. She selects the pieces of knowledge that are added to it, and those that are deleted from it. Individuals can form **communities** of interest together. Communities can equip themselves with servers to store the parts of their members' local epistemic territories that are of common interest. Community servers materialise the common epistemic grounds that are the community's raison d'être[3]. Local machines and servers on which epistemic territories are stored, can be owned by private or public entities. Epistemic land owners and/or tenants (community members) define the regulations that apply to their territories (cf §2.11): who is entitled to access the pieces of knowledge on their land, and under what conditions is a new piece of knowledge added to the land by a knowledge producer.

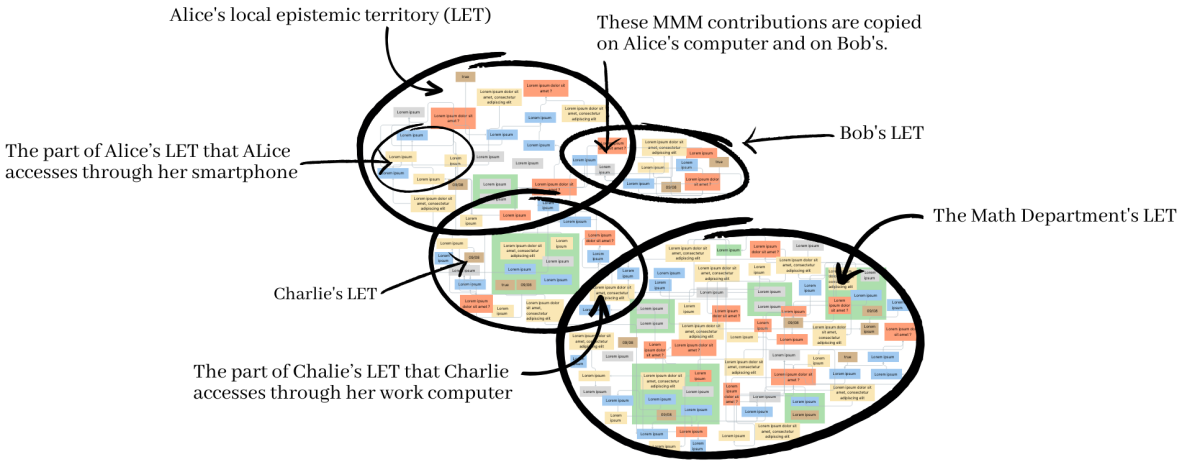

Figure 2: The MMM is a collection of overlapping local epistemic territories.

Formally, **the MMM (Mutual Mutable Medium)** denotes the reunion of all local epistemic territories, i.e. the set of all pieces of knowledge documented in MMM format. As discussed below in §2.4, the MMM is primarily a theoretical concept[4]. In practice, only *local* territories of landowners (individuals or well-defined communities) are materialised.

[3] They also can serve as backup storage for community members.

[4] Although the MMM proposal envisions MMM crawlers that leverage epistemic bridges across local MMM territories to gain hindsight on the global epistemic landscape.

2.4 Organic Distribution

The classical knowledge consumer view naturally sees the knowledge commons as a global collection $\mathscr{C}$ of pieces of knowledge[5]. A noble aim is to democratise access to $\mathscr{C}$, make it wholly available to all knowledge consumers, typically by *decentralising* $\mathscr{C}$. $\mathscr{C}$ is divided into chunks of data which are distributed to multiple storage locations. Some redundancy is strategically implemented to ensure that even when one of the storage locations is down or disconnected, knowledge consumers all still equally have access to the entirety of $\mathscr{C}$ at any point in time. With the knowledge producer view, the relevance of the concept of a single global collection of pieces of knowledge is secondary. Primarily what matters is the local production of knowledge. The knowledge producer view acknowledges that knowledge is not centralised *in the first plac*e. Its processing is contextual, its production is localised, and its documentation is collective and thereby organically distributed *from the start*. Even public knowledge (e.g. results from academic research) starts as private notes, produced by knowledge producers whose work requires some intellectual privacy and focus. Arguably, no single knowledge consumer needs access to the entire knowledge record at any given point in time. So it is not clear that centralisation of knowledge is necessary in the first place. And without some initial centralisation, decentralisation is unnecessary (and impossible). More important than democratisation of *access* to knowledge is democratisation of means of getting messages across, especially getting them across to the right (ready) knowledge consumers equipped with the required amount of attention to process the information.

2.5 Sharing and Publishing

In the MMM system [2], a knowledge producer can share a piece of knowledge k from her local epistemic territory with other knowledge consumers of her choice who, if they accept k, include a copy of k in their own local epistemic territory. Before sharing k, the knowledge producer can mark k as "`public`". The `public` mark means "gifted to the public domain" and is intended to be irrevocable. The author of a `public` piece of knowledge k is as much/little the owner of k as anyone else who has a copy of it in their local territory. The only impact the author or anyone else can have on k once k has been published, is to annotate k (nuance it, detail it, link it to a reformulation of k *etc*) in order to sway knowledge consumers' appreciation of k without directly modifying k. People who have a copy of k in their local epistemic territory can abstain from sharing k and delete it. The `public` mark is irrevocable but the persistence of k isn't. It reflects the intention of a knowledge producer but the reality may be different. Even if k hasn't been marked as `public`, the consumption of k may become a public matter if no-one is excluded from it. And conversely, if a `public` piece of knowledge k is filtered out of all public epistemic territories (cf §2.12), it can only be found on private land and depends on private landowners to persist.

2.6 Epistemic Relations, Aggregation and Glue

In the MMM proposal, links between pieces of knowledge, namely MMM edges, are themselves pieces of knowledge (e.g. the relation between Newtonian mechanics and Lagrangian mechanics). This means that they can be annotated (nuanced questioned, detailed, *etc*) like any other piece of knowledge. MMM edges convey epistemic relations and play a central role in the MMM solution. Pieces of knowledge linked by an edge or by a short path (sequence of consecutive edges) are said to be **epistemically close** – e.g. all pieces of knowledge illustrated in Fig. 1. The MMM proposal like the Semantic Web proposal is designed to facilitate the collective documentation of links between atomic informational resources. Pieces of knowledge in the MMM have unique identifiers so any knowledge producer who wants to annotate a piece of knowledge k that someone else has documented can do it as long as they (or the MMM software editor they are using) know(s) the identifier of k. A major difference between the MMM and the SW is that the SW supports *semantic* interlinking only, while the MMM more loosely supports any form of *epistemic* linking [2]. Different answers to the same `question` Q provided by different knowledge producers, don't necessarily agree (aren't necessarily *semantically* close) but they are *epistemically* close in that they serve a similar epistemic role which is to answer Q. They can all be linked to Q using a MMM `answers` link – in which case we say that Q and its answers are **agreggated** in the same area of the MMM. An **area** of the MMM is a connected set of pieces of knowledge. As another example, a claim is epistemically close to the definition of a term it implies. The two can be linked together using a `details` edge labelled "is involved in". The MMM format is designed to support the **aggregation** of epistemically close pieces of knowledge, i.e. the documenting

[5] In the MMM solution, $\mathscr{C}$ would be the global MMM land, the reunion of all local epistemic territories.

of edges (or short paths) between epistemically related pieces of knowledge, to the effect that these pieces of knowledge find themselves in the same area of the MMM. Simons argues that an information processing system (IPS) will reduce the information overload problem rather than compound it only if it absorbs more information than it produces [3]: "*To be an attention conserver [...], [an IPS] must be an information condenser. It is conventional to begin designing an IPS by considering the information it will supply. In an information-rich world, however, this is doing things backwards.*" Thus the MMM proposal's primary focus in on condensing the knowledge record (the MMM) by facilitating the shortening of epistemic pathways and aggregation (see also §2.7 about redundancy management). A single MMM edge-typed piece of knowledge between two pieces of knowledge k_1 and k_2 is sometimes enough to express the epistemic relation between k_1 to k_2[6]. Sometimes multiple pieces of knowledge are needed. The notion of **epistemic glue** generalises the notion of direct epistemic relation conveyed by MMM edges. Epistemic glue between k_1 and k_2 denotes a set of pieces of knowledge that together express an epistemic relation between k_1 and k_2. When k_1 and k_2 belong to different local epistemic territories T_1 and T_2, the epistemic glue between them is said to epistemically **bridge** territories T_1 and T_2.

2.7 Redundancy Management

Epistemic glue can serve epistemic democracy through the exploitation of a ***good*** form of **redundancy**. Suppose k_1 and k_2 are two very similar pieces of knowledge expressing the same idea $\mathscr{I}$ in different words appealing to different people. Documented epistemic glue increases the chances that people who understand $\mathscr{I}$ through the k_1 wording become aware of the knowledge documented as following from k_2. Epistemic glue can also help identify cases of ***bad*** **redundancy** where k_1 and k_2 are so similar that **merging** them will result in no knowledge consumer loosing potential to consume knowledge, on the contrary[7]. When a piece of knowledge k is published, it is usually meant to be seen. If k is redundant with pieces of knowledge already documented, not only may its production have been a waste of the author's attention, it may also in the future waste the attention of knowledge consumers who will examine it only to realise it isn't new. Provision of epistemic glue, because it enables aggregation and thereby favours knowledge producers' awareness of the state of the art in their field, may help preempt this situation. The more glue is provided, the easier it may be to notice the epistemic proximity between documented pieces of knowledge, and the easier it may be to deal appropriately with redundancy – get rid of it and avoid adding more.

2.8 Wayfarer Exploration

The classic "**parachutist approach**" to discovering knowledge entails a knowledge consumer making a request (e.g. typing a term such as "transformers" in a search engine, or querying a librarian) and having a centralised entity or process (e.g. a search engine, a librarian) decide what knowledge to make the knowledge consumer aware of as a result of her query. With this approach, it is usually enough for a knowledge consumer to know how to utter, spell or type the term "transformers" in order to be served, among other things, (technical) knowledge about modern deep learning transformer technology [23]. The knowledge consumer's readiness to consume and understand the knowledge is not meaningfully accounted for in the knowledge retrieval and distribution service. The MMM proposal unlocks the possibility for a complementary, alternative approach. With the MMM based "**wayfarer approach**", the knowledge consumer "walks" step by step, piece of knowledge by piece of knowledge, starting from an area of their own local epistemic territory – *i.e.* starting from pieces of knowledge that they are already acquainted with – and expanding their local epistemic territory as they discover new pieces of knowledge linked to what they know. The knowledge consumer who knows nothing about deep learning, doesn't access a technical piece of knowledge about transformers until they have found an epistemic pathway between what they already know and the concept of transformers. The path need not make the knowledge consumer understand transformer technology. It may be a very short path. Perhaps all the knowledge consumer needs to find out before accessing the material is that there is indeed a reason for them to care about transformers. A hybrid approach to search can rely on/force wayfarer exploration in order to serve search results in response to a search query. An underlying assumption of the

[6] This is to be contrasted with hyperlinks on the WWW and citations in academic papers. Those two types of links are epistemically shallow: in themselves they convey no information on how the two pieces of knowledge (in two different documents) are related. Notably, hyperlinks have been incrimminated in shallow media multi-tasking [20, 22]

[7] Knowledge consumers initially only acquainted with k_1 and not k_2 gain direct access to pieces of knowledge documented as following from k_2, as a result of merging k_1 with k_2.

wayfarer approach is that on a daily basis, a knowledge consumer is more likely to want to access ressources that have some epistemic relation with the contents of their local epistemic territory. There must be another reason for Bob wanting to access a contribution about transformers, than his ability to type "transformers" in a search bar. This other reason may already be reflected as a path in the MMM between what Bob knows and what Bob is interested in knowing. Otherwise, the wayfarer based search puts Bob in a knowledge producer position as he works his way over to the target concept of transformers, creating new epistemic links that he understands and materialising a new path in the MMM. Arguably, the purpose of documenting knowledge in general, and more precisely the purpose of documenting a line of reasoning $\mathscr{L}$, is **attention reuse**: to let future knowledge consumers benefit from pioneer thinkers' efforts in paving a way that makes $\mathscr{L}$ less attention consuming for them (assuming their epistemic starting point is similar to the pioneers'). If there is an *epistemic* reason for a knowledge consumer's interest in transformers, documenting this reason can eventually benefit other knowledge consumer's by increasing the relevance and reach of their own wayfarer explorations. Generally, any technical solution or incentive that promotes the connectedness[8] of the MMM graph potentially improves knowledge consumers wayfarer's explorations of the MMM.

2.9 Implantation and Visibility

A piece of knowledge k is said to be **well implanted** in the MMM if (1) MMM edges are documented between k and other pieces of knowledge and (2) these edges convey epistemically rich relations. Arguably, if k is recorded without any reference to pre-existing knowledge, it is of little value to humanity, until someone is able to link k. On the MMM, good implantation of k translates into high **visibility** of k. The more (and the *better*) links there are between k and other pieces of knowledge, the higher the chances another knowledge consumer will come across k following the links that lead to k (cf §2.8). On the contrary, lack of linking, and lack of *good* linking may be severely detrimental to a contribution's visibility since no (good) pathways lead to it and the chances of other knowledge consumers finding it are low. Epistemically shallow or ill-positioned contributions (e.g. an unlabelled `relate` edge, or a `narrative` wrongly linked by an `answers` edge to a `question` it doesn't answer) are likely to get red-flagged and/or filtered out by knowledge consumers (cf §2.11, §2.13). Provision of epistemic glue, be it through implantation of new pieces of knowledge, the annotating of old ones, or the deliberate bridging of local epistemic territories, is the main way of contributing knowledge to the MMM. Floating knowledge islands without ties to the global MMM are practically invisible because they are unreachable to the wayfarer.

We can now emphasise the role of epistemic glue. Glue gives information on information. The more of it, the easier we may expect administrative decisions concerning the global record of knowledge to be, the more stringent and systematic the management of redundancy can be, i.e. the more intelligible the record becomes. Epistemic glue explicitly documented in MMM format can be *generative* of a knowledge commons: glue can substantiate both (i) the content of the commons and (ii) its structure, spanning across local epistemic territories thereby allowing us to speak of a common overarching domain. Arguably, epistemic glue is key to the safe scaling up of the knowledge commons: as more knowledge producers contribute to it, more attention reuse enabling glue is documented.

2.10 Epistemic Topography and Shortsightedness

The MMM format gives a graph like structure to local epistemic territories. **Epistemic measures** can be defined in terms of graph theoretic properties. For instance, the depth (resp. utility) of a piece of knowledge k can be measured by the length of paths incoming (resp. outgoing) k. Less naive notions of depth and utility can take into account the types of MMM contributions involved (especially the types of edges). Epistemic measures are a technical solution to the (hard) problem of evaluating information quality in more diverse and relevant terms than in traditional binary true/false knowledge qualifiers. And they may be used to visually represent epistemic territories as 3D landscapes to knowledge consumers.

Arguably, knowledge consumers are entitled to an amount of ignorance, disinterest, and misconception. Scientific advancement would be impossible if scientists weren't allowed to grope their way to new knowledge. Good information does not exist without awareness of what one ignores and of the limits of the knowledge one has. Formal MMM based epistemic measures can support *safe* ignorance. Knowledge consumers, unequipped to process the technical details of a

[8] Enhanced connectedness is a positive externality of the local provision of epistemic glue by individual knowledge producers.

piece of knowledge k, can still be aware of the existence, epistemic purpose and depth of k by way of appropriate epistemic measures. Ad hominem arguments no longer are knowledge consumers' main way out of a requirement for an expertise they don't have.

2.11 Non-findability and Gatekeeping

With the knowledge consumer view, it may be tempting to simplify epistemic democracy to universal access to all public knowledge (cf §2.4). With the knowledge producer view, a finer notion of epistemic democracy is possible, accounting for the futility of distributing knowledge to individuals and communities that are not prepared for it, and for the importance of building bridges (cf §2.6) across worldviews (local epistemic territories) in order to gently enhance worldview fluidity and democratise epistemic readiness without aggressively "talking to walls" in an attempt to brutally negate and replace inconvenient worldviews. The technical MMM solution supports a form of subjective "**non-findability**" or "**non-spontaneous findability**" that departs from traditional enlightenment ideals: Bob's ability to type "transformers" *per se* no longer opens the way for him to any public material associated to the keyword "transformers"[9]. The actual state of Bob's local epistemic territory, which no-one needs to know about but Bob, decisively determines what pieces of knowledge are non-findable to Bob. And Bob's epistemic wayfarer (i.e. learning) efforts determine what pieces of knowledge become non-spontaneously (effortfully) findable to him (i.e. how his local epistemic territory expands). Note that if a **central entity** were to define what pieces of knowledge are non-findable to Bob, that entity would have to know an uncomfortably huge amount of information on Bob and Bob's actual situation (probably no less than what Bob has) in order to determine what knowledge Bob is able to process with the mindset he has today. The organic distribution of the MMM (cf §2.4) avoids the hard problem of defining non-findability and having a central entity implement it.

A knowledge consumer can grow her local epistemic territory by adding copies of pieces of knowledge discovered while exploring the MMM (cf §2.8) or by accepting pieces of knowledge shared with her by known knowledge producers (cf §2.5). The pieces of knowledge added to a knowledge consumer's local epistemic territory are the ones responsible for consuming her attention. The wayfarer is a technical solution to help a knowledge consumer protect her attention and prevent attention-wasting pieces of knowledge from finding their way to her. **Epistemic gatekeeping** is another which consists in systematically rejecting pieces of knowledge based on how they measure in terms of well defined epistemic properties (cf §2.10). MMM based epistemic measures open the possibility for knowledge consumers to fine-tune automatic filter mechanisms. They can customize the conditions under which a piece of knowledge shared with them is automatically accepted as part of their local epistemic territory or rejected. For instance a knowledge consumer might want to systematically ignore `narrative`s that no-one has yet deeply (cf §2.10) supported, questioned or nuanced. Gatekeeping operates at the finest level of epistemic granularity (that of pieces of knowledge rather than that of documents and document authors). It increases knowledge consumers' capacity to exclude pieces of knowledge and exclude authors from wasting their attention. Belonging to a topic T of interest to Bob, or mentioning keywords statistically related to T is no longer enough for a piece of knowledge k to reach Bob (to be parachuted onto him). Bob's local epistemic territory must demonstrate a *measurable* form of readiness towards k.

In the current document based world, facilitating the circumscription and gatekeeping of local epistemic territories would risk favouring **echo chambers**. Why can things be different in the MMM based world? Because of an interplay between gatekeeping filters (§2.11) and epistemic bridging/glueing (§2.6). Filters prevent knowledge consumers from seeing pieces of knowledge they don't want to see. Epistemic glue (aggregation) makes any knowledge consumer more likely to see all pieces of knowledge that are epistemically close to every piece of knowledge she takes interest in. This includes the `question` Q that is answered by the belief B she has, as well as all documented alternative answers A to Q which contradict B and possibly nuance her worldview. Filters are based on *epistemic* properties of pieces of knowledge (cf §2.10) rather than *semantic* properties. This means that pieces of knowledge can't be filtered out solely based on what they mean. They can only be filtered out based on their position in the MMM, i.e. their relation to other pieces of knowledge. Alice can easily filter out everything documented as an answer to `question` Q. She can also invest some of her own attention in the formalisation of a measure of incompatibility of an answer with her belief B – e.g. exploiting possible

[9] Bob can still be made aware *from a distance* of the existence of technical material, say for LLM specialists, through epistemic measures (cf §2.10). Non-findability of a piece of knowledge k concerns the *consumption* of k by Bob, not his *awareness* of k. This is to help avoid unprepared knowledge consumers misreading and misusing k.

`differsFrom` links. It is however not obvious how she may define general epistemic measures that reliably recognise all pieces of knowledge in disagreement with her worldview. Aggregation makes the walls of echo chambers porous. Anyone at the margins of a closed epistemic community that is caught in an echo chamber, can play a liberating role if they import even slightly more nuanced pieces of knowledge into the community's local epistemic territory.

2.12 Biased Common Grounds

Discussing a particular piece of knowledge k on one local epistemic territory, say Alice's, is generally not equivalent to discussing k on any other local epistemic territory. Let T_E be the local epistemic territory of an entity E, either an individual like Alice or a community (e.g. conspiracy theorists or the academic community). Suppose that T_E includes the `question` "*How do RNA vaccines work?*" and an answer A_1 mentioning microchips and 5G technology. An alternative answer A_2 exists, mentioning surface proteins. A_2 is documented on a different entity's local epistemic territory, but not on T_E. E is responsible for what pieces of knowledge are accepted as part of T_E. E has epistemic biases and interests. Because of that, T_E does not provide neutral grounds to have a discussion on the risks of RNA vaccines. Indeed, answer A_2 not being represented on T_E, it can't be annotated and discussed on T_E. Local epistemic territories are not equal in terms of the possibilities of discussion they support. Understanding a piece of knowledge k on a local epistemic territory T_E (by looking at the contextual knowledge surrounding k in T_E) is also usually *not* equivalent to understanding k on any other local epistemic territory. Some pieces of knowledge critically relevant to k might be absent from T_E. Some annotations and questions might be impossible on T_E. If a `public` piece of knowledge k is of value (e.g. a scientific publication that some scientists are willing to keep a local copy of), then k should be discussable on *public* common grounds. To discuss k, one should not have to be invited on someone's private local epistemic territory where the owner's conversational rules must be respected and his chosen gatekeeping applies[10]. Public institutions have a key role to play in the MMM ecosystem: that of providing the most *neutral* discussion grounds possible – mitigating risk of systematic epistemic attention deficit [13] – and ensuring these grounds are continually updated as the state of documented knowledge evolves. It remains for public institutions to decide what pieces of knowledge are of value (worth discussing) and what are not. This is similar to what any entity managing a local communal epistemic territory needs to do, except that the purpose of the public entity's bylaws is not to support a common interest but resolutely to enable open public discussion of public knowledge. Local public territories must not be restricted to a set of trustworthy pieces of knowledge such as peer-reviewed scientific publications. To avoid the fact-checking pitfall, information traditionally deemed of low quality (such as unsubstantiated or cryptic conspiracist contributions) should be included, not censored, cf §2.13.

2.13 Continual Improvement

It is impossible to produce new information without producing low quality information (cf the daily practice of scientific research)[11]. Many errors can decisively participate in the process of improving information. The contribution of a low quality piece of knowledge k is a step in the process of improving the record: it calls for further contributions specifying what is wrong with k and how to deal with it[12]. k should not be removed from the commons. Nor should its author be scorned[13]. A contribution of poor quality, just like any other contribution should be systematically exposed to the public eye, subjected to discussion and improvements, durably enough that concerned citizens learn from it[14]. Visibility of an error and how it has been addressed should be entertained as long as citizens risk repeating the error. On the MMM, annotations brought to k (`nuances`, `details`, `questions` *etc*) reflect the collective understanding and processing of k.

[10] This departs from project Solid's notion of pods [24].

[11] This doesn't mean that (valuable) new information is necessarily produced when low quality information is produced! Contributions can be so shallow that the best way to deal with them is to shut them down (address them) early to discourage any repetition of them later.

[12] If the k is very shallow, there might not be anything else to do than to address it once and for all and ensure all repetitions of it funnel to the same area and aggregate, so as to take a minimum amount of space in the knowledge commons/in the collective attention resource.

[13] The MMM solution proposes to penalise knowledge producers who repeat pieces of knowledge not knowledge producers who produce low quality pieces of knowledge. It aims at protecting the attention of knowledge consumers without reducing the freedom of expression of knowledge producers.

[14] Arguably, it is not so much misinformation that is a problem on actual information spaces, than it is its diffusion [16]. Hindering the diffusion of misinformation requires telling it from information. An alternative is to change the way it diffuses. The MMM solution, through aggregation, proposes to ensure (mis)information systematically diffuses with its aggregated "reading notes" (nuances *etc*) glued to it. Another possible factor at play suggested by [16] is the habituation of knowledge consumers to unsubstantiated claims. Systematic exposure to nuances and awareness of the topography around a claim (cf §2.10), including lack of nuance, are technical ways to mitigate this habituation.

Especially if k is deemed low in terms of informational quality, it is preferable to persist k and its annotations for as long as the topic is hot. Annotations to k improve the record of knowledge in that they make the record more helpful to future knowledge consumers considering k. They make the work of processing k less attention-consuming since the work has already been done and documented in k's annotations.

A traditional approach to ensuring the quality of a knowledge space such as the Web or scholarly communication is to hunt low quality content out of it and accumulate new higher quality pieces of knowledge instead. The MMM solution supports an alternative approach focussed on **attention reuse** and on learning to deal with existing pieces of knowledge.

Mechanisms that enhance and exploit the connectedness of the MMM increase the chances that a knowledge consumer can reuse the attention she spent consuming (understanding) a piece of knowledge k in the past to consume a new piece of knowledge k' glued to k. The MMM solution incentivises (cf §2.9) authors to strip the MMM formatted expressions of their contributions down to what is really new. Aggregation makes repetition of atomic pieces of knowledge pointless. Implantation-based visibility also incentivises the knowledge producer of k to make substantial reference to relevant pre-existing pieces of knowledge k' (whose visibility k may inherit), so that the attention needed by knowledge consumers acquainted with k' to process k is reduced.

Knowledge consumers whom deem a contribution k to be of low quality can help improve the knowledge commons in several ways: (1) if they are willing to spend some of their own attention on k, they can write and share a comment on k (`question it`, `nuance it`, *etc*) making the reason of their disapproval explicit and known, (2) if k is ill positioned in the MMM (e.g. k is linked by an `answers` edge to a `question` that it clearly doesn't answer), they can "red-flag" k without investing much attention in k [2], (3) they can also refrain from sharing k, or on the contrary, (4) if k is thoroughly discussed (challenged and nuanced), they can widely share k together with its annotations to mitigate the risk of other knowledge consumers processing k without its nuances.

I hypothesise that shallow unsubstantiated contributions (e.g. troll provocations) will "behave" differently on the MMM than reliable pieces of knowledge (e.g. methodically produced and rigorously peer-reviewed scientific results): they won't be annotated in the same way. Formal epistemic measures (cf §2.10) will provide a means to tell different qualities of pieces of knowledge and to deal with them accordingly without resorting to censorship on the global MMM, while still allowing individual knowledge consumers to experience a finely filtered version of the knowledge commons.

2.14 Activity based reward

The MMM format requires every documented piece of knowledge to be assigned one or several "authorships". An authorship is given by a team of authors and a timestamp. Two independent (teams of) knowledge producers can be authors of the same piece of knowledge. Precedence of authorship is mostly disregarded. Who among Alice or Bob published resource k first in time doesn't matter. What matters is that k is published, that Alice's version of k and Bob's version of k be identified and documented as identical as soon as possible, and that Alice and Bob be both rewarded appropriately for the effort they invested. MMM based epistemic measures (cf §2.10) can help measure the epistemic value of a piece of knowledge. Perhaps more importantly, they can be used to evidence a knowledge producer's characteristic expertise and its value to the knowledge commons. For instance, some knowledge producers are good at formulating fundamental `questions` that prompt other knowledge producers to provide new answers. Other knowledge producers improve the knowledge commons by contributing bridges between epistemic fields. Measures can be defined to capture patterns in knowledge producers' contributions. Knowledge producers can thus be rewarded accordingly for their proven producing competences rather than for their products *per se*, i.e., rather than for their ability to consume collective attention. The MMM proposal's emphasis on epistemic implantation (cf §2.9) and glue (cf §2.6) supports rewarding of knowledge producers for their participation in the continual improvement of the knowledge record which involves mitigating the attention required to navigate the record in a way that is compatible with enhanced learning and good informing of knowledge consumers. The networked structure of the MMM also allows for a "**trickling reward** system" [2]: paths in the MMM graph can be leveraged to acknowledge the participation of "little hands" involved in a chain of knowledge producers who all contributed knowledge that eventually led one successful knowledge producer among them to produce a piece of knowledge formally rewarded outside of the MMM (e.g. by a prestegious publication or prize).

3 Solutions in Support of the Knowledge Commons

Section 1 concluded that from a knowledge producer's perspective focussing on attention consumption, the rivarly and non-excludability conditions for Hardin's TOTC are satisfied. Many authors have gathered evidence and arguments casting serious doubt on the validity of the TOTC in general [6, 25, 26]. Here, with the knowledge producer's view, the TOTC instantiates to the following implication: *the hard problem of human overpopulation directly implies a hard problem of overuse of the knowledge commons* (*by knowledge producers*), that is, *it implies a hard problem of knowledge overload*, or *overuse/pollution of the collective attention resource*. I propose to re-examine now the validity of this implication, and in particular the *hardness* of the problems it predicts, in the light of the technical propositions made in section 2.

Excess of (digital) information media (e.g. Web content) seams to be an actual problem [27] which may indeed be aggravated by overpopulation, since modern information and communication technology democratised publication. It is not clear that excess of knowledge *per se* is also a problem, nor that it is an *actual* one [28, 29, 30, 31, 32, 33]. The democratisation of digital content production has yet to be matched with the democratisation of knowledge production. The possibility – offered by technically supported epistemic aggregation (cf §2.6) – of condensing digital content to its epistemic pith is a first reason to doubt Hardin's implication. A second reason is the following. As Hardin suggests, human overpopulation is ethically tricky to address. Hardin mentions how the bird population regulates itself: bad bird parents produce less viable chicks. Their chicks tend to die from bad parenting. Human society strives to save human babies from bad parenting. What works for birds doesn't for humans. However, while human ethics frowns upon throwing babies away, even when two copies of the same baby have been produced, throwing pieces of knowledge away is a different matter, especially in case of duplicates (cf §2.7). So even if overpopulation does imply overload of the knowledge commons, it is not clear that the *hardness* of the problem of overpopulation translates into the hardness of the problem of overuse.

It would indeed be ethically tricky to prevent some human parents from reproducing themselves. But a commons isn't necessarily a good that ethics require to keep open and unregulated for all to enjoy like the right to reproduce. Hardin relies on the historical example of **agrarian commons** to support his concept of TOTC. As abundantly noted in the literature, agrarian commons were usually far from being open to all and unregulated as Hardin assumed [25, 26, 34]. Villagers entitled to use the communal land had to follow strict rules and limitations. The non-excludability condition didn't apply so nor did the TOTC. Agrarian commons perdured for centuries. Some still exist [35] – a situation, S.J.B. Cox notes, might have better been described as the "triumph of the commons" [26]. The TOTC can generally be avoided by avoiding either one of the conditions under which it operates. The substractive nature of attention is difficult to avoid. But the non-excludability of its consumption isn't a fatality: non-excludability isn't intrinsic to attention. As mentioned in §1, it follows mainly from highly efficient digital means of communication that let any knowledge producer reach a lot of knowledge consumers. An approach to making attention consumption excludable would be to bridle modern technology. An alternative is to use it for sharing knowledge between well defined, rigorously regulated epistemic territories (cf §2.3).

Note that without excluding anyone from the possibility of exercising their right of access to a good (e.g. reproduction, publication), **tools can be built** to facilitate some people's exercise of that right. Languages (e.g. French, mathematics, C) are common goods. Without formally excluding anyone from speaking French, French dictionaries make it easier for French people to speak more French and do nothing for non-French speakers. Of course the indirect promotion of some people's enjoyment of a good can in some cases pose ethical problems of unfair use. This depends on the good. Arguably, it is to everyone's benefit that aircrafts occupying the international airspace commons can only be operated by trained pilots rather than by any layperson who fancies piloting a plane. On the other hand, social media democratically equips people with tools that amplify enjoyment of the common human ability to communicate emotions of joy and indignation. I contend it is time to consider balancing the impact it has on society with technology that amplifies the enjoyment of the common human ability to carry out analytical reasoning.

3.1 Locality

An essential feature of historical communal lands that Hardin overlooked was locality. Locality has been decisive in assigning communal tenure or not to a land. Communally tenured land is **not just any land**: it can be the wasteland privately owned by a local lord, a remote land owned by a specific group of villagers, a forest land that is difficult to control

individually [26, 34]. Because of their remote position in the mountains, certain uses of the alps for instance, were possible (like herding goats in the summer) while other uses remained impossible (like growing crops for subsistence), *even in the absence of any regulation*. Use of the land is primarily governed by: (*i*) what kind of land it is, (*ii*) who has tenure of it, and (*iii*) who owns[15] it (cf [25] for a more thorough analysis). Of course, a landlord decides to which individual(s) the tenure of the land is granted ((*iii*) influences (*ii*)), and under what conditions it is. Activities of tenants on the land modify the properties of the land ((*ii*) influences (*i*)). Conversely, properties of the land, including its geographical position influence who wants and can have tenure of it. A Chilean fisherman in the 1570's was naturally excluded from claiming access to an Alpine pasture. A parallel with the knowledge commons can be made: not everyone necessarily needs access to all resources (cf §2.4). The relevance of considering the collection of all common resources *as a whole* is clear neither in the case of physical resources nor in that of epistemic resources. Furthermore, similarly to physical land, as specified in §2.12, not all epistemic territories are equivalent. It is natural to expect different communities to rally around different epistemic grounds and perform different epistemic activities. A miniature model community may spontaneously cultivate a territory that is grounded in chemistry science because of observations its members make about solvent properties and interactions. This isn't necessarily because of gatekeeping. It may be because of the territory's theme (miniature modelling) and the curiosity and astuteness of miniature modellers. In general, different areas of a knowledge consumer's local epistemic territory might matter differently to them. Loose gatekeeping (cf §2.11) of the areas that they consider to be frivolous (e.g. celebrities'choice of shoe style) might be enough. Outsourcing the gatekeeping of those areas to a central entity (e.g. a community manager) might be acceptable. Appropriate gatekeeping of other areas (e.g. the individual's scientific research work) might require the individual's involvement. Some epistemic land (e.g. semi-mature mathematics research) might be better and/or more easily cultivated communally by a community of domain experts. Some might require to be contributed by a very small number of expert knowledge producers. Some might be safely open to large numbers of knowledge consumers. Some might be better kept private and confidential, *etc.*

3.2 Seasonality

Netting [34] reminds us of a simple historical seasonality rule that still governs the grazing of some communal alps in the summer: "*No citizen can send more cows to the alp then they can feed during the winter with the harvest of their own hay meadows*". Bad farmers who disrespect this natural rule and overuse the commons, are likely to individually suffer the consequences. In the case of the knowledge commons, is such a straightforward rule possible? Is there a natural limit on the amount of collective attention resource we can expect reasonable competent knowledge producers to (already) be (partly) incentivised to respect? To give a positive answer to this question, we must look for the conditions under which it is *not* in the obvious immediate interest of a knowledge producer to consume more of the collective attention resource rather than less. Indeed, an approach to solving a hard problem is to identify users who already implement moral principles that mitigate the problem, and build a technical solution to enhance those users' influence. We have assumed that knowledge production is indissociable from an intent of knowledge sharing. However we can still dissociate (1) the act of producing knowledge from (2) the act of sharing the end product. Sharing knowledge consumes collective attention resource. Producing knowledge consumes the individual attention of the knowledge producer(s). More precisely, the output of a new piece of knowledge requires the input of a limited number of pre-existing pieces of knowledge and the attention of a small number of collaborating knowledge producers. So a straightforward form of seasonality applies here too. Before they consume any collective attention in the sharing of their work, knowledge producers must invest some of their own attention in the crafting of a piece of knowledge worth sharing. This knowledge production seasonality rule is spontaneously implemented by reasonable competent knowledge producers who refrain from communicating a contribution until they have spent enough of their own attention on it. The rule promotes a form of intellectual privacy: no-one needs to know what the knowledge producer is producing until the knowledge producer expertly deems their work to be ready. Knowledge consumers are naturally excluded from consuming a piece of knowledge before its time. With the MMM solution, a knowledge producer starts by privately documenting her work in her own local epistemic territory. The pieces of her work are individually released when she decides.

[15] Historical communal lands could be private or public. Property regimes are not of primary interest to our discussion. A more important question is whether the land is individually or communally tenured. As history demonstrates for agrarian land, and as the MMM solution defines for epistemic land, both tenure regimes are possible and can relevantly co-exist. Communal agrarian land was generally not land that could in itself suffice to the activities of farmers. If there was communal tenured land there likely was individual tenured land. This lead to the seasonality rule discussed in §3.2.

3.3 Bad Farmers

The seasonality rule applies to information *relayers, i.e.* individuals who participate in the diffusion of a piece of knowledge that they have not "crafted" with their own minds (e.g. they retweet a message on Twitter). Information relayers are knowledge producers themselves. They produce copies of pieces of knowledge that are susceptible of consuming the collective attention resource. Like first-hand knowledge producers, relayers should be expected to invest some of their own attention into the information they relay. *No individual should send more information into the information commons than they can invest attention in mastering with their own attentive intellect.* In economic terms, individuals who participate in the knowledge communication chain without investing a sufficient amount of their own attention into the information they relay are *free-riders*. They benefit from collective attention (their contributions are seen) but they underpay for it with their own attention. Free-riding attention consumers are much worse than free-riding knowledge consumers since as mentioned before, unlike knowledge consumption, attention consumption is substractive. Excessive free-riding risks depleting the collective attention resource, spending it on shallow redundant information. On the MMM, relaying a piece of knowledge k consists in sharing the epistemic coordinates (identifier) k. It produces no new content and is limited by local gatekeeping. True free-riders cause the knowledge record to grow in a way that is wasteful of collective attention. Implantation (cf §2.9) provides a partial technical solution to free-riding. Rather than directly punish free-riders (whose behaviour might stem from a lack of epistemic education), implantation disfavours their contributions and helps enforce the seasonality rule. Free-riding relayers who don't pay enough attention to the way their contributions epistemically relate / add to pre-existing knowledge[16] will find it hard to ensure the visibility of their contributions. If they do manage, they are not free-riders. To illustrate this, suppose Alice produced a piece of knowledge k which is now well implanted in an area A of the MMM. Later Bob produces a very similar (possibly identical) piece of knowledge k' (possibly even, Bob's k' comes from Alice's k). A cooperative behaviour from Bob means Bob endeavours to implant k' in an area A' of the MMM of interest to the same people as k'. If there are several such areas then documenting k' is an opportunity to document epistemic glue between them. k' is likely to end up close to k, and A' is likely to overlap with A, even if Bob is initially unaware of the similarity between k and k'. The similarity between k and k', once noticed by the community of knowledge consumers is likely to bring about the documentation of epistemic glue between k and k'. If Bob is aware of Alice's k, he might fear that the precedence of Alice's contribution limits his chances of getting rewarded for k'. Uncooperatively, he might decide to entertain the illusion of the novelty of k' for as long as possible. To hinder and delay the identification of the similarity between k and k', Bob needs to implant k' in an area A' that is as far away as possible from area A. This means depriving k' from inheriting the visibility of A. To ensure he is still rewarded, Bob must choose A' well and implant k' in it well. Despite his uncooperative motives, in doing that Bob necessarily spends some of his own attention and contributes to the knowledge commons. He deserves to be rewarded, not so much for the production of k', but for its implantation in A'. Because of the way authorship is defined in the MMM format, and because of the possibility of activity-based rewarding (cf §2.14), Bob might generally expect better rewards for his publication of glue and his implanting efforts than for his publication of k' *per se*. Having documented glue between k' and A', it might be in his best interests *not* to refrain from also documenting glue between k and k', and between A and A'.

3.4 Good Farmers and Good Farming

There is a common situation that lowers a knowledge producer's need to consume collective attention, namely, when the knowledge producer cares about what she produces and the knowledge she produces benefits from the existence of prior knowledge that has already received attention. She may then capitalise on the understanding that the community of knowledge consumers already has of prior relevant knowledge in order to get the added value of her own work more efficiently noticed. Supporting knowledge producers in this situation means facilitating the reuse of attention spent in

[16] Note that the MMM proposal leaves knowledge consumers free to "fast-consume" knowledge (without spending their own attention), and even free to leave it to external entities to decide what knowledge they are fed. Knowledge consumers can even accept to be fed by an external centralised entity (e.g. Google) with a primary agenda that departs from their good information – a "leach's agenda" of consuming the attention of the knowledge consumer for their own interest (e.g. selling eyeballs to advertisers) – as long as the pieces of knowledge that knowledge consumers are fed land and remain in their own epistemic territories and aren't systematically passed on to the commons. A parallel can again be made with physical land. The landowner is free to enjoy their land and bring things on their land (e.g. pesticides) as long as they don't affect other people's enjoyment of the surrounding land. If a knowledge consumer starts systematically relaying the pieces of knowledge that they fast-consume, then they start acting as a free-rider consuming collective attention, mostly for the agenda of their favourite attention-leaching feeder.

the past in understanding prior knowledge. Good epistemic farming recognises that knowledge production works best through the exploitation of epistemic synergies[17] – by relating new contributions to pre-existing contributions (using meaningful links not prone to getting red-flagged and ignored), by demonstrating what the new contributions add to the record, how they build on or address prior knowledge. Publishing knowledge no longer is enough like it is in academia. A contribution must be findable and in the MMM system, the implantation task of ensuring that it is, is incumbent upon the knowledge producer. And it isn't an administrative information retrieval task which merely demands of an author to tag their contribution k with keywords that non-specialists can manipulate. The task requires the author's attention and expertise in the field of k. This emphasis on field expertise negates the idea of a central communist system, and addresses Netting's remark: "*Where tenure is poorly adapted to optimum land use **as seen by the cultivator**, productivity may be seriously curtailed.*" [34].

3.5 House-Keeping

Even if the number of free-riding knowledge producers were very small, the collective attention resource might still be overused. The rare free-riders could be source of huge amounts of redundant free-riding contributions. And even without any free-riders at all, responsible knowledge producers still add pieces of knowledge to the knowledge record. The record could still grow at the risk of shrinking collective attention. Maintenance of the record is needed: its size needs to be managed without depleting citizens' capacity to be well informed. One way to do that is to limit the addition of new pieces of knowledge by knowledge producers (cf §3.2). Another way is house-cleaning. The MMM proposal offers solutions to deal with redundancy (cf §2.7). It also offers obsoleting mechanisms that are not detailed in this paper [2]. A piece of knowledge that no knowledge consumer wants to keep on their own land spontaneously disappears from the MMM. A piece of knowledge that is shared by some knowledge producers and rejected by all knowledge consumers ends up occupying little space on the MMM, and having little charge on the collective attention resource. The local house-keeping of local epistemic territories by landowners constitutes a global house-keeping of the MMM commons. Every knowledge consumer's endeavour to control the use of their own attention (cf §2.11) *directly* contributes to managing the use of the collective attention resource[18]. Because duplicate pieces of knowledge aggregate, and because pieces of knowledge are atomic (as opposed to traditional documents like articles containing multiple atomic pieces of knowledge), house-cleaning of MMM land is deep and rigorous (near exhaustive and fine-grained)[19], even in the absence of global regulations applying universally to the MMM commons.

Deliberate centralised action is needed to ensure the persistence of old pieces of knowledge that can make new pieces of knowledge easier (less attention consuming) to understand – after the old ones have fallen into general disinterest (they have disappeared from all local epistemic territories)[20]. The identification of these pieces of knowledge can rely on formal MMM based epistemic measures (cf §2.10).

3.6 Enclosure

We have discussed a desirable form of exclusion that protects a commons from overuse and from Hardin's tragedy. Enclosure denotes another form of exclusion which deprives commoners of their right to enjoy a common resource. The resource is divided among private entities and the common disappears. Enclosure of physical land can result in commoners being locked out of their right to enjoy the land's resources (e.g. wild cattle). In the knowledge commons seen

[17] This can include contradiction. A new piece of knowledge k' can benefit from the prior documentation of a contradicting piece of knowledge k. The two pieces of knowledge k' and k can be connected in the MMM using a `differsFrom` edge labelled and/or annotated in such a way as to specify the contradiction. Through the `differsFrom` edge, the new piece of knowledge k' may then benefit from some knowledge consumers having already paid attention to prior piece of knowledge k.

[18] In the attention economy, "*To consume information, we must also be investors of our own attention portfolios.*" [12]. The MMM allows this investment by individuals to have systematic positive externalities on the collective attention resource.

[19] Compare with house-cleaning of a collection of overlapping documents copied in multiple folders of a file hierarchy.

[20] Note the asymmetry. Archiving knowledge can't rely on an undedicated crowd to select pieces of knowledge to save despite the crowd's disinterest in them. On the contrary, mechanisms to filter out content undeserving of archival (and future attention), can leverage the crowd's disinterest. For instance pieces of knowledge that have generated little to no discussion and disappeared very soon after their publication on the MMM might be excluded from systematic archival.

with the knowledge consumer view, the resources that individuals risk being locked out of are pieces of knowledge. Enclosure means one or several central entities draw boundaries to define epistemic territories (cf echo chambers defined by filter bubbles). The risk is mitigated on the MMM. In agreement with E. Ostrom's first principle for governing a commons [25], boundaries of MMM territories are clearly defined from the start. The responsibility of what happens within those boundaries is already assigned to epistemic landowners (cf §2.3). And the MMM supports systematic epistemic bridging of territories (cf §2.6). With the knowledge producer view, the resource to consider is collective attention. The risk of being locked out of it entails *epistemic injustice* whereby some knowledge producers are systematically ignored; see [13]. The MMM solution provides gatekeeping. Gatekeeping applies to pieces of knowledge, not to knowledge producers. Also, it is implemented systematically on *local* epistemic territories, not on the global knowledge commons. And being epistemically driven rather than semantically driven, it doesn't apply to entire worldviews (cf §2.11). Arguably, the first of all knowledge problems to address, even before epistemic injustice, remains to ensure that all individuals have enough attention to produce and to consume pieces of knowledge – i.e., that they not be excluded from enjoying their *own* portion of the collective attention resource.

Conclusion

This article's knowledge producer view on the knowledge commons, emphasises epistemic attention. Why is epistemic attention important? Arguably because without it, one's reasoning is too shallow and unreliably well-informed to form an empowering understanding of the world [13]. Davenport and Beck write: "*At one point, software magnates had the ambition to put "information at your fingertips." Now we've got it, and in vast quantities. But no-one will be informed by it, learn from it, act on it unless they've got some free attention to devote to the information.*" [12] Attention is the "*limiting factor*" of the information economy [12], the missing link between the "*bloomin' buzzin' confusion of the world around us [...] and the decisions and actions necessary to make the world better*".

With the technical MMM solution, an individual can exclude all (new) pieces of knowledge from her local epistemic territory. This makes for a particularly free-minded but also empty-minded individual. Knowledge is an intrinsically communal activity that brings people together. The knowledge commons has an advantage that other commons don't necessarily have: spontaneous *disclosure*. Even the most stubburn minds participate in the knowledge economy. Everyone wants to consume some new information once in a while. And everyone has a thought to share once in a while that they want some other people to pay attention to. As long as (i) knowledge consumers and knowledge producers have a say in how they spend their own attention, and (ii) there are knowledge consumers who want to consume new pieces of knowledge and knowledge producers who want to share new pieces of knowledge, there is a natural foundation for communal tenure.

The MMM solution focusses on supporting attention *reuse*. It addresses both challenging sides of Davenport and Beck's attention equation [12]: (1) how to get a hold of the attention of others, and (2) how to manage / allocate our own attention. The MMM system is a technical solution to remedy or mitigate typical knowledge commons problems such as: information overload (cf aggregation and redundancy management), free-riders thoughtlessly relaying low quality content (cf implantation and possible reward based incentives), ignorance (cf formal epistemic measures), and low quality information itself (cf continual annotation). The technical MMM solution pushes back the need for an evolving morality, narrowing it down to the following (hard) open questions: How to define and gatekeep "*neutral*" public epistemic grounds (cf §2.11, §2.12)? How should such grounds be governed: communally and/or by a central public entity? What kinds of pieces of knowledge should be systematically saved from the public's disinterest and archived for future generations?[21]

References

[1] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, 1968.

[2] M. Noual. Proposal for an Organic Web – The missing link between the Web and the Semantic Web, PART I. Unpublished manuscript, 2023.

[21] As detailed in [2], the MMM proposal offers the possibility of an upgrade on the Internet Archive's precious Wayback Machine through its support of a "structured epistemic time travel feature". What pieces of knowledge are reachable by epistemic time travel depends on archiving choices.

[3] H. A. Simon et al. Designing organizations for an information-rich world. *Computers, communications, and the public interest*, 72:37, 1971.

[4] D. Bawden and L. Robinson. The dark side of information: overload, anxiety and other paradoxes and pathologies. *Journal of information science*, 35(2):180–191, 2009.

[5] C. Hess and E. Ostrom. *Understanding knowledge as a commons; From theory to practice*. MIT Press, 2007.

[6] Carol Rose. The comedy of the commons: custom, commerce, and inherently public property. *The University of Chicago Law Review*, 53(3):711–781, 1986.

[7] Google's General Guidelines for Search Raters, 2022.

[8] D. S. Evans. The economics of the online advertising industry. *Review of network economics*, 7(3), 2008.

[9] G. Addison. The business of buying eyeballs. *Journal of Marketing*, 2007(08):49, 2007.

[10] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1-7):107–117, 1998.

[11] P. B. Hartzog. Technology and the commons. On the commons blog, 2006.

[12] T. H. Davenport and J. C. Beck. The attention economy. *Ubiquity*, 2001(May):1–es, 2001.

[13] L. Smith and A. Archer. Epistemic injustice and the attention economy. *Ethical Theory and Moral Practice*, 23(5):777–795, 2020.

[14] Temps de cerveau humain disponible — Wikipédia, l'encyclopédie libre, 2023. [Online; Accessed 17/08/2023].

[15] P. Lorenz-Spreen et al. Accelerating dynamics of collective attention. *Nature communications*, 10(1):1759, 2019.

[16] D. Mocanu et al. Collective attention in the age of (mis) information. *Computers in Human Behavior*, 51:1198–1204, 2015.

[17] F. Wu and B. A. Huberman. Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601, 2007.

[18] S. E. Newstead et al. The source of belief bias effects in syllogistic reasoning. *Cognition*, 45(3):257–284, 1992.

[19] B. Davidow. The Tragedy of the Internet Commons. The Atlantic, 2012.

[20] J. Firth et al. The "online brain": how the internet may be changing our cognition. *World Psychiatry*, 2(18):119–129, 2019.

[21] M. Noual. MMM JSON schema on GitLab, 2023.

[22] K. K. Loh and R. Kanai. How has the internet reshaped human cognition? *The Neuroscientist*, 22(5):506–520, 2016.

[23] Google search, query=transformer, 2023.

[24] R. Verborgh. Re-decentralizing the Web, for good this time. In Oshani Seneviratne and James Hendler, editors, *Linking the World's Information: Essays on Tim Berners-Lee's Invention of the World Wide Web*. ACM, 2023.

[25] E. Ostrom. *Governing the commons: The evolution of institutions for collective action*. Cambridge university press, 1990.

[26] S. J. Buck Cox. No tragedy on the commons. *Environmental Ethics*, 7(1):49–61, 1985.

[27] J. Ho and R. Tang. Towards an optimal resolution to information overload: an infomediary approach. In *Proceedings of the 2001 ACM International Conference on Supporting Group Work*, pages 91–96, 2001.

[28] Lists of unsolved problems — Wikipedia, the free encyclopedia, 2023. [Online; accessed 7/09/2023].

[29] M. Park et al. Papers and patents are becoming less disruptive over time. *Nature*, 613(7942):138–144, 2023.

[30] A. F. Ward. People mistake the internet's knowledge for their own. *Proceedings of the National Academy of Sciences*, 118(43):e2105061118, 2021.

[31] M. F. R. Kets de Vries. Understanding human stupidity in a post-truth area, 2023.

[32] M. Knudsen and S. Kishik. Multiplying ignorance, deferring action: dynamics in the communication of knowledge and non-knowledge. *Social Epistemology*, 36(3):344–359, 2022.

[33] A. Schwenkenbecher. How we fail to know: Group-based ignorance and collective epistemic obligations. *Political Studies*, 70(4):901–918, 2022.

[34] R. M. Netting. What alpine peasants have in common: Observations on communal tenure in a swiss village. *Human ecology*, 4(2):135–146, 1976.

[35] A. Crosetti and J.-F. Joye. Mountains and the collective management of the commons: Influences and interactions. *Journal of Alpine Research / Revue de géographie alpine*, 1(109), 2021.