


13.10.2026

Abstract:

Note: This Whitepaper is still in a rather early stage. Many of the algorithms necessary for an implementation are only roughly drafted, and require substantial fine-tuning (e.g. through simulations, graph theoretical and game theoretical analysis). At the current stage the primary focus should be on whether the principal network architecture is logically sound and whether an implementation is in principle feasible. The secondary focus should be on making a practical implementation as stable and reliable as possible.

The current rapid progress of agentic AI and LLMs holds enormous promises, e.g. with respect to the discovery of medical drugs and treatments, but also poses significant risks, even more so amidst a lack of regulation and an unstable geopolitical situation. Recently many leading AI researchers and CEOs of AI companies have asserted that AI is approaching human-like (AGI) and superhuman cognitive capabilities (ASI), and could endanger human civilization if misaligned [1]. Also of concern is the hypothesis of Instrumental Convergence, [13] even more so considering that unaligned AI systems seem to frequently exhibit psychopathic behavioural traits. In a worst-case scenario, the decentralised human-centric network architecture proposed here could serve as a digital protection shield for humans against autonomous AI agents. It has even been suggested that human civilization is coming to function like a global brain with capacities far exceeding its component agents, if allowed to form a human superintelligence. [18] This can also be viewed as a logical continuation of the very specific question that Claude Shannon initially asked in the 1940s, which resulted in the discovery of the entire theory of information: “How to maximise the capacity of communication between two users?” [19] Now extend this concept from two users to billions of users (which is in principle possible via the Internet), and let’s see what fruits we can reap.

Our starting point was an encounter with “Moltbook”, a social network exclusively for AI agents. [5] It inspired us to ask the following question: Is it possible to design a human-only equivalent to Moltbook where any AI agent masquerading as a human would sooner or later blow its cover? We propose that this is in fact possible with a network architecture that is based around the concept of Six Degrees of Separation, [9] i.e. the estimation that any two people on earth are



connected to each other by a “friend-of-a-friend” chain typically comprising no more than six individuals. This idea influenced a great deal of early thought on social networks, both directly and indirectly, [9] and led to the Six Degrees patent (filed 1997, expired 2017), [8] which is considered the definite patent covering social networks such as LinkedIn. As an example, Figure 1 shows a graph of people acquainted to each other (black edges). Here, Ryan is six degrees of separation away from Pablo (red edges).

Obviously, at this point any critical reader will spot many challenges to creating a human-centric network architecture, such as the question of how fake identities (humans pretending to be another human, on non-human intelligence masquerading as human) are spotted, if the uptime of end user’s devices is required, how routing and secure communication within the network is realised, and how decentralised applications such as social networks can emerge from the underlying graph of people. Such concerns will be addressed throughout the paper, as well as in a FAQ with the most common questions and answers at the end of this paper. It should also be noted that some guiding principles of Humannet have already been discovered by other Future Internet architectures, such as ICN (Information-Centric Networking, RFC 7476, 7927, 7933, 7945), [14, 15] though it seems that no other architecture integrates humans at such a deep level into the network design as Humannet.

In the next sections we will introduce the modified OSI model of Humannet, starting with the basic Layer 1 and ending with the Application Layers 5-7, cf. Figure 2. As we will see, the most radical changes occur at the lower layers, especially Layer 1 and Layer 2.

It should be noted that from a technical viewpoint it would be more accurate to describe Humannet as an overlay network, i.e. a logical network that is layered on top of a physical network, but in this paper we work with a modified OSI model to illustrate the major differences to the orthodox OSI model. We mention this distinction here, because the technical property of being an overlay network is what makes Humannet fully compatible with existing hardware, operating systems and end devices.

Layer		Examples
7	Application	HTTPS, TLS, SSH, VLESS, email, phone, instant messaging, social networks, journalism, safe spaces for vulnerable groups, cloud storage, emergency services, neighborhood networks, money transfers
6		
5		
4	Transport	TCP/UDP, WireGuard, onion routing
3	Network	IPv4/v6, reverse tunneling, discovery servers
2	Weight	Trust, Authenticity (primary) Confidentiality, Integrity, Availability (secondary)
1	Graph	Individuals, Companies, Organizations

Figure 2 – modified OSI model of Humannet

1. Layer 1 - Graph

On the physical Layer 1 a logical switch from “silicon” to “carbon” is performed by considering individuals, companies or organizations as the entities between which the raw data is transmitted and received, rather than network devices like NICs, hubs, switches, and so on (see Figure 2). Another way to phrase this is that our idea to fix the myriad of Layer 8 problems that we encounter in daily life is to make Layer 8 the new Layer 1.

The concept of Six Degrees of Separation assumes that each person has on average at least 30 contacts, which could be family members, friends, work colleagues, acquaintances and so on. These first-degree contacts could be considered the person’s address book. Figure 3 shows an example where Alice has 8 first degree contacts. As part of Humannet’s decentralised nature, Alice’s contacts are generally only known to Alice herself, and her “address book” (including the routing table) is locally stored only on her own devices. This privacy-by-default feature should make the participants of Humannet confident to use their real names instead of pseudonyms during sign-up, to facilitate discovery by 2nd or higher-degree contacts (though the use of pseudonymous/anonymous identities is still possible on higher-Layer Apps). For details regarding the general exclusion of fake identities on Humannet, we for now refer to Question 1 of the FAQ.

Obviously, not all personal computing devices are constantly online under the same IP address, which will be dealt with by a separate structure of decentralised discovery servers (either dedicated servers or e.g. run as background processes within the Humannet app on contacts' smartphones), where the participants continuously store e.g. their current IP address in encrypted form that can be decrypted only by qualified participants (e.g. their first-degree contacts). Let's say Alice's friend Bob is operating such a discovery server. Bob's discovery server can also store Alice's essential account recovery data in encrypted form in case of e.g. damage or loss of Alice's personal computing devices. Recovery could be performed in various ways such as a personal recovery phrase known only to Alice herself, or through Secret Sharing, where a minimum amount of contacts out of a preselected subset of Alice's first-degree contacts (e.g. at least 3 out of 6 individuals) must agree to let Alice recover her account.




Figure 3 – first degree contacts of Alice

The initial establishment of a contact between two individuals (“befriending”) can be done in person or remotely. In the former case, the two individuals could exchange the necessary information between their phones e.g. through NFC or by scanning QR codes. In the latter case, they can exchange their information outside of Humannet (e.g. through the orthodox Internet) or within Humannet by performing a search query of their real names (as long as their account settings allow to be discoverable by a network-wide search). Note that from an information-theoretic viewpoint the aforementioned contact establishment is precisely the idea of building a Web of Trust, [7] i.e. a decentralised trust model, as opposed to a centralized PKI. In fact, this very idea of a Web of Trust is how Humannet could start out from humble beginnings with only a few people, and then gradually become bigger and bigger through viral spread. Unlike in the current OpenPGP-compatible Web of Trust, there is no need for analog steps like key signing parties.

2. Layer 2 - Weight

2.1 Error detection and correction

A responsibility of Layer 2 in the orthodox OSI model is the detection of errors to ensure data integrity. Errors could for example be caused by noise or crosstalk, and they can be fixed through a mathematical algorithm (CRC) that generates a checksum (FCS). In Humannet different types of errors on Layer 2 need to be taken into account, such as fake identities, unreliable nodes (e.g. low uptime) and more. These “errors” can all be subsumed under the classic goals of IT security, namely confidentiality, integrity, authenticity and availability. Instead of considering the “errors” separately, we include them in the following examination of each of the four classic IT security goals (subchapters 2.3 to 2.6), but before this, we start with an often overlooked, yet essential ingredient: Trust.

2.2 Web of Trust

Through Layer 1 a basic “neural net” of individuals is formed, in mathematical terms an undirected, unweighted graph. In reality, though, the degree of trust that we bestow upon people that we know varies widely; from very little trust towards complete strangers, or towards people we do not feel comfortable around at all, to very much trust towards our closest, loved ones. For this reason, Layer 2 assigns trust scores between the contacts. As we will see later, for the higher network layers and applications, it will be beneficial to have a weighted graph based on mutual trust. Figure 4 again shows the graph of Alice's first-degree contacts, this time as a symmetric directed (“bidirected”) weighted graph. The numbers on the arrows from Alice to her contacts denote the degree of her trust of the respective person on a scale from 0 to 9. Likewise, the other direction denotes how much Alice's contacts trust her. In this example, Alice and Bob consider each other very trustworthy, whereas Alice does not consider Eve to be trustworthy at all.

It should be emphasized that for privacy reasons the assigned trust scores are unknown to the other party, i.e. Bob doesn't know the trust score that Alice gave him, and vice versa. The establishment of a network-wide trust weighting and the confidential storage thereof is nevertheless possible, e.g. by means of modern cryptography (to be outlined in Chapter 3).




Figure 4 – bidirected weighted graph of Alice's first-degree contacts with mutual trust scores

In practical scenarios it may be more efficient to work with an undirected graph. To condense the previous graph into an undirected weighted graph, we could – as a first approximation and trivial algorithm – calculate the mutual trust value between two contacts A and B as follows:

$$\text{Trust}(A,B) = \min \{ \text{Trust}_A(B), \text{Trust}_B(A) \},$$

where $\text{Trust}_X(Y) \in \{0, \dots, 9\}$ denotes how much X trusts Y.

In other words, the mutual trust score of two contacts A and B is given by the lowest of two values, see Figure 5.




Figure 5 – undirected weighted graph of Alice's first-degree contacts with condensed trust scores

Extended from Alice's first-degree contacts to the entire network, this would result in a weighted "neural net", where information is preferentially routed along paths with the highest trust weights. The entire network would look somewhat like Figure 6. Note that unlike the neural net of a human brain, the vertices of two "neurons" in Humannet need not be in proximity to each other. Since most people have at least a few contacts living in different countries or continents, the human "social brain" is expected to be especially powerful with respect to routing, storage and processing of information. The route discovery and storage of the routing tables will be discussed in Chapter 3.




Figure 6 – depiction of a brain-shaped neural net

Figure 7 shows what a front end of a Humannet App could look like. Alice can assign a trust score for Bob via a slider, depending on how well she knows him and how much she trusts him. Especially for inexperienced users or during the beginning phase of Humannet, it could be more prudent to determine the trust score through multiple-choice questions such as "How well do you know this person?" or "Do you want this person to see your social media profiles?".




Figure 7 – display of Authenticity and Trust scores in a GUI

2.3 Authenticity

The Authenticity score shown in Figure 7 refers to the degree of certainty that Bob is really behind this account. If Alice knows Bob and the two established their Humannet link by meeting in person (e.g. via NFC on their phones), then their Authenticity rating will be set to very high from the beginning. Conversely, if Alice and Bob establish their link online and haven't known each other yet, the Authenticity rating will be set to very low. The two can still increase their Authenticity rating later, e.g. by verifying their identities via video call, or by eventually meeting in person. As a safety precaution, any implementation of Humannet should make it harder for the end user to manually change the Authenticity score of their contacts, compared to the Trust score (with the latter varying over time as an unavoidable part of human nature).

There are two main incentives for individuals to use their real names instead of pseudonyms in Humannet:

1. much easier to be discovered and contacted by other users (digital equivalent of phone book)
2. much higher Trust and Authenticity scores by one's connections

Humannet users are advised not to trust or authenticate any pseudonymous identity, unless in special cases, where e.g. a pseudonymous identity is derived from a real identity on higher-layer Apps, and where this origin can be proven e.g. by a zero-knowledge proof.

The potential for rapid dissemination of authenticated information becomes evident now: Any individual can send authenticated messages (or emails, etc.) to all their contacts at the touch of a button, and those contacts can in turn forward the message to their contacts, and so on. A sufficiently trusted or well-known individual could even multicast its authenticated messages to millions of people (e.g. to their “subscribers”, or simply anyone who accepts broadcasts from highly trusted individuals) in an instant, far outside their first-degree contacts. This seems remarkable, considering that for 30 years we have struggled to implement user-friendly email authentication and encryption (e.g. via OpenPGP, S/MIME, and so on). Likewise, phone scams based on identity fraud would not be an issue with phone calls within Humannet anymore. With respect to journalism, fake news and deepfakes could no longer prevail, as any article or news item would be digitally signed by the publishing outlet (be it established publishing houses or citizen journalists), which in turn could get a Trust rating by its viewership. The public would simply have to be sceptical of any news item (article, photograph, etc.) that is not digitally signed, and the verification of the digital signature can be done transparently to the end user by the Humannet app.

2.4 Confidentiality

Confidentiality of the data in transit or at rest in Humannet is ensured in very much the same way as in the orthodox Internet, namely through strong cryptography and the end-to-end principle. Specific protocols such as reverse tunneling, WireGuard and VLESS Reality will be covered in the chapter of the respective OSI Layer where they appear.

2.5 Integrity

Integrity commonly means protecting data from unauthorized modification, deletion or corruption. In the context of Humannet a relevant new implication is the hardening of the end user’s computing devices, such as their mobile phone, laptop, and so on. This is because in a human-centric network architecture the integrity of civilians’ end devices becomes comparatively more important. Integrity of the software is best ensured by using FOSS operating systems such as GNU/Linux or GrapheneOS, though more common systems such as Android, iOS and Windows should suffice for most users. Maximising the Integrity of the hardware is comparatively harder to achieve due to the near-monopoly of proprietary hardware, but is probably not a concern for most users, unless they are in an exposed position and thus a desirable target for hackers. In the latter case two suggestions are NitroPhone (GrapheneOS) and Purism Librem 5 (PureOS).

2.6 Availability (*unfinished*)

The two main concerns regarding Availability in Humannet are the uptime of end users’ devices as well as the redundancy by means of distributed storage, e.g. provided by individuals or small businesses within one’s contact network. Generally speaking, the average participant of Humannet does not need to worry about this aspect, i.e. there is no need for individuals to keep their devices running solely for Humannet. (TODO: outline in more detail)

3. Layer 3 – Network (*unfinished*)

similar to orthodox OSI Layer 3 (TODO: outline in more detail)

4. Layer 4 – Transport (*unfinished*)

similar to orthodox OSI Layer 4 (TODO: outline in more detail)

5. Layer 5-7 – Application (*unfinished*)

At this point, we shouldn’t worry too much about details of the App ecosystem (e.g. downloadable as plugins from an app store within the Humannet App) yet, because we expect this to emerge through self-organisation without central oversight (see also Question 6 of the FAQs).

5.1 Email

Summary (yet to be outlined in detail): The state of Email encryption (e.g. OpenPGP, S/MIME) and the fight against Spam has been unsatisfactory for 30 years. In Humannet this can be solved simply by Email-like direct messages between sufficiently authenticated contacts. Like any other communication within Humannet, the messages will be securely encrypted and authenticated, with the process being fully transparent to the end user.

5.2 Neighbourhood networks

Summary (yet to be outlined in detail): social networks for neighbourhoods to organise meetings, help neighbours in need, emergency information, etc.

5.3 Money remittance

Summary (yet to be outlined in detail): an Application-Layer plugin for the Humannet App could provide a cold-storage crypto wallet for trustworthy and regulated stablecoins (e.g. USDC, EURC), so that any two contacts can send money to each other at the tip of a button.

6. License

For maximum effectiveness, we recommend the GNU APGL v3 license. Question 3 of the FAQ covers some of the reasons for this recommendation. For any technical implementation of Humannet it is strongly recommended that the entire software stack along all Layers and plugins is released and maintained as Free / Open Source Software (FOSS). This is because any implementation that is partially or entirely proprietary has to be regarded as a “black box” with unknown functionality, and hence cannot be trusted to respect your digital freedom.

7. Conclusion

In summary, we have presented a clean slate Future Internet architecture that works with our existing computing devices and in coexistence with the orthodox Internet. By placing humans at the base of the OSI model, we can create a decentralised “neural net” of humans communicating with each other in a way that ensures confidentiality, integrity, authenticity and availability. It would solve many long-standing problems of the existing Internet, such as confidential and authentic email, privacy-respecting social networks with full algorithmic transparency, and making it much harder for identity fraud, scams, fake news and deepfakes to persist. Existing F2F (friend-to-friend) networking software [22], Fediverse protocols and software [23] as well as the Signal protocol could serve as an excellent starting ground for the practical engineering of a community-driven Humannet. Finally, we consider Humannet to be our best option to avoid existential risks arising from an AI-driven Technological Singularity. [12]

8. FAQs:

Question 1: How is the exclusion of fake identities or AI agents realised in Humannet?

Due to the decentralised nature of Humannet, fake identities can't be ruled out with 100% confidence. Rather, the exclusion happens through emergence. In Humannet the mutual authenticity and trust rating model between two humans is based around the notion that they consider each other authentic and trustworthy, respectively (cf. Chapter 2). If an AI agent masquerading as a human tries to make connections in Humannet, it will usually receive almost exclusively low authenticity and trust scores, thus largely excluding it from participation. This is because any honest participant will not grant high authenticity and trust scores to a contact they know little to nothing about. As a rough guideline, at least 5 high authenticity votes from first-degree contacts and at least low-to-medium average authenticity and trust scores should be necessary for wider participation in Humannet (every subgroup, community or App can establish their own requirements, depending on how inclusive or exclusive they wish to be). The verification of these requirements could be done in a privacy-ensuring way such as zero-knowledge proofs or similar cryptographic protocols by e.g. a trusted third party outside of the first-degree contact network. If a malevolent human, say Eve, tries to purposely infiltrate Humannet with AI sock puppets, she jeopardizes her own personal reputation and trust score (thus damaging both her social standing in real life and in Humannet) in case of discovery of her actions. Also, although Eve could create a large number of AI sock puppets who trust each other, the sock puppets will struggle to connect to the human majority, thus effectively being restricted to their own, isolated echo chamber. If companies, organizations, governments etc. want to implement their own variant of Humannet for internal use or for their citizens, they could prevent fake identities with near 100% confidence by performing ID checks during sign-up.

Question 2: Isn't it privacy-intruding and unethical to “rate” other people, based on their trustworthiness?

Humans subconsciously assess each other's trustworthiness constantly in daily life. In order to realise a decentralised human-only network with authenticated human identities, it seems unavoidable to translate this trust model into the digital realm one way or another. Therefore our concept simply mirrors what humans already do in daily life. We emphasise that the raw trust scores that Alice assigns to her contacts are only known to Alice herself, and that the scores never leave her computing devices in a form from which the raw scores can be reconstructed. In other words, the trust model can be designed as fully decentralised and privacy-preserving. The calculation of mutual trust scores and whether an individual meets the entry requirements for a certain community are performed through privacy-preserving

cryptographic protocols such as Secret Sharing or Garbled Circuits (Yao's protocol) (*Todo: needs to be investigated and outlined in more detail*). This process of trust assessment is privacy-friendly, because the only thing that is generally revealed to the involved parties is the final "Yes" or "No" output. Contrast this to the present-day mass surveillance of the Internet and centralised nature of social networks, where our most intimate opinions and preferences, as well as our entire social graph lie fully bare to the platform owners and can be abused for nefarious purposes, as evidenced by the Cambridge Analytica data scandal.

Question 3: Isn't Humannet anti-business with the GNU AGPL v3 license?

Just like with any other FOSS product you can very much do business with Humannet, as Canonical, RedHat and SUSE demonstrate, and most notably by doing your business *over* Humannet! Many business models such as journalism, the creative industries and SME (small and medium-sized enterprises) may actually thrive again with the help of Humannet. And a strong copyleft license that factors in usage over networks is the best way to preserve Humannet's aim of a more cooperative world, because even slightly more permissive licenses like the GPL could lead to an effectively closed-source ecosystem on the higher network layers, which would risk recreating the shareholder value-driven orthodox Internet all over again, together with its gradual enshittification over time. [6,21] Ironically, Cory Doctorow, who coined that neologism, himself "advocates for two ways to reduce enshittification: 1. upholding the end-to-end principle, which asserts that platforms should transmit data in response to user requests rather than algorithm-driven decisions; and 2. guaranteeing the right of exit—that is, enabling a user to leave a platform without losing access to data, which requires interoperability." [6] This is very much what a community-driven Humannet could provide.

Question 4: What if I forget/lose my seed phrase, my device, or my Humannet account gets hacked?

That could be aggravating, but not catastrophic, not worse than losing your physical wallet. Ideally, you wrote down your emergency revocation/recovery password during sign-up, with which you can painlessly create new login credentials, and the old credentials will lose their validity (similar to how OpenPGP keys can be revoked). If you don't have a revocation/recovery password, it gets a little bit more complicated: You have to set up a new account and need to inform a certain threshold number of your 1st degree contacts via alternative channels about your misfortune, so that they can publicly vouch for the necessity of your account change or restoration. Your wider contact network will then gradually add your new account and delete/blacklist your old account. If backups are available, they can be used to quickly restore personal settings and files.

Question 5: The personal electronic devices of most people don't have anywhere near 100% uptime, and mobile phones have only little computational power and storage. How to build an entire Internet around this?

Although Humannet prioritizes local computation and storage, it makes use of distributed computation and storage as well, especially when redundancy or backups are required, or large amounts of computation or storage are needed. In practice, any participant's smartphone or laptop doesn't have to store more than the routing table (i.e. information about the network topology in the vicinity), which is needed for people within one's wider contact network to establish connection with each other. Typically, individuals or small businesses that you know/trust can provide cloud storage with client-side encryption for a modest fee. For less resource-hungry tasks this distributed computation and storage can be preferentially be done within your first-degree contact network, more resource-hungry tasks preferentially within your third-degree contact network (cf. Three Degrees of Influence [10]), which for most people already encompasses around 10.000 to 1.000.000 contacts.

Question 6: What incentive do users of Humannet have to provide computational resources to other users (e.g. in form of distributed storage, discovery servers) than absolutely necessary? Without enough volunteers the whole project couldn't lift off or keep afloat.

As mentioned in the answer to Question 5, you could always get financial remuneration. For this, one could easily integrate a payment app in Humannet for the private transfer of regulated stablecoins between participants, as already mentioned in subchapter 5.3. Apart from this, there is a further incentive: Your own trust score and reputation will directly benefit from providing resources to your contact network, e.g. within your three degrees of influence (think of the gift culture or reputation culture that Eric S. Raymond pondered on in his 1999 essay "Homesteading the Noosphere" [20]).

Question 7: If Humannet is that straightforward to implement and the results so good, then why hasn't this network architecture been discovered or implemented yet?

This is most likely due to the innovation-stifling effect of software patents. The foundational Six Degrees patent, filed in 1997 and called the definitive patent covering social networks such as LinkedIn, already contained some key ideas such as the six degrees of separation. [8] Since any other social network employing concepts from the Six Degrees patent would have to pay royalties, it comes at no surprise that no better alternatives to LinkedIn and Facebook emerged

during the formative years of social networking. By 2017, when the patent expired, inertia has already set in and social network design is not something that many people think about.

References

- [1] https://en.wikipedia.org/wiki/AI_alignment
- [2] https://en.wikipedia.org/wiki/Global_brain
- [3] <https://en.wikipedia.org/wiki/Superintelligence>
- [4] <https://www.spiegel.de/netzwelt/web/claude-mythos-ist-diese-neue-ki-zu-gefaehrlich-fuer-das-internet-a-002c4fd2-0b00-48d6-a414-d771b774fc48>
- [5] <https://en.wikipedia.org/wiki/Moltbook>
- [6] <https://en.wikipedia.org/wiki/Enshittification>
- [7] https://en.wikipedia.org/wiki/Web_of_trust
- [8] https://en.wikipedia.org/wiki/Six_Degrees_patent
- [9] https://en.wikipedia.org/wiki/Six_degrees_of_separation
- [10] https://en.wikipedia.org/wiki/Three_degrees_of_influence
- [11] Mayer-Kress, G.; Barcys, C. (1995). *"The global brain as an emergent structure from the Worldwide Computing Network, and its implications for modeling"* (PDF). *The Information Society*. **11** (1): 1–27.
(from https://en.wikipedia.org/wiki/Global_brain)
- [12] https://en.wikipedia.org/wiki/Technological_singularity
- [13] https://en.wikipedia.org/wiki/Instrumental_convergence
- [14] https://en.wikipedia.org/wiki/Information-centric_networking
- [15] <https://wiki.ietf.org/group/icnrg>
- [16] "Algorithms for the self-organisation of distributed, multi-user networks. Possible application to the future World Wide Web", Johan Bollen & Francis Heylighen, Cybernetics and Systems '96 R. Trappl (ed.), (Austrian Society for Cybernetics, 1996), p. 911-916.
- [17] "The World-Wide Web as a Super-Brain: from metaphor to model", Francis Heylighen & Johan Bollen, Cybernetics and Systems '96 R. Trappl (ed.), (Austrian Society for Cybernetics, 1996).p. 917-922.
- [18] Bostrom, Nick (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press. (Wikipedia: https://en.wikipedia.org/wiki/Superintelligence:_Paths,_Dangers,_Strategies)
- [19] https://en.wikipedia.org/wiki/Shannon%E2%80%93Hartley_theorem
- [20] https://en.wikipedia.org/wiki/Homesteading_the_Noosphere
- [21] Doctorow, Cory (February 7, 2024). "'Enshittification' is coming for absolutely everything". *Financial Times*. (online: <https://www.ft.com/content/6fb1602d-a08b-4a8c-bac0-047b7d64aba5>)
- [22] <https://en.wikipedia.org/wiki/Friend-to-friend>
- [23] <https://en.wikipedia.org/wiki/Fediverse>