



Edgar Degas, *At the Stock Exchange* (1878–1879).
Musée D'Orsay, Paris, France

Cody Moser

The Ties That Bind Us

THE WISDOM OF CROWDS

We live in a world that is smaller than ever. In certain portions of the world, access to high-speed internet is more readily available than is access to water, with recent reports suggesting that over 6.5 billion people own a smartphone, more than the 5.8 billion that have access to clean drinking water. With as many people online as there are today, one would expect the intelligence of the human collective to have settled many of the divisions of the modern age, but we do not seem to be any the wiser. In his book *World Order*,¹ American geopolitical strategist Henry Kissinger argued that the Internet Age would not bring about the broadening of our wisdom, but indeed the opposite, noting, ‘the computer supplies tools unimaginable even a decade ago. But it also shrinks perspective. Because information is so accessible and communication instantaneous, there is a diminution of focus on its significance, or even on the definition of what is significant.’

Researchers in the field of ‘collective intelligence’ examine the role that our social behaviours and the arrangement of our social structures play in the ability of human groups to make intelligent decisions. The argument goes that just as humans have an individual intelligence which can be cultivated, enhanced, and used to solve problems in their environment, so too do our societies. While this phenomenon has been described in writings going back as far as those of Saint Augustine of Hippo² and are similarly addressed in Adam Smith’s concept of the division of labour, the most popular modern

formal inquiry into what would come to be known as the study of collective intelligence came from the father of behavioural statistics, an early eugenicist, by the name of Francis Galton. While visiting a livestock fair near his home in England, Galton discovered a popular contest whereby locals were tasked with guessing the weight of an ox on display; whoever won the contest would then receive a monetary prize. While there would at times be an individual who emerged as the clear and accurate winner, Galton found that no matter how near or far the winner might be, the average calculated from aggregating all the guesses of the individuals who entered the contest, 542 kilograms, was remarkably close to the true weight of the ox, 543 kilograms.³ This phenomenon, now known as the ‘wisdom of crowds’, indicates that human collective intelligence works in ways to allow us to solve problems that would be difficult or even impossible for any single human brain to solve.

Nevertheless, this phenomenon of the wisdom of crowds does not apply to all crowds, and some are wiser than others. Critically, the way in which groups are structured, arranged, and the norms which govern their interactions all have noted effects in a group’s ability to solve problems. In models of this phenomenon, one guiding principle seems to unify the ability of a population to be able to solve the complex problems it is presented with, that of *transient diversity*, or the ability of a community to explore multiple theories. Within collective intelligence, this form of diversity has been shown to be preserved or maintained in a number of ways.

PHOTO: PAINTING / ALAMY STOCK PHOTO

THE STRUCTURE OF OUR INTERNET ARCHITECTURE

Although an obsessive focus on diversity initiatives has become a pervasive defining facet of modern Western institutions, recent research has shown diminishing returns on the sort of diversity our institutions are searching for, which is increasingly that of ‘surface-level diversity’ on the level of race, gender, or ethnicity. As it turns out, these facets have little effect on our ability to innovate, whereas deep-level diversity, our differences in norms, values, and external attitudes, is more significant in governing our collective outcomes⁴ (one can consider the fatal assumption here being that many organizations which select for surface-level diversity assume it comes with deep-level diversity). As it turns out, while modern culture and major tech corporations pay a great deal of lip service to diversity, their practices are doing very little to preserve and foster the very thing they are seeking to optimize.

While the internet has brought people from disparate parts of the world (and thus also formerly isolated fringe ideologues) closer together, this has come at some cost. Mass polarization in some countries has been attributed to the internet in some circumstances (although the extent to which this can be specifically attributed to the internet has been hotly debated), and at the very minimum, extremist beliefs and worldviews tend to propagate with rapid speed in social circles online. Even worse though, I would argue, is that because of the internet we tend to agree *more*: polarization is a problem, but this is only so if it extends to every facet of our online lives. A deep, online disagreement between two individuals arguing over which book of a certain author should be considered his magnum opus is very far from the level of us-or-them polarization that we are seeing online; the level at which our disagreements over more serious facets of life descend to, and are packaged together with, our disagreements over the mundane. My argument would be that although it is true that people are engaging in discourse more than ever online, the diversity of things over which they are disagreeing is decreasing

rapidly, meaning that *fewer and fewer* critical or innovative discussions are happening.

Much of this comes as a natural consequence of the fact that as the internet has developed, its structure has become more *centralized*; as our networks become closer, the information we are collectively exposed to becomes less variable, and our disagreements in turn become fewer in kind at the same time as they become greater in number. Generally, when one talks about the centralization of information, images of the circulation of *Pravda* in the Soviet Union, the Xinhua News Agency in China, or the use of ‘official’ fact checkers in the United States are what immediately comes to mind. But such forms of centralization are very different from the types of centralization we are globally experiencing online. In the prior case, centralization is defined as information being disseminated from one core information provider, giving these providers control of the narrative at hand. In the case of the internet, centralization is very different, bringing people who were previously further away in geographic distance but think about similar topics closer and closer together. In turn, they bring with them not only their opinions, but also their topics of interest in the case of Twitter, video preferences in the case of YouTube, movie ratings in the case of Netflix, and product choices in the case of Amazon.

On the micro level, the internet, through its increasing connectivity, is merging and averaging out each person’s unique set of preferences to bring them closer to the average preferences of the people in the networks to which they belong. On a larger scale, their preferences are being averaged to the preferences of people who are not *quite* like them, but similar enough. And on an even larger scale, they are being averaged to the preferences of people who are not like them at all. While people in some countries such as the United States are becoming more divided as a consequence of being online, we are nonetheless more closely connected than ever to people who are actually *not like us at all*. In the language of network science, our average path lengths, defined as the average number

of connections between an individual and the furthest person from them in a network, have decreased. Such hypothetical connections have previously been characterized by games like ‘Six Degrees of Kevin Bacon’, but in recent years the phenomenon has become far more real than we might like to believe.

HUMAN NETWORKS, THE PERISH OF THE PERIPHERY, AND EXPLORE-EXPLOIT DYNAMICS

This leaves us with several important questions. Is it not better that we are all connected? Is this why the internet is not lauded as the revolutionary technology that it is? Should being closer not allow us to develop new technologies and ideas? Superficially, of course, we do experience an increase in the overall *amount* or *quantity* of information in the world, and information certainly gets to us faster. News from around the world often breaks on Twitter long before journalists are able to draft headlines, and new social trends in the form of memes and music come to us faster than lightning.

One of the primary, and certainly one of the most explored, manners in which diversity is leveraged and maintained is through the structural arrangement of our group networks. For most networks, the most efficient network (more technically, the network which can reliably communicate information with fewer steps to the farthest constituent nodes belonging to it than any other) is one in which each individual is connected to one another. Because in these networks there is no ‘middleman’ between information flowing from one node to another node in the network, the network’s information can be considered to be fully reliable. On the other hand, as networks begin to become less connected, the number of nodes through which information must be passed to get to another person increases and the network becomes less efficient at transmitting information.

Unsurprisingly though, an increase in connectivity and efficient information allocation comes with a trade-off. In the world of collective problem-solving and complex systems, we call this the ‘explore-exploit’

trade-off.⁵ This trade-off represents a problem common to all of us in our daily lives: do we continue to explore the range of options that may be available to us in order to find more information to make better decisions, or do we finally stop and do something with (or exploit) the information we have already collected? In network science, this explore vs exploit trade-off has been described in terms of the ability of networks to behave. For example, more connected or fully connected networks tend to be optimized for exploitation of the information they have collected (one can imagine in these cases an efficient factory with no input from the outside world). As drivers of consensus and having little disagreement between units in the system, these networks are rapid: in teams where the task is clear and the information applicable to the task is known, these networks are extremely efficient at accelerating the velocity of problem-solving. On the other hand, less connected and more inefficient networks are by definition better at exploring the entire range of information which may be available (but not yet discovered) to the system; progress, if it happens at all, is slow, but a more accurate picture of the world the network is placed in is built during this long period of time.

At the risk of veering into the mundane, I must emphasize that the internet has made our world more connected; and by this I mean that the world is now made up of a smaller number of more efficient and better connected networks than it was previously. As the internet has evolved, it has developed from a niche information sharing device for hobbyists with arcane interests to the primary means through which essentially all information and most work (or at least that of a non-artisanal nature) is disseminated and assimilated on our planet. This connectivity has allowed us to access any facts or ideas at the click of a button, yet this utility is not without its risks. At risk primarily is the loss of less centralized and group-dominant information at the niche peripheries of our network, capable of the exploration of solutions which the less peripheral, better-connected mass was not privy to. The centralization

of the ‘informational network’ in Mexico following the Spanish Conquest of the Aztec Empire led to the death of certain forms of information, such as Aztec, Mixtec, and Mayan tribal knowledge, that seem to have been efficient at curing local diseases, and producing architectural structures of the kind we would struggle to build today. Considering more recent history, the death of specialized forums as a consequence of the slow but steady absorption of these internet microcultures into the privately owned company Reddit, or specialized blogs killed by Substack and Twitter, or real-world communities being replaced by Facebook groups, all represent consequences of our ever-increasing connectivity.

As fringe information is lost, so too are specialized, non-mainstream topics: you find more and more pundits, scientists, and commentators becoming COVID experts, Afghan policy strategists, and general information podcasters simply because this is where the information demand has led them. In other words, our information is *stagnating* as we are guided to talk about information which is considered appealing to the average internet user. Where clusters of specialized information used to aggregate; where we once had hashtags delimiting our information clusters on Twitter, they are a rarity for posters; Tik-Tok sorts videos by preference, not by novelty; and in the space where movie recommendations for VHS tapes came from clusters of people we knew in real life or had to seek out, Netflix now aligns our interests to other people we will never meet who are algorithmically grouped alongside us in datasets that will never be publicly available (and yet with limitless options on Netflix or YouTube, everything is so familiar that paradoxically nothing seems to be available). Eccentric expert discussion, novelty, and niche hobbies are increasingly unobtainable not just for the average person, but for the very networks which were actually *established* as homes for such self-aggregated clusters. Any person cognizant enough to realize their position as merely another undistinguishable node within our current global techno-

capitalistic network is forced to accept that they are what they have always been on the global cluster: an average.

ALGORITHMIC AUGMENTATION

Over the past several years, we have begun to see signs that our collective networks online have not only begun to circulate information faster, but also that the rate by which information circulation accelerates within these networks is accelerating too. Note that every meme, YouTube trend, TikTok recipe, prosocial game (think of Wordle, Animal Crossing, and Pokemon Go), and geopolitical incident arises, is exploited, and is burnt out or ‘run through’ almost instantaneously. An incident as terrible as the Itaewon Halloween Crush in South Korea, in which over 150 people died under horrific circumstances that were broadcast on Twitch livestreams and YouTube for the world to see, is forgotten in less than a week. The speed at which popular memes originating from obscure online forums are exploited and passed around, passing as they do so through a seemingly infinite number of iterations, is staggering. This is the highly connected network exploiting what little novelty it has been provided with. That is no accident—a highly efficient network *should* and indeed *must* circulate information with such alacrity that any surprising joke or meme we find funny will repeat itself so fast and so frequently that we become bored within mere minutes of first seeing it.

For those who may be aware of this rapid process of content stagnation by repetition, you will relate to the opinion that interfacing with the internet was not always this way. In 2008, Facebook reported that the average distance between two users, as measured by the people between them, was 5.28 degrees or people; by 2011, it had decreased to 4.74; and to a meagre 3.57 people between you and any one of Facebook’s 1.59 billion users in 2016.⁹ Given that Facebook has only reported its first drop in overall users for the first time in 2022, we can assume that since 2016, such numbers have only decreased. While the speed of information has increased due

to the connectivity of our social networks, the recommendation algorithms which tech companies employ play a major role, too. In the framework of explore-exploit dynamics, we consider social networks as *computational* in their ability to find, exploit, and transfer information between constituent nodes. Similar to visions of Elon Musk’s Neuralink augmenting the human brain with artificial intelligence, in this task of exploration and exploitation, these networks are augmented by the artificial intelligence provided by recommendation algorithms. As an artificial intelligence, the role of the recommendation algorithm is to more efficiently bring information to users which matches their preference, making predictions so they will not have to choose, and keeping them satisfied with the information presented to them.

‘We are nonetheless more closely connected than ever to people who are actually not like us at all’

The algorithms work with our networks to exploit information in at least two ways: first, they ‘recommend’ topics, movies, and products to individual users by creating profiles which treat each individual as an aggregate of individuals *like* them given the information available to the algorithm (currently, an ever-increasing amount). This allows information to more efficiently be allocated to individuals that may be likely to consume it. Second, in the case of social media, they affect exploitation in the network by altering the network’s properties. In most cases this operates by muting outlier opinions and by boosting ones which are popular by definition (notice the disappearance of the archetypal ‘snob’ in nerd sub-cultures, and his replacement with more agreeable Reddit boards), making products which do not yet have extensive reviews difficult to find, and even hiding people who are not close to an individual’s level of attraction on dating apps

to keep less desirable partners away from you and you away from the more beautiful. In the most extreme cases, Facebook and Twitter recommend new followers and ‘friends’ to you whom you would otherwise never have had contact with either in real life or in cyber space (because more people like you, on average, want to be friends with them).

While these algorithms help our networks become more efficient in exploiting the information we are given, we must remember that, by design, they are fated to do poorly with regards to novelty. Imagine a version of your favourite book which showed you only the paragraphs which other readers algorithmically grouped together with you wanted to read. Or, even more strikingly, imagine a pair of glasses which only showed you what you wanted to see in the world around you—perhaps rendering an ‘unpleasant ocular intrusion’ such as the presence of a homeless man sleeping on a park bench into a delightful bouquet of flowers, all for your visual delight. By collecting personal data from every human action or interaction ever conducted online (a feat of informational mastery entirely without precedent in the history of the universe), and by building algorithms that interpolate our preferences and desires upon the basis of this data, averaging out from the sum total of all variables to calculate what for human beings is the ‘most okay’ or ‘least dislikeable’ content, tech monopolies such as Google and Amazon have set us on a path where the diversity of information is likely doomed. In another sense: you are constrained by being shown those things you have not seen but that the ‘average you’ wants to see, not by what the real you cannot yet see. Rather than being exposed to a truly diverse range of informational outputs in response to the textual inputs we type into a search engine, some of which (by definition) will be bad, while others are better, it seems more probable that the information available to us as individuals, and indeed to the collectives to which we belong, will *shrink*.

While it feels as if these algorithmic systems have always been employed on the internet (and for many younger people, this is true),

their use is a relatively recent addition to the online ecosystem. Twitter simply presented all incoming tweets from those you followed in reverse chronological order, only switching to recommended content in 2015. One of Facebook's most valuable assets, Instagram, only added a recommendation algorithm in 2016. In fact, most of the internet's modern recommendation architecture was only constructed in the last five or six years. A cursory examination of tech giant history reveals that during one short three-year period in particular, companies which did not have these systems began to adopt them, and companies which already possessed them began to dramatically revamp them (frequently employing more advanced machine learning techniques in their implementation):

- Amazon: 2003 (the first), major changes in 2014;
- Google: introduction of RankBrain in 2015;
- Twitter: started in 2015;
- Tinder: started in 2012, revamped late 2015;
- YouTube: 2008, changes in 2016 (DeepMind);
- Instagram: started in 2016;
- Facebook: 2009 like button, major changes in 2013, 2017;
- Netflix: 2006, major changes in 2017.

It is also not as if these algorithms cannot be altered to assist in the exploitation of information, either, but in their current iterations, they are not designed to do so, especially on a large scale. The way our algorithms currently work is not by exploiting novelty, but by exploiting the familiarity of past preferences. They work by trimming aggregated content to averages and squashing preferences at the edges of the distribution of content, and in the corners of networks. These network models, broadly defined, are a unique case of applied AI—artificial intelligence not in the sense of a sentient mechanical being with life and consciousness, but in the sense used by software engineers and computer scientists of a machine that intelligently learns from what we teach it.

We often hear thought experiments about the ostensibly disastrous consequences that AI will have on the human psyche as if artificial

intelligence itself is not already here. Not only is AI here, but it is also learning from our *collective* cognition, and as individuals are targeted by the search process of our algorithms and big data systems, there are unavoidable side effects on our own cognition. Again, the speed with which novelty is destroyed, with which nuance disappears, and with which we find consensus (whether for good or bad) is rapidly increasing. The first exploitation of human cognition by artificial intelligence is not Neuralink, as some futurists suggest, but rather the squashing of the collective discovery and intellectual autonomy of individual human beings by algorithms which are blind to innovation at the edges.

DIVERSITY, AMONG ALL THINGS

In 1869, the Transcontinental Railroad was completed in the United States, unifying for the first time the West and East coasts of North America with a golden rail spike to commemorate the event. While most of the nation celebrated, sceptics in small areas around the country expressed fear at the idea that something would be lost in the culture they had developed through the importation of ideas from urban areas around the coast. In the modern era, the disappearance of local and regional dialects in the United States has spoken to the power of the radio and television set to unify a nation in the perceptual domain. While one would argue that the loss of these cultural idiosyncrasies is on some level irrelevant in light of the immense progress experienced over the last two hundred years, recent events have shown us that we have genuine reasons to express the same fear that those Americans experienced 150 years ago. I suspect that, as we continue down this path, as people are brought to the same level of fundamental agreement around the world due to algorithmic simplicity and hyperconnectedness inaugurated by network optimization, as the local becomes global, we will cease to find alternative solutions and schemata for dealing with a world of fundamental uncertainty that is always changing even if our algorithms do not.

Some networks have been shown to be swift in both locating optimal solutions and in preserving the diversity which allows for optimal solutions to be found. These tend to operate in one of two ways: allowing for individual groups to be highly connected to themselves and loosely connected to one another, or instead by allowing a highly-consolidated and connected core to exploit information without totalizing the spread of this information to the periphery. In both of these cases, knowledge diversity is preserved, and the group as a whole is able to act on this diverse information without driving a global consensus. In other words, these networks attempt to actively balance the explore-exploit trade-off. In the first case of 'cliqueish' networks, independent groups are in close agreement with one another—one could envision a world of not just one or two, but infinite and exclusive 'filter bubbles' which would only loosely work with one another. In the second case of core-periphery networks, we could imagine a highly connected internet whose architecture is opt-in; loose interactions prevail between unrelated or distal subcommunities. In either hypothetical world, exclusivity allows for the preservation of diversity by extreme means. A more serious solution would be the slowing of the centralization of the internet through platforms like Google and Facebook, which could potentially be facilitated by actively promoting alternatives and actively curtailing the use of recommendation algorithms for fast and heavy-handed categorization of the internet's users.

At this point, any solution is critical, as the problem of internet centralization has long been ignored. At the same time, we may have moved too fast to hope for a serious conversation about the speed with which we are currently moving—an ironic situation to find ourselves in. While AI safety researchers continue to debate the dangers of artificial general intelligence (AGI) systems, the relevant conversation about the dangers of collective intelligence systems is happening (or will happen) too late.

As noted by the economist Friedrich A. Hayek,⁷ *'If we can agree that the economic problem of society is mainly one of rapid adaptation to changes in the particular circumstances of time and place, it would seem to follow that the ultimate decisions must be left to the people who are familiar with these circumstances, who know directly of the relevant changes and of the resources immediately available to meet them. We cannot expect that this problem will be solved by first communicating all this knowledge to a central board ...'*

While Hayek's point was made in the context of central planning and the inefficiencies of that system, it is fundamentally the case that in any process whereby knowledge flows from the centre outward to the periphery, rather than vice versa, we are resigned to the gradual and all-encompassing inhibition of information and knowledge (both in breadth and in diversity) available to the system we inhabit and depend upon.

NOTES

¹ Henry Kissinger, *World Order* (Penguin Books, 2015).

² Augustine notes the following in Book 7 of *The City of God*, '... one vessel, in order that it may go out perfect, passes through the hands of many, when it might have been finished by one perfect workman. But the only reason why the combined skill of many workmen was thought necessary, was, that it is better that each part of an art should be learned by a special workman, which can be done speedily and easily, than that they should all be compelled to be perfect in one art throughout all its parts, which they could only attain slowly and with difficulty.'

³ Francis Galton, 'Vox Populi (The Wisdom of Crowds)', *Nature*, 75/7 (1907), 450–451.

⁴ J. Wang, G. H. L. Cheng, T. Chen, and K. Leung, 'Team Creativity/Innovation in Culturally Diverse

Teams: A Meta-analysis', *Journal of Organizational Behavior*, 40/6 (2019), 693–708.

⁵ An early exploration of this phenomenon can be found in Schumpeter's *Theory of Economic Development*, but the term arises from March 1991: 'Exploration and Exploitation in Organizational Learning', *Organization Science*, 2 (March 1991), 71–87.

⁶ S. Edunov, C. Diuk, I. O. Filiz, S. Bhagat, and M. Burke, 'Three and a Half Degrees of Separation', *Research at Facebook* (2016), 694.

⁷ Friedrich Hayek, 'The Use of Knowledge in Society', *The American Economic Review*, 35/4 (September 1945), 519–530.