Eng Int Syst (2024) 4: 385–395

Engineering
Intelligent Systems

Collective Intelligence and the Creation of Knowledge Transfer

Christian Wagner

City University of Hong Kong, Kowloon, Hong Kong

The keynote explains the enabling role of IT in unleashing collective intelligence, which thereby created a promise of widespread access to knowledge. It then points out, that the capture collective intelligence leads to attempts to algorithmize this knowledge most recently through powerful artificial intelligence models. These models disintermediate the former knowledge providers from their competitive knowledge advantage. The keynote predicts that the threats of disintermediation will ultimately bring us to the recognition of artisan intelligence that cannot be captured easily by artificial intelligence models, and a knowledge sharing model where people will share knowledge in return for micro-payments.

Keywords: Generative artificial intelligence, collective intelligence, knowledge, creativity, wiki

1. INTRODUCTION

1.1 Collective Intelligence

The world has always relied on collective intelligence. It is what we do all the time and is well recognized by researchers. For example, as individuals, we have employed collective intelligence in scenarios as a simple as joining the queue when people are lining up in front of a store. After all, there must be something good inside, right? However, traditionally collective intelligence has been rather local, because it was frequently propagated through Word of Mouth. So, when in 1993 the World Wide Web appeared, this suddenly changed everything, as we now had communication and collaboration costs really going down so dramatically and suddenly that we could spread our intelligence throughout the world, thereby bringing our fragmented knowledge together. This realization then created a vision of the democratization of knowledge where everybody could partake, thus emphasizing the importance of collective intelligence even more.

This keynote will thus describe collective intelligence and exemplify how it actually works, that is, the underlying mechanisms of knowledge aggregation and curation that distill fragmented bits of knowledge into meaningful aggregates. This will be illustrated with some examples. The examples will also identify that there are multiple collective intelligences and that there is a form of collective intelligence that may function quite differently, namely collective creativity. From there, the presentation will point out that it's an almost inevitable path from human collective intelligence, towards computerization. And that's where we move from collective to machine learning, to artificial intelligence.

1.2 Artificial Intelligence

The move to artificial intelligence then seemingly takes all the advantage we have or we have had away from us, as it captures the collective knowledge and draws inferences from it better than many people can. This, of course, makes our collective intelligence obsolete, which is a dramatic result. We then must ask ourselves whether the recent advances in generative AI (G/AI), lead to our knowledge obsolescence. And right up front, the simple answer to this is "No", but knowledge specialization will matter much more in the future than it does now and so is the prioritization of our knowledge. If we shared freely before, we may have to more selectively do so in future, guarding what only we know, so that it cannot be easily commoditized. With this as an overview, let us now move to the paper's main argument, going back timewise to the 1990s, where widespread knowledge sharing emerged and

Figure 1 AI Rendering (Midjourney).

evoked feelings of a beginning of knowledge democratization. One of the milestones in thought leadership around this theme was the book *Cluetrain Manifesto* (1999).

1.3 Knowledge Democratization—Promise and Reality

This famous work outlined the logic of free access to knowledge and its implications. The famous Fathers of the Internet and Worldwide Web outlined 95 theses of how this new world operates. The Cluetrain Manifesto had a strong focus on consumer versus industry, but its theses also hold for consumer-to-government interactions.

One of the most profound among the theses is that *markets are conversations*, truly a novel idea at the time, and possibly even now. Markets are traditionally seen as clearing mechanisms between the supply and demand for goods or services. Now markets are conversations. The explanation goes as follows "the internet is enabling conversations among human beings that were simply not possible in an area of mass media". To note, around this time, in the early 1990s, Intel had brought out the Pentium processor chip, whose coprocessors mis-calculated. And this was unbelievable, right? Surprisingly, for some people reported that in the bulletin boards. And so that was the first time when such messages which were posted on the bulletin boards, noteworthy by one Professor Nicely. Initially Intel said, no, no, no, it happens once every 9 billion calculations, or such. And then people were posting their findings on bulletin boards, and Intel had to recall the chip. And so the Internet came to demonstrate the knowledge and thus the power of the many.

Networked conversations are enabling powerful new forms of social organization and knowledge exchange, because the communication and collaboration costs are dramatically driven down. And hence the Cluetrain Manifesto also states that "people in networked markets have figured out they get far better information and support from one another then from vendors", a fact which all of us implicitly acknowledge these days as when we watch YouTube or similar videos instead of reading the manufacturer's information to learn about a product.

"There are no secrets. The networked market knows more than companies do about their own products". This was the expectation or maybe just a hope, that the Internet would be the great knowledge equalizer, thus empowering consumers vis-à-vis companies. The vision was a bit blue-eyed, of course, because the Internet is a tool, after all. Consumers were possibly quicker in learning how to use the new tool called the Internet, companies were figuring it out as well—and they have a "bigger hammer".

And so we move fast forward 2023. With the recent emergence of G/AI, content creation, real and "fake" is becoming ever more possible. Figure 1, for instance, shows a fabricated likeness of President Biden and Kim Jong Un sitting at a table on a beach at a sunset, created by instructing the Midjourney software to create a "realistic photo of President Joe Biden and Kim Jong Un sitting together on a beach near sunset drinking Matais". 60 seconds later, the software renders the image, filling in additional meaningful elements, such as a table and the chairs. This capability is both wonderful and scary, raising the question whether we are now better off than before or whether the move from capturing collective intelligence and then encoding it into AI disintermediates us from our knowledge and thus lowers our value in the marketplace.

People are also realizing that technology is not necessarily the bringer of truth sharing and that suddenly there appears to be no single truth but instead people holding on to "their truth" reinforced in like-minded echo chambers. Let us try to find an answer to these questions in a short thought journey from collective to artificial and from there to what I call *artisan intelligence*.

2. COLLECTIVE INTELLIGENCE—A DETAILED VIEW

To illustrate the promise of collective intelligence let us look at an example from Bangladesh. Bangladesh has traditionally been a less developed country and had significant food shortages, as well as and difficulties with its harvesting. Sugar cane is one of their main crops. A lot of subsistence farmers must bring their sugar cane to the sugar mills, where it is squashed and squeezed and then the sugar sap is extracted. Farmers cannot go to the mill anytime, but only once they have received a time slot from the mill. This is not a trivial issue, because how would the subsistence farmers know when

Figure 2 e-Purjee.

to bring their cane to the sugar mill to get this job done. Improving this process was one of the projects that the United Nations looked at in terms of country development. As you will know, the United Nations undertakes a lot of development projects, sometimes successful and sometimes less so. The UN Development Agency sometimes may come in with, you might say, overwhelming force, bringing resources and a lot of know-how from the outside. But once the project ends and the UN leaves, nothing much is left behind and the project dies for lack of sustainability. We can read this in the UN's own reporting—many projects are not sustainable, as sustainability is oftentimes not considered in detail as part of the selection and design process. And so this is a typical thing that projects oftentimes either they are too small and not worth well or they there must be a way found to make them. Project continuity after the external funding goes away then depends on making the project fit the local environment, which can be difficult. The experts come from afar may know how to solve problems when properly resourced, but without resourcing, a different solution may need to be found, one that is embedded in the local context. How do we know about the local context? By capturing local expertise.

One such example that was developed in Bangladesh that is the e-Purjee (e-ticket) system. As mentioned, sugarcane farmers can only bring their sugar cane to the mill when specifically called to do so. They need to get a ticket, which traditionally was a paper ticket, delivered to them by delivery persons traveling across Bangladesh. The process was often troubled. Delivery people were late, or brought the tickets to the wrong recipient, or dropped tickets at some local dispensary where the farmer never received it. This created a lot of uncertainty for farmers, because they had to manage their harvest, and it often took a few days to cut the sugar cane and deliver it to the moil. So, a delayed ticket meant missing the delivery window. Being at the mill too early meant having to wait and losing sugar cane weight, in the strong outside heat.

The unreliability of the delivery system affected smaller farmers more than larger ones. The larger farmers had ways of finding out their time slot through other channels or could convince the mill to accept their sugar cane even when delivered at the wrong time. So, when the development experts looked for a solution, they sought local insights and created a localized system based on text messaging. Farmers would be informed about their delivery time slots by SMS (e-tickets), receivable even on the least expensive phones on minimal phone plans—not on the worldwide web. Even the least well-off farmers would have access to such phones and a limited data plan to receive the messages.

Converting to SMS-based tickets created an inexpensive, reliable ticket delivery system. With it came a dramatic reduction in lost, misplaced, or misdirected tickets, which increased economic benefits to the subsistence farmers and lowered the economic uncertainty. Interestingly, when we asked the subsistence farmers about the main benefits of the new system, they did not immediately mention the economic advantages, but the improved transparency and fairness in the system—earlier we referred to this as knowledge democratization. Reduction of decision uncertainty was also a mentioned benefit. Famers would make comments such as "not every day I have to worry whether I have to go to the mill now", or "previously I didn't know if someone stole my paper ticket". Transparency created procedural and distributive fairness in that farmers now had equal opportunities to deliver their harvest to the mill. This was made possible through a system that relied on technology appropriate for the environment, created based on the input of local informants.

As such, this is a useful example of how collective intelligence works to create appropriate innovation for the specific regional context. Simple and inexpensive, and delivering not only economic benefits but also procedural and distributive fairness.

3. DEFINING COLLECTIVE INTELLIGENCE

Collective intelligence is a form of intelligence that emerges from the collaboration and cooperation (and sometimes even competition) of many. Collective intelligence appears in a wide variety of forms of consensus decision-making not just among human beings, but also various other forms of life all the way down to bacteria, and computers as well. Collective intelligence arises from, and this is wherein its essence lies, aggregation of fragmented knowledge of multiple informants. These informants must have independent sources of information and their insights must be shaped by individual information acquisition and transformation.

Let us illustrate this a bit more, with an example for everyone who is in a university teaching capacity, especially at graduate level. Assume you find yourself in a classroom with students who are accounting managers or IT specialists, or whatever your audience may be. Importantly, many of your students will possess expertise in their own right. As an academic with mostly "book knowledge", you may wonder "what can I teach these people? They know a lot more than I do." In other words, even as an expert professor, your

Figure 3 NFL Pick'em Results.

classroom, collectively, challenges your expertise. And yet we know that the professor still has a role to play, namely in bringing the knowledge of the class together, extracting it from students and curating it to let the entire classroom know more than any individual. This is one of the important aspects of collective decision making, namely the role of the coordinator. There are people who bring in the content or raw knowledge, and there are others who aggregate that knowledge, make it more understandable and to abstract it into principles. This is one of the important roles of professors—they are less content creators and more knowledge shapers in the graduate student classroom. Knowledge shaping is a higher-level skill, usually requiring multi-domain knowledge, high-level abstraction ability, and the ability to recognize patterns. Therefore, knowledge shapers are less prevalent than knowledge providers, which is fine as one shaper can manage the contributions of many knowledge providers.

3.1 Collective Intelligence Exemplified

We do not always need highly skilled knowledge shapers, sometimes even simple aggregation of the fragmented knowledge of many can provide exceptional results. Let us consider for instance an example from *Yahoo!Sports's NFL Pick'em* challenge created by *Yahoo!* about a decade ago. In this Pick'em challenge, regular participants were allowed to pick the results of major US Football games throughout the season. Their guesses of who would win or lose were aggregated, shown to several experts, and then the experts were asked for their guesses about the game outcomes. Experts could agree with collective guesses or deviate. The experts were journalists, former coaches and such, people who should really know. The situation was clearly biased in favor of the experts, since the experts would know the collective's guesses before they had to provide their own.

Surprisingly then, year after year, the collective would outperform most, if not all experts, as illustrated for instance in Figure 3.

In Figure 3 we see that the collective made 165 correct guesses and xx incorrect guesses, with none of the experts being able to match this winning tally. The best expert had 162 correct guesses, the worst 157. Figure 3 shows some old data, as Yahoo!Sports is not running this challenge anymore. In fact, while the experts have the advantage of full information, if they simply picked identical to the crowd, they would seemingly not be adding any expertise. So, the experts must go against the crowd at least on some games, but it turns out, this is to the experts' disadvantage. So don't go against the crowd.

The question then arises, how can the crowd know more than these sports experts? And the answer to the question stems from what I mentioned before, namely fragmented knowledge. If you are familiar with your local region, you may know some special knowledge, e.g., if it just rained an hour ago in your city, or whether the coaching staff had some turmoil. This local information may shift the outcome away from objective odds to situational odds, which the experts would not be aware of. And that situational contextualization could then create an advantage for the crowd. Oftentimes such information is negative knowledge, where the crowd does not know the specific outcome, but may hold strong beliefs about which outcome will not occur, i.e., "Seahawks cannot win in San Francisco", or "California teams cannot win in Philadelphia during the Winter."

3.2 Knowing What Is Not

A very powerful example of collective intelligence is the Millionaire Game. This game has been played in over 200 countries and regions around the world and may to this day still be played in about 180 markets. The dame premise is that an expert must answer a sequence of questions with the stakes doubling roughly every round. As the expert wins, the stakes either double, or the expert can cash out. If the expert loses, all money is lost. If the expert is undecided, one of his or her life lines is to "ask the audience". Interesting though, the questions are frequently not so easy. So, if you're the expert and you don't know, why should the audience know? Scarily though, the audience frequently knows, maybe not perfectly, but good enough to guide the expert. Figure 4 shows an example, with the expert having to identify which of the four listed persons was NOT a member of the original Three Tenors. The options are: Luciano Pavarotti, Andrea Bocelli, Jose Carreras, and Placido Domingo, ABCD.

Much like the expert, audience members may not know the right answer, but may possess enough fragmented (small detail) knowledge to exclude one or two choices. Let's assume, for simplicity, each audience member can eliminate two choices correctly, and the audience consists of only 6 individuals, as illustrated in Figure 5.

Out of our 6-person audience, Agnes and Alfred have similar knowledge. Both eliminate C and D as choices, but believe the right answer is either A or B. Not being certain whether A or B, they might "flip a coin" leading Agnes to vote

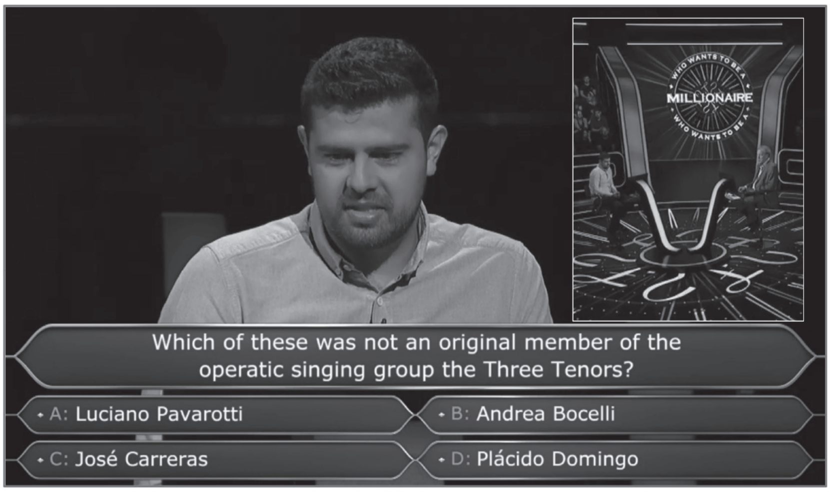

Figure 4 Millionaire Game Scenario.

Audience Member	A (L. avarotti)	B (Andrea ocelli)	C (Jose arreras)	D (Placido omingo)
Andrea	×	✓	×	×
Alfred	✓	×	×	×
Bob	×	✓	×	×
Betty	×	×	✓	×
Celia	×	✓	×	×
Chris	×	×	×	✓
Total ✓	1	3	1	1
Total ✓ or ×	5	3	5	5
Percentage	1/6 = 16.67%	3/6 = 50%	1/6 = 16.67%	1/6 = 16.67%

Figure 5 Ask the Audience.

for A and Alfred to vote for B. Bob and Bertie, same thing. They would rule out A and C but be undecided between B and D. Bob would then vote B and Bertie for D. Celina and Chris, following the same pattern, would eliminate A and D, and choose (flip) between B and C. Tallying up their votes, We see that most tenors only receive one vote each, but B would amass 3 votes. In statistical terms, Pavarotti receives 16% votes. But Bocelli, the B answer, captures 50% of the votes, with Carrera and Domingo also claiming 16% of votes (1-of-6). In the end, it looks like 50% of the crowd knew the right answer, but the crowd never really So It's not that 50% of the crowd knew the right answer, but 100% of the crowd knew the half-right answer. Statistically these two are equivalent, and practically the outcome leads the expert to the right answer, namely Option B. In the real game, Option B only obtained 40% of votes, but statistically this would have been already highly significant!

It turns out that following the crowd is the superior strategy. Evidence from actual game results show that 92% of the time the crowd is "right", or as mentioned earlier, that the crowd can separate the most likely choice from less likely ones. Furthermore, the crowd does not have to be able to eliminate 50% of all options. Even if each crowd member can only eliminate one option, the best result will reveal itself, given enough (independent) audience members. Interestingly also, no facilitator or knowledge shaper is needed in this scenario, simply tallying up the votes leads to the best answer.

The world, however, frequently poses more difficult problem-solving scenarios than the one described, scenarios which require some planning and a process. is typically more complex decision making so oftentimes we need a process to facilitate that right. Individual problem solvers may possess micro-expertise, but do not understand the "big picture" required to solve the problem. To exemplify, assume you enter your young child's room and find the scenario shown in Figure 6. You immediately instruct your child to "make up your room". Two hours later you return, only to find the room unchanged. Now envision instead giving your child your micro-instructions such as "put all the teddies in one corner, remove all the plates from the bed, hang all the clothes in the wardrobe" and so on. After you return and see your child having executed those instructions you say, "from now on we will call all these activities together 'make up your room'". In other words, there is now a program (make up your room) with the micro-instructions as described.

If we can do this to instruct a child to make up a room, we can do this also with collective intelligence.

3.3 Why Wikis Work

A good example is the world of wikis and the logic of why wikis work. Wikis are read-write webs, enabling anyone to edit the contents of the shared knowledge. Wikis have had

Figure 6 "Make Up Your Room"—The Need for Micro-instructions.

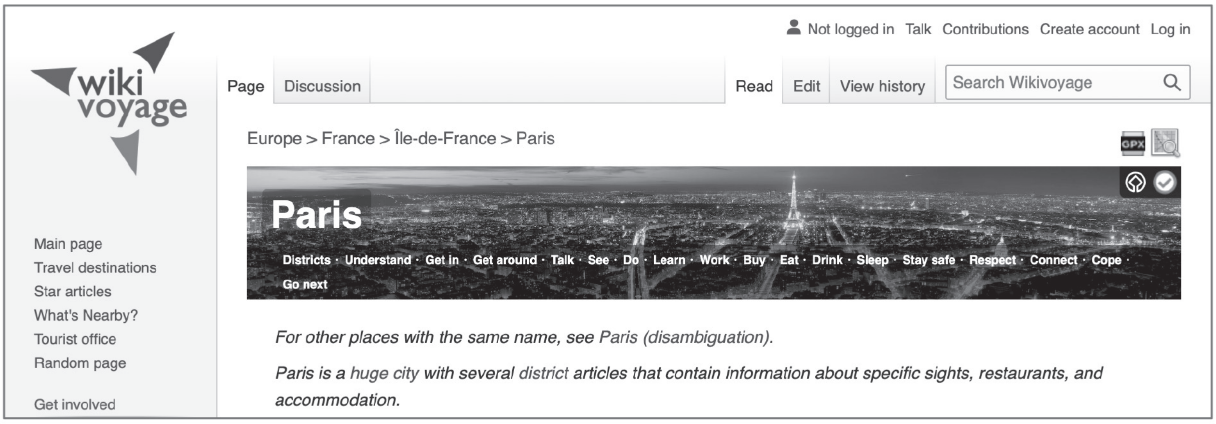

Figure 7 Micro-instructions to Add Content to Wikivoyage.

a profound impact on knowledge sharing, illustrated most impressively by Wikipedia, the world's largest encyclopedia. For our discussion here, let us look at a different wiki, namely Wikivoyage, which is a very large wiki itself with about 170k pages, focused on the domain of travel and tourism. On Wikivoyage we can look up most cities in the world, or even parts of cities. Or we can add knowledge as well, as Wikitravel also has an edit button on every page.

Let's assume we access Wikivoyage and find our hometown missing. What an opportunity to contribute, instead of just consuming the contribution of others. But where does one start? The challenge of writing a meaningful wiki page about one's hometown may be just as daunting as to make up one's room, without proper micro-instructions. Or worse, if I don't have significant knowledge for the entire city, but only have fragmented knowledge, can I contribute at all? Maybe I just know a few good places to eat or drink in Paris, can I add those and thereby contribute? Again, clearly not without micro-instructions. Interestingly though, Wikivoyage comes with pre-built micro-instructions in the form of content headings, which read like "see—do—eat—drink—sleep" and so on (Figure 7). Under each of the sub-headings, I can now add any fractional knowledge I might possess about the topic and thus make a small, but value-adding contribution. With the structure, content can be added incrementally, with the page remaining well-organized. And even if the pages get messy, other individuals may jump in and re-edit the content, just like the professor who organizes the contributions of students in the classroom. These knowledge shapers usually do not bother about adding knowledge, but instead focus on making the existing content more readable and better structured to enable further additions. With that, wikis really enable a functioning symbiosis of people who jointly create knowledge constructs. It's quite remarkable and sometimes referred to as wiki magic.

With wikis we now see collective knowledge construction in a different way, quite different from the earlier voting example. But in wikis there is also implicit voting, because every reader can edit existing content. So, if I read content and do not change it, I gave an implicit vote for the validity of that content. This is important, because it is quite possible that users provide incorrect content, either accidentally or willingly. Then others can jump in and make corrections. This makes wikis largely troll-immune and reliable in their content. And for those who wonder about how collective wikis are, we know that Wikipedia for instance, depends largely on the contributions of its 5% elite contributors, who provide

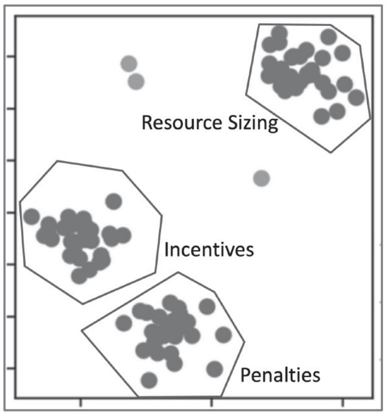

Figure 8 Clustering of Responses to a Creative Problem Solving Task.

over 70% of the effort. There are hundreds of thousands of Wikipedia contributors, most of them with less than five contributions to their username.

But wikis are not the only model of peer production in the spirit of collective intelligence. In addition to the wiki model, there is for instance the YouTube model which also brings collective knowledge together. Individuals can make contributions in forms of short videos, which is curated by user views and a provision algorithm that emphasizes popular videos and de-emphasizes less popular ones. It works through viewer statistics and server algorithms. A viewer who wants to learn about a new topic, whether hair styling, personal finance or automobile DIY, can search and will be provided with the most "authoritative results" first, as demonstrated by views, likes, or subscriptions.

Another interesting collective effort is the TikTok model. TikTok may appear to be the same as YouTube, just with shorter videos, but it is not. Just like YouTube, TikTok collects micro-contributions, but with a much cleverer delivery algorithm. First, videos are much shorter, so their narrative is more concise. Next, the viewer finds out quickly whether the video is useful, thus wastes little time on unwanted content. Next, the interface is swipe-right, to easily discard a video and receive a new system suggestion. A very tight algorithm feeds only videos aligned with prior expressed preferences, thus creating a high utility experience. This is the beauty of the TikTok—delivery according to user preference.

3.4 Collective Creativity

Collective creativity is frequently discussed together with collective intelligence, but it is a different set of principles at work than in collective intelligence knowledge creation. Whereas in collective intelligence we frequently aggregate and sum up the contributions of many to build *consensus* knowledge, the opposite is true for collective creativity. After all, creativity focuses on *uniqueness*, not consensus. We seek one-off solutions—so how can collectives do that? In the realm of creativity, collectives must operate a little bit differently. They must be looking at the margins, the tail-ends of distributions, or responses that do not fall into recognized clusters.

For example, we asked a group of experimental subjects to generate ideas on how airlines can overcome the airline no-show problem, where customers make flight reservations but then do not show up for the flight. The typical ways for overcoming this problem are well known, such giving people a better price if they buy a non-refundable ticket or offering other benefits such as miles or upgrades for showing up. Many respondents will offer answers that describe incentives. Others may suggest varying the plane size, using smaller or larger planes to accommodate the actual flyer number. In the end, many answers will fall into a few clusters, as depictured in Figure 8, with only a few standing alone as unique and thus potentially more creative answers.

The collective, based on its size, has a better chance to unearth such outliers, and can also abstract from all responses to understand their underlying principles and then push these principles to find more creative answers. A knowledge shaper within the collective can ask the group to think of more extreme form of incentives (e.g., free travel to the airport), or more plausible forms of plane resizing (e.g., swapping heavy cargo for empty seats). Overall, collectives do play a role in creative problem solving, but not through consensus, but through the leveraging of the collective's idea diversity.

4. FROM COLLECTIVE TO ARTIFICIAL INTELLIGENCE

Having illustrated several principles that explain why collective intelligence works, let us now explore why collective intelligence almost inevitably leads to artificial intelligence. A good illustration is the credit approval process. So let us dial back a few decades to the early days of credit approval, where approvals were processed manually by human appraisers. You would approach your bank or credit card company with a multi-page application asking you how much money you earned, how long you had lived at your current address, how many other cards you had, and so on. The loan officer then would look at the information and make a holistic judgment on whether you should receive that credit cards. This was much more a bespoke process than today where we receive invitations for credit cards every day and the whole process is

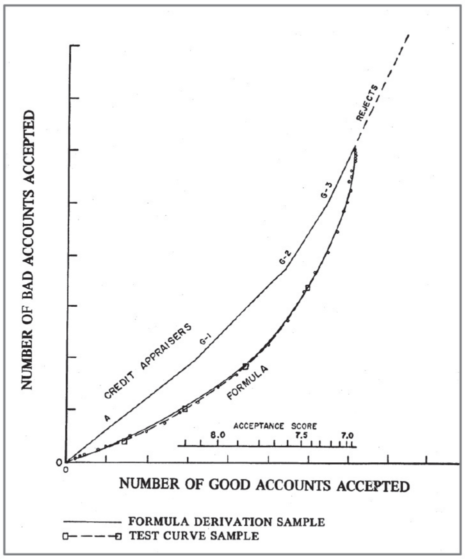

Figure 9 Loan Appraisers vs. Algorithm Performance.

automated, and database driven. This is almost inevitable, as organizations will look for ways to create efficiencies, process transparency, and result replicability.

In the traditional, manual process, loan officers based their judgment on holistic judgment, considering quantitative and qualitative data in an often-ill-defined aggregation process. The decisions, of course, had their base in logic and reasoning, considering tangible criteria to come to a decision. So, the loan officer would usually judge based on factors such as one's ability to pay back loans and the willingness to do so. But interestingly, they may not have had a specific algorithm, or at least not one they were cognitively aware of.

The question then arises, whether we can we build a system that bootstraps that decision making, even if it appears case-based instead of algorithmic. We have a process where with known input data and known decisions. This would suggest some form of statistical method, such as regression analysis would reveal a consistent pattern if there were one. The human analysts might have stated that they decide holistically and that there are no specific rules. Yet, if we treat this as a numerical analysis problem, it may turn out that there is an algorithm that can described the decision logic of these experts. Even more interesting is that once we extract the model, and then use that model to make decisions instead of the human decision makers, the model performs better than the experts from which it was extracted. Figure 9 depicts this phenomenon, comparing the performance of loan appraisers against a model derived from their decisions. First, the experts do perform better than random decision making, which would be represented by a diagonal matching good decisions and decision failures. A function that is bent more to the right and down, as we see here, identifies a higher ratio of good decisions vs. bad ones. Yet here we also see the algorithm outperforming the experts, with a curve that is bent even more down and right.

The process of extracting an algorithm from past data and then using the algorithm in lieu is referred to as bootstrapping. It is the essence of data mining efforts. But why would the bootstrap be better than the experts? It turns out that when experts make exceptions in their decision making, these exceptions do not pay off. Sticking to the rule may result in Type II errors (not giving credit where it is due) but avoids Type I errors of not giving undue credit, thus avoiding deadbeat loans or similar negative outcomes.

Here we see the essence of the transition from collective to artificial intelligence, where a bootstrapping algorithm replaces the human decision making it is extracted from, with better results. While linear regression is not artificial intelligence, but it is certainly a technique used to enable machine learning. A more widely used G/AI mechanism is the use of artificial neural networks, which serve as concept categorizers and as predictor models.

An artificial neuron is based on the logic of a human neuron and can be seen as an organic on/off switch. If the input data is negative, the output will be NO (0 or near 0), if the input data is positive, the output will be YES (1 or near 1), and if the input data is mixed, the output will be undecided by either leaning towards NO, or towards YES, following an s-shaped (sigmoidal) function. This is logic decision-making that can deal with uncertainty better than traditional prepositional logic models.

Neural networks don't consist of just one neuron. In fact, G/AI models, or specifically large language models such as ChatGPT, are very large neural networks, with billions of nodes. The larger the model, the more differentiations can be made, and the better the model can predict. But at the same time, the whole system and its training process becomes significantly more complex. As Figure 10 depicts, there are now several of these very large neural networks, and one of them, of course, GPT. We also see that the complexity of these models is exponentially growing, GPT-3 now including 170 billion nodes. This added complexity comes at a price, namely the energy required for training. GPT-3 consumed about 1 .3 gigawatt hours of energy. GPT-4, which is a current GPT model (not shown in the Figure), is estimated at about 50 to 60 gigawatt hours, similar to the energy consumptions of all

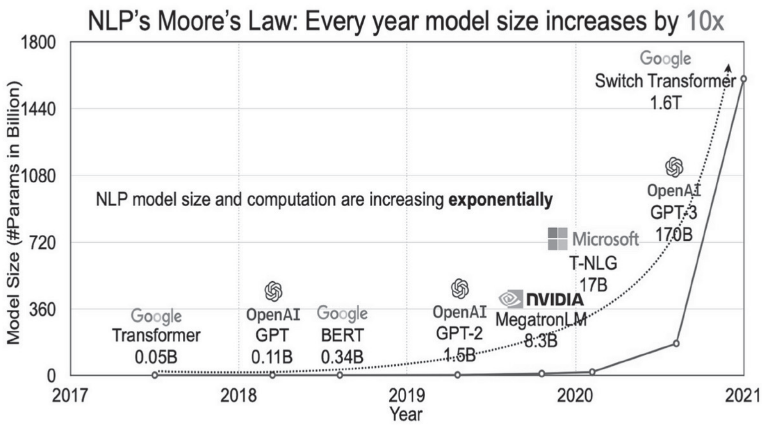

Figure 10 Growth of Large Language Models.

Q: If you collected $1 from every person on earth who earns less than $500 per year, how much money would you have?

GPT-3.5: To calculate the total amount of money collected, we need to estimate the number of people who earn less than $500 per year and multiply it by $1.
According to the World Bank's definition of extreme poverty, individuals living on less than $1.90 per day (approximately $693 per year) are considered extremely poor. However, if we consider individuals earning less than $500 per year, we can make a rough estimate.
As of 2021, the global population is approximately 7.9 billion people. It is challenging to determine the exact number of individuals earning less than $500 per year, as income data can vary significantly across countries and regions. However, for the sake of estimation, let's assume that 10% of the global population falls into this category.
10% of 7.9 billion is 790 million people.
Therefore, if you collected $1 from every person on Earth who earns less than $500 per year, you would have approximately $790 million.

Figure 11 ChatGPT3.5 Q&A.

of Hamburg's 1million households for 45 to 50 days. Clearly this points to the limits of growth for large language models, at least while they are being trained on the current paradigm. But this is not our focus here.

To note, large language models are "dumb" in that they do not have awareness of their reasoning in a way humans do. But it turns out, a lot of question can be answered very well without such awareness. Specifically, when GPT-3.5 was asked "if you collected one dollar from every person on Earth who earns less than 500 dollars a year, how much money would you have?" the system not only came up with a credible answer, but also provided the underlying logic that explains the result (Figure 11). An earlier experiment with human subjects using a similar question resulted in many non-responses or poor responses.

5. ARTISAN INTELLIGENCE

This brings us to the realization, that with knowledge capture and bootstrap through artificial intelligence, a new differentiation among categories of knowledge may emerge. Much of what we know might be described as *commodity* knowledge. As everyone knows, commodity products are worth very little and thus also difficult to monetize. Hence this knowledge cannot generate much impact and thus is freely shared. Next is a zone where artificial intelligence is rather strong, thereby competing for jobs with competent individuals, those who carry out knowledge work, but at a more general level. We can think of translators, journalists, image creators, entry level legal workers, programmers, or such. These are people operating in structured domains, especially those with well-codified professional rules sets (programming rules, laws, accounting principles). These are areas where the predictive mechanisms of G/AI produce high quality results, thereby making these professions highly vulnerable to AI replacement.

With these competencies taken away by G/AI, what is left for humans? One area will be a competency we might describe as *Artisan Intelligence*. Artisans are highly competent in the execution of their practice, while their practice also has a creative element. Hence the practice is characterized by specialization and by uniqueness, but also by a lack of codification of the practice. Artisan work is thus demonstrated in its outcome, while ingredients and processes are not widely shared. Artisan intelligence as such then relies on tacit knowledge, knowledge that is not codified, plus tangible, new-to-the-world outcomes as work product.

As such, the intension of an artisan's knowledge and skill is only told through its extension, in the outcome. One of the examples might be the work of violin maker Stradivarius. People still today are wondering how a Stradivarius violin aged 300 years can still sound exceptional. People have equally wondered how it is possible that the glues that hold

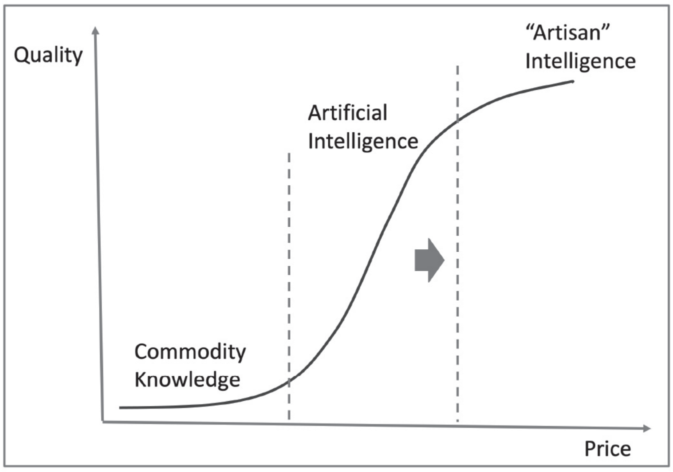

Figure 12 Intelligence Continuum (hypothesized).

the instrument together have not dried up or become brittle in all that time? Noteworthy then, Stradivarius knew even centuries ago, that the glue and varnish recipes were part of their competitive advantage and thus did not share.

Another example of artisan intelligence can be seen in Warren Buffett's exceptional investment expertise. Buffett has been able to beat the the stock market for most of the last 50 years, having yielded 150 times the S&P 500's performance. But it is not well understood what factors Buffett considers in his decision making nor what decision rules he uses. Otherwise, data mining algorithms would by now have matched his performance. Again, the knowledge is only revealed in the outcome, the actual stock picks. Yet people are willing to subscribe to this knowledge as illustrated by the large number of followers and Berkshire (BRK) stockholders.

Artisan knowledge is thus the highest level of knowledge in the continuum shown in Figure 12. This premium knowledge generates exceptional results but is also expected to command a high price, thus separating itself from lower levels of knowledge.

6. CONCLUSION

With this observation we have actually come full-circle. The discussion began with collective intelligence and the ability of collectives to perform at a higher level than individuals, based on aggregation of fractional knowledge, process and repetition. Then we saw that artificial intelligence, with its ability to capture and bootstrap that collective knowledge, can codify and parameterize these insights, thereby appropriating this expertise and disintermediating us from our knowledge, beating us on speed and efficiency. As you will know, ChatGPT-4 has absorbed much of the public WWW, including Wikipedia. So GPT knows everything that Wikipedia knows, plus a lot more. Measured in IQ points, GPT and other G/AI models by now approach the performance of highly intelligent individuals in these tests. The AI may not be as smart as the humans but can perform as if it were that smart.

With the insight that all shared knowledge may quickly become near-commodity knowledge and thereby undermining the competitiveness of the knowledge creator, this leads to only a few sustainable outcomes. First, those able to do so may engage in artisan knowledge creation. Next, everyone who is creating knowledge must from now on guard that knowledge and limit its entry into the public domain. Copyrights are not good enough but have to be replaced by lightweight idea rights—patents are too clumsy for this. Finally, the new world of knowledge sharing will require micro-payment mechanisms for the transfer and use of others' knowledge. Publicized ideas, other than common knowledge, that are used by others, must be monetizable. So if ChatGPT were using your ideas in its reasoning, you should receive a micro-payment, just like the creator of a song or other recognized intellectual property. Without it, economic rents for the provision of knowledge will inappropriately be shifted to G/AI engines, to the detriment of human civilization and human thought progress.

7. Q&A

Question: With ChatGPT, the large language models, we now have a tool to create things that seem to be true, and we can spam the entire system. Is there a solution for that?

Wagner: Thank you for pointing it out. So let me first say something about the academy. We as researchers, are one of the less affected groups of people, because we are operating at the artisan level. As a researcher, I may ask generative AI to produce some supporting research facts for me, add some background for papers, maybe to provide an overview. GPT may give a college-level overview, but it does operate at the level that I would want to share with fellow researchers. Well, at least it is not yet trained for that. And if you know these large language models, you know that they are trained on a vast corpus of knowledge, and then we add a front or add some fine tuning. The front end is not yet ready for the scholarly research task. After all, when it comes to knowledge creation, academic researchers are the outliers who create bespoke knowledge, right? This doesn't mean that large language models aren't useful in academic research. They can clearly provide a base level of knowledge from which to build. At the same time, the spam that may affect AI generated content in other knowledge creation environments will therefore not affect us academic researchers as much—yet.

Question: In your role as researcher in Hong Kong, do you see a problem with disinformation?

Wagner: Simply said, for researchers coming from an engineering science discipline and focusing on technologies supporting decision making and problem solving, I am less concerned with disinformation. We are individuals who are seeking solutions that are, technology-based advancements for the world. If you were a humanist or working in political science, I think your life is now more complicated than it used to be.

Question: What do you see as the ideal role for the next level of AI?

Wagner: Where I see the next step of the AI, is in creating models for us. You may ask your AI collaborator–the generative AI—which models should apply to this problem? And with that, you might work instantly with 50 or 60 alternate explanatory models, suggested to you by the AI, whereas traditionally, a researcher may have just considered a handful of explanations. This will offer a dramatic advantage to researchers in the future.

Question: Where do you see the truth in social media in 10, 20 years? How will this develop?

Wagner: I view the question of truth with scepticism. It is becoming obvious that with generative AI, that any kind of medium can be recreated, audio, image, or even video. Soon there will be AI-generated videos based on just descriptions of the topic and maybe some dialogue. Hence, you'll almost not be able to trust anymore anything that you see or hear or read unless you are a specialist in social media, with the ability to independently validate.

Question: Is there any research regarding where the collective intelligence, where the feed needs to come from, where the knowledge snippets should come from? Is there anything there? And have you heard any research about the quality or decrease of that quality for ChatGPT, depending on who is interacting with the system?

Wagner: In a way, GPT, harvests already much of what is publicly available. GPT-3's training data set Webtext 2 was supposedly about 45TB in size, and GPT-4 should be an order of magnitude larger. With all this knowledge having been farmed, creators of proprietary information, now have an advantage, because that proprietary knowledge will be the next farming ground for the AI technologies, in order to reach that next level of reasoning performance. As we all know, when it comes to actionable information, everything you can get for free as information is typically not particularly useful to give you a competitive advantage. The information which gives you the advantage is the rare one that only a few possess.

As to possible future improvements, I would see new capabilities arising from causal reasoning, because these models are currently quite shallow. They do not appear shallow in their behavior, but they are. They are operating conversationally and are predicting what text should be following based on prior conversation. This has worked very well. But it comes to a limit. And now we need hierarchical decomposition, breaking problems apart into smaller problems and then finding answers and then bringing those out. That, I think, is the next level to come, mimicking more closely what the human does, because we are doing everything that these technologies do in a much more efficient way.

REFERENCES

1. Jiang, L., & Wagner, C. (2022). A machine-learning approach for research model creation in theory building research. *Issues in Information Systems, 23*(3).
2. MacCrimmon, K. R., & Wagner, C. (1994). Stimulating ideas through creative software. *Management Science, 40*(11), 1514–1532.
3. MacCrimmon, K. R., & Wagner, C. (1991). The architecture of an information system for the support of alternative generation. *Journal of Management Information Systems, 8*(3), 49–67.
4. Majchrzak, A., Wagner, C., & Yates, D. (2013). The impact of shaping on knowledge reuse for organizational improvement with wikis. *MIS Quarterly*, 455–469.
5. Wagner, C. (2000). The Illusion of Knowledge: Or Why Data Is More Important than Knowledge. In Intellectics and Computational Logic: Papers in Honor of Wolfgang Bibel (pp. 347–359). Dordrecht: Springer Netherlands.
6. Wagner, C. (1996). Creative behavior through basic inferences: Evidence from person-computer interactions. *Journal of Creative Behavior, 30*(2), 105–118.
7. Wagner, C., & Suh, A. (2014, January). The wisdom of crowds: impact of collective size and expertise transfer on collective performance. In 2014 47th Hawaii International Conference on System Sciences (pp. 594–603). *IEEE.*
8. Wagner, C., & Jiang, L. (2012). Harnessing the power of social media for creativity support: A three-pronged approach. *Knowledge Management & E-Learning, 4*(2), 174.
9. Yates, D., Wagner, C., & Majchrzak, A. (2010). Factors affecting shapers of organizational wikis. *Journal of the American Society for Information Science and Technology, 61*(3), 543–554.