Handbook of Collective Intelligence

Welcome!

Terms and conditions. Please see following section on using and contributing to this handbook.

If you wish to contribute, please create an account by clicking here Thank you.

This Handbook provides a survey of the field of collective intelligence, summarizing what is known, providing references to sources for further information, and suggesting possibilities for future research.

The handbook is structured as a wiki, a collection of on-line pages, editable by their readers.

The handbook is hosted by the MIT Center for Collective Intelligence, but we hope that researchers and others from around the world will contribute to it. The process of creating this handbook could itself be an example of collective intelligence.

In parallel with this Wiki, we are developing a database of bibliographic references at citeUlike. The wikis here link to these references at citeUlike.

Contents

[hide]

Goals

This Handbook provides a survey of the field of collective intelligence, summarizing what is known, providing references to sources for further information, and suggesting possibilities for future research.

Unlike Wikipedia (an encyclopedia created by many people editing a wiki), this handbook is not structured as a “flat” collection of interlinked, but conceptually independent, short articles. Instead, this handbook is intended to provide a conceptual framework for the whole field of collective intelligence.

We hope that the top level of this framework will be written in a way that is understandable to a very broad general audience. At many points in this framework, however, we expect to have links to separate pages that elaborate on specific points or issues. Often, these separate pages will include specialized material and references for one or more of the many different fields related to collective intelligence. Even though some of this specialized material may be understandable only by specialists, it should be introduced so that other readers will at least be able to understand the basic issues involved.

Using and contributing to this handbook

Participation in the Handbook of Collective Intelligence is completely voluntary and participation will be subject to terms and conditions that will be added at a later date.

Thank you for your participation in this very exciting and groundbreaking project

Please bear in mind we expect to experiment with several different approaches to creating this Handbook.

Table of contents

Why study collective intelligence now?

Collective intelligence has existed for at least as long as humans have. Tribes of hunter-gatherers, nations, and modern corporations all act collectively with varying degrees of intelligence.And, from some perspectives, even collections of bacteria, bees, or baboons can also be viewed as collectively intelligent (e.g., Bloom 1995).

But this ancient phenomenon is now occurring in dramatically new forms. With new communication technologies—especially the Internet—huge numbers of people all over the planet can now work together in ways that were never before possible in the history of humanity (e.g., Malone 2004). It is thus more important than ever for us to understand collective intelligence at a deep level so we can create and take advantage of these new possibilities.

This document tries to help in that process by organizing what we already know and identifying what we don’t.

What is collective intelligence?

Our current working definition of collective intelligence is: *Groups of individuals doing things collectively that seem intelligent.* Details here.

Collective Intelligence:

- Old concept, been around since hunter-gather societies
- What’s different now are web technologies
- How can people and computers be connected so that collectively they act more intelligently than any individual, group, or computer has ever done before?

Examples of collective intelligence

We could classify the examples collected under different perspectives (see below).

Business organizations

See here for many more examples.

- **InnoCentive** [1], **NineSigma** [2], **YourEncore** [3]. NineSigma and InnoCentive are similar. The latter is started by Eli Lilly in 2001 and posts engineering and science problems by subscribing corporations such as Boeing, DuPont, and Procter & Gamble. The cash rewards for solutions range from thousands to hundreds of thousand dollars. It claims that as of 2006, 30% of the problems posted have been solved. YourEncore works with retired scientists and engineers. An important difference between set ups like these and our idea of collective intelligence is that these have individual problem solvers. However, it might be interesting to see what ideas about these settings might also apply when there are multiple solvers. Some of the lessons from these settings are in Lakhani et al. (2007) [4] who, in examining participants in Innocentive, observe that participants are more likely to actually solve the required problems when:
  - the prize is bigger
  - participants are more intrinsically motivated
  - participants have more free time
  - participants are non-experts in the field
  - participants are not participating due to career concerns, social motivations, or to beat others.
- **Sermo** [5]] closed community of health professionals (71,000+ U.S. physicians)
  - Individuals vetted as real physicians before gaining access; once online, operate via a pseudonym
  - Can ask any question to the community and answer any question; community of peers rank value of questions and answers
  - Sermo also provides incentives to post; several say they participate for fun or to learn
- **Collective Intellect** [6] aspires to summarize viewpoints in blogs and other web pages for applications in finance and marketing.
- **Current TV** [7], VH1's Web Junk 20 [8] and its sister company iFilm [9]: broadcast programming done by users.
- **LiveOps** [10] says it has 10,000 free agents signed up to become virtual call centers for corporations of many different industries.

Computer science and artificial intelligence

See here for more examples.

- **Image inspection**. Neufeld et al. (2003) adopt the paradigm of a partnership between artificial and human intelligences, in solving the problem of detecting edges. More examples here.
- "NASA Clickworkers" [11] where tens of thousands of volunteers worked to identify craters in Mars images, producing results that were indistinguishable from those produced by PhD image coders working within NASA.
- **Collecting user inputs**. Verbosity; see von Ahn et al. (2006) [12]. This is set up as a game, in which a Narrator gets a word from the system and sends hints to Guesser about what that word is. Both players form a pair to compete with other pairs, and once a Guesser guesses right, both Narrator and Guesser get a point. The key is that the hints issued by Narrators are sentences that are true of the word. More examples here.
- **Secure computing**. In general, this area incorporates "AI-complete" problems to boost security. Gentry et al. (2005) [13] propose that the payment for human computation can be an online currency. More examples here.
- **Collaborative filtering**. Spamnet [14] and Razor [15] use human votes to determine if an email message or source is spam. Zhou et al. (2003) [16], in an experiment involving 9,589 non-spam and 14,925 spam messages, find that the error rate of this human-computer system detecting false positives is negligible, at 10 to the power of 8. More examples here.
- **Distributed tasks**. The Distributed Proofreader Project [17] is originally set up to reduce OCR errors in book scanned under Project Gutenberg [18], but is now a stand-alone entity. Volunteers offer to correct these books in several passes, matching OCR pages against scanned pages. There is no commitment on the part of volunteers. More examples here.

- **P2P technology**. [19]
- **Swarm intelligence**. [20]

Biology

See here for many more examples.

Mallon et al. (2001) [21] report that bee scouts who have done their dances stand aside to observe dances by bees for other candidate sites; they do not continue dancing to do more promotion, nor do they interfere with competing dancing bees. This is a form of unprejudiced opinion polling. The disadvantage is that bee scouts basically lose information about sites they promoted. The advantage is that the bees as a whole do not waste too much time in endless debates. More examples here.

Computer-supported collaborative work

See here for many more examples.

Engelbart, motivated by the need for people to collaborate in order to cope with accelerating rate of change, proposes Augment, a "co-evolutionary environment capable of handling simultaneous complex social, technical, and economic changes at an appropriate rate and scale" [22]. He calls our capacity to cope with change our "collective IQ" (See section on Engelbart here).

Prediction markets

See here for many more examples.

- A number of firms [23], such as Hewlett Packard (called Behaviorally Robust Aggregation of Information Networks, or BRAIN [24]), IBM, Microsoft, and Ford.
- As a failed example, Incentive Markets closed in Feb 2005 [25]. It was designed to gather public guesses on success of pharmaceutical products in the pipelines of big pharma firms.

A taxonomy of collective intelligence

In order to make progress in understanding collective intelligence, it is useful to have a way of classifying different kinds of collective intelligence into categories.

One effort to do this is summarized here.

Measuring collective intelligence and the factors that affect it

One example of how collective intelligence might be measured is to consider how group performance has been measured (see section on group performance.

- **Task type**. For example, Steiner (1972) [26] proposes that group performance can be measured relative to some benchmark, so that is may be classified as process loss or process gain (sometimes also called the "assembly bonus effect"). For Steiner, group performance depends on the type of task: divisible, optimizable, or combinable (or unitary). The last may be subdivided into:
  - **additive**: Ringlemann (1913) [27]reports how the tonnage a group of people can pull using a rope. He reports that, unsurprisingly, a group can pull a greater weight than any individual can, but interestingly, each individual pull less hard, so there is decreasing returns to group size. He attributes the latter to coordination loss (people pulling at different times) and motivational loss (social loafing).
  - **compensatory**: Performance increases with group size (if guesses are independent?). Tziner and Eden (1985) [28] report that, among 208 three-man military crews, ability and motivation have compensatory (the authors say "additive" but they mean "average," rather than "total") effects.
  - **conjunctive**: Performance is that of the weakest member, so group size again reduces group performance. Conversely, if less skilled

members increase their effort, group performance increases (the Kohler effect, see Kohler (1926) [29]).
    - **disjunctive**: group size increases group performance. Taylor and Faust (1952) [30] report that 4-person groups do better than 2-person pairs, which in turn do better than individuals, in answer "20 questions" puzzles. This is when measuring absolute output (time needed to get right answers). But when measuring man-minutes, when the reverse is true (put another way, it is cheaper to get the job done, if the time needed to get it done is not a constraint).

- **Task-relevant information**, or a frame or reference given to group members. SeeLaughlin, et al. (1999).

- **Cohesiveness** (see the meta-analyses in Evans and Dion (1991) and Mullen and Cooper (1994)). Lea at al. (2001) suggest that electronic groups might develop stronger identities and norms, contradicting the prediction by the Social Identity Model of Deindividuation Effects (SIDE) that visual anonymity reduces group attraction. Others have found this relationship from cohesivenss to performance, but moderated by other variables:
    - Group norms. Langfred (1998) reports a similar result, but uses data from a 1995 survey of Danish military. He also finds that the link from cohesiveness to performance is moderated by the direction of norms (i.e., task-oriented or not).
    - Commitment to the group's goal, Podsakoff et al. (1997)
    - Technology mediation. Straus and McGrath (1994) report that in a study of 72 3-person groups, the face-to-face groups do significantly better than the computer-mediated ones, for the more complex tasks.

- A **transactive memory** of who knows what. Liang et al. (1995) report that subjects in a task of assembling transistor radios, groups that undergo training together to develop transactive memory do better. Moreland and Myaskovsky (2000) suggest that this can be achieved through group training that focuses on members' competencies, rather than group identification or motivation. See also early work by http://www.citeulike.org/user/rlai/article/1397545 Wegner et al. (1985)] and Wegner et al. (1991).

- **Team leadership**. Hackman (2004) and Hackman and Wageman (2004) argue that leaders can spell the difference between success and failure because they determine what kind of team is created and how the team is structured and coached. They also suggest that leaders often make suboptimal decisions.

- **Deployment** of suitable members for group tasks. Some studies identify general traits, such as loquacity, use of reason to influence, dominance, see Littlepage and Mueller (1997), Littlepage et al. (1995).

- **Training**. Conversely, configuring groups to better recognize and deploy expertise, such as giving explicit instructions to share information and exploit expertise, getting feedback, or investing in working together (e.g., Henry (1995)).

- **Brainstorming**. However, the consensus today is that nominal groups outperform brainstorming groups, Paulus and Dzindolet (1993).

- **Stress**. The literature on individual psychology, the consensus is that stress has an inverted-U shape relationship with performance—e.g., Jamal (1984). The relationship also slides rightward as individuals adapt to stress. Further, the increase in performance is biased toward quantity and against quality. The results seem to be also supported among groups, see Karau and Kelly (1992). For example, Kruglanski and Webster (1991] report that within a group, deviating opinions are rejected more with time pressure while conforming ones are accepted more. Stress works through several mechanisms in its link to performance:
    - Focusing resources. Karau and Kelly (1992) suggest that stress, in the form of time pressure, forces groups to focus their resources.
    - Need for closure. Kruglanski et al. (2006) summarizes their long line of research on how stress increases the need for closure, and therefore greater conformity in the form of a "group mind."

Consciousness as a Measure of Collective Intelligence

Consciousness and Collective Intelligence

If we think of consciousness as the ability to be aware of the external environment and our presence within it one of the potential advantages of Collective Intelligence would be an increased awareness of more elements of our environment and a wider range of potential options for how to interact with it.

The real power of CI in relationship to consciousness goes deeper. One of the areas that most scientists and thinkers struggle with in the field of consciousness is explaining how it is created. At this point the best we can figure is that it is an emergent phenomena that arises from the interaction of the comparatively simple elements that underlie it. In the case of a human brain it is the collective interaction of neurons. In the case of a bee hive it is the collective interactions of all the hive's members. The key to how CI enhances consciousness, resides in the mystery of the synapse.

The Importance of the Synapse in Creating Consciousness

Here synapse does not refer strictly to the neural structure, rather it is a metaphor for the larger the larger phenomenon of data changing states as it is transmitted from one unit to another. For our purposes a good working definition is:

Synapse: Any point at which data changes the medium in which it is encoded

Examples of synapses:

• Electrochemical data in the brain translated into speech or writing before being converted into electrochemical data in another brain

• Electrochemical data in a bee translated into movement before being converted into electrochemical data encoded in other bees

• Ant knowledge of food sources converted into chemical trails that are picked up by other ants

The importance of the synapse has been overlooked by many thinkers, and yet it is the most common element in any conscious system be it an ant hill, a mammalian brain, or a football team. According to the above definition even the interaction of a computer with its hard drive could be accurately described as a synapse. It, however, is a rudimentary synapse that lacks many of the features of synapses that appear to correlate with consciousness. The key characteristics of synapses that appear to lead to the creation of consciousness:

• Delay in time

• Ambiguity of the newly encoded data

• Ability of the new data form to stimulate multiple receptors in a synapse with a single emission

• The potential for emissions from multiple incoming data sources to strengthen or weaken the reaction of the receptors

Collective Intelligence increases the consciousness of both its members and the group as a whole not merely by pooling more data but by exponentially increasing the number, and sometimes, quality of the synapses in the system.

Measuring Consciousness

The increase of the consciousness of a CI system can be measured by:

• Speed by which it detects changes in the external environment of its constituent members

• The diversity of potential solutions that are posed within the system

• The quality of options chosen as determined by the degree to which the solutions permit the constituent members to thrive within the external environment.

Predicting Levels of Consciousness

I would further argue that the consciousness of a given CI system or community can be predicted by looking at the synapses that hold it together according to the following criteria:

• Synapse density—How many receptors/emissors come together in a particular synapse

• Firing intensity—How reactive the emissors are to various external stimuli

• Synapse receptivity—How quickly and vigorously do the receptors react to the stimulus both by triggering actions and by becoming emissors in other synapses.

Predictive measures of collective intelligence

See also the section on predicting individual intelligence.

Cattell (1948) [31] argues that a fair amount of research into individual psychology might apply to group psychology. The latter might be analyzed at three levels: (1) some statistics describing individual members of the group, (2) the structural relationships among members and between groups (for example, group membership might overlap), and (3) the syntality or behavior of the group. He also provides 7 theorems that determine the syntality of groups:

1. dynamic origins of groups
2. vectorial measurement of synergy
3. subsidiation in the syntal lattice
4. subsidiation in the personal lattice
5. hierarchies of loyalty from the law of effect
6. synergic constancy in a system of overlapping groups
7. isomorphism of syntality change and personality change.

Wegner (1987) [32] suggests that the converse, that the complementarity and specialization shown in group work might have facilitated the development of individual minds that have complementary and specialized functions.

See other ideas on predicting collective intelligence here.

What factors facilitate collective intelligence?

We list some factors drawn from different perspectives, as a working list. See the different perspectives for a complete listing.

Diversity

Diversity among members of a collective is a recurring theme from several perspectives.

For example, one view in cognitive neuroscience is that diversity (beyond one's relatives, but of an invested SPS) is an essential, not a luxury, to procreate our collective gene pool.

Consistent with this, the idea of homosocial production in the social psychology is that the lack of diversity increases the bias in collective decision-making.

Crowdsourcing rests some of its perceived value on the Diversity Prediction Theorem:

```
Crowd Error = Average Individual Error – Diversity Among Individuals
```

Crowd Error: For each question of interest, first calculate the distance between what the crowd of individuals predict (i.e., crowd consensus) from the actual outcome or “external truth”, then square this value. A distance of zero represents no error between what the crowd predicts and external truth. After squaring the distance for each question of interest, then add together these values to produce a sum of squares error for the crowd vs. external truth. Of note, squaring the calculation of distance is important so that positive and negative error values do not cancel out each other (this is a common practice in statistics).

Average Individual Error: For each question of interest, first calculate the distance between an individual (i) predicts from the actual outcome or “external truth”, then square this value. After squaring the distance for each question of interest for a given individual (i) vs. external truth, then add together these values to produce a sum of squares error. Repeat for all individuals in the crowd and compute the average sum of squares errors for an individual vs. external truth.

Diversity Among Individuals: For each question of interest, first calculate the distance between an individual (i) predicts from the consensus of the crowd, then square this value. A distance of zero represents no diversity (or different) between what the individual predicts and what the crowd predicts. After squaring the distance for each question of interest for a given individual (i) vs. consensus of the crowd, then add together these values to produce a sum of squares error. Repeat for all individuals in the crowd and compute the average sum of squares errors for an individual vs. consensus of the crowd.

(For more info see: Page, Scott. (2007) The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies.)

The beauty of the Diversity Prediction Theorem is that it works every time – there’s no way for it mathematically not to work. Specifically, if the crowd error differs from what the average individual error predicts vs. external truth, then there must be differences between what individuals predict amongst themselves (i.e., not perfect consensus). As such, if there are differences between what individuals predict amongst themselves, there also must be diversity in the different perspectives of individuals vs. the consensus of the crowd.

Viewed a different way, if one collapses the diversity in the different perspectives of individuals vs. the consensus of the crowd, one equally collapses the difference between the crowd error and the average individual error by forcing little or no difference between individuals. At zero diversity among individuals, everyone predicts the exact same thing, so the average individual error equals the crowd error at such a point.

So what does all of this mean for collective intelligence?

First, it means that if one wants to improve the accuracy of what a crowd predicts vs. external truth, one wants a diverse group of participants with different perspectives. Increasing the diversity in what they predict will reduce crowd error vs. external truth.

Second, while the Diversity Prediction Theorem is great, it does not guarantee in and of itself that unknown outcomes at a future date, better decisions will be made. For that, empirical research is required.

- With crowds: diversity trumps ability
- Yet clearly there are times when crowds are smart and also times when crowds are dumb, so why does diversity influence outcomes?
- Consider: individuals have particular perspectives on a problem, paying attention to some aspects and filtering out others

Specifically:

- Learned perspectives may limit the search space any one individual uses to reach an answer, even for “smart” individuals
- Multiple individuals with varying perspectives expand the search space employed
- A diverse crowd has more “tools” to apply

That said:

- This is not to say that diversity does not have drawbacks
- Diversity works if everyone has same goal of getting the answer right, and values this goal
- If goal-related values of different groups are not shared, crowd may splinter into factions

Formal and informal structure

There is a very large literature in organization theory on how organization structures affect performance. While we do not necessarily think of formal organizations when we think of collective intelligence, work on the former suggests at least interesting parallel research questions for the latter.

The results from organization theory might also be interesting reference points for us in collective intelligence. For example, Burns and Stalker's (1971) [33] argue that organic non-hierarchical, informal structures are more suitable for innovation in unstable and dynamic environments, because they address the informational and social requirements. Mechanistic structures are more suitable otherwise. Allen (1977) [34] provides empirical support for this with a study of R&D teams.

Interestingly, Hinds and McGrath (2006) [35] argue the opposite might be true for distributed rather than collocated teams: in a study involving 455 individuals in 49 teams within a multinational firm, they find that hierarchical (albeit informal ones) do better than network structures. The reason is that network or organic structures are "difficult to address through electronically mediated exchange"; see also Hackman and Morris (1978) [36] and Nohria and Eccles (1992) [37]. Ahuja and Carley (1999) [38] find the same in a ethnographic 4-month study of a virtual organization.

Modularization of tasks

Another relevant result from organization theory is that modularized work reduces interdependency among group members, and can improve group performance. Baldwin and Clark (1999) [39] suggest porting ideas about modularity in computer software engineering to design and manufacturing work. They also caution that modularization of work also requires managers who are comfortable controlling less, and whose knowledge is more focused on their modules. In Baldwin and Clark (2003) [40], they argue that modularity enhances the private benefits of contributing to open software projects, making the equilibrium robust. Olson and Olson (2000) [41] also conclude that groups that are loosely coupled do better. Galbraith (1973) [42] contend that modularity reduces the need for information sharing. See also Krishnan et al. (1997) [43]. Modularity, of course, has its own issues. For example, Grinter et al. (1999) [44] note that in their setting, it contributes to more handoffs, leads to later detection of software bugs, results in divergent and isolated views among collaborators, and under-leverages expertise of distant group members.

Dense communications structure

It is arguably the consensus that dense communications structures enhance group productivity; see the pioneering work by Back (1974) [45]. One mechanism through which structures enhance productivity is by improving group identification and trust; see Portes and Sensenbrenner (1993) [46] on communications structures that do not have structural holes.

Incentives for contribution

There are many types of incentives that seem to work.

One example is economic incentives, along the lines of monetary prizes, such as those in InnoCentive [47], NineSigma [48], and YourEncore [49]. NineSigma and InnoCentive are similar. The latter is started by Eli Lilly in 2001 and posts engineering and science problems by subscribing corporations such as Boeing, DuPont, and Procter & Gamble. The cash rewards for solutions range from thousands to hundreds of thousand dollars. It claims that as of 2006, 30% of the problems posted have been solved. YourEncore works with retired scientists and engineers. An important difference between set ups like these and our idea of collective intelligence is that these have individual problem solvers. However, it might be interesting to see what ideas about these settings might also apply when there are multiple solvers. Some of the lessons from these settings are in Lakhani et al. (2007) [50] who, in examining participants in Innocentive, observe that participants are more likely to actually solve the required problems when:

As another example, recent research along the lines of human computation suggests that intrinsic motivation work. In the ESP Game [51] described in von Ahn and Dabbish (2004) [52], work (indexing images using tag words) is re-cast as a game (in which pairs of players who strive to quickly come up with matching tag works, compete with other pairs). See other section on this and related examples and counter-examples.

A different line of attack is to reduce the cost of contribution. For example, Google Earth supports the keyhold markup language (KML) for users to geo-code information onto the latitude-longitude grid. Other users can access published KML information. See section on geo-coding.

Shared vocabulary and other infrastructure

Schwartz (1995) [53] argues that groups will always perform worse than individuals unless there is some collaborative creation of a visual representation.

Shilad et al. (2006) [54] investigate how people develop, share, and use tags. They find that these processes depend on our current investment in tags already being used, the degree we want to conform with the tagging community. The community's tags in turn depend on what we use and the tag selection algorithm. They also find that subjective tags are used most for self-expression, factual tags for learning and finding, and personal tags for organizing.

Lee (2006) [55], in a study of del.icio.us tags, find that users who perceive greater degrees of social presence are more likely to annotate their bookmarks with information that could facilitate the sharing and discovery of bookmarks for other del.icio.us users.

Englebart and Ruilifson (1998) [56] conjecture that the tool system needs to be structured to improve the human system. Advancements in the tools, called outposts, need to be introduced as experiments, and in configurations, not as isolated gadgets

Awareness

Dourish and Bellotti (1992) [57] examine how mechanisms to share awareness in Quilt (see Fish et al., 1988 [58]), PREP (see Hiltz and Turoff, 1993 [59]), GROVE (see Conklin, 1987 [60]), and ShrEdit (see Olson et al., 1992 [61]) enhance collaborative work. These mechanisms may be classified as:

- Informational. For example, Quilt's authoring system is annotated with email messages.
- Role-restrictive. For example, the author role might allow editing, while a reviewer role does not.

Learning

In prediction markets, Adams (2006) [62] and Ottaviani and Sorensen (2006) [63] show that with multiperiod learning, the market price converges to traders' mean beliefs, aggregating all private signals.

The Power of the Edge

- Traditional organizations exist to harness worker output, usually employ a top-down approach to maintain efficiency
- Problem: top-down hierarchies restrict opportunities for bottom-up information flow from the edge of an organization

Specifically:

- As the environment external to an organization changes, more “eyes” are on the edge than at the top seeing these partial elements of these changes
- If the organization doesn’t encourage bottom-up percolation of these insights as changing environment, high probability that the the top could lose relevancy despite larger scope

That said:

- This is not to say that the edge does not have drawbacks
- Top-down hierarchies are common as they are great at controlled consistency
- Bottom-up approaches include overlapping and potentially conflicting efforts, lacking consistency across an organization; but web technologies can help address this

(For more information see Bray, David and Prietula, Michael. Exploration and Exploitation: Managing Knowledge in Turbulent Environments (2007). International Conference on Information Systems. Best Paper KM Track.)

The Power of an Ecosystem

- Having just one view, or one highly regimented view, limits situation awareness
- A single view misses unknown elements

Specifically:

- Where any one individual is looking may no longer be relevant; whereas multiple individuals can better scan an unknown environment
- Through chatter and conversation, multiple individuals can swarm around interesting features in the external environment

That said:

- This is not to say that an ecosystem does not have drawbacks
- Unlike highly structured or single expert approaches, ecosystems are less efficient than other approaches when the environment is relatively stable (nothing new is occurring)

(See March, James. (1993) Exploration and Exploitation in Organizational Learning. Organization Science. Also see Bray, David (2008). Knowledge Ecosystems: Technology, Motivations, Processes, and Performance.)

Others

Apart from academic research, a number of writers also suggest factors that facilitate collective intelligence.

Bloom (2000, p. 42-44) [64] says that the following five elements cause a group to be intelligent (a “collective learning machine”):

- Conformity enforcers – mechanisms that cause similarities among most members of the group
- Diversity generators – mechanisms that cause some differences among members of the group
- Inner judges – mechanisms that cause individual members of a group to reward themselves for successes and to punish themselves for failures.
- Resource shifters – mechanisms that shift resources (such as money, admiration, and influence) to members of the group who are successful and away from those who aren’t.
- Intergroup tournaments – competitions between subgroups (such as games, corporate competitions, and wars)

(Note: Logically, it seems that a system could work without elements (3) and (5).)

Surowiecki (2004, p. xviii – xix, and chapters 2-4) [65] says that three conditions are necessary for a group to be intelligent (for a “crowd to be wise”):

- Diversity – the group includes members with a wide diversity of knowledge or abilities (and the ability to recognize successful and unsuccessful outcomes)
- Independence – group members use their own knowledge and abilities without being overly influenced by others.(When group members influence each other too much, various kinds of bad outcomes can results. See section on “Groupthink and Informational Cascades” below.)
- A particular kind of decentralization – group members’ actions are aggregated in a way that finds the right balance between (a) “making individual knowledge globally and collective useful”, and (b) “still allowing it to remain resolutely specific and local” (p. 72).

What factors inhibit collective intelligence?

Again, we list some factors drawn from different perspectives, as a working list. See the different perspectives for a complete listing.

Biases

Many of the above seem to also appear at the **group** level. For example, Argote, et al. (1990) report evidence that groups suffer from the representativeness fallacy. Hinsz and Indahl (1995) report that groups are also susceptible to anchoring, such as juries who anchor their damage awards to the amount requested by plaintiffs. Sunstein et al. (2002) even argue that group juries exhibit more anchoring than individual jurors.

Kerr, et al. (1996) use Davis’ (1973) theory to suggest when groups or individuals might be more biased. Hinsz et al. (2007) suggest that aggregating individual biases to the group level may attenuate or accentuate the biases.

In addition, there are biases that are unique to groups (see review in Hinsz, et al. (2007)):

- In-group bias. Individuals favor their own group members. The bias appears to be reduced with the size of the group. See Sherif's (1936).
- Out-group homogeneity bias. Individuals see members of other groups as more homogeneous than those in their own groups. The bias seems to be independent of the size and number of groups, and is not due to the relatively less interaction between the individuals and out-group members.
- Groupthink, bandwagon effect, herd behavior. This is the tendency to do what others do. See sociological perspective.
- Facilitation and loafing . A version of this is the contrast effect. For example, a person placed next to a less appealing one is viewed as more appealing than normal, but when placed next to a more appealing one, is viewed as less appealing than normal. See the section on Social facilitation and loafing
- Group polarization. This is the tendency for groups to adopt more extreme or riskier positions after discussion. Indeed, the positions are often more extreme than each individual would wants. One probable cause is the desire to conform. Another comes from [persuasive argument theory,] in which members suggest more and more reasons to distinguish the options, so that the final option chosen is backed by a lot more reasons than if there were no discussion.
- Biased use of information and the common knowledge effect. See the section on Common knowledge effect.
- Risky shift. Zajonc, et al. (1968) and others documented that groups tend toward more risky decisions. In a trial in which groups and coacting individuals choose between getting 0.75 cents if a left bulb lights up and 3 cents if the right lights up, and the subjects know that the left lights up with a probability of 80% and right, 20%, more As Davis (1973) [66] explains, the standard explanations for this include: (1) riskier members are more persuasive, (2) increased familiarization through group discussions lead members to riskier choices, (3) diffused responsibility also leads members to riskier choices, (4) there is [cultural value inclining subjects toward risk,] (5) unlike the above, the observed group outcome might be due to the social decision scheme rather than shifts in individual preferences (Vinokur (1969)). All but the last two have been doubted, especially because a converse cautious shift is also in the literature (Dion, et al., (1997)). Cartwright (1971) document that 75% of observed group decisions might be due to the social decision scheme, rather than individual preference shifts.
- Distortions in multi-level group decisions. Davis (1973) [67] suggests that when there are multi-level group decisions, such as in a democratic political process, the "people’s preference may be very distorted if we use a fair majority social decision scheme. In practice, such distortions might be corrected with minority reports and interest groups. But we are aware of no data that permit a test of the distortions-by-levels argument."

Social capital

de Tocqueville (1835) [68] observes that the norms of tolerance, reciprocity and trust are important determinants of formation of communities and associations. Putnam (2000) [69] contends that national stockpile of 'social capital'--our reserve of personal bonds and fellowship--is seriously depleted [70] and initiates the Saguro Seminar series to restore social capital in America; see his new book Putnam (2003) [71]. Perkins, et al. (1990) [72] also creates a sense of community index, which measures membership, influence, integration, and shared emotional connection. Another effort is the Asset-based Community Development [73] Institute at Northwestern University.

Underlying all these is the claim that social capital enhances collective intelligence, or at least sustain collective intelligences.

Narrow bandwidth among members

For example, Cramton (2001) [74] suggests that narrow bandwith, such as that caused by geographically distributed teams, results in difficulty in sharing mutual knowledge. She further classifies this failure into 5 types:

- failure to communicate and retain contextual information
- unevenly distributed information
- difficulty in communicating and understanding the salience of information
- differences in speed of access to information
- and difficulty interpreting the meaning of silence.

Cultural boundaries

Blau (1970) [75], among others, contend that a main barrier to collaboration may be the difficulty in achieving agreement when diverse viewpoints exist. This can make effective decision-making more difficult. Even if collaboration members do manage to agree they are very likely to be agreeing from a different perspective. This is often called a cultural boundary. For example:

A culture where rank or job title is important makes it hard for a lower rank person who may be more qualified than their superior for the job it had to collaborate. The lower rank person is told what to do. This is not collaboration "stranger danger"; which can be expressed as a reluctance to share with others unknown to you "needle in a haystack"; people believe that others may have already solved your problem but how do you find them "hoarding"; where people do not want to share knowledge because they see hoarding as a source of power "Not Invented Here"; the avoidance of previously performed research or knowledge that was not originally developed within the group/institution.

Self-interest and the free-rider problem

Olson (1971) [76] propose that individuals' self interest might hinder the emergence of collective action. He believes that group size reduces the cost of not participating (free-riding, see also Simmel and Wolff (1964) [77]). Further, group size reduces the share of benefit but increases the cost of participation (the interest of the individual is less aligned to the group average, and there is higher cost of organizing a large group). See details in the section on [Political_philosophy_perspective_on_collective_intelligence#Logic_of_collective_action | the logic of collective action.]]

Implementation issues

Bikson and Eveland (1996) [78] describe CSCW (computer-supported collaborative work) implementation issues which might also be pertinent to implementation of collective intelligence systems. Among their key findings of their qualitative study at the World Bank:

- Change in social and technical systems (the former like work groups, jobs, interdependencies; the latter like hardware, software, networks)
- Implementation has as strong an influence as technology on outcomes
- Outcomes evolve over time.

Grudin (1988) [79] list some lessons from CSCW implementation that might be instructive for collective intelligence systems:

- Divergence in incentives between users and maintainers. Perhaps more accurately, for everyone, the benefit must outweigh the cost—e.g., Carasik and Grantham (1998) [80] attribute the failure of implementing the Coordinator at Pacific Bell to the high cost of training to the benefits that could be derived.
- Breakdown of intuitive decision-making. Decision makers may over- or under-state the cost-benefit ratio for segments of user populations that they are not familiar with.
- Underestimation of difficulty in evaluating CSCW.

Market failure

Market failure is especially important in inhibiting the proper function of prediction markets.

1. Overly restrictive or unclear specification of event to be predicted. For example, a 2006 Tradesports contract[81] on whether North Korea conducts a missile test specifies that the US Department of Defense as a confirmation source. But on this event, the DOD does not confirm the incident, even though it has been widely reported on in the media. In another case, a user alleges[82] that Tradesports announces an outcome of a game in a random number of seconds after the specified time of the event. Wolfers and Zitzewitz (2004)[83] conjecture that play money might facilitate the introduction of loose definitions, enlarging the set of questions that might be amenable to prediction markets.
2. System downtime during critical betting periods. For example, one user alleges that Tradesports' site is down during the last one minute of a crucial SMC/USF game on Feb 19, 2007.
3. Involvement of biased parties, and cornering the market. Rhode and Strumpf (2003)[84] studied bets on outcomes of presidential elections between 1868 and 1940, and conclude that there is little evidence in the historical record that wagering compromised the integrity of elections despite the active involvement of political parties (pg. 1). In contrast, Friedrich Dürrenmatt's The Visit (1956, 62)[85] is a play about rich lady Claire Zachanassian who visits the town of Güllen. She puts a big bet that a resident Alfred Ill, a former lover, will be killed. The beneficiary of the bet is the impoverished town. Although the townsfolks initially found the bet repulsive, one by one, they begin to spend beyond their means, as if the town will someday be rich. In the end, Alfred Ill dies mysteriously.
4. Predictions of extreme events. The favoriate-longshot bias has been documented in TradeSport, Chen et al. (2006)[86] and Fair (2006)[87]
5. Long-lived contracts? SimExchange[88] uses infinite-life contracts (mimicking equities in the stock market) in a prediction market for video game sales. Some observers[89] suggest that this might work only for play money.
6. Affiliation bias. Koleman (2004)[90] shows evidence that New York betters in the Iowa prediction market favor the Yankees. Forsythe, et al. (1999)[91] show that traders favor bets of their own political parties.

We also conjecture that many factors that inhibit healthy development and functioning of traditional asset markets, such as the stock market, might inhibit CI. Such factors include:

1. manipulation. However, Hanson, et al. (2006)[92] argue that at least in an experimental market they studied, manipulation did not last and the non-manipulators compensate for the bias in offers from manipulators by setting a different threshold at which they are willing to accept trades.
2. insider trading. Hanson (2007)[93] also argues that there are ways to curb excessive insider trading, such as requiring elite traders (insiders) to disclose their trades ahead of time.
3. thin markets--e.g., O'Hara (2003)[94]
4. herding--e.g., Banerjee (1992)[95]. Plott and Sunder (1982)[96] and Plott and Sunder (1988)[97] show that in experiments, bubbles seem to occur. On the other hand,
5. poor rule of law--e.g., Shleifer and Vishny (1997)[98]
6. limits to arbitrage, if traders are agents for other principals--e.g., Shleifer and Vishny (1997)[99]

Perspectives on collective intelligence

One way to integrate the different perspectives on collective intelligence is to identify the:

1. commonalities: how each maps onto the generic components (see Components) common in collective intelligence, and
2. discipline-specific contributions: what are the key relevant concepts and theories.

Hopefully, the commonalities spur deeper thinking about the foundations of collective intelligence, and the discipline-specific contributions suggest cross-fertilization of ideas across disciplines.

The following table provides examples from each perspective:

Perspective	Actors	Resources	Actions and their bases	Results	Evaluation and measures	Factors enhancing CI	Factors inhibiting CI	Techniques for enhancing CI
Sociology	Humans	Status, power	Persuasion, conformity	Motivation, learning, conflict	Experiments, sociometry	-	-	-
Economics	Firms,	Wealth, goods	Trade, pricing,	Efficiency,	Welfare,	-	Arrow's	-

	consumers		incentives	monopolistic power	monetary units		impossibility theorems, information asymmetry, contractual incompleteness, Prisoner's dilemma	
Cognitive neuroscience	Neurons, codons	Weight of connections	Neural firing	Mental representations, mind-body problem	Two-photon microscopy, calcium imaging	-	-	-
Organization theory	People	Resources	Resource allocation	Productivity	Productivity	-	Differentiation and integration	-
Computer science and artificial intelligence	Computers	Computational cycles, storage	Computation	Computational solutions	Speed and accuracy of results, transmissions	-	Byzantine Failure, impossibility of distributed consensus	-
Biology	Organisms	Energy, free will, sensory capacity	Decision-making	-	-	-	-	-
Political philosophy	Humans	Status, power	Decision-making	Motivation, learning, conflict	Experiments, sociometry	-	-	-

In addition to the above discipline-specific perspectives, some other perspectives have multi-disciplinary origins, but are becoming perspectives in their own right, and we list them here:

Perspective	**Actors**	**Resources**	**Actions and their bases**	**Results**	**Evaluation and measures**	**Factors enhancing CI**	**Factors inhibiting CI**	**Techniques for enhancing CI**
Social psychology	Humans	Status, power	Persuasion, conformity	Motivation, learning, conflict	Experiments, sociometry	-	-	-
Computer supported collaborative work	Humans	Productivity	-	-	-	-	-	-
Prediction markets	Humans	Information, money	Votes	Voting result	Count of votes, variances	-	-	-
Economic sociology	Humans and organizations	Trust, identity, organizational demography	Relationships as contraints	Node-level profitability, mortality	Profits, hazard rates	Diversity	Strong ties	Brokerage over weak ties

Techniques for enhancing collective intelligence

We list some factors drawn from different perspectives, as a working list. See the different perspectives for a complete listing.

Carefully decomposing work and distributing expertise

Grinter et al. (1999) [100] examine 4 methods product development organizations in Lucent use to integrate multi-site teams (all of which suffer 2 problems: consequences of unequal distribution of project mass, and finding expertise):

- locate each expertise at only one site. The benefits include having scale with a pool of expertise, better load balancing and development of that expertise. The cost is that we now have to manage projects across sites.
- partition product development according to product structure, and locate the components at different sites. The benefit is the independence in operating environments. The cost is in integration testing of the components.
- partition process steps. The benefit is closer proximity to customers. There is also better use of resources, such as test labs. The cost is in managing temporal dependencies and in handoffs.
- customize products that are produced out of one baseline model made at one central site. This also enhances proximity to customers. There is also good division of labor for code ownership. The cost is the need to build trust across product makers, and compatibility issues, since the customizing sites tends to need tools that match those at the baseline site. There is also the need for coordinating processes; and the authors find a lack of documentation that hampers this process.

Aligning interests

How do large groups overcome the problem of collective action. One way is to be smaller, such as in corporate spin-offs. See details in the section on the logic of collective action.

Enhancing social capital and networks

Another approach is from Gould (1995) [101], who argues that dense housing and social networks in Paris result in more individuals being co-opted into the French Revolution, even if these individuals might have, on their own, are particularly enthusiastic about being revolutionaries.

A different way is to highlight the privilege of membership, as in Sierra Club calendar you get for joining the organization or the inevitably canvas tote bag that public television hawks during its seemingly interminable pledge drives (see these from University of Chicago [102]).

Enhancing social interaction

Kollock (1998) [103] and Lee et al. (2001) [104] are examples of studies in how we can increase the level of social interaction in online communities.

Girgensohn and Lee (2002) [105] more recently describes the features used in the CHIplace (website for 2002 ACM CHI conference) and Portkey (website for IBM interns) websites:

- Establishing common ground: folklore and trivia about CHI's history, pre-publication of research papers for discussion, tips provided by predecessor interns, photos provided by current interns
- Increasing awareness: listing of what has changed, chronological display of threads, selection of today's active forums
- Enhancing social interaction mechanisms: polls, writing and voting of trivia, discussion forums
- Making "place"; see Harrison and Dourish (1996) [106] and Kollock (1998) [107]: discussion forums which become places for discussants to set norms about what should be posted

See also

The Cooperation Commons, http://www.cooperationcommons.com/resources

Libert and Spector (2007), We are Smarter than Me [108]

Fortune's Imeme Conference http://www.timeinc.net/fortune/conferences/imeme/imeme_home.html

Community 2.0 Conference http://www.community2-0con.com/

Gloor (2006), Swarm Creativity : Competitive Advantage through Collaborative Innovation Networks [109]

Gloor P, Cooper S (2006), The New Principles of a Swarm Business [110]