

Collective Intelligence in Organizations: Tools and Studies

Introduction

Antonietta Grasso¹ & Gregorio Convertino²

¹Xerox Research Centre Europe, 38240 Grenoble, France (E-mail: antonietta.grasso@xrce.xerox.com); ²Palo Alto Research Center, 3333 Coyote Hill Rd, Palo Alto, CA 94304, USA (E-mail: Gregorio.Convertino@parc.com)

"Metaphorically, I see the augmented organization or institution of the future as changing, not as an organism merely to be a bigger and faster snail, but to achieve such new levels of sensory capability, speed, power and coordination as to become a new species – a cat."

(Engelbart D.C., The Augmented Knowledge Workshop, 1984)

A new generation of technologies is entering organizations. This follows the success of a new class of tools that have emerged in the consumer domain and the public web (i.e., Web 2.0 tools) (O'Reilly 2007). Crowds of users can now perform new functions collectively: they can create and leverage networks of friends or professional peers (using social networking platforms), create and share knowledge (using tools such as wikis, blogs, sites for sharing photos, videos, tags, or bookmarks, forums, and microblogging sites), and perform complex tasks collectively (using tools such as argumentation systems; crowdsourcing platforms; and prediction markets). Overall, these tools and the collaborative processes that they enable have been referred to as examples of Collective Intelligence (CI), which can be broadly defined as "the capacity of human collectives to engage in intellectual cooperation in order to create, innovate and invent" (Levy 2010, p.1).

Recent research on early successes of CI 'in the wild' (i.e., that emerged on the Internet) has started to highlight what preconditions, in terms of both technology and people, should be in place for such collective intelligent behaviors to emerge (e.g., see research on Wikipedia, online communities, open source communities, and crowdsourcing platforms). The growing attention of the research community to this broad area is also indicated by new academic centers and publication venues around CI: for example, the centers at MIT (<http://cci.mit.edu>, funded in

2006) and the University of Ottawa (Levy 2010), international workshops on topics such as crowdsourcing or CI in organizations (see next section), and the “Collective Intelligence” conference launched in 2012.

Below, we introduce the area of CI in Organizations (CIorg) by giving some context in terms of publication venues that led to this special issue and prior foundational literature. We outline the current research agenda on CIorg and, finally, we introduce the four papers of the special issue as exemplars of contributions in this area.

1. The research context of *CI in organizations* (CIorg)

This special issue reports on CI tools and processes that are specific to organizations, either commercial or not (e.g., business enterprises, government or military institutions, or local government and civic organizations), hereon referred to as CI in Organizations. This publication emerges as a follow-up to two past workshops on CI in Organizations, at two international conferences, ACM CSCW 2010 (Savannah, GA, USA) and ACM Group 2010 (Sanibel Island, FL, USA) (Convertino et al. 2010). In parallel with the preparation of this issue, two other workshops have been organized, targeted towards large-scale deliberation tools and processes, at the ACM CSCW 2012 (Seattle, WA, USA) and COOP 2012 (Marseille, France) conferences. These workshops have been occasions to bring together leading researchers and designers who are studying or developing CI tools aimed at people in organizations. Specifically, the first two workshops have prepared the ground for this special issue by discussing the state of the art and a research agenda for CI in Organizations. During the discussions it became apparent that new tools for networking, sharing knowledge, or make collective decisions are entering organizations, mostly as part of a process of diffusion of Web 2.0 tools from the consumer space to the organizational space (see the phenomenon of ‘Consumerization’ of IT for organizations in Moschella et al. 2001). However, *the research field of Social and Collaborative Computing is still lacking a body of studies and tools of reference that can guide the design of future CI tools for organizations* (Convertino et al. 2010).

Before characterizing the narrower area of CIorg, we first ground it in the prior literature on CI. The concept of CI precedes information technology. For example, in psychology Wechsler (1971) defined CI by differentiating it from mere collective behavior or pooling of common resources toward a task: according to Wechsler, CI occurs only when there is a cross-fertilization among those participating in the shared task. Moreover, even the vision of enabling CI through collaborative technology is not new. Early research on collaborative computing or groupware during the Seventies had already announced the possibility - at the time still visionary - of allowing large groups of dislocated people to carry out complex tasks by collaborating and coordinating each others’

activities (Johnson-Lenz and Trudy 1980). Engelbart, who had anticipated the need for CIorg research, had proposed the vision of enabling high-performance organizations through the augmentation of groupware technologies (Engelbart 1992). Although these pre-Web 2.0 visions were in line with the current concept of CI, they were limited to technologies for small groups, rather than larger collectives, and were often confined to work settings. Over the last decade, after the diffusion of Web 2.0 tools, CI has indeed become a reality at unprecedented scales and in various life settings. New research and definitions of CI have appeared. Some define CI as the study of intelligence that emerges when multiple people share ideas, interact, collaborate or compete in the same environment (Malone et al. 2009). Others, from a more technocratic perspective, define CI as the capability to engineer mechanisms and tools that can empower people's thinking and actions. To this end, CI research explores how people's interactions can be re-interpreted and used by other people or machines to infer and mine data, suggest or improve people's choices, change or enhance collective actions (Alag 2008).

CIorg emerges in this context of Web 2.0 tools and studies about them. These tools possess characteristics that are attractive to companies as a means to support large-scale collaboration and CI in organizations. Extending the analysis of Web 2.0 tools on the public Internet by O'Reilly (2007), McAfee examined specific opportunities and challenges that Web 2.0 tools can create for the enterprise (see his neologism Enterprise2.0). He argues that they could be appropriate to support work processes because of three features these tools present: easy communication and interaction, lack of imposed structure, and mechanisms to let the structure emerge (McAfee 2009). The growing interest from the enterprise is witnessed by the growing number of Web 2.0 tools that are now part of commercial offerings for organizations and the increasing research into these technologies and their use. Early examples include tools such as IBM's social bookmarking software Dogear (Dugan et al. 2007), and wiki platforms such as Wikispaces or TikiWiki. Emulating the success of social network systems for consumers such as Facebook or LinkedIn, other examples include new tools such as IBM's Social Blue (or Beehive, which was the research prototype) (DiMicco et al. 2009), Novell's Pulse, or Yammer. Moreover, several commercial web platforms targeted at enterprises have been extended to include social functions that can be considered Web 2.0 tools, such as IBM's LotusLive, Microsoft's SharePoint, Jive's Social Business Software, and SAP's NewWeaver.

On the basis of this first wave of adoption some studies are trying to assess the interest and benefits of them in a more systematic way. Some of these are broad industry assessments based on interviews, where key stakeholders are asked about objectives, potential and already achieved results. For example, Andriole has interviewed 15 senior managers "to measure the impact and business value of wikis, blogs, podcasts, folksonomies, mash-ups, social networks, virtual worlds,

crowdsourcing and RSS filters” (Andriole 2010). This study reports that there are still current key concerns around these technologies that refer to security, intellectual property leakage, privacy and control of the information and, more importantly, it shows that companies have not yet a well-formed way to integrate these tools with their business objectives and assess their value in order to deploy them strategically in the organizations. While this was a horizontal cross-industry and cross-technology study, other investigations aim at a more vertical understanding, typically focusing on just one of the possible Web 2.0 technologies and on a particular work setting. For example, Holtzblatt and colleagues (Holtzblatt et al. 2010) studied the deployments of wiki platforms inside the MITRE organization, reporting various issues that currently affect the full exploitation of these tools. This study clearly indicates that two elements would be required to successfully adopt these tools in the organization: clarity about the organizational purpose, and organizationally defined governance policies which are still flexible and adaptable by the users. The Web 2.0 tools, due to their low cost, ease of use, and wide adoption outside the enterprise, are often introduced in work organizations lacking precisely these two elements.

2. Research agenda for CIorg

Several of the core problems that the CIorg agenda needs to address are not new to the broader Social and Collaborative Computing (or CSCW) community. Already in 1992, a study of Lotus Notes illustrated multiple barriers to introducing collaborative technology in support of knowledge sharing and reuse in a large work organization (Orlikowski 1992). Orlikowski (1992) argues that changes in structural properties are generally required in an organization where the premises underlying the design of a newly introduced CSCW system are incongruent with those of the organization's culture, policies, and reward systems. But, while such changes are difficult to accomplish, and usually encounter some resistance, without these changes the deployment of the technology will most likely fail because the unchanged structural elements of the organization will act as barriers. This is, we believe, a foundational observation not only for traditional CSCW systems, but also for Collective Intelligence tools in organizations. This aspect is exacerbated with Web 2.0 tools, due to the hype that surrounds them.

We argue that the existing CSCW approach, grounding technology design on the basis of a deep understanding of the work practices, is beneficial for the organizational situation, where the objectives and changes required are not sufficiently understood prior to the adaptation and deployment of the tools. In fact, organizations are likely to pose distinctive requirements and constraints. For example, the experience of the Web has shown that letting behavior emerge is a winning strategy when a critical mass of users is in place that can more “naturally” allow order, as well as quality control, to materialize from the bottom.

Instead, organizations such as business enterprises, government, educational, or military institutions have structures that are different from the Web: the scale is different, a top-down control structure is already in place, and the employees have specific motivations, skills, and duties.

As a new context for Web 2.0 tools, enterprises such as private organizations or public institutions have distinctive properties (Convertino et al. 2010):

People and roles. The people using technology are workers, who are assigned roles and perform specific and more or less stable jobs. Generally, they know each other (even indirectly), are paid to work, and are part of a reporting structure. The contributions of the individuals are monitored and evaluated, while the job roles and reporting structure are the means for the organization to distribute responsibilities, keep accountability, and monitor the performance. Moreover, in the organization, there is a growing emphasis on utilizing teams, task forces, and communities of practice as a strategy to improve performance (see the first paper in this issue by Matthews et al.). Such work units are not formed on the basis of personal interest, but are assembled by the management, given the functions needed and the experts available. Typically, the workers have to coordinate with (and rely on) others for the organization to be productive as a whole.

Tools and tasks. Both the tasks performed and the tools used (email, phone, content editing tools, databases) are non-discretionary. Typically, they are assigned by the management and not chosen by the workers. Also distinctive is that the tasks are increasingly information and interaction intensive and workers in enterprises increasingly need to make sense of large amounts of information from multiple channels or information tools. This existing ecology of systems creates barriers in fitting in new tools (e.g., Tolmie 2010), for which often there is a conflict between who gets the benefits and who bears the costs of using the tools (e.g., Grudin 1994).

Goodness criteria. The criteria used to evaluate and predict the success of technology are also different. In the consumer space, these are mainly the utility to the user, quality of user experience (e.g., simplicity and fun), and social benefits. In organizations, the key criteria over and above those for the consumer space are the worker's productivity (e.g., worker's output and workload), the organization's productivity (i.e., ROI) and compliance to internal and external SLAs, political returns, security, and compatibility with the legacy infrastructure (e.g., Andriole 2010). Therefore additional constraints for Web 2.0 tools in organizations include compatibility with prior tools, security, costs of maintaining new tools, and the fit with the current management structure, work practices, and motivational mechanisms.

Starting from these observations, we developed a set of research questions that we used to frame the contributions and the discussions held at the various above-mentioned workshops (Convertino et al. 2010):

- What defines the forms of large-scale collaboration that emerge in specific organizations? How are they different from other forms that are observed in the consumer Web and what are the business opportunities and the work requirements they need to subscribe to?
- What are the organizational processes that are best suited to bottom-up emerging collaboration?
What is the degree of domain modeling that the tools need to present to be adopted effectively?
- What are available traces from previous activities and how they can be exploited for the current activity and to organize the dynamic knowledge being created?
- What visualizations and abstractions can help to monitor and make sense of the activities of others?
- How do quality, customer input, and timing affect work outcomes in organizations?
How do factors such as trust, motivation, attribution, and traceability affect information and activity flows in organizations? How can these factors be ‘designed in’ Web 2.0 tools?
- What mix of research methods, such as field studies and logs analysis, are suitable for CI research and design?

At the time of the workshops that motivated this special issue, when we framed the research questions listed above, our intent was to assess the state of the art in research and design about this structured set of problems. Then, at the stage of publishing this special issue, we have learned that most of these questions remain only partially answered, while we believe that the set of questions remains generally useful to characterize the object of study of CIorg. In the rest of this introduction we present the papers, and some first attempts to have a more systematic body of answers to the questions. In particular, we will highlight how we learned from the first two selected contributions the ways in which organizational complexity is growing along various dimensions such as process or ensemble of tools. This is particularly related to the first questions about the forms (and related issues) of large-scale collaboration that emerge in the organizations. On the other hand, from the two latter contributions we learned that emerging technologies such as Natural Language Processing, Machine Learning, and mixed initiative interfaces are becoming mature and can effectively support the refinement of the functionalities of Web 2.0 to help manage the above-mentioned growing complexity and requirements of an organizational setting. This is particularly relevant to the questions about domain modeling and trace reuse feasibility and benefit.

3. CIorg contributions in this special issue

3.1. CIorg studies pointing to emerging needs in modern organizations

If we consider the combination of a global market and a knowledge economy as a highly competitive environment, we can also view business organizations as organisms that are forced to become more “intelligent” to survive in order to increase their fitness in such a competitive environment. A general way for a company to be more collectively effective is to become more adaptive and flexible in the work arrangements of its employees when they collaborate. For example, recent research points to important changes in the way current organizations divide and manage labor (Malone et al. 2011).

In the first paper of this special issue, **Matthews and colleagues** study the phenomenon of dynamic teams and multi-teaming. That is, workers are asked to simultaneously participate in multiple workgroups and they join and leave them ‘as needed’. While the practice itself has already been documented, prior research has not yet analyzed the implications of this growing practice on the behavior and workload of individuals: i.e., how individuals manage the additional complexity of multi-teaming, how fundamental team processes such as goal setting and group maintenance are achieved in such a new dynamic setup for cooperative work in organizations.

In their field study within a large technology company, Matthews and colleagues systematically characterize the types of collaboration that are currently prevalent and the interrelations among these. The company adopted high level of matrix management and the surveyed participants managed an average of 16 collaborations. In this context, they observed six distinctive types of collaboration (or collaboration profiles). More importantly, they observed an interesting adaptive behavior of the individuals: as they multi-task and take part in multiple collaborations of different types, they also leverage productive relationships or links among their collaborations. For example, the participation in a community of practice (secondary collaboration) is exploited to support the processes of recruiting and group maintenance within dynamic project teams (mission critical collaborations) in which the same individuals are involved. Therefore, in this case study, developing productive interrelationships among collaboration represents a new adaptive behavior for the individuals, while their organization develops new adaptive arrangements of collaborative work through dynamic teams and multi-teaming. This type of contribution, in defining common types of collaborations and productive interrelation among them, informs the design of both technology and work process because it helps to show in what ways individual collaborators can become more efficient and the organization more effective.

The second paper of this special issue by **Rooksby and Sommerville** makes a related contribution. Although the reported study is about a different class of

organization, a government department, it characterizes the challenges in using various social networking technologies in such an organization.

There have been a few studies of social network site use within organizations, but most have studied the use of private internal sites in large technology corporations such as IBM, HP, and Microsoft. Very little is known about the use of social network sites in public organizations or governmental institutions. Moreover, the majority of these studies, by focusing on the use of a single technology, tend to miss the wider context of use: people in real settings work and communicate across multiple sites, which include internal and public software platforms. The challenge of managing multiple existing and emerging social networking sites is relevant to most modern organizations (see 'Consumerization' of IT for organizations in Moschella et al. 2004). This is perhaps exacerbated in institutions that must interact regularly with a large population of clients, such as the citizens of an entire country, because the multiple sites of relevance may vary over time, among age cohorts, or across geographic regions.

In their paper, Rooksby and Sommerville, after acknowledging the opportunities that Web 2.0 technologies such as social network sites offer for improving the workings of government organizations, focus on the above-mentioned challenge and give a broad, factual account of what is happening around the use of social networking sites in a specific government organization, the UK Home Office. This is a large government department that, as part of the UK Civil Service, employs 24000 people dispersed over multiple locations and performs functions such as identity and passport services or crime reduction.

They describe the various uses of seven social networking sites: two internal or bounded within the intranet and five public or hosted by commercial providers (Facebook, Twitter, Bebo, Habbo, and LinkedIn). They observed that the users had to perform their regular work functions while, among the multiple sites, there were sometimes conflicts or overlaps in the functionality, differences in the level of control allowed to individual groups, and new restrictive policies about internal or public-facing uses of specific sites (e.g., Twitter used to broadcast messages but not to reply to citizens' messages).

As qualitative findings drawn from their observations, the authors isolated five key problems that the studied organization faced in managing the multiple sites: boundary problems (internal-external, work-personal uses), limited control (quite limited overall and varying across sites), limited visibility (difficult for managers to monitor how the sites are used), overlaps in the functionality (among social networking sites or between these and other technologies), and ongoing change (changes in site popularity, practices, and platforms). Based on these findings, the

authors draw some implications for CIorg design and research. They invite the developers of CIorg tools to look beyond the design of individual technologies and address the complexity of the broader ‘ecosystem’ of several independent technologies that current organizations must manage. To make organizations more adaptive, the field of CIorg needs to develop appropriate strategies for helping organizations in managing their specific ecosystem of social technologies and for helping individual workers to choose and use the right technology.

A first generalization can be attempted based on the first two papers. Modern organizations impose additional complexity that individual workers need to manage; in some cases this is at the cost of making the organization more adaptive, in others it is just a cost imposed by unregulated diffusion of technologies in the organization. Despite the differences in the domains and tools looked at, the first two papers characterize two types of complexity that workers need to manage: managing interrelationships among their multiple ongoing collaborations (see the first paper) and managing multiple social sites or tools while efficiently performing the same set of core functions for the organization (see the second paper). Interestingly, in both papers the authors adopt the metaphor of managing a complex ‘ecosystem’. An implication of these findings is that by better understanding the current challenges for individual workers, designers of CIorg tools can indeed aim at improving the ecosystem.

3.2. CIorg prototypes responding to emerging needs in modern organizations

The third and fourth papers of this issue describe exemplars of prototype systems that support CIorg in specific settings for collaboration. Both papers propose novel support for idea management in order to augment the current functions of collective discussion and deliberation: i.e., the sensemaking and discourse undertaken by a community of practice (or a task force) in analyzing a corpus of documents in an organization; a large online community doing deliberation to solve a complex problem.

The third paper by **De Liddo and collaborators** contains the important contribution of connecting prior literature on sensemaking and computer-supported discourse to the current problem of building tools for CIorg. Based on this explicit connection, they propose a platform to support CIorg in specific conditions for collaboration in an organization, where a document corpus needs to be analyzed, while divergent perspectives, inconsistencies, ambiguity and disagreement create debates around this analysis. Specifically, De Liddo and collaborators first specialize the concept of CI for this narrower context by calling it Contested Collective Intelligence (CCI). They define it as an emergent capability that depends on being able to pool and connect people’s interpretations, comments and debates around issues, where these are often anchored in diverse documents. For example, a task force of analysts in an organization is put

in charge of analyzing a large collection of documents in order to define the state of the art in research and development in a particular domain or the main business trends in a particular market of interest.

Then they focus on the problem of designing infrastructure (i.e., a software platform) that can augment this organizational capability. As a solution, the authors propose a platform that can support CCI by providing an environment that combines automatic annotation and human annotation functions. Ultimately, by relating and sharing interpretations they generate an interactive representation of the social discourse in the organization. In the last part of this paper, the authors present the system in action through the results of a case study that involved analyzing a corpus of 125 annotated documents, using human and machine annotations.

Beyond the specific prototype, this paper also contributes a design vision for how future CIorg tools could augment the discourse and reasoning abilities of an organization. In current organizations, collective discourse and its narratives are woven not only using conventional documents and interpersonal interactions but also through digital archives and social media. However, current platforms provide only simple quantitative analytics about the organization-wide discourse (e.g. online participation levels, social networks). Instead, the vision proposed in this paper is about a CIorg platform that by adding new layers of metadata allows the display and management of higher-order aspects such as the quality of the discourse, the intellectual structure of an emerging topic, what are the key antagonist or controversial ideas, how the social network relates to a proposed course of action, how different stakeholders frame problems, what are the claimed gaps in understanding, what are the assumptions being questioned, and what diverse styles of reasoning are being adopted (e.g., technical, commercial, political).

The fourth paper by **Klein** is informed by prior research on argumentation systems. The class of systems being discussed belongs to a new generation of deliberation tools that have emerged in the Web 2.0 era. These deliberation technologies are not bounded anymore to supporting co-located groups, as often happened for group decision-support systems, but can now enable collective deliberations in large distributed communities to inform organizational decision making on a global scale. These technologies can empower organizations facing highly complex and controversial challenges – including issues such as climate change, the spread of disease, international security, scientific collaborations, product development, and so on – to rely on new tools that allow large numbers of experts and stakeholders to deliberate collectively on a global scale.

The paper makes two valuable contributions. First, the paper contains a thorough analysis of the various strengths and limitations of Web 2.0 technologies

when used in support of large-scale deliberation. Klein describes how Web 2.0 tools, while succeeding in allowing interactions on a massive scale, still have not realized their potential in allowing people to deliberate effectively at large scales. The author starts from the main strengths or opportunities that Web 2.0 systems introduce as deliberation tools, such as the phenomena of idea synergy and many eyes. Then, based on anecdotal evidence from actual systems, the author analyzes the strengths and limitations of three categories of Web 2.0 systems that differ in how they structure content: time-centric (e.g., web-forums), topic-centric (e.g., wikis), and argument-centric (e.g., argumentation systems) tools. In the context of this analysis, Klein argues that the argument-centric approach represents a more promising system design to address the known weaknesses of current Web 2.0 tools, when these are used as deliberation technologies. He bases this thesis on both simulations and the direct experience from various deployments and evaluations of the Deliberatorium, a web-based system that implemented an argument-centric approach (e.g., Klein and Iandoli 2008).

The other and perhaps most central contribution of this paper is the design of metrics to enable attention mediation in large-scale deliberations. The problem addressed is the challenge of attention allocation, which is a major obstacle in current deliberation systems. The metrics are proposed as an enhancement of a large-scale argumentation system, such as the Deliberatorium: they consist of automated algorithms that provide personalized suggestions on what part of the argument map the user can contribute to, and why. The system detects exceptions or deviations with respect to a pre-specified normative model of the deliberation process (goals to be achieved, the steps for achieving these goals). Participants are notified based on the model that the system builds of their roles and interests about the exceptions. If the suggestions are relevant and taken into account, the emergent effect of the metrics should be that the collective intelligence of the community is increased (or better exploited) because each user contributes where they can do the most good.

A second generalization can be attempted based on the last two papers. That is, new opportunities for CIorg are uncovered by building platforms that extend the functionalities of Web 2.0 tools with those of new tools. For example, natural language processing techniques for semi-automatically annotating large amounts of content, or machine learning techniques for modeling (people, tasks, and goodness criteria) and making recommendations. The design proposed in both papers also suggests that mixed initiative interfaces can allow the human (individual or collective users) to effectively coordinate between (traditionally open-ended and unstructured) Web 2.0 tools and machine intelligence (or algorithms). These augmentations of Web 2.0 can support new organizational activities such as sensemaking and discussion by distributed task forces that analyze large corpora of documents (see the third paper), or large-scale

deliberation by distributed communities tackling complex problems (see the fourth paper).

In summary, in this special issue researchers and designers can find an introduction to the new area of CI in organizations and four exemplary contributions on CIorg. Clearly, research and design on CIorg are still at a very early stage of development. This issue represents a small step in the direction of establishing a multidisciplinary science of CI in organizations: i.e., pointing to prior foundational work, proposing a research agenda, and presenting exemplary contributions on CIorg.

1. References

- Alag, Satnam (2008). *Collective Intelligence in Action*. Manning Publications Co.
- Andriole, Stephen J. (2010). Business impact of Web 2.0 technologies. *Communications of ACM*, vol. 53, no. 12, pp. 67–79.
- Convertino, G.; Grasso, A.; Millen, D.R.; De Michelis, G.; Chi, E. H. (2010). Clorg: collective intelligence in organizations. In Proceedings of the 16th ACM international conference on Supporting group work (GROUP '10). ACM, New York, NY, USA, 355–358. <http://www.parc.com/ciorg>.
- DiMicco, Joan M., Werner Geyer, Casey Dugan, Beth Brownholtz, and David R. Millen. (2009). People Sensemaking and Relationship Building on an Enterprise Social Networking Site. In *Proceedings of the 42nd Hawaii International Conference on System Sciences (HICSS '09)*, pp. 1–10.
- Dugan, Casey, Michael J. Muller, David R. Millen, Geyer Werner, Beth Brownholtz and Marty Moore (2007). The Dogear Game: A Social Bookmark Recommender System. In *Proceedings of GROUP 2007. Conference on Supporting Group Work, Sanibel Island, Florida, USA, November 4–7, 2007*. New York: ACM Press, pp. 387–390.
- Engelbart, Douglas C. (1992). Toward High-Performance Organizations: A Strategic Role for Groupware. Bootstrap Institute. June 1992 (AUGMENT,132811). In *Proceedings of the GroupWare '92 Conference, San Jose, CA, Aug 3–5*. Morgan Kaufmann Publishers. <http://www.dougelbart.org/pubs/augment-132811.html>
- Grudin, Jonathan (1994). Groupware and Social Dynamics: Eight Challenges for Developers. *Communications of the ACM*, vol. 37, no. 1, pp. 92–115.
- Holtzblatt, Lester J., Damianos, Laurie, and Daniel Weiss (2010). Factors impeding wiki use in the enterprise: A case study. In *Proceedings of the Collective Intelligence Workshop in conjunction with CSCW 2010, 6 February 2010, Savannah, Georgia, USA*. <http://www.parc.com/ciorg>. Accessed 4 April 2012.
- Johnson-Lenz, Peter and Trudy (1980). Groupware: The Emerging Art of Orchestrating Collective Intelligence. In *Proceedings of the First Global Conference on the Future, Toronto, Canada, 20–24 July 1980*.
- Klein, Mark and L. Iandoli (2008). Supporting Collaborative Deliberation Using a Large-Scale Argumentation System: The MIT Collaboratorium, in Directions and Implications of Advanced Computing. In *Proceedings of the Conference on Online Deliberation (DIAC-2008/OD2008)*. Berkeley, California, USA, 26–29 June 2008.

- Levy, Pierre (2010). From social computing to reflexive collective intelligence: The IEML research program. *Information Sciences*, vol. 180, no. 1–2, pp. 71–94.
- Malone, Thomas W., Robert Laubacher, and Chrysanthos Dellarocas (2009). *Harnessing Crowds: Mapping the Genome of Collective Intelligence*. MIT Sloan Research Paper No. 4732–09 (February 3, 2009). MIT Center for Collective Intelligence. Available at SSRN: <http://ssrn.com/abstract=1381502>. Accessed 4 April 2012.
- Malone, Thomas W., Laubacher R.J., Johns T. (2011) The Big Idea: The Age of Hyperspecialization, *Harvard Business Review*, 89 (7/8): 56–65, July/August 2011
- McAfee, Andrew (2009). *Enterprise 2.0: New collaborative tools for your organization's toughest challenges*. Harvard Business Press. Boston, MA.
- David Moschella, Douglas Neal, Piet Opperman, and John Taylor (2004). *The 'Consumerization' of Information Technology*. Position Paper at Computer Sciences Corporation, UK.
- O'Reilly, Tim, (2007). What is Web 2.0: Design Patterns and Business Models for the Next Generation of Software. *Communications & Strategies*, vol. 65, no. 1, p. 17.
- Orlikowski, Wanda J. (1992) Learning from Notes: Organizational Issues in Groupware Implementation. In *Proceedings of CSCW '92, the ACM conference on Computer-supported cooperative work, 31 October – 4 November, 1992, Toronto, Canada*. ACM Press, New York, NY, USA, pp.362-369.
- Tolmie, Peter (2010) 'I'll talk to someone': The work of bid management as a test domain for CI. In *Proceedings of the Collective Intelligence in Organizations (CIorg): Tools and Studies Workshop in conjunction with Collective Intelligence ACM Group 2010, 7–10 November 2010, Sanibel Island, FL, USA*.
- Wechsler, David (1971). Concept of collective intelligence. *American Psychologist*, vol. 26, no. 10, pp. 904–907.