Beyond the Human-Computation Metaphor

Yiftach Nagar

MIT Center for Collective Intelligence

MIT Center for Collective Intelligence Working Paper No. 2011-5

October 2011

MIT Center for Collective Intelligence
Massachusetts Institute of Technology
http://cci.mit.edu

2011 IEEE International Conference on Privacy, Security, Risk, and Trust, and IEEE International Conference on Social Computing

Beyond the Human-Computation Metaphor

Yiftach Nagar
MIT Center for Collective Intelligence and Sloan School of Management
Massachusetts Institute of Technology
Cambridge, USA
ynagar @ mit.edu

Abstract— **Two assumptions have become dominant in the field of social computing and crowdsourcing - the computational view, and the assumption of a human-only crowd. In this paper, I address those assumptions. I trace their origin in the human-computation metaphor, and argue that while this metaphor is instrumental in facilitating novel developments, it also constrains the thinking of designers. I discuss some of the limitations this metaphor might impose, and offer that additional perspectives, such as an organizational design perspective and the distributed cognition perspective can help us think of novel possibilities of organizing work with crowdsourcing. I call for extending the conversation among computer-scientists and organizational researchers, and propose that the metaphor of 'information processing' might serve as a 'boundary-object' around which the dialogue among these communities can thrive.**

Keywords-human-computation; crowdsourcing; collective-intelligence; human-computer interaction; computer-supported-collaborative-work

I. INTRODUCTION

In recent years people have started using the web to 'harness the wisdom of crowds', through various crowdsourcing endeavors. Notable examples include the recent success of a team from the MIT Media Lab to solve a challenge that was initiated by DARPA, and locate 10 red balloons that were dispatched by DARPA in different locations over the US, within hours [1]; Soylent – a word-processor plugin that lets document authors crowdsource proof-reading, word-smithing and paragraph-shortening [2]; crowdsourcing product comparisons and the creation of technical articles using CrowdForge [3]; and reCAPTCHA [4], in which people help machines recognize about a hundred million [1] difficult-to-decipher words and characters in scanned texts, to name but a few.

Indeed, these novel examples, as well as many others, have rightfully captured the attention of researchers, designers and practitioners interested in crowdsourcing and social computing, as each one of them demonstrated powerfully how human intelligence can be utilized in the framework of computational processes to perform tasks that artificial-intelligence has yet to do successfully, and how novel ways of connecting people – even lay people – can yield surprisingly high-quality outcomes. The role of software in these cases was mostly limited to facilitating task provisioning and distribution, and collating of input from people.

II. THE HUMAN COMPUTATION METAPHOR AS AN ENABLER OF INNOVATIVE DESIGN

What do these examples, and many other, have in common, beyond using crowds? Underlying all these examples is the conceptual metaphor of 'Human Computation'. While the notion of *human computation* dates back to the 19th century [5, 6], Quinn and Bederson [6] trace the modern incarnation and usage of the term to Luis von Ahn's doctoral dissertation [7] which bears that title, and they document the rise in its dominance and effect on different fields and disciplines.

In the human computation metaphor, as captured in various works of von Ahn, as well as many others (See [6] for a non-exhaustive list of definitions), humans are seen as agents, much like computer agents, that can perform certain well-defined, highly structured tasks. The role of humans is to perform certain subroutines, or 'function calls', as part of a larger algorithm that may be carried out by another human, or in a computer program.

Indeed, von Ahn posited human computation as "…a paradigm for utilizing human processing power to solve problems that computers cannot yet solve". 'Paradigm' is a bit of a loaded term, so I here follow Morgan's [8] deconstruction of Kuhn's use of the term, into paradigm, metaphor and puzzle-solving activities. Following that deconstruction, 'human computation' can be viewed as a (conceptual) metaphor, which makes an analogy between the activities in the process, and the computations carried in electronic computers.

Conceptual metaphors are powerful. As Lakoff and Johnson [9] have argued, human thought processes at large and most of our ordinary conceptual system are metaphorical in nature. Metaphors turn our imagination in ways that forge an equivalence or identity between separate elements of experience [10], and help us orient our thinking and develop insights. Accordingly, researchers have devoted considerable attention to the role metaphors play in the development of science and social thought. As Morgan [8] notes, *"The research work of these different theorists contributes to a view of scientific inquiry as a creative process in which scientists view the world metaphorically, through the language and concepts*

[1] Reported by Luis von Ahn during a talk at Carnegie Mellon, April 3, 2011. Also see http://www.google.com/recaptcha/faq (accessed 2011-06-20)

978-0-7695-4578-3/11 $26.00 © 2011 IEEE
DOI

which filter and structure their perceptions of their subject of study and through the specific metaphors which they implicitly or explicitly choose to develop their framework for analysis." This realization is, of course, true seven times over when it comes to *design* sciences, like HCI and related fields which aim not only to understand how the social world works, but to actively construct systems that interact with, and shape, social reality.

Morgan further notes that the creative potential of metaphor *"depends upon there being a degree of difference between the subjects involved in the metaphorical process… Metaphor is thus based upon but partial truth; it requires of its user a somewhat one-sided abstraction in which certain features are emphasized and others suppressed in a selective comparison."*

The human computation metaphor indeed relies on a couple of abstractions. In making the analogy between humans and electronic computers, it strips our notion of computers of the physical implementation of silicon, electronic circuitry, etc., as well as off of the implementation of algorithms using software, thus focusing our attention on the computation that needs to take place. Computation, therefore, is viewed as an abstract activity of manipulation of symbols or representations. And it strips humans of many of their properties as social animals, focusing our attention on some (but not all) of their cognitive abilities.

Relying on such abstractions, the human computation metaphor proved powerful indeed, unleashing creative designs of the kind I mentioned above. It carried with it associated terms and induced additional metaphors from the world of software design such as the notions of *transaction, scalability, parallel processing, iterations,* etc., used (some explicitly and some implicitly) by the designers of TurkIt [11, 12]; and the notion of *software design patterns* [13]. For example, the designers of CrowdForge [3] borrowed an established software design pattern, *MapReduce*, to organize the work of people, while the designers of *Soylent* [2] came up with a new design pattern, *Find-Fix-Verify*. All these draw directly from the metaphor of human computation.

III. The Human Computation Metaphor as an Inhibitor

"In creating computer-based systems, we work within a perspective that shapes the design questions that will be asked and the kinds of solutions that are sought" – Winograd [14]

At the same time that a metaphor enables and facilitates our thoughts, every metaphor also constrains our thought, and, as a result, the designs we can come up with. Paraphrasing McLuhan [15], we can say that we shape our metaphors, and they then shape us. I would like to highlight here two ways in which the human computation may impede the design of new forms of organizing work over the web – by putting the focus on the words 'human', and 'computation'. Later in the paper I offer other perspectives that can help.

A. *People as Cogs: the Untapped Potential*

While the examples brought above got a lot of well-deserved accolade among the HCI and social computing communities, and indeed are novel in how they manage to orchestrate the work of many people, they all share one other characteristic which perhaps is a pity: people are used as cogs, doing relatively small, simple tasks. By taking that approach we can, indeed, achieve some tasks in an efficient manner, as the designers and users of those tools have demonstrated. This proves an especially powerful approach in performing large tasks that are essentially made of a huge collection of micro-tasks from a single class (e.g. photo-tagging). But significant cognitive potential remains untapped. The human mind is capable of doing so much more than tagging pictures or correcting grammatical errors. The programming metaphor constrained our thought in that it assumes that function calls need to be prescribed explicitly and accurately, ignoring the power of human intelligence. In treating people as individual cogs we risk losing several things:

1) *The ability of people to intelligize with others: The Role of Conversation*

Perhaps the most crucial capability we lose is the ability of people to work *together*, engaging in conversation to discuss goals and the ways to achieve them. All those systems mentioned above, and many other, do not allow for any communication among the people who perform the task. Sure, one might argue that there is information that is flowing in the system. E.g. in the CrowdForge examples, the MapReduce pattern makes the output of one 'layer' of people be the input for the next layer. And the same applies for the Find-Fix-Verify pattern featured by Soylent. Yet, each person works, basically, alone. In some examples of 'games with a purpose', a person plays with one other person. But there is not process of collective deliberation of the kind that is so crucial for fostering sensemaking, creativity and innovation.

At the end of the day, computation is computation - not intelligence. Some cognitive scientists have offered computational models that aim to account for how individual human minds reason about the world [16, 17]. But these accounts still do not explain how minds come up with creative solutions for tough problems. In the real world, big problems are solved through the work of many minds. Rome, Apollo 11, Windows and Wikipedia were all designed and built by many people intelligizing together, through mindful, elaborate interactions of minds that were aware of the ultimate goal; that purposefully strove to make sense of problems, situations, constraints, opportunities, features of material and socio-technical environments, and their affordances; who actively and creatively engaged other cognitions (human or not) in trying to find solutions.

The web enables us to connect minds and machines in scales we could not imagine 30 years ago. It allows us to distribute small and simple tasks to many people and thus, by breaking large tasks into many tiny pieces, achieve impressive productivity gains. Much like the industrial revolution introduced mass-production of tangible goods, we are now able to mass-produce information goods. Yet to address problems that we do not yet know how to solve - individually, and more importantly, collectively - we need to extend our thinking beyond the current focus on work-breakdown-structure.

Groups of people cognizing together can achieve much more than individuals – tackling bigger problems, and finding better solutions. Such processes of collective-intelligizing require people to engage in conversations. As Stahl [18] notes, discourse is the embodiment of the group's cognitive process. Group deliberation also facilitates learning, which can promote the ability of the group to address challenges in the future. Preventing group deliberation therefore cuts the group's ability to cognize and learn.

2) Motivation

It is difficult to motivate people to perform minute tasks which hardly require any skill. Von Ahn has acknowledged that, and has therefore come up with the idea of *games with a purpose*. This proved to be a well-crafted, clever solution, which binds small tasks in games such that success in the game is correlated with the quality of performing the task. Games are an activity that people enjoy doing, and as reCAPTCHA has shown, this works very, very well.

But it is not clear that every task can be broken into tiny pieces and be implanted in a game, and, in the process of embracing the human computation metaphor in the crowdsourcing world, something slipped in the way: many of the systems no longer strive to foster intrinsic motivation in the people who do the work. They simply involve paying people. Yet we know very well from the psychology literature that a sense of intrinsic motivation is much more important than extrinsic motivation such as money – especially so when things come to cognitive tasks.

Treating people like cogs makes it very hard to induce a sense of intrinsic motivation. As Hackman & Oldham [19] note, conditions for inducing internal work motivation require that tasks are made of whole and meaningful pieces of work, that people feel responsible for the outcome, and that they have knowledge of the results. These considerations are missing from the much of the current work on crowdsourcing, arguably due to focus on computational metaphors.

3) Memory, understanding, and relationships

Computer processors do not have memory. Every time we ask them to do something, we have to repeat our request in a precise manner. Human systems, however, learn. After a few months of working with a new marketing product manager, the R&D team learns what she means when she asks them for a feature. They learn to read some signs of what is more, and less important, and how to negotiate terms and find common grounds. The same is true for organizations and their suppliers, etc.

In the current model of human computation, the human cogs, like computer processors, have no memory. In fact, the current model is more akin to a 'cloud computing', as we get new (humans) processors with every transaction. Contrast that with a service provider in the 'real world'. As described above, we can develop a long-term relationship with that service provider that will help mutual understanding. This can save a lot of communication and coordination efforts.

B. Emphasis on the Human side

I focused on the 'computational' metaphor, explaining how a focus on computation may distract our attention from two important ways in which humans are different from computers: they need to be motivated (which designers have addressed, but in ways that can still be improved), and they can, through deliberation, tackle much larger problems than those they are given now.

But the 'human computation' metaphor also includes the 'human' part. By directing our focus at asking people to perform tasks which computers cannot yet solve (as von Ahn has posed human computation, and as many others have followed), we are essentially trying to overcome a problem of artificial-intelligence (yet another metaphor). Indeed, von Ahn concludes his dissertation with the following clause: *"We believe the techniques in this thesis present an opportunity for researchers and game designers to contribute to the progress of Artificial Intelligence"*.

By classifying tasks to those which computers do not do well yet, and those that they do, we often inadvertently also allocate tasks to humans, or computers, accordingly. What we might miss, are opportunities to design ways in which humans and computers can work together, as part of a collectively-intelligent system. I bring here a few examples of such systems and then speculate on how some other systems that are based on the human computation metaphor might be enhanced to incorporate computers.

My first example is, in fact, reCAPTCHA. In reCAPTCHA whose primary goal is the transcription of scanned books, the bulk of the OCR work is in fact done by computers. Only those words that are not recognized by the computer are crowdsourced. So reCAPTCHA perfectly achieves an optimal division of cognitive labor among people and computers.

The second example is FoldIt [20]. It addresses the problem of predicting protein structures, by crowdsourcing parts of the search to humans. FoldIt, relies on the human computation metaphor and extends it to a hybrid human-computer prediction and optimization framework through a large-scale online multiplayer game.

The third example, was introduced at the Computational Social Science and the Wisdom of Crowds workshop at NIPS 2010. Authors Nagar and Malone described a set of experiments in which humans and artificial-intelligence agents created collective predictions of future events by trading in hybrid human-machine prediction markets, and reported that those predictions were more accurate and more robust than those that were created by humans, or by computers [21].

The fourth example is Wikipedia. While the substantial editing work in Wikipedia is done by humans, bots perform a lot of 'house-maintenance' jobs such as checking the validity of hyperlinks, and enforcing various rules. This in itself, can be an example of a hybrid system, but recently, Sauper and Barzily presented software that can automatically generate Wikipedia articles [22]. These articles are then also edited by people.

Common to all these examples is the seamless manner in which cognitive efforts of humans and computers are integrated. Thus, the paradigm they employ is that of a hybrid/collective-intelligent socio-technical systems, rather than systems of pure *human* computation or pure AI systems.

It should be equally intriguing to see systems like Soylent, CrowdForge, etc., incorporate AI to work together with human crowds. Humans and computers do excel at different things. Combining human and machine intelligence can be powerful, as seen in the examples. Right now, those examples are few. Platforms like Mechanical-Turk, and the layers above it open up a range of possibilities for combining human and machine work. But to achieve the benefits, we need to think differently on the possibilities at hand.

IV. Considering Additional Perspectives and Metaphors

My goal in this paper, if this is not obvious, is not to argue against any metaphor. *Any* metaphor is just a metaphor, and any metaphor we use, *because* it focuses our attention on some things, will deprive our attention from other things at the same time that it helps us think and achieve progress. In Morgan's words, *"any one metaphorical insight provides but a partial and one-sided view of the phenomenon to which it is applied"* [8].

My goal here, therefore, is to argue for using *additional* metaphors when we think about the new possibilities the web opens for us to organize work and to harness the cognitive efforts and intelligizing capabilities of people, and machines.

To that end, I bring here two additional perspectives - that of organizational design, and that of distributed cognition, which I believe can be instrumental in facilitating the thinking of researchers and designers of collective-intelligence systems.

Let it be clear, that this is not to be understood in any way as a mere note about semantics, terminology, categorization or taxonomies. Phenomena that include using the web in novel ways of organizing information and people have been classified in many ways and named many ways, including human computation, social computing, crowdsourcing, collective intelligence, wisdom of crowds, web 2.0, etc., and it is not my intention nor goal to create order in that namespace. My call to consider additional metaphors when thinking about ways of organizing work - mainly cognitive work - over web platform is aimed to trigger and foster thinking and innovative design, and to promote and facilitate the discourse among different communities of practice.

A. The Organizational Design Perspective

The organizational design perspective offers that there are problems or challenges to solve, or certain goals to achieve. And people can organize themselves, or be organized, to address those challenges and goals. Over the years, the study of organizations has taught us about various organizational design patterns, and their properties. We now have a solid understanding of how certain organizational design patterns are suited to achieving different goals, and of some of the relative advantages, disadvantages and tradeoffs involved with various designs (for example, see [23-26]). I offer that we should think of crowdsourcing as a set/family of options for organizing work. This perspective can help us in addressing big challenges.

In contrast with the human-computation perspective, tasks and goals are not perceived as necessarily computational in nature. Rather, organizations deal with a wide variety of shifting goals, moving targets, many of which are unstructured. Humans have found ways to deal with such challenges, under many constraints, by building organizations that provide robustness and agility, balancing personal and organizational goals.

An Organizational Design perspective puts primacy on people, the organization members, who need to accomplish tasks, and considers the resources they have available. Considerations of motivation, facilitating interaction among members, etc., are native to this perspective.

B. The Distributed Cognition Perspective

As its name suggests, the distributed cognition perspective is a cognitive approach. But unlike traditional cognitive science, this approach does not assume cognition to take place solely in the individual mind. Rather, cognition is understood to be distributed across several minds, as well as artifacts in the socio-technical environment. Proponents of this perspective mainly come from developmental psychology, e.g. [27, 28] and cognitive anthropology, e.g. [29, 30], and it has affected research in HCI and CSCW.

As a cognitive approach, interest is in the flow of information, and the manipulation of symbols and representations. This is closer to the computational approach. But the distributed cognition approach does not assume people to be identical to computers.

If we shift our attention to think about organizational designs, and allow ourselves to think of radical designs that go beyond what Mechanical Turk and similar applications now enable, we might think of workers forming guilds, groups, or bands [cf. 31], who get much more complex tasks. A distributed cognition approach can help understand the processes of information flow and coordination of representations in such settings [32], especially in hybrid human-machine collaborations.

C. Bridging the Gap

One reason for gaps in thinking about crowdsourcing and social computing, is that many of the people who now work on crowdsourcing platforms and solutions were trained in computer science. Hence, the use of computational metaphors that lead the design is not at all surprising. But as I argue above, the design of novel crowdsourcing solutions can gain from a continuous conversation between computer scientists and organizational researchers. The problem is how to facilitate this conversation. As Bechky [33] notes, discourse and knowledge sharing across different of communities of practice is filled with misunderstandings rooted in differences in their language, the locus of their practice, and their conceptualization of goals. Boundary objects [33-35] can help achieve common grounds between the groups.

Boundary objects are “flexible epistemic artifacts” that “inhabit several intersecting social worlds and satisfy the information requirements of each of them” [35 cited by Bechky]. Star and Griesemer further write: *"Boundary objects are objects which are both plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to maintain a common identity across sites. They are weakly structured in common use, and become strongly structured in individual-site use. They may be abstract or concrete. They have different meanings in different social worlds but their structure is common enough to more than one world to make them recognizable means of translation."*

Metaphors can therefore serve as boundary objects. One metaphor that organization designers have long ago borrowed from computer scientists is that of information processing [24]. As any metaphor, the metaphor of information processing can inspire as well as inhibit thinking. But one advantage of this metaphor is that it is already shared by the communities of organization research, cognitive scientists, and computer science at large. The information processing view therefore can be one useful tool to bridge those disciplines and facilitate a fruitful discourse around crowdsourcing.

V. Conclusion

When we think about the ways of organizing people and their work, we should think about whether what we want to reify the industrial revolution, with its associated piecemeal work, re-engaging ourselves and our research in Taylorism, or whether we want to take real advantage of the tools and technologies we have at hand to facilitate participation, inclusion and diversity, fostering processes of collective intelligizing.

Systems like TurkIt, Soylent and Crowdforge have advanced the way we can organize crowdsource work in important ways: They enable the creation of hierarchical structures and hierarchical flows of information. They help with breaking work into pieces and better coordinating inputs and outputs. They help in achieving and guaranteeing quality of work; and – importantly – they make crowdsourcing more accessible. This is akin, to some extent, to the new ways of organizing production work in assembly lines that were introduced during the industrial revolution and the following period. Indeed, there are many gains we can accrue using this approach.

However, we can gain much, much more from organizing work over the web, and it is time we think more deeply about how to do just that. Wikipedia and several other initiatives are the exception right now. To really reap the fruits of human intelligence, we should think more of ways to encourage and foster collective production and collaborations

Asking people to solve problems that computers can't yet solve but humans have no problem solving, is, really, a simple thing, isn't it? To make my rhetorical question more bold, it can be taken ad-absurdum, if you just replace the word *computers*, with, say, *pigeons*.

In a way, pursuing 'artificial artificial-intelligence' is a back-route we try because we acknowledge that the capabilities of state-of-the-art, silicon-computer implementations of artificial-intelligence are not anywhere near those of humans. But if our goal is that of designing purely-artificial intelligent systems, than that is not the path. And if our goal, rather, is that of creating highly-intelligent systems, that can "act more intelligently than any individuals, groups, or computers have ever done before" [36], then should our focus be on those smaller problems that humans have no trouble with?

Let it be clear that I have no intention here to diminish any of the work on Human Computation. While there is nothing special about a person tagging a picture of a dog, tagging the *billions* of pictures on the web *is* a big, computationally-intensive task to, and so, at least until picture recognition technology catches up, it *is* a big deal. In that sense, games with a purpose are indeed a powerful solution. While editing paragraphs, making them more readable and concise, is within the capabilities of individual human minds, crowds can indeed help authors by improving quality under time constraints. Thus, human-computation is, without a doubt, a powerful paradigm. Translating the entire English Wikipedia to another language, for free, as Duolingo hopes to do, is a very impressive fit, whether they do so in 80 hours with a million users, in 5 weeks with 100,000 users, or even in a year, regardless of the number of people involved. It is a challenge for which we haven't found any other viable solution so far.

But humanity faces far larger challenges, like global health, designing cities, overcoming hunger and climate change, etc. Addressing these challenges mandates innovative thinking, way beyond what any artificial intelligence will be able to achieve in foreseeable time. It requires the ingenuity that right now only human minds, working together can produce. The same infrastructure that enables human computation and crowdsourcing, namely: the Web, is the infrastructure that already enabled us to write Wikipedia, collaborate on the task of deciphering the human genome, operate marketplaces that match seekers of scientific solutions with potential solvers, etc.

In this paper I tracked the use of the ‘human computation’ metaphor and related metaphors by showing how they affected design choices. In tracking and exposing the role of metaphors used in the field, and how they advance, and at the same time, inhibit, thinking of designers, researchers and practitioners, I go beyond the taxonomical work of Quinn and Bederson [6].

It is no wonder that the human computation metaphor has become prominent in the discourse around crowdsourcing in the social-computing community, for ways of thinking are mediated by social millieu [8], and the language of computation is native for many of the people who have brought much progress to this field.

Rather than asking computer scientists to learn social sciences (though, that can be encouraged too), this is a call for the attention of organizational designers, and researchers of distributed cognition alike, to devote more systematic thinking to the options enabled by the web; and it is also a call to enhance the conversation among disciplines in that context. To that end, I offered that the 'information processing' view is one perspective (of course, there may be other) that organizational researchers and designers, computer scientists, and many cognition researchers can relate to. Thus, this can serve as a 'boundary object' around which we can start anchoring and cultivating those cross-disciplinary conversations. I strongly believe that through such conversation, we can design collective-intelligence systems with which we can achieve giant leaps for mankind.

ACKNOWLEDGMENT

I thank Rob Miller for letting me join Crowdfoo meetings, where discussion is always courteous and enlightening. In these meetings I've had the privilege of talking with some of the authors whose work I discuss above, and I suppose not all will agree with what I argue, but I'm sure they will listen. My advisor, Tom Malone, provides a constant source of inspiration. Wanda Orlikowski helped enhancing my critical thinking. Patrick Winston taught me to use '*I*' when *I* write a paper. None of them should be blamed for any of my ideas.

REFERENCES

[1] J. C. Tang, M. Cebrian, N. A. Giacobe et al., "Reflecting on the DARPA Red Balloon Challenge," Communications of the ACM, vol. 54, no. 4, pp. 78-85, 2011.

[2] M. S. Bernstein, G. Little, R. C. Miller et al. "Soylent: a word processor with a crowd inside," Paper presented at the 23nd annual ACM symposium on User interface software and technology (UIST), New York, NY, 2010.

[3] A. Kittur, B. Smus, and R. E. Kraut, CrowdForge: Crowdsourcing Complex Work, Technical Report. CMUHCII-11, 2011.

[4] L. Von Ahn, B. Maurer, C. McMillen et al., "reCAPTCHA: Human-based character recognition via web security measures," Science, vol. 321, no. 5895, pp. 1465, 2008.

[5] D. A. Grier, When computers were human: Princeton University Press, 2005.

[6] A. J. Quinn, and B. B. Bederson. "Human computation: a survey and taxonomy of a growing field," Paper presented at the CHI, 2011.

[7] L. Von Ahn, "Human computation," PhD Thesis, School of Computer Science Carnegie Mellon University, Pittsburgh, PA, USA, 2005.

[8] G. Morgan, "Paradigms, metaphors, and puzzle solving in organization theory," Administrative science quarterly, pp. 605-622, 1980.

[9] G. Lakoff, and M. Johnson, Metaphors we live by: Chicago London, 1980.

[10] G. Morgan, "More on metaphor: why we cannot control tropes in administrative science," Administrative science quarterly, vol. 28, no. 4, pp. 601-607, 1983.

[11] G. Little, L. B. Chilton, M. Goldman et al. "Turkit: Tools for iterative tasks on mechanical turk," Paper presented at the ACM SIGKDD workshop on human computation, 2009.

[12] G. Little, L. B. Chilton, M. Goldman et al. "Turkit: human computation algorithms on mechanical turk," Paper presented at the 23nd annual ACM symposium on User interface software and technology, 2010.

[13] E. Gamma, R. Helm, R. Johnson et al., Design patterns: elements of reusable object-oriented software: Addison-wesley Reading, MA, 1995.

[14] T. Winograd, "A Language/Action Perspective on the Design of Cooperative Work," Human–Computer Interaction, vol. 3, no. 1, pp. 3-30, 1987.

[15] M. McLuhan, Understanding media: the extension of man: Routledge & Kegan Paul, 1964.

[16] E. Téglás, E. Vul, V. Girotto et al., "Pure Reasoning in 12-Month-Old Infants as Probabilistic Inference," Science, vol. 332, no. 6033 pp. 1054-1059 27 May, 2011.

[17] J. B. Tenenbaum, C. Kemp, T. L. Griffiths et al., "How to Grow a Mind: Statistics, Structure, and Abstraction," Science, vol. 331, no. 6022, pp. 1279, 2011.

[18] G. Stahl, Group cognition: Computer support for building collaborative knowledge, Cambridge, MA, USA: MIT Press, 2006.

[19] J. R. Hackman, and G. R. Oldham, "Motivation through the design of work," Work redesign, pp. 71-97, Reading, MA: Addison-Wesley., 1980.

[20] S. Cooper, F. Khatib, A. Treuille et al., "Predicting protein structures with a multiplayer online game," Nature, vol. 466, no. 7307, pp. 756–760, 2010.

[21] Y. Nagar, and T. W. Malone, "Combining Human and Machine Intelligence for Making Predictions," MIT Center for Collective Intelligence Working Paper No. 2011-02, 2011.

[22] C. Sauper, and R. Barzilay. "Automatically generating wikipedia articles: A structure-aware approach," Paper presented at the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2009.

[23] J. R. Galbraith, Designing complex organizations, Reading, MA: Addison-Wesley, 1973.

[24] J. R. Galbraith, "Organization Design: An Information Processing View," Interfaces, vol. 4, no. 3, pp. 28-36, 1974.

[25] T. W. Malone, K. Crowston, J. Lee et al., "Tools for inventing organizations: Toward a handbook of organizational processes," Management Science, vol. 45, no. 3, pp. 425-443, 1999.

[26] J. R. Galbraith, Designing organizations: an executive guide to strategy, structure, and process: Jossey-Bass Inc Pub, 2002.

[27] R. D. Pea, "Practices of distributed intelligence and designs for education," Distributed cognitions: Psychological and educational considerations, Learning in doing: Social, cognitive, and computational perspectives G. Salomon, ed., pp. 47-87: Cambridge University Press, 1993.

[28] G. Salomon, "No distribution without individuals' cognition: A dynamic interactional view," Distributed cognitions: Psychological and educational considerations, Learning in doing: Social, cognitive, and computational perspectives G. Salomon, ed., pp. 111-138: Cambridge University Press, 1993.

[29] E. Hutchins, Cognition in the Wild: MIT press Cambridge, MA, 1995.

[30] J. Hollan, E. Hutchins, and D. Kirsh, "Distributed cognition: toward a new foundation for human-computer interaction research," ACM Transactions on Computer-Human Interaction (TOCHI), vol. 7, no. 2, pp. 196, 2000.

[31] T. W. Malone, The Future of Work: How the New Order of Business Will Shape Your Organization, Your Management Style, and Your Life, Boston, MA: Harvard Business School Press, 2004.

[32] M. S. Ackerman, and C. Halverson, "Organizational memory as objects, processes, and trajectories: An examination of organizational memory in use," Computer Supported Cooperative Work (CSCW), vol. 13, no. 2, pp. 155-189, 2004.

[33] B. A. Bechky, "Sharing Meaning Across Occupational Communities: The Transformation of Understanding on a Production Floor," Organization Science, vol. 14, no. 3, pp. 312, 2003.

[34] S. L. Star, "The structure of ill-structured solutions: boundary objects and heterogeneous distributed problem solving," Distributed Artificial Intelligence, L. Gasser and M. N. Huhns, eds., pp. 37-54, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1990.

[35] S. L. Star, and J. R. Griesemer, "Institutional ecology,translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39," Social studies of science, vol. 19, no. 3, pp. 387-420, 1989.

[36] T. W. Malone. "What Is Collective Intelligence, and What Will We Do About It? Edited transcript of remarks at the official launch of the MIT Center for Collective Intelligence," December 1, 2008; http://cci.mit.edu/about/MaloneLaunchRemarks.html. Permanent URL: http://www.webcitation.org/5K9ZqKlVU.