Distributed Intelligence: Extending the Power of the Unaided, Individual Human Mind

Gerhard Fischer
University of Colorado, Center for LifeLong Learning and Design (L^3D)
Department of Computer Science, 430 UCB
Boulder, CO 80309-0430 – USA
http://l3d.cs.colorado.edu/
gerhard@colorado.edu

ABSTRACT
The history of the human race is one of increasing intellectual capability. Since the time of our early ancestors, our brains have gotten no bigger; nevertheless, there has been a steady accretion of new tools for intellectual work (including advanced visual interfaces) and an increasing distribution of complex activities among many minds. Despite this transcendence of human cognition beyond what is "inside" a person's head, most studies and frameworks on cognition have disregarded the social, physical, and artifactual surroundings in which cognition and human activity take place.

Distributed intelligence provides an effective theoretical framework for understanding what humans can achieve and how artifacts and tools can be designed and evaluated to empower human beings and to change tasks.

This paper presents and discusses the conceptual frameworks and systems that we have developed over the last decade to create effective socio-technical environments supporting distributed intelligence.

Categories and Subject Descriptors
H.5.3 Group and Organization Interfaces

General Terms
Design, Experimentation, Human Factors

1. INTRODUCTION
Inventing and creating good visualization is an important objective for numerous intellectual activities: thinking, working, learning, and collaboration. Advances in human cognition and intelligence have been made by powerful representations [Larkin & Simon, 1987] [Donald, 2004] — of which visualizations are a specific class (including such specific examples as Napoleon's March to Moscow [Tufte, 1983], Simon's analysis of strengths and weaknesses of different representations [Simon, 1996], and a collection of papers using vision to think [Card et al., 1999]).

Distributed intelligence is a conceptual framework that includes visualizations but has additional dimensions. Over the last decade, the Center for LifeLong Learning and Design [L3D, 2005] has developed a research agenda focused on distributed intelligence and lifelong learning. Our fundamental assumptions and objectives relevant to the framework of this paper are:

- Focusing on *Intelligence Augmentation (IA)* rather than on Artificial Intelligence (AI) by empowering human beings rather than replacing them [Fischer & Nakakoji, 1992; Terveen, 1995];
- Exploiting the power of the *human visual system* [Boecker et al., 1986] [Boecker et al., 1991] by creating task- and user-specific visualization;
- Providing support not only to individuals but to groups and communities, and thereby exploiting the power of *social creativity* based on informed participation [Fischer et al., 2005];
- Contextualizing generic systems to person- and task-specific environments to account for a "universe of one" by supporting *meta-design, customization, and end-user development* [Fischer & Giaccardi, 2006];
- Transcending *"gift-wrapping" and "techno-determinism"* as isolated and one-sided design objectives for new media by pursuing co-evolution among (i) new media; (ii) new theories about working, learning, and collaborating; and (iii) the creation of a new learning organization in a synergistic approach [Brown & Duguid, 2000; Fischer, 1998];
- Exploring *"computing off the desktop"* in different directions [Fischer & Konomi, 2005] by
  - *going small:* socio-technical environments supported by personalized, portable devices and wireless communication that afford information and communication among people as they move around in the world — the specific application context being the *CLever* project [Carmien et al., 2005];
  - *going large:* large computational tables that allow people from diverse backgrounds to access, contribute to, and interact with information in an inherently social manner to support collaborative work among others in shared physical locations — the specific application context being the *Envisionment and Discovery Collaboratory* [Arias et al., 2000].

2. DISTRIBUTED INTELLIGENCE
In most traditional approaches, *human cognition* has been seen as existing solely "inside" a person's head, and studies on cognition have often disregarded the physical and social surroundings in which cognition takes place. *Distributed intelligence* [Hollan et al., 2001; Salomon, 1993] provides a theoretical framework for understanding what humans can achieve and how artifacts, tools,

AVI '06, May 23-26, 2006, Venezia, Italy.

and *socio-technical environments* [Mumford, 1987] can be designed and evaluated to *empower human beings* and to *change tasks*. Distribution can take place

- *among people* — the rationale for it being: (1) "symmetry of ignorance" (being knowledgeable in some domains and not in others); (2) exploring new divisions of labor and developing new specializations; and (3) engaging stakeholders in collaborative learning and working efforts;
- *between humans minds and artifacts* — the rationale for it being: (1) complementing humans with computational capabilities; (2) developing "things that make us smart"; (3) creating powerful external representations (including advanced visualizations); and
- *as the integration of these two dimensions of distribution* — the rationale for it being: (1) creating tools for collaboration; (2) enhancing social creativity; (3) supporting reflective practitioners.

The critical importance of *externalizations* (and *oeuvres* [Bruner, 1996]) are: (1) they produce a record of our efforts, one that is "outside us" rather than simply in memory; (2) they make our thoughts and intentions more accessible to reflective efforts [Schön, 1983]; (3) they produce situations that talk back to us (e.g., in the form of critiquing systems [Fischer et al., 1998]). Externalizations that represent *works-in-progress* are of special interest, especially when all stakeholders are empowered to make active contributions to the evolution of artifacts [Fischer & Giaccardi, 2006], thereby producing and sustaining creativity with shared and negotiable ways of thinking in a group.

There are two major perspectives [Norman, 1993] conceptualizing distributed intelligence:

- *The personal point of view:* In this perspective, distributed intelligence *changes the nature of the tasks* that human beings have to do by creating new divisions of labor (e.g., the work that check-out clerks in a supermarket or airplane pilots have to do today compared to 30 years ago);
- *The system point of view:* In this perspective, the combined *"human(s) and artifacts"* system [Engelbart, 1995] is more powerful, is more reliable, and can achieve tasks that none of the components could achieve by themselves. As Einstein remarked: "My pencil is cleverer than I" (e.g., socio-technical environments for people with cognitive disabilities, modern cockpit (including pilot and computers) of an airplane).

Numerous claims and arguments can be found to indicate the necessity for a distributed intelligence perspective, including:

- human mental activity is neither solo nor conducted unassisted, even when it goes "inside the head" [Bruner, 1996];
- "how the mind works" is itself dependent on the tools at its disposal (analogous to "how the hand works" not being fully appreciated unless one takes into account whether it is equipped with a screwdriver, a hammer, or a pair of scissors) [Bruner, 1996];
- *brain-culture symbiosis*, which states that the human brain cannot realize its potential unless it is immersed in a distribution network [Donald, 2004]. This brain-culture symbiosis can lead to "higher intelligence" with the support of "mind tools" that allow people to think previously unthinkable thoughts by integrating the raw intellectual power of the human brain with appropriate technologies; and
- *material culture*, which externalizes memory and greatly amplifies the permanence and power of distributed cognition, and frees the symbolization process from the limitations of biological memory [Donald, 2004].

2.1 Distributed Intelligence in a Historical Context

Distributed intelligence is a fundamental framework by which to marry the raw intellectual power of the human mind with appropriate technologies [Donald, 2004]. Skills related to literacy [Ong, 1982] are the most important contribution. There is no doubt that information and communication technologies have created fundamental new design possibilities, but their true strengths and weaknesses need to be further explored. Figure 1 provides a graphical illustration of some of the major stepping-stones toward increasing the power of the collective human mind aided by technology (as an extension to Figure 1, the brief analysis of "Faustian Bargains" in section 5 is a reminder that new technologies will not necessarily lead to some monotonically increasing power of the collective human mind).

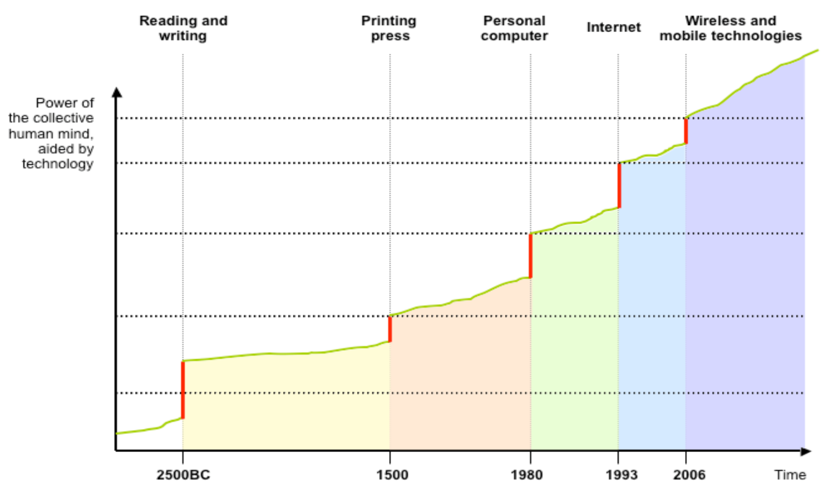

Figure 1: Beyond the Unaided Individual Human Mind

2.2 Tools for Learning and Tools for Living

An important specific classification within a distributed intelligence framework is to differentiate between tools for living (which remain part of our environment and our activities) and tools for learning (which often serve as a scaffolding mechanism and are designed and employed to fade away over time) [Carmien & Fischer, 2005; Pea, 2004].

Tools for learning support people in learning a new activity with the objective that they will eventually become *independent* of the tool. Tools for learning afford an internalization of what was (if it existed previously at all) an ability supported by external mechanisms. Examples of tools for learning are bicycles with training wheels or toddlers' walkers.

Tools for living are artifacts that empower human beings to do things that they could not do by themselves. They support distributed intelligence. Examples of tools for living include eyeglasses, the telephone, screen readers for blind people, visualization tools, and adult tricycles. No matter how many times people use the phone to talk to friends across town, their native ability to converse over long distances unaided remains the same. Tools for living allow people with disabilities to perform tasks that they would not able to accomplish unaided, and therefore allow these people to live more independently.

Whether a tool is a tool for learning or a tool for living is in many cases not an attribute of the tool itself, but is determined by *the use context and the objectives of the user*. Wizards used in many computational environments, spelling correctors, and handheld calculators can serve both purposes with different trade-offs. Learning to live and act without a tool will create an independence of the tool and may lead to a deeper understanding of the activity itself, but this will often come at considerable costs, including time and effort in learning the activity and then ultimately executing it in a possibly more error-prone and time-consuming way compared to using the tool.

2.3 "Eyeglasses" for the Mind

Distributed intelligence approaches support the vision and object that anatomy and cognitive abilities are not destiny. An important intellectual or philosophical grounding of this mission is provided by Postman [Postman, 1985]: *"The invention of eyeglasses in the twelfth century not only made it possible to improve defective vision but suggested the idea that human beings need not accept as final either the endowments of nature nor the ravages of time. Eyeglasses refuted the belief that anatomy is destiny by putting forward the idea that our minds as well as our bodies are improvable!"* The observation that "our minds are improvable" through media and technologies by constructing *"eyeglasses for the mind"* has led to the following objectives for our research with people with cognitive disabilities:

- the assertion that the cognitive abilities of all of us are limited (the most convincing example being the limitations of our memories addressed by the invention of reading and writing);
- the development of computational media that provide unique opportunities to "improve our minds" (and especially the minds of people with cognitive disabilities), leading to fundamental research challenges in media as extensions of humans [Engelbart, 1995; McLuhan, 1964; Norman, 1993].

3. SOCIO-TECHNICAL ENVIRONMENTS BASED ON A DISTRIBUTED INTELLIGENCE FRAMEWORK

This section describes two major development efforts at the Center for LifeLong Learning and Design. These evolving socio-technical environments are grounded in a distributed intelligence framework and at the same time contribute to the further development of that framework.

3.1 Cognitive Levers: Improving Minds with Distributed Intelligence

The CLever ("Cognitive Levers: Helping People Help Themselves" [CLever, 2005]) research project develops *socio-technical environments* [Mumford, 1987] to support persons with cognitive disabilities and their caregivers. These environments are designed to allow people with disabilities to perform tasks that they would not be able to accomplish unaided. The objective is to make people more *independent* by assisting them to live by themselves, use transportation systems, interact with others, and perform a variety of domestic tasks. CLever's goal is to create more powerful media, technologies, and communities to support new levels of distributed intelligence. Research in CLever includes the following specific environments: (1) human-centered public transportation systems, and (2) end-user development environments for prompting systems needed in this environment. The technologies developed in the CLever project will be broadly available as dual-use technologies applicable to a variety of different application areas.

Mobility-for-All: Human-Centered Public Transportation Systems. Public transportation systems are among the most ubiquitous and complex large-scale systems found in modern society. For those unable to drive, such as persons with cognitive disabilities, these systems are gateways for participation in community activities; they increase their opportunities for socialization, and they can provide a level of independence from other human beings. The *"Mobility-for-All" project* [Carmien et al., 2005] embedded in CLever is creating mobile architectures and prototypes to support persons with cognitive disabilities and their caregivers. This research has broad implications for designing more human-centered transportation systems that are universally accessible for other disenfranchised communities, such as the elderly [National-Research-Council, 2004], nonnative speakers, and infrequent users of public transportation systems.

Our field studies (analyzing current public transportation systems in several major cities) have suggested two major design strategies for creating a Mobility-for-All public transportation architecture: (1) design components that *simplify* the complex navigational artifacts encountered in public transportation systems; and (2) design architectures and components that transcend the need to understand complex artifacts and serve as a *dynamic, contextualized navigational assistant.* Our research is focused on the second design approach, and it explores specifically:

- *technologies for mobile users*: directly supporting the mobile user (specifically the person with cognitive disabilities) equipped with a PDA or mobile phone with personally relevant navigational tasks, including selecting a destination, locating the right bus, preparing to board, boarding the bus, signaling the driver where to get off, and disembarking;

- *technologies for caregivers:* when needed, initiating or facilitating communications between the mobile user, support communities, and transportation system operators; providing a "safety net" when something goes wrong.

Figure 2: Mobility-for-All — An Agent-based Prototype

Figure 2 shows an agent-based prototype environment integrating a mobile prompting device (to be used by the mobile user) synchronized with a virtual 3D display of a real-time bus system (to be used by the caregiver) [Repenning & Ioannidou, 2006].

The Mobility-for-All prototype provides personalized and contextualized assistance to address the unique "*universe of one"* [Fischer, 2001] problems of travelers with cognitive disabilities (people with disabilities more than those without form a "universe of one" in the sense that their abilities are unique). Human beings have *internal scripts* in their heads (representing tools for learning) that can be complemented by *external scripts* (representing tools for living).

Memory Aiding Prompting System (MAPS). MAPS [Carmien, 2006] is an end-user development environment supporting the creation of *external scripts* tailored to distribute cognitive tasks (such as traveling, cooking, etc.) by complementing a user's abilities. Scripts consist of memory prompts with task-specific visual and auditory stimuli and feedback. They are organized as finite state sequences with state changes triggered by the user's actions (selection from a menu, movement, etc.) or external events in the environment (arrival of a bus, passage of time, etc.).

One key design parameter in developing scripts is the granularity and specificity of a particular prompt and feedback sequence. When learning to travel independently, the prompting granularity may be very detailed, with frequent prompts and feedback. As the user gains experience, granularity of instructions can become more coarse-grained to reflect learning and be less intrusive (to fit the requirement of a tool for learning).

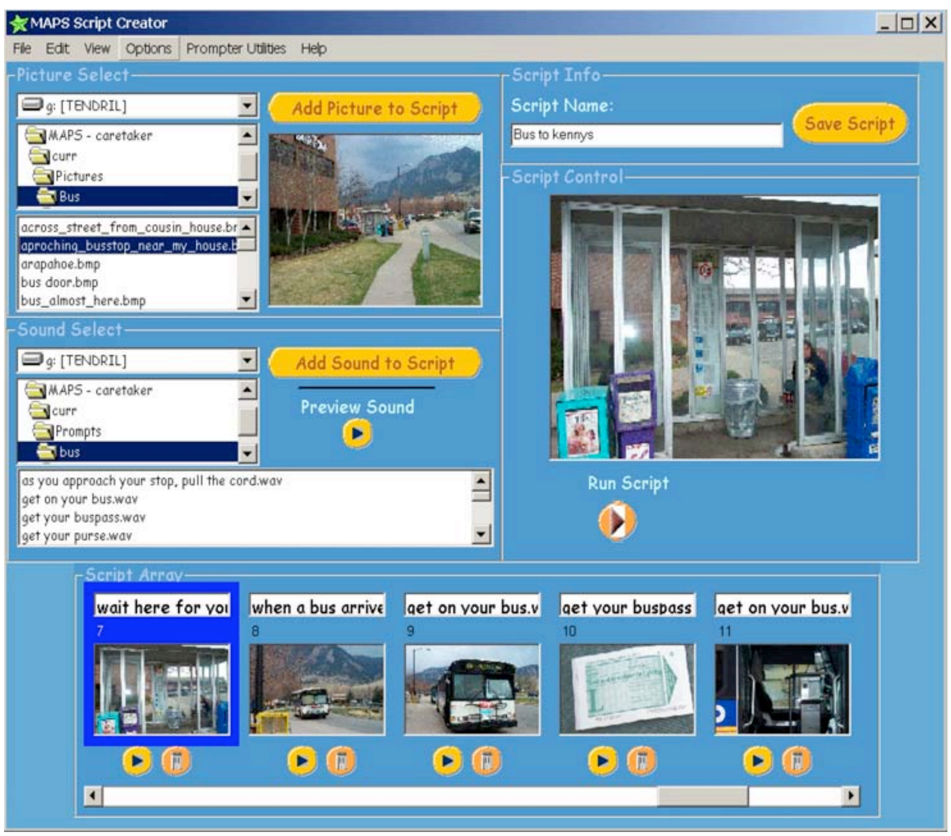

Figure 3: MAPS —An End-User Development Environment for External Scripts

The scripts needed to effectively support users are *specific for particular tasks*, creating the requirement that the people who know about these tasks (the local caregivers and not some technologists far removed from the action) must be able to develop scripts. Caregivers generally have no specific professional technology training nor are they interested in becoming computer programmers. This creates the need for design environments with extensive end-user support to allow caregivers to create, store, and share scripts [Fischer & Giaccardi, 2006]. Figure 3 shows the prototype of a caregiver configuration environment (embedded in MAPS) for creating complex, location-aware, multimodal prompting sequences. The prototype allows sound, pictures, video, and so forth to be assembled by using a film-strip-based scripting metaphor.

The design of the script editor presents interesting challenges for *meta-design*, a process for creating new media and environments that allow users to act as designers [Fischer, 2002; Fischer et al., 2004]. The caregiver script editor will support the specification of contextual tests that trigger specific interventions and guide the MAPS user back on track, or notify the caregiver when appropriate. Enabling non-programmers to create complex scripts with embedded error trapping and correction information constitutes a complex meta-design challenge, and a major emphasis in the overall script editor design.

3.2 Envisionment and Discovery Collaboratory

The *Envisionment and Discovery Collaboratory* (EDC) [Arias et al., 2000] is an environment in which participants collaboratively solve problems of mutual interest. The problem contexts explored in the EDC, such as urban transportation planning, flood mitigation, and building design, are all examples of open-ended social problems. The EDC empowers users to act as designers in *situated learning* and *collaborative problem-solving activities*. For most design problems, the knowledge to understand, frame, and solve them does not already exist, but is constructed and evolved during the solution process, exploiting the power of the "symmetry of ignorance" and "breakdowns." The EDC is an environment in which *social creativity* can come alive [Fischer, 2000].

Figure 4: Stakeholders Collaborating by Using the EDC

Figure 4 shows the EDC in use. The stakeholders (representatives from the Boulder City Council and the Regents from the University of Colorado) discuss and explore problems within a shared construction space in which they interact with physical objects that are used to represent the situation currently being discussed. As they manipulate physical objects, a corresponding computational representation is created and incrementally updated by using technologies that recognize the placement and manipulation of physical objects. Computer-generated information is projected back onto the horizontal physical construction area, creating an augmented reality environment. This physical construction is coupled with information relevant to the problem currently being discussed.

The collaborative design activities supported by the EDC rely upon the contribution and active participation of all involved stakeholders [Fischer, 2002]. Design domains consist of *ill-defined and wicked problems* [Schön, 1983], for which there are no correct answers, and in which framing problems is a major aspect of problem solving. The collaboration of stakeholders is an inherent aspect of these domains. For example, citizens, urban planners, and transportation experts who solve or are influenced by urban transportation issues are themselves an integral part of the problem context. The goal of the EDC is to bring together these stakeholders (each one or each group owning one part of the problem) to solve problems of mutual interest and to take active roles in addressing the problems that shape their lives.

Figure 5 shows the integration of the EDC with Google Earth. The participants have created sketches of new buildings with a specific height in their design, and Google Earth is used to show the impact of these actions (e.g., how much they block the view of the mountains — a major controversial issue in the City of Boulder, Colorado).

Figure 5: Integrating the EDC with Google Earth

The EDC has allowed the exploration of different dimensions of distributed intelligence, including [Fischer et al., 2005]:

- **Articulating tacit knowledge.** Often, critical aspects of the perspectives being brought to the EDC are based on knowledge not previously externalized and formalized. These are sometimes individual perspectives that have been experienced but never expressed. By allowing shared interaction through a descriptive process, the EDC supports methods whereby such perspectives can be drawn out and brought to the table, enhancing individual and collective contributions to the creative process. The overall design of the EDC, targeted toward these participants, employs the use of physical objects to create an inviting and natural interaction with the simulation, and recognizes that parallel interaction capability is essential to support this natural interaction [Eden, 2002]. The development of active critics

[Fischer et al., 1998] and virtual stakeholders [Arias et al., 1997] supports informed participation.

- **Creating bridge objects.** One danger of any model or conceptualization is that it may embody certain assumptions and perspectives that, if not questioned, can lead to an imbalance of influence within the process. These forms of *model monopoly* [Bråten, 1983] and *"group-think"* [Janis, 1972] need to be balanced by having open representations of the models and by supporting *epistemological pluralism* [Turkle & Papert, 1991] that allows for deeper understanding, experimentation, and possibly refutation. The goal is to permit a migration toward shared representations that are useful across contexts and communities as *bridge objects* [Bowker & Star, 2000]. The EDC design goals are to provide an open environment and design process that will allow these models to be developed and extended.
- **Supporting meta-design.** Creating models within the EDC requires a considerable amount of programming effort. This represents a high degree of reliance upon high-tech scribes, distancing the real designers from the medium of expression. Environments (even domain-oriented ones) that are open and easily modifiable and extensible are still elusive. While we continue to work on support for end-user development and meta-design [Fischer & Giaccardi, 2006], we are also looking at ways to harness existing tool use, integrate with existing practice, develop models (inspired by open source software development methodologies [Raymond & Young, 2001]), and empower local developers [Nardi, 1993].

Working from a meta-design perspective, we have begun to include mechanisms within the EDC to allow participants to inject content into the simulations and adapt the environment to new scenarios. The next steps include creating ways to link to existing data and tools so that participants can draw on information from their own areas of expertise to contribute to the emerging, shared model. These developments will allow us to explore further dimensions of a distributed intelligence framework within the EDC.

4. ASSESSMENT AND IMPLICATIONS

Distributed intelligence is a fundamental concept to rethink the relationship between humans and artifacts and to reinvent thinking, working, learning, and collaborating in an information-rich world equipped with computational tools and linked together with communication technologies. Numerous interesting assessment challenges and implications exist — a few of them are briefly discussed in the following paragraphs.

From Reflective Practitioners to Reflective Communities. In today's society *"the goal of creating current-day Leonardos who are competent in all of science"* [Campbell, 1969; Shneiderman, 2002] has to fail because (as argued in this paper) the individual, unaided human mind is limited. Knowledge is shifting from individual minds to a collective social product only imperfectly represented in any one mind: *"Even within disciplines, disciplinary competence is not achieved in individual minds, but as a collective achievement made possible by the overlap of narrow specialties"* [Campbell, 1969]. We need to invent alternative social organizations that will permit the flourishing of narrow interdisciplinary specialties as well as new media to support these reflective communities [National-Research-Council, 2003].

Quality of External Representations. In order to support distributed intelligence effectively, we need criteria to assess the quality, learnability, and effectiveness of external representation, including the following attributes:

- long-lasting (not ephemeral), allowing incremental improvements over time;
- produced, modified, reproduced, and evolved not only by experts or technologists, but by the owners of the problem;
- communicable over spatial, temporal, and conceptual distances;
- possessing computational capabilities (e.g., support for multimodal representations, or context-awareness of users, tasks, and situated environments); and
- exploiting and complementing the strength of humans (e.g., sometimes visualizations make a big difference and sometimes they do not).

Beyond Technology. How innovative are our ideas about the use of new technologies supporting distributed intelligence? Innovations should not be restricted to new technologies per se, but they should support the *co-evolution* of social practices, new media, and new learning organizations [Brown & Duguid, 2000]. To deeply understand the real impact of new technologies, we need to shift the emphasis from a concern about who has access to new technologies to who can use them in interesting ways for *personally meaningful tasks*.

"Faustian Bargains". All important technologies are *"Faustian bargains":* they give and take away and they produce winners and losers. Socrates argued that reading and writing will weaken human memory structures, and books will destroy thoughts. In our work in the CLever project, we always have to be aware of whether an *over-reliance on tools for living* will lead to learned helplessness and deskilling, ruining humans' native abilities by making them dependent on tools.

The material culture around us often overwhelms people with its richness. Based on the fundamental objective that the scarce resource is not information, but human attention, the real design challenges are (1) not to make the invisible visible, but to make the *important invisible* visible, and (2) not to provide information *"anywhere at anytime to anyone,"* but to say *"the 'right' information at the 'right' time, in the 'right' place, in the 'right' way, to the 'right' person."* Wireless and mobile technologies and pervasive computing environments allow us to be always connected, support distributed intelligence by making tools for living available at all times, and enable learning and using on demand in situated environments. At the same time, however, they can be intrusive and disruptive, leading to a state of continuous partial attention and destroying the moments needed for solitary reflection.

5. CONCLUSIONS

Distributed intelligence is anchored in the basic assumption that the human mind cannot reach its potential unless it is immersed and embedded in a distributed communication network (with other humans and with artifacts) supported by socio-technical environments. Visualizations have been one of the most important contributions to external representations, allowing people to understand and reason about complex activities. The last two decades have seen far-reaching changes in living, learning, working, and collaboration, fundamentally influenced by information and communication technologies that make new levels of distributed intelligence feasible and (potentially) desirable. In order to make further progress, the communities focused on advanced visual interfaces, human-computer interaction, collaborative systems, and pervasive computing should take up the challenge that the future is not "out there" to be discovered (like Columbus discovered America), but has to be invented and designed. Engaging in this challenge, we should not only create and promote new technologies, but make major contributions to fundamentally rethinking, reinventing, and redesigning the role of humans in a technology-rich world, and using our design capabilities and possibilities to improve the human condition.

6. ACKNOWLEDGMENTS

The author thanks the members of the Center for LifeLong Learning & Design (L3D) who have made substantial contributions to the conceptual framework and systems described in this paper. Special thanks go to: Stefan Carmien, Andrew Gorman, and Jim Sullivan for the development of the CLever projects; Ernesto Arias and Hal Eden for the development of the EDC, and Elisa Giaccardi for the collaborative work on meta-design.

The research was sponsored by: (1) a grant from the Coleman Institute, University of Colorado; (2) the National Science Foundation, Grants (a) REC-0106976 "Social Creativity and Meta-Design in Lifelong Learning Communities"; (b) CCR-0204277 "A Social-Technical Approach to the Evolutionary Construction of Reusable Software Component Repositories"; and (c) IIS 0456053 "SGER: Designing and developing mobile computing infrastructures and architectures to support people with cognitive disabilities and caregivers in authentic everyday tasks"; and (3) SRA Key Technology Laboratory, Inc., Tokyo, Japan.

7. REFERENCES

Arias, E. G., Eden, H., & Fischer, G. (1997) "Enhancing Communication, Facilitating Shared Understanding, and Creating Better Artifacts by Integrating Physical and Computational Media for Design." In *Proceedings of Designing Interactive Systems (DIS '97),* ACM, Amsterdam, The Netherlands, pp. 1-12.

Arias, E. G., Eden, H., Fischer, G., Gorman, A., & Scharff, E. (2000) "Transcending the Individual Human Mind—Creating Shared Understanding through Collaborative Design," *ACM Transactions on Computer Human-Interaction,* 7(1), pp. 84-113.

Boecker, H., Fischer, G., & Nieper, H. (1986) "The Enhancement of Understanding Through Visual Representations." In M. Mantei, & P. Orbeton (Eds.), *Proceedings of CHI'86 Conference on Human Factors in Computing Systems,* ACM, New York, pp. 44-50.

Boecker, H., Fischer, G., & Nieper-Lemke, H. (1991) "The Role of Visual Representations in Understanding Software." In D. Partridge (Ed.), *Artificial Intelligence and Software Engineering,* Ablex Publishing Corporation, Norwood, NJ, pp. 273-290.

Bowker, G. C., & Star, S. L. (2000) *Sorting Things Out — Classification and Its Consequences,* MIT Press, Cambridge, MA.

Bråten, S. (1983) "Asymmetric Discourse and Cognitive Autonomy: Resolving Model Monopoly Through Boundary Shifts." In G. de Zeeuw (Ed.), *Problems of Levels and Boundaries,* Princelet Editions, London/Zurich, pp. 7-12.

Brown, J. S., & Duguid, P. (2000) *The Social Life of Information,* Harvard Business School Press, Boston, MA.

Bruner, J. (1996) *The Culture of Education,* Harvard University Press, Cambridge, MA.

Campbell, D. T. (1969) "Ethnocentrism of Disciplines and the Fish-Scale Model of Omniscience." In M. Sherif, & C. W. Sherif (Eds.), *Interdisciplinary Relationships in the Social Sciences,* Aldine Publishing Company, Chicago, pp. 328-348.

Card, S. K., Mackinlay, J., & Shneiderman, B. (Eds.) (1999) *Readings in Information Visualization, First Edition : Using Vision to Think,* Morgan Kaufmann Series in Interactive Technologies, Orlando, FL.

Carmien, S. (2006) *Moving the Fulcrum: Socio-Technical Environments Supporting Distributed Cognition for Persons with Cognitive Disabilities*, Ph.D. Dissertation (forthcoming), University of Colorado at Boulder.

Carmien, S., Dawe, M., Fischer, G., Gorman, A., Kintsch, A., & Sullivan, J. F. (2005) "Socio-Technical Environments Supporting People with Cognitive Disabilities Using Public Transportation," *Transactions on Human-Computer Interaction (ToCHI),* 12(2), pp. 233-262.

Carmien, S., & Fischer, G. (2005) "Tools for Living and Tools for Learning." In, Proceedings of the HCI International Conference (HCII), Las Vegas, July 2005, p. (published on a CD).

CLever (2005) *CLever: Cognitive Levers -- Helping People Help Themselves,* Available at http://l3d.cs.colorado.edu/clever/.

Donald, M. W. (2004) "Is a Picture Really Worth a 1,000 Words? Review essay of David S. Staley, Computers, Visualization and History: How New Technology Will Transform Our Understanding of the Past.," *History and Theory: Studies in the Philosophy of History,* 43(3), pp. 379-385.

Eden, H. (2002) "Getting in on the (Inter)Action: Exploring Affordances for Collaborative Learning in a Context of Informed Participation." In G. Stahl (Ed.), *Proceedings of the Computer Supported Collaborative Learning (CSCL '2002) Conference*, Boulder, CO, pp. 399-407.

Engelbart, D. C. (1995) "Toward Augmenting the Human Intellect and Boosting Our Collective IQ," *Communications of the ACM,* 38(8), pp. 30-33.

Fischer, G. (1998) "Making Learning a Part of Life—Beyond the 'Gift-Wrapping' Approach of Technology." In P. Alheit, & E. Kammler (Eds.), *Lifelong Learning and Its Impact on Social and Regional Development,* Donat Verlag, Bremen, pp. 435-462.

Fischer, G. (2000) "Social Creativity, Symmetry of Ignorance and Meta-Design," *Knowledge-Based Systems Journal (Special Issue on Creativity & Cognition), Elsevier Science B.V., Oxford, UK,* 13(7-8), pp. 527-537.

Fischer, G. (2001) "User Modeling in Human-Computer Interaction," *User Modeling and User-Adapted Interaction (UMUAI),* 11(1), pp. 65-86.

Fischer, G. (2002) *Beyond 'Couch Potatoes': From Consumers to Designers and Active Contributors, in FirstMonday (Peer-Reviewed Journal on the Internet),* Available at http://firstmonday.org/issues/issue7_12/fischer/.

Fischer, G., & Giaccardi, E. (2006) "Meta-Design: A Framework for the Future of End User Development." In H. Lieberman, F. Paternò, & V. Wulf (Eds.), *End User Development: Empowering People to Flexibly Employ Advanced Information and Communication Technology,* Kluwer Academic Publishers, Dordrecht, The Netherlands, pp. 421-452.

Fischer, G., Giaccardi, E., Eden, H., Sugimoto, M., & Ye, Y. (2005) "Beyond Binary Choices: Integrating Individual and Social Creativity," *International Journal of Human-Computer Studies (IJHCS) Special Issue on Computer Support for Creativity (E.A. Edmonds & L. Candy, Eds.),* 63(4-5), pp. 482-512.

Fischer, G., Giaccardi, E., Ye, Y., Sutcliffe, A. G., & Mehandjiev, N. (2004) "Meta-Design: A Manifesto for End-User Development," *Communications of the ACM,* 47(9), pp. 33-37.

Fischer, G., & Konomi, S. (2005) "Innovative Media in Support of Distributed Intelligence and Lifelong Learning." In *Proceedings of the Third IEEE International Workshop on Wireless and Mobile Technologies in Education,* IEEE Computer Society, Tokushima, Japan, pp. 3-10.

Fischer, G., & Nakakoji, K. (1992) "Beyond the Macho Approach of Artificial Intelligence: Empower Human Designers - Do Not Replace Them," *Knowledge-Based Systems Journal, Special Issue on AI in Design,* 5(1), pp. 15-30.

Fischer, G., Nakakoji, K., Ostwald, J., Stahl, G., & Sumner, T. (1998) "Embedding Critics in Design Environments." In M. T. Maybury, & W. Wahlster (Eds.), *Readings in Intelligent User Interfaces,* Morgan Kaufmann, San Francisco, pp. 537-559.

Hollan, J., Hutchins, E., & Kirsch, D. (2001) "Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research." In J. M. Carroll (Ed.), *Human-Computer Interaction in the New Millennium,* ACM Press, New York, pp. 75-94.

Janis, I. (1972) *Victims of Groupthink,* Houghton Mifflin, Boston.

L3D (2005) *Center for LifeLong Learning and Design, University of Colorado, Boulder,* Available at http://l3d.cs.colorado.edu/.

Larkin, J. H., & Simon, H. A. (1987) "Why a diagram is (sometimes) worth 10,000 words?," *Cognitive Science,* 11, pp. 65-100.

McLuhan, M. (1964) *Understanding Media: The Extensions of Man,* The MIT Press, Cambridge, MA.

Mumford, E. (1987) "Sociotechnical Systems Design: Evolving Theory and Practice." In G. Bjerknes, P. Ehn, & M. Kyng (Eds.), *Computers and Democracy,* Avebury, Aldershot, UK, pp. 59-76.

Nardi, B. A. (1993) *A Small Matter of Programming,* The MIT Press, Cambridge, MA.

National-Research-Council (2003) *Beyond Productivity: Information Technology, Innovation, and Creativity,* National Academy Press, Washington, DC.

National-Research-Council (2004) *Technology for Adaptive Aging,* National Academy Press, Washington, DC.

Norman, D. A. (1993) *Things That Make Us Smart,* Addison-Wesley Publishing Company, Reading, MA.

Ong, W. J. (1982) *Orality and Literacy,* Routledge, London.

Pea, R. D. (2004) "The Social and Technological Dimensions of Scaffolding and Related Theoretical Concepts for Learning, Education, and Human Activity," *The Journal of the Learning Sciences,* 13(3), pp. 423-451.

Postman, N. (1985) *Amusing Ourselves to Death—Public Discourse in the Age of Show Business,* Penguin Books, New York.

Raymond, E. S., & Young, B. (2001) *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary,* O'Reilly & Associates, Sebastopol, CA.

Repenning, A., & Ioannidou, A. (2006) "Mobility Agents: Guiding and Tracking Public Transporation Users." In *Proceedings of the AVI Conference 2006*, Venice, p. (this volume).

Salomon, G. (1993) *Distributed Cognitions: Psychological and Educational Considerations,* Cambridge University Press, Cambridge.

Schön, D. A. (1983) *The Reflective Practitioner: How Professionals Think in Action,* Basic Books, New York.

Shneiderman, B. (2002) *Leonardo's Laptop — Human Needs and the New Computing Technologies,* MIT Press, Cambridge, Mass.

Simon, H. A. (1996) *The Sciences of the Artificial,* third ed., The MIT Press, Cambridge, MA.

Terveen, L. G. (1995) "An Overview of Human-Computer Collaboration," *Knowledge-Based Systems Journal, Special Issue on Human-Computer Collaboration,* 8(2-3), pp. 67-81.

Tufte, E. R. (1983) *The Visual Display of Quantitative Information,* Graphics Press, Cheshire, CT.

Turkle, S., & Papert, S. (1991) "Epistemological Pluralism and the Revaluation of the Concrete." In I. Harel, & S. Papert (Eds.), *Constructionism,* Ablex Publishing Corporation, Norwood, NJ, pp. 161-191.