

Are We There Yet? Spatial Cognitive Engineering for Situated Human-Computer Interaction

Kai-Florian Richter, Martin Tomko, and Arzu Çöltekin

Department of Geography
University of Zürich, Switzerland

{kai-florian.richter,martin.tomko,arzu.coltekin}@geo.uzh.ch

1 Introduction

Spatial cognitive engineering exploits scientific understanding of the way people conceptualize, perceive and communicate about space in order to devise computational systems that support spatial tasks and spatial decision making. It is the spatial extension of Norman's [13] concept of cognitive engineering—a kind of applied cognitive science, exploiting knowledge and experience from the cognitive sciences to the design of machines (see also [16]).

As such, spatial cognitive engineering is a form of user-centered design. It aims at reducing the human-computer interaction gaps by increasing the computer's 'understanding' of the human users. The overall aim is to get the device to adapt to the user, instead of forcing the user to accommodate to the device. It is assumed that achieving this aim will lead to a more natural interaction, better decision making, and truly intelligent spatial services [17], thus alleviating some of the 'ironies of automation' [1].

In this position paper we argue for a few key principles and design considerations that we deem crucial in spatial cognitive engineering, and outline a list of major challenges. We have already hinted at one of the principal challenges: while humans are masters in adaptation, machines are so far particularly bad at it.

2 Seamless Perfection or Perfect Seams

In 1991, Mark Weiser [22] presented his vision of ubiquitous computing, which would have computing devices disappear into the background of our everyday activities—always present, always active, but no more noticeable than street signs, billboards, or most furniture. Nearly 25 years later, this vision is far from becoming reality, even though it has begun to manifest in subtle and unexpected ways. In many parts of the world, the Internet is near ubiquitous and smart phones and wearables are in most people's pockets. People use these devices in many unintended and novel ways, for example, for socializing, entertainment or fitness tracking. These devices now offer ubiquity, but different to Weiser's vision, they are seamful and 'messy' [2, 3].

This does not necessarily mean Weiser's vision is an un-achievable ideal, but rather might indicate that 'seamless perfection' does not lead to optimal usability. The world may not be a desktop [23], but perhaps it is not necessary to hide computers from the users completely either. We draw an analogy to Montello's [11] distinction of two

components of navigation: the near automatic, subconscious *locomotion*, and the goal-directed, conscious planning component of *wayfinding*. Similarly, with assistance systems there are also aspects that are (1) important for the accomplishment of the goal, but may be executed fully automated, without the user realizing when and how they are performed; and (2) aspects that directly affect the interaction between human and system, in which humans should 'have a say'.

Consider the case of self-driving cars that do not seem to be too far from becoming a reality. There are a number of tasks which may indeed be best left to the car. For example, the car's global positioning system (GPS) receiver is certainly better at computing the current position, surpassing numerical processing capabilities of any human user. We also cannot reach the reliability of Anti-lock Braking Systems (ABS) to keep the wheels from locking in case of sudden braking. These are aspects of assistance that require complicated, fast numeric processing capabilities or involve instant safety decisions based on numerous and frequent sensor readings. In our eyes, these decisions are a parallel to *locomotion* in wayfinding as their execution needs to be near instant and can be automated. These routine decisions can be left to reliable and fast embedded hardware and software. Processes that are clearly beyond human capacity may be fully automated and delegated to assistance systems, thus remaining hidden from users.

In higher-level tasks requiring complex cognitive processing, however, there are advantages in making users fully aware that they are assisted by a computational system (cf. [14]). Such awareness may be beneficial for both the developers and users of such systems. When trying to develop more natural, human-like ubiquitous interaction interfaces we risk to face the *uncanny valley* [12], situations where computer interaction becomes awkward or subconsciously repulsive to the human user due to exaggerated lifelike appearance or behavior. Furthermore, systems that completely suppress users' responsibility and ability to override system decisions may render the users helpless during system failure (e.g., [1]).

Computers and humans are good at different things, therefore, a crucial aspect of designing meaningful interaction should be to smartly divide the tasks between humans and computers. While humans are better at tasks involving unfamiliar pattern interpretation, adaptation to changing contexts and heuristics, machines can detect, process, store and retrieve accurately and efficiently large amounts of data, interpret familiar patterns and efficiently conduct complicated calculations. Accordingly, when faced with dynamically changing traffic conditions a navigation system may be much faster and more accurate in identifying the delay to be ten minutes, while the decision whether to take a detour may be better left to the user. This is because users may (1) have different, un-quantifiable route preferences and behavioral patterns; (2) they may be better at assessing how likely it is that the delay will actually happen as predicted; and importantly (3) the users will ultimately bear the consequences of the decision.

We believe that it may be smarter to embrace these differences in the design of assistance systems, suitably reminding the users that they are being assisted by a computer. In other words, as a first principle, design assistance with *perfect seams* in mind, rather than to strive for *seamless perfection* in vain.

3 To Each Their Own and All Together Now

A corollary of the above principle is that the aim of (spatial) cognitive engineering should not be to mimic humans faithfully. Rather, the aim should be at combining the strengths of the human and computer in a complementary manner. The two should become a team where, first, users are assisted by complementing and enhancing their own capabilities with the power of the computer and its sensors, and second, the computer gets help with cognitively complex tasks for the human's own benefit. This has a range of consequences for the design process:

Model perfection, not imperfection First, system designers should take the interaction partners as what they are. Instead of pursuing a complex model of a human user with all their shortcomings, quirks, and cognitive biases, it seems more useful to exploit the predictability and reliability of computational systems and to construct a kind of 'super-human' [18]. We are seeking a system without unnecessary biases that would simulate human imperfections, forgetfulness, and misrepresentations. We envisage a system that consistently provides information when and how users need it, and if unsure, requests inputs. We do not need perfect imitation, but meaningful interaction [7]. Our self-driving car should not take a less optimal route, or learn and entrench biases just because people do so. Similarly, in user interaction, the car should acknowledge inputs provided by a user (e.g., destination input) and if this is evaluated as ambiguous, it should seek further input to resolve the ambiguity, rather than attempt autonomous interpretation. This requires having the ability to exchange information in terms that users can understand and interpret.

Do not substitute each other Second, we do not believe that strong anthropomorphism is a necessity. Supporting meaningful interaction is possible without pretending to be human. While adaptation to the human user is important, it should be offered in moderation and with a transparent and balanced task distribution. A software system trying to entirely substitute human judgment will annoy the user, whether they are reliable or whether they fail. Disempowering, patronizing systems are excellent examples of design failures.

Offer inclusive assistance Third, assistance systems should not simply assist the human user in their task execution—they should involve them in the decision making. Participation in decision making is important to keep users engaged, and supports a resilient human-computer system. It enables faster reaction to failures, since the user is able to identify the stage of the decision process where the computer failed. We have to accept that the system as well as the user may make a wrong decision because of the ambiguity of the situation, its novelty, or sensory and perceptual uncertainty. Recall the traffic delay and the destination input scenarios described earlier. In such cases, providing interaction facilities to resolve issues is a necessary design feature. Occasionally, the system may not even be aware that it is reacting to an uncertain situation, for example, if a self-driving car receives incorrect or imprecise route instructions from the user [21]. In short, opt-in and opt-out capabilities should be clearly presented to the user and where possible, adapt to the context as well as to the users' experience and abilities.

Provide transparent assistance Finally, transparency of the assisting processes should be ensured through interaction design and operation. A transparent system would, at all times, provide the user with the ability to verify the decisions, and possibly check the data based on which the decisions have been taken. While such detailed analyses will usually be unnecessary, in critical situations having available facilities to reconstruct system decisions may be crucial. This might be particularly important for safety critical applications, such as piloting a plane, managing a nuclear power plant, or indeed traveling in a self-driving car.

4 On Usefulness and Usability

Usability and usefulness are critically important considerations when designing assistance systems. The technical sciences are riddled by statements regarding the usability and user-friendliness of technology-driven products, but lack awareness of some of the fundamental principles. Schneiderman and Plaisant [19] point to the importance of gauging the skill level of the user. Invariably, *the designer is not the user*, and too often the designer is an expert. There is ample evidence that experts have different strategies in using artifacts than novices (e.g., [6]). Since spatial cognitive engineering starts off by examining the users' abilities and understanding of the problem domain, it inevitably takes a user-centered perspective—but this is only the first step towards usability.

While usefulness indicates the utilitarian value of a given artifact, i.e. whether it is fit for purpose and actually works, usability qualifies how fit it is for use [4, 10]. In that respect, usefulness helps us assess whether the artifact actually provides an added value to the user in a given situation (e.g., the car is in perfectly good condition but do you really need the car, or could you just walk?). Likewise, perfectly useful systems may not be usable in particular situations. Users may apply them in unintended ways, or design testing may be insufficient (e.g., you might very well need the car, but the acceleration and the brake functions are swapped).

Cognitive engineering requires an iterative design and testing process where concepts identified in empirical research are first implemented in prototype systems. The prototypes are then tested in realistic (or at least representative) scenarios and with a well-selected, varied, and counter-balanced set of users. Since the developed systems are based on the designers' understanding of human abilities, it is necessary to assess whether their assumptions hold. For example, in our design we can make use of the numerous standardized spatial abilities tests available to determine how well users with varied spatial abilities are supported by the designed system and whether the interaction between computer and human will work as intended. Thus acquired new insights need to be fed back into the design and tested again in a cyclical design process.

Furthermore, the adaptation process between human and computer may benefit from the machine's ability to learn. This facilitates adaptation to a given user and leads to system personalization [5]. By collecting interaction data (e.g., through user input or eye tracking), patterns in a particular user's behavior may be inferred using machine learning techniques. The system may then react to these patterns and adapt its behavior—still following the principles outlined above. For example, as individual spatial abilities strongly influence our performance in different cognitive tasks (e.g., [8]), adaptation

to these abilities could be targeted by spatial assistance systems. Specific modes of assistance benefit users differentially (e.g., those with higher vs. those with lower spatial abilities), as shown in situations including the interpretation of 3D displays or route descriptions [9, 15]. Ideally, the mode of assistance and interaction should adapt to an individual's capabilities and needs.

Finally, learning processes may also help systems in dealing with unusual or novel situations, drawing analogies from past experiences. The systems may then be better prepared for similar situations in the future to resolve these situations from their previous interactions with the human user. For this to work the human and the machine need to be perceived as a team. Then, both the machine and the user will learn about their reactions in diverse situations and may be able to adapt their behavior.

5 What is so Special about Spatial?

The properties of physical space require specific attention in designing interaction and assistance systems. Some aspects of spatial embedding may significantly simplify the design of assistance systems. Any physical object has to be at a certain, single location that cannot be concurrently shared with another object (of the same type). The layout of our spatial environments also restricts potential relationships and movement, which limits possible solutions and ambiguities while reasoning about space.

The ability to implement efficient reasoning heuristics in the machine is well complemented with the human expertise with spatial behavior and communication, which is efficient and effective despite the individual differences. People deal with space all the time, every day of our lives. However, this entails a potential risk as well: we all think that we are spatial experts. Thus, extra care is necessary when designing interaction and communication in and about space because everybody *knows* how this should be done right. The *uncanny valley* may be particularly deep here. This is further reinforced by the fundamental differences in how people perceive and conceptualize space, and how space is represented in geographic information systems. This mismatch in understanding space along with the human expertise is a major hurdle in designing spatial assistance systems that account for the guidelines laid out above.

6 Are We There Yet?

Where does this leave us?

Context may be the single biggest challenge in (spatial) cognitive engineering, and in (situated) human-computer interaction in general. Proper modeling, sensing and handling of various parameters of context in conjunction with the inference of user intent are far from being solved—and may never be. Still, the ability to provide tailored assistance requires that we at least try, as context handling remains a necessity for adaptive assistance systems. Until recently, context has been researched as a means of providing tailored information to the user. However, adaptation mechanisms must increasingly support the entire team of user and machine interacting in as natural a way as possible. Coming back to the scenario of the self driving car, it will require a model of user context to understand informal natural language interaction, such as destination

descriptions (e.g., “Take me to John’s golf club!”), or interceptions (the driver waking up during a long drive and asking “Are we there yet?”). But the user may also need to understand the car prompting for inputs en-route, for instance if the traffic conditions change or a parking lot is not available: “Do you want me to park a block further or do you prefer to pay for underground parking?”.

As capturing context is such a notoriously hard problem, we may need to start thinking about clever ways to cheat our way out of proper context handling. This may also hold true for strong personalization, i.e., personalization to each and every individual’s different needs and abilities. First, figuring out what these needs and abilities are is difficult. Psychological research has developed a range of assessment tests to determine people’s abilities and personality traits. But we can hardly ask users to first answer a large battery of questions before starting to use a system. Again, machine learning based on (large amounts of) collected observational data of system use may help to solve some of the issues of personalization. But learning takes time and repetitive use. Systems, however, need to be usable and useful from the very beginning as users would not accept them otherwise.

Second, with strong personalization, the space of design and interface options may just become too large to set up and maintain. Thus, in order to keep it manageable group-based personalization seems to be a sensible approach (e.g., [16, 20]). Still, while sometimes the classification of users into groups may be relatively straightforward, the parameters defining a group may not be simple to detect (supervised classification), or meaningful (unsupervised clustering) for the characteristics to which the system adapts.

Finally, a major challenge faced by (spatial) cognitive engineering is the transfer of cognitive findings from the lab into the real world. Psychological research, by necessity, usually is extremely controlled. It tries to manipulate a single aspect of a specific task in systematic, but often also artificial ways. Such research provides important basic insights. But they are often also too specific for real-life applicability, isolating individual cognitive responses and neglecting the interaction with other basic cognitive phenomena. The need to assemble these findings in a working system demands careful balancing of the influences of individual requirements. In the end, we need to design working systems that are useful and usable, and that do not suffer from ‘nervous breakdowns’.

So, “Are we there yet?” for spatial cognitive engineering? In short: “No”. We are only inching out of the labs to real applicability and production of cognitive systems. The long answer is that a range of major challenges remain to be overcome—or better, circumnavigate—to arrive at truly smart, cognitively supportive spatial services.

References

1. Bainbridge, L.: Ironies of automation. *Automatica* 19(6), 775 – 779 (1983)
2. Barkhuus, L., Polichar, V.: Empowerment through seamfulness: Smart phones in everyday life. *Personal and Ubiquitous Computing* 15(6), 629–639 (2011)
3. Bell, G., Dourish, P.: Yesterday’s tomorrows: Notes on ubiquitous computing’s dominant vision. *Personal and Ubiquitous Computing* 11(2), 133–143 (2007)
4. Bevan, N.: Usability is quality of use. *Advances in Human Factors Ergonomics* 20, 349–349 (1995)

5. Blom, J.: Personalization: A taxonomy. In: CHI'00 Extended Abstracts on Human Factors in Computing Systems. pp. 313–314. ACM, New York, NY (2000)
6. Çöltekin, A., Fabrikant, S., Lacayo, M.: Exploring the efficiency of users' visual analytics strategies based on sequence analysis of eye movement recordings. *International Journal of Geographical Information Science* 24(10), 1559–1575 (2010)
7. French, R.M.: Moving beyond the Turing test. *Communications of the ACM* 55(12), 74–77 (2012)
8. Hegarty, M., Richardson, A.E., Montello, D.R., Lovelace, K., Subbiah, I.: Development of a self-report measure of environmental spatial ability. *Intelligence* 30(5), 425 – 447 (2002)
9. Huk, T.: Who benefits from learning with 3D models? The case of spatial ability. *Journal of Computer Assisted Learning* 22(6), 392–404 (2006)
10. Kidwell, P.: The trouble with computers: Usefulness, usability and productivity. *IEEE Annals of the History of Computing* 18(4) (1996)
11. Montello, D.R.: Navigation. In: Shah, P., Miyake, A. (eds.) *Handbook of Visuospatial Thinking*, pp. 257–294. Cambridge University Press, Cambridge (2005)
12. Mori, M., MacDorman, K., Kageki, N.: The uncanny valley [from the field]. *IEEE Robotics Automation Magazine* 19(2), 98–100 (2012)
13. Norman, D.: Cognitive engineering. In: Norman, D., Draper, S. (eds.) *User Centered System Design*, pp. 31–61. Lawrence Erlbaum Assoc., Hillsdale, NJ (1986)
14. Olsen, K.A., Malizia, A.: Interfaces for the ordinary user: Can we hide too much? *Communications of the ACM* 55(1), 38–40 (2012)
15. Padgitt, A.J., Hund, A.M.: How good are these directions? Determining direction quality and wayfinding efficiency. *Journal of Environmental Psychology* 32(2), 164 – 172 (2012)
16. Raubal, M.: Cognitive engineering for geographic information science. *Geography Compass* 3(3), 1087–1104 (2009)
17. Richter, K.F., Winter, S.: Citizens as database: Conscious ubiquity in data collection. In: Pfoser, D., Tao, Y., Mouratidis, K., Nascimento, M., Mokbel, M., Shekhar, S., Huang, Y. (eds.) *Advances in Spatial and Temporal Databases. Lecture Notes in Computer Science*, vol. 6849, pp. 445–448. Springer, Berlin (2011)
18. Richter, K.F., Winter, S.: *Landmarks — GIScience for Intelligent Services*. Springer International Publishing (2014)
19. Shneiderman, B., Plaisant, C.: *Designing the user interface*. Pearson, Reading, MA, 5 edn. (2009)
20. Teevan, J., Morris, M.R., Bush, S.: Discovering and using groups to improve personalized search. In: *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pp. 15–24. ACM, New York, NY (2009)
21. Tomko, M., Richter, K.F.: Defensive wayfinding: Incongruent information in route following. In: *Spatial Information Theory. Lecture Notes in Computer Science*, Springer International Publishing (2015), to appear
22. Weiser, M.: The computer for the twenty-first century. *Scientific American* 265(3), 94–104 (1991)
23. Weiser, M.: The world is not a desktop. *Interactions* 1, 7–8 (1994)