

Towards On-Line Services Based on a Holistic Analysis of Human Activities

William J. Clancey

NASA-Ames Research Center and Florida Institute of Human and Machine Cognition

William.J.Clancey@NASA.gov

Abstract

Very often computer scientists view computerization of services in terms of the logistics of human-machine interaction, including establishing a contract, accessing records, and of course designing an interface. But this analysis often moves too quickly to tactical details, failing to frame the entire service in human terms, and not recognizing the mutual learning required to define and relate goals, constraints, and the personalized value of available services. In particular, on-line services that “computerize communication” (McCarthy, 2000, p. 17) can be improved by constructing an activity model of what the person is trying to do, not just filtering, comparing, and selling piece-meal services.

1. Introduction

From the customer’s perspective, the task of an on-line travel service is not merely to establish confirmed reservations, but to have a *complete travel plan*, usually integrating many days of transportation, lodging, and recreation into a happy experience. The task of the travel agent is not merely “ticketing,” but helping the customer understand what they want and providing services that will connect everything together in an enjoyable way.

The communication process required to provide a travel service is necessarily iterative, based on tentative constraints and plans, and considers long-term business relationships. Indeed, in the case of a computerized travel agent, mimicking human behavior and social customs would be a bad idea—we need to provide much better services than what people currently receive.

Using a web-based travel planning experience as a detailed example, this chapter critiques the approach to computerized interactions suggested by “Elephant 2000” (McCarthy 1998) and relates it to the response “Language for Interactions” (Cerri, et al. 2000). The ensuing discussion considers relations to Grid learning (e.g., Cerri 2002; Mostow & Tedesco 2004) and makes recommendations for using activity-based modeling to

facilitate mutual learning between computer agents and people. A broad view of educational material and learning is required to realize the potential of Grid computing to provide high-level services, as in the example of a travel agent.

2. Example: Web-based travel reservation

The following example is expressed in first-person form, to emphasize that it was an actual event and to remind the reader of the everyday nature of the situations and problems described here.

In the winter of 2001, I used the WebTravel system (a pseudonym for a free Internet service) to specify a flight from San Francisco to Page, Utah. The program told me that there were no planes flying to Page and listed the closest airports. I chose the Grand Canyon, seeing that I could arrive by 2 PM and join my friends who by coincidence would be staying there the previous night and then driving on to Page. After making the plane reservation, I then proceeded to make a rental car reservation. Surprise! There were no rental cars available at the Grand Canyon! This is remarkable given that 2 million visitors come to that area per year. Why didn’t the WebTravel program notice this? My remaining itinerary, already input to the program, clearly shows that I am flying from Grand Junction, Colorado to San Francisco—how did the program expect me to get from the Grand Canyon to Grand Junction?

The user experience is even worse than I have described. For the program does not say, “no rental cars are available”—it displays instead that “The car type requested is not available at this location through WebTravel.com at this time. Please select a different car type.”

The WebTravel program, as in the Elephant 2000 analysis, focuses not on my travel, but myopically only on reservations. I indeed now held guaranteed reservations (the emphasis of McCarthy’s logic-based approach), but they were worthless to me. My stated goal of “going to Grand Canyon” (which was not even my destination)

doesn't capture my overarching *activity*, which is visiting the Page area, driving for a week to Arches, Utah, and then finally flying home via Grand Junction. The program needed to know my plan—locations, activities at each location, timing, who was involved (could my wife and I share a ride with friends already at the Grand Canyon?), etc.

Capping this experience, when I sent an email requesting customer support from WebTravel.com, I received the following response (vendor name has been edited):

Subject: Re: computer reserves non-rental car airports (KMM2768689V6442810KM)
Date: Fri, 09 Feb 2001 18:25:58 -0600
From: WebTravel <WebTravel@travel.com>
TO: <BCLANCEY@mail.arc.nasa.gov>

Due to unexpected system errors, we are unable to accept your email at this time. We hope to have the system back up soon. If you have a question regarding a reservation you have made, please call our customer service center at 1-888-709-5983.

Thank you,
WebTravel customer service

Unfortunately, the problem is not just that WebTravel provides no real travel service, but human travel agents fail, too. Here's what happened next. I went to an airline counter at an airport to attempt to resolve my situation. I had already learned that the closest airport to Page, Utah that had rental cars was Flagstaff, Arizona. I wouldn't have to pay the \$75 change fee per ticket because the new flights would cost \$350 less than the original. However, I wouldn't receive a refund either. I had already looked up the best flights, and presented them. The agent indicated that my preferred flight (to the required stop in Phoenix) was booked, but I could get on a later flight; this would require also changing the flight from Phoenix to Flagstaff. Finally, she handed me the tickets. I remembered that I should read them over before leaving the counter. But I was flustered, trying to find the flight information on the ticket, and then trying to evaluate the revised plan, since all the flight numbers and times were different from what I had carefully determined were best. I checked that the dates and locations were correct, and thanked her. (I was given \$200 coupons for future flights as a consolation for the lost \$350.)

However, when I was able to sit down and read over the tickets, I discovered something terrible: Not only did

we not arrive in Phoenix until 2 PM, there was a wait of more than three hours until the flight to Flagstaff! And we would not arrive at Page until well into dinner time, losing the afternoon activities we had planned for nearby Lake Powell. I found a map and determined that we could drive from Phoenix to Page in less time than it would take to fly to Flagstaff and then drive to Page. Knowing all of this, I might have chosen to fly from San Francisco to other locations in Utah that were almost as close as Flagstaff. But once again, the travel agent (remember, an actual human being) engaged in no background discussion about my actual trip—where I was going, when I wanted to get there, my tradeoffs on cost and driving effort and arrival time, etc. Here I had spent more than \$500 than was necessary and was faced with a four-hour unnecessary car trip, arriving half a day later than I desired!

Ironically, the special flight coupon I received as a consolation gift prominently displays the airline's mantra: "More CARE Plus." Simply caring at all would have been good enough.

3. Holistic analysis of travel semantics

Given the realities of everyday travel planning, computational talk about promising to make a reservation (cf. McCarthy, 2000, p. 1) appears silly. Ticketing has very little to do with the customer's perspective of what making reservations is about. A travel agent, properly concerned with my trip and not just my individual reservations, would ensure the completeness of my plan (am I ensured of transportation between all points?), timing constraints (is driving faster than flying?), and my activity goals (arriving in time to enjoy Lake Powell during the day; arriving in time for dinner).

A travel agent must understand how people spend time enjoyably—maybe driving is tedious, but are there places I might enjoy photographing along the way? Travel is not all just about cost and time optimization—that's the least of it.

So how would we improve WebTravel and the human travel agent? First, we would reorient how the software system is designed and how the person works. The work system design of people, roles, software, and conversation must reflect that the goal of a service provider is to develop a travel activity plan, not merely to reserve transportation tickets. The human travel agent's focus on ticketing is very clear; she has no interest whatsoever in my vacation, which is what I am actually needing to discuss. As put by Cerri et al. (2000), what is missing is a "shared ontology" used by "communicative agents."

Here the idea of speech acts (Searle, 1969), properly

interpreted, is useful: The travel agent must interpret my requests about flights and destinations as being tentative articulations, placeholders for specifying what I really need. When I say, “I want to fly to Flagstaff,” the travel agent should at least ask, “Is that your final destination for the day?” This simple question would reveal that I don’t want to go to Flagstaff at all on that day! The lack of understanding here of my interests, which is a routine deficiency in both human and computer conversations about tickets, is astounding when you think about it. And compounding matters, travelers have been coerced by reservationists to speak using their ontology of “flights,” “Saturday stay-over,” and “Class Y ticket.” Within the software communication and the asymmetric dialog controlled by the airline reservationists, there is virtually no investigation of the traveler’s intentions and activities.

A key problem is that customers do not know what questions to ask. I never thought to ask whether Grand Canyon had rental cars. A “non-defective” performance (cf. McCarthy, 2000, p. 27) by the travel agent would involve volunteering information about atypicality (“Did you realize that the Grand Canyon has no rental car locations?”). The form of communication required here goes well beyond Gricean conventions. The truth of individual assertions about flight availability and timing are meaningless if taking such flights would violate what I am broadly trying to accomplish. Yet ironically, the first level of logical analysis is well within a computer program’s capability. WebTravel could ask users, “How will you get from the Grand Canyon to Grand Junction?”

But a radical shift in design is required. Rather than conceiving web services only in terms of logistics—accessing information, interoperability, composing meanings within different ontologies, etc.—we require methods for understanding a person’s overall activity. Crucially, this requires techniques for articulating desires and negotiating goals incrementally, within the process of constructing alternative solutions (Schon 1987). In other words, we can’t first bundle up the “specifications” for the journey, and then produce an optimized plan. Because the traveler’s desires are shaped by the travel options available—which may be unavailable, a poor fit for the emerging schedule, or too expensive—each tentative schedule provides an opportunity for the schedule to modify goals, relax constraints, and discover unexpected value (e.g., visiting an Indian reservation for the first time).

How can we design software to do what people don’t do well? The solution is of course the essence of artificial intelligence programming—model-based problem solving. The travel agent should construct a model of the traveler’s trip. Again, such a model is constructed in the

process of searching, filtering, and composing alternative services, such that the model is created, refined, and validated in a conversation that changes the goals, constraints, and how solutions are evaluated.

In summary, mapping stated goals about destinations into guaranteed ticket reservations (cf. McCarthy, 2000) is an absurd view of the problem to be solved in providing travel services. A joint construction is required, a conversational interaction that bit-by-bit draws out from me what I am trying to do and draws out from the travel agent the space of relevant options (Brown & Duguid 2000; Gasser 1991; Greenbaum & Kyng 1991; Suchman 1987; Vicente 1999; Winograd & Flores 1986; Wynn 1991). In developing a model, the travel agent learns about my activity plans: lodging, sightseeing, recreation, relaxation, and all the other people involved (children, friends, pets, etc.). A timeline is developed that includes activities of “getting settled,” “enjoying a resort,” “photographing while driving,” eating, sleeping, etc. Notice that these are facts about a social world, not just a physical world. This social world is always personal—most travel plans are customized—though they may fit norms and even highly probable routes.

A candidate representational framework for formalizing such a travel activity model is the Brahms multiagent simulation system (Clancey 1998; Sierhuis 2001; Sierhuis, et al. 2003; Clancey in press). Brahms was conceived to complement business process modeling tools by representing how work actually gets done. As a model of practice, in contrast with formal processes, Brahms simulations emphasize circumstantial interactions (e.g., how placement of people and tools affects what information is shared or how long a job takes). People are modeled as located, having multiple group identities, having behaviors that dynamically blend the conventional activities of multiple group memberships, and having perceptual interests coupled to activities.

If a travel agent (computational or person) were to construct a Brahms activity model of the customer’s trip, then it might be able to determine transportation options that truly satisfy the customer. However, the agent must use conversational techniques and representational tools (e.g., maps, timelines, and diagrams) to help the customer recognize the implications of how options interact (“No, of course I don’t want to fly to Flagstaff if I must wait in Phoenix for over three hours. I’m trying to get to Page by 3 PM!”).

The problem involves not only expressing information (e.g., available schedules) in a comprehensible way, but also instructing the traveler about how services interoperate (e.g., you can move from an airport to rental car provider using a free shuttle bus). In effect, the travel

plan has a corresponding procedure that the traveler must learn (e.g., “When you get to the gate, look for the DesertAir representative”). At the same time, the travel agent learns about the traveler’s values. As is well known in instructional research, personalization (as in student modeling, Clancey 1986) is aimed at much more than classifying the user and selecting from fixed choices to be delivered as lessons. Rather the modeling process itself is a problem-solving endeavor, leading to constructing a potentially quite complex process with causal and temporal constraints, built out of simpler components (cf. classification vs. construction, Clancey 1992). More generally, the modeling dialogue must be structured, planned, and negotiated, as revealed in a variety of research from medical diagnosis to design tools, but explored especially well in intelligent tutoring systems (Psozka, Massey, Mutter 1988).

Finally, the travel-planning problem can be framed in a yet broader way, viewing over a longer time scale what it means for someone to be a customer. The objective of the travel service is not only to sell a ticket and make a profit, but also to secure future business. The travel agent must establish a relationship with the customer. That is, the travel agent must consider his or her own identity as a service provider, as viewed by the customer. The idea of speech acts might need to be extended to take into account that service-oriented communications transcend the particular trip being discussed. Each interactive moment is part of an overarching activity in which the customer is conceiving, “You are (or are not) my preferred service provider. This is (or is not) how I will in general construct travel plans.” Thus a “successful” speech act (cf. McCarthy 2000, p. 27) is one that accomplishes the broader goal of establishing long-term trust, usually by showing that the travel agent’s capabilities are very broad and general—going beyond what is stated by the customer, surprising and delighting him or her.

In our example, by providing coupons as a form of “good will,” the airline supervisor showed that she cared about my future business. She didn’t merely recite the \$75 penalty fee rule or say that I had been amply warned by the computer program about non-refundability. She recognized that I had not been served well and had paid far too much for my flights.

4. Conclusions

Moving from the simple idea of hyperlinks on the Internet to web semantics and grid computing necessarily involves at first an emphasis on the logistics of how to

relate diverse, independently created resources. The vision is to make it easier for people to benefit from a vast network of facts, media, presentations, and computing resources. But interposing agents between people and the web replaces one access and interpretation problem by several more: Besides the problem of getting software agents to find and understand web content, we now have the problem of getting the software agents to understand people, just as people now need to understand the agents (Figure 1).




Figure 1. In the original Internet search model (top), people interpret filtered results (web pages, products) retrieved from the web. In the ideal “Grid services” model (bottom), software agents perform these search and filter functions, while providing higher-level services to people in a conversation requiring mutual learning.

The real-world examples we encounter everyday should be sobering: If trained human agents cannot easily meet our needs, how will computer systems do better? The experience with WebTravel reveals that the service must be conceived in terms of what the person is trying to do (the broader activity and its motives), not the local, and superficial aspects of web transactions. And this requires constructing a model of the person’s activity—ironically something a computer program might do more patiently and cost effectively than people today who are travel, real estate, and investment agents.

To summarize the travel agent example, making a reservation is a means, not an end in itself. A distinction must be drawn between providing confirmed tickets at the least cost (a simple view of a web-based transaction) and getting me to my destination in time for dinner with my companions, with consideration for cost, driving distance, rental car expense, etc. (my activity). The problem from the customer’s perspective is not merely to hold confirmed reservations but to have a complete travel plan that will be a happy experience. The task from the travel agent’s perspective is not “ticketing” (as it is so often described by the airlines), but helping the customer

understand what they want and providing services that will connect everything together in an enjoyable way.

Perhaps the most difficult challenge in developing an automated travel planning service is that travel constraints and preferences usually cannot be formulated up front. The process of articulating constraints is necessarily interactive and iterative. A seed goal produces a set of options and tentative plan, which are evaluated for implications that were previously unanticipated. New preferences or constraints are articulated and the options and plan are revised, etc.

What is missing from the method of relating requests and promises, the most common technical approach for coordinating goals and services, is that requests are often just tentative placeholders for unarticulated (non-verbal) beliefs, desires, and concerns. Helping the traveler to describe his or her intentions to create a good experience is the actual service required.

The interaction between the person and the travel agent can be phrased as, “Do what I will prefer (once I have all the information), not what I say I want now.” The program must be proactive to model and understand the person’s activities—what they will be doing before, during, and after the transportation event. Reservations for a trip constitute a total plan, which cannot be evaluated solely by time or cost optimization. Relating different ways of framing a situation is the essence of a holistic analysis (Wynn 1991; Clancey 1995).

Certainly, some of the problems we experience today with human agents stems from inadequate computer tools, which shape how a business views the roles of their workers (Wales et al. 2002). The human airline reservationist is using a tool not much different from WebTravel, and is coerced by its design to converse in the ontology of ticketing, not trip planning. A crucial conclusion is that the opportunity and problem for Grid computing and the semantic web community is to not merely replicate what human service providers do, but to discover and address what services people really want to buy. Accordingly, terms like “educational material” and “learning” should be viewed broadly, not just in terms of presentation or explanation, but what are people trying to accomplish in their lives by means of the present inquiry. Indeed, a good design heuristic may be to make it possible for a person to interrupt an interaction and express a conflict with what he or she is trying to accomplish and/or introduce a constraint that better defines high-level goals (e.g., “But wait, how will this get me to Page in time for dinner?”).

Perhaps the most specific lesson for today’s system developers is to beware definitions of “successful performance” that are based on narrowly framing the

problem to be solved—especially technical definitions based on legacy software that has produced unsuccessful business practices. Fulfilling reservation commitments, the emphasis of Elephant 2000, is probably the narrowest view of a travel service provider. I had high confidence that the airline would allow me to get onboard the airplanes I reserved for August 3, 2001. But I also knew that I was going to a city where I don’t want to go, that I paid \$350 too much for the flight, that I would throw my Phoenix-Flagstaff tickets into the garbage, and I would have to drive 4 hours to arrive at my desired destination, arriving too late to enjoy the day.

In artificial intelligence research, people sometimes refer to what is missing in WebTravel and the reservationists I encountered as “common sense.” But if the capability and knowledge were so common, we would not be so routinely frustrated and poorly served. Very likely some of the deficiencies in human performance at the airlines stem from an overworked, undertrained work force, and a business model that focuses on moving planes and bodies from place to place, rather than constructing travel plans for loyal customers (Wales et al. 2002). A “happy performance” (Austin, 1962) must be one that is helpful and useful—one that constructs a plan that will be a happy experience and not just satisfies legal obligations and notions of reservation “correctness.”

The analysis presented here does not mean that the mechanics addressed by the “Elephant 2000” or “Language for Interactions” papers and all of the related problems researchers are now handling to develop Grid technologies are unnecessary. Rather, a holistic analysis complements the analysis of mechanics, to consider the kind of content, the models that such computer programs must develop through interactive communication, in order to know what to search, filter, compose, etc. from available resources. To adapt McCarthy’s phrase, “computerizing a communication” should entail viewing the interaction as a process of constructing an activity model. This process is necessarily iterative, based on tentative constraints and plans, and considering larger issues of business relationships. In the case of a computerized travel agent, mimicking human behavior and social customs would be a bad idea—we need to do much better than that.

Acknowledgments

This paper was originally presented at the San Marino Summer School, June 2001, “Semantics and Pragmatics on the Web: Synergies between Philosophy and Computing.” To appear in P. Ritrovato and S. Cerri (editors), *Towards the Learning GRID: Advances in*

Human Learning Services, in the series, *Frontiers in Artificial Intelligence and Applications*, IOS Press. Charlotte Linde and Chin Seah provided helpful comments. This work has been supported in part by NASA's Computing, Communications, and Information Technology Program, Intelligent Systems subprogram, Human-Centered Computing element.

References

- Austin, J. L. 1962. *How to Do Things with Words*. Oxford.
- Brown, J. S. and Duguid, P. 2000. *The social life of information*. Boston: Harvard Business School Press.
- Cerri, S. A. 2002. Human and artificial agent's conversations on the grid. In *1st LEGE-WG International Workshop on Educational Models for GRID Based Services*, Lausanne, Switzerland. September. *Electronic Workshops in Computing (eWiC)*.
<http://ewic.bcs.org/conferences/2002/1stlege/session3/paper1.htm>
- Cerri, S. A., Sallantin, J., Castro, E., Maraschi, D. 2000. Steps towards C+C: a language for interactions. In S. A. Cerri and D. Dochev (eds.), *Artificial Intelligence: Methodology, Systems, Applications* (AIMSA 2000), Lecture Notes in Artificial Intelligence 1904, Springer Verlag. pp. 33-46.
- Clancey, W. J. 1986. Qualitative student models. *Annual Review of Computer Science*. Palo Alto: Annual Reviews Inc., pp. 381-450. Also excerpted in J. Self (Ed.), *Artificial intelligence and human learning: Intelligent computer-aided instruction*, London: Chapman and Hall, 1987, pp. 49-68.
- Clancey, W. J. 1992. Model construction operators. *Artificial Intelligence*, 53(1):1-115.
- Clancey, W. J. 1995 Practice cannot be reduced to theory: Knowledge, representations, and change in the workplace. In S. Bagnara, C. Zucermaglio, and S. Stucky (eds.), *Organizational learning and technological change*. Berlin: Springer-Verlag.
- Clancey, W. J., Sachs, P., Sierhuis, M., & van Hoof, R. 1998. Brahms: simulating practice for work systems design. *Int. J. Human-Computer Studies*, 49, 831-865.
- Clancey, W. J. 1997. The conceptual nature of knowledge, situations, and activity. In P. Feltovich, K. Ford, & R. Hoffman (eds.), *Human and Machine Expertise in Context*, pp. 247-291. Menlo Park, CA: The AAAI Press.
- Clancey, W. J. 2002. Simulating activities: Relating motives, deliberation, and attentive coordination. *Cognitive Systems Research* 3(3):471-499, September.
- Clancey, W. J. in press. Cognitive modeling of social behaviors. To appear in R. Sun, *Cognition and multi-agent interaction: From cognitive modeling to social simulation*. New York: Cambridge University Press.
- Gasser, L. 1991. Social conceptions of knowledge and action, *Artificial Intelligence*, 47(1-3)107-138., January.
- Greenbaum, J., and Kyng, M. (Eds.) 1991. *Design at work: Cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- McCarthy, J. 1998[1989]. Elephant 2000: A programming language based on speech acts. Stanford University unpublished manuscript: <http://www-formal.stanford.edu/jmc/elephant/elephant.html>, appeared: 6 Nov 1998, written 1989.
- Mostow, J. and Tedesco, P. (Eds.). 2004. *Grid learning services*. Workshop proceedings from ITS 7th International Conference on Intelligent Tutoring Systems, Maceio, Brasil.
- Potoka, J., Massey, L. D., Mutter, S. A. (Eds.). 1988. *Intelligent tutoring systems: Lessons learned*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schön, D.A. 1987. *Educating the reflective practitioner*. San Francisco: Jossey-Bass Publishers.
- Sierhuis, M. 2001. "Modeling and simulating work practice," Ph.D. thesis, Social Science and Informatics (SWI), Univ. of Amsterdam, SIKS Dissertation Series No. 2001-10, Amsterdam, The Netherlands, ISBN 90-6464-849-2.
- Sierhuis, M., W.J. Clancey, Seah, C., and Trimble, J., and Sims, M.H. 2003. Modeling and simulation for mission operations work systems design. *Journal of Management Information Systems*, 19(4), 85-128.
- Suchman, L. A. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge: Cambridge Press.
- Vincente, K. J. 1999. *Cognitive Work Analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Erlbaum.
- Wales, R., O'Neill, J., and Mirmalek, Z. 2002. Ethnography, customers, and negotiated interactions at the airport. *IEEE Intelligent Systems*, 17(5): 15-23.
- Winograd, T. & Flores, F. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Norwood: Ablex.
- Wynn, E. 1991. Taking practice seriously. In J. Greenbaum and M. Kyng (Eds.), *Design at work: Cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 45-64.