

When Is Assistance Really Helpful?

Wayne Iba*

Department of Mathematics and Computer Science
Westmont College
955 La Paz Rd.
Santa Barbara, California 93108
IBA@WESTMONT.EDU

Abstract

We regularly operate under the notion that one agent assists another when the first does something for the second. However, the story behind this is much more complicated. In this position paper, we explore two questions: How can we evaluate the quality or goodness of a particular instance of assistance? and How can we design agents to initiate and provide “good” assistance? In asking these questions, we are assuming two things. First, performing a task for another agent is not always helpful to that agent even if the task really needs to be done. But second, a helper that does not have the skill or resources to accomplish a specifically requested task may still provide real help. Our experimental work attempts to address the fundamental elements of helpful assistance. One surprising initial result reminds us of an common understanding about help: providing assistance is a two-way street.

Measuring the Helpfulness of Assistance.

Providing assistance can be one of those activities that is performed for the benefit of another agent without sufficient reflection. We want to better understand the nature of assistance that is truly helpful.

At one point or another, everyone has experienced service that was exceptional or memorable – either because it was unusually good or unusually bad. Either at a restaurant, an auto mechanic, or a health clinic, we have been “assisted” by someone who was inattentive, incompetent, unscrupulous, or even all three at the same time. Likewise, most of us have encountered situations where the service provider anticipated our needs, took care of those needs efficiently, or treated us with respect and courtesy. This common experience forms the basis of my primary claim, which we consider to be self-evident:

Assumption 1: All assistance is not created equal.

Or in other words, not all assistance is actually helpful. If we are willing to grant this assumption, then at least two questions should immediately spring to mind: How can we evaluate the quality of assistance? and How can we construct

agents that truly are helpful? Both questions have implicitly been at the core of essentially all work in the area of intelligent assistants. However, the implicit answers to the first question are sometimes overly narrow.

Question 1: How can we evaluate the quality or goodness of assistance? A traditional answer to this question focuses on the assistant alone and consists of specifying a task that needs to be done and then measuring how often or how completely an assistant accomplishes the given task. Essentially, the approach measures the *competence* of a service provider’s assistance in a given encounter. Unfortunately, this addresses only a single facet of assistance. We identify at least five other dimensions that influenced the positive or negative assessment of an interaction. Other dimensions along which we can view an assistant include: attention, anticipation, persistence deference and integrity.

Briefly, each of these other dimensions allow us to describe characteristics of assistance that contribute to our assessment of a particular engagement as being helpful or not. For instance, we expect an assistant to pay *attention* to our current situation and track changes in our needs over time. We would also hope that our assistant can *anticipate* our needs and satisfy those needs or perform tasks that contribute to our goals in a timely fashion and perhaps without being asked.¹ Once given a task, we hope that an assistant will be *persistent* in working on and completing the task. Also, an assistant should operate with some *deference* toward those it is intending to assist; this might amount to giving the goals of the recipient a higher priority than the goals of the assistant itself. Finally, an assistant should operate with *integrity*; that is, information and services must not be misrepresented and payments charged for the assistance should be fair. Each of these, together with the *competence* of an assistant, combine to determine the overall helpfulness of the interaction.

Reflection on the dimensions suggests at least two corollary assumptions to our central claim. These assumptions may provide further constraints on the design of assistants.

¹In this context, we do not wish to address the thorny issue of autonomy or initiative; this certainly requires additional consideration. Minimally, we can say that accurate anticipation is better than faulty anticipation whether or not the assistant is able to act independently on that anticipation.

*We are grateful to Westmont Provost, Shirley Mullen, and the Westmont Professional Development Committee.
Copyright © 2007, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Assumption 1.1: Competent helpers are not enough.

Good assistance requires more than simply performing a task for another agent. This is true even when the task in question really needs to be done. For example, if two sub-tasks must be completed in a specific sequence, then an assistant that arbitrarily does one task may not be helping. Or two sub-tasks may have significantly different difficulties and help is needed with the most difficult one; in this case, an assistant that performs the easier task may not be as helpful as the one that works on the more difficult task. Again, performing a task may bring a user enjoyment and having the task completed by an assistant could deprive the user of that satisfaction. Thus, we have several examples where an assistant may be competent, anticipatory, and persistent, yet may fail to be helpful. Note then that these features – competence, anticipation, and persistence – are not sufficient for helpful assistance.

But if the examples above are not helpful, what is missing? In each case, an assistant that was attentive to the goals of the one receiving the help would have avoided the problems identified. That is not to say that attention is then the key to good service; other examples involving attentive but incompetent assistants would demonstrate the non-sufficiency of attention. To sharpen this point, let us consider situations where we might be surprised to find helpful assistance.

Assumption 1.2: Incompetent helpers can be helpful.

The long tradition of apprenticeship provides adequate evidence that assistants with less than masterful skills can provide helpful service while they are developing their expertise. A child helping its parent on a project around the house can also provide real help by fetching tools or parts. Together with the previous assumption, this suggests that competence is neither a necessary nor sufficient feature of good assistance.

We claim that there are no set of necessary and sufficient features that define “good” assistance. If we want to carefully and fully evaluate the quality of service delivered by our intelligent assistants, then we need to consider all of these dimensions. However, in addition to providing a strategy for evaluating assistants, these dimensions also suggest behaviors or inclinations that assistants should display, and thus contribute to an answer to our second question.

Question 2: How do we design agents that can deliver “good” assistance? Of course, we ultimately want to design and implement assistants that can be truly helpful. The dimensions identified above may be generally useful for evaluating the assistance provided by a given agent compared to that of another. They may also suggest designs for how to implement assistants. Most simply, an assistant design should respond to each of the six dimensions. However, because they are neither necessary nor sufficient, assistants without one or more of these capabilities may still prove helpful. Nevertheless, holding all other things equal, an improvement along one dimension leads to an overall improvement in helpfulness.

Empirical Work.

We have conducted our work within the context of the MÆDEN simulated environment (Iba & Burwell 2005a; 2005b). This multi-agent environment supports a variety of tasks, agent architectures and communication protocols. Using the GARCIA agent framework (Iba & Holm 2006), we have implemented agents that ask for and deliver assistance in the midst of problem solving tasks. The agent framework addresses all of the dimensions discussed above; specifically, we can vary competence, attention, anticipation, persistence, deference and integrity.

Our previous experiments have addressed competence, persistence, and integrity. We evaluated problem solving success rates and efficiency while varying the skills that an assistant possessed. As expected, our dependent measures (number of problems solved and average cost to solve them) improved with increasing skills. We also varied parameters that control persistence and integrity with the obvious anticipated results. So far, we have mostly focused on an incompetent agent seeking help from an assistant that is more skilled to varying degrees.

A somewhat surprising result that emerged from our previous work highlights the role of the recipient in the exchange of assistance. Based on a level of patience, an agent may request help but later give up on the assistant and continue trying to solve the problem on its own (possibly asking for help again later). We observed that the agent’s commitment to the assistant positively influenced the overall success rate and efficiency. Although it is commonly understood that unteachable students cannot be helped, we were not expecting to encounter this effect in our tests.

Thus, we intend to extend our dimensions to include factors pertaining to the agent receiving assistance. At the very least, we can say the quality of the assistance provided will depend on the recipient’s willingness to accept and integrate the assistant’s work. However, we suspect that there are other dimensions that pertain to the service recipient that need to be identified. Along these lines, our current experiments vary the competence of both the assistant and the agent receiving the help; among other results, we hope to clearly demonstrate the beneficial assistance that can be provided by a helper that is less competent than the agent requesting the help.

References

- Iba, W., and Burwell, N. 2005a. Building a testbed for studying service. In *Proceedings of the Stanford Spring Symposium on Persistent Assistants: Living and Working with AI*. Palo Alto, CA: AAAI Press.
- Iba, W., and Burwell, N. 2005b. Studying service: An exploration of the costs and benefits of assistance. In *Proceedings of the 18th International FLAIRS Conference*. Clearwater, FL: AAAI Press.
- Iba, W., and Holm, J. 2006. Assistance: Is it better to receive than to give? In *Proceedings of the 2006 International Conference on Artificial Intelligence*, 84–87. Las Vegas, NV: CSREA Press.