Involving Intelligent Assistants in Active Human Communication

Donald J. Patterson
University of California Irvine

Abstract

Intelligent assistants that support human communication need to respect the difficulty of understanding the context surrounding the interchange. Rather than attempting to directly communicate for a user, intelligent assistants should support decision making on the part of the involved parties so that complex social negotiations are preserved. We describe an intelligent assistant that does this for instant messaging called Nomatic*Gaim.

Introduction

There are many examples of learning algorithms that support communication exchanges between individuals. Spam filters, grammar checkers and fundamental TCP packet routing algorithms are some examples. There are few examples, however, of AI algorithms or intelligent assistants that users rely on during real-time communication. An excellent example of an application that would benefit from such support is instant messaging (IM).

IM was originally architected in a world in which users went online in at most two places, home and work. With the arrival of IM on cell phones, the typical IM-enabled distinctions of being online/offline no longer give enough information for communication partners to understand whether or not it is appropriate to initiate a communication.

Within the old paradigm of desktop computers, even in the presence of online status indicators, availability negotiation consumed 13% of IM communications (Handel & Herbsleb 2002). As users are increasingly mobile and increasingly "always online" it is reasonable to expect that this percentage will increase.

As a result it is necessary to give users assistance in negotiating communication availability. A classic AI solution would be to craft an algorithm that would automatically learn and then set the availability status of a user. We claim, however, that this is the wrong approach. Deciding whether or not one is available is a complex social negotiation that is best handled by people. This paper proposes using an intelligent assistant to communicate context information in an IM side-channel on behalf of a user to reduce the need to engage in discussions about availability.

Palen & Dourish frame the problem of determining availability as one of boundary negotiation (Palen & Dourish 2003). It is a constant subconscious process that people do on a regular basis that requires complex calculations that are unlikely to be successfully done by an intelligent assistant anytime soon. In contrast people do this without a second-thought. The problem can be exemplified by the question of whether one should interrupt someone with a phone call during a movie. In most cases the right answer would seem to be "no," unless the call is for a doctor, unless it's about a billing question not an emergency, unless the doctor isn't actually watching a movie, but is just picking up her son, etc. Making a good decision is easy as an informed person, but crafting an effective algorithm for the same task seems daunting.

Palen & Dourish might argue that one of the reasons why it is difficult for intelligent agents to participate in automatic communication on behalf of users is because it puts them in the position of negotiating boundaries: a place where meaning and nuance dominate. In such a situation, current generations of assistive agents have no place. Instead they should be supporting the decision making process so that users can negotiate the boundary more efficiently.

A better approach is to provide the users with information about the context of the parties so that that information can be used as part of the boundary negotiation. Knowing that the doctor is at the movie theater might be all that a caller needs to know in order to make a good decision about interrupting her. By giving the responsibility of reporting context to an intelligent assistant, the human boundary management negotiation is supported but not replaced. In the case of IM, we propose Nomatic*Gaim as an example of such a solution.

Nomatic*Gaim

Nomatic*Gaim is an open-source, multi-operating system, multi-protocol IM client that is based on the open-source project gaim. It has been enhanced to support more complex presence indicators than just online/offline. Instead it takes an unusual approach by revealing the current location of the user on the presence/status line.

The input to Nomatic*Gaim is a set of Wi-Fi access points (APs) that can be seen from the user's current computing platform. The APs need not be "open" in the sense that they provide Internet connectivity without a password, they simply need to be broadcasting their existence, as most do, for Nomatic*Gaim to use them. After finding a collection of APs that are visible in the environment, Nomatic*Gaim looks up their position in a database of known APs and infers the user's current location using beacon positioning algorithms (e.g., (LaMarca *et al.* 2005; Letchner, Fox, & LaMarca 2005)). From this data, the user's approximate latitude and longitude can be obtained.

It would be possible to simply list the latitude and longitude on the presence status indicator of an IM client. However, although this is a great deal of information, it is not very useful in determining whether or not it is appropriate to interrupt a user for an IM chat. Enter the intelligent assistant.

Instead, the SSIDs of the APs, denoted as a, the current day and time, t, and the latitude and longitude, x, are used as observed features in a machine learning algorithm that uses a trained model, M, to obtain a guess about the optimal semantic label, L^*, to use to describe a user's current place (versus position, see (Hightower 2003)). This place becomes valuable information that can then be used by buddies on a user's IM buddy list to decide whether or not to initiate a communication with an individual. *Rather than deciding availability we instead support the IM users in making the decision themselves.* We use the formulation below to account for positioning error in the neighborhood around the true position, x' :

$$L^* = \operatorname*{argmax}_{L \in Places} \int P(L \,|\, a, t, x', M) P(x' \,|\, x) dx' \quad (1)$$

The training data comes from two places. First the best data is going to come from users themselves. We hypothesize that the amount of time currently spent negotiating availability in IM is enough incentive for a user to enter place information into their IM client because doing so will reduce unnecessary interruptions. When a user does this, the information is sent to a central repository for collecting training data. The user is then further incentivized to provide this information because it is only necessary to enter the information once per position. The second source comes from other users. When a user arrives at a new place it is possible to make a guess about the correct place label to use for a given position simply based on the training data provided by other people that have been at the same location. Assuming a user's model and "other's" model are independent allows the following decomposition:

$$P(L \,|\, a, t, x, M) = P(L \,|\, a, t, x, M_{user}) P(L \,|\, a, t, x, M_{others}) \quad (2)$$

Initially there is no user training data, so $P(L \mid a, t, x, M_{user})$ is uniform and uninformative and the label is chosen entirely according to M_{others}. However, it is unlikely that the data provided by other people will be as good as the data provided by a specific user, so the data should be biased toward training data provided by the user himself. If there is no training data, then no label will dominate the probability calculation. In between these extremes, an intelligent assistant could display the following behavior:

A user arrives at a brand new location and opens her laptop. Nomatic*Gaim localizes the user's position and sends the information to the centralized machine learning repository. Rather than having the assistant respond with, L^*, it can respond with the list of highest rated place labels, $(\{L^* = L_0\}, L_1, L_2 ...)$ and the associated probabilities, $(P(L_0 \,|\, a, t, x), P(L_1 \,|\, a, t, x), P(L_2 \,|\, a, t, x), ...)$ calculated according to equations 1 and 2.

The Nomatic*Gaim assistant can now take one of four options under the assumption that $\alpha > \beta > \gamma$:

- If $P(L_0) > \alpha$ then Nomatic*Gaim will automatically set the place context for the user in their status line
- If $P(L_0) > \beta$ then Nomatic*Gaim will automatically set the place context for the user and ask for the user to confirm the setting.
- If $P(L_0) > \gamma$ then Nomatic*Gaim will open a dialog box with the list of locations, $L_0 ... L_n$, ordered according to their associated probabilities in a drop down box for user selection.
- Otherwise, Nomatic*Gaim will open the dialog box for the user and allow him to manually set the current location place.

In any case, if the user corrects the assistant's guess, then that correction becomes training data which alters M_{user}, the next time the user opens their laptop in the same position.

In this way an intelligent assistant is able to support efficient communication initiation, between trusted parties on an IM buddy list, but avoid the complex problem of deciding whether or not someone is available.

References

Handel, M., and Herbsleb, J. D. 2002. What is chat doing in the workplace? In *CSCW '02: Proceedings of the 2002 ACM conference on Computer supported cooperative work*, 1–10. New York, NY, USA: ACM Press.

Hightower, J. 2003. From position to place. In *Proceedings of The 2003 Workshop on Location-Aware Computing*, 10–12. part of the 2003 Ubiquitous Computing Conference.

LaMarca, A.; Chawathe, Y.; Consolvo, S.; Hightower, J.; Smith, I. E.; Scott, J.; Sohn, T.; Howard, J.; Hughes, J.; Potter, F.; Tabert, J.; Powledge, P.; Borriello, G.; and Schilit, B. N. 2005. Place lab: Device positioning using radio beacons in the wild. In Gellersen, H.-W.; Want, R.; and Schmidt, A., eds., *Proceedings of PERVASIVE, The Third International Conference on Pervasive Computing*, volume 3468 of *Lecture Notes in Computer Science*, 116–133. Springer.

Letchner, J.; Fox, D.; and LaMarca, A. 2005. Large-scale localization from wireless signal strength. In Veloso, M. M., and Kambhampati, S., eds., *Proceedings of The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference*, 15–20. AAAI Press AAAI Press / The MIT Press.

Palen, L., and Dourish, P. 2003. Unpacking "privacy" for a networked world. In Cockton, G., and Korhonen, P., eds., *CHI*, 129–136. ACM.