

A Flexible Human Agent Collaboration (HAC) Framework for Human-Human Activity Coordination (H2AC)

Wei Chen and Vikram Manikonda

Intelligent Automation Inc.
15400 Calhoun Drive, Suite 400, Rockville, MD 20855
(wchen, vikram)@i-a-i.com

Edmund Durfee

EECS Department, University of Michigan
1101 Beal Avenue, Ann Arbor MI 48109
durfee@umich.edu

Abstract

This paper presents an innovative shared intelligence framework for flexible human-agent collaboration. By ‘flexible’, we mean the level of collaboration is selected through a negotiated and iterative human-agent collaboration (HAC) process, which ensures that the needs of a system user are balanced with the availability of suitable experts to form and maintain an expert team that supports the user’s decision-making. To enable flexible human-agent collaboration to develop expert teams, our framework suitably represents team members’ roles and capabilities, represents tasks and environments, reasons simultaneously about task decomposition and team formation, and coordinates human experts and software agents. We describe our prototype implementation as it applies to a military medical service scenario.¹

1 Introduction

High quality decisions often require the combined intelligence of multiple agents, whose knowledge, expertise, and capabilities are complementary. Cooperative decision making is hard for people who have different languages, priorities, and incentives, and is even more difficult with the participation of non-human agents with inscrutable behaviors to people. Our work is in this area of distributed intelligent systems comprised of both human and artificial agents. A fundamental challenge is in developing artificial agents that collaborate in a natural way with humans in order to utilize individual expertise to benefit the whole. Current multi-agent decision support systems expect a user to be a system expert (e.g. knowing the parameters, operation flow/steps and limitations of the system), when in reality a user (e.g., soldier in the field) most in need of support might lack such expertise and could be facing a stressful, time-critical decision problem.

In this paper, we carve out a constrained, but still widely-applicable, role for human agent collaboration (HAC) for managing human-agent teams. Our approach gets traction on the overall problem of cooperative intelligent behavior by splitting it into two pieces, where collaborative decision-making in a particular domain falls on the shoulders of human experts who work together

(possibly through agent intermediaries), but where the generation and flexible adaptation of the expert team is conducted through HAC.

We explore these concepts in the context of a simulated military medical service scenario, where a surgeon facing a potential life-altering decision for a wounded soldier (e.g., deciding whether to amputate a leg) can make a better, more confident decision after conferring with other humans with medical expertise (e.g., a rehabilitation specialist) and expertise about the soldier (e.g., a company chaplain). Our goal is to demonstrate the effectiveness and practicality of our new HAC framework and its application to other problems requiring the rapid and flexible teaming of human experts to address unusual situations.

In the next section, we summarize the wider context of different types of distributed intelligent systems to clarify where our work fits. Then we describe more of the scenario and its consequent problems to motivate and clarify the HAC problem for human-human activity coordination (H2AC) in Section 3. We present the general HAC framework with a focus on the HAC process in Section 4. Section 5 gives concluding remarks.

2 Background

Related work to HAC and H2AC can be summarized in Figure 1. The edges between H (human), A (agent), and M (mediator) are the focus of discussion. The vertical edges involve collaborative decision making among Humans and Agents possibly involving Mediators; the horizontal edges correspond to the (possibly Mediated) coordinated interactions between Humans, and between Agents. Schematically our work focuses on using the vertical edges and the lower horizontal edge as a means to provide flexible activity coordination. The upper horizontal edge, representing direct human-to-human interaction, is underpinned through a collaborative agent network infrastructure.




Figure 1. Schematic relationship among humans (H), agents (A) and mediators (M)

Human Agent Collaboration (HAC)

Human agent collaboration is the effort carried out by human experts and software agents in pursuing mutual

¹ This material is based upon work supported by DARPA SBIR program under contract # W31P4Q-07-C-0090.

goals. The HAC explored in this paper refers to the collaboration for decision-making processes at the coordination meta-level, not at the domain task level; thus our proposed effort is general and could be applied to various scenarios. HAC addresses the following questions: How should the human collaborate with software agents to solve problems that are beyond their individual capabilities? How do the human and agents converge on sufficient shared understanding of the problems that they are solving, and on the ways in which each is contributing (and expecting the others to contribute) to solving the problems? How do they monitor individual and collective progress on solving problems, and redefine tasks or redirect efforts to improve collaborative performance?

Providing technological support for human decision making extends all the way through old-fashioned tools (e.g. pencil and paper for solving an arithmetic problem) through modern conveniences (e.g. PDAs for managing tasks) and advanced technologies (e.g. decision support systems). Among the many pioneering efforts in HAC were the Pilot's Associate (Banks & Lizza, 1991 [1]) and Tactical Assistants for Interaction Planning and Execution (Durfee et al, 1997 [6]), and recent efforts in developing computational agents that work with people to solve complex problems, e.g. CALO (Berry et al, 2006 [2]).

Human-Human Activity Collaboration (with potential agent assistance) (H2AC)

H2AC concentrates on the process of connecting suitable human experts with information requesters (the horizontal dimensions in Figure 1). Human experts need to be identified, scheduled, and interfaced with users (and other experts) in response to problem needs. Software agents can act as assistants to facilitate this process. H2AC addresses the followings: How can humans working on related tasks but in different contexts (physically separated, from different backgrounds, etc.) coordinate their activities to achieve their shared objectives efficiently and effectively?

As was the case with individual humans, decisions across humans (about how they should act as individuals to pursue collective goals) have been supported by technology since the dawn of time. Language is of course a key technology that enables human activity coordination, but over time more effective technologies for coordinating human activity have been developed, ranging from the use of written messages, to the use of cell phones and wireless networks to allow people to be in touch with each other more immediately and deeply than previously possible. Pushed further, technology not only supports human activity coordination but also promotes it. A variety of agent-based systems for H2AC have been proposed and prototyped, e.g. developing algorithms and protocols to enable agents to solve extremely complicated coordination problems in a distributed manner (Wagner, 2005[7]).

3 Battlefield Medicine Application Domain

Given the above definitions, we now ground the problem of using HAC to solve the meta-level coordination problem of H2AC in the context of a specific motivating

application, which is to support people providing time-critical medical care in battlefield environments.

In general, a mobile military hospital near the front line of a battlefield will be populated by a number of humans with different expertise (physicians, surgeons, nurses, clerical staff, chaplains, etc.). Moreover, a non-local network of human specialists might be (sporadically) available for consultation as needed. Meanwhile, a number of non-human "agents" are also available for tasking at the hospital, including agents capable of sensing and monitoring patient vital signs (heart rate, blood pressure, etc.), retrieving medical information (patient history, established treatment guidelines, etc.), providing decision support (diagnostic workflows, data interpretation, etc.), and managing information flow and coordinating activities within and between the hospital and larger military theater.

A general medical response flow is as follows. A case (a wounded warfighter) is transported to the medical facility by a medical vehicle, or medics or peers. First, the case information (wound information) is gathered by a triage nurse supported by sensing and diagnostic agents. The nurse/agent team assigns a priority to the case according to injury severity score (ISS)². Based on previous experience or by enlisting the aid of decision-support agents, an initial diagnosis is made, and one or more candidate treatment strategies are developed. A treatment plan for which the necessary human agent team can be formed is selected, where the roles, capabilities and availabilities of the candidate team members are checked against a registration database. Notably the generation of a treatment plan and the formation of a mixed human agent team is iterative (see section 4.4). The medical treatment will be provided after proper preparation. If without complication, the treatment finishes successfully. The team is dismissed and the team members are available for other cases.

The HAC process for converging on human-to-human consultation plans is not restricted to any specific scenario. E.g., a warfighter chasing an insurgent into a religious compound faces the same needs for rapid consultation with experts on religious customs, appropriate phrases to say, political ramifications, etc. We focus on collaboratively forming human agent teams for meta-level coordination. The *key idea* lies in the separation of functionalities: *software agents take care of computational tasks* (e.g. enumerating and analyzing the large, possibly combinatorial, space of configurations of experts for the task, or monitoring the distributed activities of an expert team to assess the progress) *while humans do what humans are good at* (both engaging in task-level problem solving, e.g. combining their expertise to decide on the best course of action, and collaborating with the agents to manage teams better by applying their knowledge of strengths and backgrounds of experts to help identify effective teams).

² ISS is a calculation of available Abbreviated Injury Scale (AIS) scores estimating six body regions to determine the whole body injury severity.

4 Components of the HAC Framework

We have developed a HAC framework and will describe its constituent technologies as follows.

4.1 Representation of Agents' Roles and Capabilities

Our approach adopts a service-oriented architecture (SOA) for pooling expertise. In a SOA, the capabilities of agents (whether human or artificial) are described in an ontology that supports the matching of task requirements (e.g., the need to decide about how to prevent infection) with the capabilities of an agent who can accomplish the task (e.g., an infectious disease expert). It also supports the assessment of gradations of capabilities (e.g., if an infectious disease expert is not available, which of the available specialists (or generalists) is best suited to fill in for this role). This allows a task that requires a service to find the best possible matching service provider available given that some services (experts) might not be accessible 24/7, and allows service providers to dynamically expand their roles as the task needs change (e.g., an internist called upon to perform minor surgeries if all surgeons are tasked). In the SOA, each potential service provider and consumer has an associated computational agent representative in the network. The representative agent is responsible for describing the capabilities of the service provider, monitoring the current availability of the provider, and engaging in behind-the-scene negotiations with other agents to make (tentative) commitments about the service provider's filling a role in the joint pursuit of some task.

4.2 Representation of Task and Environments Using Extended Hierarchical Task Networks (EHTNs)

We use an extended version of Hierarchical Task Networks (HTNs) to represent knowledge about the partially-ordered sequences of decision-making tasks, and the capabilities expected of humans and agents performing those tasks. Traditional HTNs have been effective in practical planning systems, but are not expressive enough to represent worth-oriented goals, contingencies, or the uncertainties that arise when an agent's plans are in fact distributed over multiple agents. Chen and Decker (2006 [4]) have developed the extended hierarchical task network (EHTN) representation to overcome these limitations.

EHTNs are based on previous approaches, e.g. the plan representation for the RETSINA agent architecture (Williamson et al. 1996 [9]) and TÆMS (Decker 1996 [5]), but incorporates extensions that include new ways of modeling the interplay between tasks in achieving quality on super tasks (Wagner 2006 [8]). The main advantages of EHTN includes: (1) explicit representation of information flow and control flow of task structures; (2) a quantitative definition of a vector of measurable, utility-influencing characteristics, e.g. quality, cost, duration, and how these characteristics accumulate as actions are being executed – characteristic accumulation functions (CAFs); (3) explicit task relationships indicating how task progress affects primitive action characteristics elsewhere (even at remote agents) in task structures; (4) the introduction of human involvement in the decision making process at a meta-level; and (5) an EHTN-based task pre-planning tool, PPT, described in section 4.3. These features move the planning problem from a goal achievement view to a worth-oriented

view, and more importantly provide a basic unit over which all coordination reasoning, no matter how complex, is ultimately concerned with task relationships. A fundamental advantage is that tasks for both humans and computational agents use the identical representation language. This allows agents and people involved in the generation and instantiation of alternative strategies for making complex decisions to have shared knowledge and models to seamlessly interoperate. This is particularly relevant when a human needing consultation help and the agent tasked with the human need to collaboratively form a consultation plan. Details of EHTN and sample task representations can be found in (Chen & Decker 2006 [4]) and especially (Chen 2006b [3]) with a formal explanation.

4.3 A Pre-planning Tool (PPT) for EHTNs

The tools for traditional HTNs have significant limitations, e.g. the inability to represent user-specified task characteristics for (planning) evaluation purposes, and the inability to specify how result values of certain enabling tasks differentially achieve the preconditions of other enabled tasks. We have developed a pre-planning tool (PPT) corresponding to the extended expressiveness of the EHTNs to overcome the above limitations. The PPT and associated (Java) APIs have the following advantages.

- Facilitating pre-planning, contingency planning and re-planning efforts for HAC system users and enabling the management of networked human expertise for H2AC.
- Providing effective interfaces: GUIs for HAC system users and Java APIs for system/application developers to facilitate automated HAC and H2AC.
- Enabling a domain-independent approach.
- Enriching a powerful task/environment specification to model heterogeneous interests (e.g. represented by utility functions) for both humans and agents, e.g. task structures, tasks relationships, task features, etc.

4.4 General Human Agent Collaboration Process

The general iterated HAC process is shown in this table.

Step 1	Initially the HAC system suggests widely diverse choices (both task plans and teams) to the user
Step 2	The user selects some suggestions of interest
Step 3	The user specifies further preferences/constraints
Step 4	The HAC system updates the suggestions accordingly and presents them to the user in the next round as variations on those narrower themes
Step 5	Go to Step 2 or exit upon convergence

Following the medical scenario introduced in section 3, let us focus on the process of the generation of a treatment plan and the formation of a mixed human agent team.

A suitable task plan is determined by the nature and severity of a case. E.g., if the wounded warfighter's leg has been injured badly, a worst case scenario, losing a leg, has to be considered with these procedures: screening, diagnosis, preparation, surgical operation, etc., which further require more medical experts: nurse, surgeon, physician, specialist (neurologist and rehab specialist), etc. Instead a minor scratch in the leg just needs brief treatment from a nurse. Given that different users of the HAC system

may possess different levels of experience and expertise (from a peer warfighter with little medical knowledge to an experienced surgeon), the HAC system needs to present suggestions suitable for all users: a non-professional could simply use the suggested task plan and expert team, while an advanced user might modify the task plans and specify preferred (or not) team members. For example, a surgeon, as a HAC system user, could accurately diagnose the case by himself, sidestep the nurse's screening and jump to the surgical operation directly. Similarly the user can specify constraints based on his preferences, e.g. including or excluding a particular medical expert, specifying a particular location the key treatment should happen, etc.

The HAC system provides the user with projected outcomes for alternative response plans by calculating the task features of interest (e.g. maximizing plan quality, or meeting earliest deadline). If the user is satisfied with a projected outcome, he simply selects it for execution; the selected response plan in turn determines a set of experts to carry out the plan. If the estimated plan outcomes are not satisfactory, the user can provide new preferences or even use the PPT to edit the response task plan, and requests the HAC to calculate the new plan outcomes; this process iterates until convergence to a satisfactory estimate of plan outcome and a suitable corresponding human agent team.

The key idea is that, while the human might have an ideal consultation plan and a "dream team" of experts to work with, the practical limitations on expert availability and interaction scheduling might make this impossible. So the human and agents need to work together with mixed initiatives, where the underlying agent network develops candidate consultation teams and schedules given the evolving activities of humans on the network (thanks to the agents' computational advantages), while the human seeking consultation expresses preferences over the possibilities and tweak the options to reflect knowledge, opinions and tradeoffs not captured in automated systems.

4.5 Other Technical Components of HAC Framework Dynamic Team Formation

As described earlier, each medical case may be different and require different combinations of human experts and software agents. At the same time, resource availability varies as well: no surgeon would be available for a new case if all of them are busy with operations. We have a flexible mapping between task requirements and human agent capabilities to facilitate suitable team formation that decides which human/agent to include in a team. Dynamic team formation considers factors such as: the nature of a given case, the availability of human experts and software agents, the mapping that ensures the task requirements satisfied with human/agents' capabilities, etc.

Negotiated Processes Facilitating Human and Agents Cooperative / Collaborative Problem Solving

It is far more difficult to consider human participation: a human may reject an assigned task (particularly assigned by a piece of software) simply because he is capable of doing it; human experts have their own schedules and preferences, or just might not "feel like working today",

which unpredictably affect task allocation and team formation. It is impossible to predict a person's mind. We introduced negotiation into the above processes: a human expert can negotiate about suitable task decomposition, team formation, and task allocation based on his biases, expertise and experiences through an effective yet constrained human agent communication interface.

HAC is not a separate operation and should be integrated within the entire problem solving process: team formation, task allocation, re-grouping and re-planning (caused by the arrivals of high priority cases, exceptions, or changes to resource availability, e.g. a nurse gets sick or a hardware/software failure), group termination, etc. This project tries to develop an integrated approach for effective HAC that facilitates H2AC by employing all the aforementioned component technologies. Other HAC components, such as exception/dynamics handling and communication protocols, etc., are omitted in this paper.

5 Conclusion and Future Work

This paper studies how HAC assists H2AC at the coordination meta-level. This DARPA project has proved the proposed technologies effectively support the notion of engaging human's biases, expertise and experiences and software agent's computational capabilities to achieve better decision-making. An updated project description can be found here: <http://www.i-a-i.com/view.asp?aid=279>.

References

- [1] Banks, B. and Lizza, S. *"The Pilot's Associate: A Cooperative, Knowledge-Based System Application"*, IEEE Intelligent Systems, 6(3):18-29, June 1991.
- [2] Berry, P., K. Conley, M. Gervasio, B. Peintner, T. Uribe, and N. Yorke-Smith (2006a). *Deploying a Personalized Time Management Agent*. In Proc. of the 5th AAMAS'06, Industrial Track, Hakodate, Japan, May 2006.
- [3] Chen W. *Designing an extended set of coordination mechanisms for multi-agent systems*. PhD thesis, Computer & Information Sciences, Univ. of Delaware, Jan. 2006.
- [4] Chen W. and Decker K., *Analyzing characteristics of task structures to develop GPGP coordination mechanisms*. The 5th Intl. Joint Conf. on AAMAS06, Pages. 662-669, Hakodate, Japan, May, 2006.
- [5] K. Decker. *TÆMS: A framework for environment centered analysis and design of coordination mechanisms*. In G. O'Hare and N. Jennings, editors, Foundations of DAI, chapter 16. Wiley Inter-Science, 1996.
- [6] Durfee, E., et. al. (1997). *"TAIPE: Tactical Assistants for Interaction Planning and Execution."* In *Proc. of the 1st Intl. Conf. on Autonomous Agents*, P. 443-450, Feb. 1997.
- [7] Wagner, T. (2005). DARPA COORDINATORS. <http://www.darpa.mil/ipto/programs/coordinators/>
- [8] T. Wagner, A. Raja, & V. Lesser, *Modeling Uncertainty and its Implications to Sophisticated Control in TAEMS Agents*, in JAAMAS, Apr 2006.
- [9] M. Williamson, K. Decker, and K. Sycara. *Unified information and control flow in hierarchical task networks*. In Proceedings of the AAI-96 workshop on Theories of Planning, Action, and Control, 1996.