Modeling Individual Decision for Collaboration and Implications for Intelligent Assistance

Bowen Hui

Department of Computer Science, University of Toronto, Canada
bowen@cs.utoronto.ca

Abstract. As collaboration has proven to be beneficial in learning environments, there has been an emerging interest in automating the formation of student groups. However, existing work focuses on optimal formations from the instructor's perspective based on pedagogical criteria (e.g., maximal group productivity, avoiding *orphans* or unmatched students). In contrast, we propose a formal collaboration model from the student's perspective using formation criteria that are important to individual students. As a consequence, resulting formations may contain groups of isolated students. We believe this result opens an opportunity to utilize the model in helping designers develop better adaptive systems for single users. To this end, we derive a formal model that explains why some users choose to collaborate while others choose to work independently. We implement this model and demonstrate in simulation the various factors involved in people's choices in the context of collaborative story writing. Furthermore, we use this model as the basis of designing an adaptive story writing assistant for individual users and discuss the design implications for intelligent assistance in general.

Key words: Group formation, adaptive systems, influence diagram

1 Introduction

Collaborative systems provide interaction environments that enable multiple users to work toward a common goal. As different users have different styles of interaction and varying needs, these collaborative environments provide a rich domain to study the interaction preferences of individual users in a multiuser setting. For example, some users prefer to maintain a private space that only allows the owner of that space to see and edit information in it, while maintaining a separate, common space for shared access among all the users. On the other hand, some users simply prefer to have a single shared space. The ability to design computing environments that accommodate the needs of different users in the same application setting is crucial. To this end, most of the work in the area of computer-supported collaborative systems has focused on creating static (i.e., non-adaptive) environments to facilitate the interaction needs and preferences of different kinds of user in a multiuser collaborative system [14]. However, the development of these systems tend to be based on investigations that *assume* collaboration among the study participants. In the general case, those same participants may not necessarily choose to form groups with each other to carry out the designated task. Likewise, these systems, when studied in real-use case scenarios, may result in users choosing to work alone rather than in collaboration with others. For these reasons, we propose to investigate the way in which users make collaboration decisions. We believe that understanding people's decision making process for collaboration provides insight to designing adaptive, intelligent assistance in general.

One reason why modeling collaboration choices is important is that a system equipped with the ability to reason whether collaboration takes place is able to mimic real usage more closely. In general, multiuser systems do not decide on the groups that should be formed and do not force users to collaborate. However, when multiusers *choose* to collaborate, there is a reason indicating that the joint collaboration is (expected to be) more beneficial than working independently. Understanding this information provides insightful design criteria for developing intelligent assistance when users have to work alone.

More specifically in the domain of collaborative learning, we are interested in modeling different aspects of student behaviour in independent learning. A specific interest is the ability for a student to assess her own skills and to estimate the skills required in accomplishing domain goals. With these two sets of skills at hand, we believe that students "match" their own skills to what is required from the goal in order to estimate the *expected performance* of working alone and arrive some level of desire to collaborate with others. We believe the result of this matching process guides a student in estimating the utility of group collaboration in learning environments. In particular, every group formation affords a certain set of *benefits* that results in the collaboration, which would have otherwise been unrealized if the student were to work alone. Examples of these benefits include time savings, higher quality of achievement, and social enjoyment. At the same time, forming a group also has potential *costs* which may not have been present otherwise. For example, two people who simply do not get along may result in lower productivity and lower satisfaction. Thus, we believe that a student's ability to quantify and anticipate these benefits and costs is what drives the reasoning behind the collaboration choices.

We use this general reasoning process to structure our decision model and formalize it as an *influence diagram* [12]. An influence digram is a graphical model that captures an agent's decision making process using variables and causal dependencies. In addition to variables that model factors in the environment, an influence diagram also models a decision maker's possible range of actions and the (numeric) utility of arriving at each outcome. In other words, each decision maker has a distinct *utility function* that quantifies the goodness of every possible outcome so that the decision maker can choose the best action leading to the outcome with the highest utility. A collaborative system equipped with such a model is able to model each student's individual preferences in learning, such as finishing the goal fastest or enjoying the processing of working with others. We demonstrate how such a decision model is constructed in Section 3.

Unlike existing approaches, we frame the group formation process from the student's perspective rather than from that of the instructor. For example, the objectives of an instructor would include avoiding *orphaned* students, i.e., students working by themselves, while this criteria would not be an objective of the students. As such, we cannot guarantee that all students will be placed in the groups they most prefer, and may result in working alone. We believe these circumstances should drive the design of adaptive systems that accommodate to individual needs. In particular, we illustrate that the same collaboration model we develop can also play a key role in designing an intelligent, assistive agent that helps the user accomplishes the domain goals, in the absence of collaborators. We illustrate the flexibility of such an intelligent system in the domain of collaborating story writing in Section 4 and discuss the general design implications in Section 5. Lastly, we discuss related work and summarize our work.

2 Related Work

Since collaboration has long been demonstrated to be an effective approach to student learning, researchers have been interested in "optimal" group formation and computer-supported automatic techniques for group formation. Existing formal models of student group formation are typically represented from an instructor's point of view and the effectiveness of the automatic formation technique is evaluated in terms of productivity of the group and perceived satisfaction of the students in the group [14, 10, 11, 9]. For example, when these models compute the best way to group students together (say, via diversifying student learning styles), the objective is designed to maximize the outcome performance based on what the formed group is expected to achieve and how satisfied the group is. To our knowledge, none of the existing models are formulated from the *student's* perspective. In contrast, we propose to create a collaboration model based on the student's individual skills and preference, and model the student's decision to collaborate directly. In this way, the resulting formed groups are a mere artifact of the students' behaviours.

From the area of adaptive systems for single users, our work is most similar to methods that use decision theory as a framework to automatically estimate the user type and reason how best to adapt the interface for that type of user (e.g., [5, 4, 1, 6]). The design of these systems are generally handcrafted and/or learned from data, whereas the design in our approach is guided by a collaboration model in this work. Another distinction is that we focus on the development of a collaborative group formation model and its design implications, rather than the development of an adaptive system. Given the current set-up though, existing approaches to adaptive systems can be applied in a straightforward fashion.

3 The Choice of Collaboration

In this section, we derive a formal model of how people make decisions to collaborate with others using influence diagrams. We refer to this as the *collaboration* *model*. Our derivation adopts the assumption that people carry out some sort of matching algorithm in order to assess one's desire to collaborate. Our objective is to use this model to guide the design of an intelligent system for single users in the case when they have to work alone, which we discuss in Section 5.

Let us consider the scenario where a user has a specific goal, *Goal*, that is decomposed into a series of three tasks, $Task_1$, $Task_2$, $Task_3$, each requiring a particular set of skills needed for task completion. This decomposition is illustrated in a graphical representation in Figure 1, where nodes represent a variable and arrows represent causal influences between the nodes. The difficulty of each task influences whether a task can be completed and to what extent is that quality, *Completion Quality*. This quality yields a particular reward to the user (e.g., monetary or self-gratification) denoted as $Payoff$. Furthermore, a high completion quality from a previous task may improve the overall quality future tasks, while a low completion quality from a previous task may degrade the overall quality of future tasks. Therefore, *Completion Quality* from previous tasks can influence any future *Completion Quality*.

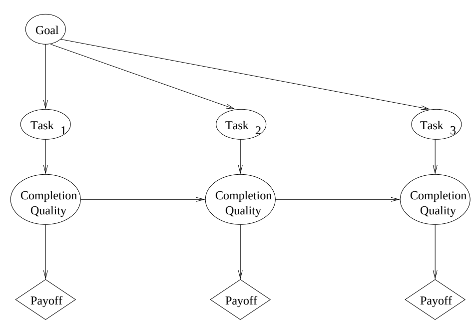

Fig. 1. Goal-task decomposition and the causal implications on personal payoff.

In addition, *Completion Quality* also depends on the user's own ability to achieve the task at hand. We introduce the variable *Core Skills* to represent the user's profile of the set of skills required to achieve the goal. Together, the user's skills and the difficulty of the tasks determine the quality of the task completion. The revised model is shown in Figure 2(a). Depending on the duration of tasks and the nature of skills, *Core Skills* may change (e.g., improve) over time. Consider a user with the goal of sending an email to each of her friends in her graduating class for a reunion. Two necessary skills for these tasks are composition skills and typing skills. The user's composition ability would likely remain static during the course of achieving this goal, while her ability to type may improve with repeated practice from typing each individual email. For this reason, we illustrate the model with dynamic skills in Figure 2(b). In this figure, the user beings with an initial set of skills, $Core\ Skills_0$, which changes to $Core\ Skills_1$ as $Task_1$ is completed, and so forth. Note that the latter model that incorporates dynamics skills is a general model that subsumes the one presented in Figure 2(a), simply by creating identical copies of the skills over time and defining those causal influences as identify functions. Therefore, we continue our discussion using Figure 2(b) only.

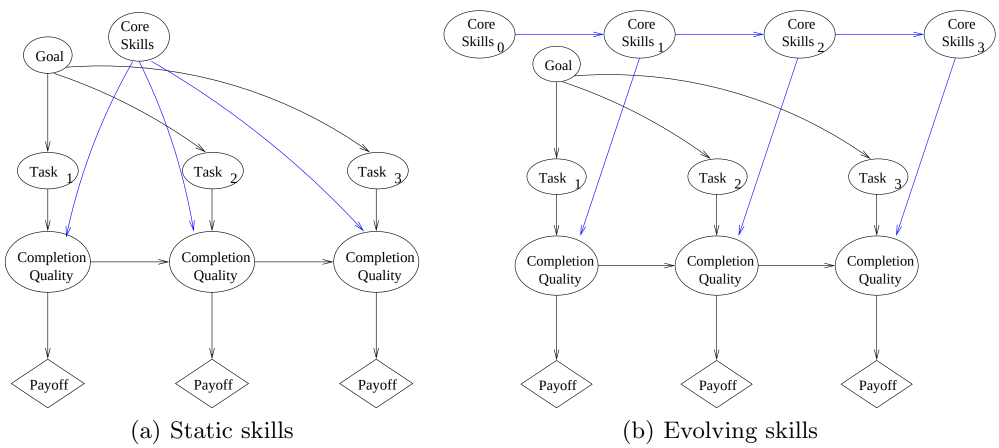

Fig. 2. The influence of the user's core skills on completion quality.

We view the user's choice to collaborate as a means to add more skills to the project in order to improve the completion quality of each task. We introduce the variable, A, as a set of actions that enables the user to choose to collaborate with others who have certain skill sets. Thus, A influences the *Core Skills* available to the overall project, which in turn influences the *Completion Quality* and the immediate $Payoff$ of each task. At the same time, collaborating with others require additional communication overhead. We represent this relationship by connecting A to *Communication*. The level of communication required also influences the quality of completion, since communication among group members may improve or degrade the completion quality depending on the specific combination of group members. Lastly, this extra *Communication* also plays a role in the user's personal payoff function, because the user may enjoy communicating with others as well as being able to complete tasks at a high level of quality. Note that although we have decomposed *Goal* into three tasks, the model in Figure 3 can be analogously derived for any number of tasks.

Note that Figure 3 illustrates a *sequential* decision model that allows for multiple decisions in the problem, each of which may have consequences for future actions. We have divided the overall decision problem into a sequence of smaller decision problems which are defined around tasks. In this way, the collaboration decision made for $Task_i$ has specific implications on the completion quality and payoff, as well as the completion quality and payoff for $Task_{i+1}$.

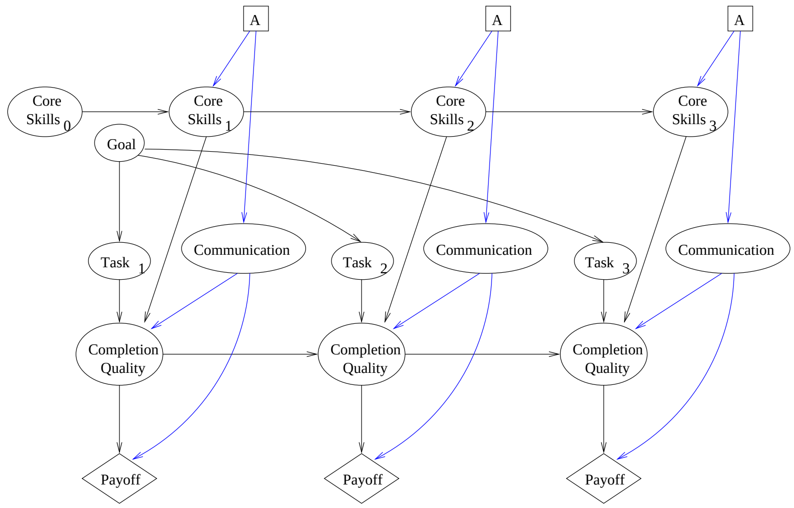

Fig. 3. A decision model of collaboration choice (assuming full observability).

3.1 Variables and Causal Dependencies

To explain the collaboration model in more details, we elaborate on the definitions of the variables and dependencies involved. We will illustrate with a small example in Section 4.

To begin, we assume there is a finite set of goals and a finite set of skills of interest. The *Goal* variable specifies which goal the user is interested in tackling. The specific goal selects the sequence of tasks involved. Each $Task_i$ is defined in terms of which skills are needed to achieve it. Each user has a profile of skills that describes how good the user is at each of the skills required in order to achieve the goal. These skills are expressed in terms of the level of expertise, e.g., low, medium, high. The variable $Core\ Skills_i$ represents the overall profile of skills that are available based on the group members involved. For example, if the user is working alone, $Core\ Skills_i$ is simply that user's skill profile. On the other hand, if there are two users involved, $Core\ Skills_i$ is the combination of the skill profiles of the two users. We define this joint profile so that it takes the highest level of expertise from all the members. In other words, if user A has a high level of expertise for skill s_1 and user B has a low level of expertise for s_1, then the joint profile for users A and B for that skill is high.

Similarly, the variable $Communication$ represents the overall communication needs of all the group members involved. In particular, each user has an individual communications profile that expresses the level of social activity, e.g., low (1) to high (5). We assume the level of communication in a group is a function the number of group members and their individual needs. Therefore, we define the joint communication profile as the sum of the individual profiles.

The quality of a finished task, *Completion Quality*, is defined as a function of $Task$, *Core Skills*, *Communication*, and the previous *Completion Quality*. Specifically, we are interested in how well the available skills from the group members, *Core Skills*, matches those needed by the current $Task$. In other words, this is the potential quality achievable by the group members involved. However, these members also have communication needs. Thus, we subtract *Communication* from the potential quality. Furthermore, we adjust the potential quality based on the *Completion Quality* achieved by the previous task because the quality of the previous task may limit what can be achieved in the current task. This gives an estimate of the quality of completing the current task.

As we mentioned earlier, *Core Skills* and *Communication* are inferred based on observations such as how well other members finish their tasks and how talkative other members are respectively. Since these observations are not definitive indications of skills or communication needs, we define a probabilistic observation model that describes the strength of indication for each observation toward each variable. Such an observation model allows us to incorporate uncertainty and noise naturally.

3.2 Individual Utility Functions

Using an influence diagram formalism to model a user's choice for collaboration, we are also able to model individual preferences in the utility nodes in the diagram. Specifically, the user's utility function is defined as follows: $Payoff = u(Completion\ Quality, Communication)$. Unlike previous group formation strategies that focus only on task performance, this definition also accommodates one's preference toward social interaction. By assuming additivity, we decompose this function as $w_1 u_1(Completion\ Quality) + w_2 u_2(Communication)$, where w_1 and w_2 are weights on the two subutility functions summing up to 1.0. Using these weights, we can model individual preferences as follows. For example, a user who prefers to achieve high quality work and spend little time communicating with group members would place a high value on w_1 and a low value on w_2. On the other hand, a user who enjoys communicating with others and also value high quality work can assign equal weights to both subutility functions. In general, any relevant factor that plays a role in capturing user preferences can be defined as an input variable to the utility function in a similar manner.

4 Testbed: Collaborative Story Writing

To illustrate its applicability, we use the collaboration model from Figure 3 to demonstrate how users choose collaborators with the general goal of writing a story. First, we use a simple example where we have three users, Alice, Bob, Colin, each with different skill profiles. In the case where skill profiles are not readily available, we may consider the use of skill assessment measures that elicit the users' abilities automatically and create the necessary profile for each user. Now, let us consider the goal of story composition, which generally requires these skills: keyboard typing (s_1), vocabulary (s_2), and creative composition (s_3). With respect to these three skills, let us suppose that each user has the skill profiles shown in Table 1 at levels ranging from low, medium, and high. Suppose the story composition goal consists of two tasks. The first is a brainstorming session with the goal of generating as many interesting scenarios, characters, and objects as possible. The second task is the actual writing of a story, based on the ideas generated from the first task. In the first task, the most important skill is vocabulary, while in the second task, the most important skills are typing and composition. We also define a level of communication needs for each of these users, indicating the social interaction needs of the user — i.e., how likely each user is likely to talk to others (or otherwise be distracted from the task). Each user's communication level is shown in the rightmost column in Table 1, with levels ranging from low (1) to high (5).

Table 1. Skill and communication profiles for each user for story composition.

Users	**Typing (s_1)**	**Vocabulary (s_2)**	**Composition (s_3)**	**Communication**
Alice	high	medium	medium	4
Bob	low	low	high	1
Colin	high	high	low	2

Using this scenario, we will consider how well each user may achieve each task independently, and what they gain or lose by collaborating with others. We define the set of possible actions for each user as a choice of collaboration. First, we consider Alice's decision making process. Alice's possible actions are: a_1 as the choice of working alone, a_2 as the choice of working with Bob, a_3 as the choice of working with Colin, and a_4 as the choice of working with both Bob and Colin. Analytically, we see that Alice can achieve the first task (brainstorming) fairly well, since her vocabulary skills are decent. However, Alice enjoys working with others a lot and places a very high weight on communication. To represent this preference, we define the weight vector $[wQ_a, wC_a]$, where $wQ_a = .1$ corresponds to the tradeoff value Alice has toward *Completion Quality*, and $wC_a = .9$ corresponds to the tradeoff value Alice has toward *Communication*. In brief, we show the resulting utility computations for Alice and the four possible collaboration choices in Table 2 (*Quality* is short for *Completion Quality*.) We see that working alone yields the worst utility for Alice, while working with both Bob and Colin yields the best utility. As well, working with Bob does not improve the task completion quality, but still results in higher utility because this option involves more communication.

On the other hand, consider Bob's choices for the brainstorming task. Let us define his preferences as being "opposite" from Alice's, with the weight vector $[wQ_b, wC_b]$, where $wQ_b = .9$ and $wC_b = .1$. The utility of each of his actions are shown in Table 3. Here, we see that working alone yields very low quality for Bob, as his vocabulary skills are low. However, because he prefers little communication and working alone requires little communication, this option is still better than

Table 2. The utility of Alice's collaboration choices for Task 1.

Action	Quality	Communication	Utility
Work Solo	4	2	3.6
Work with Bob	4	3	4.4
Work with Colin	9	6	5.7
Work with Bob and Colin	9	7	6.5

working with Alice. Finally, the best option for Bob is to work with Colin, which is very good at vocabulary and does not require much communication.

Table 3. The utility of Bob's collaboration choices for Task 1.

Action	Quality	Communication	Utility
Work Solo	1	1	0.1
Work with Alice	4	5	-0.4
Work with Colin	9	3	5.7
Work with Alice and Colin	9	7	2.5

Finally, we show Colin's choices and corresponding utilities in Table 4. We define Colin as someone who prefers to generate high quality results and does not mind a little communication, using the weight vector $[wQ_c, wC_c]$, where $wQ_c = .7$ and $wC_c = .3$. Note that since Colin is the best at vocabulary skills, the quality of any of his collaboration choices is the same. However, the communication costs change depending on whom he chooses to work with. Among the various options, we see that the best option for Colin is to work alone in this task.

Table 4. The utility of Colin's collaboration choices for Task 1.

Action	Quality	Communication	Utility
Work Solo	9	2	5.5
Work with Alice	9	6	3.9
Work with Bob	9	3	5.1
Work with Alice and Bob	9	7	3.5

Next, we carry out the same calculations for the second composition task. Recall that this task requires typing and creativity. The best options for each of the users are shown in Table 5. The best actions for Alice and Bob are still the same; for Alice, it is in her best interest to work with everyone, while for Bob, it is best for him to work with Colin. On the other hand, Colin's best action is to collaborate with Bob because his skills alone are not adequate to achieve a high completion quality for this task.

This small example illustrates the use of a formal, quantitative model to model individual user's decision making process for collaboration choices. In particular, this example demonstrates several key ingredients: (i) the representation of individual skills and communication profile, (ii) individual tradeoffs

Table 5. The best collaboration choices for each user for Task 2.

User	Best Action	Quality	Communication	Utility
Alice	Work with Alice and Bob	18	7	7.4
Bob	Work with Colin	18	3	13.8
Colin	Work with Bob	18	3	11.4

between *Completion Quality* of tasks and the cost of *Communication* represented as weights in the utility function, and (iii) the best action as determined by the option with maximum utility.

An additional aspect of the collaboration model in Figure 3 that is not illustrated in this simple example is the sequential impact of decisions across tasks. In particular, note that the model allows users to create new collaborations for every task. For example, we saw that Colin's best course of actions is to work alone for $Task_1$ and to work with Bob for $Task_2$. However, the computations in the example did not consider the impact of the completion quality that $Task_1$ has on the quality of future tasks. Therefore, in Colin's case, it is worthwhile sacrificing some utility in $Task_1$ and work with Bob, so that the potential quality of the next task is maximal.

5 Implications for Intelligent Assistance

Developing a collaboration model that explains the costs and benefits of the group enables us to design a user-specific system for students who prefer to work alone. In particular, our model highlights exactly how adaptation could help users achieve better quality work and improve communication preferences. Developing models to explain individual preferences for collaboration also provides a starting point in understanding the dynamics of group preferences. In this section, we focus on the steps needed in applying the collaboration model in an intelligent system that helps individual users with specific goals.

5.1 Adaptive Functionality to Assist Users

Often there are situations when users end up working alone, either because they prefer to work independently or they were unable to form a group with other members. In these scenarios, it is crucial to be able to develop a system that help individual users achieve their goals in the best way possible. The role of adaptive (and assistive) systems becomes especially important when users do not have the sufficient skills in achieving the goals they undertake (e.g., in the domain of assistive technologies [3, 13, 8]). These scenarios are more likely to occur with users who are novices in the domain or users who have cognitive, sensory, or motor impairment. To illustrate how the collaboration model aids in designing such systems for individuals, we continue with the testbed application of story writing as described in Section 4 in the context of single users. Intelligent systems can be designed in various ways. The purpose of this discussion is to demonstrate a systematic approach to designing an intelligent system based on what helps the user achieves the target goal while taking into account the user's preferences.

Recall that an input to the collaboration model is a specification of required skills for each task and the user's skill profile. Thus, by implementing a simple matching algorithm in the system, it is easy to determine whether the user has the necessary skills to achieve the target tasks. For example, a system that supports story writing can be designed with functionality that helps each of the required skills in the goal. Recall that the story writing goal requires keyboard typing (s_1), vocabulary (s_2), and composition (s_3). To assist a user in typing, a useful function is word prediction (e.g., as in the adaptive word prediction program in [6]). To help the user brainstorm interesting and diverse ideas, the system can provide a thesaurus to help broaden the user's vocabulary, word association lists based on what the user has brainstormed already, or pictures and scenarios consisting of objects related to the existing ideas. Finally, to support the user in the composition task, the system can provide high-level templates to help users create the plot or low-level templates to help structure individual paragraphs and sentences (e.g., as in the personalized email program in [8]).

Another input to the collaboration model is the user's communication needs. In other words, an intelligent system should also provide functionality to improve the user's social environment. Such functionality includes encoding private settings to ensure the user does not get distracted by other users (e.g., from an online chat program) or by system interruptions (e.g., adaptive notifications), playing music or changing background styles and colours, and interacting with friends remotely. Note that communicating with other users through the system is different from communicating with others in person. Thus, in a single-user system, the user's communication profile will need to include preferences toward interaction via system means (e.g., using an interaction cost model in adaptive systems [2, 7]).

After considering these design implications, an adaptive system has to ability to decide whether to provide the functions to the user or not. In other words, rather than always providing help to the user or never providing help, the system uses its collaboration model to reason the costs and benefits of its help and only offer it to the user if benefits outweigh the associated costs. More specifically, rather than modeling user actions in Figure 3, we use the variable A to model system actions. Here, system actions would be to add or delete any one of the aforementioned functions, or to do nothing. Again, adding or removing a function has influence on the user's *Core Skills* and *Communication* needs (depending on the function), which in turn influences the task's *Completion Quality*. Since the system's objective is to help the user, we can model the system's $Payoff$ function directly as the user's utility function so that maximizing the system's payoff will also increase user satisfaction. In this way, we augment the collaboration model with system actions and use the causal dependencies as the system's *reasoning* engine in choosing which function to provide (or hide) from the user.

In other words, the collaboration model enables the system to *predict* the *expected utility* of each system action based on the task description and the user's core skills and communication needs. Since the utility function incorporates the user's preferences (i.e., the weight vector that expresses the tradeoffs between quality and communication), the system can carry out assistive actions to help high achievers to produce better quality work, and/or carry out interactive or entertaining actions to help the more social users enjoy the work process.

5.2 Domain-Specific and Domain-Independent Components

Most of the knowledge required to build the collaboration model is domain-independent because it focuses on the users. For example, the definitions of $Core\ Skills$ and $Communication$ needs are independent of the domain since they describe individual users irrespective of the tasks or context. The relationship $A \rightarrow\ Core\ Skills_i$ describes how an individual user's choice for collaboration affects the overall skills available for $Task_i$. This information is general to any domain, because we have modeled tasks as a set of skills. Similarly, the relationship $A \rightarrow\ Communication$ is also domain-independent, because it describes how a user's actions affect the overall social interaction during task completion. Lastly, $Completion\ Quality, Communication \rightarrow\ Payoff$ is a relationship that defines the user's preferences, and thus, is also domain-independent.

On the other hand, $Goal \rightarrow\ Task_1, .., Task_k$ is a domain-dependent relationship since it requires knowing how goals can be decomposed into smaller tasks and what skills are required to achieve each of those tasks. Moreover, this information is expert knowledge that is not necessarily known to all users. As well, $Task_i, Core\ Skills_i, Communication \rightarrow\ Completion\ Quality$ is a relationship also requiring expert, domain knowledge. Since not all users have this kind of information available, another functionality to include in an intelligent system is a way to facilitate such expert knowledge. For example, the system could have predefined goals, tasks, and required skills so that their relationship is already specified. Alternatively, instructors or expert users may wish to add new goals and tasks for future use. The knowledge pertaining to the latter relationship can also be acquired in a similar manner.

5.3 Developing an Accurate Model

The feasibility of the proposed collaboration model hinges on how accurately it depicts the actual process of people deciding whether to collaborate with others. In Section 3, we derived a collaboration model following generic steps that apply across domains. We have illustrated the efficacy of this model using an example of collaborative story writing. However, when tested in more general settings, the model may require further development. In such circumstances, the same steps used to derive this model can be taken to extend the model structure.

The second aspect in model accuracy is the set of parameters used in defining the causal relationships between the variables. Here, quantitative data is needed to ensure that the model is accurate for the intended domain and users. First of all, for the domain-independent relationships as mentioned in Section 5.2, we can use elicitation methods to acquire the necessary data. For example, paper-based questionnaires as well as online question-answering can also be used to intelligently elicit the user's skill and communication profile. On the other hand, the domain-specific relationships mentioned in Section 5.2 require quantitative data situated in the target domain. For example, quantitative data pertaining to story writing skills and task completion quality can be used to estimate the correlation between the task and quality. In this case, the amount of data needed depends on the complexity of the model (i.e., the number of variables and the number of values per variable).

6 Discussion

In this paper, we have developed a formal model of collaboration choices using the influence diagram formalism. For simplicity, there are two aspects that we have not incorporated into the model, which we will discuss in this section.

Cost of Working Alone. When a user chooses to work alone, there is inevitably a need to model the cost of asking others for help when one's skills are not sufficient in achieving the target goal (or at least, to a desired level of quality). Although it was not demonstrated in our example, we can introduce another variable to model the query cost as part of the user's preferences in the decision making process. For example, how well the required skills of $Task_i$ and the available $Core\ Skills_i$ match could suggest that the user needs to ask an external member for help. However, some users like to work independently and prefer to not ask for help. Therefore, we need to incorporate this factor into the user's utility function and add another parameter to the weight vector. In this way, the user's utility function expresses the tradeoffs between completion quality, communication needs, and query cost. The design of an adaptive system in Section 5.1 is augmented similarly.

Partial Observability The collaboration model presented in Figure 3 assumes that the user knows the values of $Core\ Skills$ and $Communication$ with certainty. In a realistic setting, a user may not recognize her own skills or may not be able to entail the communication needs of others in new collaborations. In other words, the variables $Core\ Skills$ and $Communication$ are *unobservable* and need to be inferred. For example, a user who has previously achieved tasks of a certain difficulty level may infer that she is good at the skills required in those tasks. On the other hand, if the user has failed tasks requiring certain skills, then she may infer that she does poorly at those skills. Similarly, if the user has observed the behaviour of a talkative member, she may infer that collaboration with this member would require a lot of potentially unnecessary communication, and thus, lowering the overall task completion quality and yielding a lower payoff. Therefore, we introduce $Observations$ into the collaboration model in Figure 3. As a result, we have a *partially observable* model where the values of some of the variables are observed and the others are inferred. This modified model is illustrated in Figure 4, and this version is a generalization of the model we have used throughout the paper.

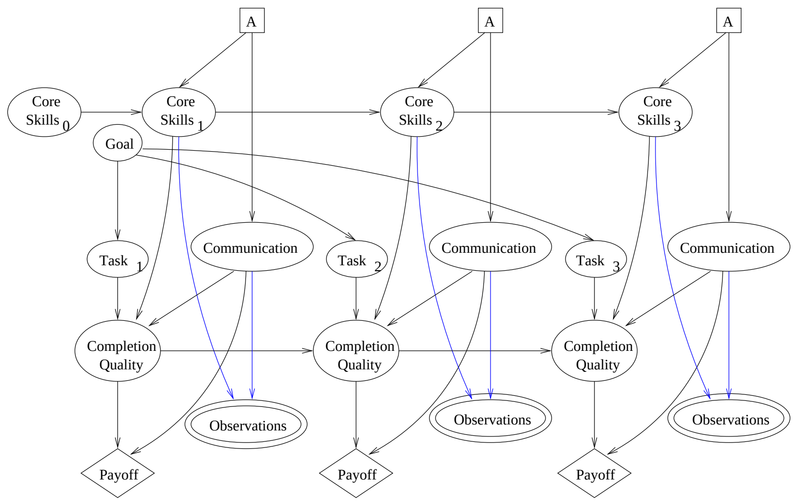

Fig. 4. A partially observable model for making collaboration decisions.

Accounting for partially observability enables us to describe a more realistic picture of the user's decision making process in choosing potential collaborators for a specific goal. Since the values of several variables are inferred, the computation of $Payoff$ is also taken in *expectation* of the probability distribution of those variables. In this way, the model provides an expected payoff (or the expected performance) of the specified group.

7 Conclusions

The focus on this paper is to develop a formal, quantitative model of one's decision to collaborate with others from the student's own perspective. Our model made several assumptions about the general process that students use to make such decisions and the input available for implementing a formal model into an adaptive system. Although our collaboration model is an initial step built using these assumptions, the principles used to derive the model for specific domains still apply. In general, we believe that understanding people's decision to collaborate can help designers develop better adaptive systems for individual users when collaboration is not possible. As we discussed, the functionality of such adaptive systems can be customized to accommodate user needs, as well as designed with automatic adaptive behaviour that assist users during their interaction. The separation between domain-specific and domain-independent components also provide a way for designers to understand which part of the system requires expert knowledge and which part can be data-driven based on specific user interaction patterns. As more psychological evidence becomes available to provide a clearer picture of people's reasoning process in their decision to collaboration, the formal model can be modified accordingly, which in turn unfolds more design issues for developing adaptive systems.

References

1. T. Bohnenberger and A. Jameson. When policies are better than plans: decision-theoretic planning of recommendation sequences. In *Proc. of Intelligent User Interfaces (IUI)*, pages 21–24, 2001.
2. K. Gajos, M. Czerwinski, D. Tan, and D. Weld. Exploring the Design Space For Adaptive Graphical User Interfaces. In *Proc. of AVI*, pages 201–208, 2006.
3. T. Hasselbring and C. Glaser. Use of computer technology to help students with special needs. *The Future of Children: Children and Computer Technology*, 10(2):102–122, 2000.
4. E. Horvitz. Principles of mixed-initiative. In *Proc. of Human Factors in Computing Systems (CHI)*, pages 159–166, 1999.
5. E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumière Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. In *Proc. of Uncertainty in Artificial Intelligence (UAI)*, pages 256–265, 1998.
6. B. Hui and C. Boutilier. Who's Asking for Help? A Bayesian Approach to Intelligent Assistance. In *Proc. of Intelligent User Interfaces (IUI)*, pages 186–193, 2006.
7. B. Hui, S. Gustafson, P. Irani, and C. Boutilier. The Need for an Interaction Cost Model. In *Proc. of Advanced Visual Interfaces (AVI)*, pages 145–148, 2008.
8. B. Hui, S. Liaskos, and J. Mylopoulos. Requirements Analysis for Customizable Software: A Goals-Skills-Preferences Framework. In *Proc. of RE*, pages 117–126, 2003.
9. E. Martin and P. Paredes. Using learning styles for dynamic group formation in adaptive hypermedia systems. In *Proc. of Adaptive Hypermedia and Collaborative Web-Based Systems (AHCW)*, 2004.
10. A. Ounnas, , and D. M. H. Davis. A metrics framework for evaluating group formation. In *Group*, 2007.
11. K. Papanikolaou, E. Gouli, and M. Grigoriadou. Accomodating individual differences in group formation for collaborative concept mapping. In *Proc. of Concept Mapping*, 2006.
12. R. D. Shachter. Evaluating influence diagrams. *Operations Research*, 33(6):871–882, 1986.
13. F. Shein, T. Nantais, R. Nishiyama, C. Tam, and P. Marshall. Word cueing for persons with writing difficulties: WordQ. In *Proc. of the 16th Annual International Conference on Technology and Persons with Disabilities*, 2001.
14. L. Soh, N. Khandaker, X. Liu, and H. Jiang. A computer-supported cooperative learning system with multiagent intelligence. In *Proc. of Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2006.