Synergistic Capabilities in Believable Agents

W. Scott Reilly
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
scott.reilly@cs.cmu.edu

Abstract

The Oz project at CMU is developing technologies that will allow artists to create dramatic, interactive environments.[1] One important aspect of such environments is the existence of believable agents. We believe that the way to create such agents is to give them a broad range of integrated capabilities, any of which may be somewhat shallow. We also believe that one of the advantages of this approach is that these capabilities will combine in useful ways that will provide an additional source of believability for these agents. I will examine this type of synergy through a number of examples.

1. Introduction

The Oz project at CMU is developing technologies that will allow artists to create dramatic, interactive environments [1]. One important aspect of such environments is the existence of believable agents. We have previously put forward the hypothesis that the way to create such agents is to give them a broad range of integrated capabilities, each of which may need to be somewhat shallow [2]. In our current architecture, called Tok, these capabilities include: goal-directed behavior, reactivity to the environment, emotions, social skills, and perception [3]. We are also exploring the integration of natural language generation and understanding.

One of the advantages of such an approach is that the various capabilities seem to combine in a synergistic fashion. In other words, we hope the agents we build will be more believable than would be expected from just having a variety of capabilities. I will look at a number of these relationships between capabilities in Tok to better demonstrate the advantage of this approach.

2. Synergistic Capabilities

Many of Tok's capabilities interact with each other in a synergistic manner. For instance, an agent with natural language generation is going to be more believable than one without, but if the agent has a broad range of other capabilities, language generation adds even more believability. There are two reasons that this occurs. First, these other capabilities give the agent more to talk about: a memory system allows an agent to talk about past experiences, an emotion system allows an agent to talk about its emotional state, and a social system allows an agent to talk about its relationships and beliefs about other agents.

[1] This article first appeared in the Working Notes of the AAAI Spring Symposium on Believable Agents, Stanford, CA, March 1994.

Second, other capabilities may affect the way language is generated. For example, emotions may cause changes in what an agent chooses to say and how they say it. Also, an agent's language may change based on social constraints, like who the agent is talking to.

Another example of synergy arises between emotion and other capabilities. As just noted, emotions and language generation combine nicely. Also, an agent with a memory system can have emotions based on past experiences, and the current emotional state may influence what memories arise. Goal-based action and social skills also combine in synergistic relationships with emotion. As these are two relationships that exist in the current Tok architecture, I will use them to provide somewhat deeper examples of synergistic combinations of capabilities in believable agents.

2.1. Emotion and Action

The emotions in Tok agents are controlled by the Em system [7]. Em generates emotions based on ideas from the cognitive emotion theory of Ortony, Clore, and Collins [6]. Em has a set of emotion generation rules that fire when certain internal or external states or events occur. A number of Em's emotion generation rules are based on the processing of Tok's action architecture, called Hap [4, 5].

Hap is a goal-based reactive architecture that processes the goals, plans, and actions in Tok agents. The emotions that depend on Hap are described briefly in table 1. For example, if an agent has a goal that fails, that agent will feel distress. If it is merely likely that the goal will fail, the agent will feel fear. Without being integrated with a goal-based action system, Em would be greatly restricted in the emotions that could be produced.

Furthermore, emotional agents are more believable if they express their emotions through their goals, plans, and actions. At the goal level, Em can affect what goals an agent has and what priority they are given. For example, Sue may have a do-not-be-hurt goal that she normally doesn't think about. When she fears that this goal may fail, however, it suddenly takes on a great deal of importance and will likely give rise to new defense goals. Emotions may also affect what plan an agent chooses to use in pursuit of a goal. For instance, when angry, Tom may choose to yell and bang on a closed door to have it opened. When calm, he may ring the bell and wait patiently. At the action level, emotions may affect the style for performing a particular action. If Mary is sad, she will probably walk slower and with worse posture than if she were happy.

Without an action system that to influence, Em would either have to resort to other methods of expressing emotions or be unable to express the emotions that were being generated. Tok agents, however, because of the integration of emotion and action, not only produce emotions and actions, they produce also produce something else: emotional actions.

A similar relationship holds between emotion and a number of other capabilities; the broader the agent, the more emotionally expressive the agent can be. For instance, adding memory allows an agent to recall past emotional experiences. Also, adding a language system allows the agent to use expressive speech, both in what is said and how it is said. An emotion system could exist in a very simple agent, but the broader the agent's capabilities, the more emotional the agent will seem to be.

In summary, there are two ways in which combining emotion and action produce more believable agents. First, the emotions an agent "feels" will be much more believable because of the presence of the action system. Second, the agent will be able to express these emotions

Emotion	Cause of Emotion
Joy	Important goal success or increased likelihood of important goal success
Distress	Important goal failure or increased likelihood of important goal failure
Hope	Likelihood of important goal success
Fear	Likelihood of important goal failure
Satisfaction	Hoped for goal success occurs
Fears-Confirmed	Feared goal failure occurs
Relief	Feared goal failure does not occur
Disappointment	Hoped for goal success does not occur
Gratitude	Other agent is responsible for important goal success
Anger	Other agent is responsible for an important goal failure
Gratification	Self is responsible for an important goal success
Remorse	Self is responsible for an important goal failure
Frustration	A plan fails but the parent goal is still achievable

Table 1: Action related emotions generated by Em.

in a variety of ways that depend on goals, plans, and actions.

2.2. Emotion and Social Behavior

There are (at least) two ways that emotions and social skills combine to make Tok agents more believable. First, the combination allows for emotions that are greater in number and variety. Second, the relationships that agents have with each other are more believable because of this integration of capabilities. I will expand on these ideas in turn.

First, social factors are very important in determining emotions. Many causes of emotion in ourselves and in interesting agents from literature arise because of others. Examples of social emotions in Em are anger, which is caused by another agent causing a goal failure, and love, which is the result of being near (or thinking about) a agent that one has a particular attitude or relationship with. Table 2 gives a more complete list of emotions based on social factors. Note that in this table, I have included hope and fear, which are not necessarily social emotions, but that can be social emotions if another agent is the cause of a goal potentially succeeding or failing. Without social knowledge and relationships with other agents, Em would have a large gap in the types of emotions that it could generate.

Second, the relationships between agents are more believable because of the integration of emotion and social skills. This is because the dynamics of social relationships often depend on emotions. If Bill is regularly mean to John, it is likely that John will get angry at Bill and eventually learn to dislike Bill. If John couldn't feel anger, he would probably continue to associate with Bill, which seems less believable. Tok agents could enter into social relationships with other agents without having emotions, but by integrating these two capabilities these agents have much richer relationships that can change over time based on emotional factors.

In summary, by integrating emotions and social skills we get more believability than from the two independently. We add a considerable number of social emotions that wouldn't be possible otherwise, and we have more believable

Emotion	Cause of Emotion
Hope(soc)	Likelihood of other agent causing an important goal to succeed
Fear(soc)	Likelihood of other agent causing an important goal to fail
Gloating	Disliked agent has goal failure
Resentment	Disliked agent has goal success
Happy-For	Liked agent has goal success
Pity	Liked agent has goal failure
Admiration	Other's action is approved of according to standards
Reproach	Other's action is disapproved of according to standards
Gratitude	Other agent is responsible for important goal success
Anger	Other agent is responsible for an important goal failure
Love	An object (or agent) is appealing according to an attitude
Hate	An object (or agent) is unappealing according to an attitude

Table 2: Social emotions generated by Em.

relationships between agents. To build either capability on its own would lead to large holes in our agents, which we are desperately trying to avoid in order to achieve believability.

3. Summary

I have touched briefly on a number of ways that capabilities can combine in a synergistic manner to create new behaviors and more believable agents. There are many more such fortuitous combinations of capabilities that have not been mentioned here. These benefits are a direct result of creating an agent architecture that includes a broad range of tightly integrated capabilities. This is one reason why creating broad agents seems to be a promising path towards the final goal of creating believable agents.

One question that remains unanswered is which set of capabilities should be included in a broad, believable agent. If we assume that it isn't feasible to build completely broad agents, we must decide which set of capabilities to build and integrate. The answer probably depends on the task the agent is designed to perform, but knowing how various tasks influence this choice would be interesting. It would also be useful to identify capabilities that seem to be generally important across tasks. When choosing a set of behaviors for an agent architecture, it also seems reasonable to try to determine what synergistic interactions might exist between various capabilities in order to get the most from this choice.

4. Acknowledgments

This research was supported in part by Fujitsu Laboratories, Ltd. The ideas in this paper were developed in conjunction with Joe Bates and Bryan Loyall. I would also like to thank the rest of the Oz group for their assistance.

References

[1] Joseph Bates. Virtual reality, art, and entertainment. *PRESENCE: Teleoperators and Virtual Environments*, 1(1):133–138, 1992.

[2] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Broad agents. In *Proceedings*

of AAAI Spring Symposium on Integrated Intelligent Architectures, Stanford, CA, March 1991. Available in *SIGART Bulletin*, Volume 2, Number 4, August 1991, pp. 38-40.

[3] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Integrating reactivity, goals, and emotion in a broad agent. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, IN, July 1992.

[4] A. Bryan Loyall and Joseph Bates. Hap: A reactive, adaptive architecture for agents. Technical Report CMU-CS-91-147, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, June 1991.

[5] A. Bryan Loyall and Joseph Bates. Real-time control of animated broad agents. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Boulder, CO, June 1993.

[6] A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.

[7] W. Scott Reilly and Joseph Bates. Building emotional agents. Technical Report CMU-CS-92-143, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1992.