Eurographics/ ACM SIGGRAPH Symposium on Computer Animation (2007)
D. Metaxas and J. Popovic (Editors)

A Decision Network Framework for the Behavioral Animation of Virtual Humans

Qinxin Yu[1,3] and Demetri Terzopoulos[2,3]

[1]Artificialife Inc., Montreal, QC, Canada
[2]University of California, Los Angeles, CA, USA
[3]University of Toronto, Toronto, ON, Canada

Abstract

We introduce a framework for advanced behavioral animation in virtual humans, which addresses the challenging open problem of simulating social interactions between pedestrians in urban settings. Based on hierarchical decision networks, our novel framework combines probability, decision, and graph theories for complex behavior modeling and intelligent action selection subject to manifold internal and external factors in the presence of uncertain knowledge. It yields autonomous characters that can make nontrivial interpretations and arrive at rational decisions dependent on multiple considerations. We demonstrate our framework in behavioral animation scenarios involving interacting autonomous pedestrians, including an elaborate emergency response animation.

1. Introduction

Creating autonomous characters with humanlike behaviors is a serious challenge. Our goal is to develop advanced behavioral systems for virtual humans. In particular, we address the level of decision-making that enables the characters to interact appropriately with their perceived environment, especially with other virtual humans. We focus on action selection; i.e., on simulating how humans decide what to do at any given time. To this end, we introduce a *decision network framework* for specifying and activating human behaviors that is easy to define and modify, scalable, and ostensibly emulates how people make decisions.

Uncertainty and complexity are characteristics of human behavior that make it especially difficult to simulate. Uncertainty has largely been ignored in prior behavior models, particularly uncertainty resulting from the natural limitations of perception, especially perception of the intentions of other people. Furthermore, no systematic approach has been proposed to deal with complexity. Our decision network framework addresses both issues. Decision networks are a generalization of Bayesian networks [Pea88], also known as probabilistic graphical models, which combine probability theory and graph theory to capture uncertain knowledge in a natural and efficient manner. An attractive feature of the decision network is that it is a powerful tool for modeling decision making under uncertainty. It provides an elegant and rigorous mathematical formalism for modeling complicated relationships among random variables and an intuitive visualization of these relationships as a graphical structure, thus facilitating comprehension and debugging. Furthermore, the modularity of a decision network facilitates the intuitive reduction of a complex behavior into manageable components.

Our work should not be misconstrued as yet another effort on so-called "crowd simulation." Our objectives differ. In particular, we are *not* interested in modeling multitudes of rather simple characters. Instead, we seek to develop complex autonomous individuals that, in addition to motor and perceptual components, include broad behavioral repertoires that are much more challenging to model. Our self-animating pedestrians can independently assess the interrelationships among all the relevant factors to make rational decisions in the presence of uncertainty. Hence, they are suitable for animating the detailed behavioral interactions of small social groups.

2. Related Work

Human modeling is a broad, multifaceted subject in computer graphics. The goal of our work is autonomous virtual humans that behave intelligently in complex, synthetic worlds [ST05]. To that end, we focus on human behavioral modeling. Since the introduction of behavioral animation by Reynolds [Rey87], researchers have pursued the ethological approach to modeling animal behaviors, where the autonomous character takes actions based on its internal state and its perceptual interpretation of external stimuli [TT94, BDI*02]. Human behavior, by far the most complex of animal behaviors, is the subject of multiple disciplines, including ethology, psychology, sociology, and anthropology. Our work addresses the human character's autonomy and interaction in its virtual environment, aside from natural verbal communication and dialog, nor do we consider the cognitive level of decision making, which concerns what a character knows, how that knowledge is acquired, and how it can be used to reason and plan actions [FTT99].

In the area of human behavioral animation, Musse and Thalmann [MT97] simulated crowd behavior using a rule-based system. Badler et al. [BAZB02] proposed a Parameterized Action Representation (PAR), which includes specifications for low-level animation concepts, and descriptions of primitive or complex actions, with action selection based on the conditions specified in the PAR structure.

Closer to our approach, Ball and Breese [BB00] encoded emotions and personality using Bayesian networks. Unlike our work, however, their emphasis was on conversational agents with speech recognition and generation. Kshirsagar and Thalmann [Ksh02] also used a Bayesian network to model personality and mood in a chat application, as did Egges et al. [EZKT03] to model mood in their conversational agent. The work closest to ours is that by Hy et al. [HABL04] who simulated simple behaviors for a first-person shooter game character by using a Bayesian network to specify finite-state-machine-like behavior selection, and to learn by imitating a human player. Unlike us, they did not simulate human interactions.

Although prior animation work and existing computer game titles have used finite state machines, fuzzy logic, neural nets, scripting, smart environments, and Bayesian networks, to our knowledge, ours is the first effort in computer graphics to develop and demonstrate a unified framework for behavioral animation based on *decision networks*. The decision network (or "influence diagram"), which was introduced by Howard and Matheson [HM81] in the area of decision analysis, combines probability theory and utility theory to provide a simple visual representation of a decision problem. Decision networks extend Bayesian networks by adding actions and utilities.

Compared with other common decision-making mechanisms such as fuzzy logic [Zad88] and neural networks [Bis95], or rule-based architectures [LNR87], decision networks offer the advantage of providing an intuitive yet rigorous way to identify and display the essential elements of the problem, including objectives, uncertainties, interpretations, and decisions, and how they influence each other, as well as the clear attribution of outcomes to the inputs that generated them. We assert that decision networks are significantly better able to simulate social interactions among autonomous pedestrians.

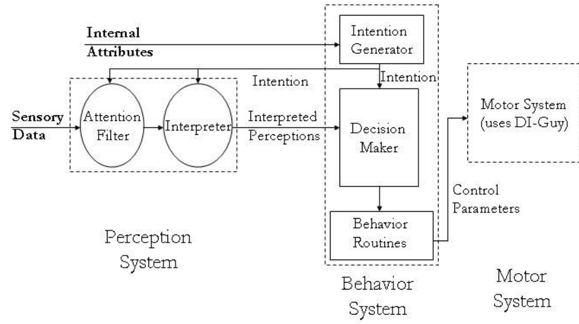

Figure 1: *Autonomous Virtual Pedestrian Model.*

3. Virtual Human Model

To evaluate our framework, we have implemented a virtual human model based on the software that was described in [ST05], including the environment model of the original Pennsylvania Train Station in New York City. Fig. 1 illustrates the architecture of our virtual pedestrian. Like real pedestrians, the synthetic humans sense their virtual environment, interpret the sensory stimuli, make decisions based on their perceptual interpretations, and act in accordance with their decisions. The important contribution of our human model is the behavior submodel, in which we exploit decision networks to simulate complex interactions between multiple pedestrians and to model the effect of different personalities on their decisions.

Each character acts autonomously within the virtual environment. To understand the structure of our system, let us consider an arbitrary character, say Jane. At any given time, Jane's intention generator assesses her current intention based on internal attributes and memory. Jane observes her surroundings to determine what objects are within her 180-degree field of view. The perceptual data Jane can gather by querying the environment model includes the position, speed, and orientation of objects in the environment, including other characters as well as their gaze directions.

Jane's attention mechanism guides her gaze depending on which objects are of interest given her current intention, or if an object attracts attention by making a sudden movement [EY97]. When Jane attends to a character she recognizes, or to a character with which she may be interested in interacting, or when there is a potential collision with some character, she draws inferences about the character using the interpreter in her perception system and decides how to interact with them. This decision making process is accomplished within our novel behavioral framework.

In particular, Jane is equipped with the new behavioral models that we develop in the next section. She also possesses a set of behavior routines that enable her to carry out primitive actions, such as walking to certain locations. Once an action selection decision is made, the relevant behavior routines are invoked to carry out the necessary actions. At a lower level, her motor system is responsible for carrying out the actual primitive movements such as walking and running. Her geometric body model and its primitive movements are provided by Boston Dynamics Inc.'s DI-Guy API.

Jane must remember the sequence of tasks she wants to perform. For this purpose, we implement a stack based memory in her behavior system, as was done in [ST05], enabling her to maintain persistence in her behaviors, while also adapting to the changing environment by storing new interim goals that attract her attention.

We have designed behavior routines to couple the decisions made at the behavioral level to the low-level DI-Guy motor system. Unfortunately, DI-Guy characters suffer limitations not just in their appearance but also in their motor skills, which restricts the possible motions that may be used in actions triggered by our decision network framework. For one character to interpret the behavior of another, it must make observations. Currently there are only a limited set of cues upon which our virtual humans can base their observations, since their facial expressions and gestures are highly constrained. The available cues include change in direction, change in speed, gaze direction [Pet05], and body orientation. Change in speed is an especially unreliable visual cue as DI-Guy characters cannot change their speed quickly.

4. Behavioral Modeling Using Decision Networks

Our new behavioral modeling approach employs decision networks as its core methodology. We have applied our decision network framework to the design of interaction models between virtual humans, guided by our commonsense knowledge of how real humans behave in similar circumstances. After motivating our framework in the next section, we present in subsequent sections four specific behavior models implemented within the framework—emergency response behavior, acquaintance behavior, partnering behavior, and collision avoidance behavior.

4.1. Decision Network Framework

A complex human behavior usually requires a sequence of assessments and decisions to be made. To model the behavior, the various contributing factors as well as their interrelationships must be identified, specified, and quantified. A decision network encodes events, represented by nodal variables, and causal relationships between them, represented by directed edges. This facilitates behavioral modeling as it is natural to think of behaviors as causal relations between events. A decision network computes rational decisions based on what the agent *wants* and what it *believes*, whereas a purely logical agent would not be able to handle uncertainty combined with conflicting goals [RN03].

We take advantage of the features of decision networks that address the key issues identified in Section 1. The uncertainties associated with various variables of interest are represented by the probability distributions encoded in the decision network. The decision network's powerful inferencing capability enables the explanation of observations made about the world, as well as predictions based on the evidence. The use of decision networks provides a convenient way to control how the character makes decisions. Adjusting the conditional probabilities and the utility functions will influence how the decision gets made. Another way to exert control is to adjust internal parameters, which will be monitored at simulation time by the network in order to make inferences and assessments.

In our application, not every decision need take all sensory and internal factors into consideration. To avoid the potential intractability of large decision networks, we build the behavioral model as a hierarchical set of relatively small decision networks. The hierarchical structure helps to ensure a manageable number of variables in each individual network and reusability of the functionality of component networks. At the lower level, a separate, smaller decision network structure is implemented for each decision item, while at higher level(s) the decision network structure at each node represents how a decision is made based on results from its children nodes. Decision network behavior models are also readily extensible, as we will demonstrate in Section 4.2.4.

Fortunately, there is no need to reinvent the wheel when applying decision networks. We use Netica (www.norsys.com), which is a commercial-quality implementation of Bayesian and decision networks along with a convenient GUI. It uses the junction tree algorithm to evaluate the networks and draw inferences. †

Applying our decision network framework, we have developed a set of networks that implement each of the four behavior models. Each network is responsible for drawing an inference or making a sub-decision that contributes to the final behavior. They are invoked only when there is a need to make the corresponding inference or decision.

4.2. Emergency Response Behavior

Our highest-level and most elaborate new behavior model is the emergency response behavior, which simulates how people might respond to an emergency situation in a public space like the train station. This behavior will serve as a concrete example to explain in detail how the decision networks are constructed, how the network parameters are adjusted and, subsequently, how we can obtain various possible decisions automatically from them and even extend them.

4.2.1. Network Construction

Our framework facilitates the application of commonsense knowledge. Regardless of the methodology that one uses to

† The role played by Netica, and DI-Guy for that matter, in our work is no different than the role played by, say, Matlab and Maya in many CG research projects. Netica serves simply as a software tool.

design a behavior model, it is natural for the designer to think in terms of what factors cause what effects, eventually leading to a decision underlying action selection. For human character animation, it often suffices to consider the various factors that contribute to making the appropriate decision in some real-life scenario. Since decision networks represent causal relationships, one can start the construction process by considering the root cause. Once we have determined the variables representing the factors we would like to consider and have understood the causal relationships among them, we effectively have designed the network topology. It then remains to complete the network by filling in the associated probability or utility values. Given sufficient data, it is possible in principle to train decision networks automatically. In modeling human behaviors, however, we often lack the necessary data to do so. Hence, we rely on intuition and common sense to tune the network parameters.

In particular, we approach the task of modeling how people respond to an emergency situation by considering their possible reactions when encountering such a situation. When someone collapses in a public place, for example, some people may approach to investigate, while others may prefer to avoid the situation and carry on with what they were doing. Of the people who do decide to approach the scene, the more concerned among them will run over, while the less concerned will walk. These will be the three possible intitial reactions from which our virtual pedestrian character can select.

Next, we consider the factors that affect this action selection decision. Three main factors come to mind: how serious the character thinks the situation is, how much the character wants to help others, and how courageous the character is. A timid character may shy away from an emergency situation in order to avoid confronting a potentially unpleasant scene at close range. Having identified these three factors, we simply use a random variable to represent each factor, and create a *chance node* (circle) for each variable, as shown in Fig. 2. The action selection decision is represented with a *decision node* (square). We also need to add a *utility node* (diamond) to indicate the character's preferences over the three input factors on the decision.

To complete the network, we must define the *prior probabilities* for the three chance nodes (upper table). The seriousness of the situation must be assessed, so its value should be set based on the interpretation result. To keep the network structure modular and hierarchical, the interpretation result can be obtained from a separate network. The prior probabilities for the Help and Courage nodes are set directly based on the character's internal factors and personality, which are represented on a continuous scale from 0.0 to 1.0. The *utility table* entries (lower table) are set based on our intuition about the character's preference over the three possible actions given the state of the input factors. For example, for a strong, helpful character, when the situation is considered serious, the character is most likely to run over and investigate, less likely to walk over and check out the situation, and least likely to ignore it. The corresponding utility values are

P(Serious)		P(Help)		P(Courage)	
yes	no	yes	no	s	w
sr	1.0-sr	he	1.0-he	n	1.0-n

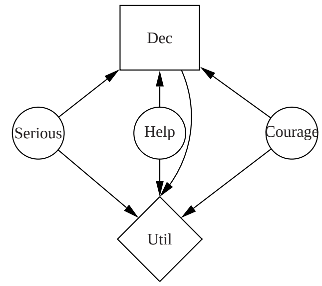

Serious=*yes*				Serious=*no*			
Help	Courage	Dec	Util	Help	Courage	Dec	Util
yes	s	ig	-10	yes	s	ig	-5
yes	s	run	8	yes	s	run	5
yes	s	obs	6	yes	s	obs	3
yes	w	ig	-8	yes	w	ig	3
yes	w	run	6	yes	w	run	2
yes	w	obs	7	yes	w	obs	4
no	s	ig	1	no	s	ig	5
no	s	run	3	no	s	run	1
no	s	obs	5	no	s	obs	2
no	w	ig	3	no	w	ig	6
no	w	run	1	no	w	run	0
no	w	obs	2	no	w	obs	2

Figure 2: *Network to determine how to respond to the emergency. Variable values s: strong, w: weak, ig: ignore, run: run over to check, obs: walk over to check.*

P(Serious)	
yes	no
0.5	0.5

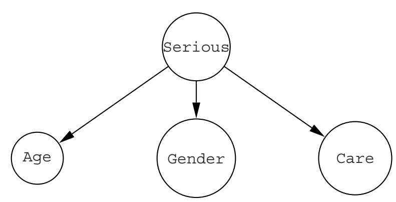

	P(Age)			P(Gender)		P(Care)	
Serious	y	p	o	male	female	yes	no
yes	0.3	0.2	0.5	0.3	0.7	0.9	0.1
no	0.3	0.5	0.2	0.7	0.3	0.1	0.9

Figure 3: *Network to assess the seriousness of the situation. Variable values: y: young, p: prime age, o: old.*

accordingly set for this input combination, with the highest utility assigned to the run action, since the network evaluation will choose the action with the highest utility value.

The Serious node represents the character's interpretation of the emergency scene that it is observing. This interpretation process is captured by a separate network (Fig. 3). Similar to the construction of the network of Fig. 2, we create a chance node for each of the cues people look for when encountering such a scenario in real life, and which are also observable in our characters given their limited DI-Guy motor system. In our current implementation, we selected the

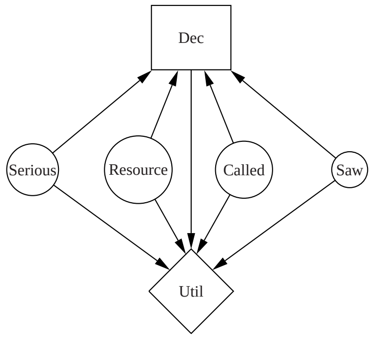

Figure 4: *Network to decide if should go fetch a police officer.*

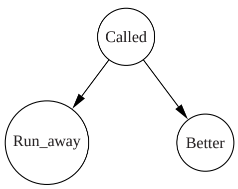

Figure 5: *Network to assess if someone else is calling the police.*

age and gender of the patient, and how caring the character is about others. The conditional probabilities are set based on the seriousness judgment and the chances that the corresponding cues will be observed.

The aforementioned procedure constructs the networks needed to decide the initial reaction to an emergency situation. Among the characters who have chosen to investigate, some may decide to seek assistance using the decision network shown in Fig. 4. This decision depends on how serious the character regards the situation and how resourceful the character is in finding help, including whether or not the character has noticed a law enforcement officer earlier and can estimate the officer's current location, and the character's assessment of whether someone else is already seeking assistance. This assessment is made by observing whether another character is running away from the emergency scene and if it seems likely that the character can summon help more quickly (Fig. 5).

Some characters may choose to leave the scene after being present for some time and unable to do much to help, a decision process implemented by the network of Fig. 6. The contributing factors include how impatient the character is, if the character is in a hurry, the time that it has been idling at the scene, if someone is already summoning help, and how altruistic the character is.

Thus, we have constructed all the networks that our virtual characters use to decide how to respond to emergency situations. The probability and utility parameters are set based on the designer's intuitions, but they may need to be fine-tuned to truly reflect how the designer wants the characters to make decisions. The next section describes the tuning process.

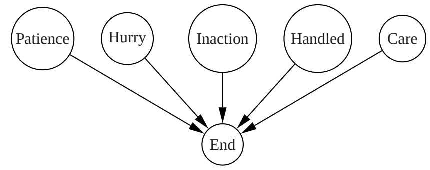

Figure 6: *Network for deciding to end the emergency response behavior.*

4.2.2. Adjustment of Network Parameters

Each network can be tested separately to ensure that its parameter settings allow it to make reasonable decisions in accordance with the designer's intentions. Netica provides a GUI for defining and evaluating the networks. Therefore, once the networks have been constructed, one can easily fine tune the parameters through the GUI by testing various values and assessing the corresponding network evaluation results in order to make the necessary adjustments.

Since our framework makes it easy to trace how a conclusion is drawn, when we see an undesirable result, we can quickly identify the cause and fix it. For example, consider the network of Fig. 3. Suppose that we initially assigned $P(\text{Care} = \textit{yes} \,|\, \text{Serious} = \textit{yes}) = 0.6$ and $P(\text{Care} = \textit{yes} \,|\, \text{Serious} = \textit{no}) = 0.4$. Trying the input combination of a male patient in his prime and a fully caring character, the situation is interpreted to be not serious. This is undesirable, since we want the caring character to consider the case serious even when the patient is a male in his prime. This evaluation result indicates that we have not given enough influence to the Care variable and need to increase the probability of its value being *yes* when the situation is serious, and similarly increase the probability of its value being *no* when the situation is not serious. Setting $P(\text{Care} = \textit{yes} \,|\, \text{Serious} = \textit{yes}) = 0.9$ and $P(\text{Care} = \textit{yes} \,|\, \text{Serious} = \textit{no}) = 0.1$ yields desirable results.

In our experience, the probabilities need not be set precisely for the decision process to work well; a range of parameter values will usually work. Fine tuning the parameters in this way is a quick and intuitive process. Given the same network structure, by modifying various conditional probability and utility parameter settings, the behavior designer can create a variety of characters that will decide to take quite different actions when faced with the same internal and external factors, even "irrational characters" if that is what the designer wants.

4.2.3. Network Capabilities

Our decision network framework includes theoretically sound inference algorithms. They make interpretations based on the evidence collected, including uncertain evidence. Once the networks have been constructed, corresponding decisions based on various input combinations are automatically generated. Regardless of the number of influencing factors and how many possible states each factor can take, for every possible combination the network is able to draw a corresponding conclusion by executing the inference algorithm.

Input factors include observations made about the character's external environment and the character's own internal factors and personality traits. Differences in these factors lead to different decisions. There can be many possible input factor combinations, but the same networks can automatically accommodate them. In the network of Fig. 2, all the input factors are binary. When we know for certain which state they are in, they are entered as *hard evidence*. However, most of the time, these factors are not strictly in one state or another. For example, the assessment of seriousness that comes from the network of Fig. 3 is represented with probability values to indicate different degrees of seriousness. In the real world, people can vary in their willingness to help others and their level of courage. In our network, these assessments can be represented on a continuous scale from 0.0 to 1.0, rather than as binary variables, and entered as the prior probability for their corresponding nodes. Different degrees of various personality traits may lead to different decisions. The network is capable of drawing inferences and making decisions incorporating these uncertainties and various other parameter inputs.

4.2.4. Network Extension

Since our framework supports a modular representation, it is easy to make extensions. When we need to add a new factor to be considered in a decision process, all we need do is add a node for it in the corresponding network and add the associated links to indicate its relationship with existing nodes. We must also add the associated conditional probabilities for the new node and make the necessary adjustments to nodes affected by the node addition.

For example, suppose we want to consider the effect of emotion on the emergency response behavior. A depressed character that is feeling low may lose interest in other things, including helping others. This would involve adding to the network of Fig. 2 a new chance node, Depress, which represents the state of depression, and links from it to the utility and decision nodes, since it is also an input factor for the decision process. Under the assumption that when the character is not depressed, it should act the same as before the inclusion of this new emotional factor, then the utility values when Depress $=$ *no* should remain the same, and we need only update the utility values when Depress $=$ *yes*. Given this extended network, with the same seriousness assessment and the same helpfulness and courage for the character, it will decide to ignore the incident instead of walk over to the scene.

The hierarchical structure of our networks keeps the scale of addition and modification to associated probability and utility settings manageable. For example, suppose that we want to model multiple emotional factors. We can reduce the number of entries that must be added to the utility table by building a subnetwork, where the overall effect from the emotional components is assessed first, and then only add that effect to the utility table. Hence, extension can be easily accomplished within our framework as the changes needed are isolated to the concerned network(s), with the addition of the corresponding nodes and links.

4.3. Acquaintance Behavior

Interpretation becomes more important in interactions between pedestrians. For example, when two pedestrians who know each other meet on the street, how they react to one another depends not only on their own intentions, but also on their interpretation of what the other pedestrian will do. We have developed an acquaintance behavior model to simulate this interaction. When two characters meet, they will choose among talking to each other, acknowledging one another without stopping to chat, or ignoring one another. For the acquaintance model, two pedestrians will not talk with each other unless both are willing to talk; hence, a reasonably accurate interpretation is needed.

Character A's interpretation of character B's intention is divided into an assessment of whether or not B is just starting to do something with A (showing an intention to talk, or to greet), or was already in the process of doing something with A, as they exhibit different cues. When B is just starting to do something with A, it would most likely be looking at A; if it wants to talk with A, it would most likely start to walk towards A and slow down. The absence of these cues is evidence that B intends to ignore A. When B is in the process of carrying out its intended action with A—for example, trying to talk with A—it may already be in the course of walking towards A, hence need not change direction at the moment. The prerequisite is that the previous interpretation already indicated that B had the intention of doing a certain action with A.

The difficulty in making such an interpretation is further complicated by the fact that B may intend to do something with A, but is temporarily interrupted by the need to avoid collision. Hence, an assessment is required first of whether or not B is in collision avoidance mode. Such an assessment is made based on change in direction, change in speed, whether B is facing some obstacles, and determination of whether B is just starting to avoid collision or is in the middle of collision avoidance. B talking with a third party will also prevent it from talking with A. Figs. 7, 8, and 9 show the networks for interpreting B's action at the moment, whether it be meeting with others, avoiding collision, or just walking.

Once this observation is made, further decisions are made based on its result. If the interpreted action for B is talking with another character, no further action is taken (in future work, we may implement the possibility of A joining the conversation). If B is in the middle of avoiding a collision, the previous interpretation for B was maintained. In the absence of any special action detection, A will continue to assess B's intention (Fig. 10). In addition, A evaluates its own intended action based on how friendly it is with B, if it is in a hurry, and how much it desires to greet or talk to B.

The decision network which determines the action to be taken is depicted in Fig. 11. The distance factor is added since the character may be uncertain about the interpretation and, hence, may take several time steps to decide. In the meantime, it is not committing to any action. However, as the two characters approach each other to within a certain threshold, a decision must be made based on information available at the time. The distance factor here indicates if such a threshold has been met. The utility function is a balance among the internal desire and the interpreted intention of the other party. The decision can take on the value of talk, greet, or ignore the other party, or uncertain, which means further observation is necessary.

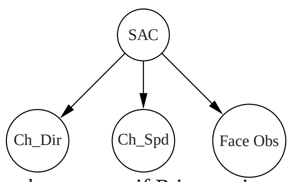

Figure 7: *Network to assess if B is starting to avoid a collision.*

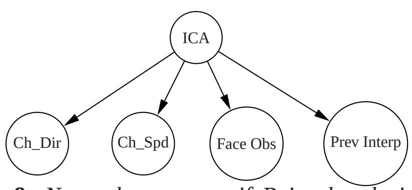

Figure 8: *Network to assess if B is already in collision avoidance.*

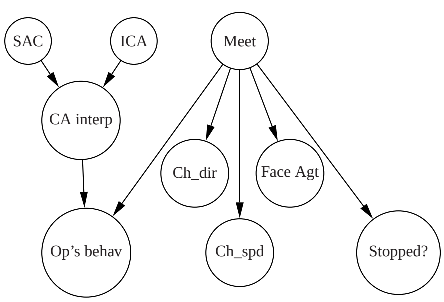

Figure 9: *Network to interpret B's behavior.*

4.4. Partnering Behavior

It is quite natural for friends meeting unexpectedly to form partnerships if circumstances permit. We have developed a partnering model to simulate such behavior. There are several possible scenarios. When one character sees another friend character ahead, it needs to decide whether to catch up and try to partner with that character; then the other character must decide whether to accept the advance. When two characters meet on the street, they may have a chat and then decide if they want to partner. In the first case, the initiating character A must first decide if it wants to catch up and form

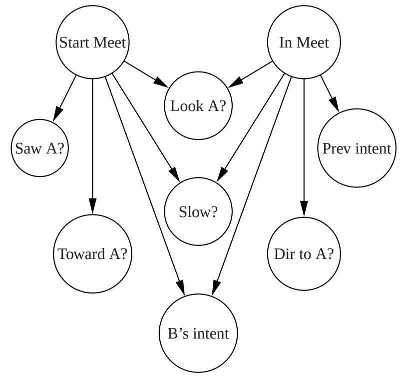

Figure 10: *Network to interpret character B's intended action with A.*

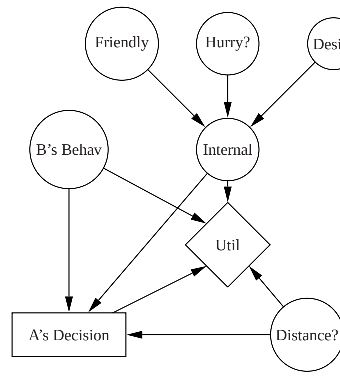

Figure 11: *Network to decide the action to be taken by A with B.*

a partnership with character B in view. This decision must take into consideration A's intention, personality, if A feels its goal matches that of B, and its perception of B's intention. Fig. 12 shows the decision network for deciding whether or not to form a partnership with the potential partner.

Once A decides to form a partnership with B in view, it will try to catch up and express its partnership request. B must then decide whether to accept this request. As shown in Fig. 13, how friendly B feels about A, if B is social, and B's desire all affect the decision. In addition, how persistent B is in pursuing its own intention and how long A has persisted in making the advance also play a role in B's reaction.

As B makes up its mind, A is observing B and tries to assess if B has accepted its advance (Fig. 14). Not all observed data have equal reliability. For example, when trying to interpret B's response to character A's request to partner, B looking at A and stopping are much stronger indications than B turning to face A. When B accepts A's request, B may or may not choose to face A directly. To accommodate this difference in reliability of the observed data, some adjustments are made in the network; hence, the additional reliability chance node (Fig. 14).

If B accepts A's request, the two characters will form a pair and proceed. Should B refuse the request, A must then

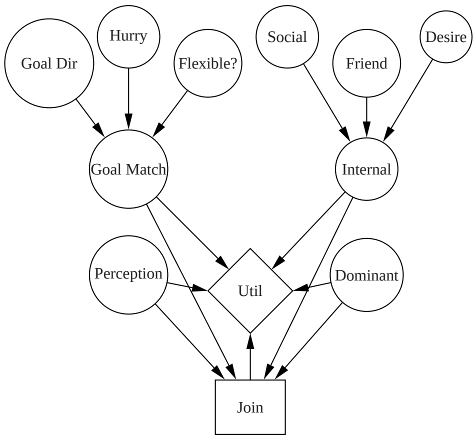

Figure 12: *Network to decide whether to form a partnership.*

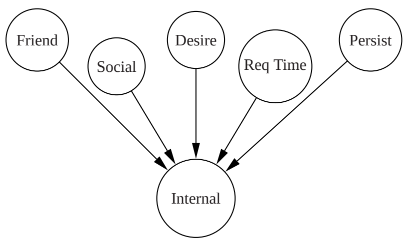

Figure 13: *Network to assess whether to accept a partnering request.*

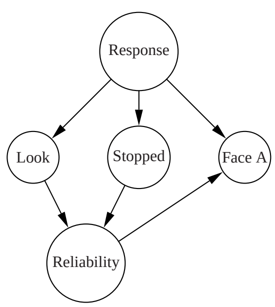

Figure 14: *Network to interpret B's response to A's pair-up request.*

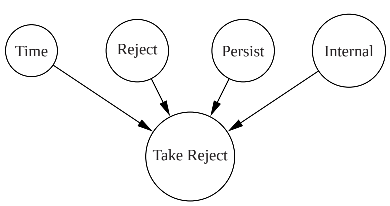

Figure 15: *Network to assess how to take rejection.*

decide how it will react to the rejection; i.e., give up or continue its persuasion effort (Fig. 15). This cycle of A making a request, B deciding whether to accept it, A's interpretation of B's response, and decision about what to do next continues until either B is persuaded to accept the partnership offer or A gives up. Once a partnership is formed, negotiation continues between the partners to maintain the partnership and handle unexpected events, such as collision threats, and making decisions together as a team.

4.5. Collision Avoidance Behavior

Our framework is also applicable to low-level collision avoidance behavior modeling. Much behavioral animation research has focused on collision avoidance. Our framework incorporates the interpretation of the opposing character's avoidance intention, which is absent in earlier models. Our model implements a local object avoidance strategy rather than global path finding.

The character first tries to avoid any static obstacles. The avoidance strategy includes the options of steering to the left, steering to the right, or stopping when in imminent danger of collision. As for dynamic obstacles, a collision possibility value is computed for each obstacle and the one with the highest value assumes priority for avoidance. The optimal avoidance strategy can be one of the following: no avoidance, speed up, slow down, stop, steer towards left, steer towards right, steer towards left with speed change, steer towards right with speed change, stop while steering towards left, and stop while steering towards right. Due to the limitations of the motor system, where a speed change command has a delayed effect, currently we do not use the speed change options, but have left them in place in case a better motor system becomes available. The utility values are weights associated with the avoidance strategy options, which are determined by our observations of how real humans balance avoidance choices under different circumstances.

The details of our collison avoidance decision networks are presented elsewhere [Yu07].

5. Animation Results

We have equipped our virtual humans with the aforementioned behavior models which take advantage of our decision network framework's ability to handle uncertainty and behavioral complexity. As described above, the collision avoidance model anticipates the opposing character's avoidance strategy and makes a decision accordingly. In addition to simulating the interpretations that a character draws about its environment and about its potential partner's intentions, the acquaintance, partnership, and emergency response models also take into account the effect that personality and internal factors have on decision making. Decisions are made based on observations and the analysis of objects that fall within the character's focus of attention. Our accompanying animations demonstrate the effectiveness of our decision network framework.

Our most complex animation is the "emergency response" animation illustrated in Fig. 16. A young female commuter

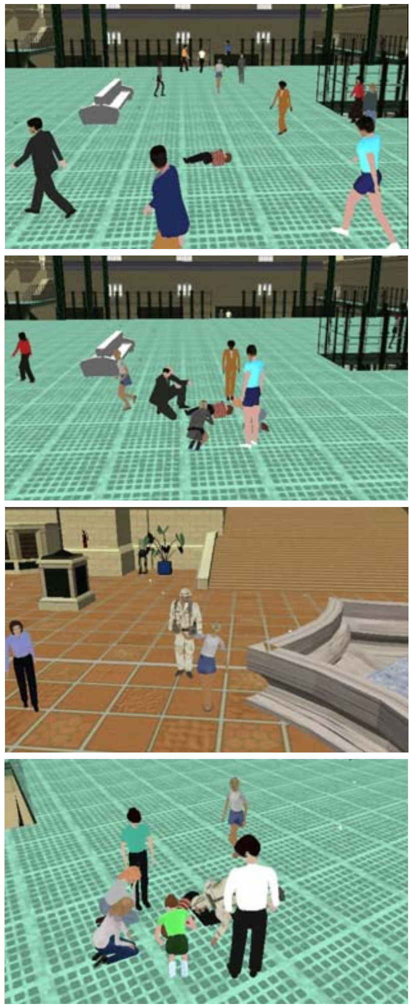

(a) (b) (c) (d)

Figure 16: *Still frames from the emergency response animation. (a) Female fainted on floor in concourse. (b) Commuters gather around to attend to patient; a woman runs for help. (c) Woman finds a security officer in the main waiting room. (d) The officer attends to the patient.*

feels ill, staggers, and then faints to the floor of the concourse area in the virtual Penn Station. Some commuters around her turn their heads to look, but keep on walking. Others walk over to see what happened. The most concerned run over to investigate. After observing the patient's situation, a female commuter recalls seeing a security officer in the main waiting room of the station and, hoping that he is still there, decides to run over to ask for help. While she is gone, some of the commuters attending to the patient decide to leave as they determine that they cannot do much more given that someone has already gone to summon assistance, while some new passers-by approach to investigate the scene. The security officer who was summoned eventually reaches the scene, examines the patient and decides to radio paramedics for further assistance.

Shorter simulations demonstrate the functionality of the partnering behavior model, differing only by internal factor and personality settings. One animation shows pedestrian A eventually giving up on the decision to partner with B, after repeatedly being shunned by B. A second animation shows a pedestrian A who notices a friend B walking ahead, catches up, and tries to partner with B. Because of an earlier quarrel with A, B initially refuses A's advance, but reconsiders upon A's insistence.

Our framework is efficient. The network with the largest utility table takes from 1 to 2 msec to build and compile in Netica and usually less than 1 msec to evaluate. The emergency response network takes about 1 msec to compile and 1 msec to evaluate. At run time, the behavior models are event triggered—i.e., whenever the circumstances trigger certain behaviors, the corresponding networks are evaluated to make the necessary decisions, but they are invoked only on a need-to-use basis. Collision detection takes place every few time steps, but the collision avoidance networks are invoked only when a potential collision is detected. These networks are invoked the most frequently. Each character independently makes its own decisions based on what it can observe of its environment and its personal internal factors. Although different execution times are needed to arrive at different decisions, this does not adversely affect the parallel simulation of different characters. On an Intel Xeon 3.2GHz PC with 1GB RAM, the behavior system for a single pedestrian usually takes under 1 millisecond to execute.

6. Conclusion

We have introduced a decision network framework for advanced behavioral animation in virtual humans. To our knowledge, ours is the first virtual human behavior modeling system and architecture that is based on decision networks, and the first use of decision networks in computer graphics. Our decision network framework offers several advantages. It has the ability to handle uncertainties, which are largely ignored in previous behavior models. Having the characters act based on the often imprecise interpretation of their surroundings rather than on exact information gained through querying the world database results in more realistic simulations that more faithfully approximate how real people behave. This also facilitates the simulation of interactions between characters.

Because decision networks offer a formal yet easy to understand representation of the world and the behavior system, our framework makes it easier to translate the character designer's commonsense knowledge into code that governs the characters' decisions. It facilitates parameter fine-tuning and debugging, and makes it easy to identify what factors contribute to the final decision. Modifications take only minutes to complete, as the structure is so clear, compact, and modular. Given the same network topology, we can have characters reasoning and acting quite differently by specifying different probability and utility settings. Decision networks also make the behavioral model flexible and able to respond sensibly to a wide range of changes in the character's internal and external environment. Our results demonstrate the effectiveness of our framework. Its hierarchical structure keeps computational complexity at bay, and its clarity provides insights into behavioral mechanisms while facilitating the design, implementation, and debugging of behaviors. The network construction and parameter adjustment took less than a day of effort for each of our behavior models.

Although our decision network framework is general and largely independent of the underlying motor system, the specific behavior models we have built rely on the limited cues that our virtual humans can observe about one another to arrive at interpretations. As better motion APIs become available that provide a richer variety of human-like visual cues, more varied and accurate interpretations will be possible and the quality of the behavior model will improve.

Decision networks are not suitable for handling domains with co-related decision factors [Jen01], as the large computational requirements for such domains would make it intractable to handle complex problems. Given its acyclic nature, the decision network is also not appropriate for modeling behaviors that involve cyclic causal chains or recursive plans. One way to resolve this is to use dynamic decision networks which handle time series of decision problems by regarding causation as a temporal phenomenon.

In future work, we plan to expand our interaction models to simulate more sophisticated coordination and cooperation behaviors among multiple characters. We would also like to develop an intuitive user interface to facilitate the graphical construction of behavior models, enabling the user to input parameters that specify personality traits and internal factors on the fly; these values will be entered as evidence to the behavior models so that the simulation will immediately reflect the changes.

Acknowledgements

We would like to thank Wei Shao for kindly providing his software system for use in our research. We also thank Hanieh Bastani, Sam Hasinoff, and Mike Daum for their voice acting for the animation audio tracks, which are dramatizations of what the characters might say were they able to speak.

References

[BAZB02] BADLER N. I., ALLBECK J. M., ZHAO L., BYUN M.: Representing and parameterizing agent behaviors. In *Computer Animation* (2002), pp. 133–143.

[BB00] BALL G., BREESE J.: Emotion and personality in a conversational character. In *Embodied Conversational Agents* (2000), MIT Press, Cambridge, MA, pp. 189–219.

[BDI*02] BLUMBERG B., DOWNIE M., IVANOV Y., BERLIN M., JOHNSON M. P., TOMLINSON B.: Integrated learning for interactive synthetic characters. In *Proc. ACM SIGGRAPH 02* (San Antonio, TX, 2002), pp. 417–426.

[Bis95] BISHOP C. M.: *Neural Networks for Pattern Recognition*. Oxford University Press, Oxford, UK, 1995.

[EY97] EGETH H. E., YANTIS S.: Visual attention: Control, representation, and time course. *Annu. Rev. Psychology 4* (1997), 269–297.

[EZKT03] EGGES A., ZHANG X., KSHIRSAGAR S., THALMANN N. M.: Emotional communication with virtual humans. In *Multimedia Modelling* (Taiwan, 2003).

[FTT99] FUNGE J., TU X., TERZOPOULOS D.: Cognitive modeling: Knowledge, reasoning and planning for intelligent characters. In *Proc. ACM SIGGRAPH 99* (Los Angeles, CA, 1999), pp. 29–38.

[HABL04] HY R. L., ARRIGONI A., BESSIÈRE P., LEBELTEL O.: Teaching bayesian behaviours to video game characters. *Robotics and Autonomous Systems 47*, 3 (2004), 177–185.

[HM81] HOWARD R., MATHESON J.: Influence diagrams. *In Howard, R. and Matheson, J., editors, Readings on the Principles and Applications of Decision Analysis II* (1981), 721–762.

[Jen01] JENSEN F. V.: Bayesian networks and decision graphs. In *Bayesian networks and decision graphs*. Springer, 2001.

[Ksh02] KSHIRSAGAR S.: A multilayer personality model. In *Proceedings of the 2nd interna-tional symposium on Smart graphics* (2002), pp. 107–115.

[LNR87] LAIRD J. E., NEWELL A., ROSENBLOOM P. S.: Soar: An architecture for general intelligence. *Artif. Intell. 33*, 1 (1987), 1–64.

[MT97] MUSSE S. R., THALMANN D.: A model of human crowd behavior: Group inter-relationship and collision detection analysis. In *Computer Animation and Simulation '97, Proc. EG Workshop* (Budapest, 1997), Springer-Verlag, pp. 39–52.

[Pea88] PEARL J.: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, San Mateo, CA, 1988.

[Pet05] PETERS C.: Foundations for an agent theory of mind model for conversation initiation in virtual environments. In *Proceedings of the AISB '05 symposium on Virtual Social Agents: Mind-Minding Agents* (Hatfield, England, 2005).

[Rey87] REYNOLDS C.: Flocks, herds, and schools: A distributed behavioral model. *Proceedings of ACM Computer Graphics 21*, 4 (Jul 1987), 25–33.

[RN03] RUSSELL S., NORVIG P.: *Artificial Intelligence - A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, 2003.

[ST05] SHAO W., TERZOPOULOS D.: Autonomous pedestrians. In *Proc. SIGGRAPH/EG Symposium on Computer Animation (SCA'05)* (Los Angeles, CA, July 2005), pp. 19–28.

[TT94] TU X., TERZOPOULOS D.: Artificial fishes: Physics, locomotion, perception, behavior. *Computer Graphics (Proc. SIGGRAPH'94)* (1994), 43–50.

[Yu07] YU Q.: *A Decision Network Framework for the Behavioral Animation of Virtual Humans*. PhD thesis, Dept. of Computer Science, Univ. of Toronto, Toronto, ON, January 2007.

[Zad88] ZADEH L. A.: Fuzzy logic. *Computer 21*, 4 (April 1988), 83–93.