

A Template-Based and Pattern-Driven Approach to Situation Awareness and Assessment in Virtual Humans

Weixiong Zhang
USC/Information Sciences Institute
Marina del Rey, CA 90292, USA
zhang@isi.edu

Randall W. Hill, Jr.
USC/Information Sciences Institute
Marina del Rey, CA 90292, USA
hill@isi.edu

Abstract

Situation awareness and assessment are fundamental components of rational virtual humans. They are central parts of models of virtual humans and critical elements of efficient reasoning and planning systems. In this paper, we describe our efforts of developing the capability of situation awareness in autonomous, synthetic virtual pilots in a military domain. After briefly describing the motivations of this research, we present an agent architecture that integrates perception, reasoning with situation awareness, and actions. We describe a representation of situations, methods for situation assessment, and applications of situation awareness in information collection, and creating focused attention. The quantitative measure from our experiment shows that by using focused attention based on the ability of situation awareness, the perceptual load of virtual pilots is significantly reduced, reducing pilots' response time by more than 50 percent.

Keywords

Situation awareness and assessment, perception, organizational and spatial relationship, templates, pattern matching.

1. Introduction and Motivation

One of the most important characteristics of virtual humans is the ability to be aware of current situations of the environment where they reside and operate. Situation awareness plays a central role in cognition, which spans the entire spectrum of cognitive activities, from perception, such as lower-level vision activities, to high-level cognition, such as understanding, reasoning, and decision making.

1.1 Why is situation awareness needed?

1.1.1 Long-term goals

Our research on situation awareness and assessment is motivated by both of our long-term and short-term goals. One of our long-term goals is to develop realistic and feasible models of virtual humans in synthetic environments for training [13], military simulations for mission rehearsal and tactics evaluation [15,9],

and other applications such as animated characters and entertainment. Specifically, we are interested in developing models of realistic, human-like characteristics, such as rationality, personality, and emotion. Such cognitive models are also potentially useful for testing computational models of psychological theories.

A model of virtual humans must implicitly or explicitly represent their conscience at any given time, which is only meaningful when put into perspectives of different situations. Situation awareness is a foundation of rational virtual humans. It provides the context under which virtual humans can identify achievable goals as well as establish a course of actions that abides by their limited resources and external constraints. Situation awareness is also an indispensable element of efficient reasoning and planning in virtual humans. A sense of the current situation can eliminate uncertainties in reasoning and planning, and provide plausible expectations of future actions.

1.1.2 Near-term goals

We have been developing autonomous, synthetic virtual pilots for military simulation for mission rehearsal and tactics evaluation for a number of years now [15,9]. Our short-term goal is to enhance such virtual pilots with rationality, an important feature of human-like characteristics. By rationality we mean the ability to reason and act based on the current situation of the world, the conditions of external environments, and internal resource constraints.

There are many different ways that situation awareness can be used in virtual pilots, such as in planning and understanding reactions from others. In this paper, however, we focus on an important issue of rationality, the focus of attention in virtual pilots that relies on situation awareness. Attention forms the nexus of perception and cognition. The cognitive system needs what the perception system offers, including visual objects with features such as location, color, orientation, and motion. Without these percepts, the cognitive system would not know about the external world. At the same time, the perception system requires the focus and control provided by attention to allocate limited processing resources to selected regions of the visual field. Without a focus of attention, the perception system can be overloaded, in that too much information needs to be processed before an action is taken, resulting in unpredictable consequences. Focus of attention can be regarded as a means of allocating limited computational resources in the process of perception. The goal is to reduce the workload of the cognitive system by focusing attention to the most important perceptual information under the current situation.

The cognitive system steers attention, focusing the perceptual resources according to goals and tasks. Conversely, attention can

be involuntarily captured by certain types of stimuli in the visual field, providing inputs to the cognitive system that were not sought after but which may have profound effects on subsequent decisions. The existing research, unfortunately, does not provide a theory or guidance on how perception and cognition actually interface as a coherent whole via the attention mechanism [9]. In this research, we propose to use situation awareness as a working interface between perception and cognition via attention. The most plausible situation will guide the perception system to the area of visual field where information should be collected. The processed visual information can be used to verify hypothesis of current situations, to generate future actions, and to expect possible reactions.

1.2 What is needed for situation awareness?

Two elements are needed to support situation awareness. The first is a representation of the situation. The representation needs to include information about relevant objects, their features and logical, organizational and spatial relationships, actions for supporting understanding the situation, and possible actions for responding to different perceptual input and external events. We call such a representation a *situation template* or *template* for short

The second element to support situation awareness is a set of tools for situation assessment. Situation assessment has at least three objectives. It is to correctly identify the relevant objects in a visual field from perceptual inputs, find association relationships among the perceived objects and create a structured representation of the objects, and map the structured representation to possible situation templates and identify the most similar ones. We call a structured representation of a set of sensed objects a *pattern*.

In this research, we develop a representation of situations, or situation template, for virtual pilots in a synthetic battlespace. The central piece of a template is a tree-structured, hierarchical representation of the organizational and spatial relationships of entities, such as tanks and helicopters, and of units, such as a company of tanks. This tree-structured representation contains grouping information of entities and units. A unique feature of this representation is that it can support different levels of abstraction, which leads to applications for controlling attention.

The tree-structured representation is also used to represent patterns of objects from perceptual inputs. We used this tree-structured representation to develop a set of algorithms for matching a pattern of perceived objects against a set of situation templates. Based on the matching results, which are mostly partial matching, different hypotheses about possible situations can be established to drive the subsequent perception and reasoning processes.

1.3 How can situation awareness be used?

Situation awareness can be used in many different aspects of virtual humans. A perceived situation can be used to determine achievable goals; select appropriate strategies and tactics; determine a course of actions; predict possible reactions and consequences; establish focused attention; intelligently allocate limited resources; provide explanation of decisions made; and reduce uncertainty and speed up reasoning processes.

Each of the possible applications of situation awareness mentioned above requires a substantial amount of research. In

this research, we focus on situation awareness and its applications to focused attention and better resource allocation mechanisms in the interface between cognition and perception, by determining what should be sensed and brought into a cognitive system's working memory for efficient reasoning.

1.4 Structure of the paper

The paper is organized as follows. In Section 2, we describe a system architecture that integrates the capability of situation awareness with perception, planning, and decision making. In Section 3, we discuss a representation for situations. We consider situation assessment in Section 4. In Section 5, we apply situation awareness to information collection for verifying situation hypotheses, and to control attention for reducing perceptual overload. A quantitative measure of the utility of situation awareness in reducing perceptual overload is also given. We discuss related work in Section 6, and conclude in Section 7.

2. System Architecture

The overall system of virtual pilots is implemented in two layers. The first layer defines the connection between virtual pilots and the synthetic battlespace as implemented in a simulator. The second layer describes the virtual pilots' cognition and its interaction with the perception that is directly supported by the simulator.

2.1 Virtual pilot and simulation environment

A virtual pilot is implemented in a distributed, interactive simulation system called ModSAF. The virtual pilot flies a synthetic helicopter and performs tactical operations with a team of other pilots. Teams of pilots are deployed along with thousands of other entities (i.e., tanks, trucks, individual combatants, airplanes, etc.) in what is known as a synthetic battlespace [15].

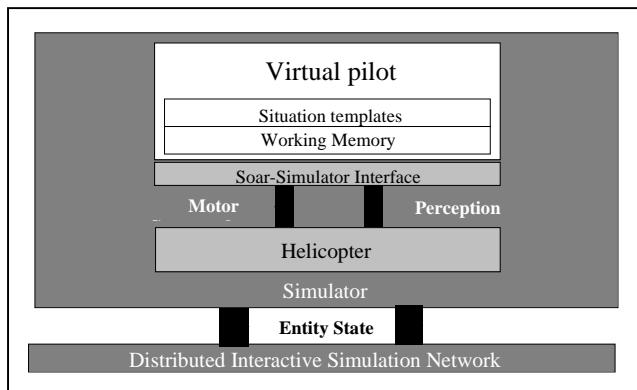


Figure 1: Virtual pilot and simulation environment.

Figure 1 shows the simulation environment, the virtual pilot and the interface between them. The perceptual and motor systems are implemented such that they provide the interface between the simulator and the virtual pilot. The simulator has a model of the simulation environment including terrain, entities in the battlespace, etc. It provides perceptual information to the virtual humans through an API. It can provide information of entities in the battlespace, such as their types, identity, locations, colors, orientations, velocity, etc. Given that the state information of the

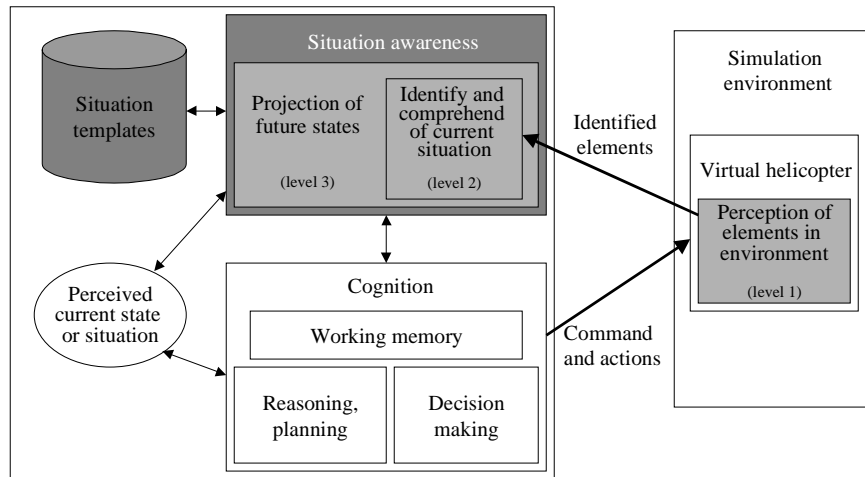


Figure 2: Architecture integrating situation awareness, perception, reasoning, and decision making.

perceptible entities is directly available to the virtual pilot, many of the standard vision problems can be finessed, so that we can concentrate on modeling and reasoning. For instance, understanding whether an entity from the perception inputs refers to a known entity, or whether it represents a newly perceived object can be done by comparing the identities of the two entities.

2.2 Architecture for situation awareness

The second layer of the system architecture integrates situation awareness, planning and decision-making, and perception. Figure 2 is a sketch of this layer. In principle, the cognition of the virtual pilot should include the modules for situation awareness, reasoning and decision-making all together. We separate the situation awareness module from the overall cognition system to emphasize the role that this module plays as an interface between perception and cognition.

Following Endsley[5], the situation awareness module has three major components, namely, the perception of the environment, the identification and comprehension of current situations, and the projection of future states. These components reside at different levels of cognition and form a chain of information of perceived objects, possible situations, and decisions.

The overall inner working process of the virtual pilot goes in cycles. It starts at the perception when the perceptual system searches the visual field and generates a set of visual entities in the visual field. The situation awareness module then processes the perceptual information and finds the most possible situation indicated by the perceptual inputs. The perceived current situation subsequently drives the process of reasoning and decision making. Subsequently, actions may be executed to affect the current situation. The inner working process then starts another cycle of perception, reasoning and action.

2.2.1 Perception in synthetic battlespace

Perception in the virtual pilot involves four distinct problems: perception of terrain, perception of messages, perception of cockpit instruments, and the perception of entities. It is critical to properly address and solve all these problems in order to fly a helicopter and to successfully carry out a mission. The perception

of entities provides the most critical information to current situations, since current situations are directly tied to the entities involved. In this paper, therefore, we focus on the issue of focusing attention in the context of perception of entities.

Entity perception is driven by the arrival of a stream of entity-state updates provided by the simulation interface. Each update characterizes the momentary state of an entity: it provides information about the identity, location, and velocity of an entity, such as a tank. These updates are filtered through models of the pilot's visual sensors to determine what information is potentially perceptible. Entities that are too far away will be imperceptible. Entities within the perceptible range of the model may still be rendered imperceptible if they are occluded by a terrain feature or an environmental factor such as smoke or dust. The sensor models also determine the resolution of the percept based on factors like distance, dwell time, and visibility. Hence, an entity may initially be recognized only as a vehicle when perceived at a great distance, but it may be identifiable as a specific tank model at a closer range.

2.2.2 Identification and comprehension of situation

A detailed description of this component is one of the major topics of this paper, which is covered by Sections 3 and 4. We give a brief account here to highlight its main steps.

The first step is to organize the perceived entities from the perception. The goal is to encode the elements into a structure, called a *pattern*, to reflect their organizational and spatial relationships. For example, according to military doctrine, ten tanks can be organized as a tank company with one commanding tank. Such a tank company takes different spatial formation when deployed under various battlefield conditions. Particularly, the formation in an attacking posture is different from that in a defending posture.

The second step is to identify the current situation based on the patterns of the sensed entities. This step involves matching the patterns against the situation templates, and identifies the closely matched template. It constitutes the core of situation assessment.

2.2.3 Projection of future states

Sitting atop the situation identification module is the module for projecting future situations to determine the actions that need to be taken. This module supports short-term planning and option evaluation in the reasoning and planning processes.

3. Situation Representation

A good representation of the situation needs to capture a number of aspects of situation: the essential entities, their organizational and spatial structures, and the possible actions that can be taken to affect internal states and external environments.

3.1 Organizational and spatial structures

Different situations involve different organizational structures of entities in which different parts play different roles. In military domains, organizational structures and spatial structures normally interleave with each other. For example, a company of tanks may take different formations in different battlefield situations, and different tanks or subunits play different roles in a formation.

We developed a tree-structured representation for both spatial and organizational structures of entities in a synthetic battlespace. This representation is based on the k-d tree data structure [8] which has been widely used in many areas such as computational geometry and computer vision. To make our discussion concrete, we now present a working example of a tank company to show the tree representation and how it is created.




Figure 3: A two-level formation of a tank company.

Figure 3 shows a two-level formation of a company of ten tanks, labeled from T1 to T10. At the low level, three tanks form a platoon on the bottom, and three platoons and a commanding tank form a Tank Company on the top. The three tanks in a platoon in Figure 3 have a line formation, and the Tank Company has a wedge formation.

To represent a formation of more than one entity, we first divide the plane that the formation occupies along the horizontal or vertical direction until a subdivided plane does not have more than one entity. The dividing direction depends on the orientations of the entities and the shape of the area they occupy. In our tree representation, a node represents a plane or subdivided plane, and two nodes form a parent-child pair if one plane is created by subdividing the other. We also allow a dividing line to go through an entity. This typically happens when the formation is symmetric along the dividing line. If this occurs, the node representing the plane to be divided will also contain the information of the entity transected by the dividing line. The overall representation is then a tree. The root node represents the

original plane and its two child nodes are the subplanes created by a division of the plane. A leaf node represents an entity and contains detailed description of its properties, such as the type, location and speed of an vehicle. An internal node represents an organization and maintains its property, such as number of members in the organization, their capabilities, etc.

To further represent the spatial relation between a plane (entity) and its parent plane (entity), we introduce a numeric code to specify the spatial relation. For instance, if the entity in the parent node is directly north of the entity in the child node, a numeric value 1 is assigned to the child node. This spatial code is used in reasoning of situations.




Figure 4: Division of a plane and formation.

Figure 4 illustrates how the formation and plane of Figure 3 are divided, and Figure 5 shows the tree representation of the formation. We would like to point out that this tree structure also nicely represents the organizational structure of a formation. This is because the entities in a sub-organization tend to be close to one another, which is captured by the tree representation.




Figure 5: Tree-structured representation of formation.

Another important feature of the tree representation is that it provides different levels of abstraction of an organization. As illustrated in Figure 5, the root node of a tree represents an entire organization, and a node deep in the tree contains information of a subunit. If a deliberation process only needs to reason about an organization at a particular level, it only needs to consider the nodes at the corresponding level of the tree. By selecting nodes at different levels, one can shift the reasoning focus among the various abstraction levels.

3.2 Actions

In addition to a tree representing entities and their relationships, a situation template also include a set of possible actions that can be applied under the situation. There are two types of actions. The

first type is for short-term planning and option evaluation in the process of deliberation of the current situation and possible future states in terms of overall objectives. These actions are represented as production rules. All applicable rules will be placed in the working memory and be selected to fire. The results may lead to a change of agent's internal state, or a command of actions sent to the actuator, which can be the motor of a robot, a sensor, or a synthetic helicopter in our case, to react to the environment.

The second type of action involves directing the actuator immediately. Such actions require very limited deliberation, and may bypass the reasoning system and go directly to drive the actuators. The number of such actions, however, is typically very small. They are often the ones used in emergencies, such as changing the moving direction of a robot before it collides with an obstacle or changing the flying course when the virtual pilot senses a missile coming in its direction.

4. Situation Assessment

There are two major steps in situation assessment. The first is to organize the sensed entities from perception. The second step is to match the observed entities against situation templates and to identify the most probable situations the observation indicates. The second step can be viewed as a classification problem in which the observation is classified into related situation templates under some closeness measure.

4.1 Organizing objects, finding patterns and situations

To facilitate a mapping of sensed entities to a situation template, the entities need to be organized into what we call an object pattern. An object pattern is nothing but an encoding of organizational and spatial relationships of a set of perceived entities. Therefore, the tree representation discussed in Section 3.1 is used to represent object patterns. Compared to other representation schemes, this tree representation provides some benefits to the later matching step. Using the same tree representation for both situation templates and object patterns, situation assessment becomes a problem of tree matching, and the tree matching problem becomes a tree alignment problem. Furthermore, the similarity between a pattern and a template can be measured based on the closeness of the alignment of the two.

We use two approaches to creating object patterns from perception inputs. The first is a well-known method but has a poor performance. The second is newly developed, and has some favorable features, such as robustness and scalability.

The first approach is based on the similarity and proximity of entities, which are commonly used Gestalt principles for grouping. In other words, objects are grouped using the nearest-neighbor method. In this method, we only put similar objects (e.g. vehicles with same type) that are close to one another into a group. In each decision cycle, newly sensed visual objects are compared to existing groups. If a visual entity is within a predefined distance to the center of mass of a group, and its attributes are similar to the existing group members' attributes, it is clustered with that group. If no such clustering occurs, a new group is formed. To form a hierarchical representation, we use a set of distances for clustering at different levels. For instance, the distance between two tanks in a subgroup in Figure 3 is smaller than the distance between two subgroups.

This similarity and proximity based grouping approach is cursed with two problems. There are many situations where more than one type of entity may appear. Grouping based on similarity can fail miserably to find correct patterns. The other problem of this approach is its sensitivity to the distance parameters for defining nearest neighbors. It is very often the case that different situations may involve entities at different distances. An approach based on generic proximity is rigid, in that it uses a set of fixed parameters to handle all different situations.

The second approach to finding object patterns uses situation templates. The process starts with a selection of possibly relevant situation templates. It then builds a set of patterns, one for each selected template, using structures, features and parameters of the template. The resulting patterns are measured against their templates, and the best ones are chosen to predict the most probable situation. The psychological rationale behind this approach is that the goal of pattern recognition is to map what has been observed to what have been known, and the recognition process uses a priori knowledge as guidance [10]. We now discuss this approach in detail.

The types of observed entities and their features are good indicators of the most probable situations. For instance, if all the sensed entities are trucks, then the situation is probably related to transportation, so only the situation templates that contain trucks need to be considered. Furthermore, the number of perceived entities can be used to further restrict the search for possible situations. For example, if ten trucks have been reported, then the templates with less than ten trucks can be ruled out. In short, the features of sensed entities can be used to pre-select possible situation templates, providing significant savings in computation.

With a set of pre-selected situation templates, the next step is to build a set of patterns of the observed objects, one pattern per template. Building a pattern based on a template is to create such a pattern that is as similar to the tree structure stored in the template as possible, and then to measure the similarity between the pattern and the tree of the template. In building a pattern, the parameters of the tree of a template, such as the distances of entities and their spatial codes, are used as parameters to cluster the sensed entities and organize them into another tree structure. The algorithm for building a pattern works in a bottom-up fashion. It first retrieves the distance of the lowest-level cluster in the tree of a template, and uses the distance as the proximity parameter for clustering the perceived entities. Each set of clustered entities is stored in a subtree, with the root representing the cluster and its children the entities in the cluster. The algorithm recursively applies itself to the clusters, creating a set of clusters of clusters. The algorithm continues until one cluster is formed.

How a pattern is similar to its corresponding template needs to be measured, and the similarities of all patterns need to be compared in order to choose the best one. In this work, we developed a domain-independent measuring method. The similarity of a pattern to its template is measured as follows. A mismatch between two nodes, one from the pattern and the other from the tree of the template, will incur a penalty. The penalty is greater if the nodes are higher in their trees, since a node higher in a tree represents an organization in a higher organizational hierarchy and a mismatch between two larger organizations will impose a bigger impact on the overall matching. In our system, we

experimentally adjust a set of penalty weights for mismatches at different depths.

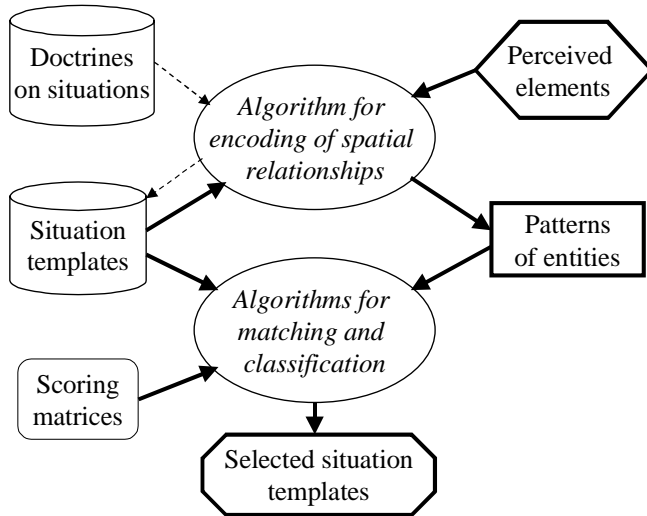


Figure 6: Situation recognition.

The overall process of processing perceptual information, organizing perceptual inputs and building object patterns, and searching for the best situation template is summarized in Figure 6. As shown by the dotted arrows in Figure 6, the situation templates are created in an off-line, preprocessing stage, which are stored in a database. During the process of situation assessment, patterns of perceived objects are built and compared against a set of candidate situation templates based on some scoring matrices. A small set of selected situation templates is then used to generate hypotheses of current situation for reasoning. Note that the algorithm for encoding the spatial and organizational of objects is used in both creating situation templates and patterns of observed elements.

5. Applying Situation Awareness

In our first applications of situation awareness, we are using it for collecting information in synthetic battlespace and creating focused attention to reduce perceptual overload in the virtual pilot. Collecting additional information may be crucial in some situations, such as where enemy's attacking units are hiding behind a hill. Creating focused attention can solve the serious problem of perceptual overload, which caused severe performance problem in virtual pilot.

5.1 Information collection

Visual perception reflects a momentary state of an environment. Perception provides information about the identity, location, and velocity of an entity, such as a tank. Entities that are too far away will be imperceptible. Entities within the perceptible range may still be rendered imperceptible if they are occluded by a terrain feature or an environmental factor such as smoke or dust. Consequently, it is frequently the case that a pattern of sensed entities will not completely match to any situation template. Indeed, in military domains, terrain features, such as hills, are usually exploited to hide attacking units.

When a mismatch occurs, additional information needs to be collected to verify whether the best-matched template should be taken for the current situation. This is important because a correct model of the current situation will affect the agent's reasoning and decision-making. To collect additional information, the entire visual field should not be scanned or randomly searched, to do so could be costly. Rather, selected regions should be examined.

The mismatch between the pattern of perceptual entities and the template of the situation to be verified provides strong cues where additional information can be found. These are the places in the visual field corresponding to the mismatches between the pattern and the template. Searching for missing objects can be confined to these possible regions. In this search process, the virtual pilot may need to fly the helicopter to a certain height or fly over a hill to avoid visual obstacles. The verification process may take a few iterations in which the pattern is refined. In short, situation templates can be used to drive the process of information collection and situation verification.

5.2 Perceptual overload: The problem

How wide and deep should the visual field be set? In other words, how much visual information should the virtual pilot receive in each decision cycle (discussed in Section 1.1.1)? The pilot's reasoning system does not impose a limit on perceptual processing. When no limit is set, the entire visual field or a large portion of it needs to be used to provide sufficient information for reasoning. In the case of entity perception, this primarily involves computing a variety of geometric relationships between the pilot and each entity. It turns out that the cost of these computations can be expensive, especially when there are many entities in the visual field. Without a way of limiting perceptual processing, overload conditions can occur, where there are hundreds of entities in the field of view and the pilot perceptually attends to all of them, in every decision cycle. When this occurs, the virtual pilot begins to lose control of the aircraft and sometimes crashes into the terrain because it cannot keep up with the perceptual processing demands while simultaneously flying the helicopter (and performing other tasks) in real-time. Without a mechanism to focus and limit perceptual processing, the pilot's decision cycle can be dominated by input.

5.3 Controlling attention: Zoom lens

As in humans, the virtual pilot's attention can be controlled endogenously, a top-down, goal-driven form of control. To control attention endogenously, the pilot normally chooses to perceive only groups of objects, which is the lower acuity mode of perception that saves the processing. Groups provide cues such as location, size, and some details about the membership that can be used for searching for specific visual objects. The pilot may be interested in focusing on visual objects with a particular set of features. For example, one of the highest priorities in the pilot's visual search is to identify enemy air defense vehicles—this is driven by a goal of survival. Using its capability of situation awareness, the pilot can focus its attention to the areas where the enemy's air defense units most likely reside.

It has been hypothesized that attention acts like a zoom lens [6]. According to this metaphor, a zoom lens has a reciprocal relation between the magnification, or level of detail, and the size of the viewing field. When the zoom lens is set to a low level of magnification, the field of view is greater, albeit with a low level

of detail. Conversely, as the power of the lens increases, the amount of detail increases, while the field of view decreases.

Applying the zoom lens metaphor to the virtual pilot, the ability to perceive at different resolutions is provided by enabling the pilot to perceive groups when attention is set at a low resolution and to perceive individual entities at a high resolution. With an attention mechanism, the pilot exercises greater control over the amount of information that is processed, ameliorating the overload problems experienced previously. At the same time, the groups perceived at low acuity provide the right level of abstraction for tracking and reasoning about groups of entities. Furthermore, groups play an important role in controlling and shifting the focus of attention.

Using situation awareness, setting the zoom lens for shifting attention in virtual pilot is supported by the hierarchical tree structures in situation templates. When the zoom lens is set to the greatest resolution, information stored in the leaf nodes of a tree is used, when the lens is set to a low resolution, information in internal nodes is extracted. In other words, the zoom lens can be adjusted by using nodes at different levels of the tree.

5.4 Controlling attention: Spotlight

Attention may also be captured exogenously or involuntarily. This occurs when a feature reaches a threshold, such as when there is a flash of light from a muzzle blast. The sudden onset of this type of stimulus causes the visual object to automatically be placed in the working memory of the pilot's reasoning system. Once in working memory, the pilot's productions fire and react to the event. In this way, visual objects can capture attention by exogenous events. Other events that could possibly be used to capture attention in this way are motion, luminance, and the sudden appearance of smoke or dust.

Such triggering events and their thresholds are stored in situation templates. They are typically properties of internal nodes that represent regions or groups of entities.

5.5 Experimental evaluation

To measure the effectiveness of situation awareness, we quantitatively evaluated its utility in the context of focusing attention in the virtual pilot. One way to perform such an evaluation is to compare the perceptual loads on the pilot with and without attention. To gather data, the pilot was sent on the same mission (a standard deep attack) two times. The mission was such that for approximately 1500 decision cycles there were only a few vehicles in the pilot's visual field at any given time while the pilot traveled to a destination. Once at the destination, the pilot would periodically pop-up and peek over a hill at 120 other vehicles, then mask itself again behind the hill so that the vehicles were no longer in view. The amount of time spent in perceptual processing per decision cycle was measured over the course of the mission. The results are shown in Figure 7. Note that the average time spent on perceptual processing per decision cycle was approximately 15 milliseconds when attention was used, with peaks around 40 milliseconds. Without attention, the average perceptual processing time was around 50 milliseconds, with peaks above 500 milliseconds (not shown due to scale of graph). From experience, we found that peaks much above 100 milliseconds will sometimes result in helicopter crashes and that an average rate of 50 milliseconds puts a high load on the pilot.




Figure 7: Time (ms) for perceptual processing per decision cycle of two identical missions, one without perceptual attention (light gray) and the other with it (dark).

6. Related Work and Discussions

In applied behavioral science, the term situation awareness has emerged as a psychological concept similar to such terms as intelligence, vigilance, stress, workload and attention. [11] gives an excellent discussion and survey on several definitions and systems of situation awareness.

The existing situation awareness models can be classified as descriptive and prescriptive. Numerous descriptive situation awareness models have been proposed for modeling, mostly in military domains [11]. Descriptive models are suitable for capturing basic issues of decision making. They usually do not, however, support a quantitative simulation of the process by which cues are processed into perceptions, situations are assessed, and decisions are made. Among the proposed descriptive models, Endsley's three-layer model [5] is probably the most succinct. This model makes connections among situation awareness, perception, and reasoning and decision-making.

A few prescriptive models have been proposed, most of which use production rules [3,15]. In such rule-based systems, situation assessment is driven by production rules, in which a situation is assessed using the rule "if events E occur, then the situation is S". Such rule-based systems are easy to implement, but lack the diagnostic strength of a causal-reasoning type of approach. The reasoning processes in such systems are simply forward chaining, and normally do not use all information of the current situation.

Our model of situation awareness is a unique combination of the descriptive and prescriptive models. Our approach takes the advantage of the modeling power of the descriptive approach and the operational strength of the prescriptive approach, and integrates them in one system. The situation awareness module of our agent system architecture in Figure 2 follows Endsley's model [5]. Furthermore, the hierarchical tree structure and other static information, such as information of the members of a group, of a situation template can be viewed as a description of a situation. The actions and possible triggering events associated with situation templates, which are also in the form of production rules, can be considered as a prescription of a situation. Therefore, in our system, quantitative situation description, such as numbers

and types of entities, can be naturally accommodated, and operations, such as flying the helicopter above the obstacle in the front of visual field, can be implemented. Moreover, our system also supports some new features of situation awareness that have not been addressed in the other systems. Most importantly, our system supports controlling of attention and creates focused attention to reduce perceptual load

There is existing work related to creating rational and believable behavior in virtual humans. Rickel and Johnson [13] and Chopra and Badler [4] generated believable gaze (overt) behaviors in virtual humans. Although these systems appear to be successful in their application domains, the gaze behaviors of these systems do not affect what the agents actually perceive or are generated to service a cognitive need for information. Aasman [1] built a virtual car driver in a simulation world. It employs a cognitive model for controlling attention via gaze and head movements. Aasman's work is potentially useful in extending our pilot's ability to control gaze, but like the others, he does not address the issue of balancing perceptual and cognitive processing.

With respect to perceptual grouping, Flinchbaugh and Chandrasekaran [7] did some of the early work using Gestalt principles of proximity, similarity and motion. Numerous AI researchers have addressed the importance of attention in vision. For example, Tsotsos et al. [16] argued that attention acts to optimize the search in solving a vision problem, and they have implemented an artificial neural network to solve vision problems where attention selects spatial regions and features of interest. It is not clear whether they have employed their model in a dynamic environment yet, or how they would integrate it with goal-driven behavior. Reece and Shafer [12] employed a model of attention in their work on (simulated) robot driving. They showed that by making assumptions about world constancy, they could avoid re-sensing relatively static aspects of a scene, thus saving computation. Their work and other AI research on active perception [2] address many of the efficiency issues that we are concerned with in this research. In our future work, we will try to extend this body of research by integrating active perception with a cognitive architecture like Soar [14].

7. Conclusions and Future Work

Situation awareness is an important component of virtual humans, and plays fundamental roles in creating believable, human-like characteristics, such as rationality, personality, and emotion. In this paper, we proposed to use situation awareness as a working interface between perception and cognition. We developed an architecture that supports the integration of perception, situation awareness, and reasoning and decision making. We developed a representation for situations, which encodes the organizational and spatial relationships of visual entities. The new representation combines the existing descriptive and prescriptive representations. We also developed pattern-matching methods to find the most possible situation based on perceptual inputs.

We applied situation awareness to virtual pilots in synthetic environment. The virtual pilot uses its capability of situation awareness to decide where to search for additional information in dynamic environments and to create focused attention to reduce perceptual overload. Our experiments show that situation awareness has a great impact on reducing perceptual overload and allocating the virtual pilot's limited computation resources. The

results show that by using situation awareness, the virtual pilot's response time is reduced by more than 50 percent.

In our future work, we will study the utility of situation awareness in reasoning and decision making.

8. Acknowledgments

The authors were supported in part by NSF under grant IRI-9619554, by DARPA via subcontract L74158 with the University of Michigan under prime contract N61339-97-K-008, and the ONR under contract N00014-98-1-0132.

9. References

- [1] J. Aasman. "Modeling Driver Behaviour in Soar," Ph.D. Thesis, Rijksuniversiteit Groningen, Leiden, 1995.
- [2] D. Ballard. "Animate Vision," *Artificial Intelligence* 48 (1991), pp.57-86.
- [3] S. Baron, G. Zacharias, R. Muralidharan, and R. Lancraft. "PROCURU: A model for analyzing flight crew procedures in approach to landing," *Proc. 8th IFAC World Congress*. 1980.
- [4] S. Chopra and N. Badler. "Where to Look? Automating Visual Attending Behaviors of Virtual Human Characters," *Proc. 3rd Intern. Conf. Autonomous Agents*. May 1-5, 1999.
- [5] M. Endsley. "Toward a theory of situation awareness in dynamic systems," *Human Factors*. 37 (1995) pp.32-64.
- [6] C. Eriksen and Y. Yeh. "Allocation of Attention in the Visual Field," *Journal of Experimental Psychology, Human Perception and Performance*, 11(1985) pp.583-597.
- [7] B. Flinchbaugh and B. Chandrasekaran. "A theory of spatial-temporal aggregation for vision," *Artificial Intelligence*, 17 (1981), pp387-407.
- [8] J.H. Friedman, J.L. Bentley, and R.A. Finkel. "An algorithm for finding best matches in logarithmic expected time," *ACM Trans. Mathematical Software*. 3 (1977) pp.209-226.
- [9] R. Hill, Jr. "Modeling Perceptual Attention in Virtual Humans", *Proc. 8th Conf. Computer Generated Forces and Behavioral Representation, May 1999*.
- [10] U. Neisser. *Cognitive Psychology*. NY: Appleton-Century-Crofts, 1967.
- [11] R. W. Pew and A. S. Mavor, editors. "Situation awareness," in, *Modeling Human and Organizational Behavior: Application to Military Simulations*. National Academy Press, Wash., D.C, 1998.
- [12] D. Reece and S. Shafer. "Control of Perceptual Attention in Robot Driving," *Artificial Intelligence*, 78 (1995) 461-505.
- [13] J. Rickel and W. Lewis Johnson. "Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control." To appear in *Applied Artificial Intelligence*.
- [14] P. Rosenbloom, J. Laird, A. Newell, editors, *The Soar Papers: Research on Integrated Intelligence*. MIT Press, Cambridge, MA, 1993.
- [15] M. Tambe, et al., "Intelligent agents for interacted simulation environments," *AI Magazine*, Spring (1995) pp. 15-37.
- [16] J. Tsotsos, and et al., "Modeling visual attention via selective tuning," *Artificial Intelligence* 78 (1995) 507-545.