Virtual Human Animation Based on Movement Observation and Cognitive Behavior Models

Norman I. Badler, Diane Chi, Sonu Chopra
Center for Human Modeling and Simulation
University of Pennsylvania
Philadelphia, PA 19104-6389

Abstract

Automatically animating virtual humans with actions that reflect real human motions is still a challenge. We present a framework for animation that is based on utilizing empirical and validated data from movement observation and cognitive psychology. To illustrate these, we demonstrate a mapping from Effort motion factors onto expressive arm movements, and from cognitive data to autonomous attention behaviors. We conclude with a discussion on the implications of this approach for the future of real-time virtual human animation.

1 Introduction

Automatically animating virtual humans with actions that reflect real human motions is still a challenge. Skilled animators are able to create effective and compelling human characters through labor-intensive scripting of every subtlety of motion. For example, in producing the movie *A Bug's Life*, Pixar animators manipulate huge spreadsheets of animation parameters [41]. As we expect more human-like *real-time* behaviors, we cannot afford the temporal indulgence and expert skill required by off-line animation. Even if we are provided with real-time movement modules that can animate walking or gesturing in real-time, goals and parameters for them must be provided. This contextual information comes from at least three sources:

- the external world,
- the objective being undertaken by the character, and
- the internal model of what it means to be human.

In order to address the joint problems of enhancing "human-ness" while insisting on real-time synthesis, a new approach is needed. In this paper we explore the distinction between the external and internal motivators that ultimately create more believable movements.

We present a framework for animation that is based on utilizing empirical and validated data from movement observation and cognitive psychology. To illustrate these, we demonstrate a mapping from Laban's *Effort* motion factors [11, 20] onto expressive arm motions, and from cognitive data to autonomous attention behaviors. These two approaches to animation exemplify our philosophy:

- Human animation requires an understanding of empirical data on human movement and cognitive properties.
- Human movement cannot simply be captured and replayed, but must be parameterized and re-usable in various contexts.
- Human reactions to the world must be combined with task objectives.
- Human sensing provides individualized models of the local context which can then be combined with internal models of personality and emotion to produce appropriate actions and reactions.

These desiderata have some consequences:

- Formal approaches to animation such as genetic algorithms [26, 43, 32], are not likely to converge to natural human motions without extraordinary care to link the fitness or evaluation functions to known human motions, thus obviating this approach.
- Reactive approaches for navigation and pursuit [27, 31, 39] do not readily accommodate human task objectives, sensing costs, and cognitive principles.
- Humans exhibit a wide variety of expressive actions which reflect their personalities, emotions, and communicative needs [6, 8, 40]. These variations often influence the performance of simpler gestural or facial movements.

In order to focus this presentation, we will restrict our scope to two problem areas among the several that we are currently investigating: expressing the qualitative aspects of movement and controlling visual attention. For Effort, we use the collected empirical wisdom and descriptive taxonomy of Laban Movement Analysis to create movement models. This parameterization not only yields a convenient user interface that covers a wide range of human expressions, but also appears amenable to driving animation from a character's inner drives and temperaments, emotions, or personalities. For attention and motions dependent on attention such as reaching and walking, sensing events and interleaving consequent actions is crucial. Attention is a reaction to what is sensed; action selection is mediated by cognitive models.

2 Animating Qualitative Aspects of Movement

A system that allows users to customize basic movements based on a character's personality, mood, and attitudes is the first step towards simplifying the development of a repertoire of characters with a wide range of expressiveness. By selecting a general human movement description language to customize motions, such a system can also lead to the generation of virtual characters from different cultures. Although certain gestures are culture-specific, the descriptions of the movements of individuals with the same emotions and intent are often similar. A playful child moves with free, indirect abandon; an aggressor makes strong, direct, and sudden attacks; and a soldier marches in bound, sustained strides. We present a technique where the user customizes movements through a qualitative description – the Effort component of Laban Movement Analysis. This is the first step towards a system where a user creates character by specifying attitudes and intentions, which in turn may eventually lead to the automatic generation of appropriate movements from speech text, a storyboard script, or a behavioral simulation.

2.1 Background: Laban Movement Analysis

Rudolf Laban (1879-1958) made significant contributions to the study of movement, observing human movements in dance, combat, performing industrial work tasks, conducting cultural rituals, and just in every day situations. His theories on movement and its extension by his students and colleagues have resulted in a rich vocabulary for describing and analyzing movement, leading to the development of Laban Movement Analysis (LMA) [5, 24, 11, 23], which is promoted by the Laban/Bartenieff Institute of Movement Studies (LIMS)[1]. LMA has evolved into a comprehensive notation system being used in dance, nonverbal research, psychoanalysis, dance therapy, and many movement-related fields.

Laban Movement Analysis has four major components – Body, Space, Shape, and Effort. We only use the Effort component of Laban Movement Analysis for the purposes of controlling animation. Effort describes the qualitative aspects of movement. Effort comprises four motion factors: Space, Weight, Time, and Flow. Each motion factor ranges between *indulging* in and *fighting* against the quality. These extremes are seen as basic, "irreducible" qualities, meaning they are the smallest units of change in an observed movement. Fig. 1 shows Laban's Effort Graph of the Effort elements – the extremes for each motion factor. The dotted line represents the separation of the indulging and fighting qualities.

The Space motion factor (sometimes referred to as Focus) refers to one's attention to the surroundings. Space dynamics vary from Indirect to Direct. Indirect is seen as flexible, meandering, wandering, and with multi-focus attention. Direct motion has a single focus and is channeled and undeviating. Indirect movements include waving away bugs, slashing through plant growth to get through the woods, surveying a crowd of people, and searching a room for misplaced keys. Examples of Direct movements include pointing to a particular spot, threading a needle, and describing the exact outline of an object.

The Weight motion factor (also referred to as Pressure) describes one's attitude towards gravity and the impact of one's movement. Weight dynamics range from

[1]Laban/Bartenieff Institute of Movement Studies, 11 East 4th Street, 3F, New York, NY 10003; http://www.artsnet.org/camt/lbi1/laban.htm

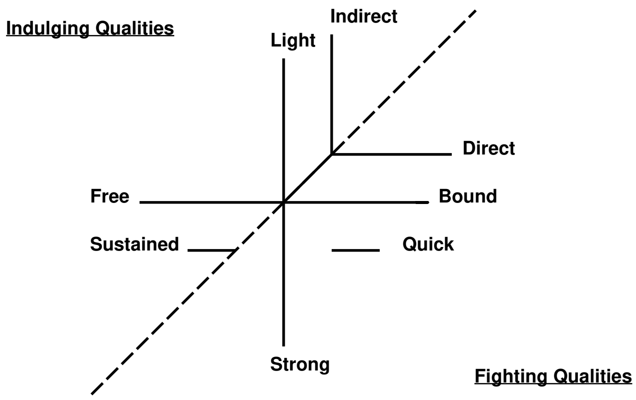

Figure 1: Graph of Effort Elements

Light to Strong. Light movements are buoyant and delicate, easily overcoming gravity and marked by decreasing pressure. Strong movements are powerful, have an impact, and involve throwing one's weight into the movement. Light motions include adding dabs of paint to a canvas, pulling a splinter out of one's finger, and describing the movement of a feather. Strong motions include punching, pushing a heavy object, wringing a towel, and expressing a firmly held opinion.

The Time motion factor involves one's lack of or sense of urgency. Sustained movements are lingering, leisurely, and indulging in time; while Sudden movements are hurried and urgent. The stretch of a yawn and the stroking of a pet can be sustained. Swatting a fly, lunging to catch a ball, and grabbing a child from the path of danger are Sudden movements.

The Flow motion factor describes the nature of control and bodily tension. Human movement displays a continuous variation in flow changes, as well as expressive fluidity. Free movements are uncontrolled and abandoned–the individual is unable to stop in the middle of the movement. Examples include waving wildly, shaking off water, and flinging a rock into a pond. Bound movements are controlled and restrained. Slow motion is Bound (and notably NOT Sustained) because the mover can stop at any moment, and the fluidness is constrained. Holding back tears and carrying a cup of hot liquid are Bound.

2.2 Using Effort for Animation

Effort descriptions of Laban Movement Analysis provide an adequate interface for controlling the expressiveness of computer animated figures, as originally proposed (but not implemented) by Badler [4]. Effort proves intuitive by providing a small number of textual descriptors as compared to the detailed, cumbersome notations or mathematical and physics-based parameters for describing movements. Also, Effort covers the range of human expressive movements since it is derived from extensive observation practice. Most importantly, Efforts provide a small number of descriptive parameters that can in turn control numerous low-level motion degrees-of-freedom.

We describe a mapping of Effort elements into low-level movement parameters derived from empirical methods. These methods include: visual analysis of the playback of the motion capture data and video of a Certified Movement Analyst[2] (CMA) trained in Laban Movement Analysis performing Effort combinations, descriptions of Effort from the literature [24, 11, 5, 23], application of traditional animation principles [22, 35], and much experimentation with feedback from our CMA.

Animations can be defined using a sequence of end-effector locations and settings for the Effort motion factors. The sequence of end-effector locations define a general movement sequence, while the Effort settings describe the desired qualitative nature of the movement. Using inverse kinematics, we compute postures for an articulated figure by determining the position and orientation of the limbs from the specification of end-effector locations. Since inverse kinematics does not specify how a figure achieves a computed posture or changes between a series of postures, we are able to modify the in-between postures. Rather than defaulting to a standard interpolation method, we use the Effort settings to determine how to generate the in-betweens. The Effort elements map into a set of low-level movement parameters which includes velocity changes, anticipation, overshoot, interpolation method, end-effector path, duration, wrist bend, arm twist, limb volume, and torso support.

Several of the low-level movement parameters are used as input into the interpolation process. We interpolate using a variation of the double interpolant method introduced by Steketee and Badler [33]. They use both kinetic and position interpolants to separate kinetic control from parametric (position) control. The kinetic interpolant specifies the relationship between time and keyframes, while the position interpolant defines the relationship between keyframes and position or some other motion parameter. Rather than requiring the user to specify the times of the keypoints to generate a kinetic interpolant, we define a parameterized keyframe to time function, where keyframe number and time are normalized between 0 and 1. For each frame, we compute time by normalizing the number of frames between keypoints to 1 and determining where the given frame lies on the normalized scale. We then use this time to compute the position values (or joint angle values for angular interpolation) for the given frame. Further, while the original method was applied to a whole movement (a series of keyframes), we specify interpolants between each keypoint. By parameterizing the position interpolant, we achieve varying velocity changes, anticipation, overshoot, and end-effector path. Strong and Sudden movements are accelerating, while Light and Sustained movements are decelerating. Strong movements begin with anticipation, while Light and Free movements overshoot their goals. Direct motions follow a straighter path between keypoints, while Indirect motions follow a more curved path. We also allow three types of interpolation: end-effector position, joint angle, and elbow position. The default interpolation method is end-effector position interpolation. Joint

[2] Janis Pforsich

angle interpolation is used for Free movements, while elbow position interpolation is used for Indirect motions.

A duration parameter changes the number of frames between keypoints. Sudden movements shorten the amount of time between keypoints, while Sustained and Bound movements lengthen the time. Wrist bend specifies the changes in the wrist throughout a movement. Bound and Strong movements keep a stiff, slightly bent wrist. Light, Free, and Indirect movements display a significant amount of wrist bends. Indirect movements also displays elbow and wrist twists. Changes in the volume of the upper arm are used to simulate biceps contraction in Strong movements. Squash and stretch of the figure's torso is used to simulate the expansion and contraction of the chest cavity.

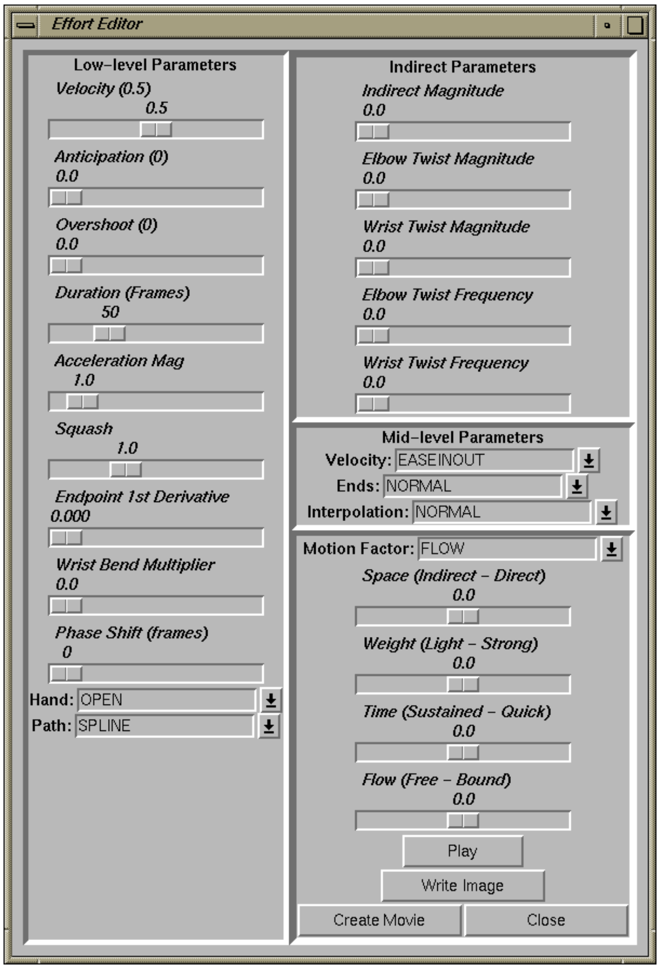

Figure 2: Effort Editor

We implement EMOTE (Expressive MOtion Engine), a character animation system

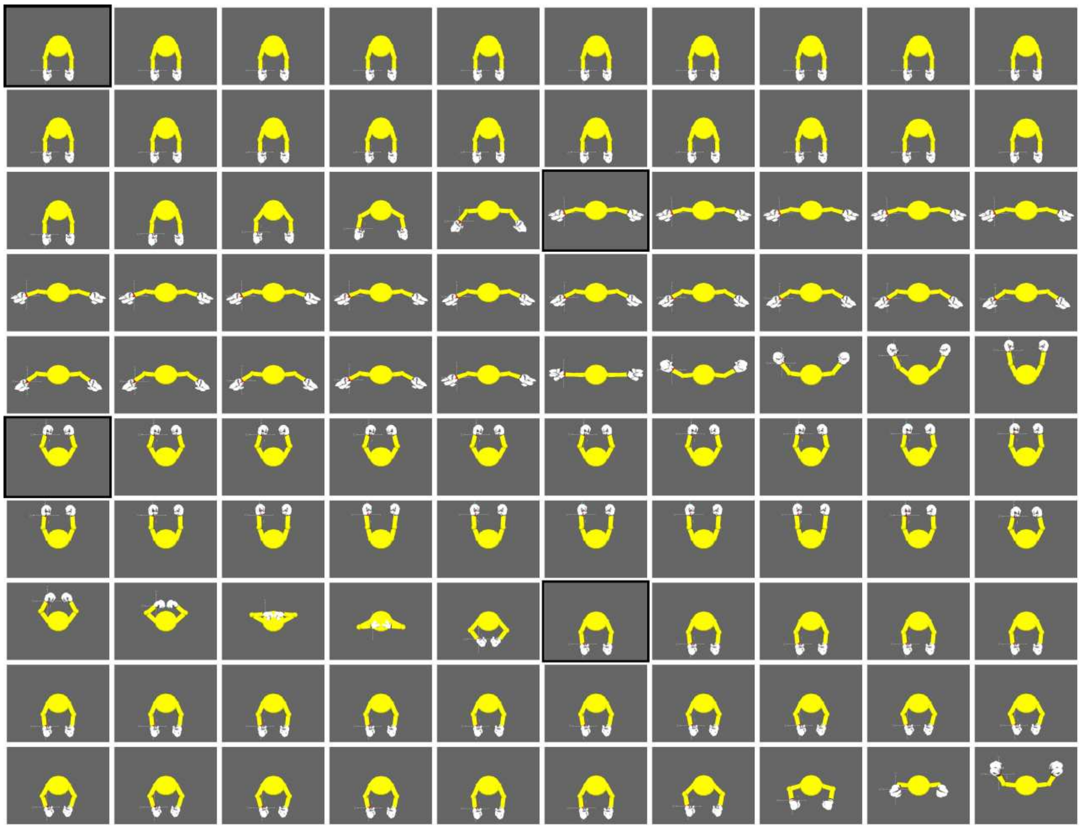

Figure 3: Frames from Strong Effort Setting (ordered left to right, top to bottom)

that utilizes these principles [9]. EMOTE captures a wide range of expressive movements, provides an easy-to-use interface, and features interactive editing. EMOTE offers mid-level control using Effort's textual descriptors, as well as more fine control using sliders for the low-level movement parameters (Fig. 2). Fig. 3 illustrates an example animation for a Strong movement derived from five keypoints, which are outlined in black.

Effort provides a systematic method for specifying expression. While an animator could conceivably control the individual lowest level degrees-of-freedom independently, the advantage of the Efforts is in organizing the whole set around four meaningful high-level scales. Varying the intensities of the Effort elements and combining different elements produce a rich language for expressive movements.

3 Animating Attention

Synthetic humans should exhibit the appropriate looking or attending behaviors relevant to the activities they are engaged in. Since gaze is significant in communication and behavioral representation, random or uncontrolled looking behavior is both misleading and disconcerting. Characters for which motion alone is animated while gaze remains fixed appear robot-like or mechanical. Also, *real-time* performance is necessary for interactive simulation, pedagogy, games and 3D chat. Avatars in cyber communities should respond to events such as someone entering a virtual chat room or responding to objects in their path just as real participants might do (they will either acknowledge the presence of an individual, alter their motion to avoid an obstacle and sometimes fail to notice an event because their attention is otherwise engaged). The STEVE system from the Information Sciences Institute makes good use of a pedagogical agent with attention capabilities [16]. Other simple attention models have been used for conversational agents [8, 36].

Given a high level script (possibly interactive) that an agent should follow, how do we animate details of the script with the appropriate behavior? The mapping between motor tasks and the corresponding motion is well understood, but attending behavior is often not specified and is *emergent*: where an agent looks changes due to interactions between simultaneous tasks and in response to the dynamics of the environment. Further, motor actions may be modified by input from the attentional system (e.g., if an agent notices an object bearing down him, he will step out of the way).

We implement a psychologically-motivated framework, called *AVA* (Automated Visual Attending) [10], in which simultaneously active gaze behaviors compete to direct an animated agent's line of sight or visual attention. Deliberate behaviors, known in the cognitive psychology literature as endogenous [44], are associated with broad categories of motor and cognitive activity. Involuntary behavior, indicating attentional capture by exogenous effects [17, 12], is implemented in AVA as a peripheral motion sensor. Since humans are not constantly monitoring task related objects (for example, an agent does not continuously look at his destination or the ground in front of him while walking to a goal), we associate memory uncertainty thresholds with certain classes of monitoring activities, of which locomotion is a specific example. Such activities require glancing at specific objects or locations when the memory uncertainty limit for a task object's state is reached [25]. In the absence of demands ¿from deliberate or exogenous effects, we implement a type of idling behavior known as spontaneous looking [18].

When engaged in several simultaneous tasks or in the presence of exogenous distractors, an agent's performance will degrade by increasing response time to deliberate activity targets [13, 29]. AVA recognizes increasing cognitive load (since the number and identity of objects that require attention are maintained) and animates time to complete eye movements accordingly.

Since agent animation in 3D chat or networked virtual games is of necessity unscripted and immediate, our system builds behaviors that query the graphics scene graph whenever possible. In order to sense peripheral motion, for example, rather than use a vision processing or optic flow filter, we instead determine those objects that fall into a character's peripheral field of view and which have moved interframe (we can do this by determining the location of figures that are visible to an agent, fall into the periphery and then query the location of such figures in successive frames). In the absence of any task information, we employ a simplified image processing technique modified from vision approaches [19, 38].

3.1 Arbitration Mechanism

An arbitrating process called the Gazenet (Fig. 4) determines where an agent looks by selecting from the three levels of behavior: deliberate, exogenous and idling. Two queues are maintained: an **IntentionList** that stores sites or objects that need to be attended due to the demands of current activities and a **Plist** that indicates objects in agent's peripheral field of view that are moving. When both queues are empty, a spontaneous looking or idling behavior is activated. Behaviors are implemented in our technique as parallel, executing finite state machines [37].

Fig. 5 illustrates *AVA's* architecture. Users enter task requests as text input. A task manager process for each agent consumes such requests and generates the appropriate eye gaze or looking behaviors for an action (some activities such as walking and monitoring may be requested in parallel). The motions which correspond to motor tasks are also generated. When the memory uncertainty threshold for a monitoring activity is reached, the corresponding eye behavior adds relevant sites to an **Intentionlist**. (E.g., the locomotion eye behavior will add the goal destination or ground at particular intervals indicating that those locations should be attended). A peripheral motion sensor behavior is active for each agent and updates the **Plist** as needed.

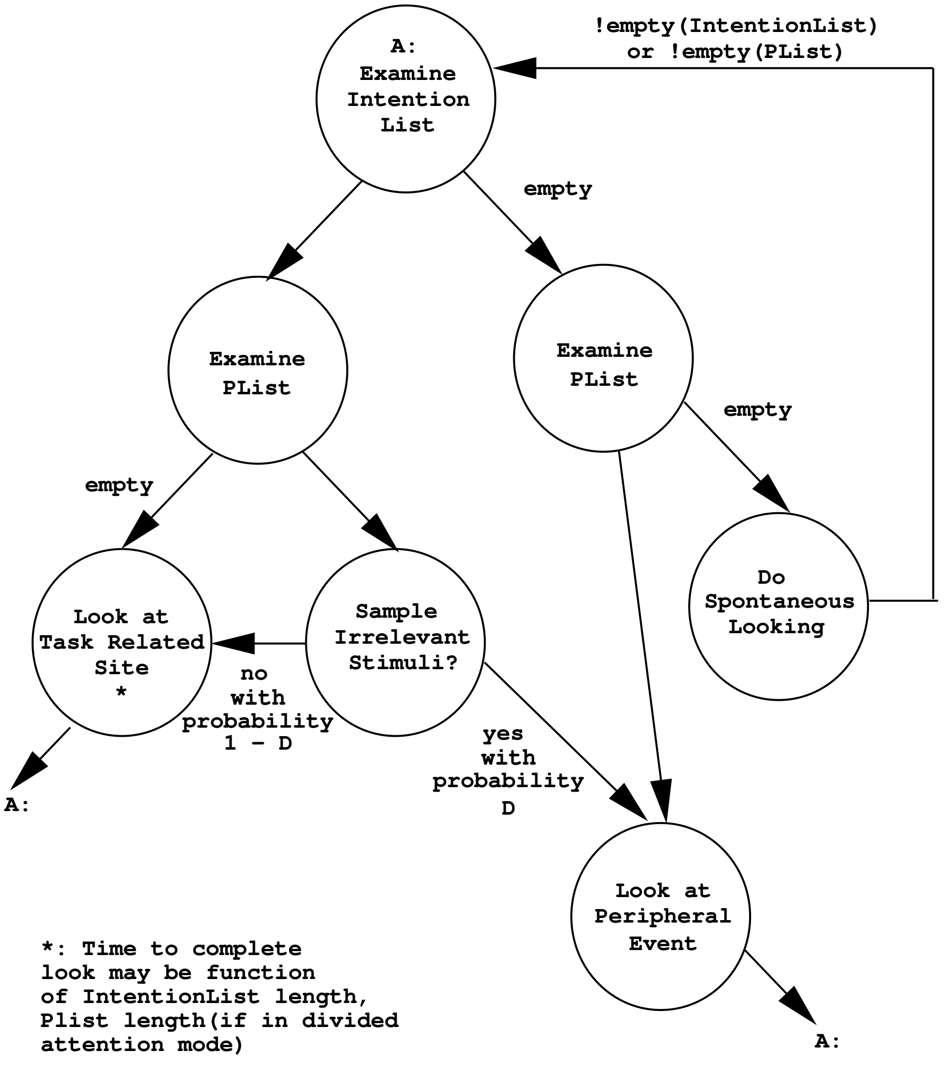

Figure 4: GazeNet

3.2 Deliberate Behaviors

The following behaviors, reflecting endogenous activity, are implemented in AVA and add objects, locations or relative angles to **IntentionList**.

1. Visual Search – We model visual search by first determining the angle between the center of fixation and the target. We generate a sequence of intermediate positions (visual angles) that move the eye ¿from its current position to the target location. If the target is not present in the environment or occluded by another object, a sweep of the field of view is performed [30].

2. Visual Tracking – An agent periodically (with high uncertainty) glances at the tracked target. If the target passes behind an occluding figure, AVA calculates the estimated reappearance location (based on the last perceived target heading).

3. Monitoring Behaviors Monitoring tasks (locomotion being a general case) use uncertainty thresholds [25] that relate how often a signal, event, or goal should be glanced at in order to maintain an accurate view of the signal's state in memory. When the uncertainty threshold for a given monitoring task is reached in AVA, the relevant site is added to **IntentionList**.

   While walking, for example, an agent in under AVA control looks toward the horizon or destination and occasionally glances at the ground [34] . This is an example of a monitoring task with low uncertainty thresholds.

4. Limit Monitoring Monitoring may also be associated with limit conditions [25]. As a signal's state approaches a critical or cautionary level, it will occasion more frequent eye fixations. For example, when crossing the road, an agent will more frequently glance at the light or crossing signal if it is yellow rather than green.

5. Reach and Grasp The eye is supposed to establish targeting for the hand [1]. When initiating a reach and grasp motion, we generate eye movement toward the relevant grasp site by adding it to **IntentionList**.

3.3 Involuntary Behaviors

A motion sensor net continuously checks the environment for objects that change position between frames. Such objects are added to **PList** if they fall within an agent's periphery. If a moving object is noticed by the agent and its velocity and heading indicate a likely collision, deliberate visual tracking is performed and locomotion is updated or activated (e.g., to alter speed or heading). A *distractability* parameter is set in the simulation indicating likelihood of noticing peripheral motion while consumed in deliberate activity. Such a parameter tends to be age dependent and may also reflect impairments in oculomotor control [15, 21].

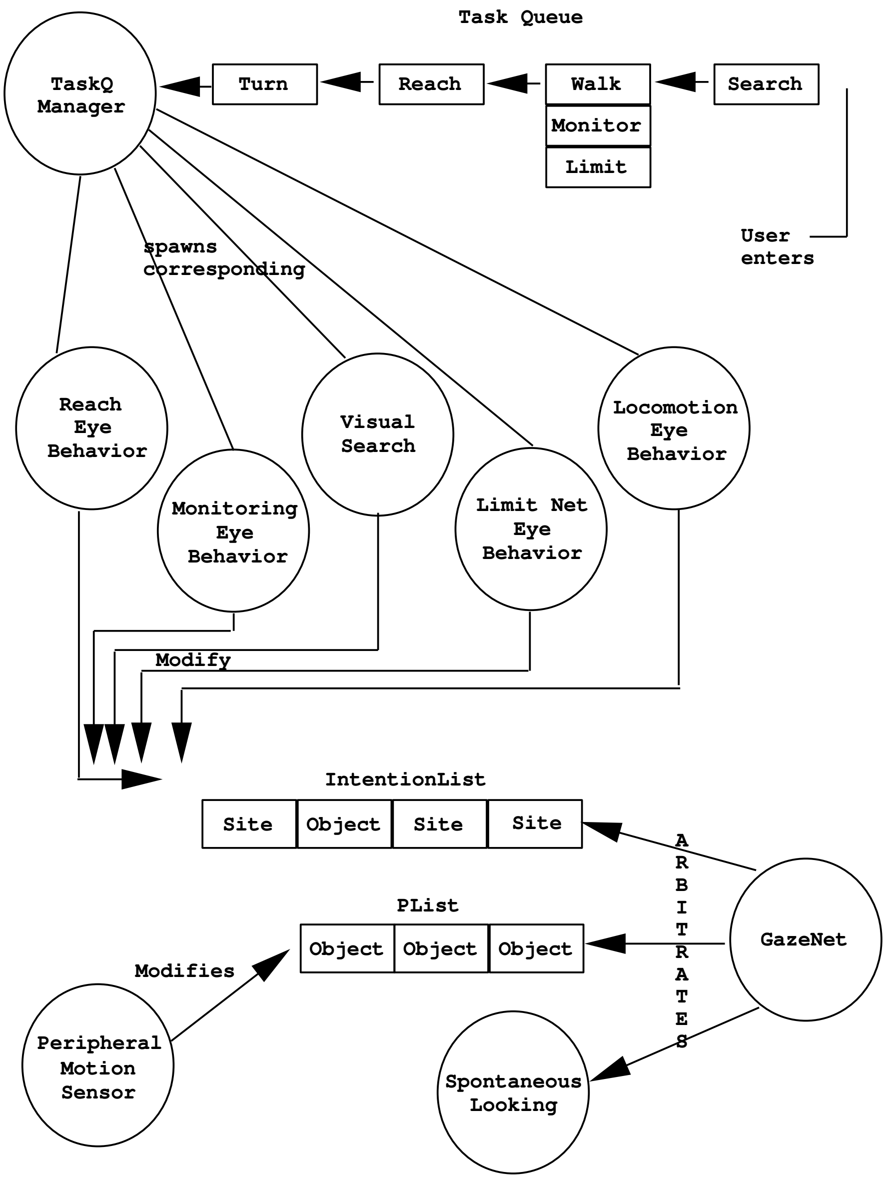

Figure 5: Method Architecture

3.4 Idling Behavior

In the absence of a specific goal or task, attention follows patterns of spontaneous looking [18]. Attention is drawn to items in the environment that are likely to be informative or significant. Psychologists argue this is due to a need to reduce uncertainty about our surroundings.

Since we wish to generate real-time eye behavior, we use a simplified novelty measure. AVA copies a 1000x400 snapshot of the agent's field of view into a pixel buffer. We select those pixels whose color values are the furthest from their neighbors in RGB space. We convert the location of these pixels back into 3D world by inverting and applying the graphics pipeline rendering transforms.

Figure 6 shows rays intersecting those locations in an agent's field of view that are the most locally conspicuous.

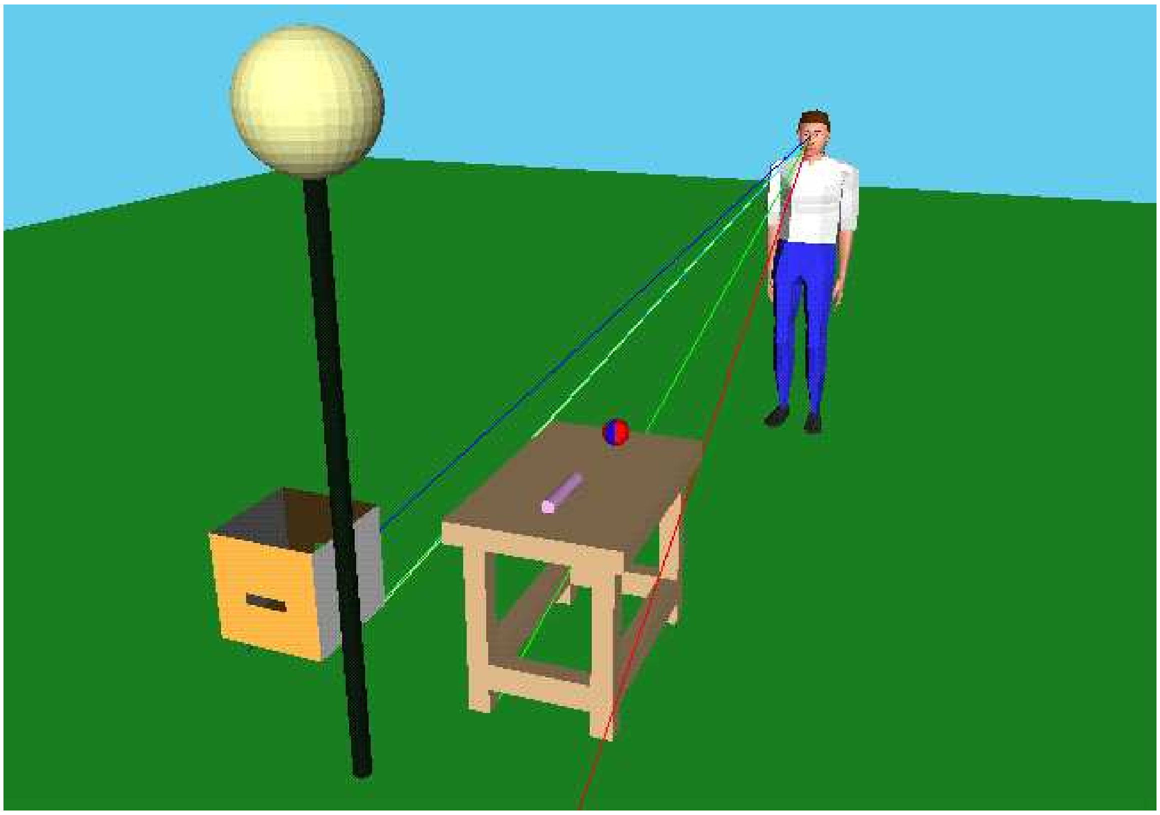

Figure 6: Spontaneous Looking - Rays Intersect Features with Local Contrast

4 Discussion

We have seen that Effort, as defined by Laban Movement Analysis, provides a systematic method for specifying expression. Further, varying the intensities of the Effort elements and combining different elements produce a rich language for expressive movements that cleanly and robustly fits into a key pose framework. The resulting lower dimensional parametric controls are not only "intuitive": they are validated by decades of movement analysis and refinement that provides evidence for both movement coverage and descriptive power. The next step is connecting these Efforts to a representation of the character's inner emotion and personality states.

With the attention control, we see that tasks and the external world impose both voluntary and involuntary patterns of eye movements. As several tasks are simultaneously attempted, performance degrades. Peripheral events capture attention when the agent is engaged in a task which requires diffuse attentiveness (e.g., visual search or divided attention). In the absence of tasks or peripheral stimuli, attention follows patterns of spontaneous looking. All of these effects are mediated by human cognitive principles and empirical models. Their internal parameters are psychologically meaningful, they are accessible and may be "tweaked" for an individual, and they automatically produce human-like behaviors.

In other work not described here, our animation philosophy has been applied to face animation [28] and is being applied to human locomotion [14]. To extend this philosophy to complete human motion, we are examining parameterized motion capture [7].

One of our long-term objectives is producing real-time animations from natural language instructions [3, 42]. We are modeling the scope and richness of action expression with a *Parameterized Action Representation* (PAR) [2]. Among the semantic features of a PAR are the agent who will carry out the action and the manner in which it is done. Different agents will execute a PAR in different ways according to their capabilities, personality, and knowledge of the current situation. Attention plays a major role in initiating and sequencing the lower level movements of the agent – movements that do not appear explicitly in the instructions. Likewise, manner reflects the attitudes and personality of the agent and its approach to the task. A link between manner expressions and Effort elements is the next step in creating individualized agents.

These examples are intended to strengthen the case for the particular philosophy or paradigm of "human-ness" in animation. Both Efforts and attention are functioning, real-time, human character animation tools. As more components become integrated into the human model, we expect the emergence of a new level of animation control. An agent will have a personality and varying emotional states. Elucidating the connections between these states and the external Efforts and attentional choices of the agent is the next step. There is evidence in the scientific literature that these connections have been studied and may therefore be modeled.

References

[1] R.A. Abrams, D.E.M. Meyer, and S. Kornblum. Eye-hand coordination: Oculomotor control in rapid aimed limb movements. *Journal of Experimental Psychology: Human Perception and Performance*, 16:248–267, 1990.

[2] N. Badler, R. Bindiganavale, J. Bourne, M. Palmer, J. Shi, and W. Schuler. A parameterized action representation for virtual human agents. In *Workshop on Embodied Conversational Characters*, Lake Tahoe, CA, Oct. 1998.

[3] N. Badler, C. Phillips, and B. Webber. *Simulating Humans: Computer Graphics Animation and Control.* Oxford University Press, New York, NY, 1993.

[4] Norman I. Badler. A computational alternative to effort notation. In Judith A. Gray, editor, *Dance Technology: Current Applications and Future Trends.* National Dance Association, Reston, VA, 1989.

[5] Irmgard Bartenieff and Dori Lewis. *Body Movement: Coping with the Environment.* Gordon and Breach Science Publishers, New York, 1980.

[6] Becheiraz and Daniel Thalmann. A behavioral animation system for autonomous actors. In *Workshop on Embodied Conversational Characters*, October 1998.

[7] R. Bindiganavale and N. Badler. Motion abstraction and mapping with spatial constraints. In *Workshop on Motion Capture Technology*, Geneva, Switzerland, November 1998.

[8] J. Cassell, C. Pelachaud, N. Badler, M. Steedman, B. Achorn, W. Becket, B. Douville, S. Prevost, and M. Stone. Animated conversation: Rule-based generation of facial expression, gesture and spoken intonation for multiple conversational agents. In *Computer Graphics, Annual Conf. Series*, pages 413–420. ACM, 1994.

[9] Diane M. Chi. *A Motion Control Scheme for Animating Expressive Figure Movements.* PhD thesis, University of Pennsylvania, 1999. In progress.

[10] Sonu Chopra. *Where to Look? Automating Certain Visual Attending Behaviors of Human Characters.* PhD thesis, University of Pennsylvania, 1999. In progress.

[11] Cecily Dell. *A Primer for Movement Description: Using Effort-shape and Supplementary Concepts.* Dance Notation Bureau, Inc., New York, 1970.

[12] A. Hillstrom and S. Yantis. Visual motion and attentional capture. *Perception and Psychophysics*, 55(4):399–411, 1994.

[13] W. Hirst. The psychology of attention. In *Mind and Brain: Dialogues in Cognitive Neuroscience*, pages 105–141, 1986.

[14] H.Sun, 1998. Univ. of Pennsylvania. Personal communication.

[15] M. Johnson. Visual attention and the control of eye movements in early infancy. *Attention and Performance*, 15:291–310, 1994.

[16] W. L. Johnson and J. Rickel. Steve: An animated pedagogical agent for procedural training in virtual environments. *SIGART Bulletin*, 8(1-4):16–21, 1997.

[17] J. Jonides. Voluntary versus automatic control over the mind's eye movement. *Attention and Performance*, 9:187–203, 1981.

[18] D. Kahneman. *Attention and Effort.* Prentice-Hall, 1973.

[19] C. Koch and S. Ullman. Shifts in selective visual attention: Toward the underlying neural circuitry. *Human Neurobiology*, 4:219–227, 1985.

[20] Rudolf Laban and F. C. Lawrence. *Effort: Economy in Body Movement.* Plays, Inc., Boston, 1974.

[21] E. Ladavas, G. Zeloni, G. Zaccara, and P. Gangeni. Eye movements and orienting of attention in patients with visual neglect. *Journal of Cognitive Neuroscience*, 9(1):67–75, 1997.

[22] John Lasseter. Principles of traditional animation applied to 3D computer animation. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 35–44, July 1987.

[23] Vera Maletic. *Body, space, expression: The development of Rudolf Laban's movement and dance concepts.* Mouton de Gruyte, New York, 1987.

[24] Carol-Lynne Moore and Kaoru Yamamoto. *Beyond Words: Movement Observation and Analysis.* Gordon and Breach Science Publishers, New York, 1988.

[25] N. Moray. Designing for attention. In *Attention, Selection, Awareness, and Control: A Tribute to Donald Broadbent*, pages 53–72. Clarendon Press, Oxford, 1993.

[26] J. T. Ngo and J. Marks. Spacetime constraints revisited. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 343–350, August 1993.

[27] H. Noser, O. Renault, D. Thalmann, and N. Magnenat-Thalmann. Navigation for digital actors based on synthetic vision, memory and learning. *Computers and Graphics*, 1995.

[28] C. Pelachaud, N. Badler, and M. Steedman. Generating facial expressions for speech. *Cognitive Science*, 20(1):1–46, 1996.

[29] M. I. Posner, R. D. Rafal, L.S. Choate, and J. Vaughan. Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, 2:211–228, 1985.

[30] P. Rabbit. Control of attention in visual search. In *Varieties of Attention*, Series in Cognition and Perception, pages 273–288. Academic Press, 1983.

[31] C. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.

[32] K. Sims. Evolving virtual creatures. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics

Proceedings, Annual Conference Series, pages 15–22. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.

[33] Scott Steketee and Norm Badler. Parametric keyframe interpolation incorporating kinetic adjustment and phasing control. In B. A. Barsky, editor, *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 255–262, July 1985.

[34] M. Swain and M. Stricker. Promising directions in active vision. *International Journal of Computer Vision*, 11:109–126, 1993.

[35] Frank Thomas and Ollie Johnston. *Illusion of Life: Disney Animation*. Hyperion, New York, 1995.

[36] K. Thorisson. Real-time decision making in multimodel face-to-face communication. In *Proceedings of the Second International Conference on Autonomous Agents*, pages 16–23. ACM, May 1998.

[37] T. Trias, S. Chopra, B. Reich, M. Moore, N. Badler, B. Webber, and C. Geib. Decision networks for integrating the behaviors of virtual agents and avatars. In *Proc. IEEE Virtual Reality Annual International Symposium*, pages 156–162, 1996.

[38] J.K. Tsotsos, S.M. Culhane, W.Y.K. Wai, Y. Lai, and F. Nufflo. Modeling visual attention via selective tuning. *Artificial Intelligence*, 78:507–545, 1995.

[39] X. Tu and D. Terzopoulos. Artificial fishes: Physics, locomotion, perception, behavior. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 43–50. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.

[40] H. H. Vilhjalmsson and J. Cassell. Bodychat: Autonomous communicative behaviors in avatars. In *Proceedings of the Second International Conference on Autonomous Agents*, pages 269–277. ACM, May 1998.

[41] G. Walters, 1998. Pixar. Personal communication.

[42] B. Webber, N. Badler, B. Di Eugenio, C. Geib, L. Levison, and M. Moore. Instructions, intentions and expectations. *Artificial Intelligence J.*, 73:253–269, 1995.

[43] D. Wiley and J. Hahn. Interpolation synthesis of articulated figure motion. *IEEE Computer Graphic and Applications*, November 1997.

[44] S. Yantis. Stimulus-driven attentional capture and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, 19(3):676–681, 1993.