Chapter 7

Conclusion

The fundamental challenges in computer graphics and animation lie primarily in having to deal with complex geometric, kinematic, and physical models. The development of better software tools has advanced the state of the art in the modeling and rendering of these models. However, software tools for the automatic generation of motion are still relatively scarce. In particular, techniques are needed to animate both autonomous and user-controlled human figures naturally and realistically in response to high-level task commands.

The preceding chapters of this thesis have presented a research framework aimed at facilitating the high-level control of animated characters in interactive virtual environments. The next section briefly summarizes the key ideas contained in this thesis, and the following section contains a concluding discussion.

Summary

We approach the problem of automatically synthesizing motions for animated characters from the standpoint of modelling and controlling a "virtual robot" (an *autonomous animated agent*). We propose a general architecture for autonomous animated agents based on *planning*, *sensing*, and *control*. To test the viability of this approach, we have developed techniques for automatically generating the gross body motions for animated human figures given high-level navigation or manipulation task commands. Although we have concentrated on the animation of human-like character models, the basic ideas are applicable to other types of characters.

Goal-directed Navigation: We present a technique for quickly synthesizing collision-free motions for animated human figures given high-level navigation tasks amidst obstacles in changing virtual environments. The method combines a fast 2D path planner, a path-following controller, and uses cyclic motion capture data to generate the underlying animation. In order to provide a feedback loop to the overall navigation strategy, we incorporate an approximate synthetic vision module for simulating the visual perception of the character. It uses rendering hardware to quickly identify the objects visible to the character, along with a simple model for representing and updating a character's spatial memory. The model is efficient in both storage requirements and update times, and can be flexibly combined with a variety of higher-level reasoning modules or complex memory rules. Using such a sensing and memory model, a character can be made to explore in real-time an unknown environment, and incrementally build its own internal model of the world while it navigates toward a goal location.

Object Manipulation: We introduce a new manipulation planner designed for efficiently generating collision-free motions for single-arm manipulation tasks given high-level commands. For moving an object, the planner automatically generates the motions necessary for a human arm to reach and grasp the object, reposition it, and return the arm to rest. The planner searches the configuration space of the arm, modeled as a kinematic chain with seven degrees of freedom. Goal configurations for the arm are computed using an inverse kinematics algorithm that attempts to select a natural posture. To compute a collision-free path connecting the arm initial configuration to the goal configuration, we present an efficient, general path planning approach *RRT-Connect*. The planner grows two rapidly-exploring random trees (RRTs) in the configuration space and employs a simple greedy heuristic that aggressively attempts to establish a connection between the trees. This heuristic is demonstrated to be well-suited to generating motions for manipulation tasks commonly faced by human characters. For increased naturalness, we also present an approximate technique for coordinating the head and eye movements in conjunction with the arm motions.

For testing and evaluation purposes, we have designed an experimental implementation of the navigation and manipulation motion generation strategies and integrated them into an interactive application involving human characters. Though many improvements can be made, the generated animation looks fairly realistic.

Discussion

Nearly all known motion generation techniques developed in the graphics and robotics literature can be unified under a common framework by viewing the motion synthesis problem as one whose solution involves two fundamental tools: *model-building* and *search*. Specifically, solving a motion synthesis problem almost always involves *constructing a suitable model* and *searching an appropriate space* of possibilities. At present, no single motion synthesis strategy can adequately handle the variety of tasks that can arise. Instead, different tools and techniques can be used in combination and tailored to meet the requirements of specific tasks. We believe that assembling a library of algorithms and motion data libraries under the unifying framework of an animated agent will help to facilitate the development of autonomous characters with realistic motions and behaviors.

In the future, interactive virtual environments will be populated by lifelike autonomous characters. For offline animation applications, smarter software will alleviate the heavy burden of key-framing motions for complex characters. Next generation cinematographers will direct virtual humans in virtual scenes in much the same way as they direct real human actors. Ultimately, the dream of *virtual reality* may be fully realized, in which artificial scenes can be created that are indistinguishable from real life. In a sense, virtual reality defines a kind of "Turing Test" for animated agents (i.e. whether or not a user can distinguish animated agents from living people).

Let us consider for a moment some of the software technology and components necessary for realizing truly lifelike characters. Clearly, visual realism will involve both sophisticated graphical models as well as motions. For human characters, *expressive facial animation* using realistic skin models, and flexible models of *hair* and *clothing* that move in response to body motions will likely be important. In addition to having a realistic appearance, lifelike human characters will move fluidly and naturally, and exhibit believable behavior.

Besides the characters themselves, the virtual world they inhabit will also be lifelike. Since the real world is driven by the laws of physics, physically-based models will form an integral part of any attempt to realistically simulate humans in the real world. Even now, increases in computing power have allowed fairly sophisticated dynamic simulations to be calculated at interactive rates. Assuming that Moore's Law holds true, we can expect that future animation software will embody the laws of physics, and facilitate complex dynamic simulations. Given these advances, it is therefore natural to ask whether the fundamental problem of task-level motion control for human figures should be reformulated. Instead of trying to find a set of joint functions that solve a particular task, we should find a set of *controls* (i.e. muscle activation patterns) that solve the task. Ultimately, this could lead to a new representation of motions: *control functions* (forces and torques), instead of purely kinematic functions of joint variables.

For true believability, an autonomous character should exhibit realistic behavior. This implies that the character is capable of emulating some of the higher-level cognitive functions of real humans. This includes logical deduction and reasoning, adapting to changes in the environment, and learning from mistakes. Thus, an artificial life approach to animation that integrates perception, machine learning, and knowledge representation will likely be needed.

As mentioned in Chapter 2, the goal of building believable autonomous animated agents (virtual robots) shares much in common with efforts to build intelligent physical robots. However, a virtual robot has numerous advantages over a physical robot. While a physical robot must contend with the problems of uncertainty and errors in sensing and control, a virtual robot enjoys "perfect" control and sensing. This means that it should be *easier* to design an animated agent that behaves intelligently, than a physical agent that does so. In some sense, creating an intelligent autonomous animated agent can be viewed as a necessary step along the way towards creating an intelligent autonomous physical robot. If we cannot create a robot that behaves intelligently in simulation, it is unlikely that we will be able to create a robot that behaves intelligently in the real world. Thus, the "grand vision" for animated agents encompasses many of the long-term goals for artificial intelligence.

At the present time however, we have a more modest objective: the animation of lifelike autonomous characters for video games, animated films, and interactive simulations. Since these applications do not always require perfectly realistic motions, this more modest goal should conceivably be easier to achieve relative to the "grand vision". In any case, we believe that applications involving autonomous characters will continue to increase both in number and in sophistication as software and hardware technology improves. Clearly, many challenging research issues must be addressed before the dream of fully autonomous animated agents can be realized.