Published as: Dannenberg, Sanchez, Joseph, Saul, Joseph, and Capell, "An Expert System for Teaching Piano to Novices," in Proceedings of the 1990 International Computer Music Conference, Glasgow, Scotland U.K. September 10-15, 1990. Ed. Stephen Arnold. San Francisco: International Computer Music Association, 1990. pp. 20-23.

An Expert System for Teaching Piano to Novices

Roger B. Dannenberg, Marta Sanchez, Annabelle Joseph,
Ronald Saul, Robert Joseph, Peter Capell
Studio for Creative Inquiry
Carnegie Mellon University
Pittsburgh, PA 15213, USA
email: roger.dannenberg@cs.cmu.edu

ABSTRACT: The Piano Tutor is a computer system for teaching beginning piano students. The system is highly interactive, with an expert system to analyze student performances and a multi-media presentation system to deliver instruction. Score following, which matches performances against a model, is used as the basis for detecting student errors. The Piano Tutor gives intelligent feedback and help rather than just listing all errors that are detected. The curriculum is organized into a set of lessons that are automatically chosen by the system according to students' needs.

Introduction. The Piano Tutor is an expert system for teaching beginners how to play the piano. [1] The Piano Tutor can help students to avoid practicing bad habits as well as to teach new material when the student is ready. By adapting to students' needs, the system can teach students with different musical abilities.

In a session with the Piano Tutor, a student is first shown a video presentation covering a small amount of new material. Then, the Piano Tutor displays a musical exercise in conventional music notation on a computer monitor and asks the student to play. The student, seated at a MIDI-equiped piano, plays the exercise while the system automatically turns pages. If the student makes mistakes, the Piano Tutor analyzes the student errors, selects the error or errors that are most in need of correction, and offers a suggestion to the student using a combination of digitized voice and graphical feedback.

After the student has mastered the material, he or she may opt to perform the piece again with a poly-timbral accompaniment provided by the Piano Tutor. The Piano Tutor then selects a new lesson and the instructional process repeats.

The Piano Tutor also offers a practice mode in which the student can practice without comments from the system. Practice mode allows the student as much time as needed and offers the features of poly-timbral accompaniment, a metronome, hearing a model performance, and seeing annotations that indicate pitch and rhythmic errors.

The Piano Tutor goes beyond most computer-based instruction systems in several ways, including:

- the Piano Tutor teaches the psycho-motor skills of piano playing, going beyond music theory and ear-training that has dominated music instruction software;
- the Piano Tutor uses score-following software [2] which tracks the student's place in the score and tempo in spite of performance errors. This allows the Piano Tutor to ignore irrelevant errors and to focus on problems that need attention;
- the Piano Tutor is able to take low-level errors, such as a wrong note, and hypothesize a cause, such as difficulty playing a skip. This enables the Piano Tutor to provide

specific and helpful instruction rather than simply listing all the errors that were made;

- the Piano Tutor employs a student model enabling instruction to be tailored to the needs of a student. Lesson material is selected according to the skills a student has already mastered. Remedial instruction also can be chosen according to the level of the student;
- the Piano Tutor is a multi-media system, employing videodisc, computer graphics, digitized voice, and music synthesis.

Lesson Selection and Instructional Design. The Piano Tutor curriculum is organized in terms of what we call *lessons*, but *exercises* might also be a good term. A typical lesson consists of a presentation to the student, a performance by the student, remediation of student errors, and more performances and remedial efforts until the music has been mastered.

One of the important tasks of the Piano Tutor is the selection of lessons for the student. The Piano Tutor is organized according to principles of instructional systems design (ISD). The premises of ISD include that lessons only be taught when the student has demonstrated readiness for them. All lessons have a set of *prerequisite* skills that must be met before the student is allowed to take the lesson, and a set of *objective* skills that the lesson attempts to teach the student.

To test the integrity of our curriculum, software has been specially designed to help the instructional designer analyze the relationships in a complex network of lessons. In the Piano Tutor's case, there are now 70 lessons, each with different prerequisites and objectives, all of which must be satisfied in order to complete the curriculum. The instructional designer must account for the linking of skills to lessons such that all lessons are reachable and all skills can be taught through some sequence of lessons. This method shows great promise for future instructional systems.

The Piano Tutor models the knowledge of each student as a collection of skills the student has demonstrated. As new lessons are passed, the student model is augmented with new skills, but if the student has problems with a lesson, it is possible to remove a skill from the model.

To select a lesson, the Piano Tutor searches for lessons whose prerequisites are met and whose objectives teach something the student does not already know. If there are several alternatives, our current system gives the student a choice of what lesson to work on next.

Because the system can select lessons based on student need, the system takes an active role in the teaching process. This is in contrast to hypermedia systems in which the student is free to explore various presentations of information. We believe that in piano instruction, deciding what to teach next requires substantial knowledge of the subject. We have designed the system with this knowledge so that students will be guided in a productive manner.

Remediation. Once a lesson has been selected and instructional material has been presented, the student is asked to perform an exercise or play a given song to demonstrate mastery of the skills introduced by the lesson. It is a simple matter to capture the performance via MIDI, but it is a difficult problem to analyze the performance in order to provide feedback or update the student model.

Polyphonic score-following software is used to track the student's performance and "turn pages" on the computer monitor. A by-product of score following is the identification of wrong notes and a mapping from score time to real time. This leads to a collection of "primitive errors" consisting of discrepancies in duration, tempo, pitch, and loudness. These errors are processed

with the goal of identifying trends and higher-level explanations for the problems the student is having. For example, if a note is played late, it might be explained as dragging the tempo, playing a previous rest or note too long, or problems with repeated notes, depending upon the context in which the error was made.

The Piano Tutor then decides what error or errors should be reported to the student (reporting every error would be time consuming and hopelessly confusing). After deciding what error to work on first, the Piano Tutor must select appropriate remedial actions, again depending on context. For example, the Piano Tutor may say "you played the wrong pitch here", give specific directions like "repeat this passage", or even abandon the lesson and look for something easier.

The remediation process is aided by annotations in the internal representation of scores. For example, if a lesson introduces the use of the fourth finger, then notes to be played with the fourth finger will be annotated as such. If the student misses one of these notes, the Piano Tutor can highlight the note and generate a message such as "Be sure to use the fourth finger on this note; please try again." The use of annotations allows us to use a single remediation procedure for virtually every lesson. Even with the hints provided by annotations, the procedure is quite complex, but this minimizes the effort to add new lessons to the system and increases the consistency of the Piano Tutor's remediation.

Multimedia. A variety of input and output media are available to make the Piano Tutor interesting to use, effective in communicating, and flexible for us as developers. The output devices are a videodisc player, a bit-mapped display for text, music notation and graphics, an electronic piano, a multi-timbral synthesizer, and digital audio for voice. Input consists of the piano keyboard, a mouse, and a QWERTY keyboard (which is only used to sign into the system).

We have found this multiplicity of devices to be extremely useful for presenting information to students. The videodisc can store examples of hand position and posture. The higher resolution bit-mapped display is used for music notation. Because the display is computer generated, we can highlight notes and generate annotations in response to student input, a task that would not be possible with videodiscs. We also use the display for graphics-oriented theory lessons, user prompts and queries, and text messages. The piano, besides being an essential input device, can be played by the computer via Midi. The Piano Tutor can play exemplary performances or recordings of student performances. The synthesizer is used for accompaniments and synthesizing a metronome. Digital voice recordings are used extensively for instruction. We can store more voice than video, and by keeping voice on a separate disk, we can combine voice with still videodisc images.

The variety of media is also convenient for system development. We can use sequencers to capture Midi data and digitizers to record voice at much lower cost than producing video. When some segment of the videodisc is found to be unsatisfactory, we can supplement the video presentation with voice or graphics to avoid the delay and cost of replacing the videodisc (several days and several hundred dollars). We have implemented a number of theory lessons for tasks such as teaching notation that use computer graphics extensively.

Implementation. The Piano Tutor is written in Common Lisp and C and runs on an Apple Macintosh computer, which was selected for its combination of real-time capability and availability of Common Lisp. Unlike many efforts to use Lisp in real-time situations, we chose to implement virtually all real-time operations in C and to launch these operations from Lisp. This allows us to do some substantial real-time computation, such as the simultaneous following of a polyphonic performance, "turning pages" of a graphical score, analyzing the performance for

various errors, and synchronizing a Midi accompaniment. Lisp is used to make a detailed analysis of the performance after the student stops playing, to select lessons, and to handle the (non-real-time) graphical interface.

FrameKit [3], a software package for knowledge representation, is used throughout the Piano Tutor. For example, lessons—with their prerequisites, objectives, presentations, musical scores, and other attributes—are represented by frames. Presentations, which describe multi-media output sequences, are also represented by frames. Frame Graphics, a frame-based graphics system, is used for interactive graphics required by various "theory" lessons on notation and rhythm. A special-purpose graphical sequencer is implemented to help us "compose" multimedia presentations. The sequencer allows the designer to combine video, audio, and graphical elements into a presentation that can then be saved for use in one of the Piano Tutor lessons.

Status. The Piano Tutor concept, due to Marta Sanchez and Annabelle Joseph, has grown into a multi-year project of the Carnegie Mellon University Studio for Creative Inquiry. We have implemented about 70 lessons. Presentations are assembled from one half hour of video, over an hour of voice, and many musical scores. We are just beginning to test the system, and we plan to improve the system on the basis of our tests and experiences through mid-1991.

Acknowledgments. We are grateful to the Markle Foundation and the Studio for Creative Inquiry for their support of the Piano Tutor project.

References

1. Dannenberg, Roger B., Sanchez Marta, Annabelle Joseph, Peter Capell, Robert Joseph, Ron Saul. "A Computer-Based Multi-Media Tutor for Beginning Piano Students". *Interface* (to appear).

2. Bloch, Joshua J. and Roger B. Dannenberg. Real-Time Computer Accompaniment of Keyboard Performances. Proceedings of the 1985 International Computer Music Conference, Computer Music Association, 1985, pp. 279-290.

3. Nyberg, Eric H. *The FRAMEKIT User's Guide Version 2.0*. Computer Science Department, Carnegie Mellon University, 1988.