Chapter 3

Feedback in Musical Skill Acquisition

Implicit in the experienced musician's understanding of the relationship between action and sound, between performance gesture and musical phrase, is a more or less sophisticated internal representation of the dynamics of an instrument. This internal representation, built up over many years of practice, provides the skilled performer with an understanding of the instrument that is free from the context of any one piece. Such contextual independence provides the means for continually fresh interpretation and flexibility in performance. This chapter discusses the building of the musician's internal representation of an instrument's dynamics, drawing upon the theory of motor skill acquisition to develop a model of musical skill acquisition in which feedback plays a crucial role.

3.1 Modelling the Linkage

We learn to manipulate the tools in our environment by learning the relationship between our actions (in the form of efferent commands sent to muscles and joints) and the perceived reaction of manipulated objects (in the form of afferent feedback from our sense organs.) When we first encounter a new tool, we are very aware of its form (weight, material properties, etc.) As we become proficient users these aspects of the tool recede from our awareness and it becomes an extension of our body. In our mind, we have learned how the tool will behave and can predict for a given action the tool's response. As Loomis points out (Loomis, 1992) this process is the same regardless of the complexity of the tool or system being controlled. When a teleoperator learns to control a remote manipulator, he or she must learn the relationship between actions performed on the master and reactions of the slave. Loomis describes this learning process as "modeling the linkage" between efferent commands and afferent feedback from the remote manipulator. The operator's success, Loomis says, depends on the ease with which he or she can construct an internal representation of the physical linkage being controled and hence the degree to which the operator comes to regard the manipulator as an extension of his or her own body. (White, 1970) and (Loomis, 1992) have pointed out that our ability to project our apparent interface with the environment beyond the bounds of our body to the end of a tool (a phenomenon referred to in the literature as "distal attribution") is most likely to occur when an individual's sensory inputs (afference) are lawfully related to their motor or communication outputs (efference). Presumably it is the individual's recognition of this lawful relationship, often contingent on related past experience, that promotes the recognition of the object's identity and external location, i.e. the operator's ability to regard tools as extensions of his or her own body.

The occurrence of distal attribution, inasmuch as it represents a robust internal representation of the relationship between action on and reaction of an object or system being manipulated, can be regarded as a measure of skill. In the case of the musician, for example, the player's model of the behavior of the instrument, what Loomis refers to as the "Model of the Linkage," is their representation of the linkage (musculature and sense organs) and any physical extension (e.g. bow, etc.) intervening between efferent commands (in the form of, say, bow strokes) and afferent feedback from the instrument (in the form of its sound and, we suggest, tactile cues as well.) This representation might not necessarily be cognitive or even accessible to conscious awareness.

To the extent that the musician can “model the linkage” between input commands and output sound, either because the extension is natural and simple (e.g. a drum stick) or because the player has learned it through extensive training, the experienced performer eventually reaches the point where the linkage and the extension are transparent. The bow tip effectively becomes an extension of the hand and the perceived interface with the violin shifts from the hand/bow interface to the bow-hair/string interface.

3.2 Hierarchical Processes in Movement Control

So how does the novice player become an expert musician?

For a musician, the learning process involves committing to memory sequences of movements at many levels of abstraction. At the lowest level, there is the individual movement responsible for a single gesture – hitting a piano key, bowing a string, etc. At a higher level, sequences of these gestures are combined into learned patterns — scales, genre-dependent ornaments, etc. At the highest level of all is the organization of these learned patterns into a temporal sequence uniquely designed to execute an individual piece of music. Thus the hierarchical nature of learning involves moving between levels of more or less complex information. In this respect, learning to play a musical instrument is like learning to use any other tool. Therefore we here ask the question: How do we learn sequences of actions that allow us to execute highly complex movements? By looking at what is known about manual skill acquisition, we can cast light on skill acquisition in music.

3.2.1 Manual Skill Acquisition

Several theories of manual skill acquisition have been proposed. Common to many is the concept first proposed by Bryan and Harter (Bryan and Harter, 1897, 1899) of progression through stages of learning to higher orders of skill. Their work examined the perceptual and motor changes that occurred during the acquisition of telegraphic skills. Skill, in Bryan and Harter's view, was a process of achieving a hierarchy of "habits." The most fundamental skill was learning to discriminate perceptually and motorically between units of time, a skill that is quickly learned. In Morse code, a dot is one unit of time of continuous auditory signal, a dash is three units. One unit of no signal occurs between dots and dashes within a letter. Three continuous units of no signal denotes the start of a new letter and six units denotes the start of a new word. The next skill or "habit" was being able to recognize and transmit sequences of dots, dashes and pauses as letters of the alphabet. This skill too is learned relatively quickly. Then Bryan and Harter noticed something about the rate of improvement of some of their subjects. At times subjects seemed to reach plateaus in performance. They noted that these plateaus occurred prior to the formation of a new advanced capability. They proposed that, rather than hearing dots and dashes, telegraphers hear letters, with practice words and for the most skilled, sentences. Plateaus in performance occur because a maximum performance capability of one "habit" places a constraint on performance that is then lifted when a higher order habit is formed.

What Bryan and Harter observed in their study was the process later called "chunking." Observed first as a way of extending the capacity of short-term memory by grouping "like" elements of incoming information into meta-elements (see Simon, 1974), chunking also operates on the output from the motor system, functioning to combine multiple component movements that together execute an action into one chunk or "motor program" (Keele, 1973; Schmidt, 1976). It is not clear from Bryan and Harter's study whether the new capabilities responsible for improvements in performance, and hence a higher order habit, resulted from being able to chunk incoming auditory signal or from being able to draw on learned sequences of movements to produce outgoing code. What is certain, though, is that skilled operaters, whose average speed was between 20 and 25 words of approximately four letters each per minute (80 to 100 characters per minute) Bryan and Harter (1897) would not be able to use feedback from the auditory system to verify their output, since the minimum time required to respond to an auditory signal is around 150 msec (Phillips, 1977). Therefore, these operators must have been triggering learned sequences of movements or motor programs in response to higher order cognitive instructions to key a given letter or word. Essentially, motor programs can be thought of as choreographed actions in a notation that can be readily interpreted by the motor system. They are derived from one or more levels of an abstract goal making new information explicit at each level. At their lowest level they are sufficiently abstract to allow the motor system freedom to compute a movement adaptive to the current context.

Motor programs serve two important functions: Firstly, they enable the execution of actions that occur too close together to be processed individually by the sensorimotor system. Kalmeman (1973) examined the effect on the “inter-response interval” (the interval between responses of the motor system to incoming stimuli) of presenting incoming stimuli closer together in time. He found that as the time between stimulus onsets decreased, the interval between individual responses to these stimuli also decreased, but only to a certain point. Provided the stimuli are not grouped and processed similtaneously, no matter how small the interval between them, approximately 200 msec occurs between responses. This finding suggests that when a stimulus arrives a response to it is generated. If another stimulus is presented soon afterwards, indicating that the system should do some other action, the second action must wait for at least 200 msec before it can be initiated. If there were no way for the motor system to overcome this limitation, we would be unable to produce even repetitive movements at speeds greater than 5 Hz. And yet pianists can play trills at speeds of up to 20 Hz and can play repeated notes with the same finger at speeds of 10 Hz. Motor programs are the system’s way of overcoming the need for movements that are separated by less than 200 msec by generating a response that in itself is complex and involves many movements in rapid succession. The motor program is triggered as one response. It is presumed that according to this general theory, these programmed outputs occur in discrete bursts, separated by at least 200 msec. These discrete elements are difficult to view directly, however, because the muscles and limbs smooth out the transitions between elements, giving the impression that response is continuous. Building motor programs to execute sequences of actions is therefore a key component of acquiring a manual skill, particularly when the skill requires rapid movements.

The second important function of motor programs is to allow for levels of supervisory control, freeing up the central processing system for other work. We are able to walk down a street and carry on a conversation with a friend only because we do not have to consciously dictate the intricate sequence of commands to muscles and joints necessary to bring about walking. Evidence for such supervisory control of movement is found in many domains of motor skill. Typists who were instructed to type a single familiar word repeatedly, for example, were unable to switch to typing a second word without completing the first, indicating that they could not interrupt the motor program generating its sequence of key strokes (Shaffer, 1991).

3.2.2 Motor programs in music

The hierarchical nature of music performance, though, requires more than just an ability to execute sequences of actions on demand. A skilled musician will not and probably cannot play the same piece in exactly the same way twice. Though low level motor programs may be executed to produce individual notes and patterns of notes, their expressive contour can be continually shaped and reshaped in real time. How can learned sequences of actions, complete with their own temporal structure, be overridden on demand?

Central to the modern theory of motor skill is the concept of a central system that contains models of the subsystems it controls. Thus the motor system contains a model of the musculature and the cognitive system contains a model of the motor system and of the environment in which it acts.

Planning an action such as the execution of a phrase of music therefore entails constructing a mental representation of its goals relative to a cognitive model and these representations in turn serve as a basis for producing the sequence of movements that give rise to the notes themselves. Indeed, modern theories of motor behavior suggest that the division between cognitive "schema" and motor programs is not cut and dried (Schmidt and Lee, 1999). This is borne out by the fact that people can play expressively at sight and can modify performance on the fly, suggesting that, although much of the low-level response is governed by motor programs set in place when techniques for playing an instrument are learned, higher-level cognitive processes can still intervene to modify low-level control. As noted by Shaffer (1993), a player may give a very different performance of a piece having had weeks to study it than if sight-reading. Moreover, the circumstances of an individual performance may influence the way a piece is played, reshaping expressive gestures and thereby influencing the execution of learned sequences of movements.

What then are the cognitive processes that link musical expression to movement control? Musical expression is difficult to quantify. It is that component of performance that cannot yet be generated by even the most sophisticated computer algorithm working from a detailed coded score complete with its dynamic, tempo and articulation markings. Performers express their interpretation of a piece in musical phrases, moving between tension and relaxation by manipulating timing and dynamic contours. Shaffer's results indicate that attempts to define context-free or even context-limited rules for expression are of limited value. There is unlikely to be a simple mapping between musical patterns in a score and expressive devices in a performance. A performer does not simply reproduce what is on the page of a printed score. Rather, the performer mediates between a text and its performance. Thus there exists a hidden underlying variable — interpretation. We cannot assume a simple mapping from text to interpretation on the one hand and from interpretation to expression on the other.

Evidence from two sources points to the notion that, though a cognitive executive may override the execution of motor programs in music, the scope of this override is highly correlated with the boundaries of the musical phrase. The first source of evidence is from an observation on rubato in performances of a Chopin mazurka by Rosenthal (Koprowski and Barth, 2000). Rosenthal was reputed to be a somewhat erratic performer, seemingly varying his interpretation of music at will. In untangling an anomaly in catalogued recordings of Chopin's mazerkas by Rosenthal, Koprowski synchronized two recordings to compare them phrase by phrase. Then Barth noticed that, though Rosenthal varied the dynamic and timing trajectories of phrases considerably between the two performances, the two recordings were always synchronized at the ends of phrases. This would suggest that Rosenthal had some notion of an underlying temporal scaffold at the level of the phrase, a baseline for tension-relaxation trajectories. Shaffer and others (Shaffer and Todd, 1987; Kugler and Turvey, 1987) argue for the existence of an internal timekeeper, a kind of programmable clock that can modulate its pulse rate under instruction from the motor system. The fact that it can return very accurately to a particular tempo shows that this clock has stable referents. If Rosenthal's internal clock is a neutral oscillator with phrase boundaries acting as attractor points, then his rubato excursions are like departures or perturbation away from an attractor. Shaffer's work on hand independence in piano playing also indicates the presence of a hierarchy of time-keepers (Shaffer, 1993) suggesting that a musical pulse is generated at the level of abstract units while the motor system arranges the kinematics of playing individual notes in relation to this. Where two hands play very different rhythms, the motor system must be able to subdivide pulse intervals differently for each hand. Thus the executive level operates independently of motor programs that govern either hand, though it can intervene to effect the expressive goals of each.

The second source of evidence for overriding movement control at the level of the phrase comes from Carolyn Palmer's work on cognitive representation of musical structure. Palmer discovered that players were far more likely to have memory slips in transitioning from one phrase to the next than within the boundaries of a phrase. Moreover, in order to recover, players usually back-tracked to the start of the phrase in which they had stumbled rather than taking up exactly where they had stopped (Palmer, 1997).

Thus it seems musical phrases and their attendant sequences of motor programs are organized in memory in phrase-level chunks that can be reshaped according to the demands of a given performance situation. If the phrase is interrupted because the player is distracted or forgets, it is most likely that the performer will have to backtrack to the start of that chunk. These studies together tend to support the phrase as an important structural element in organizing and executing motor programs in music performance.

To summarize, acquisition of motor skill in music entails more than simply knowing how to play an instrument. The player's understanding of an instrument's behavior provides them with the ability to superimpose upon a piece a new expressive contour at will. This ability to reinterpret music at will is evidence that the cognitive and motor systems, hitherto thought to be independent layers of the hierarchy for movement production, interact to produce performances of music that are adaptive to current demands. Finally, there is considerable evidence that phrase boundaries are important anchors for the cognitive structures that generate movement production in performance. In Chapter 6, we will return to the question of hierarchical control in music performance in the context of translating performance nuance, conveyed by movement, into expressive control of computer-based musical instruments.

In the next section of this chapter, we will turn once again to low-level motor programming and consider the role of feedback in the process of movement production.

3.3 Feedback in Manual Skill Acquisition

Feedback, in the context of motor skill acquisition, can be defined as that class of sensory information that is movement related. As such it includes visual, auditory (verbal or non-verbal), haptic (tactile and kinesthetic) and proprioceptive information picked up by our sense organs. Feedback can be classified into two basic categories:

1. inherent feedback — feedback that is intrinsic to a task, e.g. balancing on a bicycle, and

2. augmented feedback — feedback that is supplementary to a task, e.g. a score

achieved for a given skill or a teacher's comments on its execution.

In this section we focus on inherent feedback, particularly as it relates to the acquisition of a motor skill. Under what conditions does the motor system rely on constant feedback from sensory apparatus and when does the central control mechanism that generates motor programs take over?

3.3.1 Feedback in Movement Control

One way of modeling motor control has been to consider sensory control of movement as analogous to control of mechanical systems. Thus like the thermostat that controls the temperature of a room, the sensory motor system executes actions with reference to a known goal, comparing feedback from the environment with a known reference of correctness. An executive control mechanism, informed of the difference between actual and desired states, may then cause the system to act to reduce this difference. Such a system, which continuously samples its environment to determine if action is necessary, is called a closed-loop system. Closed-loop systems are important in many situations, especially those that require a system to "control itself" for long periods of time.

Unlike the thermostat, though, the "reference of correctness" for movement is not a single state but a set of states that change with time. For each moment in time, the reference of correctness has a different specification for position of the moving limb. Because the reference is constantly changing, it can be matched against the feedback from the moving limb, so that errors in the movement's trajectory can be detected and corrected. This is the basis for Adams' theory of learning, according to which the subject learns a set of references of correctness that the closed-loop system is to track during a movement (Adams, 1971). As noted earlier in this chapter, such a view of motor processing has several limitations. Firstly, the system must process information very rapidly even for the simplest of movements. Secondly, all the references of correctness must be stored somewhere, and will create problems if the movement starts from a slightly different place or takes a different path, because each of the sampling points will be different.

Modern theories of motor control therefore lean toward a model where slow movements and movements that require fine control closely track sensory feedback, while rapid and learned movements are executed essentially open-loop, relying on motor programs initiated by a central control mechanism. Because motor programs are structured in advance, they are executed without reference to sensory feedback. However one possibility that is particularly relevant to music performance is that even when a player is highly skilled, sensory feedback may still function to inform the motor system of unexpected perturbation in movements, whether or not it is involved in their control. In examining the possibility that vibrotactile feedback from a 'cello might be a useful cue to a player, Chafe (1993) concluded that while most of the time the vibrations coming back through the bow were too high to be perceived by the tactile system, the bursts of noise that result from changes in bowing direction or large slips generated signals that could be felt. He conjectured that the tactile system might use these bursts of noise that represent perturbation of an otherwise stable system as cues about potential loss of control of the bow. If this is indeed the case, it raises one more important question: How can the motor system differentiate between movements that are results of its own actions and those that are caused by disturbance in its environment?

3.3.2 Feedforward Signals in Movement Control

The question of how we track our own movement relative to the movement of other objects in our environment was first studied experimentally by von Holst (1989). In explaining how we know whether our eye has moved relative to an object or whether the object has moved relative to us, he suggests that the visual perceptual system is informed about upcoming eye movement ahead of time. Copies of commands destined for end effectors in the motor system are sent to known sensory centers in the brain (Evarts, 1973). Moreover, research on roughness perception indicates that we are more sensitive to stimuli in our environment when we actively explore than when we passively observe (Lederman and Klatzky, 1997). Why should this be so? One answer is that when the motor system sends the commands to move actively, it also sends a copy of efferent commands to sensory areas in the brain to enable the feedback to be evaluated properly. When the finger is moved passively, no motor commands are issued by the muscles, hence there is no efferent copy and thus the same feedback signals from the finger are not perceived so accurately. Thus feedforward signals prepare the motor system for upcoming interactions with objects in the environment, anticipating actions and reactions on the basis of past experience.

Of direct relevance to the work presented here are a series of experiments by Bhushan and Shadmehr(Bhushan and Shadmehr, In Press) that examined the hand trajectories of subjects making reaching movements in a novel force field. They point out that, when a novice operator learns to control a novel mechanical system, the brain must solve three types of computational problems. Firstly, it must figure out how the mechanical system should behave in order to achieve a given goal; secondly, it must learn the dynamics of the mechanical system in order to be able to predict how the system will respond to a given input, i.e. it must build a so-called forward model of the system's dynamics; lastly, it must figure out the input commands necessary to bring about a desired change in the mechanical system, i.e. it must build an inverse model of the dynamics of the system. In their work, Bhushan and Shadmehr systematically modeled a reaching task, taking each of these three components in turn to discover how much of the psychophysical data from their subjects it could account for. Using the forward model alone, their controller was unable to account for situations where both the derivative of hand speed and the direction of hand velocity changed rapidly. On the other hand, a model based solely on descending neural commands and the time-delayed sensory feedback from the force field was inherently unstable. Moreover, gains in the feedback loop that generated an error signal could not be set large enough to provide an adequate model of the inverse dynamics of the distal system, i.e. the arm coupled to the force field. This was particularly true when the distal system's dynamics were changed, i.e. when a novel force field was introduced. Bhushan and Shadmehr suggest, based on these findings, that the supra-spinal controller is neither a purely feedforward system using only an inverse model, nor a purely feedback system that uses only a forward model. The behavior of the human arm, particularly about its segmentation points, they suggest, is due to the action of a supra-spinal feedback system that uses a forward model that provides an error signal that through an inverse model results in modification of descending motor commands. They concluded that these reaching movements could best be modeled by a controller that incorporated both a feedforward model of the inverse kinematics of the arm and a feedback model that predicted position and velocity of the arm from both a copy of descending motor commands and the latest sensory feedback from the moving arm.

In light of these findings, it would seem that efference copy not only allows the motor system to keep track of its own effect on the environment, but it also performs an important function in maintaining stability by maintaining a trajectory of equilibrium to which the motor system can refer if perturbed by external forces.

3.3.3 Summary

In this section we have focused on the role of inherent feedback in movement production, particularly as it applies to the acquisition of motor skill. The time taken for feedback from sense organs to reach the brain and the time required to process this feedback prohibit its use as the primary determinant of subsequent action. Inherent feedback, which informs the performer of the success of his or her actions, is therefore unlikely to be based purely on sensory input and is most probably a combination of both sensory feedback and feedforward commands (in the form of efference copy). It should be pointed out that becoming sensitive to the feedback inherent to a skill is itself part of the learning process. The novice approaches the learning of a new skill with no knowledge of what it should feel like when it is performed correctly. Initially the learner depends upon feedback from a teacher (augmented feedback) to guide his actions, reinforcing good performance and indirectly establishing a “feel” for the task.

With increasing proficiency, the pupil comes to recognize and trust inherent feedback and can judge the goodness of these actions. The goal of a good teacher, therefore, is to train the student to train himself. The final section of this chapter examines this changing role for feedback in the learning process.

3.4 Feedback and Motor Learning

So far this chapter has considered the role of sensory feedback in performance of skilled movement, particularly as it relates to the performance of music. In this final section we consider how the role of sensory feedback changes as a novice player acquires musical skill. Motor learning is defined as a set of internal processes associated with practice or experience leading to a relatively permanent change in a capability for skilled behavior, the state that Bryan and Harter termed "habit." (Bryan and Harter, 1899) Inherent in the learning process for the musician is a transition from feedback-dependent exploration of an instrument's dynamics to feed-forward dependent skilled performance. The transition from novice to expert is not smooth but, as we have already seen, is marked by distinct plateaus in performance. In the literature, this non-linearity in performance improvement with practice is described by the "law of practice" (Schmidt and Lee, 1999) which states that improvements in performance, at first large and rapid, become smaller as practice is continued.

What components of the learning process define the rate of improvement in performance of a skill? As has already been suggested, students measure success of their actions as they learn by monitoring feedback from their environment. Initially, this feedback is extrinsic, often in the form of instruction from a teacher or coach who reinforces correct performance. As learning progresses, the student relies more and more on his own judgment and comes to recognize when actions "feel" right. This transition from extrinsic to intrinsic feedback is an essential component of learning because it provides the student with his own "reference of correctness" for a task and allows him to practice fruitfully on his own. Ultimately, motor programs will be established so that skilled movements are executed independently of feedback from the environment.

Fitts described the process of learning in terms of three phases of practice (Fitts, 1964; Fitts and Posner, 1967). In the first phase, which Fitts termed the "cognitive phase," the primary concern for the learner is to understand what is to be done, how performance is to be evaluated and how best to attempt the first few trials. Here, considerable cognitive activity is required so that the learner can determine the appropriate cognitive strategy. Good strategies are retained and inappropriate ones are discarded. As a result, the performance gains during this phase are dramatic and generally larger than at any other single period in the learning process. Performance is usually very inconsistent, perhaps because a learner is trying many different ways to solve the problem. The learner is focused on the immediate task, paying little attention to anything other than sensory feedback directly related to the actions required. This is the phase when augmented feedback in the form of instructions, models and training simulators is most useful.

The second phase, which Fitts termed the "associative phase," is characterized by more consistent but smaller improvements. The learner has selected an appropriate strategy and spends time honing the skill, making subtle adjustments to improve its efficiency. The term "associative" refers to the fact that the learner now has some context for his actions. In the case of the musician, for example, decisions about interpretation are made on the basis of both past instruction from a teacher and knowledge of context such as musical genre. At this stage the learner can practice independently, relying on an internal reference of correctness to judge the success of his actions. This is the stage of motor learning that we think of as "practice" and it can last for weeks or even months.

Fitts' third stage of motor learning, which he called the "autonomous phase," is characterized by highly skilled performance. Sequences of movements have been grouped or chunked to form complex motor programs. The performer, freed from the burden of consciously controlling every movement, now operates in a supervisory capacity.

Expert pianists, for example, can perform simple arithmetic while continuing to play (Shaffer, 1993).

Because learned actions are governed to such an extent by established motor programs (feedforward control), sensory feedback is generally believed to have little effect on performance. However, as we suggested earlier, sensory feedback may function to alert the performer to sudden changes or perturbances in the system.

In summary, the acquisition of a motor skill is paralleled by a transition from extrinsic or augmented feedback, toward intrinsic feedback and finally feedforward or anticipatory control.

3.5 Summary and Conclusions

This chapter has explored the role of feedback in learning and performance within the context of musician/instrument interaction. Drawing upon the literature of motor control and learning, a model for the role of sensory feedback in musical performance has been proposed. In the early stages of learning to play a musical instrument, feedback from many sources other than the instrument's sound plays a part in helping the player select strategies for successfully achieving a goal. Where a mechanical coupling between player and instrument exists, for example, haptic cues can function to help the player build an inverse model of the instrument's dynamics. Once selected, these strategies are reinforced by the player through practice until they are eventually encapsulated in motor programs.

Figure 3.1 describes the relationship between the phases of skill acquisition and the building of the cognitive representations that support them.

In order to reach the level of proficiency of an expert performer, the brain must be able to figure out the behavior of the mechanical system it is to control — in our case the musical instrument. The motor system must also learn the dynamics of the instrument to be able to predict its response to a given action. Over years of practice,

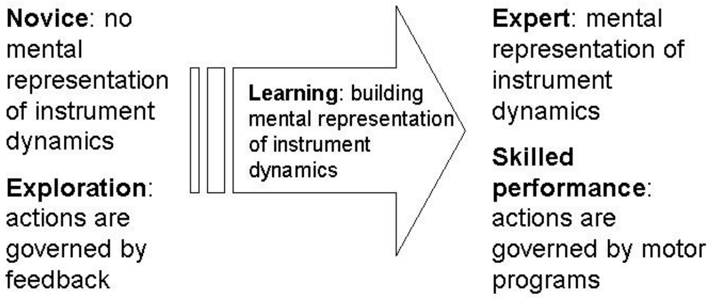

Fig. 3.1. The role of feedback in the stages of musical skill acquisition.

the skilled musician's mental representation of an instrument's behavior becomes so sophisticated that it is even possible for an experienced player to predict responses to actions never actually practiced. This is the case for skilled improvisation. Indeed, the skilled musician can hear music in his head by imagining the movements involved in playing notes on an instrument. As Caroline Palmer has shown (Palmer, 1997), anticipatory behavior increases with increasing levels of musical skill even when the music being played has never been seen before.

Lastly, in learning a new instrument, the brain must figure out the input control commands necessary to bring about a desired change in the instrument's response, i.e. the mapping between input gesture and expressive performance.

In the next two chapters of this work we present a series of experiments designed to explore the process by which a player builds an internal representation of an instrument's dynamics. Using computer-generated programmable haptic feedback we systematically control the "feel" of virtual musical instruments. Experiments I, II and III are concerned with the very early stages of learning. Here, the player explores the dynamics of a new instrument and begins to internalize these dynamics. Experiment IV explores the internal representation of the interaction between bow and string for a player with several years of experience.