Cognitive Models of Emotionally Labeled Compositional Design

James Peterson [*]

May 24, 2003

Abstract

In this paper, we discuss the design of a cognitive model for emotionally labeled compositional design using sensory to auditory, visual and associative cortex pathways and a limbic system based model of neutral emotion processing. The auditory cortex portion of the model is trained uisng data based on music (neutral, happy, sad and angry); the visual cortex, data based on three layered paintings (neutral, happy, sad and angry); and the associative cortex, data based on genetic optimization plans for job scheduling with intangible labeling (neutral, contented, demoralized and antagonistic). The outputs from the cortex models combine with the outputs from the limbic system. The training data we use provides three sets of emotionally labeled emotional outputs to constrain the output of the limbic system. The music and painting based data provides constrains the basic emotional outputs of the limbic system which we label as happy, sad and angry, while the genetic optimization data constrains higher level qualia. The full system then provides outputs which given a starting musical, painting or optimization element, autonomously generate a full musical phrase, a painting or an optimization plan for job scheduling in a chosen emotional state. There are two goals to this research: first, provide a basic model of emotionally labeled compositional design ; and second, use this model also as a fundamental building block in a more general model of cognition for use in the study of cognitive dysfunction.

1 Introduction:

In this paper, we discuss the design of musical data for use in the training and building of models of cognition which consist of many parts such as cortex (auditory, visual and associative) and emotions. In the companion paper, (Dzuris & Peterson, 2003a), the details of how this data is prepared when it is emotionally neutral in tone, have been discussed. In this paper, we focus on how we solve the problem of generating musical

---

[*] Department of Mathematical Sciences, email: petersj@clemson.edu

compositional designs of a particular emotional slant (happy, sad and angry) for use in the training of the cognitive model. There are two goals to our research: first, model the process of musical compositional design for eventual use in the development of an autonomous musical composition program with a mixture of emotional content; and second, use the musical model as a quantitative means of training the auditory cortex portion and associated emotional circuits of a general model of cognition.

In Section 2, we show how pivotal experiments from psychophysiology give us valuable clues about how to organize data arising from visual images into emotional categories. From this, we infer which musical data sets are useful to generate and we show that this paper, which concentrates on the development of the three other portions of our emotional data triangle, furnishes us with the final portions of the algorithms that allow us to generate large amounts of data for the full model of cognition and emotion.

The complete description of how we generate emotional musical data and rules which we used for our compositional elements is found in Section 3. We begin with a careful discussion of what has previously been done in the literature in 3.1. In Section 3.2, we discuss a rapid prototyping technique used in the 18^{th} century know as a Würfelspiel matrix (Cope. 2001) and show how these ideas can be phrased in abstract grammatical terms. Then, in 3.3, we discuss in detail the rules which we use to generate emotionally slanted musical designs. We show a simple alphabet choice we could use to encode our data in Section 3.7. We finish with samples of the generated compositional designs; happy data is shown in Section 3.4, sad in 3.5 and angry in 3.6.

In Section 4, we discuss how this work is placed within the more general area of cognitive modeling and how a model of musical composition provides us with a useful means of validating complicated and complex cognition models. We conclude this section with a short discussion of some of the clues and ideas from biology and cognitive science that partially shaped our thinking on how to build useful abstract models.

In Section 5 and 5.1, we lay out the basics of the cognitive modeling architectural design we use for modeling the music composition process. In essence, we need to capture why certain notes in a valid composition are preferred over other choices. The discussion is based on the use of suitable encodings of the musical notes into mathematical forms useful for model building. Since we do not actually build these models in this paper, the discussion is necessarily brief. In our brains, pathways that are useful are often given an enhanced probability of use via a process called excitation; similarly, pathways of limited usefulness are actively discriminated against using a process called inhibition. In Section 5.1, we discuss an overview of the modeling process that uses inhibition and excitation to ensure that only valid notes are selected in our neutral music samples. Hence, we train our model so that it captures how notes are generated in the data faithfully. Then, section 5.2, delineates the procedure by which we construct the transitional mappings between opening and middle and middle and closing phrases. Finally, in 5.3 and 5.4, we show how actual musical fragments can be generated by constructing the analog of sentences in this grammar. We close with a short discussion with the role the musical data plays within the context of a full cognitive model.

2 Emotion Models:

In a sequence of seminal papers , Lang et al. (1998), Codispotti, Bradley and Lang (2001) and Cuthbert, Bradley and Lang (1996)), it has been shown that people respond to emotionally tagged or affective images in a semi-quantitative manner. Human volunteers were shown various images and their physiological responses were recorded in two ways. One was a skin galvanic response and the other a fMRI parameter. Typical results are plotted in Figure 1. In this database of images, extreme images always generated a large plus or minus response while neutral images such as those of an infant generated null or near origin results.

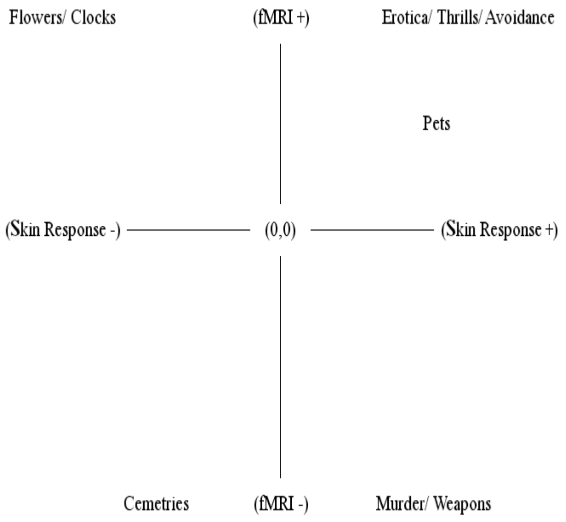

Figure 1: Human Response To Emotionally Charged Picture Data

If we followed the original intent and spirit of the Affective Image research, we would like to develop or generate Würfelspiel type arrays for each of nine primary emotional states as indicated in Figure 1. These would be emotional states that correspond to the following nine locations on the two dimensional grid:

2D Coordinates	Physiological Responses	Image Type
(High, High)	high galvanic and high fMRI response	Thrills
(High, Null)	high galvanic and flat fMRI response	
(High, Low)	high galvanic and low fMRI response	Murder
(Null, High)	flat galvanic and high fMRI response	
(Null, Null)	flat galvanic and flat fMRI response	
(Null, Low)	flat galvanic and low fMRI response	
(Low, High)	low galvanic and high fMRI response	Flowers
(Low, Null)	low galvanic and flat fMRI response	
(Low, Low)	low galvanic and low fMRI responses	Cemeteries

Clearly, the emotional tags associated with the images in the affective image database are not cleanly separated into primary emotions such as anger, sadness and happiness. However, we can infer that the center (Null, Null) state is associated with images that have no emotional tag. Also, the images do cleanly map to distinct 2D locations on the grid when the emotional contents of the images differ. Hence, we will assume that if a database of images separated into states of anger, sadness, happiness and neutrality were presented to human subjects, we would see a similar separation of response. Our hypothetical response would be captured in the emotion triangle Δ as seen in Figure 2.

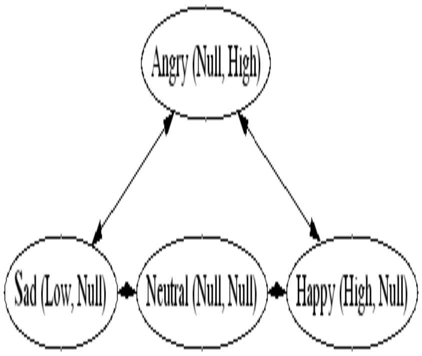

Figure 2: Emotionally Charged Compositional Data Design

In both the musical and painting compositional domain, we will therefore design Würfelspiel matrices for the four positions marked in Δ. In addition, we can identity the intangibles *antagonistic* with the emotional attribute *anger*; *demoralized* with *sad* and *contented* with *happy* and use a similar triangle to design job scheduling data. The motivations and arguments that explain these mappings are explained in our job scheduling papers (Kurz & Peterson, 2003e and 2003f).

In this paper, we will concentrate on how we generated the musical data that is interpreted as *anger*, *sadness* and *happiness* in the triangle Δ.

3 Generating Emotional Musical Data:

In (Dzuris & Peterson, 2003a), we discussed how we would design a Musicalisches Würfelspiel matrix which contained simple emotionally neutral musical compositions. In this paper, we will extend our construction process to include Musicalisches Würfelspiel matrices that are sad, angry and happy. Our review of previous work has led us to the following observations:

3.1 Emotion and Music:

Schellenberg, Krysciak, and Campbell (2000), "*Perceiving Emotion in Melody: Interactive Effects of Pitch and Rhythm*" specifically addresses music and emotion. The researchers were trying to decipher what effects two specific musical elements (pitch and rhythm) had upon the perception of listeners. Before manipulating elements, they had to establish a set of melodies that "unequivocally expressed one of three emotions: happy, sad, or scary." The authors decided that each of the three emotions: happy, sad, scary, is considered to be a "basic" emotion (see e.g., Ekman and Davidson (1994)). Further, based on a review of literature cited in their paper (Hevner, (1935), (1936), and (1937); Kratus (1993); Sloboda (1991); Terwogt and Van Grinsven (1991); Thompson and Robitaille (1992), there is a consensus that "adults from a common culture generally show broad agreement when associating such emotions with particular pieces of music. Other research shows that young children have similar associations (Cunningham and Sterling (1988); Dolgin and Adelson (1990); Giomo (1993); Kastner and Crowder (1990); Kratus (1993); Terwogt and Van Grinsven (1991))

Various attributes were used to describe melodies in the each of the three emotion categories. In *happy* melodies, there should be fast tempi (Hevner (1936); Rigg (1940); Scherer and Oshinsky (1977); Wedin (1972)); major modes (Crowder (1984) and (1985); Schere and Oshinsky (1977); Wedin (1972); Geraldi and Gerkin (1995); Gregory, Worral, and Sarge (1996); Kastner and Crowder (1990); Kratus (1993)) and staccato articulation (Juslin. 1997). On the other hand, in *sad* music, the use of lower pitches (Hevner (1936); Crowder (1985); Wedin (1972)) along with minor modes(Crowder (1984) and (1985); Schere and Oshinsky (1977); Wedin (1972); Geraldi and Gerkin (1995); Gregory, Worral, and Sarge (1996); Kastner and Crowder (1990); Kratus (1993)) and legato articulation (Juslin. 1997), are necessary to elicit the proper tone of sadness. Finally, in *scary* music, there should be a broad pitch range. Any instrument can produce staccato, legato or broad range sound. However, in the study above, listeners associated the staccato articulation produced by a guitar to a happy state; the legato articulation of a violin to a sad state; and the broad pitch range of an organ to a scared state.

Meyer, (1956), "*Emotion and Meaning in Music*", examines "meaning" in music. His arguments are based on what he labels an absolute expressionist viewpoint. This specifically means "expressive emotional meanings arise in response to music and that these exist without reference to the extra musical world of concepts, actions, and human emotional states." Expectation is a concept that Meyer considers a product that comes from natural mental processes of perception. The process involves instinctive grouping and organizing of information coming in through the senses. Applying this logic to music, Meyer states that music elicits varying responses from a listener by manipulating the expected. For example, musical progression that moves in an irregular way throughout elicits a feeling of suspense or ambiguity for listener. Why? Meyer claims the listener would begin to doubt the relevance of his own expectations. This may be true for a trained musician, but we are not so sure that the average listener is aware of having certain expectations and therefore does not consciously go through phases of doubt. Meyer's second argument makes more sense to us. It is the opposite notion that if the music is so uniform or repetitive, then the music itself has ambiguity in that it seems static, going nowhere.

Meyer links our experience of modulation (shifts in tonal center) and key changes to departures in a narrative line in a novel. These complications of the plot function like our extensions in musical phrases and is simply understood by the listener. Different musical styles are complex systems of probability. This idea seems to tie into what Cope did in later years (see Cope (1991) and (2001)). Cope entered many examples of Chopin and then used a complex system of computed probabilities to form Chopin-like pieces. We think that to know what expectation to have, based on probability, and to understand additive elements in writing and music is due to having encountered them before. Thus, we feel this is a human experience rather than some intrinsic response.

In a later section, Meyer states that an expectation must "have the status of an instinctive mental and motor response, a felt urgency, before its meaning can truly be comprehended." He does not go on to say how that line is crossed, but suggests that it is the deviation of the pure tone, exact intonation, perfect harmony, rigid rhythm, etc.. which conveys emotion.

As we have seen documented in multiple sources, there is an association between minor mode and the emotional state of sadness. Meyer adds the emotional state of suffering as well. He reasons that these types of emotion are a product of the unstable character of the mode itself. The unstable character refers to the fact that the minor mode is presented in different versions: melodic minor, natural minor, and harmonic minor. Because it is possible and likely to have a combination within one piece, it is very chromatic. Chromaticism de-emphasizes tonal centering. Since tonal centering is the basis of Western musical language, chromaticism seems unstable due to its unpredictability.

Balkwill and Thompson, (1999), "*A Cross-Cultural Investigation of Emotion in Music: Psychophysical and Cultural Clues*" attempts to answer the following question:

> "Can people identify the intended emotion in music from an unfamiliar tonal system? If they can, is their sensitivity to intended emotion associated with perceived changes in psychophysical dimensions of music [defined later as any property of sound that can be perceived independent of musical experience, knowledge, or enculturation]?"

We were very interested in the results to see if our ideas for emotionally tagged music would have characteristics that could be readily identified without cultural constraints. In this small case study, four specific emotions (joy, sadness, anger, and peacefulness) were presented using ragas of India. In this Hindustani system, there is a specific raga or collection of notes for nine individual moods. Participants were asked to rate tempo, rhythmic complexity, melodic complexity, and pitch range in addition to the four emotions.

Balkwill and Thompson give us some preliminary expectations based on other studies. Tempo is most consistently associated with emotional content. A slow pace equates to sadness and, a faster one; joy. Also, simpler melodies with few variations in melodic contour and more repetition are associated with positive and peaceful emotions. Complex melodies with more melodic contour and less repetition are associated with negative anger and sadness. Timbre plays a role as well. Fear and sadness were reported more when expressed by a violin or the human voice. Finally, timpani was associated with anger.

An interesting note they mention is that a narrower pitch range (reduced melodic contour) may be processed as one auditory stream, therefore "easy" to process which may cause positive emotional ratings. This may be linked to Meyer's idea of an instinctive mental or motor response.

The conclusions were that given music that was not culture-specific to the listeners, they were forced to rely on other, psychophysical, cues to perceive emotional content. As predicted, tempo was a strong cue used to successfully identify the ragas intended for joy, sadness, and anger. It did not work with peacefulness. Also, ratings made by expert and non-expert listeners were pretty equal in identifying joy and sadness. The only significant predictor of peacefulness in this study seemed to be timbre. A flute was highly rated as peaceful.

What happens when two real performers are put up against a computer generated performance of a piece is exactly the focus of Clarke and Windsor, (2000), "*Real and Simulated Expression: A Listening Study*". There is no solid conclusion in this paper, other than the matter will need further investigation. They do state that in this study, the simulated performance treated tempo and dynamics as elements that were correlated based on principles of energy and motion. There were minute differences in the way human performers treated repeated notes, both rhythmically and dynamically. Hence, each performance was perceived in different ways by the listeners.

Basic emotions are defined in Juslin (1997), "*Emotional Communication in Music Performance: A Functionalist Perspective and Some Data*", by the following attributes. They have distinct functions that contribute to individual survival. They are found in all cultures and are experienced as unique feeling states. Further, they appear early in the course of human development and are associated with distinct autonomic patterns of physiological cues. Further, he states that most researchers agree on at least four basic emotions: happiness, sadness, anger and fear. In this small study, three guitarists were asked play the same melody five different ways. One was to be without expression. This would correspond to our (null, null) or *neutral* fragments. Two aspects examined were whether or not emotions could be communicated to the listeners, and how the performers' intentions affected expressive cues in the performance (the psychophysical cues studied in Balkwill and Thompson). Like other studies in this commentary, they found that expressions of happiness, sadness, and anger were readily identified by listeners. The fourth emotional state, in this case fear, was a little elusive. Gender and training did not significantly effect ones ability to identify intended emotion. The study suggests that each emotion has certain characteristics as detailed below (the authors point out that in other instruments it is typical to use staccato articulation when expressing anger, but for electric guitarists, they uniformly revert to legato articulation for expressing anger):

	Loud	Quiet	Fast	Slow	Staccato	Legato
Anger	x		x			x
Sadness		x		x		x
Happiness	x		x		x	
Fear		x		x	x	

In Juslin and Madison (1999), "*The Role of Timing Patterns in Recognition of Emotional Expression from Musical Performance*", we quote from the abstract:

> "We gradually removed different acoustic cues (tempo, dynamics, timing, articulation) from piano performances rendered with various intended expressions (anger, sadness, happiness, fear) to see how such manipulation would effect a listener's ability to decode emotional expression. The results show that (a) removing the timing patterns yielded a significant decrease in listeners' decoding . accuracy, (b) timing patterns were by themselves capable of communicating some emotions with accuracy better than chance, (c) timing patterns were less effective in communicating emotions than were tempo and dynamics"

The authors acknowledge the nature of their study as preliminary and in need of further extended study. At any rate, a few hypotheses are put forth. The first is that long and short note durations may be played differently depending on the intended emotion. They found that expressions of happiness were played in shorter note values and patterns in the expressions of sadness were played in longer notes. Secondly, anger and happiness were associated with staccato articulation. This is the first instance we encountered of further distinction between the staccato articulation representing anger and the staccato articulation representing happiness. It was found that the anger expressions used uniform staccato patterns where the happiness ones were more variable depending on the positions within the phrase. It is suggested that more study of this phenomenon needs to be done, as it may be a key component to decoding happiness.

3.2 The Würfelspiel Approach:

We will start by using an 18^{th} century historical idea called The Musicalisches Würfelspiel. In the 1700's, fragments of music could be rapidly prototyped by using a matrix $\mathcal{A}$ of possibilities. We show an abstract version of a typical Musicalisches Würfelspiel matrix in Equation 1. It consists of P rows and three columns. In the first column are placed the opening phrases or nouns; in the third column, are placed the closing phrases or objects; and in the second column, are placed the transitional phrases or verbs. Each phrase consisted of L beats and the composer's duty was to make sure that any opening, transitional and closing (or noun, verb and object) was both viable and pleasing for the musical style that the composer was attempting to achieve.

$$\mathcal{A} = \begin{bmatrix} \text{Opening 0} & \text{Transition 0} & \text{Closing 0} \\ \text{Opening 1} & \text{Transition 1} & \text{Closing 1} \\ \vdots & \vdots & \vdots \\ \text{Opening P-1} & \text{Transition P-1} & \text{Closing P-1} \end{bmatrix} \tag{1}$$

Thus, a musical stream could be formed by concatenating these fragments together: picking the i^{th} Opening, the j^{th} Transition and the k^{th} Closing phrases would form a musical sentence. Since we would get a different musical sentence for each choice of the indices i, j and k (where each index can take on the values 0 to $P-1$), we can label the sentences that are constructed by using the subscript i, j, k as follows:

$$S_{i,j,k} = \text{Opening i} + \text{Transition j} + \text{Closing k}$$

Note that there are P^3 possible musical sentences that can be formed in this manner. If each opening, transition and closing fragment is four beats long, we can build P^3 different twelve beat sentences.

It takes musical talent to create such a The Musicalisches Würfelspiel array, but once created, it can be used in the process of learning fundamental principles of the music compositional process.

3.3 Emotional Music Data Design:

The underlying goal in building each matrix was to remain as basic as possible. We decided to work within a monophonic texture, meaning melody line only. Note values were restricted to quarter notes and half notes in quadruple meter. Quarter rests were also allowed, but used sparingly. All four matrices (neutral, happy, sad, angry) are structurally similar. Each consists of three columns with four fragment choices that are one measure in length.

Any fragment from column one from any of the matrices is designed to function as an opening phrase. We define an opening phrase as one that clearly establishes a tonal center. In western tonal music, the tonal focal point can be narrowed to a single tone/pitch that is known as the tonic. We have made C our tonic note in all cases. In all but one case, the opening fragment also established the mode as either major or minor. The exception is made in the angry matrix, where ambiguity is desirable.

All fragments in column two of any of the matrices are designed to function as a transition phrase. As the label implies, these transition phrases serve as connectors between a choice from column one and a choice from column three. It is in these middle phrases that movement away from the tonic is made or continued. This movement is necessary for forward progress of a melody. Therefore, each transition phrase is now highlighting a secondary pitch, one other than the tonic note established by the opening phrase.

To close our melodic lines, an ending phrase is chosen. Any fragment from column three of any of the matrices will function in the same manner. We designed each to move back to the tonic note in such a way as to produce a quality of closure to our melodic lines. This was done by approaching tonic in the most basic of ways. Using stepwise motion up or down to the tonic logically ends the melodic journey by bringing you back to the home pitch. An alternative is to return to tonic via melodic skip from the third or fifth scale degree. In the key of C, the tonic (C) is scale degree 1, D 2, etc. So, a melodic skip from the third of fifth scale degree means a skip from an E or a G note. Together with the tonic note, third and fifth scale degrees make up a tonic chord (harmony). By using either the third or the fifth to lead back to tonic, we again produce a sense of closure by reinstating tonic as the final destination of our brief melodic journey.

We used the following guidelines to design our Würfelspiel matrices of different emotional slants. We will begin by stating our emotionally neutral design choices, which are thoroughly discussed in (Dzuris and Peterson, 9, and then discuss the other emotional choices for our data.

To produce *emotion-deprived* or *neutral* fragments, individual characteristics that have been documented by researchers as being contributing factors of basic emotion in music have been neutralized to the best of our abilities. Some of the contributing factors, such as mode, are also essential in a plausible melody, and could not be removed. The use of major mode is the default with a tempo of ♩ = 45, slow to the point that the individual notes outweigh the overall sense of a melody. The rational comes from reading. The telltale sign of a new reader is the slow pace during which equal emphasis is placed on every word. A beginner will produce an emotion-deprived reading. The same result is our goal here. Further, we use even rhythms and exact note durations with the melody played on a basic computer generated sound

Likewise, fragments we intended to emotionally tag as *happy* had individual characteristics researched by others incorporated into the design. Each characteristic was chosen based on a general consensus by other researchers and authors as being a contributing factor in representing "happy" in music. Again, this entails choosing a major mode, a very quick tempo (♩ = 250) and the use of staccato. If we wish, we can use quarter rests. We could choose to present the melodies using a flute, as this particular instrument has often been linked to happiness in the literature.

Once more, individual characteristics that have been researched by others were incorporated into the design of our fragments tagged by *sad*. Each characteristic was chosen based on a general consensus by other researchers and authors as being a contributing factor in representing "sad" in music. There is a use of minor mode with a slow tempo (♩ = 70). Also, we use slurs and legato and the bass clef to put us in a lower register. We choose to present these melodies using a violin, as stringed instruments are particularly lined to sadness in the literature. Finally, there is some use of chromaticism.

To emotionally tag the fragments as *angry*, individual characteristics that have been researched by others were incorporated into the design. Each characteristic was chosen based on a general consensus by other researchers and authors as being a contributing factor in representing "angry" in music. We use a minor mode, a moderate tempo (♩ = 180) faster than used for the sad melodies, slightly slower than the tempo used for the happy melodies and increased variation of articulation (slurs, accents). Further, there is the incorporation of larger leaps. We choose to present these melodies using a trumpet, as brass and percussion are often linked to anger in the literature. There are also more repeated notes and the use of an ambiguous fragment where the mode is not clearly established in opening phrase.

3.4 *Happy* Musical Data:

Following the outline above, we have designed the *Happy* musical data as shown in Figure 3.

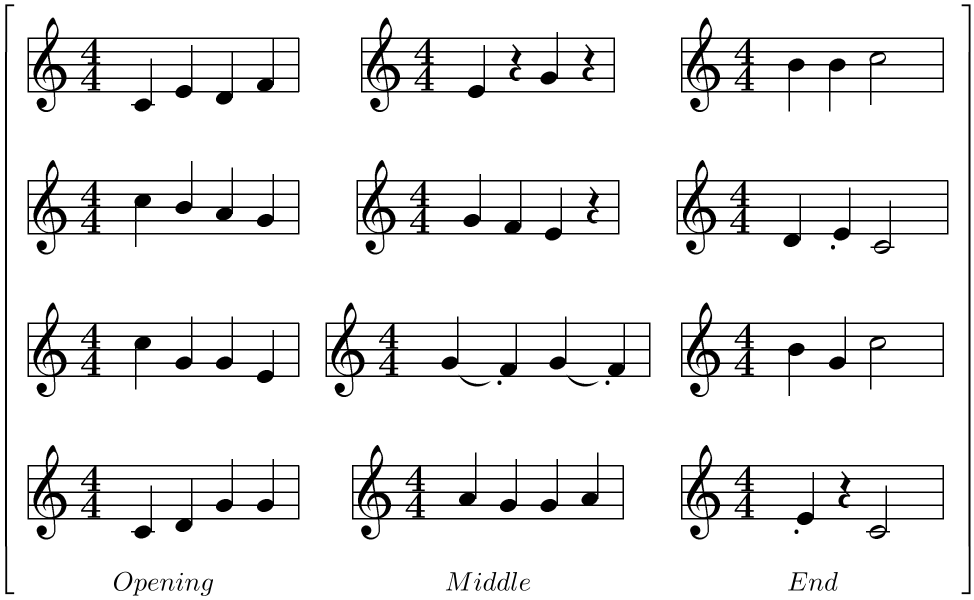

Figure 3: The Happy Music Matrix: To emotionally tag these fragments as ”happy,” individual characteristics that have been researched by others were incorporated into the design. Each characteristic was chosen based on a general consensus by other researchers and authors as being a contributing factor in representing ”happy” in music. These include the use of major mode, a very quick tempo, use of staccato, use of quarter rests, melody played on a flute

From our 4 × 3 Musicalisches Würfelspiel matrix, we can generate 64 musical selections of twelve beats each. In Figure 4, we show the selections generated using opening one, all the possible middle phrases and the first cadence phrase.

Then we show all the selections for opening two, verb one and all the closings in Figure 5. This still only shows eight of the sixty four possible pieces, of course. We invite you to sit at the piano and see how they sound. You should hear how that they are distinctly happy. Our interpretation of the *core* meaning of a *happy* musical fragment is based on ideas from two separate disciplines. First, our reading of the relevant literature has given us guidance into the choices for notes, tempo and playing style as outlined above; and second, the psychophysiological studies of Lang et. al. (1998) as outlined in many papers has given us a pseudo-quantitative measure of the affective content on an emotionally charged image. Musical studies have shown that if a composer deliberately attempts to convey a given emotional content in their music, queries of their audience show that the desired emotional flags have been set.

Figure 4: Happy musical fragments that have been generated using the Happy Musical Matrix using the first opening phrase, all the middle phrases and the first ending phrase. The first column of the figure provides a label of the form xyz where x indicates the opening used; y, the middle phrase used; and z, the ending phrase. Thus, 131 is the fragment built from the first opening,the third middle and the first ending.

3.5 *Sad* Musical Data:

Now, we move toward the design of musical data that is intended to be sad. In Figure 6, you can see the musical data that we designed to have an overall tone of sadness.

From our 4 × 3 Musicalisches Würfelspiel matrix, we can generate 64 musical selections of twelve beats each. In Figure 7, we show the selections generated using opening one, all the possible middle phrases and the first cadence phrase.

Then we show all the selections for all the openings, the first middle phrase and the second ending in Figure 8. This still only shows eight of the sixty four possible pieces, of course. We invite you to sit at the piano and see how they sound. You should hear a sense of sadness in each.

3.6 *Angry* Musical Data:

From the angry 4 × 3 Musicalisches Würfelspiel matrix, as usual, we can generate 64 musical selections of twelve beats each. In Figure 10, we show the selections generated using opening one, all the possible middle phrases and the first cadence phrase.

Figure 5: Happy musical fragments that have been generated using the Happy Musical Matrix using the second opening phrase, the first middle phrase and all the ending phrases. The first column of the figure provides a label of the form *xyz* where *x* indicates the opening used; *y*, the middle phrase used; and *z*, the ending phrase. Thus, 213 is the fragment built from the second opening,the first middle and the third ending.

Then we show all the selections for opening two, the fourth middle and all the endings in Figure 11. This still only shows eight of the sixty four possible pieces, of course. We invite you to sit at the piano and see how they sound. You should be able to hear a sense of anger in each selection.

3.7 Alphabet Selection:

As you have seen in Sections 3.4 to 3.6, the emotionally tagged musical data uses a richer set of notes and *articulation* attached to the notes to construct grammatical objects. We can think of the added articulation as punctuation marks. *Slurs* (one note and multiple note), *staccato* and *marcato* accents are attached to various notes in our examples to add emotional quality. Our design alphabet can be encoded as $\mathcal{H} = \{c,\ d,\ e,\ f,\ g,\ a,\ b,\ r\}$ where each note in this alphabet is now thought of as a *musical object* with a set of defining characteristics. Here r is rest. For our purposes, the attributes of a note are choices from a small set of possibilities from the list $\mathcal{A} = \{p,\ b,\ s,\ a\}$. The index p indicates what pitch we are using for the note: -1, denotes the first octave of pitches below middle C; 0, the pitches of the middle C octave and 1, the first octave of pitches above middle C. The letter b tells us how many beats the note is held. The length of the

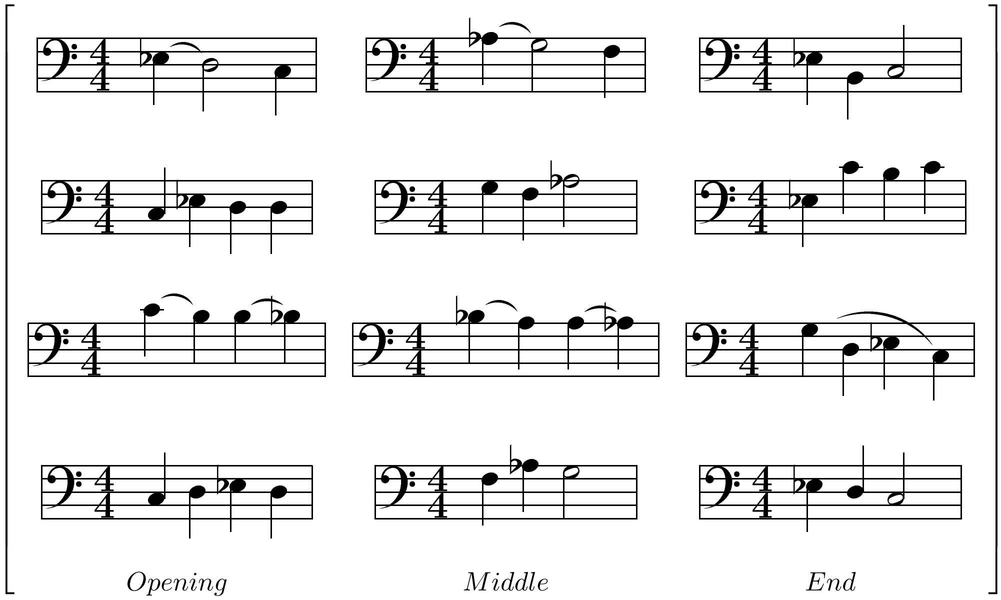

Figure 6: The Sad Music Matrix: To emotionally tag these fragments as ”sad,” individual characteristics that have been researched by others were incorporated into the design. Each characteristic was chosen based on a general consensus by other researchers and authors as being a contributing factor in representing ”sad” in music. These include the use of minor mode, a slow tempo the use of slurs and legato, the use of bass clef to put us in a lower register, the melody played by a violin

slur is given by the value of s and a denotes the type of articulation used on the note. We choose to treat slurs as entities which are separate from the other accent markings for clarity. For these examples, we have slurs that range from zero to three in length, so permissible values of s are taken from the set {0, 1, 2, 3}. This could easily be extended to longer slurs. The beat value b is either one of two as only quarter and half notes are used. There are many possible articulations. An expanded list, for marks either above or below a note for effect, might include *neutral*, no punctuation ($a = 0$); *pizzicato*, a dot ($a = 1$); *marcato* or *sforzando*, a > ($a = 2$); *staccato* or *portato*, a −, ($a = 3$); *strong pizzicato*, an apostrophe ($a = 4$); and *sforzato*, a ˆ ($a = 5$). A given note n is thus a collection which can be denoted by $n_{p,b,s,a}$ where the attributes take on any of there allowable values. A few examples will help sort out this out. The symbol $d_{1,2,2,1}$ is the half note d above middle d with a pizzicato articulation which is the start of a two note slur which ends on the second note that follows this middle d. The rest does not have pitch, articulation or slurring. Thus, we set the value of pitch, slurring and articulation to 0 and use the notation $r_{0,1,0,0}$ or $r_{0,2,0,0}$ to indicate a quarter or half rest, respectively. Our alphabet is thus $\{\mathcal{H}$ which has cardinality 8. Each letter has a finite set of associated attributes and each opening, middle or closing phrase is thus a sequence of 4 musical

Figure 7: Sad musical fragments that have been generated using the Sad Musical Matrix using the first opening phrase, all the middle phrases and the first ending phrase. The first column of the figure provides a label of the form xyz where x indicates the opening used; y, the middle phrase used; and z, the ending phrase. Thus, 131 is the fragment built from the first opening,the third middle and the first ending.

entities. Within this alphabet, an angry middle phrase such as shown in Figure 12(a), can be encoded as $\{e_{0,1,0,2},\ d_{0,1,0,2},\ c_{0,2,0,0}\}$. This would be then written as a matrix

$$
n_1 \quad = \quad \left[\begin{array}{|c|c|c|c|c|c|c|c|}
\hline
0 & 0 & \{0,1,0,2\} & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \{0,1,0,2\} & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
\{0,2,0,0\} & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
\end{array}\right]
$$

If we had a fragment with a slur such as shown in Figure 12(b), this would be encoded as $\{c_{0,1,1,0}, d_{0,1,0,0}, d_{0,1,0,2}, c_{0,1,0,0}\}$. In matrix form, we have

$$
n_1 \quad = \quad \left[\begin{array}{|c|c|c|c|c|c|c|c|}
\hline
\{0,1,1,0\}0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \{0,1,0,0\} & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \{0,1,0,0\} & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
\{0,1,0,0\}0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
\end{array}\right]
$$

Figure 8: Sad musical fragments that have been generated using the Sad Musical Matrix using the first opening phrase, all the middle phrases and the first ending phrase. The first column of the figure provides a label of the form *xyz* where *x* indicates the opening used; *y*, the middle phrase used; and *z*, the ending phrase. Thus, 131 is the fragment built from the first opening,the third middle and the first ending.

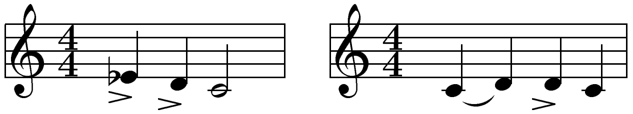

(a) Middle Phrase (b) Opening Phrase

Figure 12: Some Angry Phrases

These matrices indicate which musical object is used in a sequence. The four opening phrases in a Würfelspiel music matrix can thus be encoded into matrices that are 2×8 (both notes in the phrase are half notes) to 4×8 (all notes are quarter notes). Each of these matrices has the special property that a row can only have one nonzero entry. A given middle phrase will have a similar structure, making only some of the possible middle phrase matrices acceptable.

Note that encoding music in this way generates a compact data representation. However, we need to model the data so that each possible musical entry is encoded in a unique way. The first seven entities in $\mathcal{H}$

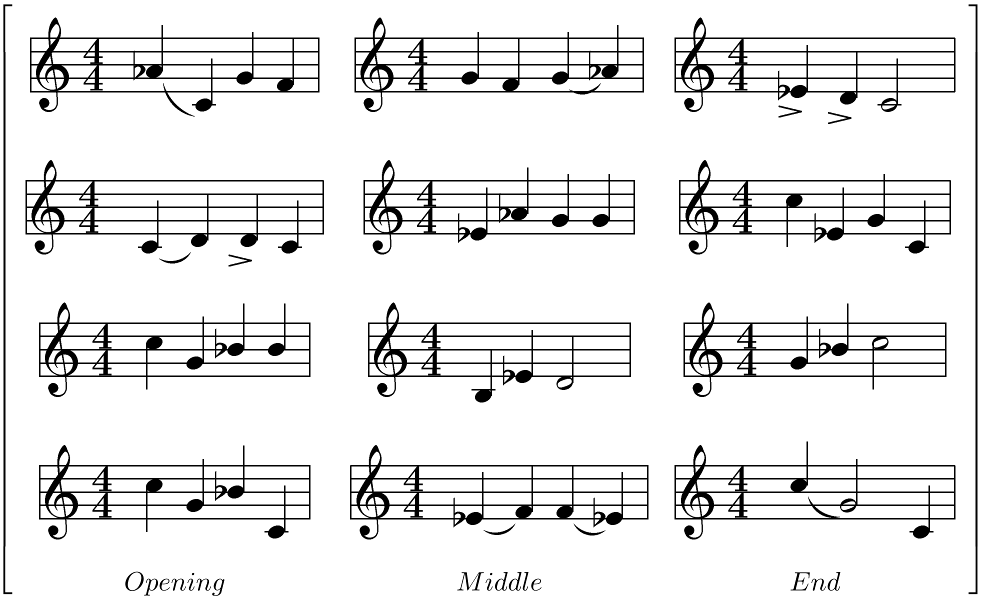

Figure 9: The Angry Music Matrix: To emotionally tag these fragments as "angry," individual characteristics that have been researched by others were incorporated into the design. Each characteristic was chosen based on a general consensus by other researchers and authors as being a contributing factor in representing "angry" in music. These include the use of minor mode, a moderate tempo (faster than used for the sad melodies (slightly slower than the tempo used for the happy melodies), increased variation of articulation (slurs, accents), incorporation of larger leaps, melody played by a trumpet, more repeated notes and ambiguous fragment (mode not clearly established in opening phrase)

come in a total of 144 distinct states: three pitches, two beats, four slur lengths and six articulations. The rest comes in only two states. Hence, a distinct alphabet here has cardinality $7 \times 144 + 2$ or 1010. The size of this alphabet precludes showing an example as we did with the compact representation, but the matrices that encode these musical samples still possess the property that each row has a single 1. Of course, with an alphabet this large in size, we typically do not use a standard matrix representation; instead, we use sparse matrix or linked list techniques. The data representation we used in the neutral data paper, (Dzuris, et al., 2003a), had a much lower cardinality because the neutral data was substantially simpler.

4 Cognitive Modeling:

We are developing models of abstract compositional design in the three separate domains of music composition, painting composition and job scheduling. All three share similar structure despite great differences in culture and background. We use variants of the historical compositional prototyping mechanism from the

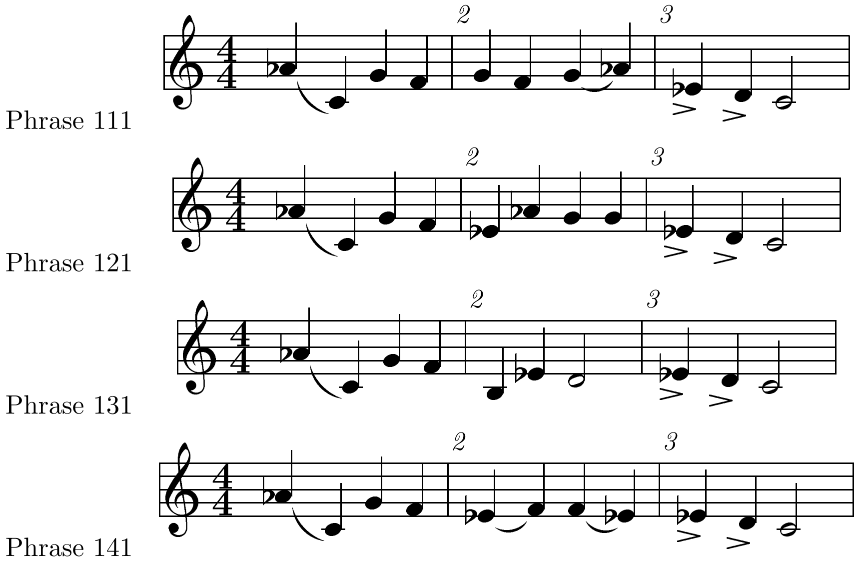

Figure 10: Angry musical fragments that have been generated using the Angry Musical Matrix using the first opening phrase, all the middle phrases and the first ending phrase. The first column of the figure provides a label of the form *xyz* where *x* indicates the opening used; *y*, the middle phrase used; and *z*, the ending phrase. Thus, 131 is the fragment built from the first opening,the third middle and the first ending.

18^{th} century known as Toss of the Dice (Würfelspiel) matrices to construct training data for examples of good compositional design in which any particular sample of training data is equally valid as a choice. This can be done in a number of emotionally distinct flavors following the affective image psychophysiological literature although the notion of emotional attribute must be extended to more general notions of intangibles in job scheduling using genetic algorithms.

Our models of cognitive processing are based on distributed models of computation in which ensembles of interacting computational modules are linked to create larger functional units. The linked units are then used to model cognitive functions such as emotions and, more importantly, disturbances in emotional processing.

Cognitive models are built by finding an appropriate level of abstraction for known cellular biological and neurobiological data that allows us to find generic principles of biological information processing. One of the hardest problems we face in our attempt to develop software models of a high level cognitive process is validation. We mean validate in this sense: the model should behave in the way we expect the cognitive function to behave. This is not generally the way a cognitive model is evaluated. If one builds a model of memory (hippocampus) with a lot of low level biological detail, one might have 10000 neurons, each in several interconnected modules. We know that if you take a slice through the hippocampus, we can get that

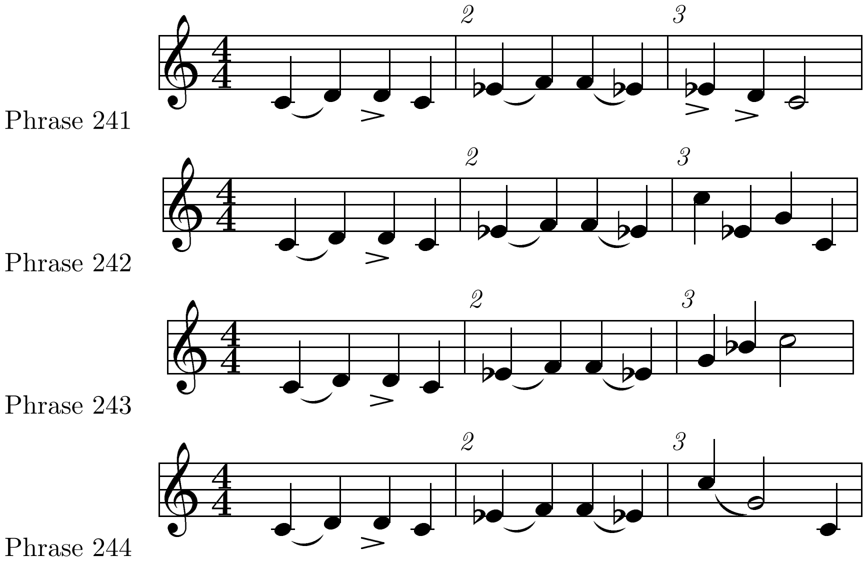

Figure 11: Angry musical fragments that have been generated using the Angry Musical Matrix using the second opening phrase, the fourth middle phrases and all the ending phrases. The first column of the figure provides a label of the form *xyz* where *x* indicates the opening used; *y*, the middle phrase used; and *z*, the ending phrase. Thus, 241 is the fragment built from the second opening, the fourth middle and the first ending.

slice to live in the artificial environment of a petri dish for some time. Moreover, we can take measurements from an ensemble of reading points in that tissue – perhaps 100 or more – and get voltage and/ or current maps of that local patch of tissue. One type of validation is to get the model just built to match these maps. However, there is a lot of debate as to what this means. Those sorts of local readings are clearly not what that complicated system of interrelated modules actually does. There will always be questions as to whether or not matching these readings really is validation of the model. There is a large and unexplained gap between validation of the sort proposed by matching voltage and current maps to validating that the computations performed by the modules are to the external human witness similar to what we would see in a real human performing tasks. Further, validation of the type performed by the human witness is usually done from the perspective of psychology or psychobiology, but most of these validations are more qualitative than quantitative.

To address this problem, we decided that in addition to developing cognitive models that are measured in traditional ways, we would also develop models that can be assessed by experts in other fields for validity. There are three such models we are attempting to build: one is a model of music composition in which short stanzas of music that are emotionally colored or tagged are generated autonomously; second, a model of painting composition in which primitive scenes comprised of background, foreground and primary visual elements are generated autonomously that also have emotional attributes; and third, a model of job scheduling design using genetic algorithms in which the optimal solutions are colored by intangibles which are similar to emotional attributes and could therefore be classified as meta level outputs of an emotional subsystem. This paper will focus on the important task of designing emotionally tagged musical data. In other papers that are in preparation, we address neutral data generation for painting (Peterson & Dzuris, 2003c) and job scheduling (Kurz & Peterson, 2003e). The task of generating neutral musical data is presented in (Dzuris & Peterson, 2003a) while the problems of generating emotionally labeled paintings and the intangibles of job scheduling are presented in (Peterson & Dzuris, 2003d) and (Kurz & Peterson, 2003f), respectively.

4.1 Clues To Our Models From The Literature:

Clues as to how to set up these abstract models can be found in a variety of sources from the open literature. The evolution of nervous systems and in particular, the large order structures of the brain are very informative. Discussions in Kornack, (2000), "*Neurogenesis and the Evolution of Cortical Diversity: Mode, Tempo, and Partitioning during Development and Persistence in Adulthood*", Redies and Puelles (2001), "*Modularity in vertebrate brain development and evolution*", and Catania, (2000), "*Cortical Organization in Insectovoria: The Parallel Evolution of the Sensory Periphery and the Brain*", help us understand the modularity of the underlying neural structures for information processing. This has influenced our software design by helping us determine the minimal module architecture which will allow interesting cognitive response. These issues are also explored in Coltheart, (1999), "*Modularity and Cognition*" and Deacon, (1990), "*Rethinking Mammalian Brain Evolution*".

Motivations for the modeling of the core emotional and cognitive dysfunction engines are varied. A proper model of the computational outputs we interpret as emotional states or qualia requires a concomitant model of how cognitive processes develop flaws. Hence, an understanding of models of cognitive dysfunction such as depression is closely connected to models of emotional processing. Key resources include those that are neurophysiological in nature such as detailed in Davis and Lee, (1998), "*Fear and Anxiety: Possible Roles of the Amygdala and Bed Nucleus of the Stria Terminalis*" and Deadwyler and Hampson, (1997), "*The Significance of neural Ensemble Codes During Behavior and Cognition*". Some of the actual circuitry that may permit emotional responses to be generated are discussed in Drevets, (2000), "*Neuroanatomical Circuits in Depression: Implications for Treatment Mechanisms*". These potential neural architectures lead us to ask whether or not there are specialized places within the brain which correspond to specific emotions and/ or cognitive states. For example, there has been a long standing debate about whether or not language capability is hardwired into neural circuitry or it is molded from generic neural wiring via subjective experience. We gain some insight into these questions by doing neuroimaging studies such as are detailed in Horwitz, Tagamets and McIntosh (1999), "*Neural modeling, functional brain imaging, and cognition*". These techniques are used to target questions relevant to cognitive dysfunction in Honey, Fletcher and Bullmore (2002), "*Functional brain mapping of psychopathology*" and Drevets, (2000), "*Neuroimaging Studies of Mood Disorders*". It is also important to realize that everything we know about neural architecture is based on indirect evidence via the analysis of key substances that are measurable in some way. Some of these tools are discussed in Kobbert et al. (2000), "*Current Concepts in Neuroanatomical Tracing*". Hence, the quest to find neural correlates for cognitive functions is quite hard. These experiments then help us to discuss intelligently the issue of neural correlates as detailed in Frith, Perry and Lumer (1999), "*The neural correlates of conscious experience: and experimental framework*". There are many questions still, of course. We have difficulty understanding the data that the imaging studies show us as is shown in Heller and Nitschke (1998), "*The Puzzle of Regional Brain Activity in Depression and Anxiety: The Importance of Subtypes and Comorbidity*".

There is also a rich literature in the field of psychophysiology which provides a way to assign quantitative measures to certain kinds of emotional or affective outputs. The psychophysiological literature includes acoustical and visual studies. The acoustic studies concern the physiological responses of listeners to certain specific auditory tones or probes. This is not a response to music per se, but it gives valuable data as to normal human responses. These studies include Bradley and Lang, (2000), "*Affective reactions to acoustic stimuli*" and Cuthbert, Schupp, Bradley, McManis and Lang (1998), "*Probing affective pictures: Attended startle and tone probes*". There are also studies that measure the response of human subjects when briefly exposed to samples of pictures from a carefully selected database of images whose emotional content is varied. For example, the results of these studies are included in Codispotti, Bradley and Lang, (2001), "*Affective reactions to briefly presented pictures*"; and Junghöffer, Bradley, Elbert and Lang (2001), "*Fleeting images: A new look at early emotion discrimination*". In addition, the responses of children to such a database of affective pictures has been examined in McManis, Bradley, Berg, Cuthbert and Lang (2001), "*Emotional reactions in children: Verbal, physiological, and behavioral responses to affective pictures*". We can also examine patients with cognitive dysfunction such as schizophrenia via imaging tools as in Lang, et al. (1998), "*Emotional arousal and activation of the visual cortex: An fMRI analysis*".

With the data obtained from low level and high level experiments as outlined above, a variety of general models of emotion have been presented. These include the popular treatment of LeDoux, (1996), "*The Emotional Brain: The Mysterious Underpinnings of Emotional Life*, and more the more technical discussions of Levenson,(1999), "*The Intrapersonal Functions of Emotion*", and Simons, Detenber, Roedema and Reiss (1999), "*Emotion Processing in three systems: The medium and the message*". An older quantitative view is presented in the classic of Ortony, Clore and Collins (1988), "*The Cognitive Structure of Emotions*". However, Ortony is not based on our current biological understanding of these processes.

The development of software models of emotion must utilize all of this information to varying degrees. Some of the issues that would arise in such software models are discussed in Armony, Servan-Schreiber, Cohen and LeDoux, (1997), "*Computational modeling of emotion: explorations through the anatomy and physiology of fear conditioning*".

There are many models of emotional processing that completely bypass biological structure and constraints and instead simply develop a software architecture whose outputs are seen by viewers as valid emotional states. These usually utilize 3D virtual environments where the emotional output is transmitted to the viewer in the facial expressions of a character. However, there is no real underlying biological model encompassing neural circuitry, neurotransmitters and so forth. These high level cognitive models are typically based on software agent technologies. Some of this research is detailed in the virtual actor studies of Badler, Reich, and Webber (1997), "*Towards Personalities for Animated Agents with Reactive and Planning Behaviors*" and Petta and Trappl (1997), "*Why To Create Personalities for Synthetic Actors*". The use of animated characters is shown in Blumberg, (1997), "*Old Tricks, New Dogs: Ethnology and Interactive Creatures*". Models of the software agent type are presented in Wright, (1997), "*Emotional Agents*".

Finally, there is little in the physiological literature which is directed Towards how a human responds to music. An overview which summarizes some of the relevant physiology is given in Peretz (2001), "*Listen to the brain: a biological perspective on musical emotions*". However, there is little said specifically about musically correlated neural cognates.

4.2 The Basic Model:

Our basic neural model is based on abstractions from neurobiology. A model of isocortex is motivated by recent models of cortical processing outlined in (Raizada & Grossberg, 2003). This article uses clues from visual processing to gain insight into how virgin cortical tissue (isocortex) is wired to allow for its shaping via environmental input. Clues and theoretical models for auditory cortex can then be found in the survey paper of (Merzenich. 2001).

The first layer of auditory cortex is bathed in an environment where sound is chunked or batched into pieces of 200 mS length which is the approximate size of the phonemes of a person's native language. Hence, the first layer of cortex develops circuitry specialized to this time constant. The second layer of cortex then naturally develops a chunk size focus that is substantially larger, perhaps on the order of 1000 mS to 10000 mS. Merzenich details how errors in the imprinting of these cortical layers can lead to cognitive impairments such as dyslexia. As processing is further removed from the auditory cortex via mylenated pathways, additional meta level concepts (tied to even longer time constants) are developed. While it is clear that a form of Hebbian learning is used to set up these circuits, the pitfalls of such learning are discussed clearly in the literature (McClelland. 2001). For example, the inability of adult speakers of Japanese to distinguish the sound of an *ell* and an *r* is indicative of a bias of Hebbian learning that makes it difficult to learn new strategies by unlearning previous paradigms. Hence, we will not use strict hebbian learning protocols; instead, we will model auditory and visual cortex with three layers using modified Grossberg processing with Hebbian learning. Our third layer of cortex is then an abstraction of the additional anatomical layers of cortex as well as appropriate mylenated pathways which conduct upper layer processing results to other cognitive modules. We will use the musical data to imprint the first two layers of our model of auditory cortex, the painting data to imprint the first two layers of our visual cortex models and the job scheduling data to enable the imprinting of the third layer meta processing.

The individual neural objects in our cortex will be abstractions of neural ensembles, their behaviors and outputs gleaned from both low level biological processing and high level psychopharmacology data. The low level information must include enough detail of how inputs are processed (spike train generation) to be useful and enough detail of second messenger pathways to see clearly that the interactions between a pre and a post neural ensemble are really communications between their respective genomes. Clearly, this implies that an appropriate abstract second messenger and genome model is needed. Appropriate details on low level modeling can be obtained from many sources. Our current abstract model generates a low dimensional feature vector output rather than an analog action potential pulse (Peterson & Khan, 2003a). This feature vector is shaped by many inputs.

The feature vector output of a neural object is due to the cumulative effect of second messenger signaling to the genome of this object which alters the action potential and thus feature vector of the object by altering its complicated mixture of ligand and voltage activated ion gates, enzymes and so forth. For example, the G protein-linked receptor superfamily second messenger system would consist of a receptor with seven transmembrane regions with links to G proteins that uses a second messenger system activated by an enzyme – cAMP and PI and because a second messenger system is used, response to a signal is delayed and hence, these are *slow response* systems. Another family uses receptors with four transmembrane regions. In this family, the ion channel is surrounded by multiple copies of multiple different receptors and ion flow is directly controlled by a given particular mixture of neurotransmitters and receptors. Clearly, there are many control possibilities that arise here due to the combinatorial nature of this family's channels. It is difficult to find appropriate high level descriptions of such events so that the algorithmic structures are evident and not obscured by the detail. An important resources that have been prime influences and have helped us develop generic models of neural objects which have feature vector outputs which can be shaped by such pharmacological inputs is Stahl (2000).

5 Data Set Design:

The emotionally tagged music data set contains examples of equally valid solutions to a compositional design process. In other papers, we discuss musical data that is neutrally tagged (Dzuris et al., 2003a), neutral and emotionally labeled paintings (Peterson et al., 2003f and 2003g) and job scheduling data (Kurz et al., 2003b and 2003c). Each of these papers discuss the generation of 64 examples of solutions to compositional design tasks that are equally acceptable according to some measure. So, can we *learn* from this kind of data how the experts that designed these data samples did their job? Essentially, buried in these data sets are important *clues* about what makes a great design. How do we begin to understand the underlying compositional design principles?

Each data set that is encoded into a Würfelspiel matrix, whether using music, art or abstract optimization languages, therefore contains crucial information about equally valid examples of data in different emotional/ intangible modalities. From this data, we can build mappings that tell us which sequences of choices are valid and which are not from the perspective of the expert who has designed the examples. In a very real sense, a cognitive model built from this data is automatically partially validated from a psychological point of view.

Each data set has an associated alphabet with which we can express noun, verb and object units. For our purposes, let's say that each of the nouns, verbs or objects is a finite list of actions from an alphabet of R symbols, where the meaning of the symbols is of course highly dependent on the context of our example. In a simple music example, the list of actions might be a sequence of four beats and the alphabet could be the choices from the C major scale. Thus, we will assume each of the P nouns consists of a list of length

L from an alphabet of R symbols. Since the general alphabet has R symbols, each element of a noun can be thought of as a vector of size R whose components are all 0 except for a single 1. Let the letters in the alphabet be denoted by a_0 through a_{R-1}. Then a noun has components 0 to $L-1$ where component $n[j]$ is the letter a^n_j. The letter a^n_j is then encoded into a vector of length R all of whose entries are 0 except for the entry in slot j.

5.1 Inhibition and Excitation:

The raw inputs we see as the noun vector n are thus normally processed into a specialized vector for use in algorithms that model the compositional process. For example, in the generation of music, we write an opening phrase in standard notation (the kind we would see on sheet music). This is the raw input n. We can choose to encode this data in a form more amenable to computation and data processing in many ways. In our work, we have encoded the raw data into an abstract grammar, thereby generating the feature vector N. In general, the input nouns are preprocessed to create output noun states denoted by N. In a similar fashion, we would preprocess verb and object inputs to create verb and object output states denoted by V and O, respectively. The preprocessing is carried out by mappings f_n, f_v and f_o, respectively. For example, the mapping from raw data to the feature vector form for nouns is represented by Equation 2:

$$\begin{bmatrix} a^n_0 \\ \vdots \\ a^n_{L-1} \end{bmatrix} \rightarrow_{f_n} \begin{bmatrix} N_0 \\ N_1 \\ \vdots \\ N_{L-1} \end{bmatrix}$$

Our association of a noun n with the feature vector N is thus but one example of the mapping f_n. Now, a primitive object in our purported compositional grammar has length L. We can denote such an input noun object as n_i and a corresponding output noun object as N_i. Our mapping problem is thus to determine the rule behind the mapping from the noun feature vectors N to the verb feature vectors V, g_{NV}, and from the verb feature vectors V to the object feature vectors O, g_{VO}. We can express this mathematically by Equation 2:

$$\begin{bmatrix} N_0 \\ N_1 \\ \vdots \\ N_{L-1} \end{bmatrix} \rightarrow_{g_{NV}} \begin{bmatrix} V_0 \\ V_1 \\ \vdots \\ V_{L-1} \end{bmatrix} \text{ and } \begin{bmatrix} V_0 \\ V_1 \\ \vdots \\ V_{L-1} \end{bmatrix} \rightarrow_{g_{VO}} \begin{bmatrix} O_0 \\ O_1 \\ \vdots \\ O_{L-1} \end{bmatrix} \tag{2}$$

We can combine these processing steps into the diagram shown in Figure 13.

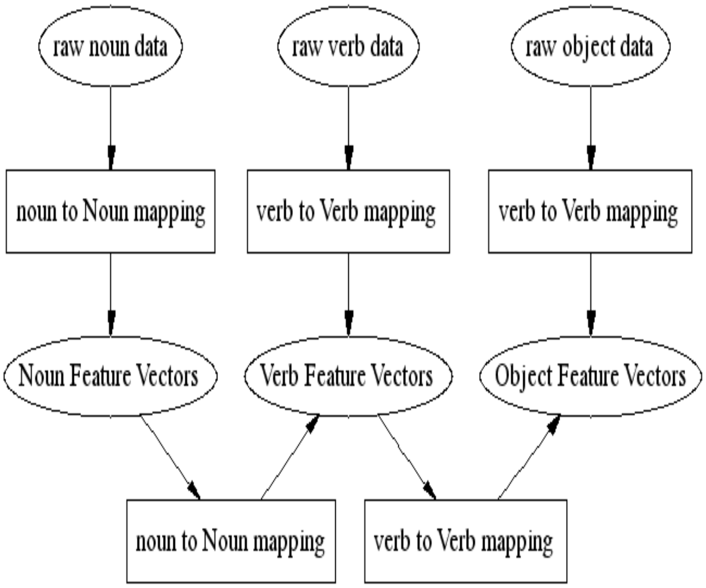

Figure 13: Raw Sentence to Feature Vector Processing

5.2 Noun to Verb Processing:

The data we are given is order dependent. For example, if we are given a noun, n_i of form $\{a^n_{i0}, \ldots, a^n_{i,L-1}\}$, then we attend to the letters of this noun sequentially as $a^n_{i0} \rightarrow a^n_{i1} \rightarrow \ldots \rightarrow a^n_{i,L-1}$. We are given that each noun n_i is associated with a set of possible verbs, $\{v_j\}$ equal to $\{a^v_{j0}, \ldots, a^v_{j,L-1}\}$ for $0 \leq j < P-1$ and one task is thus to understand the noun to verb mapping. However, another task is to understand how to generate the original noun sequence. Why are some noun sequences useful or pleasing in this context and others are not? To generate a noun sequence n equal to $\{a^n_0, \ldots, a^n_{L-1}\}$ means we choose a random start letter a^n_0 and then from that preferred sequences are generated while non-interesting words are biased against. Hence, we think of a mapping, the Noun Generating Map or NGM as accepting an input, a^n_0 and generating a preferred second note a^n_1. Then a^n_1 is used as an input to generate a preferred third note, a^n_2 and so on until the full string of letters is finished. To model this mapping, we start by using the information about useful noun strings we have. Given letter a^n_{i0}, we know that a^n_{i1} is preferred.

We embed the original data into an analog vector by converting each letter a^n_{i0} of the noun, which is a 0 or a 1 in our initial encoding, into a real number ξ^n_{i0}. To set the value of the real number ξ^n_{i0}, we choose a tolerance, ϵ, in the interval $(0, 0.25)$ and choose a real number y randomly from the interval $[-0.5\epsilon,\ 0.5\epsilon]$ and then set the value of ξ_{i0} to be $\epsilon \pm y$. Hence, the value of ξ^n_{i0} lies in the interval $[0.5\epsilon,\ 1.5\epsilon]$. For example, if ϵ was chosen to be 0.20, then for all indices in the binary encoding of the letter a^n_{i0} that are 0, we would randomly choose y from $[-0.1, 0.1]$ generating ξ^n_{i0} values that lie in $[0.10, 0.30]$. We will call the number ϵ our analog threshold. The entry in the binary encoded letter that corresponds to a 1 will be randomly chosen from $[1 - 1.5\epsilon,\ 1 - 0.5\epsilon]$. Hence, for $\epsilon = 0.2$, entry with a 1 will be assigned a real number in the interval $[0.7, 0.9]$. Hence, the raw binary noun, verb and object data is mapped into a new analog representation in which each entry is a real number chosen as above.

We know that only certain letters should follow a^n_{i0}. Hence, only certain analog states ξ^n_{i1} are permissible given a start state of ξ^n_{i0}. We infer from this that there is an unknown mapping h^{01} which maps the analog encoding of letter 0 to the analog encoding of letter 1, $h^{01}(\xi^n_{i0}) = \xi^n_{i1}$. This mapping has special characteristics: our data tells us that only certain letters that can follow letter 0. Each acceptable second letter is a vector in R dimensional space whose components are analog zero except the one the corresponds to the second letter. That component is an analog one. We have at most P examples of acceptable second letters. Hence, we have $R - P$ second letters that are not acceptable. In other words, the preferred output for a given noun is a R row matrix formed from the acceptable verbs that has at most P columns.

We can do this for all of the nouns in our data set. Hence, we will have at most P first letter choices and each of these will have at most $R - P$ unacceptable second letters. Let T and T' denote the set of all acceptable and unacceptable outputs respectively. We model the mapping h^{01} as a chained feed forward network, Peterson (1998), with feedforward and feedback connections between artificial neurons. This mapping takes the first letter of a noun and outputs a set of acceptable second letters. Training is done by matching input to output using excitation and inhibition. We know which elements in the analog output vector should be close to one and which should be close to zero for the analog input. We initialize all of the tuneable parameters to be small positive if the connection from component k in the input to component j on the output is between two analog ones. All other connections are initialized to small negative numbers. For each first letter we have data for, we do the following: pick the initial first letter in our data set and compute the relevant output for the first associated second letter. Increase the connection weights on any path between a high input and a high output and decrease the connection weight on any other paths. Cycle to the next second letter and redo until all possibilities are exhausted. We thus continue this process until every input generates an output with a high component value in the location that corresponds to the index for the second letter. At this point, we say we have trained our nonlinear mapping h^{01} so that first letters in our data noun sequences are biased to connect to their corresponding second letters. A second letter is then chosen randomly from the set of acceptable second letters via an additional input line which is a sense is a coarse model of *creativity.*

If we let the set of all the generated weights be the matrix W^{01}, we note this is an $R \times PR$ size matrix. We can develop a similar mapping for the second to third letter, h^{12} with weights W^{12}, the third to fourth letter, h^{23} with weights W^{23}, and finally, the mapping from letter $L-1$ to letter L, $h^{L-2,L-1}$ with weights $W^{L-2,L-1}$.

The procedure for creating a valid noun sequence can now be given. Choose a valid starting letter for a noun, a^n_{i0} and we map it to its analog form, ξ^n_{i0}. Then, applying the first to second letter map, we find an acceptable second letter by the computation $h^{01}(\xi^n_{i0}) = \{\xi^n_{i1}\}$. This second letter can be used as the input into the next map, generating an acceptable third letter. Hence, the composite map, $h^{12}h^{01}$ takes a valid first letter and creates the three letter analog sequence defined by Equation 5.2:

$$\begin{bmatrix} \xi^n_{i0} \\ \xi^n_{i1} \\ \xi^n_{i2} \end{bmatrix} \in \begin{bmatrix} \xi^n_{i0} \\ h^{01}(\xi^n_{i0}) \\ h^{12}(\, h^{01}(\xi^n_{i0})) \end{bmatrix} \tag{3}$$

The analog sequences are then mapping into three letter sequences by assigning an analog value to either a one or zero using a threshold tolerance τ. This means we map a component whose value is above τ to 1 and one whose value is below τ to 0. This can of course generate invalid sequences as we are only supposed to have a single 1 assigned from any analog sequence. We do have to make sure that our developed map does not allow this. For example, for $\tau = .6$, the vector

$$\begin{bmatrix} 0.83 \\ 0.55 \\ 0.35 \end{bmatrix} \rightarrow \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

and the example below generates an invalid binary sequence:

$$\begin{bmatrix} 0.83 \\ 0.55 \\ 0.35 \end{bmatrix} \rightarrow \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$$

which we would not know how to interpret as part of a noun. Nevertheless, despite these obvious caveats, the procedure above learns how to generate all acceptable three letter nouns given an initial start letter.

There will be at most P^2 possible second and third letters in this set. Since this set of possibilities will grow rapidly, after generating the letter two set of possibilities, we randomly choose one of the columns of the letter two matrix as the second letter choice and apply the h^{12} mapping to that letter. We then randomly choose one of the columns of the letter three matrix as the third letter choice.

We then extend this procedure to the generation of all P letters by the concatenation $h^{P-2,P-1} \cdots h^{12}h^{01}$ which we will denote by the symbol H^n, where the superscript indicates this is the mapping we will use for noun sequences. The mapping H^n is the Noun Generating Map or NGM that we seek. It generates a set of P^{P-1} letter two to letter P sequences. By making a random column choice at each letter, we generate one random P letter noun sequence for each initial letter we use. We can do something similar for the verb and object data generating the Verb and Object Generating Maps H^v and H^o, respectively.

These three mappings are the noun, verb and object generator mappings we were seeking. Then we need to connect nouns to verbs and verbs to objects. The mapping from N to V is where the real processing lies. The Würfelspiel matrix training data approach tells us that it is permissible for certain nouns to be linked to certain verbs. While we could memorize a look-up table based on this data, that is not what we wish to do. We want to determine underlying rules behind these associations as emergent behavior in a complex system of interacting agents. Thus, each output noun N_i in the the collection of P nouns $\{N_i : 0 \leq i < P\}$ should activate any of the P verbs $\{V_j : 0 \leq j < P\}$ via the map g_{NV}. Further, each verb V_j in $\{V_j : 0 \leq j < P\}$ should activate the output nouns $\{O_i : 0 \leq i < P\}$ by the action of the map g_{VO}.

To build the mappings g_{NV} and g_{VO}, we will use a more sophisticated model than a simple chained feedforward network with feedback and feedforward connections. Our computational model will consist of four fully connected abstract neural objects capable of receiving input as is shown in Figure 14.

In such a cluster of four nodes, organized in a square fashion there are six possible node to node connections. Each of these neural objects will consist of nine individual abstract neurons as shown in Figure 15. These nine node internal clusters will also be fully interconnected. The output of a given internal neuron is a feature vector whose exact values can be shaped by pharmacological inputs so as to phase lock with either none of the other neurons or some subset of the nine internal nodes. Since there are nine internal neurons, there are many possibilities for phase locked subclusters. Each of the four ensemble nodes has three interconnections to other ensemble nodes. Rather than simply weighting such interconnections with a scalar value as is usual in naive connectionist models, we instead will determine whether a connection is active and its weight based on which phase locked subclusters are available. One such possible mapping can be set up as follows: there are 130 phase locked clusters of size 6 to 9: if such a phase locked cluster

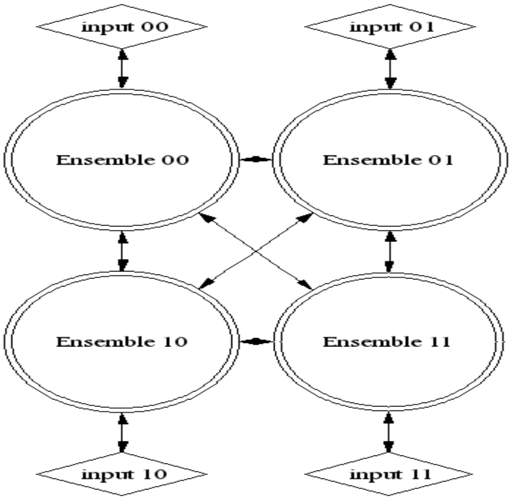

Figure 14: Four Neural Ensemble Structure

is available, this activates connection one; there are 252 phase locked clusters of size 4 and 5; any such cluster activates connection two; there are 130 remaining phase locked clusters: any of these being available activates connection 3. Hence, shared activity chooses activation of a connection pathway and the weight of the ensuing synaptic connection can be set via standard Hebbian protocols.

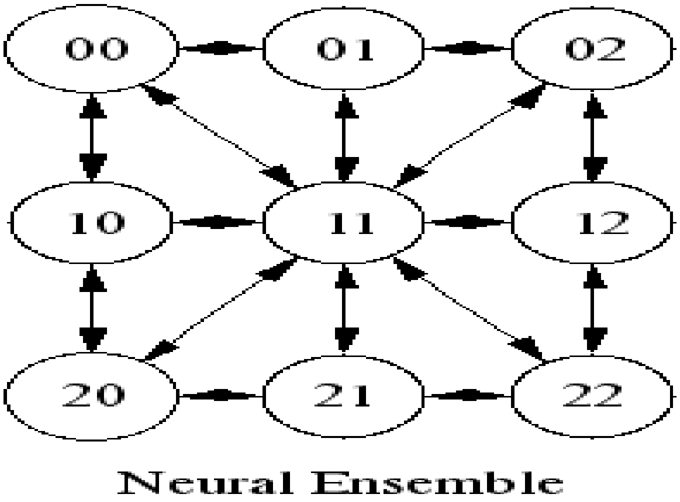

Figure 15: A Typical Neural Ensemble

5.3 Sentence Construction:

There are then two ways to create a valid musical sentence. The first does not use the mappings g_{NV} and g_{VO}. A random choice of starting phrase or noun is chosen to begin the sentence selection process. This input generates a valid noun or opening phrase. Then, a randomly chosen starting note for the middle phrase is then used to generate a valid verb. Finally, a random starting note for the closing phrase generates the final four beat sequence of our valid neutral phrase. The output of the cognitive module is a short sentence of the type we have discussed. The second method is more interesting. The randomly generated noun N generates a valid verb $g_{NV}(N)$ and the valid object is generated by the concatenation $g_{VO}(g_{NV}(N))$. Thus, the composite map $g_{VO}\,g_{NV}$ provides a sentence generator. Note that we can build these mappings for each emotional modality in our musical data.

Once we can generate sentences, we note that we can move to the generation of streams consisting of sentences concatenated to other sentences once we create an Object to Noun mapping. This is done in a way that is similar to what we have done before using a Würfelspiel array approach. Other possibilities then come to mind; for example, in the context of music composition, since key changes are logical transitions, we can create arbitrarily long musical streams punctuated by appropriate key changes by using the Würfelspiel array approach to model which key changes between given key signatures are pleasing between two musical streams.

5.4 The Cognitive Models:

The model we have generated so far, creates valid musical phrases with emotional tags. In other papers, we address the manner in which we generate valid neutral musical phrases and neutral and intangibly tagged optimization algorithms following the design of the triangle shown in 2. While the full discussion of how a cognitive model of emotions is built from this type of data is relegated to other papers (Peterson & Charlesworth, 2003d and 2003e), it is worthwhile to mention the process briefly for completeness. Hence, if we use the superscripts α and β to denote emotional modality and data type, respectively, we can label the mappings in the form $\{g^{\alpha\beta}_{NV}, g^{\alpha\beta}_{VO}\}$. We let $\alpha = 0, 1, 23$ denote neutral, happy (contented), sad (demoralized) and angry (antagonistic) and $\beta = 0,\ 1,\ 2$ indicate music, painting and genetic algorithm optimization. We thus have a collection of mappings

$$\left[\begin{array}{|c|c|c|c|c|}\hline g_{NV}^{00} & g_{VO}^{01} & g_{VO}^{02} & g_{VO}^{03} & \text{music} \\ \hline g_{NV}^{10} & g_{VO}^{11} & g_{VO}^{12} & g_{VO}^{13} & \text{painting} \\ \hline g_{NV}^{20} & g_{VO}^{21} & g_{VO}^{22} & g_{VO}^{23} & \text{GA} \\ \hline \end{array}\right]$$

for the generation of musical, painting and optimization compositional designs.

For each emotional modality, our musical, painting and job scheduling Würfelspiel data provides 64 equally valid data points. Consider the 64 “sad” data points for music. We know as humans that this data is “sad”. Each such “sad” data point provides auditory cortex training data and 64 examples of the “sad” emotional attribute. The painting data gives us 64 examples of visual cortex training data as well as 64 additional “sad” emotional attributes. Hence, we have 128 examples of “sadness” split between the sensory pathways of hearing and vision. In this way, we build 128 examples of each emotional attribute from music and painting split equally between hearing and vision. The intangibly tagged data “demoralized” from job scheduling provides 64 examples of a higher level cognitive output. We know from studies of neural processing, that the front of the auditory and visual cortex is closely aligned to sensory data and as you progress into more interior layers of cortex, neural ensembles begin to respond to progressively higher and more abstract patterns. For example, in the auditory cortex, initially the nerve cells respond to simple phonemes of perhaps 20 mS duration and higher levels are responsive to words, then sentences and so forth. We can make similar analogies to processing in the visual cortex. Outputs from primary sensory cortex are fed into higher level associative cortex where more abstract processing is performed. The intangible “demoralized” tagged data set, then provides 64 examples of a particular type of valid compositional design for the training of the outputs of the higher level associative cortex. Hence, our data provides a validating pathway for two types of primary sensory cortex as well as a primitive model of higher level associative cortex. We can design algorithms to train our cortical tissue models using laminar cortical processing via on-surround excitation/ off-surround inhibition and Hebbian based connection strengthening.

As discussed for emotional data in the context of music, given a random starting note say from column one of a Würfelspiel music matrix of given emotional modality, our model will generate an entire valid musical composition. Note that this output should actually be interpreted as two separate pieces of information: one, as a musical composition and two, as an emotional state. Our data thus provides training for the correct output of a model of emotions as well as we have a total of 128 happy, sad and angry input/ output training samples. In addition, we have 64 input/ output samples each for contented, demoralized and antagonistic emotional states. We can thus generate valid compositional designs that have a specific emotional tag for both the auditory, visual and associative cortex pathways.

We will start with an emotional model which outputs two parameters (loosely based on the psychophysiological data experiments). The first is actually a skin conductance parameter and the second is a complicated computed value that arises from certain fMRI measurements. The interesting thing about these values is that in experiments with human subjects, when people saw pictures with emotional contents such as "sad", "angry" and so forth, the two parameters mentioned above determined a x-y coordinate system in which different emotional attributes were placed experimentally in very different parts of this plane. For example, "sad" images might go to quadrant 2 and "angry" images might be mapped to quadrant 3. This is an over simplification of course, but the idea that images of different emotional attributes would be separate in the plane is powerful. Our 128 examples of each emotional attribute thus give us 128 data points which should all be mapped to the same decision region of this two dimensional plane. Thus, we have data that unambiguously gives us a desired two dimensional output for our emotional model.

The emotional model we will use will be based on four fully connected software agents. Each software agent will consist of nine artificial neurons that accept neuro transmitter inputs that can modulate their outputs in the three time frames we have mentioned (200 mS, 1 - 10 S and days) via some interesting second messenger interactions. The outputs of these neurons can phase lock. Depending on how many phase lock, the output of the software agent can be directed to be high toward another software agent. For example, if anywhere from six to nine neurons are phase locked, this would indicate that the output signal from the software agent would be sent to its neighbor to the east. It is easy to devise many strategies for agent connection that are based on cooperating ensembles of neural activity. This mimics many aspects of what we know about motor and behavior control.

The overall output of the four interacting software agents that model our emotional engine is then trained via Hebbian principles to match the data we have provided. At this point, we have a model that given an abstract auditory and visual input stream from the data will generate both a musical and painting composition and a two dimensional emotional attribute vector. We can then turn the system around if you like, by noting that each data set corresponds to a certain decision region in "emotional space". Further, recognize that we have a coupled model of sensory processing and emotional computation. We have an auditory and visual agent which given input from a known emotion decision region, generates a musical and artistic stream of a given emotional attribute. We model this as a three software agent construction: auditory, visual and emotional. Each of these agents accepts inputs from the others. We have enough data to develop a first pass at a model which given a two dimensional emotional input and a visual and auditory random start, will generate an emotional tagged auditory and visual steam. We can then *validate* this model easily by just letting anyone listen or look at our output and tell us if it is good. Hence, we are validating the whole model instead of just the equivalent of a local patch of hippocampal tissue in a slice. Indeed, it should also be possible to include validation in the learning algorithm.

6 Conclusions:

The design for the generation of emotionally tagged musical data presented here provides us with as much data as we wish for the training of the models of cognition that we have described in outline. We note that the way this data is designed captures primitive elements of good musical compositional design in a convenient matrix form that is ideally suited for use in the training of a variety of artificial neural models that can be used to build larger models of cognition. Further, this data, in conjunction with neutral musical data (Dzuris & Peterson, 2003a), allows us to train both the auditory cortex portion and some of the emotional pathways of the cognitive model. We can leverage these models into two directions: the first, a model of the musical compositional process that is based on fundamental models of cognition and not based on statistical analysis of a composer's work and pastiches such as presented in (Cope. 1991 and 2001); the second, true models of cognition proper in which the musical data provides just one part of the overall model. The data generated from our work with art primitives used to assemble painting compositions, detailed in (Peterson & Peterson, 2003c and 2003d), trains the visual cortex and the data generated from our work with the intangibles associated with complex optimization algorithms trains the higher level associative cortex (Kurz & Peterson, 2003e and 2003f). Together, these are all necessary components to complete our model.

7 Acknowledgments:

This research was partially supported by the National Science Foundation grant DBI 0119171, "*Asynchronous Methods on Heterogeneous Computer Networks for Abstracting High Level Biological Meaning*", from the Division of Biological Infrastructure; Biological Databases and Bioinformatics. In addition, we are grateful for the many conversations on the interface between music, art, mathematics, computer science and cognition that we have had with our research colleagues M. Kurz and S. Charlesworth which have helped shape this work.

8 References:

1. Agnetis, A., Pacifici, A., Rossi, F., Lucertini, M., Nicoletti, S., Nicolo, F., Oriolo, G., Pacciarelli, D.and Pesaro, E. "*Scheduling of flexible flow lines in an automobile assembly plant*". European Journal of Operational Research, Vol. 97, 1997, 348 - 362.

2. Armony, J., D. Servan-Schreiber, J. Cohen and J. LeDoux. "*Computational modeling of emotion: explorations through the anatomy and physiology of fear conditioning*", Trends in Cognitive Sciences, Vol. 1, No. 1, April 1997, 28 - 34.

3. Badler, N., B. Reich and B. Webber, "*Towards Personalities for Animated Agents with Reactive and Planning Behaviors*", in **Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents**, Trappl, R. and P. Petta, eds., Lecture Notes in Artificial Intelligence 1195, Springer-Verlag, 1997, 43 - 57.

4. Bean, J.C. "*Genetic algorithms and random keys for sequencing and optimization*". ORSA Journal on Computing. Vol. 6, 1994, 154 - 160.

5. Catania, K. "*Cortical Organization in Insectovoria: The Parallel Evolution of the Sensory Periphery and the Brain*", Brain, Behavior and Evolution, 2000, Vol. 55, No. 6, 311 - 321.

6. Cope, David. **Virtual Music: Computer Synthesis of Musical Style**, MIT press, 2001.

7. Cuthbert, B., M. Bradley and P. Lang. "*Probing Picture Perception: Activation and Emotion*", Psychophysiology, 33 (1996), 103 - 111.

8. Davis, L. and Streenstrup, M.. "*Genetic algorithms and simulated annealing: an overview*". In **Genetic Algorithms and Simulated Annealing**. Ed. L. Davis. London: Pitman, 1987, 1 - 11.

9. Dzuris, L and Peterson, J. "*Data Abstraction In Cognitive Models for Compositional Design in Music*", Submitted to *Music Perception.*

10. Dzuris, L and Peterson, J. "*Emotionally Tagged Models for Compositional Design in Music: Data Abstraction*", Submitted to *Journal of New Music Research.*

11. Frith, C., R. Perry and E. Lumer. "*The neural correlates of conscious experience: and experimental framework*", Trends in Cognitive Sciences, Vol. 3. No. 3, March 1999, 105 - 114.

12. Holland, J.H. **Adaptation in Natural and Artificial Systems**. Ann Arbor: The University of Michigan Press, 1975.

13. Juslin, P. "*Emotional Communication in Music Performance: A Functionalist Perspective and Some Data*", Music Perception Summer 1997, Vol.14, No.4, 383-418.

14. Kobbert, C., Apps, R., Bechmann, I., Lanciego, J. Mey, J. and Thanos, S. "*Current Concepts in Neuroanatomical Tracing*", In **Progress in Neurobiology**, 62, 2000, 327 - 351.

15. Kurz, Mary E. and Ronald G. Askin. "*Note on An adaptable problem-space-based search method for flexible flow line scheduling*". IIE Transactions, Vol. 33, No.8, August 2001, 691 - 693.

16. Kurz, Mary E. and Ronald G. Askin. "*Scheduling Flexible Flow Lines with Sequence-dependent Setup Times*". In revision for European Journal of Operational Research, submitted July 2001, revisions begun December 2002.

17. Kurz, Mary E. and Ronald G. Askin. "*Real-Time Scheduling of Flexible Flow Lines with Sequence Dependent Setups and Parallel Machines*". International Journal of Production Economics. In Press, accepted February 2003.

18. M. Kurz and Peterson, J. "*Compositional Design in Genetic Algorithms with Intangibles: Data Abstraction*", in preparation.

19. Lang, P., M. Bradley, J. Fitzmmons, B. Cuthbert, J. Scott, B. Moulder and V. Nangia. "*Emotional arousal and activation of the visual cortex: An fMRI analysis*", Psychophysiology, 35 (1998), 199 - 210.

20. Levenson, R. "*The Intrapersonal Functions of Emotion*", Cognition and Emotion, 1999, 13 (5), 481 - 504.

21. McClelland, M. "*Failures to Learn and Their Remediation*", in **Mechanisms of Cognitive Development: Behavioral and Neural Perspectives**, edited by J. McClelland and R. Siegler, Lawrence Erlbaum Associates, Publishers, 2001, 97 - 122.

22. Merzenich, M. "*Cortical Plasticity Contributing to Child Development*", in **Mechanisms of Cognitive Development: Behavioral and Neural Perspectives**, edited by J. McClelland and R. Siegler, Lawrence Erlbaum Associates, Publishers, 2001, 67 - 96.

23. Norman, B.A. and J.C. Bean. "*A genetic algorithm methodology for complex scheduling problems*". Naval Research Logistics. Vol. 46, 1999, 199 - 211.
24. Ortony, A., Clore, G. and Collins, A., (1998). **The Cognitive Structure of Emotions**, Cambridge University Press.
25. Pehlivan, Sertac and Mary E. Kurz. "*Creating Schedules for Flexible Flowlines with Sequence-dependent Setup Times in Random Keys Genetic Algorithms*" Industrial Engineering Research Conference Proceedings 2002, Orlando FL.
26. Peterson, J. (1998). General Feed Forward Architectures: Volume 1: Theory. Available at http://www.ces.clemson.edu/∼/petersj/NeuralCodes/ffntheory.
27. Peterson, J. and Charlesworth, S. "*Multi-Agent Cognitive Models of Compositional Design*", in preparation.
28. Peterson, J. and Charlesworth, S. "*Multi-Agent Cognitive Models of Emotionally Tagged Compositional Design*", in preparation.
29. Peterson, J. and Khan, T. "*Voltage Clamp Action Potential Classification Using Intra-Cellular Models With Diffusion*", with T. Khan, in preparation.
30. Peterson, J. and Khan, T. "*Voltage Clamp Action Potential Classification Using Generalized Diffusion Fit Conductance Models*", with T. Khan, in preparation.
31. Peterson, J. and Khan, T. "*Toxin Recognition From Voltage Clamp Action Potential Classification Using Generalized Diffusion Fit Conductance Models*", with T. Khan, in preparation.
32. Peterson, J. and Dzuris, L. "*Cognitive Models for Compositional Design in Painting: Data Abstraction*", in preparation.
33. Peterson, J. and Dzuris, L. "*Emotionally Tagged Models for Compositional Design in Painting: Data Abstraction*", in preparation.
34. Petta, P. and R. Trappl, "*Why To Create Personalities for Synthetic Actors*", in **Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents**, Trappl, R. and P. Petta, eds., Lecture Notes in Artificial Intelligence 1195, Springer-Verlag, 1997, 1-9.
35. Piramathu

36. Raizada, R and S. Grossberg. "*Towards a Theory of the Liminar Architecture of Cerebral Cortex: Computational Clues from the Visual System*", Cerebral Cortex, Jan. 2003, 100 - 113.
37. Redies, C. and L. Puelles. "*Modularity in vertebrate brain development and evolution*", Bioessays, 23: 2001, 1100 - 1111.
38. Rowe, Robert. **Machine Musicianship**, MIT Press, 2001.
39. Simons, R., B. Detenber, T. Roedema and J. Reiss. "*Emotion Processing in three systems: The medium and the message*", Psychophysiology, 36 (1999), 619 - 627.
40. Sivrikaya-Serifoglu, F. and G. Ulusoy. "*Parallel machine scheduling with earliness and tardiness penalties*". Computers and Operations Research. Vol. 26, 1999, 773-787.
41. Snyder, B. **Music and Memory: An Introduction**, MIT Press, 2000.
42. Snyder and Daskin
43. Stahl, S. **Essential Psychopharmacology: Neuroscientific Basis and Practical Applications**, Second Edition, Cambride University Press, 2000.