Live Soundscape Composition Based on Synthetic Emotions

Klaus C. Wassermann, Kynan Eng, and Paul F.M.J. Verschure
Institute of Neuroinformatics, University of Zurich and Swiss Federal Technical University

Jônatas Manzolli
University of Campinas, Brazil

We conceived *Ada: Intelligent Space* exhibit as an artificial organism, integrating a large number of sensory modalities, and let it interact with visitors using a multitude of effector systems. Ada used a language of sound and light to communicate its moods, emotions, and behaviors. Here we describe the mechanisms behind Ada's sound communication, its real-time performance, and its interpretation by human subjects.

Imagine communicating with an artificial being that expresses synthetic emotions based on how you interact with it. This is just what visitors experienced with *Ada: Intelligent Space*, an exhibit presented at the Swiss National Exhibition Expo.02. Ada became operational on 12 May 2002 and catered to 553,700 visitors until 20 October 2002. Ada—named after the 18th-century computer pioneer Lady Ada Augusta, Countess of Lovelace—was an interactive space embedded in an exhibit representing a synthetic organism.[1] We divided the exhibit into five regions (see Figure 1), providing visitors with a didactical sequence of an introduction to Ada, the Ada experience, and background information on the Ada concept.

Ada perceived its environment and behaved coherently to achieve a set of behavioral goals. The Ada space sensed its world through 10 moveable and 4 static video cameras, 6 microphones, and pressure sensors in nearly 400 floor tiles. Ada's behavioral output included light effects generated by 20 moveable lights and color light elements in each floor tile, real-time computer graphics projected by 12 video projectors on a 360-degree canvas surrounding the space, called Bigscreen (see Figure 2), and the real-time music composition and performance system Roboser.[2] (To hear examples of Ada using Roboser, visit the official Ada Web site, http://www.ada-exhibition.ch.)

We integrated Ada's sensory information and controlled her behavior with a hybrid software architecture based on concepts derived from the computational neurosciences. This software architecture comprised a spectrum of implementation methods, ranging from standard procedural to agent-based methodologies and large-scale neuronal simulations.[3]

A central component of Ada's control system was a model of emotions operating in real time. We gave Ada the means to communicate via computer graphics and a musical soundscape. The goal was to use visual and auditory cues to induce visitors to adjust their behavior to Ada. This article focuses on the application of the Roboser music system to express Ada's behavioral modes and emotional states.

Ada's emotions

The starting point for developing Ada's emotional system was the distributed adaptive control (DAC) architecture that provides a neuronal model of the paradigms of classical and operant conditioning.[3,4] (We discuss other models for creating emotions in the "Natural and Synthetic Emotions" sidebar on p. 84.) DAC was originally developed using mobile robots, and is based on the assumption that these two fundamental forms of learning (classical and operand) result from three tightly coupled layers of control, called reactive, adaptive, and contextual control.

The reactive control layer provides basic predefined reflexive responses to low-complexity sensory events via internal state representations labeled as either aversive or appetitive. These internal state representations provide an interface to a first level of learning provided by the adaptive layer. The adaptive control layer constructs representation of sensory events that predict the occurrence of aversive and appetitive events and allows their sensory events to trigger the associated behaviors of avoidance and approach, respectively.

Hence, the DAC architecture lets a behaving system bootstrap itself from a purely predefined reflex level of behavior to one of actions triggered by acquired sensory representations, culminating in the use of acquired plans by the contextual layer. Researchers have translated basic elements

1070-986X/03/$17.00 © 2003 IEEE

Published by the IEEE Computer Society

of the DAC architecture into models of learning in the auditory cortex,[5] and they've shown that DAC displays basic properties observed in human decision making.[4] The aversive/appetitive value system of DAC formed the starting point for developing Ada's emotional system, with a design emphasis on communication (rather than learning).[2] With her emotions intact, Ada was ready to interact.

Ada's control architecture

Ada displayed numerous behaviors when interacting with visitors. The system organized her behavior around a number of basic functions, including

- identifying and tracking visitors,
- encouraging explorative behavior from visitors,
- guiding visitors through the space,
- gathering visitors in groups, and
- playing games.

We grouped these behaviors into modes following a cycle inspired by circadian rhythms in animals and humans. The space continuously evaluated how well it was doing in achieving behavioral goals by comparing current state values with target values. This self-evaluation process represented the basis for synthesizing Ada's moods and emotions.

Ada's behavioral modes

We organized Ada's behaviors in a cycle that lasted for 4 to 7 minutes, depending on visitor interaction. In this cycle Ada moved through phases of activity and rest, which were organized into six behavioral modes:

- sleep,
- wake up,
- explore,
- group,
- play, and
- fatigue.

Each mode had its own behavioral repertoire

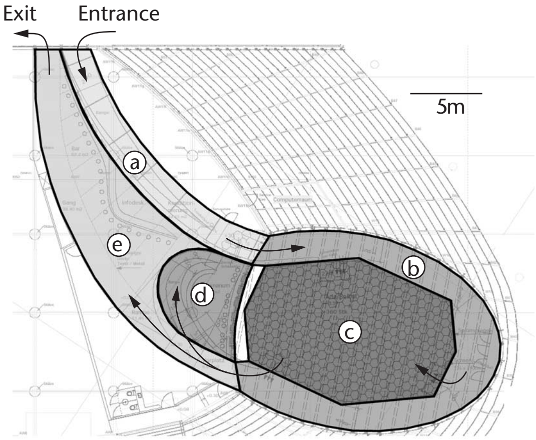

Figure 1. Floor plan of the Ada exhibit. (a) Conditioning tunnel: visitors were introduced to Ada's components. (b) Voyeur corridor: semitransparent mirrors allow a view of what happens inside Ada. (c) Ada self: the Ada main space for visitor interaction. (d) Brainarium: room with six monitors providing information and real-time graphical displays showing the current dynamics in Ada's control system. (e) Explanatorium: visitors are provided with background information on the exhibit. The arrows indicate visitor flow.

Photo by Tobi Delbruck

Figure 2. The Ada main space (from Figure 1c). The hexagonal floor tiles are pressure sensitive and display colored patterns dependent on Ada's behavior modes and on visitor interactions. The walls are made of semitransparent mirrors and allow visitors in the voyeur corridor (see Figure 1b) to view what happens inside Ada. Above the mirrors a circular projection screen displays real-time animated graphics that, similar to the music, represent Ada's current behavior and emotional state.

Natural and Synthetic Emotions

What good are emotions? In a functional sense, moods and emotions serve to bias perception and behavior to evaluate situations quickly and to generate appropriate and fast behavioral responses (such as escape behavior).[1,2] In social contexts, emotional communication establishes interpersonal relations and guides group behavior. A crying individual might facilitate nurturing behavior in a group, a single person expressing fear might lead to collective panic.[2]

It's generally agreed that music can convey emotions. However, not until recent years has scientific interest in the relationship between music and emotion emerged.[3] For example, in music psychology, the question of how music communicates emotional content has only recently been investigated empirically. Gabrielsson and Juslin[4] described a number of parameters used by musicians to convey emotional content in music performances and showed how effective these parameters were in enabling listeners to decode the emotional content intended by the musicians. In their experiments performers used cues like tempo, volume, and articulation (legato versus staccato) to interpret a melody with different emotional coloring.

Based on the concept of a bipolar affective space,[5,6] Schubert[7] used a 2D emotion-space model to empirically evaluate the emotional content of musical performances. Using a two-parameter joystick interface, subjects continuously reported on the valence (happy/sad) and the arousal (busy/sleepy) content of musical performance recordings.

A number of approaches relevant to the Ada concept have also been taken in the fields of artificial intelligence and robotics. Taking the work of Gabrielsson and Juslin[4] as a starting point, Bresin and Friberg[8] developed a set of rules to define a music program called Director Musices. Director Musices performs piano pieces from a digitized score to generate interpretations with different emotional coloring. In their software the rule set was used to vary musical parameters like sound level, the interval between note onsets, and note duration contrast.

Some examples of the role of synthetic emotions in man–machine interaction can be found in the experiments with the humanoid robot head Kismet.[9] Kismet interacts with humans, showing a set of emotional facial expressions and vocalizations. Its control architecture involves a model of a 3D affect space with the parameters of arousal, valence, and stance (social openness). Kismet, however, doesn't aim to directly influence the behavior of human observers, and hasn't been exposed to the kind of systematic experimentation performed with the Ada exhibit.

References

1. N.H. Frijda, "Emotions Are Functional, Most of the Time," *The Nature of Emotion: Fundamental Questions*, P. Ekman and R.J. Davidson, eds., Oxford Univ. Press, 1994, pp. 112-122.
2. R.W. Levenson, "Human Emotion: A Functional View," *The Nature of Emotion: Fundamental Questions*, P. Ekman and R.J. Davidson, eds., Oxford Univ. Press, 1994, pp. 123-126.
3. P.N. Juslin and J. Sloboda, eds., *Music and Emotion: Theory and Research*, Oxford Univ. Press, 2001.
4. A. Gabrielsson and P.N. Juslin, "Emotional Expression in Music Performance: Between the Performer's Intention and the Listener's Experience," *Psychology of Music*, vol. 24, 1996, pp. 68-91.
5. J.A. Russell, "Affective Space is Bipolar," *J. Social Psychology*, vol. 37, 1979, pp. 345-356.
6. J.A. Russell, "A Circumplex Model of Affect," *J. Social Psychology*, vol. 39, 1980, pp. 1161-1178.
7. E. Schubert, "Measuring Emotion Continuously: Validity and Reliability of the Two-Dimensional Emotion-Space," *Australian J. Psychology*, vol. 51, no. 3, 1999, pp. 154-165.
8. R. Bresin and A. Friberg, "Emotional Coloring of Computer-Controlled Music Performances," *Computer Music J.*, vol. 24, no. 4, 2000, pp. 44-63.
9. C. Breazeal, *Designing Sociable Robots*, MIT Press, 2002.

that defined the possible interactions between Ada and the visitors. It also had its own computer graphics and ambient light and soundscape aesthetics. The duration of single modes depended on how visitors acted inside the space. In terms of music expression, six different composition styles represented the behavioral modes.

Ada's goals

Because Ada was a public exhibition in an entertainment-oriented context, we chose to define three high-level behavioral goals relevant to visitor interaction behavior that Ada needed to achieve.[1] These three goal variables constituted the core of a homeostatic process (see Figure 3).

Survival. To achieve the survival goal, Ada needed to invite visitors to enter and participate when the space was empty, while it had to make visitors leave when the visitor capacity exceeded an ideal set value (in the exhibition this was 20 people). We derived the basis of this measurement from a people-tracking system that used the pressure sensors in the floor tiles. The survival value expressed the absolute difference between the set value and the actual number of visitors.

Recognition. We defined recognition as the ability of the space to locate and track individual visitors over extended periods of time. In addition, recognition took into account Ada's ability

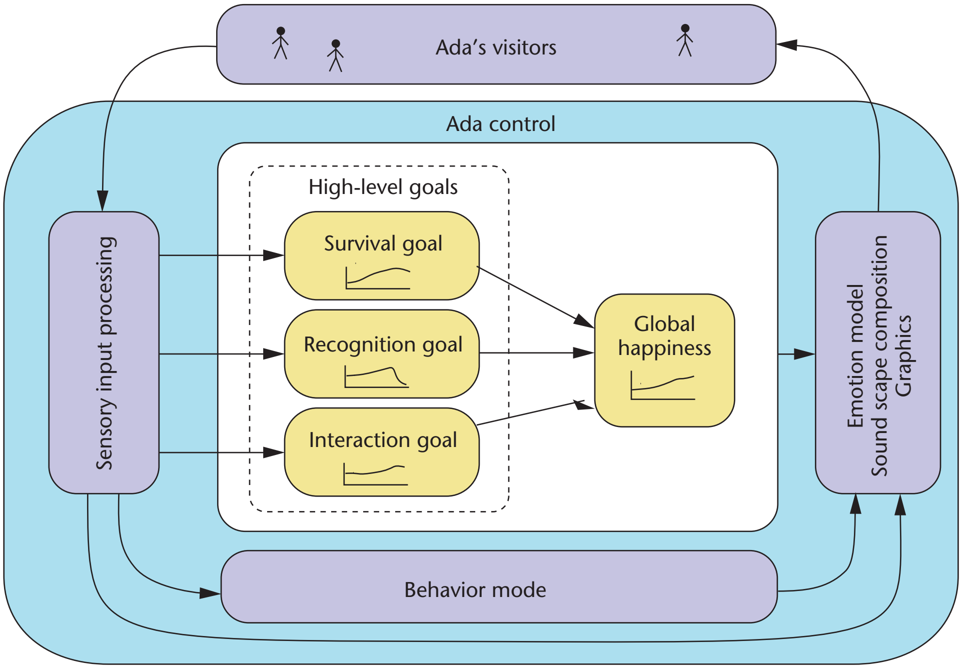

Figure 3. Scheme of Ada's control architecture and the goal functions of survival, recognition, interaction, and global happiness. Ada processes and uses sensor data generated by the actions of visitors to decide on local actions, bias behavioral mode selection, and to update the goal functions. The behavior mode and the status of the goal functions set Ada's emotional parameters that are the basis for synthesizing the musical output.

to correctly detect specific sonic events like hand claps, whistles, or speech events. We considered the detection of a visitor correct if a sonic event's estimated location matched the location of a person the floor tracking system had identified.

Interaction. In the course of the person-tracking process, Ada probed single visitors for attention and compliance by generating visual cues using the floor tile lights. In play mode, the system deployed a game in which visitors were invited to participate. We quantified cue compliance and game participation to set the interaction parameter value.

Survival, recognition, and interaction together defined Ada's global happiness (*H*). Ada tried to maximize *H* by generating interactive behaviors to achieve at least one of the three behavioral goals. As a first approximation, we can summarize the relation between the four goal variables as

$$H = f\,(g_s, g_r, g_i)$$

where *H* is global happiness, g_s is the survival goal, g_r is the recognition goal, and g_i is the interaction goal. Ada used her status of goal achievement to set her emotional parameters at the next processing level.

Ada's moods and emotions

To distinguish between Ada's moods and emotions, we adopted a model where mood changes on a time scale of hours or days, whereas emotions occur on a time scale of seconds to minutes.[6] Because the time course of Ada's behavior was rather compressed in comparison with humans, we set mood activity to change within tens of seconds to minutes and set emotions to occur within seconds.

Ada's mood system was defined by the two parameters of *arousal* and *valence*. The current behavior mode set the arousal parameter, resulting in low arousal for sleep and high arousal for play. The valence parameter represented the status of Ada's *H* goal achievement. Low happiness led to low valence and high happiness resulted in high valence.

The system synthesized Ada's emotions on input from the three high-level goals of survival, recognition, and interaction (see Figure 4, next page). Joy was set by the goals of survival and interaction in the sense that joy was high if survival or interaction approached maximum achievement. The sadness parameter was raised in case either recognition or interaction decreased from maximum achievement. Anger was excited in case survival decreased from maximum achievement. Surprise was triggered by a sudden increase in recognition.

The Roboser music system

Roboser is a real-time music composition and performance system that accepts input from a variety of sources to guide a composition process. Input

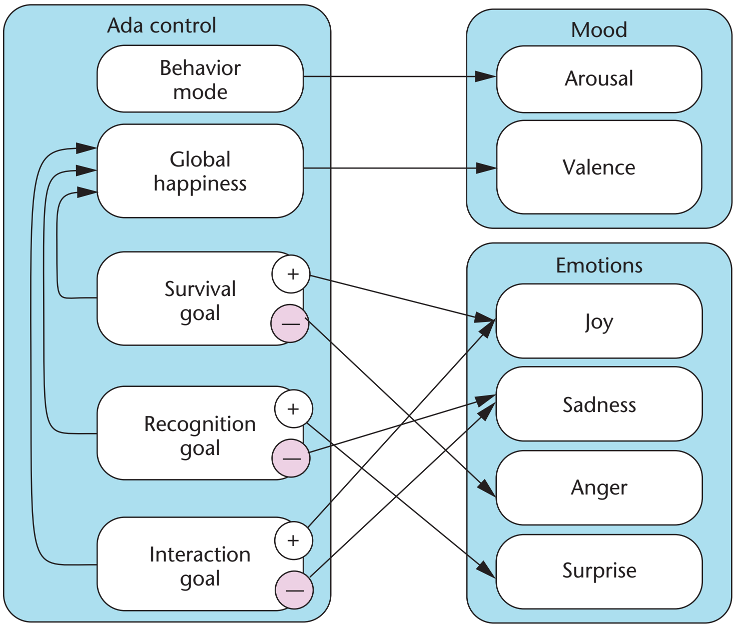

Figure 4. Ada's mood and emotion synthesis. The current behavior mode sets arousal, whereas global happiness sets valence. The emotions joy, sadness, anger, and surprise are set by the status of achievement for the high-level goals survival, recognition, and interaction (as the arrows indicate). The symbol (+) signifies approaching maximum goal achievement, while (–) signifies moving away from maximum goal achievement.

activity ranges from raw sensory data (for example, video images, audio events, and floor load) to high-order control parameters like biased neural oscillators, models of circadian rhythms, or models of behavior control. In Ada, Roboser expressed parameters like arousal, joy, and surprise, and gave sound feedback to simple events detected inside the space like footsteps, whistles, or hand claps.

Technically, the Roboser composition engine synthesizes a stream of MIDI data upon simulated neuronal input.[2] Roboser composes music on up to 12 performance tracks. This is similar to tracks of a multitrack tape recorder, the difference being that the performance of every track is synthesized as well as performed in real time. Musical parameters that are interactively and independently controlled in each track include the MIDI parameters for instrument, velocity, volume, pitch bend, tempo, and articulation. In addition, predefined fragmented note sequences, rhythm lines, and note onset dynamic sequences are interactively selected for each single performance track.

Each track's output is delivered on a single MIDI channel. During performance in Ada, the outputs of the Roboser tracks were performed using a sampler, resulting in a complex soundscape. Following the distinction between moods and emotions in Ada's control model, the music performance was synthesized on two separate layers: mood and emotion.

Mood layer

Two Roboser performance tracks continuously played a mood composition that expressed the two mood parameters of arousal and valence. The arousal parameter set the performances' tempo, volume, and octave register. As arousal increased, the number of note onsets per second, note overlaps, and volumes increased while note pitches were shifted upward in octave steps. The valence parameter generated pitch material that changed from dissonance to consonance, moving from semitone clusters for low valence toward pitch material from a harmonic series for high valence.

Emotion layer

For the emotion layer, we used a set of eight performance tracks to express the four emotions of joy, surprise, sadness, and anger. We assigned two tracks to each emotion. With the onset of an emotion, the volume of the respective two voices gradually increased from zero to maximum, fading the emotion compositions in and out on top of the mood layer composition. The emotion compositions' musical features are based on an extension of the scheme outlined by Gabrielsson and Juslin.[7] However, we extended this scheme by introducing the emotion of surprise. Because of this extension we slightly changed the original scheme to increase contrast between the four emotions.

The sadness composition used slow tempo and mellow timbre sounds, and the scales comprised low registers and minor or diminished chords. To express joy we used major pentatonic scales, rather bright timbres, rhythmical lines, and moderately fast tempo. We represented anger with rough and distorted sounds, semitone scales, fast tempo, and mostly staccato articulation. Surprise entailed a very fast tempo, high volume, bright timbres, and mostly augmented chords (see Table 1 for more information).

Results

To communicate behavior modes, goals, moods, and emotions effectively, Ada's control and music systems had to operate in real time. In addition, the music expressing the four emotions needed to be intuitively understandable to an audience that had never encountered Ada before. In this section we present data showing that we were successful in fulfilling both of these requirements.

Real-time performance

Figure 5 (on page 88) shows recordings from Ada's control system and music output. All data were recorded simultaneously during a single behavioral cycle. Figure 5a displays the normalized activity of six simulated neurons representing the six behavioral modes. One cycle consists of the respective behavior mode neurons being activated in a sequence from sleep to wake up, explore, group, play, and fatigue, returning to sleep at the end. The dynamics of Ada's moods and emotions show that arousal remains at a practically constant level after the system "woke up" (see Figure 5b). Ada's valence, however, shows a more intermittent activation due to specific visitor interactions.

We represented the behavior mode with shifts in key as well as differences in harmonic structure and articulation. You can see these in the MIDI piano roll plot (see Figure 5d) and in the sonogram (see Figure 5e), both representing the musical output. As Ada woke up, arousal, valence, joy, and surprise neurons became transiently active, leading to a pronounced change in musical communication. The music's pitch range widened involving higher notes, brighter sounds, and rising tempo. As the behavioral cycle progressed to explore, the mood and emotion neurons changed their activities according to visitor behavior. Surprise activity consistently led to a transient rise in pitches throughout the cycle (see Figures 5c, 5d, and 5e). As we can see in the MIDI and sonogram displays (see Figures 5d and 5e), the soundscape was densest in play mode with high tempo and high-register pitch material, immediately decreasing in volume, tempo, and pitch after the mode changed to fatigue. Fatigue mode is represented by a decrease in the arousal parameter of Ada's mood and involves downward pitch bends in the music (see Figure 5e) to give the impression of tiredness.

Table 1. Musical parameters used to express Ada's emotions.

Emotion	Parameter	Ada's emotion expression	Gabrielsson and Juslin
Sadness	Scale material	Minor, diminished	—
	Timbre	Mellow	—
	Tempo	Slow	Slow
	Sound level	Low or moderate	Low or moderate
	Articulation	Legato	Legato
	Time deviations	Moderate	Moderate
Surprise	Scale material	Major, augmented	—
	Timbre	Bright	—
	Tempo	Very fast	—
	Sound level	Loud	—
	Articulation	Staccato	—
	Time deviations	None	—
Anger	Scale material	Semitone	—
	Timbre	Harsh, distorted	Harsh, distorted
	Tempo	Fast	Fast
	Sound level	Loud	Loud
	Articulation	Mostly nonlegato	Mostly nonlegato
	Time deviations	Moderate	Moderate
Joy (happiness)*	Scale material	Major, pentatonic	—
	Timbre	Mellow to bright	Bright
	Tempo	Moderate or fast	Fast
	Sound level	Moderate	Moderate or loud
	Articulation	Rhythmical	Airy
	Time deviations	Moderate	Moderate

* Where available, we used the respective category outlined by Gabrielsson and Juslin[7] for comparison (wherein happiness applies instead of joy).

Audience comprehension

To assess whether an audience could decode Ada's emotions, we conducted an experiment that involved 35 subjects (15 living in Switzerland, 20 living in Spain) listening to a set of 27 Roboser music pieces outside the context of the exhibit. These compositions represented the four emotions of sadness (nine pieces), surprise (seven pieces), anger (four pieces), and joy (seven pieces). After listening to a composition, the subjects had to score it for each of the four emotional categories, on a scale of not applicable, low, medium, and high. We then converted these responses into an ordinal integer scale from 0 to 3, with 0 corresponding to not applicable and 3 corresponding to high.

In three of the four cases the subjects consistently identified the intended emotional expression (see Figure 6, p. 89). For the categories of sadness, anger, and surprise, the score for the intended emotional category was significantly higher than for the other three categories (with a t-test, where $p < 0.001$). The subjects identified anger compositions the best (the average score was 2.23—see Figure 6c), followed by the surprise pieces (with an average score of 1.63—see Figure 6b), and the compositions expressing sadness came in third (with an average score of 1.60—see Figure 6a). The result for the joy compositions is less clear (see Figure 6d). The subjects in the Swiss-resident group interpreted them as expected (an average score of 1.25, where $p < 0.001$), but

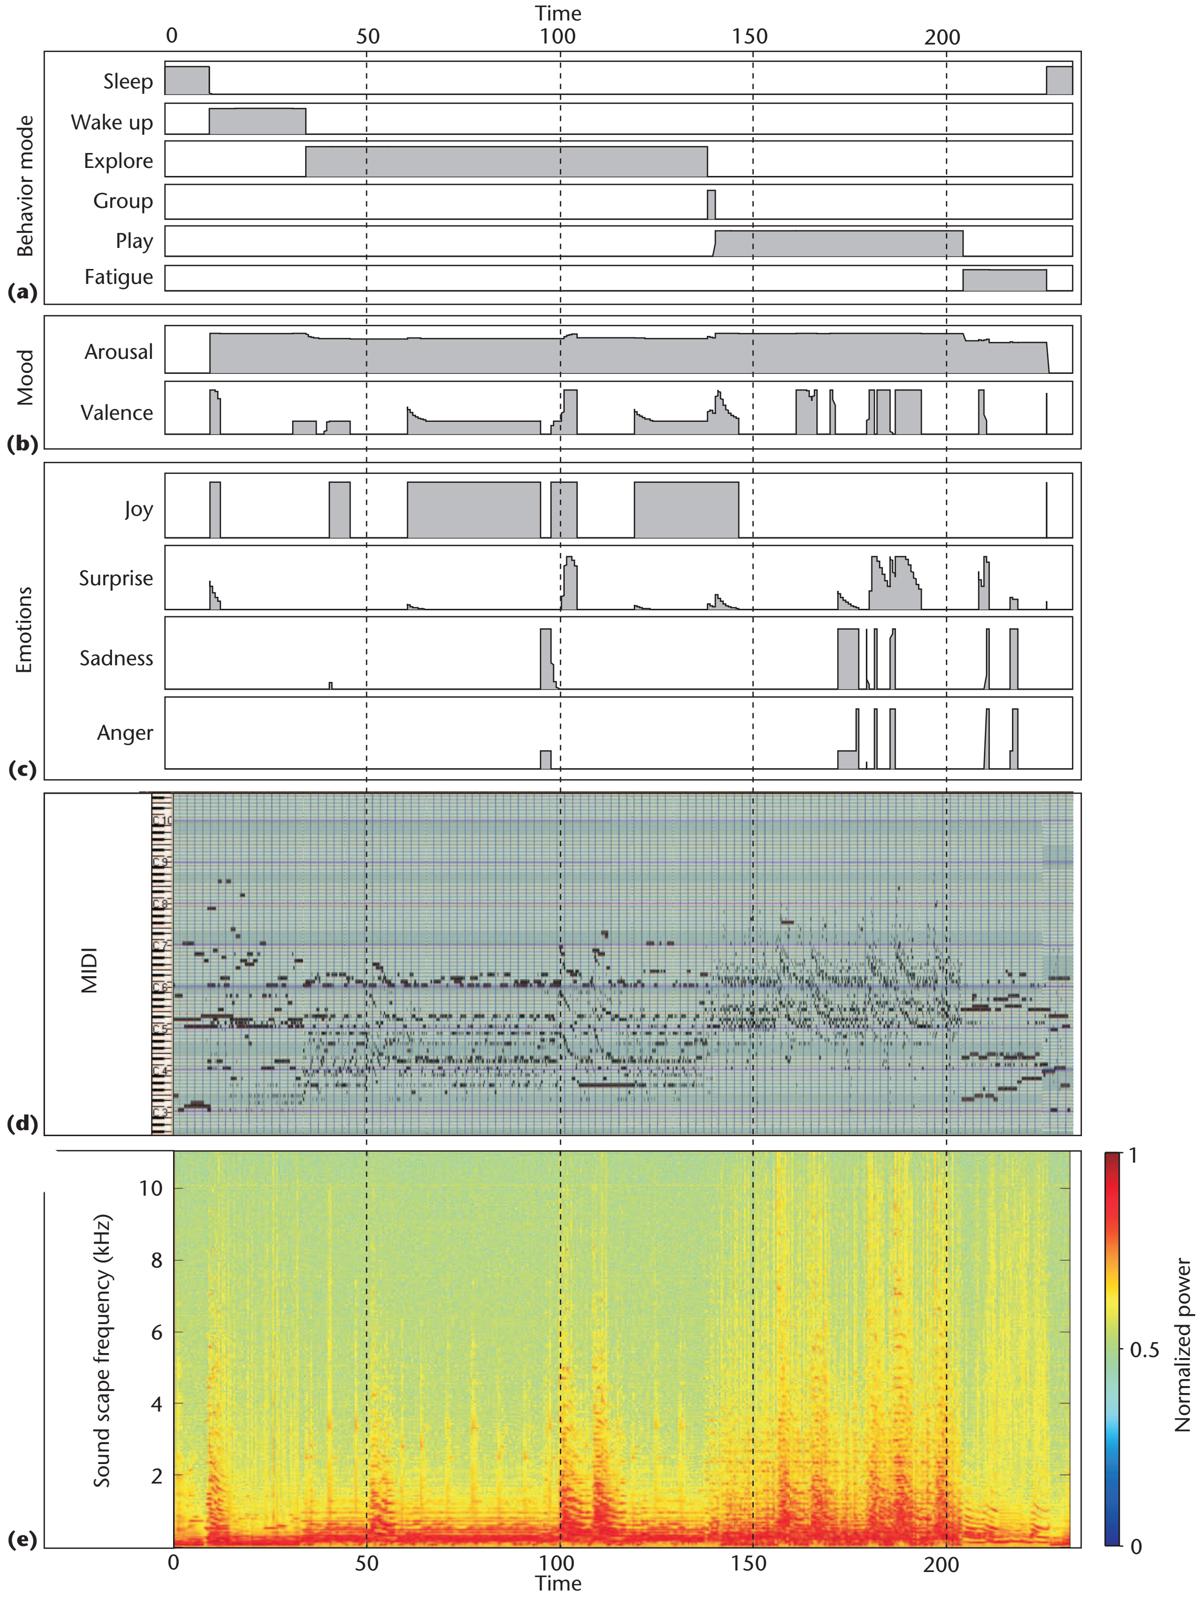

Figure 5. Simultaneous recordings of Ada's behavior modes, emotions, and music communication during a single behavior cycle. All plots share the same time scale. (a–c) Normalized displays of the cell activity of simulated neurons coding for behavior mode, mood, and emotions. (d) MIDI piano roll plot of the resulting Roboser music composition. (e) Sonogram of the actual sound output.

the subjects in the Spanish group tended to interpret these pieces in terms of sadness instead ($p < 0.01$). This suggests that cultural effects may play a role in interpreting the emotional content of musical pieces.

Overall, these results show consistency between the designed emotional qualities and the test subjects' interpretations. In particular, the surprise category, which we have introduced here in our extension of the Gabrielsson and Juslin scheme, gave a consistent result. The one inconsistent case was joy. Further work will deal with an interpretation of these results at the level of specific musical parameters, the role of cultural bias, and investigations into how the interpretations of the music are affected by other modalities (such as the visual information in Ada).

To evaluate the effect of Ada's emotional sound communication we measured the visitor response to Ada in the presence and absence of sound. Visitors reacted differently to Ada, in particular by making fewer sounds themselves, and attributed less importance to Ada's Bigscreen.

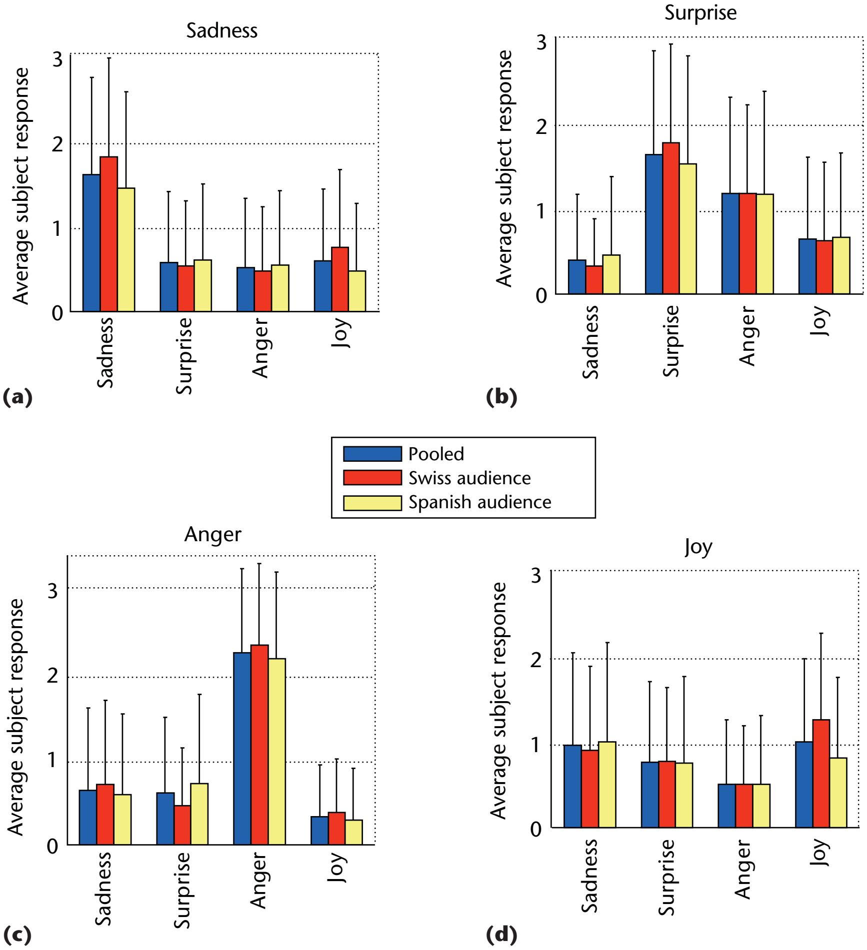

Figure 6. Results of an experiment in which 27 Ada emotion compositions were played to 35 subjects. Each of the four plots (a–d) represents the emotion that we wanted the composition to express. The vertical scale indicates the average subject responses on a scale of 0 = not applicable, 1 = low, 2 = medium, and 3 = high. Error bars indicate one standard deviation. Scores are shown in percentage for the following music types: (a) sadness, (b) surprise, (c) anger, and (d) joy.

Discussion and conclusion

Emotions are a core component of our subrational cognitive system. They assign positive and negative value to situations we find ourselves in, enhancing the prediction of what will happen in the near future, and optimizing the choice of appropriate behavioral responses. To maintain a communication process like the one in Ada, it's essential to give feedback immediately and interactively. Therefore, real-time performance is crucial. Another consideration was making a language comprehensible to an audience when they hadn't experienced it beforehand.

On aesthetic grounds we were successful in creating music with an appeal to the general audience and to specialists alike. Music generated by Ada was used for *Brainworkers*, a documentary film that the Swiss Federal Institute of Technology (ETH) in Zurich commissioned to introduce the Ada concept to the public. *Brainworkers* was awarded Best Swiss Commissioned Movie of the Year 2002, winning in four categories (including sound design). Unbeknownst to the jury, the soundtrack consisted (to a large extent) of Ada pieces. This result led us to the conclusion that Ada passed the musical Turing test.

The Ada exhibit attempted to define a large-scale interactive and communicative real-world artifact. Relying on an implicit language of sound and light rather than on explicit natural language, we found that Ada successfully communicated its emotions and (indirectly) its behavioral goals. This success points toward a domain of application and further development beyond the area of explicit and command-based user interfaces, employing a more implicit, intuitive, and multimodal communication between humans and machines.

The Roboser sound communication we've described was one element of the overall Ada exhibit. Our systematic analysis of both visitor behavior and their interpretation of the multimodal experience Ada offered shows that visitors were willing to consider Ada as an entity. An observer's interpretation of this real-world artifact is influenced by a large number of factors that are related to his or her demographics and personality as well as the different communication media and functions of Ada. We're currently further analyzing the complex interaction between humans and the largest artificial organism ever constructed. **MM**

Acknowledgments

Ada was a large, interdisciplinary team project involving many people. We would like to acknowledge the vital contribution of all of the team members who have made the realization of Ada possible. The Ada project was supported by the ETH, University of Zurich, Swiss National Exhibition Expo.02, Manor, Velux Foundation, and the Gebert Rüf foundation.

References

1. K. Eng et al., "Design for a Brain Revisited: The Neuromorphic Design and Functionality of the Interactive Space Ada," *Reviews in the Neurosciences*, vol. 14, nos. 1-2, 2003, pp. 145-180.
2. K.C. Wassermann et al., "Roboser: An Autonomous Interactive Composition System," I. Zannos, ed., *Proc. Int'l Computer Music Conf.*, Int'l Computer Music Assoc., 2000, pp. 531-534.
3. T. Voegtlin and P.F.M.J. Verschure, "What Can Robots Tell Us About Brains? A Synthetic Approach Towards the Study of Learning and Problem Solving," *Reviews in the Neurosciences*, vol. 10, nos. 3-4, 1999, pp. 291-310.
4. P.F.M.J. Verschure and P. Althaus, "A Real-World Rational Agent: Unifying Old and New AI," *Cognitive Science*, vol. 27, no. 4, 2003, pp. 561-590.
5. M.A. Sanchez-Montanes, P. König and P.F.M.J. Verschure, "Learning Sensory Maps with Real-World Stimuli in Real Time Using a Biophysically Realistic Learning Rule," *IEEE Trans. Neural Networks*, vol. 13, no. 3, 2002, pp. 619-632.
6. P. Ekman, "Moods, Emotions, Traits," *The Nature of Emotion: Fundamental Questions*, P. Ekman and R.J. Davidson, eds., Oxford Univ. Press, 1994, pp. 56-58.
7. A. Gabrielsson and P.N. Juslin, "Emotional Expression in Music Performance: Between the Performer's Intention and the Listener's Experience," *Psychology of Music*, vol. 24, 1996, pp. 68-91.

Klaus C. Wassermann is a biologist, composer, and performing musician who works as a research associate at the Institute of Neuroinformatics, ETH-University, Zurich. His interests include neural systems and bio-inspired robotics. Wassermann earned an MS in biology and epistemology from the University of Vienna.

Jônatas Manzolli is a composer and mathematician and currently heads the Interdisciplinary Nucleus for Sound Studies (NICS) at the University of Campinas (UNICAMP), Brazil. He acted as musical adviser for the Ada project. Manzolli received his PhD in computer music from the University of Nottingham.

Kynan Eng joined the Institute of Neuroinformatics in 2000 to work on the software development, system integration, and live operation phases of the Ada project. Eng is currently pursuing a PhD in the field of neuroinformatics. Eng received MS degrees in mechanical engineering and computer science from Monash University, Australia.

Paul F.M.J. Verschure is group leader at the Institute of Neuroinformatics, as well as project leader of *Ada: Intelligent Space*. He works on biologically realistic models of perception, learning, and problem solving applied to real-world artifacts (ranging from flying robots to large-scale interactive spaces). Verschure received his MA and PhD in psychology and has pursued research at institutions in the US and Europe. He is an IEEE associate member.

Readers may contact Paul F.M.J. Verschure at pfmjv@ini.phys.ethz.ch.