The Three Sirens

Die Musen des Jenseits (1992 bis 2005)

Nicolas Anatol Baginsky

nab@baginsky.de

Keywords:
Robot
musical instrument
neural network
T.Kohonen, blues

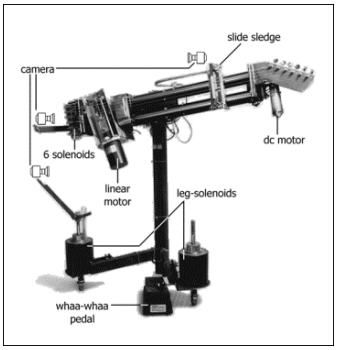

www.baginsky.de/agl/
www.the-three-sirens.info

Die Musen des Jenseits (1992 bis 2005)
Originally, three robots formed the band The Three Sirens, those were Aglaopheme (slide guitar), Peisinoe (bass) und Thelxiepeia (drum). By now, a few other artificial intelligences have joined the group: Aciilyzer (Vocals), LynxArm (Percussion) and some *off the shelf* electronic instruments.

The Three Sirens are not conventional musical instruments but autonomous, self learning robot musicians. They are not controlled by humans. Artificial neural networks control every aspect of the robot's activities. This means that the authorship for their music does not belong to any living being (no predefined material or musical knowledge is available to the system), all musical material is generated by self organizing learning processes in an improvising manner.

Aglaopheme
As an example I will now describe one member of the band in more detail - Aglaopheme the slide guitar robot.

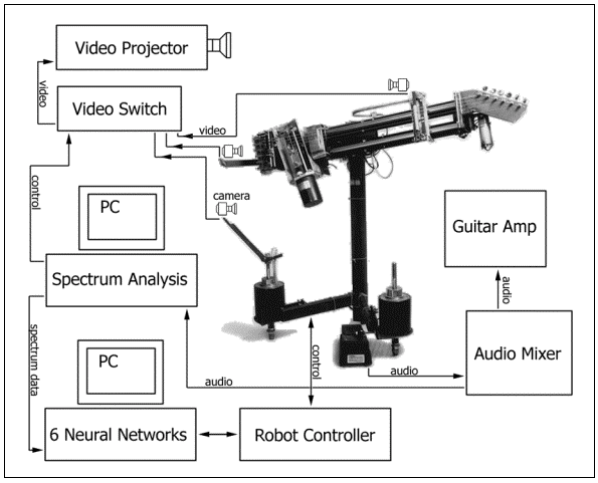

Aglaopheme schematic

Nicolas Anatol Baginsky

nab@baginsky.de

Keywords:
Robot
musical instrument
neural network
T.Kohonen, blues

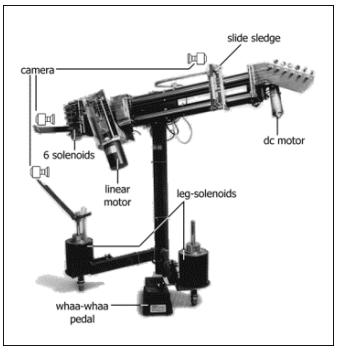

www.baginsky.de/agl/
www.the-three-sirens.info

The Body

Aglaopheme's body is both, a robot with framework, actuators and position sensors and a musical instrument (slide guitar) with six strings, machine heads and a guitar pickup. Six solenoids are used to either dampen and release or to pluck the strings. A short piece of cable tie serves as a very durable plectrum that is moved across the strings by a linear motor to pick selected strings. The entire plectrum-drive assembly can be pulled away from the strings by a solenoid to allow the plectrum to drive to the desired strings without hitting any others. A dc motor drives a sledge up and down the guitar neck. This sledge carries the slide and a small solenoid that pushes the slide against the strings if necessary. The sledge also carries a small video camera.

Two big solenoids serve as legs to the robot. They can push up the entire framework in two different directions and thereby make the robot rock. This rocking motion also moves the whaa-whaa pedal up and down. Two more video cameras are attached to the machine. One is located at the lower end of the guitar neck. The other one is mounted on top of a leg-solenoid.

The Brains

Aglaopheme's main behaviour control is composed of six artificial neural networks that control one guitar string each. The neural nets compete in being in control of the robot's mechanics. The network that, at a given time, has the most definite answer to the current sensory input is allowed to generate a command to be sent to the robot controller.

When the robot first started playing in December 1992, the six neural network that control the machine's behaviour were randomly initialised. Today there are several different sets of networks available for different modes of operation (different speeds and tunings). All these sets are descendants of the primal neural nets from 1992. This means that the robot system now has the experience of about ten years of playing. Not constantly but regularly.

The network model used in this system is a variant of the so called Self Organizing Maps (SOM) originally introduced by Tuevo Kohonen [Ko]. A weak inhibitory mechanism was added to the network model in order to avoid excessive feedback.

The Circulation

The sounds generated by the six guitar strings are converted to electrical signals by a guitar pickup, which are then fed into a whaa-whaa pedal that is operated by the motion of the robot's body. The resulting signal then runs through a distortion effect box and an audio mixer before being digitised by a computer that performs a spectrum analysis of the guitar sounds. The resulting spectrum information is transmitted via ethernet to another computer where it serves as sensorial input to the six neural networks (see The Brains). The neural network in command now sends it's response to the robot controller where the mechanics are instructed to make a movement that will result in new sound being generated. The control loop is finally closed and the system can begin to learn by playing.

Nicolas Anatol Baginsky

nab@baginsky.de

Keywords:
Robot
musical instrument
neural network
T.Kohonen, blues

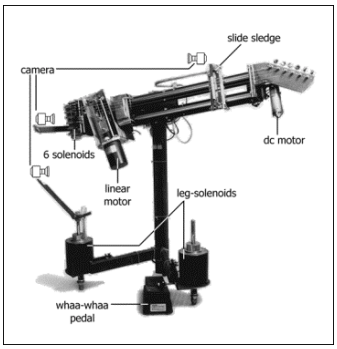

www.baginsky.de/agl/
www.the-three-sirens.info

Observations
When listening to the robot's play, it soon becomes apparent that the machine has a strong preference for the blues (in the Hendrix's sense). This first seems to be surprising but by carefully studying the basic conditions of the system one may find good reasons for this observation.

First of all the tuning. For slide guitar players it is very common to use so called open tunings. This means that the strings are tuned to a chord (preferably something with maj.7 in the end), something you don't do for normal guitar playing. I have experimented with various tunings and the one that the system seems to like the most (or rather, the one where I like the system's performance most) is the chord D#maj.7. This chord clearly is a blues chord as it contains the typical blues harmony.

A second prerequisite for the system's affinity with blues is the guitar string itself. By looking at the first few harmonics in the overtone spectrum of a vibrating string one finds that the first and third harmonic (unequal the fundamental) are the fifth (=tone 5 above fundamental) and that the seventh harmonic (unequal the fundamental) is the forth (=tone 4 above fundamental). Both intervals are essential for the blues scheme. Another interesting candidate is the fourth harmonic (unequal the fundamental). This tone turns an ordinary chord into a …major7 chord when added. This type of chord is heavily used in blues music.


Peisinoe, bass

The other Members of the Band

Peisinoe,
the bass player is a system very similar to the slide guitar robot described above. It is only much more simple in that it only has one string and therefore only one neural network to control it. In order to vary the pitch, the robot uses one strong solenoid to adjust the tension of the string. The machine has two mechanisms to cause vibration of the string: a motorized bow and a slapping device.

Nicolas Anatol Baginsky

nab@baginsky.de

Keywords:
Robot
musical instrument
neural network
T.Kohonen, blues

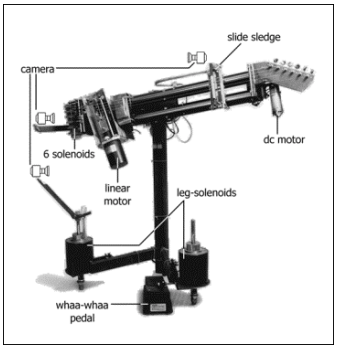

www.baginsky.de/agl/
www.the-three-sirens.info

Aciilyzer,
the voice of The Three Sirens is still somewhat bodyless. Again, there is a neural network that takes audio spectrum data as sensorial input and that is in control of an sound generating device. But in case of this band-member the visual appearance of the machine is very boring.
The sound generating device is a database filled with vocal sound samples of phonemes. Recordings of my own voice were used to set up the phoneme collection. The neural network is trained by allowing it to explore the outputs of the database. In a later state of learning the vocal sounds fed into the network are mixed with, or even replaced by, audio signals from the other members of the band. Intermixing the sounds of the different instruments makes the individual systems responsive to each other. In other words: they hear each other.
The percussionists in the band are operated differently from the melodically instruments. The rhythm machines are basically just hooked up to the system update frequency. If for example the band plays at 100 bpm, and the guitar- and the bass-system perform a learning cycle every quarter note (roughly 6 analysis per second), the bass drum would make a kick once per measure. All percussion devices can be programmed to act with different repetition rates within this basic quarter note structure.

Thelxiepeia,
the drum was build around an existing rototom, a drum that was fashionable in the seventies. On this drum the player could change the pitch during playing by turning it. I motorised the rotational activity and added a mechanic drumstick.

LynxArm
is a small five axis robot arm kit that I bought during a robotic exhibition in Vienna. After assembling the kit I changed one mechanic function and attached a small camera to what originally was the gripper. I also attached some sound pickups to the framework. The robot now knows two modes of operation. It either follows directly the spectrum analysis of the guitar sounds or knocks the camera against a small wooden board at system update tempo divided by four.

The Questions
The questions I am asked most often after concerts or during exhibitions of The Three Sirens is:
1. How can it …?
2. If it learns, what are the criteria?
3. How does it know what is good and what is bad?
4. The names, where do they come from?

My answers are:
1. I don't know, it simply does it.
2. + 3. No criteria, no predefined musical knowledge, the acoustic world is the teacher. The musical systems that we humans developed in the last couple thousand of years did not come out of thin air. It is all derived from the physical constellations of our world and of our hearing apparatus.
4. Greek mythology. The sirens that tried to distract Odysseus have individual names. Aglaopheme for example is one of them. The name translates "the one with the bright voice".

N.A.Baginsky, Hamburg, March 2002 (c) 2005NAB
[Ko] Tuevo Kohonen: "Self-organization and Associative Memory", Springer, Berlin Heidelb. (1984)