

Sansa: A Modified Sansula for Extended Compositional Techniques Using Machine Learning

McLean J. Macionis
California Institute of the Arts
24700 McBean Parkway
Valencia, California
mcleanmacionis@alum.calarts.edu

Ajay Kapur
California Institute of the Arts
24700 McBean Parkway
Valencia, California
akapur@calarts.edu

ABSTRACT

Sansa is an extended sansula, a hyper-instrument that is similar in design and functionality to a kalimba or thumb piano. At the heart of this interface is a series of sensors that are used to augment the tone and expand the performance capabilities of the instrument. The sensor data is further exploited using the machine learning program Wekinator, which gives users the ability to interact and perform with the instrument using several different modes of operation. In this way, Sansa is capable of both solo acoustic performances as well as complex productions that require interactions between multiple technological mediums. Sansa expands the current community of hyper-instruments by demonstrating the ways that hardware and software can extend an acoustic instrument's functionality and playability in a live performance or studio setting.

Author Keywords

Sansula, thumb piano, hyper-instrument, Wekinator, machine learning, force-sensitive resistor, accelerometer

CCS Concepts

•Applied computing → Sound and music computing; •Computing methodologies → Machine learning; •Hardware → *Sensors and actuators*; Tactile and hand-based interfaces;

1. INTRODUCTION

The sansula¹ is a modern adaptation of the kalimba, a type of thumb piano that originates from southern Africa. The instrument consists of several metal tines mounted to a wooden block that is attached to a hollow resonating membrane. Throughout its existence, musicians, composers and amateur music-makers have enjoyed the sansula for its wide variety of uses including composition and performance tool.

Two key features of the sansula make it an attractive choice for music technologists seeking to develop hyper-instruments. First, the instrument's simple interface makes it incredibly easy to pick up and play. This means that users of all musical backgrounds can explore the instrument

with little to no prior experience using it. Second, the instrument's frame is hollow and therefore an ideal location to install sensors and micro-controllers.

Sansa relies on these features to extend the functionality of the sansula into the digital domain and realize novel ways of playing the instrument.

2. RELATED WORK

Several related works involving extended thumb pianos inspired the design and development of Sansa: the first is *Kalimba Mocante* (2016) by Meng Qi;² the second is *Ember* (2013) by Colin Honigman et al.[9]; the third is *El-Lamellophone* (2014) by Shawn Trail et al.[13]; and the fourth is *Kalimbo* (2017) by Rob Blazey[1]. Each of these interfaces uses real-time data streams from micro-controllers and sensors to expand the capabilities of homemade kalimbas.

Sansa is also motivated by a number of works within the field of machine learning, audio and gesture recognition, and multimodal analysis. Jordan Hochenbaum et al. have used these techniques within the context of drum stroke computing[6] and musical pedagogy[7]. Alongside Matt Wright, they also used these techniques to classify and distinguish between the playing styles of multiple musicians[8]. Similarly, projects by Parag Chordia et al.[2], Kameron Christopher et al.[3], Arne Eigenfeldt et al.[4], Alex Tindale et al.[12][11], and Manj Benning et al.[10] rely on machine learning and multimodal analysis to recognize musical sounds and gestures for the purpose of enhancing musical performances.

Sansa draws on these examples and shows that acoustic instruments are capable of uses far beyond what is expected of them. In some cases users may not need to play these instruments at all, instead relying on a combination of vocalizations and gestures in order to create music. Additionally, Sansa demonstrates that there are numerous opportunities to incorporate machine learning into the design of future hyper-instruments.

3. INSTRUMENT DESIGN

Sansa was created with the intention of extending the sansula's functionality and playability with minimal impact to the original design and use of the acoustic instrument.

3.1 Instrument Body

The instrument itself is a Sansula Elektra.³ It contains a 3.5mm mini-jack, a piezo transducer and a condenser microphone, all of which are used to collect and amplify the

¹<https://www.hokema.de/en/products/sansula/>



Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).

NIME'18, June 3-6, 2018, Blacksburg, Virginia, USA.

²Meng Qi <https://www.mengqimusic.com/>

³<https://www.kalimbamagic.com/shop/hokema-kalimbas/hokema-sansula-elektra>

instrument's acoustic signal. Connecting to an audio interface provides the onboard amplification system with 48V of phantom power and passes the audio signal into any digital audio workstation (DAW) or audio programming environment with an analog-to-digital converter. Figure 1 shows the instrument's body including the built-in hardware.




Figure 1: Front and back view of Sansa.

3.2 Micro-Controller and Sensors

Sansa contains a Teensy micro-controller,⁴ an XBee wireless data transmitter,⁵ a battery and a series of sensors, specifically a 3-axis accelerometer⁶ and two force-sensitive resistors (FSR).⁷ Each of the hardware components is located within the frame of the instrument and configured according to the diagram in Figure 2. There are several reasons for this design configuration. First, the sansula frame contains a hollow cavity which is suitable for housing small electronics with minimal effect to the sound and feel of the instrument. Second, installing the hardware within the frame of the instrument removes it from the user's line of sight, which serves to limit distractions during use. Lastly, locating the FSRs within the frame provides users with easy access to apply pressure.




Figure 2: Circuit diagram.

While Sansa is in operation, the sensors generate a data stream based on the user's interactions with the instrument. The data is then transmitted wirelessly via the onboard XBee to a computer. ChuckK,⁸ a music-oriented programming language, is used to parse and distribute the data to different applications.

4. MACHINE LEARNING

Sansa provides users with precise control over compositional form and the arrangement of sound and music in real-time

⁴<https://www.sparkfun.com/products/13736>

⁵<https://www.sparkfun.com/products/11215>

⁶<https://www.sparkfun.com/products/13926>

⁷<https://www.sparkfun.com/products/9375>

⁸<http://chuck.cs.princeton.edu/>

through its application of machine learning. This is accomplished using Wekinator,⁹ a machine learning program developed by Rebecca Fiebrink[5]. Wekinator is free and open source and allows individuals to build interactive systems through the use of supervised learning algorithms. A number of different learning algorithms are included with the software such as AdaBoost, dynamic time warping (DTW), decision tree and nearest neighbor.

Networking between Sansa and Wekinator is done using Open Sound Control (OSC).¹⁰ Sansa sends the sensor data as a continuous stream of OSC messages from ChuckK to Wekinator. These messages are used to tell Wekinator to alter its behavior based on the chosen learning algorithm and training data. Using Wekinator's Input Helper allows the user to filter the incoming data using averages, buffers, first and second order difference equations and custom mathematical expressions. Sansa uses the Wekinator's built-in algorithms to great effect, cuing new musical sections, processing audio, adding or removing instruments from an arrangement and driving a performance forward.

5. MODES OF OPERATION

Sansa draws on its unique configuration of hardware and software to realize new ways of playing the sansula. This section describes five experimental modes used during live performances. Figure 3 shows the instrument's signal flow and lists several possible applications of the data.




Figure 3: Signal flow diagram.

5.1 Conductor

Sansa is capable of organizing and conducting an entire ensemble of instruments by means of OSC and sensor mapping. Under one configuration, the sensor data is used to trigger scenes in Ableton Live.¹¹ Tilting the instrument to the left or right moves the arrangement backward or forward while tilting the instrument to the front or back starts or stops the music. The FSR sensors are used to apply effects to the stereo mix or individuals tracks.

The conductor mode is also useful for score following. By using specialized gestures and Wekinator's dynamic time warping algorithm, Sansa is able to move through a score and shift between user modes. Figures 4 shows one possible score and Figure 5 shows how a performer might move through that score using simple gestures.

5.2 Advanced Filtering Unit

Another mode of operation involves using all of the available sensor data to shape external audio. In this mode, the user is not required to play the instrument at all. Instead, their gestures are used to shape and apply filters to a sound. Using Wekinator's Input Helper allows the user to apply

⁹<http://www.wekinator.org/>

¹⁰<http://opensoundcontrol.org/>

¹¹<https://www.ableton.com/en/live/>

filters to the raw data, which can then be passed via OSC to any number of sound shaping programs such as Ableton Live or Reaktor.¹² An example of this configuration can be seen in Figure 6.




Figure 4: One possible score incorporating several of Sansa's user modes.




Figure 5: One possible configuration for the conductor mode.




Figure 6: One possible configuration for the advanced filter mode.

5.3 Audio Processor

Sansa's internal microphone is able to capture the audio from external sources when placed near the source. This feature is especially useful for capturing vocal sounds or audio from other acoustic instruments. One application of this mode involves speaking or singing into the microphone

¹²<https://www.native-instruments.com/en/products/komplete/synths/reaktor-6/>

to take advantage of any filtering or processing tools that are already affecting the tone of the instrument. Another application of this mode involves using Wekinator's DTW system and an audio feature extractor for timbre or speech recognition. The signal flow of this application is shown in Figure 7.




Figure 7: Signal flow for audio input through Sansa and Wekinator.

5.4 Percussive Instrument

By taking advantage of the internal microphone again, the user is able to treat Sansa like a percussive instrument. Sansa's microphone is capable of amplifying any number of percussive interactions on the membrane of the instrument, such as tapping or scratching. This feature is especially useful within the context of loop-based live performances in which a performer records multiple layers of sounds to create a cohesive track. By sending percussive audio from Sansa to a looping device such as a loop pedal or Ableton Live's "Looper" plugin, a performer can record and overdub a track thereby generating a rhythmic foundation to play over.

5.5 Visualizer

In addition to the aforementioned musical applications, Sansa's sensor data is also capable of driving visualizations. Many visual programming languages, such as Processing¹³, openFrameworks¹⁴ and TouchDesigner¹⁵, accept OSC messages. This means that creating audio-reactive or gesture-reactive visualizers is as simple as sending OSC messages between ChuckK or Wekinator and the target visual programming environment. Musical performances take on more complexity and depth when combined with additional sensorial information such as those from visualizers and Sansa is capable of driving both an auditory and visual performance with ease.

5.6 Acoustic Performance

It is worth noting that Sansa is capable of purely acoustic performances. No modifications to the instrument were

¹³<https://processing.org/>

¹⁴<http://openframeworks.cc/>

¹⁵<https://www.derivative.ca/>

made that impede on the user’s ability to explore the instrument without the use of the hardware.

6. DISCUSSION

6.1 Challenges and Limitations

Several challenges emerged during the development phase of Sansa. First, mounting electronics within the frame of the instrument dampened the overall sound and shortened the decay of each note. One possible solution for restoring the original timbre of the instrument involves mounting the electronics away from the wooden resonator in some capacity. However, this solution requires new paneling for the back of the instrument, which may create additional changes in tonality and difficulties holding the device.

Another challenge that emerged during production was the existence of accelerometer data spikes after playing each note. These data spikes occurred due to the accelerometer’s location directly on the wooden resonator, which vibrates after each note is plucked. This resulted in unwanted sounds and changes to the arrangement. A combination of data averaging and a gentle playing style helped to minimize the effect of the data spikes but some presence still remains.

Lastly, striking the instrument for percussive purposes occasionally caused the internal microphone to malfunction resulting in loud amplified buzzing. Reinforcing the connection between the microphone and the mini-jack is necessary in order to resolve this issue.

6.2 Future Work

Future work on Sansa will focus on two distinct areas of research. The first is minimizing the impact of the hardware on the overall sound and playability of the instrument. The most obvious means of accomplishing this is to substitute larger micro-controllers and sensors for smaller ones, effectively reducing the mass of the hardware on the resonator. Another possible way to reduce the impact of the hardware on the instrument is to create a dedicated panel for the electronics that is isolated from the resonator. However, this runs the risk of impairing playability if the panel obstructs the user’s hands from easily gripping the instrument.

The second area of research involves carrying out rigorous tests of Sansa’s software and hardware system. A greater understanding of the audio and control capabilities related to latency, distortion and connectivity is needed to better understand the instrument. Sansa relies primarily on Wekinator’s dynamic time warping and nearest neighbor algorithms to drive each of the user modes. Initial testing found these to be the most successful algorithms for real-time performance. However, more analysis of Wekinator’s machine learning models is needed to optimize the instrument for future performances.

Finally, one possible expansion of this project involves adding capacitive touch capabilities to the metal tines. In this scenario, playing the instrument using the tines would produce MIDI or OSC messages. This modification would further extend the functionality of the sansula and create a one-to-one relationship between playing a note and a resulting output.

6.3 Conclusion

Sansa is a powerful and novel instrument that is capable of both simple acoustic performances as well as those that involve complex interactions between multiple technological mediums. Machine learning further extends the functionality of the instrument by providing users with a wider range of possible interactions and resulting outputs. In this way, Sansa offers a peak at the sansula’s future while simultane-

ously honoring and preserving the instrument’s past.

7. REFERENCES

- [1] R. Blazey. Kalimbo: an extended thumb piano and minimal control interface. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 501–502. NIME, May 2017.
- [2] P. Chordia and A. Rae. Tabla gyan: A system for real-time tabla recognition and resynthesis. In *Proceedings of the International Computer Music Conference (ICMC 2008)*, August 2008.
- [3] K. Christopher, J. He, R. S. Kapur, and A. Kapur. Kontrol: Hand gesture recognition for music and dance interaction. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 267–270. NIME, May 2013.
- [4] A. Eigenfeldt and A. Kapur. Multi-agent multimodal performance analysis. In *Proceedings of the International Computer Music Conference (ICMC 2008)*, August 2008.
- [5] R. Fiebrink and P. R. Cook. The wekinator: A system for real-time, interactive machine learning in music. In *Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2010)*, August 2010.
- [6] J. Hochenbaum and A. Kapur. Drum stroke computing: Multimodal signal processing for drum stroke identification and performance metrics. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. NIME, May 2012.
- [7] J. Hochenbaum and A. Kapur. Toward the future practice room: Empowering musical pedagogy through hyperinstruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 307–312. NIME, May 2013.
- [8] J. Hochenbaum, A. Kapur, and M. Wright. Multimodal musician recognition. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 233–237. NIME, May 2010.
- [9] C. Honigman, A. Walton, and A. Kapur. The third room: A 3d virtual music paradigm. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 29–34. NIME, May 2013.
- [10] A. Kapur, M. Benning, and G. Tzanetakis. Query-by-beat-boxing: Music retrieval for the dj. In *Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2004)*, October 2004.
- [11] A. Tindale, A. Kapur, and I. Fujinaga. Towards timbre recognition of percussive sounds. In *Proceedings of the International Computer Music Conference (ICMC 2004)*, November 2004.
- [12] A. Tindale, A. Kapur, G. Tzanetakis, and W. Schloss. Indirect acquisition of percussion gestures using timbre recognition. In *Proceedings of the Conference on Interdisciplinary Musicology (CIM 2005)*, March 2005.
- [13] S. Trail, J. Snyder, D. MacConnell, G. Tzanetakis, L. Jenkins, and P. Driessen. El-lamellophone - an open framework for low-cost, diy, autonomous lemellophone based hyperinstruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 537–540. NIME, June 2014.