



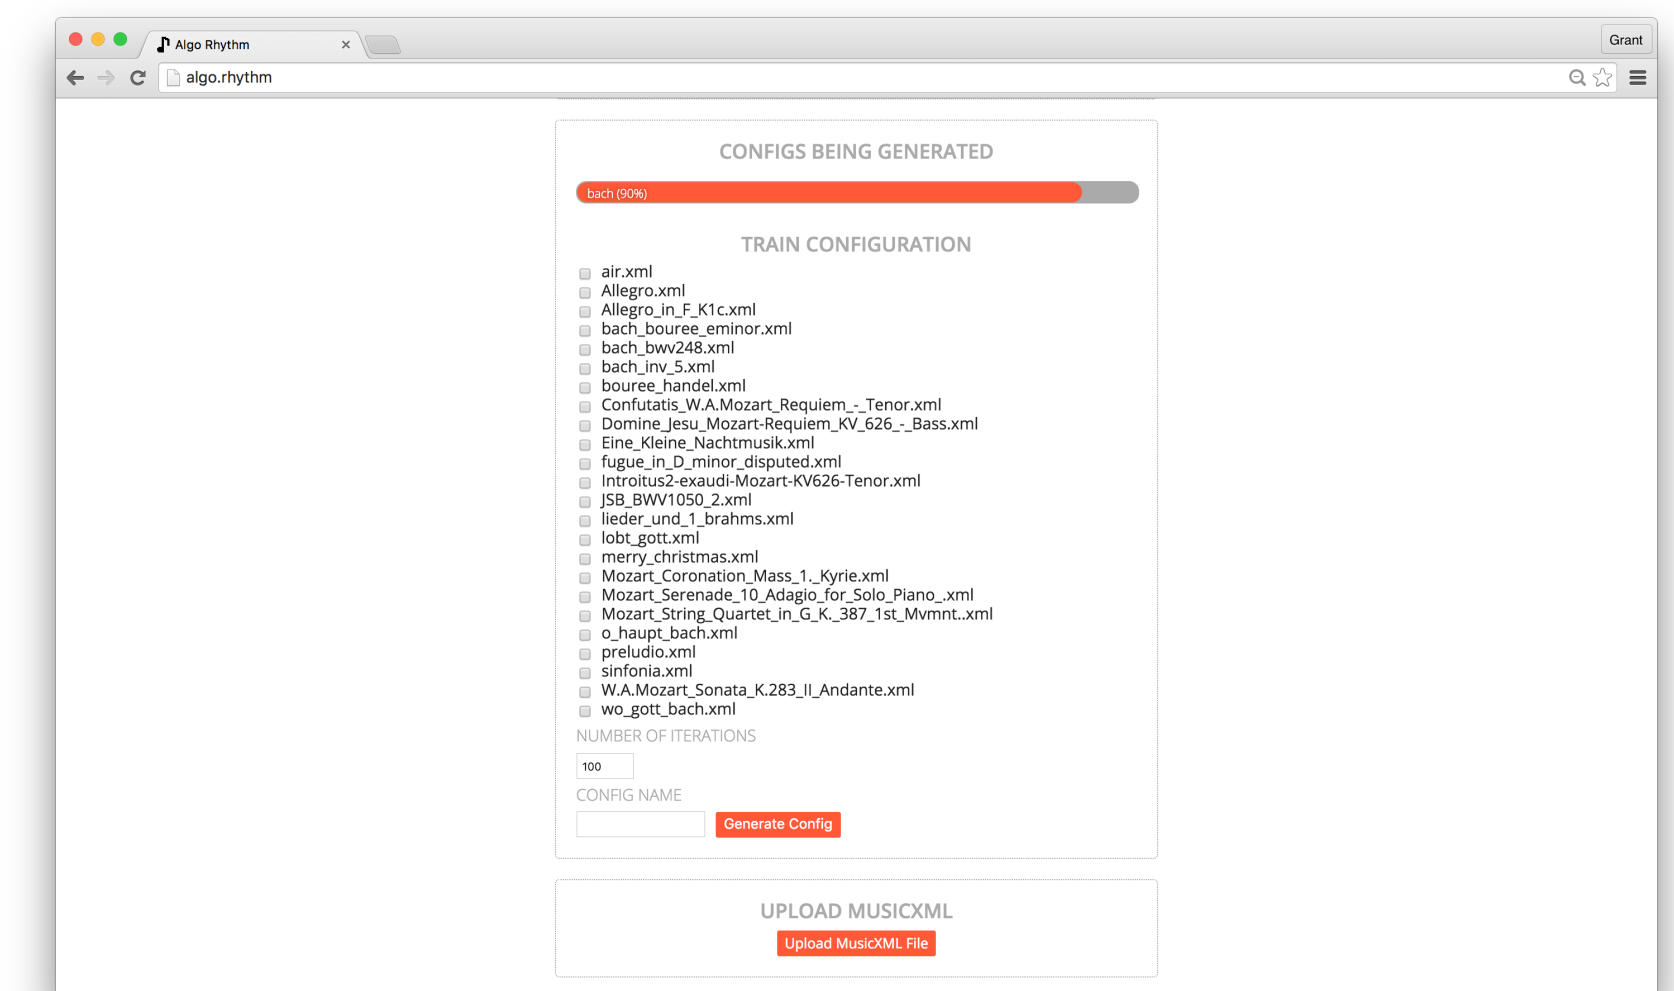
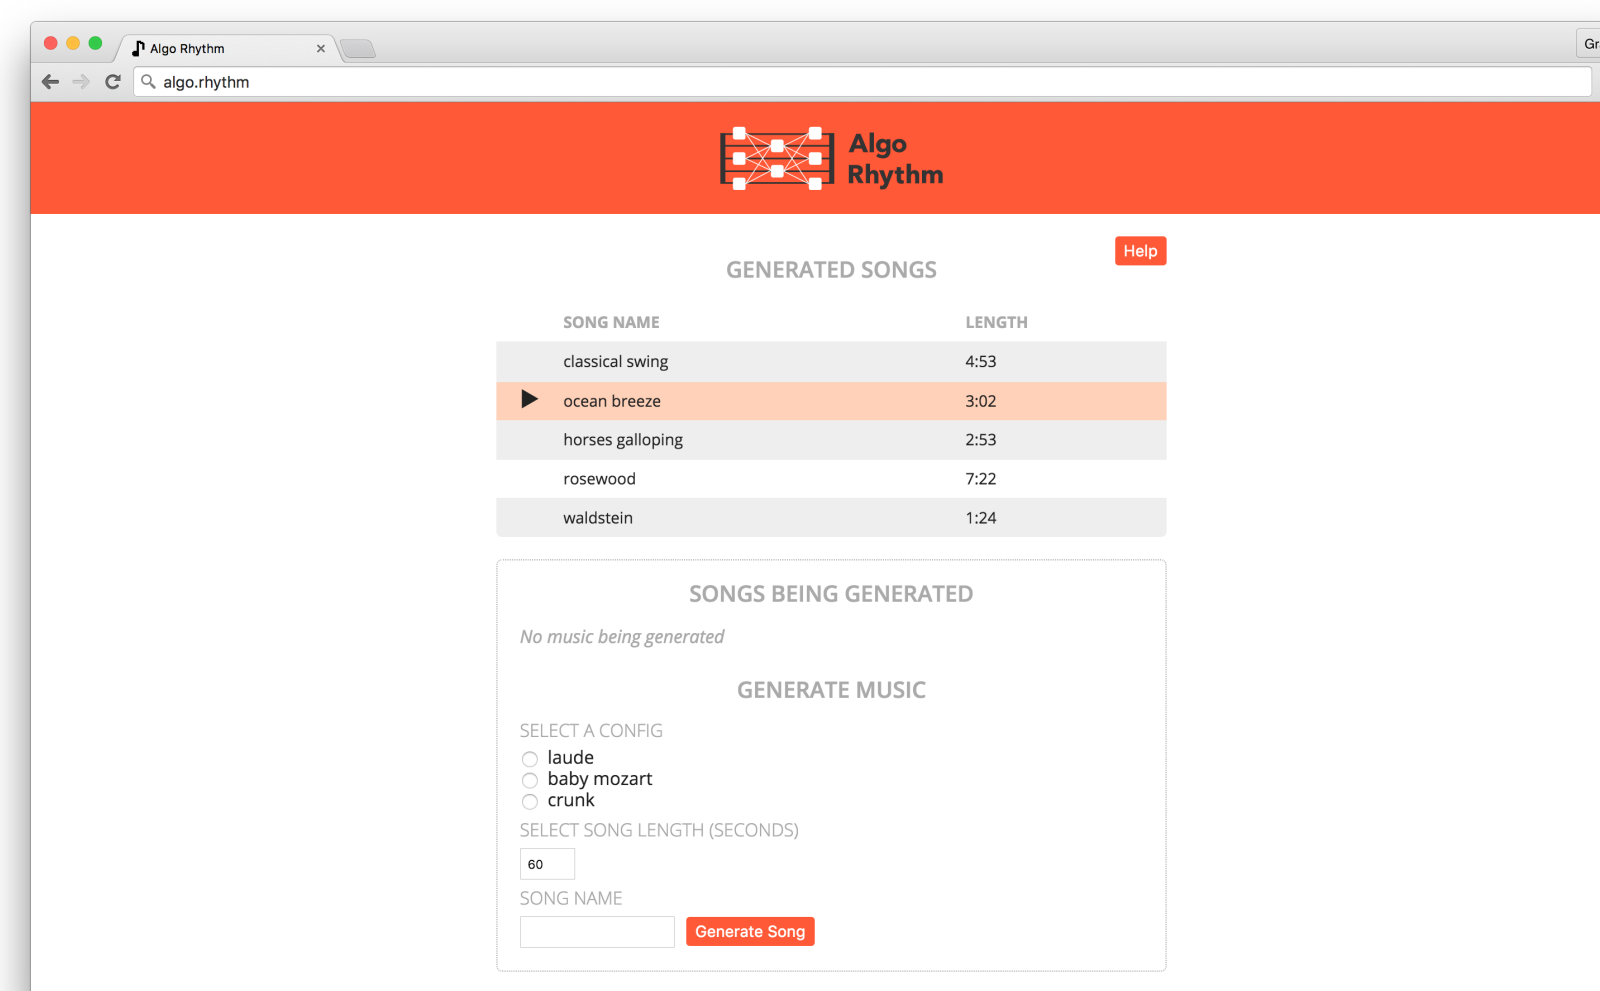
A Web-Based System for Algorithmic Music Composition Using Recurrent Neural Nets.



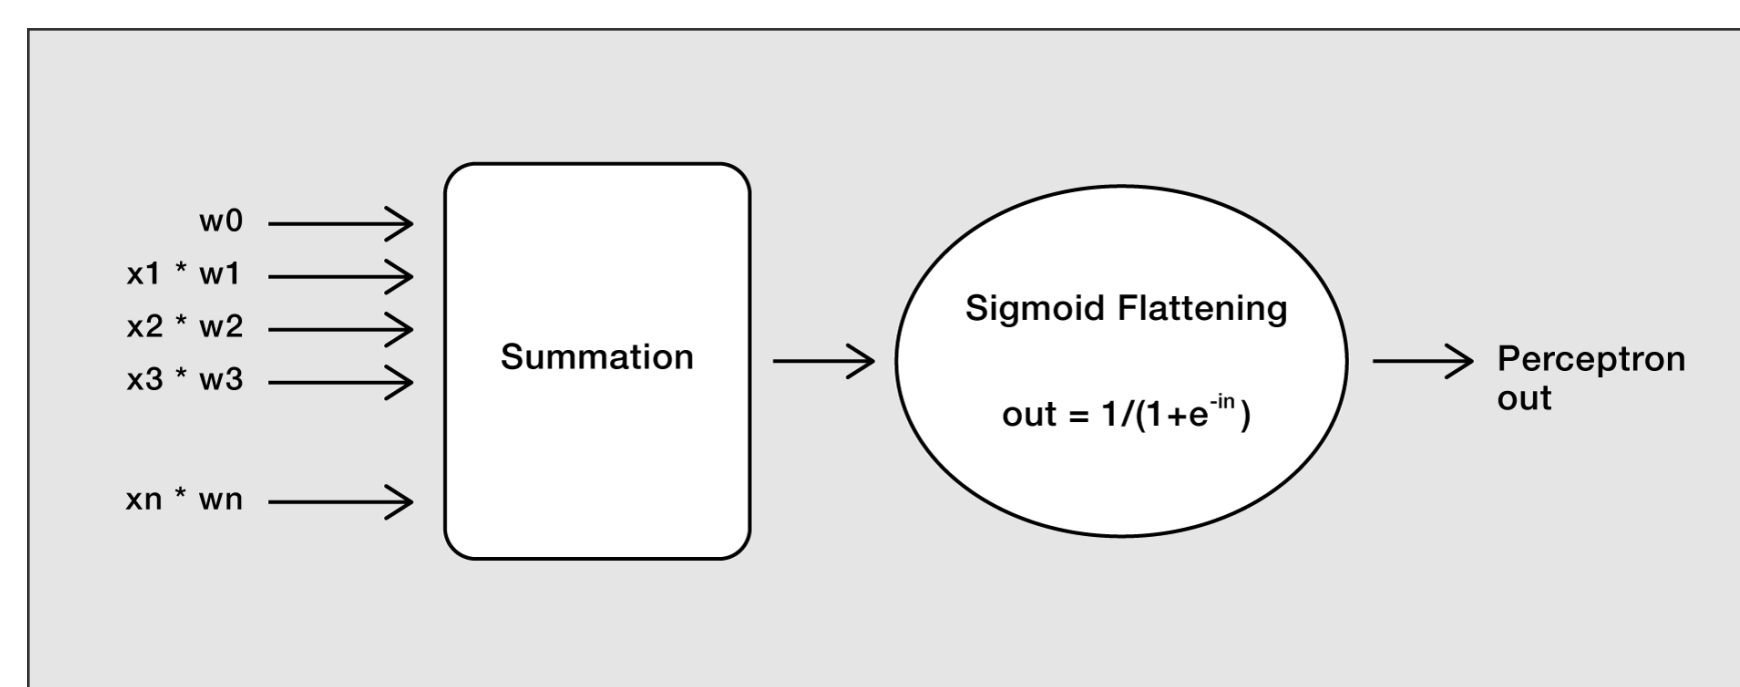
Web Interface

Algo Rhythm trains artificial neural networks to quickly write music in the style of sample data using GPU acceleration. It comes with a web interface to allow users to easily upload custom training data and generate music.


1. User uploads pieces of music in XML.
2. The user then presses the "train" button and waits for training to finish.
3. User selects a config and presses "Generate Song" to generate a song of N seconds.



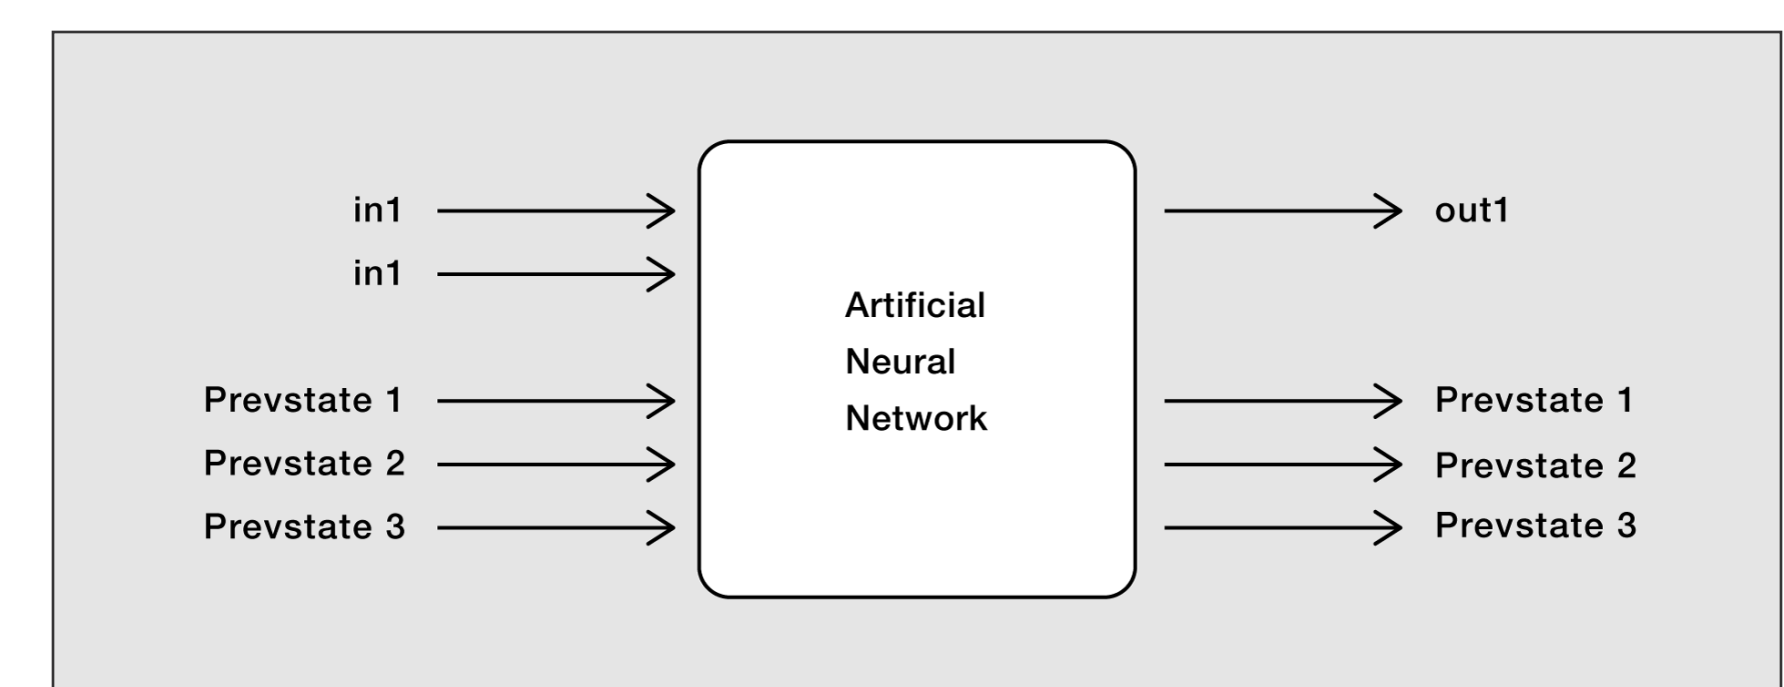
Artificial Neural Network (ANN) / Recurrent Neural Network (RNN)



A Perceptron represents a function from n real inputs to a single real output, all numbers in the range (0, 1). The function is determined by n + 1 parameters.



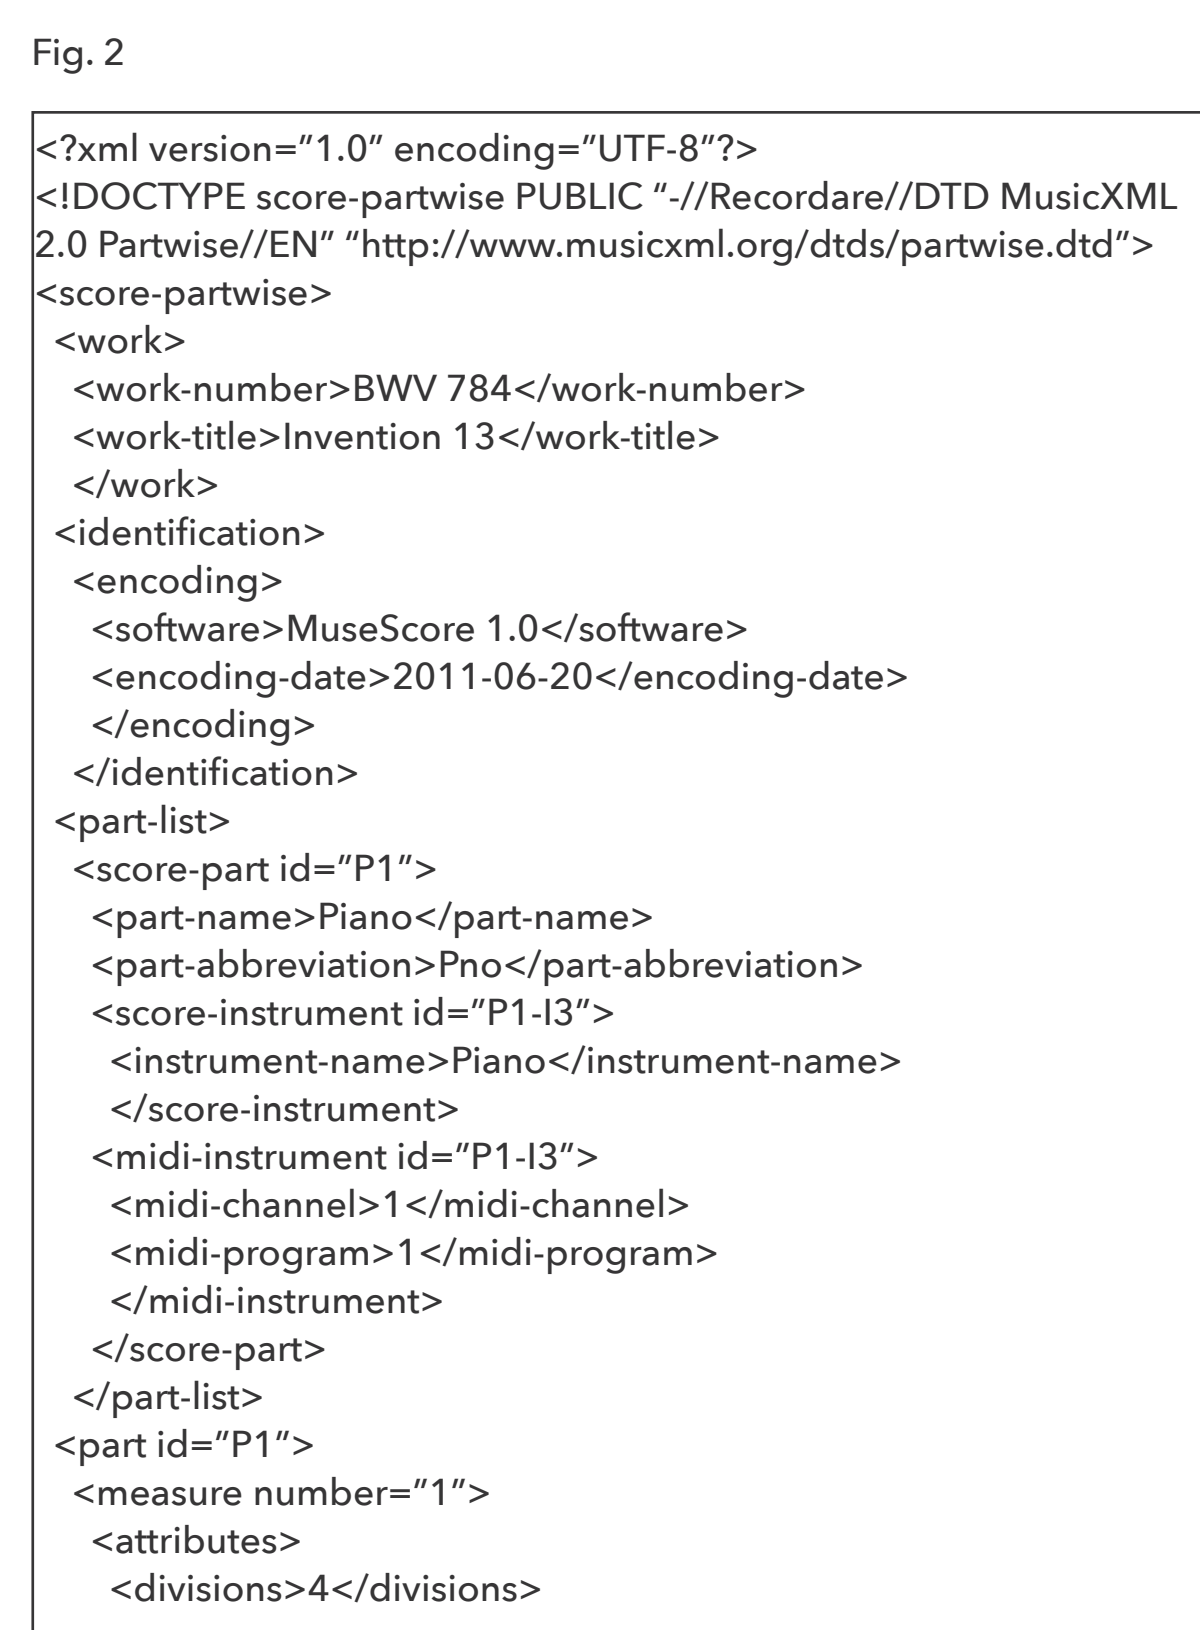
ANNs can be constructed to map m inputs to n outputs by linking groups of perceptrons in layers. Given training data, the parameters can be adjusted using training algorithms to "learn" many different functions, ranging from image recognition to music generation.



RNNs are perceptron networks in which layers feed their outputs back into themselves, creating a facsimile of memory. They allow training on (and generation of) sequences rather than single input/output pairings. This lends itself well to music generation.

Music XML

- Music XML is an XML format used to encode sheet music.
- The XML contains a standard format for valuable song information such as tempo, note pitch, duration, octave, et cetera.
- Music XML can be converted to graphical scores (e.g. fig 1) or to an audio format such as MIDI.



GPU

The training of a neural net, as well as the evaluation of a neural net after training (referred to as "classification") involves a large amount of parallelizable floating point computation. This computation is amenable to being accelerated by graphics processing units, or GPUs.

Running time:


Intel Core i5



Jetson TX 1



Jetson TX 1 board is a development board for the NVidia Tegra APU chip (accelerated processing unit) that combines CPU and GPU functionality in a single unit. The Tegra APU is targeted at the embedded systems market, providing high floating point throughput while consuming very low power.



The Tegra GPU consumes under 10 watts. During our project was measured to be 10 times faster than an Intel i5 Core processor at performing neural net training.