Little Ludwig, an evolutionary learning machine for musical composition

EARL BELLINGER, State University of New York at Oswego

This paper introduces a genetic artificial intelligence program that learns to compose music via the symbolic manipulation of existing melodies. Sample compositions, evidence of learning, and the process by which the machine determines a melody to be fit are presented.

Categories and Subject Descriptors: J.5.7 [**Computer Applications**]: Arts and Humanities—*Music*

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Algorithmic composition, Common Lisp, computer music, genetic algorithm, JFugue, learning machine, search heuristic

ACM Reference Format:
Bellinger, E. P. 2011. Little Ludwig, an evolutionary learning machine for musical composition. J. ACM 1, 1, Article 1 (May 2011), 7 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

The Little Ludwig is a Common Lisp program written for the purpose of musical composition. The Little Ludwig evolves random melodies into pieces containing elements that mimic those found in known music. On average, each population of generated compositions tends to contain better melodies, which suggests that the Little Ludwig learns. Not limited to music, the machine works with any list of symbols, such as a sequence of numbers, the alphabet, or a paragraph of words.

2. BACKGROUND

The famous Greek philosopher Pythagoras of Samos takes credit for the discovery of mathematically quantified harmony. Vibrating musical objects sounded good to Pythagoras when played along with other musical objects vibrating at frequencies of simple ratio multiples of the first. “In the teachings of Pythagoras and his followers,

The author would like to thank Craig Graci.
Author’s address: E. Bellinger, Computer Science Department, State University of New York at Oswego.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

music was inseparable from numbers, which were thought to be the key to the whole spiritual and physical universe. So the system of musical sounds and rhythms, being ordered by numbers, was thought to exemplify the harmony of the cosmos and corresponded to it" [Grout and Palisca 1996].

Once the seed of mathematics was planted, many algorithmic procedures were devised in attempt to emulate human creativity. For example, in 1787 Wolfgang Amadeus Mozart devised *Musikalisches Würfelspiel* ("Dice Music"), a musical game to compose Minuets using a stochastic process [Ruttkay 1997]. A person seeking to employ the musical game for composition would roll dice in order to piece together musical fragments and wind up with a dance or a march.

The approaches typically employed for algorithmic composition include knowledge-based systems, grammars, mathematical models, evolutionary methods, and systems which learn [Papadopoulos and Wiggins 1999]. Knowledge-based systems operate on a rule set to make justified decisions by employing knowledge acquisition and inference mechanisms. Grammars come from the examination of music as a language, and give rise to programs that generate music corresponding to that grammar. Because composition is difficult, many of the aforementioned approaches are often utilized simultaneously. The Little Ludwig is a hybrid system that implements a coordination of mathematical, evolutionary and learning methods.

3. RELATED RESEARCH

The Musical Weighted Synchronous Calculus of Communicating Systems (MWSCCS) and the Cybernetic Composer are two examples of implementations of mathematical models using equations and random events to generate music [Ames and Domino 1992]. MWSCCS uses a process algebra for stochastic composition [Ross 1995]. The Cybernetic Composer deduces rhythms of a melody by employing the chain process developed by Andrey Markov for probabilistically generating future states based on the present state [Ching and Ng 2006]. In contrast, Gibson and Byrne [1991] created NEUROGEN, a program that produces "small diatonic, western-type, four-part harmony compositions" derived from musical fragments by training a set of neural networks and holding partial fragments of music in genetic algorithms.

David Cope's Experiments in Musical Intelligence (EMI) software [1989] is perhaps the most famous example of algorithmic composition. It has produced music in many genres and has even composed full operas. Although he later changed direction, Cope began his project by coding rules of part-writing. This approach is one that the Little Ludwig avoids entirely. Instead, the Little Ludwig operates on a similar premise to those found in later developments of EMI. They are similar in the respect of operating on the analysis of existing music.

4. APPROACH

Many problems of life are hard to solve directly and require a search in order to find the best solution. Often the solution domain is far too vast to weigh all options, so searches must be cleverly designed to rapidly find optimal candidates. Genetic algorithms are heuristics that model the processes of nature for the purpose of expediting search. In this case, the search is for music.

Several facets characterize evolution. Heredity is the passing of characteristics from parent to offspring. Selection is when individuals are assessed by their fitness and chosen for breeding. Mutation is when one characteristic of an individual randomly changes. To model music, a melody is an individual, and its characteristics are its constituent notes. A population is a collection of melodies.

The Little Ludwig undergoes two phases. The first is the "inspiration phase" and consists of the Little Ludwig building a relational vector for the input inspiration piece. A relational vector is simply a list of each symbol's probability of being n positions away from any other symbol.

The next phase is the "composition phase." It begins with the generation of an initial population containing melodies whose notes are randomly chosen. Each melody is assessed for fitness. In order to determine the next generation, a small, random group of these melodies are picked out, and the most fit within this small group either copies itself or mates to produce offspring. After the copy or child randomly mutates, it waits in the next population while the selected group returns to the current population. This selection process continues until the next generation is full of individuals, and then the

procedure repeats [Sivanandam and Deepa 2007]. It is easy to see how this procedure facilitates the production of fitter individuals.

The fitness of each composed melody is determined by the probability of its note arrangement occurring given the compositional influence provided by the user. When the Little Ludwig composes music, it operates on the assumption that context dictates the note placement within a piece. In other words, a note fits within a melody if it has a high probability of occurring when the notes surrounding it also occur. More specifically, the fitness of an individual is the dot product between the relational vectors of the inspiration piece and that individual.

5. KNOWLEDGE REPRESENTATION

The Little Ludwig can operate on any list of symbols. In fact, the program was constructed with Common Lisp, a list processing language. For musical composition, individuals are melodies modeled as lists of notes. The symbols used for notes can be those of JFugue, an open source Java library for music programming [Surhone et al. 2010]. In this representation, the notes are written as their pitch, such as C, D, E, etc., but have the capacity to represent complex musical notions by simple concatenation of more information. For instance, to play C for a whole beat, one uses the symbol Cw. To put that note in the fifth octave, one writes C5w. The first bar of "Twinkle Twinkle Little Star" may be represented as this list: (C5q C5q G5q G5q A5q A5q G5h). A list of any other symbols works just as well. For example, consider this list of words: (this list of words). The Little Ludwig is able to process that list and generate a new list of words with similar positional qualities.

6. RESULTS

Due to its nature, many hundreds of thousands of pieces composed by the Little Ludwig exist. Figures 1 and 2 show two examples of such compositions in the form of JFugue music strings. In both of these cases, Bach's Crab canon was used as the inspiration piece. Table 1 illustrates that the Little Ludwig learns over time by presenting the fitness of the least and most fit individuals as well as the mean fitness of the entire population for each generation in one sample run of the program. For

aural examples, please visit the YouTube channel for the Little Ludwig, found at http://www.youtube.com/user/TheLittleLudwig.

D5H G5I E5H E5Q A5Q G5I E5I F5Q F5I A5Q D6Q G5Q C#5H E5I RQ C#5Q A5I E5H BB5I AB5H D5Q A3Q F5I EB5Q F5I E5H G5Q F5H C#6I RQ G5I E5H F6I F5H C#5Q BB5I BB5H A5I F5H A3Q D6I A5I E5H B5I A5I F#5H E5I D6I BB5I EB6I BB5I F6I RQ C#6I EB6I E5I B5I A5Q EB6I G6I F6I EB6I EB5Q A3Q F5I E6I EB5Q A5I F5H G6I EB6I RQ F5H A3Q B5I A5I A5H E6I E5I A5Q E6I D5Q B5I C6I C#6I D5H BB5H BB5I D6I D6Q C6I E5H D5Q

Fig. 1: The most fit melody composed in the first generation of one sample run of the Little Ludwig when using Bach's Crab canon as inspiration. This melody scored a fitness of 0.091179.

D5H A5Q E5H A5Q A5Q G5I A5I F5Q F5I A5Q G5I G5Q A5I F5I G5I A5I A5I A5I BB5I A5I A5I G5I F5I A5I F5I A5I A5I A5I A5I E5I G5I G5I A5I A5I A5I BB5I A5I A5I A5I F5I D6I A5I A5I A5I A5I G5I E5I A5I BB5I G5I BB5I A5I A5I C#6I A5I BB5I A5I G5I A5I A5I BB5I D6I E6I D6I D6I C#6I A5I E6I D6I D6I C#6I G6I A5I A5I BB5I D6I E6I E6I A5I D6I E6I F6I D6I E6I D6I D6I F6I F6I A5Q F6I E6I D6I D5Q

Fig. 2: The most fit melody composed in the 99th generation in the same sample run. This melody scored a fitness of 0.555080.

7. DISCUSSION

The true fitness of any aesthetic object is determined by the reaction of humans. The project was a success not only in the sense that it produced music, but also in that it has received praise from many humans. Within its short lifespan (on the order of days), the music produced by the Little Ludwig has gathered tens of thousands of hits and many expressions of approval and acclaim on the internet.

8. FUTURE WORK

Although the outlined tasks were satisfied, the Little Ludwig is far from being satisfactorily finished. The most immediate future plan for the Little Ludwig is composition based on several inputs simultaneously. This will produce unique music that no longer serves as just an approximation of the input piece. Additionally, perhaps another fitness metric will be applied in tandem with the present one in order to gauge some qualities that are specifically musical as opposed to positional.

Table I. Least fit, mean fitness and most fit scores for populations 0 through 19 and 80 through 99 of one sample run of the Little Ludwig, which serves to demonstrate that over time the Little Ludwig produces melodies with greater fitness.

Population	Least Fit	Mean Fitness	Most Fit
0	0.0632931	0.08677963	0.12921695
1	0.08464871	0.108017646	0.14593005
2	0.09996905	0.12623337	0.15985763
3	0.120705664	0.14439183	0.17579697
4	0.1448468	0.16161251	0.183225
5	0.1536676	0.1730749	0.19312906
6	0.16728567	0.18336274	0.1999381
7	0.17517796	0.19051532	0.20535438
8	0.18276076	0.19673166	0.20999691
9	0.18446301	0.20458217	0.21959145
10	0.19127205	0.21049212	0.23166202
11	0.20612814	0.21844167	0.23800682
12	0.21077067	0.22817549	0.24388734
13	0.224853	0.23700093	0.2458991
14	0.22887652	0.24135098	0.2547199
15	0.23166202	0.24527392	0.25998142
16	0.23939957	0.2527886	0.26586196
17	0.24110183	0.25811204	0.27158773
18	0.25162488	0.26425874	0.2821108
19	0.2534819	0.26911482	0.28567007
...	...	...	...
80	0.63463324	0.6603884	0.6895698
81	0.6482513	0.67160475	0.6895698
82	0.6590839	0.6792015	0.6895698
83	0.66279787	0.68247604	0.69452184
84	0.66976166	0.6853482	0.6960693
85	0.6660477	0.68686014	0.7011761
86	0.6744042	0.6904596	0.70504487
87	0.6753327	0.6954751	0.7124729
88	0.6816775	0.6980857	0.71711546
89	0.6799752	0.7050665	0.727948
90	0.68678427	0.711074	0.73506653
91	0.69823587	0.71727175	0.73723304
92	0.69947386	0.72307336	0.7389353
93	0.70643765	0.72943825	0.7482204
94	0.7175797	0.73623025	0.7596719
95	0.7172702	0.742038	0.7596719
96	0.7250077	0.7481058	0.76617146
97	0.72872174	0.7537728	0.76617146
98	0.7321263	0.75746673	0.78675336
99	0.7426493	0.76360726	0.7930981

9. CONCLUSIONS

The Little Ludwig is able to abstract qualities from inspiration pieces and compose new pieces that improve over time with respect to its fitness metric. Some humans have validated the aesthetic qualities of these pieces, which suggests that the fitness metric somewhat mimics the manners in which some humans gauge the quality of music. All in all, the project was a success.

ACKNOWLEDGMENT

The author would like to thank Craig Graci.

REFERENCES

AMES, C. AND DOMINO, M. 1992. *Cybernetic composer: an overview*. MIT Press, Cambridge, MA, USA, 186–205.

CHING, W. AND NG, M. 2006. *Markov chains: models, algorithms and applications*. International series in operations research and management science. Springer.

COPE, D. 1989. Experiments in musical intelligence (emi): Non-linear linguistic-based composition. *Interface 18*, 117–139.

GIBSON, P. AND BYRNE, J. 1991. Neurogen, musical composition using genetic algorithms and cooperating neural networks. *Second International Conference on Artificial Neural Networks*, 309–313.

GROUT, D. AND PALISCA, C. 1996. *A History of Western Music* 5 Ed. W. W. Norton and Company, New York.

PAPADOPOULOS, G. AND WIGGINS, G. 1999. Ai methods for algorithmic composition: A survey, a critical view and future prospects. *AISB Symposium on Musical Creativity*, 110–117.

ROSS, B. 1995. A process algebra for stochastic music composition. *Proceedings of the 1995 International Computer Music Conference*.

RUTTKAY, Z. 1997. Composing mozart variations with dice. *Teaching Statistics 19*, 18–19.

SIVANANDAM, S. AND DEEPA, S. 2007. *Introduction to genetic algorithms*. Springer.

SURHONE, L., TENNOE, M., AND HENSSONOW, S. 2010. *Jfugue*. VDM Verlag Dr. Mueller AG & Co. Kg.

Received May 2011; revised ???; accepted ???