A Neural Probabilistic Model for Music Prediction

Srikanth Cherla
Supervisors:
Dr. Artur Garcez
Dr. Tillman Weyde

BCS Doctoral Consortium
May 16, 2013

CITY UNIVERSITY LONDON

Outline

Introduction: Language Modelling & Music Prediction

Background: Computing Music

Approach: Learning pitch sequences with RBMs

Initial Results: Improved pitch prediction

Next

Introduction: Language Modelling & Music Prediction

Background: Computing Music

Approach: Learning pitch sequences with RBMs

Initial Results: Improved pitch prediction

Statistical Language Modelling

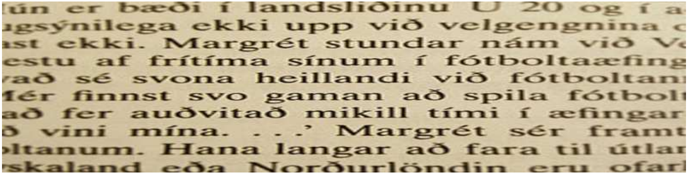
ún er bæði í landsliðinu U 20 og í a
gsýnilega ekki upp við velgengnina
ast ekki. Margrét stundar nám við V
estu af frítíma sínum í fótboltaæfing
að sé svona heillandi við fótboltann
ér finnst svo gaman að spila fótbolt
að fer auðvitað mikill tími í æfingar
ð vini mína. . . .' Margrét sér framt
ltanum. Hana langar að fara til útlar
skaland eða Norðurlöndin eru ofarl

- Modelling sequences of words in text.
- Understanding properties of language through documents.
- Applied in speech recognition, document analysis & classification.

Music Prediction

- Modelling sequences of *musical events* (musical pitch, note duration, etc.).
- Understanding properties of music through scores.

Relevance

Scientific:

- Analysis of
  - Compositional practices
  - Musical style
- Music Education & Training
- Organizing music data
- Aiding music transcription
- Music & emotion

Creative:

- Generative music
- Compositional assistance

Next

Notable Approaches

- Markov models to carry out a “musical conversation”.
- Rule-based systems to imitate classical composers.
- Neural networks to play like Charlie Parker.
- Genetic Algorithms to improvise Jazz solos.
- Hidden Markov models to harmonize chorales like Bach.

Next

Introduction: Language Modelling & Music Prediction

Background: Computing Music

Approach: Learning pitch sequences with RBMs

Initial Results: Improved pitch prediction

Restricted Boltzmann Machine

- A bipartite graphical model for learning probability distributions
- Adjusts weights according to data while learning.
- Data presented at visible layer. Features learned in hidden layer.
- Generative model.

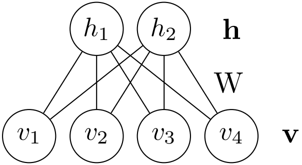

Why RBMs?

RBMs have some desirable qualities for learning from music data

- Reasonable predictions for unseen cases.
- Non-exponential scaling with sequence length.
- Efficient learning algorithms.
- Scalable to deeper architectures.
- Generative properties.

Neural Probabilistic Music Prediction

A simple model based on the Restricted Boltzmann Machine for learning sequences of musical pitch.

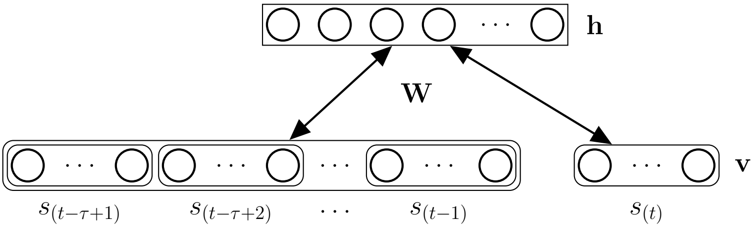

The present approach motivated by previous work in language modelling.

Next

Introduction: Language Modelling & Music Prediction

Background: Computing Music

Approach: Learning pitch sequences with RBMs

Initial Results: Improved pitch prediction

Sequence Length

Makes use of information in longer pitch sequences.

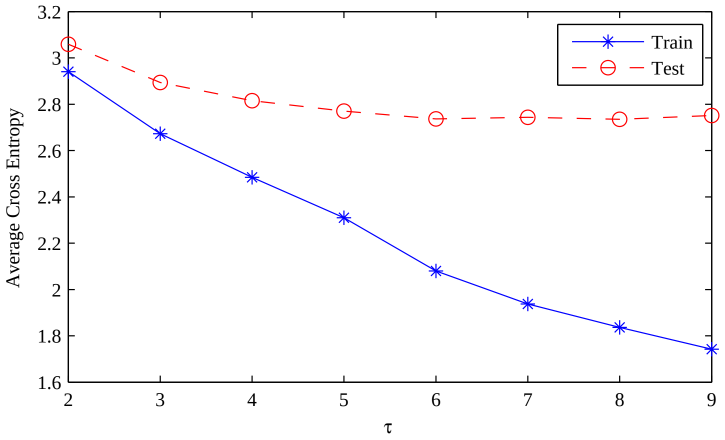

Prediction Performance

Better cross entropy estimates.

ID	Markov-3	Markov*	RBM-2	RBM-3	RBM-4	RBM-5	RBM-6	RBM-7	RBM-8	RBM-9
0	2.885	2.861	3.062	2.895	2.812	2.787	2.783	2.770	2.772	2.793
1	2.544	2.444	2.732	2.554	2.510	2.465	2.442	2.442	2.455	2.465
2	3.110	3.115	3.174	3.030	2.935	2.894	2.852	2.865	2.848	2.849
3	2.791	2.721	2.886	2.738	2.666	2.672	2.641	2.703	2.741	2.802
4	3.192	3.010	3.202	3.069	2.974	2.867	2.822	2.826	2.838	2.829
5	3.385	3.340	3.475	3.286	3.203	3.140	3.107	3.117	3.026	3.070
6	2.596	2.428	2.740	2.597	2.519	2.461	2.432	2.411	2.379	2.383
7	3.083	3.105	3.196	2.980	2.899	2.878	2.815	2.812	2.817	2.817
Av.	**2.948**	**2.878**	**3.059**	**2.894**	**2.815**	**2.771**	**2.737**	**2.743**	**2.734**	**2.751**

Case Study:Folk Melody Classification

Overall accuracy of 59.93%

	Nova-Scotia	Alsace	Yugoslavia	Switzerland	Austria	Germany	China
Nova-Scotia	74.83%	5.96%	0.00%	2.65%	4.63%	4.63%	7.28%
Alsace	13.19%	29.67%	7.69%	10.99%	19.78%	16.48%	2.19%
Yugoslavia	10.08%	10.08%	42.02%	14.29%	14.29%	8.40%	0.84%
Switzerland	6.45%	17.20%	6.45%	35.48%	22.58%	7.53%	4.30%
Austria	10.58%	19.23%	4.81%	12.50%	41.35%	9.62%	1.92%
Germany	11.27%	11.74%	2.35%	9.86%	8.92%	53.99%	1.88%
China	6.75%	1.27%	0.00%	0.84%	3.80%	0.00%	87.34%

Conclusions & Future Work

Conclusion

- RBMs are a good starting point for music prediction.
    - Use longer contexts.
    - Handle unseen contexts.
    - Scale gracefully.

Some interesting directions for future work

- Extensions to harmonic sequences.
- Predicting other musical dimensions.
- Learning higher-level structure.
- Other applications.

Acknowledgements

Dr. Marcus Pearce (QMUL)
Son Tran (City)

Thank you!

Questions?