

MELODY MODELLING WITH NEURAL NETS

Srikanth Cherla^{1,2}, Artur Garcez¹ and Tillman Weyde^{1,2}

{ SRIKANTH.CHERLA.1, A.GARCEZ AND T.E.WEYDE }@CITY.AC.UK

¹MACHINE LEARNING GROUP, ²MUSIC INFORMATICS RESEARCH GROUP



CITY UNIVERSITY
LONDON

INTRODUCTION

- Learning sequential information in Melodies.
- Working with symbolic music data (MIDI, GP4, etc.).
- Neural networks predict probability of the pitch of a note given information about the preceding notes.
- Applications: computational musicology, music creation, generating musical stimuli, music education.


MULTIPLE VIEWPOINTS REPRESENTATION OF MELODY [2]

- An event-based representation of music.
- *Viewpoint type* (feature) sequences extracted from score.




Viewpoint	Transformed sequence									
pitch	67	69	71	72	69	72	64	67	72	69
int	⊥	2	2	1	-3	3	-8	3	5	-3
onset	0	2	5	6	9	10	12	15	16	20
ioi	⊥	2	3	1	3	1	2	3	1	4
int ⊗ ioi	⊥	2, 2	2, 3	1, 1	-3, 3	3, 1	-8, 2	3, 3	5, 1	-3, 4

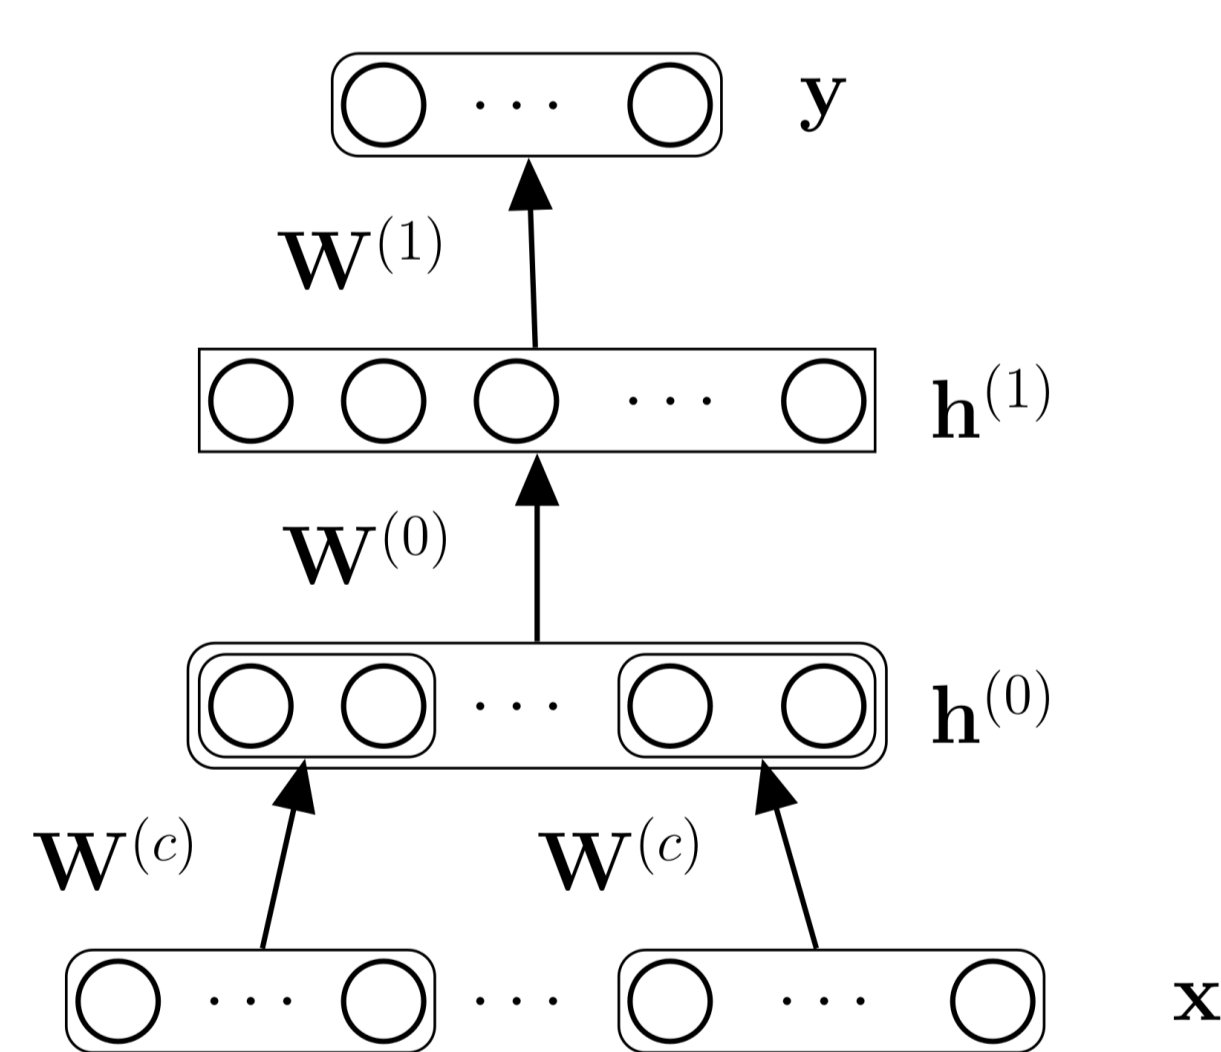
PREDICTION MODELS



(a) Restricted Boltzmann Machine [3]



(b) Feed-forward Neural Network [5]




(c) Neural Probabilistic Language Model [1]

- Input: Concatenated one-hot vectors
- Hidden: *logsig* or *tanh* units
- Output: Softmax

EVALUATION

- Evaluated using cross entropy on 185 Bach chorales (total 9224 notes).
- Models compared with variable order Markov Models [4].
- Each model performs better than VOMM for a context length of 1.
- Feed-forward networks with more layers perform worse.



FUTURE WORK

- Best way to represent missing values in context.
- Unbounded context models - recurrent neural networks.
- Adapting predictions by learning as melody progresses.
- Representing polyphony with multiple viewpoints.

REFERENCES

- [1] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A Neural Probabilistic Language Model. *Journal of Machine Learning Research*, 3:1137–1155, 2003.
- [2] Darrell Conklin and Ian H Witten. Multiple viewpoint systems for music prediction. *Journal of New Music Research*, 24(1):51–73, 1995.
- [3] Hugo Larochelle and Yoshua Bengio. Classification using discriminative restricted Boltzmann machines. In *International Conference on Machine Learning (ICML)*, pages 536–543. ACM Press, 2008.
- [4] Marcus Pearce and Geraint Wiggins. Improved methods for statistical modelling of monophonic music. *Journal of New Music Research*, 33(4):367–385, 2004.
- [5] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. *Learning representations by back-propagating errors*. MIT Press, Cambridge, MA, USA, 1988.