Markov Networks for Free Improvisers

Stefano Kalonaris
Sonic Arts Research Centre
Queen's University, Belfast
Skalonaris01@qub.ac.uk

ABSTRACT

This paper discusses the use of probabilistic graphical models (PGMs) for initiating dynamical human musical interactions, in the context of free improvisation. This study proposes the model of Markov Networks and it speculates how they may serve for forming dynamical sepsets amongst players, based on their reciprocal beliefs, expressed as Bayesian inference. The prior is an assigned, private, musical personality. The players communicate their affinity preferences over a computer network using a graphical user interface. The conclusion is that Markov Networks viewed as dynamical Bayesian games are employable in the context of free improvisation and distributed creativity, providing a useful (and conceptually dissimilar) alternative to other structures that have been employed in music improvisation, such as graphic scores and idiom-based improvised forms.

Keywords: Probabilistic Graphical Models, Markov Networks, Music Games, Free Improvisation.

1. INTRODUCTION

The purpose of this research is to apply models of dynamical structural organisation based on probabilistic graphical models (PGMs) and Bayesian inference to game-based approaches in free improvisation. Although models derived from statistical, economic and computational sciences have been employed successfully in the areas of composition [1], algorithmic composition [2], [3] and machine improvisation [4], [5], there seems to be a shortage of studies that have addressed real-time interactions inspired by such mathematical and computational models in the context of free improvisation. Moreover, even in such cases, the conceptual framework has been that of a free improviser playing along/with an intelligent artificial counterpart. This paper proposes a model in which all players are human, a model that retains the performers' agency in the musical output, and where the machine is used for interfacing tasks only, as to provide a communication network within which the players operate. I claim that there is little if no historical precedent in this direction as all examples of Markov Networks (MNs) applied to music have been and are to be found in the areas of artificial intelligence and machine learning, be it applied to automatised generation of musical material (often *in the style of*) [6], [7], [8], [9], modelling musical structure [10], [11], statistical methods for audio processing [12], [13] and music information retrieval [14], amongst others. In contrast to these applications, I propose an abstraction for and between human players which is realised in real-time and, ultimately, with an associated freely improvised output.

2. FREE BAYES

Free music improvisation entails a high degree of dynamical shuffling of roles, which allows the participants to shape interactions in real-time and to react to unforeseen circumstances with split-second decision-making wizardry. This requires both the ability to make sense of the information available to them at any given time as well as the capacity to store and edit such information and beliefs in order to respond to their best. Such responses are based on the evaluation and the inferential analysis of the contextual evidence players are presented with. Such evidence is not immutable and static but, on the contrary, malleable and dynamic. Put simply, any improviser at any given point in time is actuating musical strategies that result from what she believes is happening or is going to happen in the near future. As soon as the player is provided with new evidence, she will adjust her response accordingly. This is analogous to what, in probability theory, is defined by the Bayes rule.

$$P(cause|observation) = \frac{P(observation|cause)P(cause)}{P(observation)} \tag{1}$$

The above reads: the probability of a cause, given the observation of an event, equals to the probability of the event given the cause, times the probability of the cause, all divided by the probability of the event.

3. MARKOV NETWORKS

Markov Networks (MNs) are undirected and possibly cyclic graphs. In the context of the natural interactions between free improvisers, MNs are more appropriate than directed graphs (Bayesian Networks), as they allow influences and inferences to flow in both directions. A formal definition can be stated as follows: a Markov Network is a random field *S*, which is a collection of indexed random variables (either discrete or continuous) where any variable $\boldsymbol{F_i}$ is independent of all other variables in *S*. Such a network also satisfies the Markov property, which states that no matter what path the system took to get to the current state, the transition probability from that state to the next will be independent from such path.

$$p_{(i,j)} = P(X = (n+1) = j \mid X(n) = i, \quad X(n\text{-}1), \ldots, X(0)) \tag{2}$$

The simplest class of MNs is the pairwise MN, an example of which is depicted in the following example:

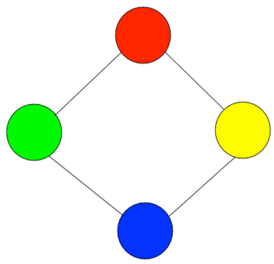

Figure 1. A pairwise Markov Network

In my musical implementation, the nodes of above graph represent four players, which, by virtue of playing together, influence each other. Since there is no strictly conditioning and/or conditioned variable, as one would have in a Bayesian Network, the notion of factor, hereinafter indicated as φ, will come in handy for defining the interactions between the nodes (players). Factors also go under the names of affinity functions, compatibility, soft constraints, and they generalise the idea of the local predisposition and willingness of any pair of nodes to take a joint assignment.

$\varphi_1[Red, Yellow]$

R^0	Y^0	30
R^0	Y^1	5
R^1	Y^0	1
R^1	Y^1	10

Table 1. Example of a local distribution amongst player Red and player Yellow

The above are arbitrary values and are chosen for illustration purposes only. The binary superscript (either zero or one) for the Red and Yellow players, in the respective columns, indicate their willingness to undertake a joint assignment with the other, or not. In the above example, the strongest factor indicates that neither Red nor Yellow would prefer to cooperate, talk to each other, play with each other, etc.
Similarly, one can imagine the other three local factors φ_2, φ_3, φ_4 and the probability distribution over the depicted pairwise MN would be:

$$\tilde{P}(Red, Green, Blue, Yellow) = \varphi_1(Red, Yellow)\varphi_2(Yellow, Blue)\varphi_3(Blue, Green)\varphi_4(Green, Red)$$

The above is not a proper probability distribution, since the sum over all the marginal distributions does not equal to one. In order to obtain a proper distribution, one needs to normalise by diving by the partition function Z.
The partition function is expressed as follows:

$$Z = \sum_{x \in \chi} \exp\left(\sum_k w_k^T f_k(x_{\{k\}})\right) \tag{3}$$

After having obtained the probability distribution, one can observe that the local preferences are no longer represented, as they have all been affected by the propagation of beliefs of all players over the network.
Put simply, even a four-player structure as this one, ends up in a complex aggregate of all the different factors that compose the MN. This is in contrast to what occurs in Bayesian Networks, where it is possible to inspect the probability distribution and retrieve a local factor. Pairwise MNs are not fully expressive and they are insufficient and inappropriate for representing all possible interactions. A more expressive model, used in my musical translation for improvisers, is the induced MN. In this model, each general factor ϕ has a scope that might contain more than two variables (as opposed to pairwise interactions).
A Gibbs Distribution is parameterised over a set of factors Φ, where

$$\Phi = \{\phi_1(D_1), \ldots, \phi_k(D_k)\} \tag{4}$$

The un-normalised probability distribution will be:

$$\tilde{P}_{\Phi}(X_1, \ldots, X_n) = \Pi_{i=1}^{k}\, \phi_i(D_i) \tag{5}$$

Whereas the normalised probability distribution will be expressed by:

$$P_{\Phi}(X_1, \ldots, X_n) = \frac{1}{Z_{\Phi}} \tilde{P}_{\Phi}(X_1, \ldots, X_n) \tag{6}$$

Where

$$Z_{\Phi} = \sum_{X_1, \ldots, X_n} \tilde{P}_{\Phi}(X_1, \ldots, X_n) \tag{7}$$

It is now possible to express a much wider range of scenarios, which might involve factors over three or more variables.

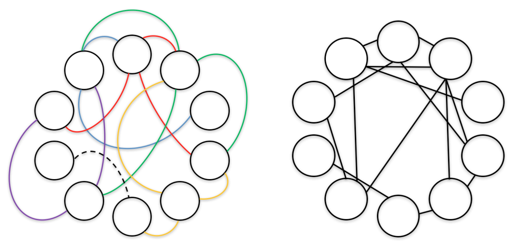

Figure 2. From local factors to induced Markov Network

Simply put, two or more variables (players, in this case) are connected whenever they appear in the same scope of a given factor. However, it would be impossible to infer the factorisation from the graph. In this sense, influence can flow along any active trail/edge. I find this model an exquisite abstraction of a typical interactive and dynamically assigned scenario amongst music improvisers, where alliances and joint assignments are formed, undertaken, updated, abandoned, in continuous real-time. Having understood the workings of an induced MN, I will now present my original rendition in musical form, as a dynamical model of interactions amongst free improvisers.

4. MN FOR FREE IMPROVISERS

4.1 Motivation

The decision of employing MN as a model for improvised musical interaction follows on previous experiments of mine, carried out in regards to focal points, Schelling's salience [15] and Markov Chains. The aforementioned experiments[1], pointed at the need to move towards an increased complexity of inter-relations and a decreased complexity of the instructions/constraints, as a step in the direction of allowing for more prompt and reactive environments for the player to operate in. Unlike the literature that has dealt with models based on either probabilistic graphical models or automata theory, Markov Network for Free Improvisers (hereinafter MN4FI) is an abstraction for and between human players and no musical output is generated by the machine. The hypothesis to be tested is to whether this given model provides alternative dynamical and interactive opportunities and modalities to groups of free improvisers, while maintaining freedom and flow in the performance.

4.2 System Design and Interaction Model

MN4FI follows directly from the example above regarding an induced MN. It formally maps players to a type and each type to a set of weighted strategies, or affinity preferences. The potential of this model lies in the fact that local distributions are not reflected or retrievable by the global graph and in that each of the players' screen might depict a different locality of connections.

The interaction model is described by a graph with up to ten vertices, each representing a different player. Each player is assigned a musical personality, what in Bayesian terms would be referred to as a *type*. Such type is private information and it is not shared amongst the players. This very fact implies that particular care needs to be taken when deciding on the spatial physical distribution of the players, in that they ought not to be able to see each other's screen.

Each of these four different types has an optimal local pairwise counter type, and ideally each player will try to infer the others' type in order to achieve such optimal joint assignment. Each player's degree (the number of other players he/she is connected to) is capped to *n-2* with *n* being the total number of players. Players can only musically interact with players they are connected to. The structuring principle consists in that each local graph might differ from any other, and the induced MN will not be common knowledge, nor will the local factor be retrievable should one be able to observe the resulting induced MN. Players are free to revoke one connection at a time, thus regaining the faculty of initiating a different one, if they wish. They instantiate and revoke edges according to their beliefs about the players they are connected or want to connect to. Additionally, players can also trigger a stochastic change of their own type.

The table below describes the affinity preferences for any local pair of players.

	Cooperative	Non-cooperative	Chaotic	Solipsistic
Cooperative	100	10	40	60
Non-cooperative	10	70	60	80
Chaotic	40	60	100	80
Solipsistic	60	80	80	100

Table 2. Inter-type local strategy matrix

It is important to note that the above numbers are arbitrary and the table is clearly not normalised.

4.3 Implementation and individual modules

MN4FI is realised in the programming language Max (`http://www.cycling74.com`). It follows a centralized design, consisting of one module for the players and one module which acts as a hub, receiving and despatching requests over a custom network, using the OSC protocol. There are *n* workstations, one for each player. At present, MN4FI accommodates up to ten vertices. In the presence of more than ten players, two or more of them can cluster around one workstation, thus sharing the screen, the assigned type, and the responsibility of connecting and removing edges.

In terms of actual coding, the core of the player's interface is realised using JavaScript within Max. This allows for a dynamical instantiation of the graph, depending on the number of players performing. Players connect and disconnect to and from a vertex by means of the numerical keyboard or using their GUI. They can also operate a trigger, which randomly reassigns their type. It is worth reminding that such type is private information.

The player patch has been compiled into a standalone application, in order to ensure that all players can run the interface, regardless of whether they have Max or not. Such consideration stemmed from the necessity to widen participation beyond limits of economical nature, Max being proprietary software. Each player's node will appear in red on their respective graph, and each node they are connected to will be coloured green. Else, the disconnected nodes appear in yellow. The GUI shown in fig.3 is what any given player sees and interacts with.

[1] partially available online at:
http://www.ransompaycheque.com/the-brazilian-games
http://www.ransompaycheque.com/finite-state-machines

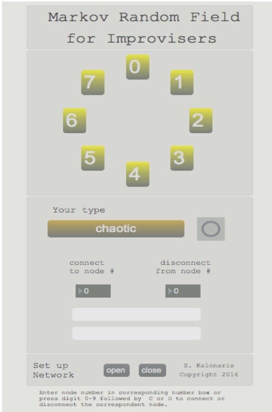

Figure 3. The player's GUI

4.4 Evaluation

MN4FI was first played during the visit of Amsterdam-based duo Shackle at the Sonic Arts Research Centre (SARC) in Belfast on 03.12.2015, during which I had the opportunity to try the model out. At the time, MN4FI was implemented rather differently and it did not accommodate more than four players. MN4FI was subsequently reworked and tested with some of the members of QUBe, the resident experimental music ensemble of SARC. This time, MN4FI was played by eleven players. Both sessions have been recorded in audio and video format, and can be found online at the following website:

`http://ransompaycheque.com/markov-random-fields`

Fourteen players completed an evaluation form, as well as participating in short focus group discussions, over the course of the two instances. These were both valuable tools for obtaining feedback and suggestions for improvements, with respect to aesthetic, artistic and technological considerations. Given the small size of the sample (fifteen players in total), this paper can by no means claim to be conclusive or statistically significant.
The results obtained are, rather, a way to inform the next steps for the development of MN4FI.
The evaluation form is divided into three sections, each containing multiple questions to which the player can answer categorically on a Likert scale in 5 levels (from 'strongly disagree' to 'strongly agree', re-coded to 1-5). Players' proficiency was also reported in 5 categorical levels (from 'none' to 'expert'). The three sections present questions that address the degree of freedom experienced within the model, the degree of satisfactory output perceived, and how appropriate the design of the GUI is deemed, respectively. The answers collected pointed to the need of rehearsal time dedicated to familiarise and operate the GUI whilst maintaining the flow of musical improvisation. This is particularly true with respect to players who do not normally include electronics or other interfaces in their artistic practice. Proficiency levels were almost exclusively distributed between 'good' and 'expert' with 38.5% and 46.1% respectively. The remainder was evenly split between 'none' and 'proficient'. No player self-reported their proficiency as 'fair/basic'.

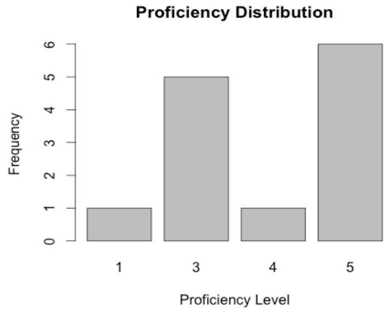

Figure 4. Proficiency levels

Levels of freedom experienced in playing this model were evenly distributed amongst level 3, 4 and 5, at 0.308, 0.385 and 0.308 respectively, leaving out the categories 1 and 2, which would correspond to a lower perceived freedom.

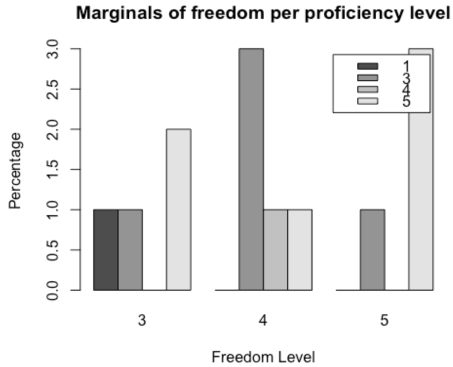

Figure 5. freedom with respect to proficiency levels

It appears clear that more experienced players had better chances to navigate the model with a higher likelihood of experiencing flow and un-hampered creativity in their performance. Overall, there was a consensus of the positive experience that all participants had of the piece.

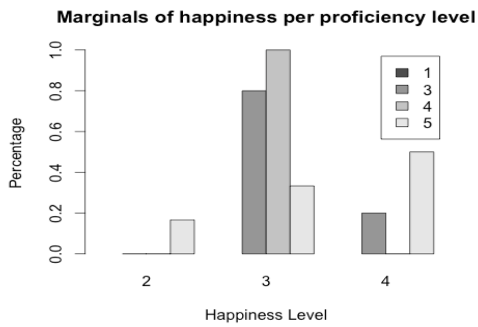

Figure 6. happiness with respect to proficiency levels

With regards to the evaluation in terms of inspiration for new ideas and interactions, the following are the percentages:

2	3	4	5
0.23076923	0.23076923	0.46153846	0.07692308

Table 3. Marginals for novel interaction

As seen from above, nearly 70% of the participants felt that the model suggested new and non-normative ideas (level 3 and 4). Furthermore, over 84% of the participants reported being happy and satisfied with the musical outcome.

From examining the correlation matrix for some variables in both the freedom and the output section, one can notice that, unsurprisingly, the strongest correlations are between freedom and constraint, proficiency and constraint, and between freedom experienced and the willingness to play again according to the model.

	Proficiency	Freedom	Novelty	Constraint	Play again
Proficiency	1.0	0.319	0.323	-0.412	-0.104
Freedom	0.319	1.0	-0.319	-0.795	0.532
Novelty	0.323	-0.319	1.0	0.022	0.234
Constraint	-0.412	-0.795	0.022	1.0	-0.334
Play again	-0.104	0.532	0.234	-0.334	1.0

Table 4. correlation matrix

The most valuable finding was, however, to be had during the focus group discussion, where it was reported that MN4FI encouraged a type of behaviour that was atypical, with respect to the simultaneous focus on both inner clusters of musical interaction and the global musical outcome.

" *I think it encourages a lot more interaction, like whenever we just play free improvisation people tend to go in their own wee world sometimes whereas with this, it kind of focuses you more on the fact that there are other people around you, also playing, and you have to listen to them".*

"Yes, it forms subgroups within something larger that is going on".

(QUBe members, focus group discussion, 23.02.2016)

5. CONCLUSIONS

In this work I have shown an equivalence between probabilistic graphical model based structures and Bayesian games in the context of a real-time interaction network amongst free improvisers. By implementing and testing a Markov Network as the determining structure for forming or abandoning musical local relationships amongst the performers, I have been able to show that insights from one area (PGMs) may be applied to the other (musical free improvisation) to provide an alternative and artistically valid and satisfactory modus operandi. I believe that this result is particularly exciting as it opens up numerous possibilities of intersection between free improvisation and paths that have so far been exclusive to the domains of decision theory, propositional logic and artificial intelligence. I claim to have employed a methodology that asserts the real-time human interaction as paramount, much in contrast to the uses of Markov Processes that have so far informed the discourse around musical improvisation and artificial intelligence/machine learning. In the latter cases, the Markov and Bayesian processes are employed to train an intelligent and autonomous artificial agent that either interacts with the human performer or generates music *in the style of.* Future work includes extending the model to allow players to send local assignments to their sepset and/or adapting the network to include more complex rules, for example in the form of Markov Logic Network, by the introduction of first order logic.

REFERENCES

[1] Xenakis, I., 1992. *Formalized Music: Thought and Mathematics in Composition*, Pendragon Press.

[2] Cope, D., 1991. *Computers and musical style*, A-R Editions.

[3] Cope, D., 2000. *The Algorithmic Composer*, A-R Editions.

[4] Lewis, G.E., 2000. Too Many Notes: Computers, Complexity and Culture in Voyager. *Leonardo Music Journal*, 10(2000), pp.33–39.

[5] Rowe, R., 2004. *Machine Musicianship*, MIT Press.

[6] Pachet, F., 2002. The Continuator: Musical Interaction With Style. *Proceedings of the International Computer Music Conference*, (September), pp.333–341.

[7] Pachet, F., Roy, P. & Barbieri, G., 2011. Finite-length Markov processes with constraints. *IJCAI International Joint Conference on Artificial Intelligence*, pp.635–642.

[8] Allan, M. & Williams, C.K.I., 2004. Harmonising Chorales by Probabilistic Inference. *Advances in Neural Information Processing Systems*.

[9] Assayag, G. & Dubnov, S., 2004. Using factor oracles for machine improvisation. *Soft Computing*, 8(9), pp.604–610.

[10] Mavromatis, P., 2005. The Echoi of Modern Greek Chant in written and Oral Transmission: A Computational Model and Its Cognitive Implications.

[11] Mavromatis, P., 2008. Minimum Description Length Modeling of Musical Structure. Journal of Mathematics and Music, 00(00), pp.1-21.

[12] Pearce, M.T. et al., 2010. Unsupervised statistical learning underpins computational, behavioural, and neural manifestations of musical expectation. Neurolmage, 50(1), pp.302-313.

[13] Pearce, M. & Wiggins, G, 2004. Improved Methods for Statistical Modelling of Monophonic Music. Journal of New Music Research, 33(4), pp.367-385.

[14] Cemgil, A.T. et al., 2008. Bayesian Statistical Methods for Audio and Music Processing, pp.1-45.

[15] Schelling, T., 1960. The Strategy of Conflict, Literary Licensing, LLC.