

What were you expecting? Using Expectancy Features to Predict Expressive Performances of Classical Piano Music

Carlos Cancino-Chacón^{1,2}, Maarten Grachten², David R. W. Sears², and Gerhard Widmer^{1,2}

¹ Austrian Research Institute for Artificial Intelligence
carlos.cancino@cfai.at

² Department of Computational Perception, Johannes Kepler University Linz
{maarten.grachten,david.sears,gerhard.widmer}@jku.at

Abstract. In this paper we present preliminary work examining the relationship between the formation of expectations and the realization of musical performances, paying particular attention to expressive tempo and dynamics. To compute features that reflect what a listener is expecting to hear, we employ a computational model of auditory expectation called the *Information Dynamics of Music* model (IDyOM). We then explore how well these expectancy features – when combined with score descriptors using the Basis-Function modeling approach – can predict expressive tempo and dynamics in a dataset of Mozart piano sonata performances. Our results suggest that using expectancy features significantly improves the predictions for tempo.

Keywords: Musical expression, Information theoretic features, IDyOM, RNNs

1 Introduction

Computational models of musical expression can be used to explain the way certain properties of a musical score relate to an expressive rendering of the music [12]. However, existing models tend to use a combination of high- and low-level *hand-crafted features* reflecting structural aspects of the score that might not necessarily serve as perceptually relevant features. An example of such a model is the Basis-Function modeling approach (BM) [7].

To examine the relationship between the formation of expectations during music listening on the one hand, and the realization of musical performances on the other, Gingras et al. [4] employed the *Information Dynamics of Music* model (or IDyOM) [10], a probabilistic model of auditory expectation that computes information-theoretic features relating to the prediction of future events. In their study, these information-theoretic features were shown to correspond closely with temporal characteristics of the expressive performance, which suggests that the performer attempts to decrease the processing burden on listeners during perception by slowing down at unexpected/uncertain moments and speeding up at expected/certain ones.

Here we present preliminary work to support the claim that expectancy measures can inform predictions of expressive parameters related to tempo and dynamics. We extend the work in [4] in two ways. First, rather than simply demonstrating that expectancy measures are related to expressive performances, we show that the use of expectancy features improves the predictive quality of models using other score descriptors, thus providing a more comprehensive framework for the modeling of expressive performances in music of the common-practice period. Second, as opposed to *fitting* the expectancy features to each performance (i.e. training and testing the model on the same performance), the models presented in this paper are evaluated by measuring their prediction error on unseen pieces.

The rest of this paper is organized as follows: Section 2 presents our formalization of expressive parameters, describes the score and expectancy features employed in this study, and finally outlines the regression model used to predict the expressive parameters. Section 3 describes the empirical evaluation of the proposed approach, the results of which are discussed in Section 4. Finally, conclusions are stated in Section 5.

2 Modeling expressive performances

In this section we provide a brief description of the proposed framework. First we describe how expressive dynamics and tempo are encoded. Second, we describe the expectancy and score features. Finally we describe the recurrent neural network (RNN) models used to connect the input features to the expressive targets.

2.1 Targets: Expressive Parameters

An *expressive parameter* is a numerical descriptor that corresponds to common concepts involved in expressive piano performance. We take the local *beat period ratio* (*BPR*) as a proxy for *musical tempo*. We average the performed onset times of all notes occurring at the same score onset and then compute the *BPR* by taking the slope of the averaged onset times (in seconds) with respect to the score onsets (in beats) and dividing the resulting series by its average beat period. For *dynamics*, we treat the performed MIDI velocity as a proxy for loudness. We take the maximal performed MIDI velocity per score onset, divided by 127. This expressive parameter will be denoted *VEL*. To explore how well the expectancy and score features describe the *relative* changes in *BPR* and *VEL*, we also calculate their first derivatives, denoted by BPR_d and VEL_d , respectively. Furthermore, including the derivative time series allows us to compare our findings with the results obtained in [4].

2.2 Features: Multiple Viewpoints

Expectancy Features IDyOM provides a conditional probability distribution of a musical event, given a preceding sequence of events, i.e. $p(v_n | v_{n-1}, v_{n-2}, \dots)$.

Following [4], we use IDyOM to estimate two information-theoretic measures representing musical expectations:

1. **Information content** (IC). The IC measures the unexpectedness of a musical event, and is computed as $IC(v_n) = -\log p(v_n | v_{n-1}, v_{n-2}, \dots)$.
 - (a) IC_m . The information content for each melody note. This value is computed using a model that is trained to predict the next chromatic melody pitch using a selection of melodic viewpoints, such as pitch interval (i.e. the arithmetic difference between two consecutive chromatic pitches, measured in MIDI note values), and contour (whether the chromatic pitch sequence rises, falls or remains the same). IDyOM performs a step-wise selection procedure that combines viewpoint models if they minimize model uncertainty as measured by corpus cross entropy [11].
 - (b) IC_c . Estimation of the IC computed for the combination of pitch events (a proxy for harmony) at each score onset. IDyOM predicts the next combination of vertical interval classes above the bass (see Score Features 1b).
2. **Entropy** is a measure of the degree of choice or uncertainty associated with a predicted outcome. The entropy can be computed as $H(v_n) = \mathbb{E}\{-\log p(v_n | v_{n-1}, v_{n-2}, \dots)\}$.
 - (a) H_m . Entropy computed for each chromatic pitch in the melody.
 - (b) H_c . Entropy computed for the combined pitch events at each score onset.

Score Features Following [7], we include low-level descriptors of the musical score that have been shown to predict characteristics of expressive performance.

1. **Pitch.**
 - (a) $(pitch_h, pitch_l, pitch_m)$. Three features representing the chromatic pitch (as MIDI note numbers divided by 127) of the highest note, the lowest note, and the melody note at each onset.
 - (b) (vic_1, vic_2, vic_3) . Three features describing up to three vertical interval classes above the bass, i.e. the intervals between the notes of a chord and the lowest pitch, excluding pitch class repetition and octaves. For example, a C major triad (C, E, G), starting at C_4 would be represented as $(pitch_l, vic_1, vic_2, vic_3) = (\frac{60}{127}, \frac{4}{11}, \frac{7}{11}, 0)$, where 0 denotes the absence of a third interval above C_4 , i.e. the absence of a fourth note in the chord.
2. **Metrical position.**
 - (a) $b_{\phi,t}$. The relative location of an onset within the bar, computed as $b_{\phi,t} = \frac{t \bmod B}{B}$, where t is the temporal position of the onset measured in beats from the beginning of the score, and B is the length of the bar in beats.
 - (b) (b_d, b_s, b_w) . Three binary features (taking values in $\{0, 1\}$) encoding the metrical strength of the t -th onset. b_d is nonzero at the downbeat (i.e. whenever $b_{\phi,t} = 0$); b_s is nonzero at the secondary strong beat in duple meters (e.g. quarter-note 3 in $\frac{4}{4}$, and eighth-note 4 in $\frac{6}{8}$), and b_w is nonzero at weak metrical positions (i.e. whenever b_d and b_s are both zero).

2.3 Recurrent Neural Networks

RNNs are a state-of-the-art family of neural architectures for modeling sequential data. Following [1, 6], we use bidirectional RNNs as non-linear regression models to assess how well the features described above predict expressive dynamics and tempo. In this work, we use an architecture with a composite bidirectional hidden layer with 5 units, consisting of a forwards and backwards long short-term memory layer (LSTMs).

3 Experiments

We perform a 5-fold cross-validation to test the accuracy of the predictions of three models trained on different feature sets for each expressive parameter: a model trained only using expectancy features (**E**), a model trained only using score features (**S**), and a model trained on both expectancy and score features (**E+S**). Each model is trained/tested on 5 different partitions (folds) of a dataset, which is organized into training and test sets, such that each piece in the corpus occurs exactly once in the test set.

For this study we use the Batik/Mozart corpus, which consists of recordings of 13 Mozart piano sonatas (39 movements) by Austrian pianist Roland Batik performed on a computer controlled Bösendorfer SE [2]. Melody voices were identified and annotated manually in this dataset. For each fold, we use 80% of the pieces for training and 20% for testing. The parameters of the models are learned by minimizing the mean squared error on the training set³. We evaluate model accuracy with the coefficient of determination R^2 and Pearson’s r .

4 Results and Discussion

The results of the 5-fold cross-validation are shown in Table 1. To examine the differences between the R^2 values of the **E**, **S**, and **E+S** feature sets we performed a separate one-way ANOVA for each expressive parameter (BPR , BPR_d , VEL and VEL_d). These differences were statistically significant in all cases at the $p < 0.05$ level as measured by Fisher’s F ratio. The same trend emerged for most expressive parameters, with **E+S** outperforming the other models, although post-hoc pairwise comparisons using Tukey’s HSD only revealed a significant difference for BPR_d . These results therefore suggest that the models including both expectancy and score features better predict expressive tempo than expressive dynamics. Furthermore, although not directly comparable, the values for R^2 and r in Table 1 seem to be on par with those reported on Chopin piano music using the BM approach [6].

The fact that the use of expectancy features improves model performance for expressive tempo but not for dynamics might be due to the relation of expressive tempo to structural properties of the music, such as phrase-final lengthening, such as the final *ritardando* at the end of a piece [8]. Since expectation features

³ A repository containing the code is available online: https://github.com/neosatrapihereje/mml2017_expression_expectation.

Feature Set	Tempo				Dynamics			
	BPR		BPR_d		VEL		VEL_d	
	R^2	r	R^2	r	R^2	r	R^2	r
E	0.038	0.201	0.067	0.259	0.234	0.496	0.185	0.429
S	0.065	0.289	0.105	0.326	0.299	0.569	0.244	0.494
E + S	0.072	0.288	0.124	0.351	0.312	0.574	0.230	0.477

Table 1. Predictive results for expressive tempo and dynamics, averaged over all pieces on the Batik/Mozart corpus. A larger R^2 and r means better performance.

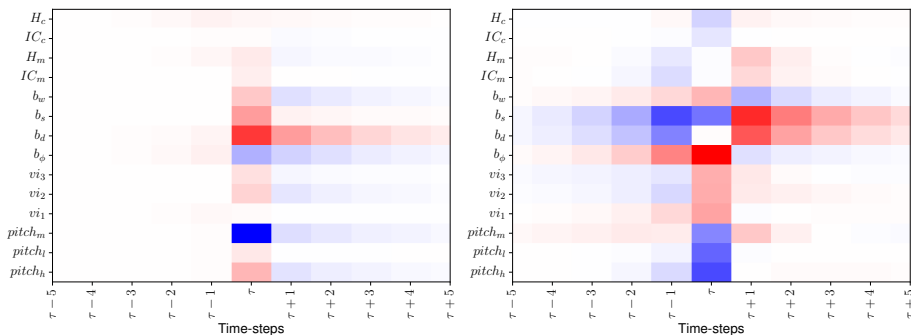


Fig. 1. Sensitivity plots for BPR (left) and BPR_d (right). Each row in the plot corresponds to an input feature and each column to the contribution of its value at that time-step to the output of the model at τ (the center of each plot). Red and blue indicate a positive and negative contribution, respectively.

also relate to music structure in the sense that music tends to be more unpredictable at boundaries between musical segments than within segments [9], this may in part explain why the models are better at predicting changes in expressive tempo BPR_d .

Figure 1 shows 2D differential sensitivity maps that examine the contribution of each feature to the output of the model trained on all features (**E+S**). Although these plots show that the score features have a more prominent role in predicting expressive tempo, as suggested by the results in Table 1, we will focus here on the contribution of the expectancy features. On the one hand, the plots suggest a tendency for the performer to slow down if the next melodic events are unexpected or uncertain (see the reddish hue in H_m and IC_m for time-steps $> \tau$ in the right plot), and to speed up if the previous melodic events were unexpected or uncertain (the bluish hue in H_m and IC_m for time-steps $< \tau$ in the right plot), which is consistent with the findings reported in [4]. On the other hand, while a passage consisting of uncertain harmonic events contributes to an overall slower tempo (the reddish hue in row H_c in the left plot), there is a tendency to speed up if the current harmonic event is unexpected or uncertain (the bluish hue in H_c and IC_c at τ in the right plot).

5 Conclusions

In this paper we presented a model for predicting expressive tempo and dynamics using a combination of expectancy and score features. Our results support the view that expectancy features, as reflecting what a listener is expecting to hear, can be used to predict the way pianists perform a piece. The sensitivity analysis also found some evidence relating to well-known rules/guidelines for performance [3, 4]. Future work may include the use of expectancy features in combination with larger sets of score descriptors (such as those in [5, 1]), and derive expectancy features from deep probabilistic models trained directly on (polyphonic) piano-roll representations.

Acknowledgements This work has been funded by the European Research Council (ERC) under the EU's Horizon 2020 Framework Programme (ERC Grant Agreement No. 670035, project CON ESPRESSIONE).

References

1. Cancino-Chacón, C.E., Gadermaier, T., Widmer, G., Grachten, M.: An evaluation of linear and non-linear models of expressive dynamics in classical piano and symphonic music. *Machine Learning* 106(6), 887–909 (2017)
2. Flossmann, S., Grachten, M., Widmer, G.: Expressive Performance with Bayesian Networks and Linear Basis Models. In: *Rencon Workshop Musical Performance Rendering competition for Computer Systems*. pp. 1–2 (Mar 2011)
3. Friberg, A., Bresin, R., Sundberg, J.: Overview of the KTH rule system for musical performance. *Advances in Cognitive Psychology* 2(2-3), 145–161 (2006)
4. Gingras, B., Pearce, M.T., Goodchild, M., Dean, R.T., Wiggins, G., McAdams, S.: Linking melodic expectation to expressive performance timing and perceived musical tension. *Journal of Experimental Psychology: Human Perception and Performance* 42(4), 594–609 (2016)
5. Giraldo, S.I., Ramírez, R.: A Machine Learning Approach to Discover Rules for Expressive Performance Actions in Jazz Guitar Music. *Frontiers in Psychology* 7, 194–13 (Dec 2016)
6. Grachten, M., Cancino-Chacón, C.E.: Temporal dependencies in the expressive timing of classical piano performances. In: Lessafre, M., Maes, P.J., Leman, M. (eds.) *The Routledge Companion to Embodied Music Interaction*, pp. 360–369. New York, NY (2017)
7. Grachten, M., Widmer, G.: Linear Basis Models for Prediction and Analysis of Musical Expression. *Journal of New Music Research* 41(4), 311–322 (Dec 2012)
8. Honing, H.: Computational Modeling of Music Cognition: A Case Study on Model Selection. *Music Perception* 23(5), 365–376 (Jun 2006)
9. Pearce, M., Müllensiefen, D., Wiggins, G.A.: A Comparison of Statistical and Rule-Based Models of Melodic Segmentation. In: *Proceedings of the Ninth International Conference on Music Information Retrieval*. Philadelphia, PA, USA (2008)
10. Pearce, M.T.: *The Construction and Evaluation of Statistical Models of Melodic Structure in Music Perception and Composition*. Ph.D. thesis, City University, London (2005)
11. Sears, D.R.W.: *The Classical Cadence as a Closing Schema: Learning, Memory, & Perception*. Ph.D. thesis, McGill University, Montreal, Canada (Sep 2016)
12. Widmer, G., Goebel, W.: Computational Models of Expressive Music Performance: The State of the Art. *Journal of New Music Research* 33(3), 203–216 (Sep 2004)