

# CONTROLLING SURPRISAL IN MUSIC GENERATION VIA INFORMATION CONTENT CURVE MATCHING

Mathias Rose Bjare<sup>1</sup>

Stefan Lattner<sup>2</sup>

Gerhard Widmer<sup>1,3</sup>

<sup>1</sup> Institute of Computational Perception, Johannes Kepler University Linz, Austria

<sup>2</sup> Sony Computer Science Laboratories (CSL), Paris, France

<sup>3</sup> LIT AI Lab, Linz Institute of Technology, Austria

mathias.bjare@jku.at

## ABSTRACT

In recent years, the quality and public interest in music generation systems have grown, encouraging research into various ways to control these systems. We propose a novel method for controlling surprisal in music generation using sequence models. To achieve this goal, we define a metric called Instantaneous Information Content (IIC). The IIC serves as a proxy function for the perceived musical surprisal (as estimated from a probabilistic model) and can be calculated at any point within a music piece. This enables the comparison of surprisal across different musical content even if the musical events occur in irregular time intervals. We use beam search to generate musical material whose IIC curve closely approximates a given target IIC. We experimentally show that the IIC correlates with harmonic and rhythmic complexity and note density. The correlation decreases with the length of the musical context used for estimating the IIC. Finally, we conduct a qualitative user study to test if human listeners can identify the IIC curves that have been used as targets when generating the respective musical material. We provide code for creating IIC interpolations and IIC visualizations on <https://github.com/muthissar/iic>.

## 1. INTRODUCTION

In music generation, controlling the generation process with user inputs is essential for creating flexible systems that support a creative human/machine co-creation process [1]. Typically, controls are based on low-level features with a direct musical interpretation, for instance, the pitch of a generative synthesizer [2], or meter, harmony, and instrumentation for symbolic generation [3]. A high-level musical feature that has received little attention in generative composition systems is musical surprisal — how surprising a musical event is to a listener, given the past musical context. The surprisal tends to be high when the

music is complex, when a pitch deviates from the prevailing tonality, or when there is a variation in rhythm [4, 5]. As such, musical surprisal shares similarities with musical complexity, however, it is importantly also affected by *learning*: Repeating complex musical content can lead to decreased surprisal on the repetitions as a result of learning [6]. In contrast, the musical content and the complexity remain unchanged across repetitions.

Studies suggest that the amount of musical surprisal needs to be balanced for music to be deemed preferable [7, 8], which is typically achieved by balancing regularity and novelty [9]. Being able to control surprisal in generated music might help users create compositions that balance regularity and novelty and thus suit listeners' preferences. In addition, if this can be controlled, rather low surprisal could be used indirectly to induce repetitions in machine-generated music and high surprisal to produce novel parts, possibly with high perceived complexity.

In [10], it was proposed to quantify the surprisal of a musical event by its Information Content (IC) conditioned on past musical events. For that, a sequence of musical events is modeled as a stochastic process, where the conditional distribution and, hence, the conditional IC can be estimated. As such, a surprising event is an event that is unlikely to occur under the estimated distribution given the past musical context. In the works of [11], the authors find correlations between the IC of a variable-order Markov model (called IDyOM) [12] and perceived surprise in a controlled pitch anticipation experiment. A correlation between high IC and tonal and rhythmic complexity was shown in [4, 5].

This indicates that the IC of trained sequence models can be used as a proxy for human perception of musical surprisal and that its measurement can identify musical complexity and regularities. This paper proposes a novel framework for generating music with user control over the IC. Specifically, we define an *Instantaneous Information Content* (IIC) measure, which can be calculated at any time point based on the IC of musical events in the recent past and approximates a causal information density. We use the IIC as a fitness score to direct a beam search toward generating samples following a given IIC target curve. Our sampling strategy can be used with any pretrained autoregressive generative music model. We demonstrate our approach in symbolic classical music generation using a



pretrained PIA model [13] and show quantitatively that our approach can generate samples that follow IIC curves extracted from real data. We conducted a qualitative study to test if humans can identify simple IIC curves used for generation. Finally, we analyze relationships between IIC and harmonic, rhythmic, and note density complexity.

## 2. METHODS

In the following, we describe a method for IC-controlled token sequence generation. Let  $\text{IC}^*(t)$  be a target curve with support in the time interval  $[0, T]$ , representing the desired information content over time of a generated sequence of tokens  $\mathbf{x} = x_1, x_2, \dots, x_n \in \mathcal{X}$ , with a duration of  $T$  seconds. ‘Tokens’ are not necessarily individual notes or note onsets but can be any token type commonly used in Transformer-based music generation systems (e.g., [13–15]). Also, note that we operate on the physical time dimension, not symbolic (score) time measured, e.g., in beats or number of tokens.

Furthermore, let  $q$  be a generative sequence model and  $p$  an autoregressive *critic model*, used for estimating the the  $i$ ’th token’s conditional token information content

$$\text{IC}(x_i|x_{<i}) = -\log p(x_i|x_{<i}), \quad (1)$$

where  $x_{<i} = x_1, x_2, \dots, x_{i-1}$ . In our context,  $p$  will be a Transformer model. The proposed method creates new samples using  $q$  with an information content that matches the target curve as measured by  $p$ . Our method works as follows: Firstly, we define the *Instantaneous Information Content* (IIC) – a mapping from a (temporally irregular) token sequence and its information content values to a function representing the musical surprisal in the continuous time domain. Secondly, we define an *IC deviation* – a metric for comparing the similarity between a sequence’s IIC curve and the target curve. Finally, we devise a method for generating token sequences with  $q$  that minimize the *IC deviation*.

### 2.1 Instantaneous Information Content

#### 2.1.1 Temporal Localization of IC Estimates

To align the information content of musical events, measured on sequence tokens, with the time-domain target IC ( $\text{IC}^*$ ), we face a challenge: IC is calculated on sequence elements, while  $\text{IC}^*$  pertains to the time domain. Our solution involves assigning each token a temporal position using a mapping function  $f$ , effectively “temporally localizing” or aligning tokens within the musical timeline. Note that  $f$  can be constructed by analyzing the specific detokenization method associated with  $\mathbf{x}$ ’s tokenization that involves turning a sequence of tokens into a time-based music representation like MIDI<sup>1</sup>. In section 3.1, we present an example of such  $f$  using the tokenization of [13].

Temporal Localization allows us to map IC tokenizations to their respective time points in the music. This is crucial, especially for analyzing tokenizations of symbolic

music commonly used with Transformers [13–15], where the decoded musical events do not uniformly align in time. Through this approach, IC measured on tokens can be directly compared with the time-domain  $\text{IC}^*$ , facilitating a coherent analysis across different domains of musical representation.

#### 2.1.2 Interpolation

Let  $f : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$  be a localization function, mapping the  $i$ ’th token of sequence  $\mathbf{x} \in \mathcal{X}$  to the time domain. The IIC at time  $t$  in a piece (represented by token sequence  $\mathbf{x}$ ), is a real number computed by a time interpolation of  $\mathbf{x}$ ’s token ICs:

$$\text{IIC}(t, \mathbf{x}) = \sum_{f(i, \mathbf{x}) < t} \lambda(t - f(i, \mathbf{x}), i) \cdot \text{IC}(x_i|x_{<i}). \quad (2)$$

$\lambda(t, i)$  defines a weighting of the information of the  $i$ ’th token and the constraint  $f(i, \mathbf{x}) < t$  ensures causality. As a result, the IIC at any time step  $t$  is a weighted sum of IC values of past events, using a weighting kernel  $\lambda$ .

The choice of the critic model  $p$  in combination with the weight function  $\lambda$  defines different perceptual models of the instantaneous information content. We propose to choose  $\lambda$  so that the recent past is weighted higher than the remote past. More specifically, we define  $\lambda$  as a window function centered around  $t$  and equal to zero at time steps greater than  $t$ . In this initial work, we chose a Hann window for the following reasons: As it is (half) bell-shaped, it is insensitive to inaccuracies in the temporal localization of recent events. It is smooth at the boundaries, preventing sudden drops as events “leave” the window.

Using the IIC, we quantify the *segment surprisal* of segment  $[t_1, t_2]$  by the  $L^1$  norm of the IIC with support restricted to  $[t_1, t_2]$  by calculating:

$$\|\text{IIC}|_{t_1}^{t_2}\|_1 = \int_{t_1}^{t_2} |\text{IIC}(t, \mathbf{x})| dt. \quad (3)$$

In section 3.5, we compare segment surprisal with segment-based complexity metrics.

### 2.2 IC Deviation

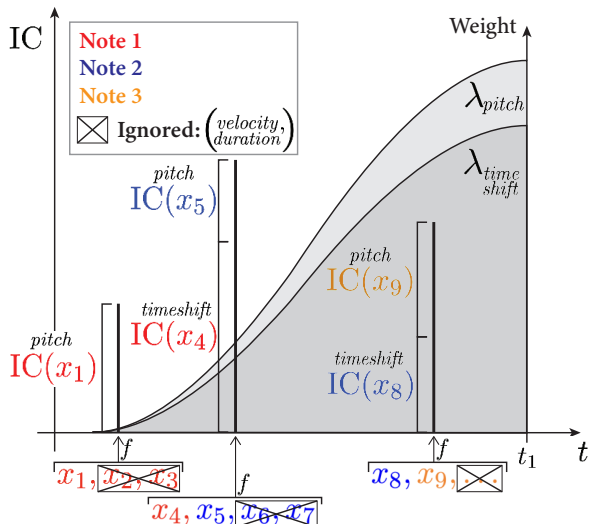
Given a sample  $\mathbf{x}$ , the *IC deviation* of  $\text{IIC}(\cdot, \mathbf{x})$  from the target  $\text{IC}^*$  is defined as the  $L^1$  norm of their function difference:

$$\|\text{IC}^* - \text{IIC}\|_1 = \int_0^T |\text{IC}^*(t) - \text{IIC}(t, \mathbf{x})| dt. \quad (4)$$

Which is equal to zero if  $\text{IC}^* = \text{IIC}(\cdot, \mathbf{x})$  almost everywhere, implying that minimizing eq. (4), aligns the target curve  $\text{IC}^*$  with the IIC curve. In practice, we compute eq. (4) by the Riemann sum:

$$\|\text{IC}^* - \text{IIC}\|_1 \approx \sum_{i=1}^m |\text{IC}^*(t_i) - \text{IIC}(t_i, \mathbf{x})| \Delta t, \quad (5)$$

<sup>1</sup><https://midi.org/midi-1-0-detailed-specification> where  $m\Delta t = T$ .



**Figure 1.** The temporal localization function  $f$  and the weight function  $\lambda$ , involved in computing the IIC of  $x_1, x_2, \dots$ , a sequence of three notes, at time  $t_1$ .

### 2.3 Information Content Conditioned Sampling

We can now rank sequences of different lengths according to their proximity to the target  $IC^*$  using eq. (5). We use this to guide a beam search to follow the target curve. The beam search is done in iterations. At each iteration, we generate  $k$  continuations of the best-performing sample from the last iteration (initially the empty sequence) in parallel. We stop expanding the continuation when the duration of the newly generated content exceeds a predefined step size  $t'$ . We then evaluate eq. (5) and keep only the continuation with the lowest IC deviation for the next iteration<sup>2</sup>. We stop when the generation’s duration is  $T$ .

## 3. EXPERIMENTS

### 3.1 Model and Data

All experiments are performed with a PIA Transformer model [13], a symbolic music generation system pretrained on expressive classical piano performances. The model was trained on data consisting of 1,184 MIDI files of expressive music recorded with high precision on a Yamaha Disklavier [16], as well as a larger dataset of 10,855 MIDI files containing automatically transcribed piano performances [17]. For evaluation, we use the dataset of [18], consisting of performances of 36 Mozart piano sonata movements. The midi files are tokenized using a *structured MIDI encoding* [13], where midi notes, sorted by their onset times, are serialized successively using four tokens *Pitch*, *Velocity*, *Duration*, *Timeshift* in that order. Therefore, every fourth token represents the same token type. *Pitch* is an integer describing the 88-note pitch values on the piano. *Velocity* is an integer describing the 128 possible midi velocity values. *Duration* is an integer representing quantized note duration in seconds:

<sup>2</sup> Practically, in beam search iteration  $i$ , we evaluate the integral of eq. (4) from 0 to  $it'$ .

$\{0.02, 0.04, \dots, 1.0, 1.1, \dots, 5.0, 6.0, \dots, 19.0\}$ . *Timeshift* is an integer encoding the inter-onset intervals (IOI, i.e., the time durations between subsequent note onsets). *Timeshift* is quantized similarly to the duration token, with the addition of an extra symbol representing a time shift of zero, allowing the model to understand that notes less than 0.02 seconds apart are to be played concurrently. In contrast to the PIA model described in [13], which does non-causal inpainting, we use a causal Transformer based on the Perceiver IO architecture [19] and do continuation generation<sup>3</sup>. We make these modifications such that the IC calculations ignore future observations. We use the same pretrained model both as the *generator* model  $q$  and the *critic* model  $p$  and leave the exploration of other critic models for future work.

### 3.2 IIC

The elements involved in computing the IIC are given in fig. 1. For IIC calculations, we choose to consider only the surprisal of *Pitch* and *Timeshift* tokens, such that the token’s IC represents the surprisal of pitches and IOI. We ignore *Velocity* and *Duration* tokens because they contribute less to the perception of surprise, being mostly related to the performance dimensions dynamics and articulation. This is achieved in the IIC calculation by setting  $\lambda(t, i) = 0$  for  $i = 2, 6, 10, \dots$  and  $i = 3, 7, 11, \dots$  in eq. (2). We choose  $f$  such that the pitch token contributes to the surprisal function at its note onset time<sup>4</sup>, and the timeshift token contributes to its surprisal at the onset of the following note (as an IOI is perceived at the onset of the next note). The remaining weights are then defined by the scaled half Hann window

$$\lambda(t, i) = \begin{cases} c_i \frac{1}{L} \cos^2\left(\frac{\pi t}{L}\right) & \text{for } 0 < t < \frac{L}{2}, \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where  $c_i$  is a weight that takes on two different values for the pitch and timeshift tokens, respectively.  $c_i$  is used to weigh the IC of pitches and timeshifts, respectively. For both token types to have equal importance, we estimate a normalization constant empirically by calculating a mean IC over all tokens of the evaluation dataset. The window length is chosen to be  $L = 4$  so that the weight is zero after 2 seconds.

### 3.3 Beam Search Parametrization Study

Using the beam search strategy described in section 2.3, we run initial experiments to determine the effect of parameters associated with the beam search on the similarity between generated samples and target curves  $IC^*$  extracted from real music. Specifically, we randomly select 400 snippets of the MIDI files (10 seconds long) and create the IIC curve associated with those snippets. Then, we generate four new samples using our beam search and evaluate the IC deviation between the IIC curve induced by the

<sup>3</sup> The model generates sequences using an initial context of real music.

<sup>4</sup> The note onset times are found by accumulating the time values associated with previous *Timeshift* tokens.

real data and the generated data. We discretize the integral in the IC deviation (see eq. (5)) with  $\Delta t = 0.1s$ .

To investigate the effect of the step size  $t'$ , we fix the number of continuations generated in parallel to  $k = 16$  to reduce computation. To investigate the importance of the number of parallel generated samples  $k$ , we use a fixed step size of  $t' = 0.3s$ .

We find that in cases where the generation model  $q$  and the critic model  $p$  are the same ( $p = q$ ), it is challenging to sample a single continuation  $x_i, x_{i+1}, \dots, x_{i+m}$  (using  $q$ ) that has a high segment surprisal  $\|\text{IIC} \left| \frac{f(x_{i+m})}{f(x_i)} \right\|_1$  (measured by  $p$ ), precisely because the probability of sampling such a continuation is low.

To sample low-probability tokens more efficiently, we propose a heuristic that alters the entropy<sup>5</sup> of the generating distribution  $H(q)$  using a temperature parameter dynamically set using the IIC. Specifically, in iteration  $i-1$  of the beam search, we measure  $\text{IC}^*(it')$ , the target IC at the time where the generation of the continuations halts next time, and calculate a target entropy:

$$H_{target} = \min \left( \frac{\text{IC}^*(it')}{C_H}, H_{max} \right), \quad (7)$$

where  $C_H$  is a constant parameter to be estimated and  $H_{max}$  is the entropy of the uniform distribution. We then fix  $q$ 's entropy to the target entropy  $H_{target}$  by searching for a temperature  $r$  such that  $H_{target} = H(q) = H(\text{softmax}(l/r))$  with binary search, where  $l$  are the logits of the neural network. Note that temperature is only used for the generator  $q$  and not for the critic model  $p$ .

### 3.4 Qualitative Evaluation

We conducted an online user study to investigate if the IIC curves computed on generated and real music correspond to users' experience of being musically surprised.

Firstly, we present the participant with a musical section generated by our method using one of five target curves. The participant is then tasked to select the IIC curve that best describes their perceived surprise when listening to the section. Secondly, we present the user with a segment of real music and IIC curves extracted from real music, one of which corresponds to the music segment. The user's task is to identify the corresponding curve.

The experiment is conducted on a website that, after an initial experiment description, asks the user for their years of musical training (more or less than five years). Then, it shows an example of a generated piano music section and the surprisal curve used as a target for the generation (together with a textual explanation).

The participant is then presented with five pages, like the one in fig. 2. Each presents a musical section generated using one of five simple target curves. The participant is asked to identify which of the five curves they think has been used to generate the section. The final page contains a 10-second segment of real piano music from the evaluation set and two IIC curves, one corresponding to the piano

music and the other to a randomly selected 10-second segment from the evaluation dataset.

The samples for the first five pages are generated as follows: As contexts for the model, we select the first 13 measures of Mozart K.331, 1st mvt. and the first 16 measures of K.332, 2nd mvt. from the evaluation dataset and generate 200 samples for every combination of the two musical contexts and the five IIC curves shown on the page, with  $C_H = 50$ ,  $t' = 0.3s$  and  $k = 128$  (i.e., the optimal beam search parameters, as shown in table 1). For each combination, we then select the 25 samples with the lowest IC deviation for the user study. For the final page of the user study related to real performances, we select 300 different 10-second segments from the evaluation dataset and compute the IIC curves. The results of the user study will be presented and discussed in section 4.2

### 3.5 Analysis of IIC

As discussed in the introduction, IC and surprisal might be related to aspects of musical complexity, but learning effects may lead to a decrease in surprisal in passages with repeated musical content. To investigate these relationships, we designed an experiment to determine if the IIC correlates with harmonic complexity, as quantified by *tonal tension (cloud diameter)* [20], where the IIC is calculated using progressively larger segments of musical context. Tonal tension is calculated for a segment of music by considering its most dissonant pitch class interval, where an interval dissonance is measured as the distance between the interval pitches embedded in a specific Euclidian space where the position is based on the circle of fifths [21].

We extracted one-second segments centered on the onsets of notes in the evaluation dataset. For  $i = 1, \dots, 1000$ , we then compute the Pearson correlation coefficient between the tonal tension and the segment surprisal (see eq. (3)) of the first  $i$  segments within every performance.

In addition, we investigate complexity in terms of note density, i.e., the number of notes per segment. To do so, we use the same setup as for tonal tension but count the number of notes within one-second segments.

Finally, we investigate rhythmical complexity using the *IOI histogram entropy* of measures [22]. We choose this measure over other structural rhythmical complexity measures [23–28] since it does not assume the rhythm to be cyclic. We follow the same procedure as mentioned above, but instead of selecting fixed-sized segments centered around note-onsets, we select segments of one measure based on the measure annotations [18]. More specifically, we match the notes of the performance with its score notes and extract for each measure: 1) the normalized entropy of the score notes IOI histogram and 2) the segment IIC of the measure normalized with the length of the measure. The segment boundaries are estimated by the mean onset time of the first and last note in subsequent measures.

<sup>5</sup> Entropy is the expectation of IC.



Figure 2. Example page of the user study with a generated musical section and five target curves to choose from.

$t'$	0.1s	0.2s	0.3s	0.4s	0.5s	0.6s	0.7s	0.8s	1.0s	2.0s
IC dev.	3.63	2.71	<b>2.61</b>	2.72	2.69	2.93	3.03	3.11	3.31	3.90
$k$	1	2	4	16	32	64	96	128		
IC dev.	8.41	5.90	4.36	2.61	2.14	1.89	1.76	<b>1.69</b>		
$C_H$	10	20	30	40	50	60	70	80	120	No
IC dev.	8.33	4.11	2.67	2.21	<b>2.15</b>	<b>2.15</b>	2.25	2.45	3.12	2.61

Table 1. IC deviation between target curves IC\* extracted from real music, and IIC curves from continuations generated with different beam search parameters.

## 4. RESULTS

### 4.1 Beam Search Parameter Study

In table 1, we report the mean IC deviation of samples generated with different beam search step sizes  $t'$ , numbers of continuations generated in parallel  $k$  and  $C_H$ , constants used for setting the softmax temperature dynamically. Bigger step sizes create longer continuations with high IC deviation variance, resulting in worse performance. The lowest values ( $t' = 0.1s, 0.2s$ ) also worsen IC deviation, likely because sampled notes exceed the timestep, causing inaccuracies in the next beam search iteration. For the number of continuations generated in parallel  $k$ , we find that the IC deviation always decreases with higher  $k$ . This is not surprising as the model has more candidate continuations to choose from. The decrease flattens out as seen by the small IC deviation differences when  $k \geq 64$ . For the dynamic temperature, we find that  $C_H = 50, 60$  reduces the IC deviation compared to using no temperature scaling (marked with "No" in table 1).

### 4.2 Qualitative Results


The user study results reported as a binary classification of finding the correct curve, among the curves described

in section 3.4, are presented in Table 2. 29 users participated, 23 participants had more than 5 years of musical training, and 6 participants had less than 5 years of experience. 152 generated samples and 21 samples of real music were classified in total. Due to the imbalance in the number of untrained and trained participants and since we found little difference in the classification performance between the groups, we combined their results in the table.

The overall  $F_1$ -score was reported as 0.52 for generated data and 0.71 for real data, which is reasonably above the proportions 0.2 and 0.5, being the  $F_1$  scores of random classifiers, with 5 and 2 classes respectively. The results for the individual curves show difficulty differences in classifying the different curve types, with RAMP\_DOWN having the lowest  $F_1$ -score of 0.41 and STEP\_UP having the highest  $F_1$ -score of 0.71. We therefore investigate the confusion of curves in Figure 3. We find that the confusion of CONSTANT is evenly distributed on all curves, except for STEP\_UP, which is reasonable since CONSTANT does not share any characteristics with the other curves. We furthermore find that generations that start with the same IIC value, either high or low, are confused. This is seen by the confusion of RAMP\_DOWN with STEP\_DOWN and the confusion of STEP\_UP and RAMP\_UP.

	IIC Curves					Gen. all curves	Real
	CONSTANT	RAMP_DOWN	STEP_UP	RAMP_UP	STEP_DOWN		
$F_1$	0.53	0.41	0.71	0.48	0.49	0.52	0.71
#True	36	34	29	29	24	152	21
#Pred	36	33	31	30	22	152	21

**Table 2.** Results from the user study reported as  $F_1$ -score of identifying the: IIC curve used for generation, the IIC curve of real music.



**Figure 3.** The confusion matrix for users identifying the IC\* curves used to generate the music examples.


### 4.3 Analysis of IIC

The correlations between IIC and the tonal tension  $tt$ , note density  $d$ , and the IOI histogram entropy  $he$  were calculated on the first  $n$  segments of the 36 evaluation data performances as described in section 3.5 and reported in fig. 4. We report the results for IIC calculated using *Pitch* only, *Timeshift* only, or both token types. For  $tt$ ,  $d$ , and  $he_{Timeshift}$ , the correlations reported were found significant using a significance level of 0.05, whereas for  $he_{Pitch}$  and  $he_{Both}$ , the correlations are not significant.

The results show a moderate to high correlation of IIC with all metrics at the beginning of the performances (when  $n$  is small). However, these correlations decrease in later parts of the performances (when  $n$  is high), likely due to “learning” (simulated by longer context) over time.

The highest correlations are found for note density  $d$ . This may be explained by the definition of IIC (see eq. (2)) as a weighted sum of token ICs since more tokens per segment simply lead to higher sums.

Considering the different token type combinations, we find that  $tt$  is most correlated with IIC calculated using only *Pitch* tokens and  $he$  using only *Timeshift* tokens. This is reasonable, considering that very dissonant segments and very complex rhythms tend to be associated with *Pitch* and *Timeshift* tokens, respectively, which are infrequent in the training dataset, resulting in a high token IC. Interestingly,  $tt$  is also correlated with IIC calculated using only *Timeshift* tokens (encoding IOIs), which



**Figure 4.** Correlation between IIC and tonal tension  $tt$ , note density  $d$ , and IOI histogram entropy ( $he$ ).

might stem from the critic model facing greater uncertainty in predicting any token type when confronted with a highly harmonic complex context that is infrequent in the dataset.

The curves of *Both* and *Pitch* follow each other closely for total tension and note density, indicating that using both *Timeshift* and *Pitch* tokens does not significantly reduce the complexity correlations. For rhythmical complexity, using *Both* tokens instead of *Timeshift* tokens alone decreases the correlation more.

## 5. CONCLUSION

In this study, we introduced a novel framework for controlling musical surprisal through Instantaneous Information Content (IIC), which maps token-based surprisal to a continuous time-domain function. Using a beam search algorithm, we demonstrated that our approach can generate music that closely follows predefined IIC curves, effectively aligning generated and target surprisal curves.

Our user study confirmed that participants could reasonably identify target IIC curves from generated music, indicating that our method captures perceptible aspects of musical surprise. Furthermore, our analysis showed that IIC correlates with measures of musical complexity such as tonal tension and note density.

Future work will explore alternative critic models, like personalized models, trained on music that is familiar to the user or models with smaller context windows to more directly control local musical complexity.

## 6. ACKNOWLEDGMENTS

The work leading to these results was conducted in a collaboration between JKU and Sony Computer Science Laboratories Paris under a research agreement. GW's work is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement 101019375 ("Whither Music?").

## 7. REFERENCES

- [1] C. A. Huang, H. V. Koops, E. Newton-Rex, M. Dinulescu, and C. J. Cai, "Human-ai co-creation in song-writing," in *ISMIR*, 2020, pp. 708–716.
- [2] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, "Neural audio synthesis of musical notes with wavenet autoencoders," in *International Conference on Machine Learning*. PMLR, 2017, pp. 1068–1077.
- [3] D. von Rütte, L. Biggio, Y. Kilcher, and T. Hofmann, "FIGARO: Controllable music generation using learned and expert features," in *ICLR*. OpenReview.net, 2023.
- [4] S. A. Sauv e and M. T. Pearce, "Information-theoretic modeling of perceived musical complexity," *Music Perception: An Interdisciplinary Journal*, vol. 37, no. 2, pp. 165–178, 2019.
- [5] M. R. Bjare, S. Lattner, and G. Widmer, "Exploring sampling techniques for generating melodies with a transformer language model," in *ISMIR*, 2023, pp. 810–816.
- [6] —, "Differentiable short-term models for efficient online learning and prediction in monophonic music," *Trans. Int. Soc. Music. Inf. Retr.*, vol. 5, no. 1, p. 190, 2022.
- [7] M. T. Pearce and G. A. Wiggins, "Auditory expectation: The information dynamics of music perception and cognition," *Top. Cogn. Sci.*, vol. 4, no. 4, pp. 625–652, 2012.
- [8] T. E. Matthews, M. A. Witek, O. A. Heggli, V. B. Penhune, and P. Vuust, "The sensation of groove is affected by the interaction of rhythmic and harmonic complexity," *PLoS One*, vol. 14, no. 1, p. e0204539, 2019.
- [9] I. Zioga, P. M. C. Harrison, M. T. Pearce, J. Bhattacharya, and C. D. B. Luft, "From learning to creativity: Identifying the behavioural and neural correlates of learning to predict human judgements of musical creativity," *NeuroImage*, vol. 206, 2020.
- [10] L. B. Meyer, "Meaning in music and information theory," *The Journal of Aesthetics and Art Criticism*, vol. 15, no. 4, pp. 412–424, 1957.
- [11] M. T. Pearce, M. H. Ruiz, S. Kapasi, G. A. Wiggins, and J. Bhattacharya, "Unsupervised statistical learning underpins computational, behavioural, and neural manifestations of musical expectation," *NeuroImage*, vol. 50, no. 1, pp. 302–313, 2010.
- [12] D. Conklin and I. H. Witten, "Multiple viewpoint systems for music prediction," *Journal of New Music Research*, vol. 24, no. 1, pp. 51–73, 1995.
- [13] G. Hadjeres and L. Crestel, "The piano inpainting application," *CoRR*, vol. abs/2107.05944, 2021.
- [14] C. A. Huang, A. Vaswani, J. Uszkoreit, I. Simon, C. Hawthorne, N. Shazeer, A. M. Dai, M. D. Hoffman, M. Dinulescu, and D. Eck, "Music transformer: Generating music with long-term structure," in *ICLR (Poster)*. OpenReview.net, 2019.
- [15] Y. Huang and Y. Yang, "Pop music transformer: Beat-based modeling and generation of expressive pop piano compositions," in *ACM Multimedia*. ACM, 2020, pp. 1180–1188.
- [16] C. Hawthorne, A. Stasyuk, A. Roberts, I. Simon, C. A. Huang, S. Dieleman, E. Elsen, J. H. Engel, and D. Eck, "Enabling factorized piano music modeling and generation with the MAESTRO dataset," in *ICLR*. OpenReview.net, 2019.
- [17] Q. Kong, B. Li, J. Chen, and Y. Wang, "Giantmidi-piano: A large-scale MIDI dataset for classical piano music," *Trans. Int. Soc. Music. Inf. Retr.*, vol. 5, no. 1, pp. 87–98, 2022. [Online]. Available: <https://doi.org/10.5334/tismir.80>
- [18] P. Hu and G. Widmer, "The batik-plays-mozart corpus: Linking performance to score to musicological annotations," in *ISMIR*, 2023, pp. 297–303.
- [19] A. Jaegle, S. Borgeaud, J. Alayrac, C. Doersch, C. Ionescu, D. Ding, S. Koppula, D. Zoran, A. Brock, E. Shelhamer, O. J. H enaff, M. M. Botvinick, A. Zisserman, O. Vinyals, and J. Carreira, "Perceiver IO: A general architecture for structured inputs & outputs," in *ICLR*. OpenReview.net, 2022.
- [20] D. Herremans and E. Chew, "Tension ribbons: Quantifying and visualising tonal tension," in *Proceedings of the International Conference on Technologies for Music Notation and Representation - TENOR2016*, R. Hoadley, C. Nash, and D. Fober, Eds. Cambridge, UK: Anglia Ruskin University, 2016, pp. 8–18.
- [21] E. Chew, "The spiral array: An algorithm for determining key boundaries," in *ICMAI*, ser. Lecture Notes in Computer Science, vol. 2445. Springer, 2002, pp. 18–31.
- [22] A. A. Moles, *Information theory and esthetic perception*. The University of Illinois Press, Urbana and London, 1966.

- [23] H. C. Longuet-Higgins and C. S. Lee, “The rhythmic interpretation of monophonic music,” *Music Perception*, vol. 1, no. 4, pp. 424–441, 1984.
- [24] L. M. Smith, H. Honing *et al.*, “Evaluating and extending computational models of rhythmic syncopation in music,” in *ICMC*, 2006.
- [25] M. Keith, *From polychords to poly: adventures in musical combinatorics*. Vinculum Press, 1991.
- [26] J. Pressing, “Cognitive complexity and the structure of musical patterns,” in *Proceedings of the 4th Conference of the Australasian Cognitive Science Society*, vol. 4, 1999, pp. 1–8.
- [27] F. Gómez, A. Melvin, D. Rappaport, and G. T. Toussaint, “Mathematical measures of syncopation,” in *Renaissance banff: Mathematics, music, art, culture*, 2005, pp. 73–84.
- [28] A. I. Mezza, M. Zaroni, and A. Sarti, “A latent rhythm complexity model for attribute-controlled drum pattern generation,” *EURASIP J. Audio Speech Music. Process.*, vol. 2023, no. 1, p. 11, 2023.