
Deep Generative Models of Music Expectation

Ninon Lizé Masclef
contact@ninonlizemasclef.com

T. Anderson Keller
t.anderson.keller@gmail.com

Abstract

A prominent theory of affective response to music revolves around the concepts of surprisal and expectation. In prior work, this idea has been operationalized in the form of probabilistic models of music which allow for precise computation of song (or note-by-note) probabilities, conditioned on a ‘training set’ of prior musical or cultural experiences. To date, however, these models have been limited to compute exact probabilities through hand-crafted features or restricted to linear models which are likely not sufficient to represent the complex conditional distributions present in music. In this work, we propose to use modern deep probabilistic generative models in the form of a Diffusion Model to compute an approximate likelihood of a musical input sequence. Unlike prior work, such a generative model parameterized by deep neural networks is able to learn complex non-linear features directly from a training set itself. In doing so, we expect to find that such models are able to more accurately represent the ‘surprisal’ of music for human listeners. From the literature, it is known that there is an inverted U-shaped relationship between surprisal and the amount human subjects ‘like’ a given song. In this work we show that pre-trained diffusion models indeed yield musical surprisal values which exhibit a negative quadratic relationship with measured subject ‘liking’ ratings, and that the quality of this relationship is competitive with state of the art methods such as IDyOM. We therefore present this model as a preliminary step in developing modern deep generative models of music expectation and subjective likability.

1 Introduction

The fields of psychology and musicology have identified expectation as a crucial factor that conveys meaning in music [14], and especially affective response [12, 21, 26]. The malleability of musical experience underscores that listening to music is not merely a passive activity, but an active learning process in which expectations are formed that shape our emotional responses [22]. Indeed, prior work has found that the sweet spot of expectation which maximizes information learning and reward system-related responses is that of intermediate complexity [18]. Wilhelm Wundt [30] proposed an inverted U-curve to describe the relationship between stimulus intensity and pleasant feeling. This idea was operationalized decades later and applied specifically to aesthetic pleasure by Berlyne [4]. Since then, the *Wundt effect* has dominated psychological research on music preference for more than two decades [7]. Specifically, Berlyne [3] found that arousal is a predominant factor in aesthetic preference, with three types of variables determining the level of arousal of a stimulus input. One in particular, collative variables, encompassing novelty, complexity, and uncertainty, has been shown to contribute most to musical liking [2]. Among these collative variables, predictability has been empirically shown to contribute to music preference [29], although this finding is not unanimous [17]. Furthermore, there is a growing body of research on serendipity as a new evaluation metric for music recommendation systems [31, 23, 5], linking surprisal and pleasure of music.

1.1 Models of musical expectancy

Meyer (1956) first theorized musical anticipation by applying principles of Gestalt psychology to the auditory domain to form an "emotional syntax of music" [14]. Influenced by Meyer's work, Narmour (1990) [15] developed the Implication-Realization (I-R) model of melodic expectation, which emphasized cognitive evaluation of the music rather than musical analysis of the stimuli. Instead, Narmour explained that music listening is the product of two expectation processes, one top-down and the other bottom-up [16]. A whole tradition of predictive models of music listening was developed later in the lineage of Meyer's work. Pearce created the Information Dynamics of Music (IDyOM) model of melodic expectations [20], but replaced the rules with statistical distribution learning. IDyOM simultaneously predicts the pitch and onset time of the next note using probabilistic prediction, and calculates the final probability of the note using the joint probability of these two predictions [21]. IDyOM was effective at predicting liking from surprise, reproducing the *Wundt effect* along a U-shaped curve [6]. Unlike Narmour's model, IDyOM is not built from symbolic rules of musical knowledge. However, it still relies on hand-crafted features such as pitch duration. Contrary to recent diffusion models, which have billions of parameters learned from large datasets, IDyOM has only a few trainable parameters. In contrast to Narmour, Pearce argued that only top-down processing occurs, rather than both top-down and bottom-up [21]. Furthermore, it is conditioned on the input of stimuli encoded in MIDI format. Therefore, musical features based on timbral or dynamic changes are not considered by the model. Instead, IDyOM focuses on syntactic features, following Meyer's argument that those features present in symbolic music are primary to the formation of musical style [21, 14]. While symbolic music can convey emotion through expressive velocity, some affective cues are inherent to the live recording of a piece, such as *rubato* [25]. The D-REX model was designed to address this issue with a model of surprise working on raw audio from an ecologically valid perspective of studying the music listening experience [1]; however, the features used to compute the surprisal values in the D-REX model are restricted to learned linear transformations of the input, and therefore limited in their modeling complexity.

Nowadays, predictive coding theories are becoming dominant in modeling human expectations over time [19]. The basic premise of predictive coding is that the brain operates as an expectation machine [9], constantly generating and updating predictions. Namely, humans embody a Bayesian generative model that maps beliefs to observations. As a consequence, modern generative models, such as diffusion models, can be seen as instantiations of predictive coding. Predictive coding theories claim that nervous system activity represents a process of matching internally generated predictions with external stimuli [8]. This has been demonstrated at a variety of spatial and temporal scales in the brain including music processing [26], which has been crystallized in the predictive coding of music (PCM) model [27].

2 Methods

In this work, our goal is to leverage modern deep probabilistic generative models as measures of expectancy for music. To evaluate the quality of such a model of expectancy, we will measure the correlation structure between the output of our model — the estimated likelihood of the data interpreted as a 'surprisal' value — and a dataset of subjective human ratings of 'liking' for a set of short musical pieces. Given prior work relating music expectancy and preference, if our model is a valid model of human music expectancy, we would therefore expect to find an inverted U-shaped curve relating the surprisal and ratings.

To approach this goal, we begin with the baseline denoising diffusion probabilistic models of Ho et al. (2020) [10] as they form the foundation for much of today's state of the art in generative modeling. As we will detail in this section, due to the specific construction of these models, they admit an easily computable bound on the likelihood of the data which we will use as our approximate estimated 'surprisal'. Furthermore, there exist open source implementations and model weights trained on raw musical audio, allowing for direct evaluation of such models without retraining [24]. In the following section we give an overview of how such models operate, how to compute the approximate likelihood, as well as how we perform analysis to accurately compare this measure of surprisal with human listener ratings of liking.

Denoising Diffusion Probabilistic Models. In brief, diffusion models can be understood as latent variable generative models composed of both a forwards and reverse process. The forward process is

defined to start from the datapoint \mathbf{x}_0 and produce sequentially noisier versions \mathbf{x}_t which we call the latent variables. At the end of the forward process, the distribution of final latent variables \mathbf{x}_T is intended to match that of the prior, typically a standard normal distribution. The reverse process then inverts this procedure, mapping from the simple prior distribution $p_{\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{0}, \mathbf{I})$ to the complex data distribution $p_{\mathbf{x}_0}(\mathbf{x})$ by slowly removing noise step by step – ‘denoising’. In practice, the forwards process is Markov, with the joint distribution (factorizing into a product of conditionals) and the conditional distribution of each step defined as:

$$q(\{\mathbf{x}_t\}_{t=1}^T | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}), \quad q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_{t-1}, (1 - \alpha_t) \mathbf{I}) \quad (1)$$

where $\alpha_1, \dots, \alpha_T$ denotes the noise schedule. The reverse process is similarly Markov, and the conditional distribution which aims to remove the noise is often parameterized by a neural network with parameters θ :

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)) \quad (2)$$

In practice, Ho et al. (2020) [10] find that a reparameterization where the network is trained to directly predict the noise component of \mathbf{x}_t , ($\hat{\epsilon}_\theta(\mathbf{x}_t, t) \approx \epsilon$) is easier to optimize than predicting \mathbf{x}_t itself.

Evaluating the Likelihood of Diffusion Models. Training the parameters of denoising diffusion models can then be done by optimizing a variational bound on the negative log-likelihood:

$$\mathbb{E}[-\log p(\mathbf{x})] \leq \mathbb{E}_q \left[\mathcal{L}_T + \sum_{t=1}^T \mathcal{L}_{t-1} - \mathcal{L}_0 \right] \quad (3)$$

where $\mathcal{L}_0 = \log p(\mathbf{x}_0 | \mathbf{x}_1)$ is the reconstruction loss of the true data given \mathbf{x}_1 ; $\mathcal{L}_T = D_{KL}(q(\mathbf{x}_T | \mathbf{x}_0) || p_{\mathbf{x}_T}(\mathbf{x}))$ is the prior loss, measuring the distance of the final latent variable from the prior; and \mathcal{L}_t is the diffusion loss over all time steps defined below. Following this, Kingma et al. (2023) [13] show that the diffusion loss term can be expressed simply as:

$$\mathcal{L}_t(\mathbf{x}) = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\frac{1}{2} (1 - \text{SNR}(t-1)/\text{SNR}(t)) \|\epsilon - \hat{\epsilon}_\theta(\mathbf{x}_t, t)\|^2 \right] \quad (4)$$

where $\text{SNR}(t) = \frac{\alpha_t^2}{1 - \alpha_t^2}$. Further, following [11] we know that \mathcal{L}_T is close to zero when the noise schedule is defined such that $\alpha_T \approx 0$. In this work, we take this bound on the likelihood as our approximate measure of expectancy or surprisal for each music sample in our dataset.

Data. To evaluate the diffusion models, we make use of the dataset from Gold et al. [6] which contains 57 audio files, liking ratings from 44 subjects, and the associated output of the IDyOM model (mean duration weighted information content, mDW-IC) which we use as a baseline. Since we are working with pre-trained diffusion models which operate on 5-second audio clips, we simply break each clip into 5-second non-overlapping blocks and take the sum of the estimated likelihoods as the total likelihood value. Although this is non-ideal and introduces significant boundary effects, we find in practice it works sufficiently well for this study.

Analysis. To compare our surprisal value with liking ratings, we follow the procedure of Gold et al. [6], first fitting a linear mixed effects model to the data (python statsmodels’s mixedlm) and using this to adjust each subjects rating based on their estimated individual random effects. This allows to us account for the subjective properties of likability rating per user, and compare them to our model which only produces a single likelihood value per audio clip. We then compute the goodness of fit of a quadratic function mapping from surprisal values to adjusted \mathcal{L} rating using ordinary least squares.

3 Experiments

In this section we report the results of evaluating the pre-trained Audio-Diffusion model (*audio-diffusion-256*) from Smith [24] on the Gold et al. (2019) [6] audio dataset. The model is trained on a dataset of over 20,000 mel spectrograms of resolution 256x256 computed using a sampling rate of 22050 Hz, a window length of 2048, and a stride of 512.

Measuring the *Wundt Effect*. As a primary result of this study, we seek to discover if an inverted U-shape relationship exists between the surprisal values computed by a diffusion model and the human liking ratings from the dataset. In Figure 1, we plot the results of our analysis for both the baseline IDyOM model (left) and the Audio Diffusion model (right). We see that indeed the Wundt effect is present, with the diffusion model exhibiting a further rightward shift of peak likeability compared with IDyOM. From the literature [6], this shift is consistent with expertise and perhaps is a result of the diffusion model’s more diverse training set compared with the IDyOM counterpart.

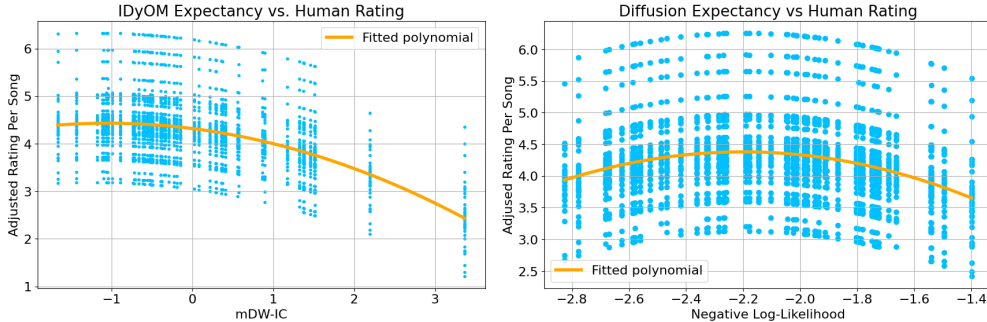


Figure 1: Plot of the observed Wundt effect for the IDyOM model (left) and our proposed diffusion model (right). Data points for the adjusted human-subject ratings of each song are plotted in blue, with the best quadratic fit model plotted in orange.

Goodness of Fit. To evaluate the quality of the above fit, in Table 1 we evaluate standard goodness of fit measures compared with the IDyOM model (such as R^2 , Log-likelihood, AIC and BIC). We see that fitted polynomial resulting from the diffusion model output achieves a higher likelihood, and better AIC/BIC values, with a lower R^2 value due to the significant inter-subject variability per song. Despite this lower R^2 , we see that the coefficient of the quadratic term of the polynomial is indeed negative and the fit is highly significant (coeff= -1.135 , 95% CI [$-1.573 -0.697$], $p < 0.0001$).

Table 1: Goodness of fit measures for the IDyOM model and the diffusion model. We see that while the IDyOM model has a higher R^2 , indicating greater variance of the data is explained, the log-likelihood of the diffusion model is markedly lower.

Surprisal Model	R^2	Log-Likelihood	AIC	BIC
IDyOM	0.240	-1869.0	3744	3761
Audio-Diffusion	0.062	-1860.5	3727	3744

4 Discussion

This work is proposed primarily as a proof of concept that modern deep generative models can be used to model music expectancy. We believe that the experiments presented in this study, while not fully exploiting the power of state of the art diffusion models, adequately demonstrate that surprise scores from diffusion models do exhibit the *Wundt effect*, and thus validate the idea that diffusion models can be used to model music expectancy. On a practical level, the potential advantages of diffusion models over existing expectation models, e.g. IDyOM and D-REX, are the larger number of trainable parameters, allowing the ingestion of larger datasets, and the use of deep feature extractors which should be able to more accurately model complex distributions. Future research endeavors should involve replicating the experiment training on a variety of datasets, especially non-Western music, and studying the interaction of unexpectedness with other collative variables such as complexity and familiarity. Furthermore, as demonstrated by the low R^2 values of our experiments, intersubject variability is significantly challenging to capture with a single model. Significant future work should be directed towards training of individual subject-level models of expectancy, either through fine-tuning or self-supervised training. As a result, an application of this work could be the construction of an individual profile of serendipitous music recommendation systems [23]. In the same way as diffusion models were used by Wang [28] for music recommendation, our measure of musical expectation could predict the user’s liking profile from the audio level instead of the user’s listening trajectory. In addition, it would be useful to compare the likelihood score of the diffusion model with biosignal correlates of surprise, such as event-related potentials (ERPs), as done in [1].

Acknowledgments

We would like to thank Ben Gold [6] for generously providing us with the musical stimuli and human ratings necessary for performing this study.

References

- [1] Ellie Bean Abrams, Eva Muñoz Vidal, Claire Pelofi, and Pablo Ripollés. Retrieving musical information from neural data: how cognitive features enrich acoustic ones, December 2022. Pages: 160-168 Publication Title: Proceedings of the 23rd International Society for Music Information Retrieval Conference Publisher: ISMIR.
- [2] D. E. Berlyne. *Aesthetics and psychobiology*. Aesthetics and psychobiology. Appleton-Century-Crofts, East Norwalk, CT, US, 1971. Pages: xiv, 336.
- [3] Daniel E. Berlyne. *Conflict, arousal and curiosity*. Martino Publishing, Mansfield Centre, CT, 1960.
- [4] Anthony Chmiel and Emery Schubert. Back to the inverted-U for music preference: A review of the literature. *Psychology of Music*, 45(6):886–909, November 2017. Publisher: SAGE Publications Ltd.
- [5] Zhe Fu and Xi Niu. Modeling Users’ Curiosity in Recommender Systems. *ACM Transactions on Knowledge Discovery from Data*, 2023. Just Accepted.
- [6] Benjamin P. Gold, Marcus T. Pearce, Ernest Mas-Herrero, Alain Dagher, and Robert J. Zatorre. Predictability and Uncertainty in the Pleasure of Music: A Reward for Learning? *Journal of Neuroscience*, 39(47):9397–9409, November 2019. Publisher: Society for Neuroscience Section: Research Articles.
- [7] David Hargreaves and A. North. Hargreaves, D.J. & North, A.C. (2010). Experimental aesthetics and liking for music. In P. N. Juslin and J. A. Sloboda (eds.), *The Handbook of Music and Emotion: Theory, Research, Applications*. Oxford: OUP, pp. 515-546. ISBN 978-0-19-923014-3., 2010.
- [8] Hauke R. Heekeren, Sean Marrett, and Leslie G. Ungerleider. The neural systems that mediate human perceptual decision making. *Nature Reviews. Neuroscience*, 9(6):467–479, June 2008.
- [9] Hermann von Helmholtz. *Handbuch der physiologischen Optik*. Allgemeine Encyklopädie der Physik. Voss, Leipzig, 1867. OCLC: 3431763.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models, December 2020. arXiv:2006.11239 [cs, stat].
- [11] Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant Diffusion for Molecule Generation in 3D, June 2022. arXiv:2203.17003 [cs, q-bio, stat].
- [12] Patrik N. Juslin. From everyday emotions to aesthetic emotions: Towards a unified theory of musical emotions. *Physics of Life Reviews*, 10(3):235–266, September 2013.
- [13] Diederik P. Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational Diffusion Models, April 2023. arXiv:2107.00630 [cs, stat].
- [14] L.B. Meyer. *Emotion and meaning in music*. Emotion and meaning in music. Chicago, University of Chicago Press, 1956.
- [15] Eugene Narmour. *The analysis and cognition of basic melodic structures: the implication-realization model*. University of Chicago Press, Chicago, 1990.
- [16] Eugene Narmour. The top-down and bottom-up systems of musical implication: Building on Meyer’s theory of emotional syntax. *Music Perception*, 9(1):1–26, 1991. Place: US Publisher: University of California Press.
- [17] Mark G. Orr and Stellan Ohlsson. Relationship Between Complexity and Liking as a Function of Expertise. *Music Perception*, 22(4):583–611, 2005. Place: US Publisher: University of California Press.
- [18] P.-Y. Oudeyer, J. Gottlieb, and M. Lopes. Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. *Progress in Brain Research*, 229:257–284, 2016.

- [19] Thomas Parr, Giovanni Pezzulo, and K. J. Friston. *Active inference: the free energy principle in mind, brain, and behavior*. The MIT Press, Cambridge, Massachusetts, 2022.
- [20] M. T. Pearce. *The construction and evaluation of statistical models of melodic structure in music perception and composition*. doctoral, City University London, December 2005.
- [21] Marcus T. Pearce. Statistical learning and probabilistic prediction in music cognition: mechanisms of stylistic enculturation. *Annals of the New York Academy of Sciences*, 1423(1):378–395, May 2018.
- [22] Mark Reybrouck, Peter Vuust, and Elvira Brattico. Neural Correlates of Music Listening: Does the Music Matter? *Brain Sciences*, 11(12):1553, November 2021.
- [23] Markus Schedl, David Hauger, and Dominik Schnitzer. A Model for Serendipitous Music Retrieval. *ACM International Conference Proceeding Series*, May 2012.
- [24] Robert Dargavel Smith. audio-diffusion, September 2023. original-date: 2022-08-07T09:11:37Z.
- [25] Neil Todd. Towards a cognitive theory of expression: The performance and perception of rubato. *Contemporary Music Review*, 4(1):405–416, January 1989. Publisher: Routledge _eprint: <https://doi.org/10.1080/07494468900640451>.
- [26] Peter Vuust, Ole A. Heggli, Karl J. Friston, and Morten L. Kringelbach. Music in the brain. *Nature Reviews Neuroscience*, 23(5):287–305, May 2022.
- [27] Peter Vuust and Maria A. G. Witek. Rhythmic complexity and predictive coding: a novel approach to modeling rhythm and meter perception in music. *Frontiers in Psychology*, 5, 2014.
- [28] Wenjie Wang, Yiyang Xu, Fuli Feng, Xinyu Lin, Xiangnan He, and Tat-Seng Chua. Diffusion Recommender Model, April 2023. arXiv:2304.04971 [cs].
- [29] Maria A. G. Witek, Eric F. Clarke, Mikkel Wallentin, Morten L. Kringelbach, and Peter Vuust. Syncopation, Body-Movement and Pleasure in Groove Music. *PLOS ONE*, 9(4):e94446, 2014. Publisher: Public Library of Science.
- [30] Wilhelm Wundt. *Principles of physiological psychology*. (e. b. titchener, trans.). leipzig, germany: Engelmann. edition, 1874.
- [31] Reza Jafari Ziarani and Reza Ravanmehr. Serendipity in Recommender Systems: A Systematic Literature Review. *Journal of Computer Science and Technology*, 36(2):375–396, April 2021.