**Pitches that Wire Together Fire Together: Scale Degree Associations Across Time Predict
Melodic Expectations**

Niels J. Verosky¹ and Emily Morgan²

¹San Francisco, CA 94117

niels.verosky@gmail.com

²Department of Linguistics, University of California, Davis

Abstract

The ongoing generation of expectations is fundamental to listeners' experience of music, but research into types of statistical information that listeners extract from musical melodies has tended to emphasize transition probabilities and n -grams, with limited consideration given to other types of statistical learning that may be relevant. Temporal associations between scale degrees represent a different type of information present in musical melodies that can be learned from musical corpora using expectation networks, a computationally simple method based on activation and decay.

Expectation networks infer the expectation of encountering one scale degree followed in the near (but not necessarily immediate) future by another given scale degree, with previous work suggesting that scale degree associations learned by expectation networks better predict listener ratings of pitch similarity than transition probabilities. The current work outlines how these learned scale degree associations can be combined to predict melodic continuations and tests the resulting predictions on a dataset of listener responses to a musical cloze task previously used to compare two other models of melodic expectation, a variable-order Markov model (IDyOM) and Temperley's music-theoretically motivated model. Under multinomial logistic regression, all three models explain significant unique variance in human melodic expectations, with coefficient estimates highest for expectation networks. These results suggest that generalized scale degree associations informed by both adjacent and non-adjacent relationships between melodic notes influence listeners' melodic predictions above and beyond n -gram context, highlighting the need to consider a broader range of statistical learning processes that may underlie listeners' expectations for upcoming musical events.

Keywords: Music, statistical learning, melody, expectation, prediction, expectation networks, idyom, corpus

1 Introduction

Listeners' experience and enjoyment of music is shaped by the ongoing formation of expectations that are to varying degrees fulfilled or undermined over time (Huron, 2006; Meyer, 1956; Rohrmeier & Koelsch, 2012; Stevens, 2012; Tillman, Poulin-Charronnat, & Bigand, 2013). Musical stimuli can establish expectations across domains such as rhythm (Hannon & Trehub, 2005; Keller & Schubert, 2011), timbre (Crowder, 1989), and harmony (Loui & Wessel, 2007; Schmuckler, 1989), but one type of musical expectation highly relevant to many musical styles is melodic expectation, which involves making predictions about upcoming pitch events in light of some recently heard melodic context. Given the central role of melody in a range of musical styles, understanding how melodic expectations are formed and what information these expectations derive from may be essential to elucidating cognitive processes underlying the perception of music more generally.

To this end, several types of statistical models have previously been put forward to account for listeners' predictions as to which continuations of a given melody are most likely. Since it is intuitively plausible that more frequently occurring melodic sequences should be more expected, models of melodic expectation that make use of transition probabilities and n -grams have been common. An exemplar of this approach is Information Dynamics of Music, or IDyOM, which uses variable-order Markov models to assign probabilities to melody continuations (Pearce, 2005; Pearce & Wiggins, 2004; Pearce & Wiggins, 2006). IDyOM is motivated by the idea that statistical learning applies across domains, and that melodic expectations are just one instance of this domain-general learning. IDyOM's predictions have been shown to correlate with listener judgments of melodic uncertainty (Hansen & Pearce, 2014) and electrophysiological responses to melodic stimuli (Pearce, Ruiz, Kapasi, Wiggins, & Bhattacharya, 2010) as well as phenomena such as melodic segmentation (Pearce, Müllensiefen, & Wiggins, 2010), expressive timing (Gingras et al., 2016), and perceived musical complexity (Sauvé & Pearce, 2019).

Stereotypical melodic sequences as captured by Markov models are by no means the only type of information previously hypothesized to contribute to melodic expectations. The Implication-Realization model (Narmour, 1990; Schellenberg, 1997) suggests several principles motivated by music theory, related to intervallic information and pitch register, that might influence listeners' melodic predictions. Temperley's Gaussian model (Temperley, 2008; Temperley, 2014) combines two such principles (a tendency for pitches to fall in a normally distributed range and a tendency for melodies to comprise small note-to-note intervals) with scale degree frequency profiles, using a Bayesian framework to predict continuations given some melodic context and training corpus. In contrast to IDyOM, Temperley's model represents melodic expectations in terms of domain-specific, Gestalt-like rules derived from music theory.

Morgan, Fogel, Nair, and Patel (2019) tested the IDyOM and Temperley models as representatives of statistical learning and Gestalt-like theories of melodic expectation respectively. Melodic expectations were measured using a musical cloze task in which listeners were prompted to sing continuations of melodic excerpts (Fogel, Rosenberg, Lehman, Kuperberg, & Patel, 2015). Using multinomial logistic regression, Morgan et al. found that both models predicted unique variance in listener-supplied melodic continuations from the melodic cloze task, with IDyOM generally having higher coefficient estimates. That result suggests that both n -gram probabilities and intervallic information may shape listeners' melodic predictions, with n -gram probabilities possibly having a larger effect. However, both models tested also left variance in human expectations unexplained and in particular tended to underestimate listeners' expectations of hearing the tonic scale degree at implied authentic cadences.

Closely related to the idea of expectation in music is that of *association*, insofar as musical elements constituting more expected upcoming events tend to be more closely associated with elements constituting recent context on some level. Associations between musical elements have sometimes been modeled using principles of spreading activation and decay, inspired by neural learning. These

models have used neural networks (Bharucha & Todd, 1989) and self-organizing maps (Collins, Tillman, Barrett, Delbé, & Janata, 2014; Toivianen & Krumhansl, 2003) to infer associations between chords, keys, and scale degrees, with Tillmann, Bharucha, and Bigand (2000) showing that implicit tonal associations learned by self-organizing maps can predict empirical effects involving judgments of harmonic and melodic similarity and expectedness.

Recently, expectation networks have been put forward as a method for learning scale degrees' generalized tonal expectations across a training corpus using principles of activation and decay (Verosky, 2019). In expectation networks, the connection from one scale degree to a second is strengthened in proportion with the tendency for the second scale degree to appear in close succession to the first (though not necessarily *immediately* after) in the training corpus. Thus, pitch classes (represented as chromatic scale degrees) are associated according to their tendency to appear in close temporal succession, in a direct analogy to neural learning that can be summarized colloquially by saying that “pitches that fire together wire together.” Verosky (2019) shows that scale degree associations learned from a corpus of Schubert vocal melodies better predict empirical ratings of pitch similarity (Krumhansl, 1979; Krumhansl 1990) than either bigram probabilities from the same corpus or the positions of the individual pitches in the Krumhansl-Kessler tonal hierarchy (Krumhansl & Kessler, 1982).

Expectation networks represent a type of statistical learning underpinned by activation-based associations formed between events proximate in time. Converging evidence for the plausibility of this type of statistical learning comes from Endress and Johnson (2021), which shows that networks based on Hebbian learning of temporal correlations between events reproduce statistical learning phenomena observed in previous studies of both adjacent (Saffran, Aslin, & Newport, 1996) and non-adjacent (Endress & Bonatti, 2007) transition probabilities. Those results further suggest a cross-domain relevance of statistical learning in the form of generalized, activation-based associations between temporally proximate events, beyond a possible role in musical expectations. Unlike the type of

statistical learning captured by bigrams, that captured by expectation networks represents the association between one musical scale degree and another not as the tendency for the second to appear immediately adjacent to the first but as the tendency for the second to follow the first in the near future, with the association between the two increasing the fewer intervening events there tend to be.

The current work shows how generalized scale degree associations learned by expectation networks can be combined with a computationally simple algorithm to predict melody continuations using principles of activation and decay. We extend Morgan et al. (2019) by comparing the predictions generated by expectation networks with those from the IDyOM and Temperley models using multinomial logistic regression. Under this experimental design, a significant effect of expectation networks suggests that the tonal associations of recently heard pitches may shape melodic expectations above and beyond the n -gram or intervallic context that those pitches are embedded in.

2 Methods

2.1 Melodic Prediction with Expectation Networks

Expectation networks provide a computationally simple framework for learning the expectation of encountering a given scale degree followed by another given scale degree in the near, but not necessarily immediate, future. An expectation network comprises a 12x12 set of connections corresponding to expectations between all possible pairs of scale degrees. The expectation network learning algorithm works as follows: each chromatic scale degree is assigned an initial activation of 0. Training melodies are read through one note at a time, and a scale degree's activation is set to 1 when an instance of that scale degree is encountered. At every time step (i.e., every subsequent note), the expectation network connections from all scale degrees to the current scale degree are increased by the respective activations of those scale degrees, and all activations are then divided by a decay constant. Because more recently encountered scale degrees have higher activations, the relative connection strength from one scale degree to another represents the tendency of the latter scale degree to occur following the former closely in time.

Verosky (2019) discusses the expectation network learning algorithm in depth and notes that although expectation networks can be expressed as weighted sums of n -grams, the learning algorithm has linear time complexity and requires maintaining only a 12x12 expectation matrix and a length-12 activation vector. Expectation networks infer the expectation of encountering any scale degree followed in close succession by any other scale degree, and these learned expectations were shown in Verosky (2019) to be a good fit with listener ratings of pitch similarity (where the similarity between two pitches is not symmetrical in time). The result of the learning algorithm is a 12x12 matrix E where $E_{i,j}$ represents the relative strength of the expectation of encountering chromatic scale degree i followed in the near future by chromatic scale degree j .

However, these generalized scale degree expectations can also be combined to predict continuations of some melodic context. Given a 12x12 matrix E of learned expectations between chromatic scale degrees, a decay constant λ (set to 1.618 here, as in Verosky, 2019), and a melodic sequence M , let the activation A of scale degree sd at time step t in the melody be the following:

$$A_{sd}^t = \frac{A_{sd}^{t-1}}{\lambda} + E_{M_t, sd}$$

where A is initialized as an all-zero vector and $E_{M_t, sd}$ represents the learned expectation from the scale degree at time step t of the melody to scale degree sd . Under this setup, the probability distribution of expected melodic continuations at time $t + 1$ is simply the normalized activation vector after time t .

Intuitively, activations for all possible scale degrees are updated at each time step to reflect the learned scale degree expectations associated with the current melody note. Activations decay over time, so more recent melodic events exert an exponentially stronger influence on overall expectations. In terms of time and space complexity, this algorithm for generating melodic expectations is equivalent to the algorithm for learning generalized scale degree associations, as can be seen in the Python implementation included in the supplementary materials. Thus, the expectation network algorithm

provides a computationally simple method for anticipating upcoming melodic events by combining the generalized scale degree expectations associated with recently heard notes.

The learned scale degree expectation $E_{M_t, sd}$ can be written as a weighted sum of n -grams, where the expectation is proportional to the number of n -grams in the training corpus that begin with M_t and end with sd , prioritizing shorter n -grams and excluding n -grams that contain additional instances of M_t . Because the learned expectation $E_{M_t, sd}$ can be written as a weighted sum of n -grams, so can the probability assigned to encountering sd at time step $t + 1$ in the melody:

$$\begin{aligned}
 P(sd|M_0, \dots, M_t) &= \frac{\frac{A_{sd}^{t-1}}{\lambda} + E_{M_t, sd}}{\sum_{sd'=1}^{12} \frac{A_{sd'}^{t-1}}{\lambda} + E_{M_t, sd'}} \\
 &= \frac{\sum_{t'=0}^t \lambda^{t'-t} E_{M_{t'}, sd}}{\sum_{sd'=1}^{12} \sum_{t'=0}^t \lambda^{t'-t} E_{M_{t'}, sd'}} \\
 &= \frac{\sum_{t'=0}^t \lambda^{t'-t} \sum_{n=2}^N \lambda^{2-n} |M_{t'} \rightarrow sd|_n}{\sum_{sd'=1}^{12} \sum_{t'=0}^t \lambda^{t'-t} \sum_{n=2}^N \lambda^{2-n} |M_{t'} \rightarrow sd'|_n} \\
 &= \frac{\sum_{t'=0}^t \sum_{n=2}^N \lambda^{t'-n} |M_{t'} \rightarrow sd|_n}{\sum_{sd'=1}^{12} \sum_{t'=0}^t \sum_{n=2}^N \lambda^{t'-n} |M_{t'} \rightarrow sd'|_n}
 \end{aligned}$$

where N is the length of the longest excerpt in the training corpus and $|i \rightarrow j|_n$ is the number of n -grams in the corpus that begin with i , end with j , and have no intervening instances of i (in other words, the number of skip-grams from i to j with no instances of i in between; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013). Therefore, sd will have a higher expectation if it has higher values of $|M_{t'} \rightarrow sd|_n$ for scale degrees heard more recently in the melody (greater t') that more closely precede sd in the training corpus (smaller n). While probabilities generated by expectation networks can be expressed as weighted sums of n -grams, the networks do not separately compute those constituent n -grams, hence their computational simplicity.

The decay parameter used here ($\lambda \approx \phi$, *the golden ratio*) was suggested in Verosky (2019) on the basis that it assigns the previous event equal weight to its two preceding events combined. Verosky further shows that using ϕ as the decay parameter bounds the network's total activation between 1 and $\phi + 1$ during the learning phase, resulting in similar associations as when the decay parameter is varied dynamically in proportion to the network's total activation. By analogous reasoning, setting $\lambda = \phi$ in the melodic prediction phase presented here results in a maximum network activation of ϕ^2 times the strongest learned scale degree association in the network, which would trivially bound network activation between 0 and $\phi + 1$ if the network were rescaled to have a maximum scale degree association of 1 (as in the learning phase, these conclusions follow from the facts that $\phi^0 + \dots + \phi^{-12} \approx \phi^2$ and that $\phi^n = \phi^{n-1} + \phi^{n-2}$; Verosky, 2019). No attempt is made here to optimize the decay parameter to the melodic expectation task specifically. However, Endress and Johnson (2021) test how different rates of “forgetting,” operationalized as decay parameters, influence a Hebbian network's ability to learn adjacent and non-adjacent transition probabilities. Their results suggest that an intermediate level of forgetting is optimal insofar as very high decay parameters prevent learning of associations between events and very low decay parameters lead all events to be associated indiscriminately. Since setting $\lambda = \phi$ causes scale degrees' activations to decline by 38 percent with each subsequently heard scale degree, it can be seen as representing an intermediate level of forgetting.

An additional modification that can be made to the melodic prediction algorithm is to discount repetitions. Specifically, we might choose to reset the expectation for a scale degree to 0 when an instance of that scale degree is encountered because the expectation of encountering that scale degree has been fulfilled. This version of the melodic prediction algorithm can be represented as follows:

$$A_{sd}^t = \frac{A_{sd}^{t-1}}{\lambda} + E_{M_t, sd} \text{ if } sd \neq M_t$$

$$A_{sd}^t = E_{M_t, M_t} \text{ otherwise}$$

After time step t , the activation for the most recent scale degree in the melody will be equal only to the expectation E_{M_t, M_t} of that scale degree to itself because the previous melodic context $\frac{A_{sd}^{t-1}}{\lambda}$ has been erased. The expression for sd 's probability in terms of n -grams remains the same when repetition discounting is applied except that scale degree associations in the current melody are summed from $t_r \leq t' \leq t$ in the numerator and $t_{r'} \leq t' \leq t$ in the denominator rather than $0 \leq t' \leq t$, where t_r and $t_{r'}$ are the most recent occurrences of scale degrees sd and sd' .

In the experiment that follows, we will generate expectations by combining the discounting and non-discounting versions of the algorithm – which is equivalent to partially discounting repetitions. This approach is motivated by a lack of prior evidence as to which version of the algorithm best models listeners' melodic expectations and the hypothesis that listeners likely treat repetitions somewhere in between the two extremes – in other words, that hearing a highly anticipated scale degree may somewhat attenuate but not entirely erase the expectation of again hearing that scale degree.

2.2 Comparison Models

We compare expectation networks' melodic predictions with those of two models tested in Morgan et al. (2019), Temperley's Gaussian model and IDyOM. Melodic predictions by both comparison models are available in the supplementary materials to Morgan et al. Temperley's model predicts the probability of a melodic continuation as the product of its probabilities under three other distributions: a normally distributed pitch range around the central pitch of a melody (a range profile, with the central pitch itself coming from a normal distribution), a normally distributed pitch range around the most recent melodic note (a proximity profile), and a distribution of zeroeth-order pitch-class frequencies (a key profile). IDyOM uses variable-order Markov chains over some "viewpoints," or set of musical features to be considered. Following Hansen and Pearce (2014), the IDyOM predictions from Morgan et al (2019) use a linked viewpoint consisting of scale degree and pitch interval. In other words, every event (note) in a melody is represented by a combination of its scale

degree and the pitch interval it completes. This configuration of IDyOM is analogous to Temperley’s model in that it takes into account a melody’s constituent scale degrees and note-to-note intervals, but the model itself differs in that this information is used to learn n -grams that are then interpolated into melodic predictions. Probabilities for upcoming pitches result from combining n -grams with n less than some order bound that changes according to context, giving more weight to higher-order n -grams that have a greater number of occurrences relative the number of event types that have previously completed the n -gram (Pearce, 2005).

2.3 Training Corpora

We train the expectation networks and comparison models on two separate corpora: the Essen Folksong Collection of 6,249 mostly European folksongs and the Pearce-Wiggins corpus, consisting of 185 soprano lines from chorales by J.S. Bach (1892), 152 Nova Scotian songs (Creighton, 1932), and 566 German folksongs from the Fink subset of the Essen collection. These corpora, available in `***kern` format from Stanford University’s Center for Computer Assisted Research in the Humanities (KernScores, 2001), were previously used to train the Temperley (2008) and IDyOM (Pearce & Wiggins, 2006; Hansen & Pearce, 2014) models respectively and to compare those two models in Morgan et al. (2019).

2.4 Human Data

As in Morgan et al. (2019), we test melodic predictions on data from a musical cloze task developed by Fogel et al. (2015). The musical cloze data can be accessed through the Open Science Framework project containing supplementary materials to Morgan et al. In the original Fogel et al. task and data, 50 musically trained listeners (minimum five years’ musical experience in the previous ten years) heard 45 pairs of monophonic melody openings, or “melodic stems,” and were instructed to “sing the note you think comes next.” Melodic stems, which were written for Fogel et al.’s experiment in a tonal style by composer Jason Carl Rosenberg, are available in the supplementary materials for Morgan et al. (2019). Each pair of melodic stems included a stem ending with an implied authentic

cadence (AC condition) and an analogous stem with no implied cadence (NC condition). Paired melodies were matched in number of notes, contour, and rhythm, but a small number of notes in each melody were manipulated either to strongly suggest an upcoming tonic chord (AC condition) or a more ambiguous range of possible harmonic continuations (NC condition). AC stems ended with scale degree 2, 5, or 7 setting up an implied dominant harmony while corresponding NC stems ended with scale degree 1, 3, 4, 5, or 6 setting up an implied non-dominant harmony. Responses from the two experimental conditions therefore reflect listeners' expectations under varying levels of certainty—higher certainty in the AC condition and lower certainty in the NC condition.

To address the possibility that responses were influenced by a desire to provide musical closure, Morgan et al. performed a follow-up experiment in which 50 musically trained listeners (according to the same musical training criteria) were prompted to “complete [the melody] by singing up to a few notes.” By analyzing the first note of these longer continuations, they could ensure that these expectations were not constrained by a desire for closure. They found that results were in fact remarkably similar across the two versions of the experiment.

We use the data from these two experiments to compare the predictions of expectation networks with predictions from the IDyOM and Temperley models. Further details regarding the data and implementation of the comparison models can be found in Morgan et al. (2019), which tested melodic expectations generated by the IDyOM and Temperley models for the 45 melodic stem pairs as predictors of listeners' expectations on both versions of the melodic cloze task. We now extend Morgan et al.'s analysis by adding expectation networks as a predictor.

2.5 Analyses

2.5.1 Multinomial Logit Modeling

Multinomial discrete-choice logit modeling is used to test the IDyOM, Temperley, and expectation network predictions as predictors of human responses on the melodic cloze task. Under this type of model, the probability of a specific outcome category (in this case, a scale degree for the next

note of the melody) is given by a softmax function of a linear combination of predictors (here, log probabilities generated by the models being tested). The coefficient estimate for a given predictor indicates how a one-unit change in that predictor affects the log of the ratio of probabilities between two categories, also known as the log odds (for example, how the log of the ratio of the probabilities of scale degree 1 and scale degree 5 in the human data changes with a one-unit change in IDyOM's predictions for those scale degrees). The logit model is fit using maximum likelihood estimation, where a single observation comprises a categorical dependent variable denoting whether a given participant chose a particular outcome (scale degree) to continue a given melody on a single trial. Continuous independent variables for each melodic expectation model (IDyOM, Temperley, and expectation networks) denote the model's prediction for the given scale degree as a continuation of the melodic context. Since predictions from all three models tested here are represented as log probabilities, the resulting coefficient estimates between predictors (i.e., melodic expectation models) are directly comparable, with larger coefficients indicating larger effects. In logit models with multiple predictors, a statistically significant effect indicates that a predictor (model) explains unique variance in the human data above and beyond other predictors (Wurm & Fisicaro, 2014; see also Agresti (2013) for an overview of multinomial logistic regression and Morgan et al. (2019) for a discussion of its application specifically to this task).

We use the `mlogit` R package (Croissant, 2019) to perform the discrete-choice logit modeling. We consider four models of melodic expectation: IDyOM and Temperley with the parameter settings previously used in Morgan et al. (2019), and expectation networks with and without repetition discounting. Tests performed include multinomial logistic regression on individual models of melodic expectation as standalone predictors of listeners' expectations (Analysis 1), multinomial logistic regression on all three models of melodic expectation in combination (Analysis 2), and post hoc consideration of possible variations in the analytical design (Analyses 3 and 4).

2.5.2 Analysis 1

In Analysis 1, we fit discrete-choice logit models for all three types of melodic expectation models independently as predictors of human expectations:

human $\sim$ IDyOM

human $\sim$ Temperley

human $\sim$ expectation network with repetition discounting + expectation network without repetition discounting

For these and all further regression models, we include random slopes and intercepts at the subject level. As in Morgan et al. (2019), we include only in-key human responses and collapse scale degrees across octaves. Likewise, we fit separate models for the combined AC+NC and NC-only conditions for both versions of the melodic cloze task and the two training corpora. Fitting to data from the combined AC+NC conditions allows models' performance to be evaluated across melodies that set up varying degrees of certainty in expectations. Meanwhile, fitting to only the NC condition excludes from consideration AC melodies, where the IDyOM and Temperley models have previously been shown to underestimate the probability of the tonic scale degree (Morgan, 2019), allowing performance to be evaluated for the subset of melodies that tend to set up weaker expectations.

The coefficient estimates obtained for the two versions of the expectation network algorithm entered as separate predictors allow us to interpolate, based on the human data, a version of the expectation network algorithm that partially discounts repetitions. Let $\overline{P}_i$ and $\underline{P}_i$ be the predictions for scale degree i under the expectation network models with and without melodic discounting respectively, and let $\overline{B}$ and $\underline{B}$ be the coefficient estimates obtained under the multinomial logit model fit with these predictors. The interpolated expectation P_i for scale degree i with partial discounting is then:

$$P_i = \frac{\exp(\overline{B}\overline{P}_i + \underline{B}\underline{P}_i)}{\sum_{j=1}^{12} \exp(\overline{B}\overline{P}_j + \underline{B}\underline{P}_j)}$$

2.5.3 Analysis 2

Next, in Analysis 2, we account for shared variance among models of melodic expectation by fitting the following multinomial logit model:

human $\sim$ IDyOM + Temperley + expectation network with partial repetition discounting

Multinomial logit modeling is performed as in Analysis 1. However, Analysis 2 includes all three models of melodic expectation in combination as predictors. Thus, the results of this analysis in particular indicate how much unique variance in melodic expectations each model explains while controlling for other models. In other words, Analysis 2 tests whether each model of melodic expectation has a significant association with listeners' expectations above and beyond other models, and the coefficient estimates from Analysis 2 allow for a comparison of each model's effect size while accounting for the other two models.

2.5.4 Post Hoc Analyses

In Analyses 3 and 4, to better understand the results obtained, we perform six post hoc variations of the above test (Analysis 2) with modified experimental designs. Analysis 3.1 uses expectation network interpolation parameters learned from the AC+NC condition to fit the multinomial logit model to data from the AC condition, which is intended to check whether the results are robust to interpolating the expectation network predictions using coefficient estimates learned across experimental conditions. Analysis 3.2 controls for scale-degree-to-scale-degree transition probabilities by adding bigram probabilities as a fourth predictor in the regression model. We then address the possibility that excluding out-of-key predictions may have influenced the results by including a separate out-of-key outcome category, with predictions represented as log probabilities and raw probabilities in Analyses 3.3 and 3.4 respectively. Analysis 3.5 drops the Temperley model as a predictor for a direct comparison between IDyOM and partially discounted expectation networks.

Finally, Analysis 4 examines how meaningful the difference in coefficient estimates between IDyOM and expectation networks is by combining the two into a single predictor under the

multinomial logit model. Probabilities from IDyOM and partially discounted expectation networks are interpolated with a softmax function, as the probabilities for expectation networks with and without repetition discounting were, but with the difference that the interpolation weights for IDyOM and expectation networks are set equal. This setup is equivalent to fitting the multinomial logit model while constraining the coefficient estimates for IDyOM and expectation networks to be equal. We then use log-likelihoods resulting from maximum likelihood estimation to obtain a likelihood ratio for the full model relative to the constrained model and a p -value according to the likelihood-ratio test statistic.

3 Results

3.1 Analysis 1

In Analysis 1, predictions generated by IDyOM and the Temperley model both explained significant variance in human predictions when considered as standalone predictors in separate regression models, with $p < .001$ across all combinations of training corpora (Essen vs. Pearce-Wiggins), conditions (AC+NC vs. NC), and datasets (versions 1 and 2 of the melodic cloze task). In the regression model that included expectation networks with and without repetition discounting as independent variables, expectation networks with repetition discounting were significant at the $p < .001$ level in all cases while expectation networks without repetition discounting were significant in some but not all cases (Table 1). Coefficient estimates were consistently greater for expectation networks with repetition discounting. As outlined above, coefficient estimates for expectation networks with and without repetition discounting were used to interpolate predictions with partial repetition discounting for subsequent model comparisons. Log-likelihoods for individual models are provided in Supplementary Table 1.

Table 1

Condition	Training Corpus	Model	Dataset 1			Dataset 2		
			Coefficient Estimate	Standard Error	z-value	Coefficient Estimate	Standard Error	z-value

AC + NC	Essen	Expectation network with repetition discounting	1.42	.095	14.93***	1.17	.088	13.25***
		Expectation network without repetition discounting	.31	.120	2.55*	.24	.113	2.08*
	Pearce-Wiggins	Expectation network with repetition discounting	1.47	.098	14.98***	1.21	.091	13.31**
		Expectation network without repetition discounting	.77	.127	6.06***	.62	.120	5.13***
	NC	Expectation network with repetition discounting	.81	.124	6.51***	.68	.119	5.76***
		Expectation network without repetition discounting	.21	.148	1.43	.13	.139	.97
	Pearce-Wiggins	Expectation network with repetition discounting	.80	.125	6.34***	.67	.120	5.61***
		Expectation network without repetition discounting	.43	.153	2.83**	.32	.145	2.19*

Table 1: Multinomial logit model using expectation networks with and without repetition discounting as predictors

3.2 Analysis 2

Tables 2 and 3 show coefficient estimates, standard errors, z -values, and p -values for Analysis 2 (multinomial logistic regression using all three models as predictors), with human expectation dataset 1 (Table 2) and dataset 2 (Table 3) regressed on the Temperley, IDyOM, and interpolated expectation network models in combination. Corresponding log-likelihoods can be found in Supplementary Table 1. All three models of melodic prediction explain significant unique variance in human expectations, with coefficient estimates higher for IDyOM than the Temperley model, as in Morgan et al. (2019). However, coefficient estimates are highest for expectation networks. Since predictions by all three models are expressed here in the same units (log probabilities), this result indicates that expectation networks have a larger role in predicting the human data than IDyOM or the Temperley model.

Table 2

Condition	Training Corpus	Model	Coefficient Estimate	Standard Error	z -value	p -value
AC + NC	Essen	Temperley	.21	.038	5.45	$5.1 * 10^{-8}$
		IDyOM	.53	.031	17.39	$< 2.2 * 10^{-16}$
		Expectation networks	.58	.035	16.69	$< 2.2 * 10^{-16}$
	Pearce-Wiggins	Temperley	.23	.035	6.69	$2.3 * 10^{-11}$
		IDyOM	.45	.032	13.86	$2.3 * 10^{-11}$
		Expectation networks	.61	.030	20.00	$< 2.2 * 10^{-16}$
NC	Essen	Temperley	.30	.053	5.67	$1.4 * 10^{-8}$
		IDyOM	.33	.037	8.88	$< 2.2 * 10^{-16}$
		Expectation networks	.53	.078	6.85	$7.5 * 10^{-12}$
	Pearce-Wiggins	Temperley	.22	.048	4.51	$6.4 * 10^{-6}$
		IDyOM	.37	.038	9.86	$< 2.2 * 10^{-16}$
		Expectation networks	.53	.073	7.29	$3.1 * 10^{-13}$

Table 2: Multinomial logit model fit to Fogel et al. (2015) data (dataset 1) with all three models of melodic expectation as predictors

Table 3

Condition	Training Corpus	Model	Coefficient Estimate	Standard Error	z-value	p-value
AC + NC	Essen	Temperley	.09	.038	2.38	.017
		IDyOM	.51	.028	18.30	$< 2.2 * 10^{-16}$
		Expectation networks	.60	.039	15.41	$< 2.2 * 10^{-16}$
	Pearce-Wiggins	Temperley	.15	.034	4.45	$8.6 * 10^{-6}$
		IDyOM	.39	.029	13.27	$< 2.2 * 10^{-16}$
		Expectation networks	.63	.034	18.34	$< 2.2 * 10^{-16}$
NC	Essen	Temperley	.18	.051	3.51	.0005
		IDyOM	.34	.036	9.52	$< 2.2 * 10^{-16}$
		Expectation networks	.54	.093	5.86	$4.7 * 10^{-9}$
	Pearce-Wiggins	Temperley	.15	.048	3.05	.002
		IDyOM	.36	.039	9.10	$< 2.2 * 10^{-16}$
		Expectation networks	.51	.088	5.72	$1.0 * 10^{-8}$

Table 3: Multinomial logit model fit to Morgan et al. (2019) Experiment 2 data (dataset 2) with all three models of melodic expectation as predictors

Table 4 shows correlations between melodic predictions by expectation networks, IDyOM, and the Temperley model as used to predict the melodic cloze data. Correlations are near perfect between versions of the expectation networks using interpolation parameters from different datasets, high between versions of each model using different training corpora, and low to moderate between different models, consistent with the idea that the three models represent distinct types of melodic information.

Table 4

Condition	Training Corpus	Model	IDyOM	Expectation networks (Dataset 1)	Expectation networks (Dataset 2)	Across Corpora
AC	Essen	Temperley	.65	.47	.47	.92

NC	Pearce-Wiggins	IDyOM		.46	.45	.83
		Expectation networks (Dataset 1)			1.00	.97
		Expectation networks (Dataset 2)				.97
		Temperley	.62	.55	.55	
		IDyOM		.52	.51	
		Expectation networks (Dataset 1)			1.00	
	Essen	Temperley	.46	.44	.43	.90
		IDyOM		.31	.30	.73
		Expectation networks (Dataset 1)			1.00	.98
		Expectation networks (Dataset 2)				.98
		Temperley	.48	.40	.39	
		IDyOM		.39	.38	
		Expectation networks (Dataset 1)			1.00	

Table 4: Correlations between melodic predictions for the Temperley model, IDyOM, and expectation networks with interpolation parameters learned from dataset 1 or 2. The “Across Corpora” column provides correlations between versions of each model trained on the two corpora.

3.3 Post Hoc Analyses

Finally, Supplementary Tables 2-6 show all coefficient estimates, standard errors, and z -values for the five variations considered in Analysis 3. Rather than the results of any single analytical variation, the most relevant finding here is the pattern that emerges across the five tests. In particular, coefficient estimates remain highest for expectation networks across all five experimental modifications. Thus, the results appear to be robust to changes in experimental design such as using expectation network interpolation parameters learned across conditions, controlling for bigram

probabilities, including an out-of-key category, representing predictors on a linear rather than log scale, and dropping the Temperley model as a predictor to allow for a direct comparison between IDyOM and expectation networks.

Table 5 shows log-likelihoods, likelihood ratios, and results of likelihood-ratio tests for the full multinomial logit model from Analysis 2 and the variation considered in Analysis 4, in which IDyOM and expectation networks were combined and assigned a single coefficient estimate. Likelihood ratios and corresponding *p*-values indicate that the difference between expectation networks' and IDyOM's coefficient estimates may be more meaningful for the combination of cadential and non-cadential melodies (AC + NC), with 3 out of 4 model comparisons significant, than for non-cadential melodies only (NC), where only 1 out of 4 model comparisons is significant.

Table 5

Condition	Training Corpus	Model	Dataset 1	Dataset 2
AC + NC	Essen	Full Model	-3130.2	-3387.7
		Constrained Model	-3139.2	-3388.9
		Likelihood Ratio	8103.1	3.3
		<i>p</i>-value	$2.2 \times 10^{-5}***$	.12
	Pearce-Wiggins	Full Model	-3077.1	-3377.9
		Constrained Model	-3088.9	-3390
		Likelihood Ratio	1.3×10^5	1.8×10^5
		<i>p</i>-value	$1.2 \times 10^{-6}***$	$8.7 \times 10^{-7}***$
NC	Essen	Full Model	-1723.8	-1790.2
		Constrained Model	-1727.7	-1792.1
		Likelihood Ratio	49.4	6.7
		<i>p</i>-value	.005**	.051
	Pearce-Wiggins	Full Model	-1721.2	-1790.7
		Constrained Model	-1723.1	-1792.3
		Likelihood Ratio	6.7	5.0
		<i>p</i>-value	.051	.07

Table 5: Log-likelihoods, resulting likelihood ratios, and likelihood-ratio test p -values comparing the full multinomial logit model with separate coefficient estimates for all three predictors (Analysis 2) against the model constrained to have equal coefficient estimates for IDyOM and expectation networks

4 Discussion

To better understand how different types of information might factor into listeners' melodic expectations, we tested three computational models of melodic expectation as predictors of listener responses on a musical cloze task. Temperley's model captures intervallic expectations, drawing on domain-specific rules posited by music theorists regarding how melodic notes tend to be spaced in terms of pitch register. IDyOM implements domain-general statistical learning principles in a musical context, capturing information about stereotypical pitch sequences. Finally, expectation networks represent an alternate set of statistical relationships present in musical melodies, capturing generalized associations as to which scale degrees tend to occur close together in time. Expectation networks have recently been suggested as a method of inferring tonal associations for individual scale degrees (Verosky, 2019), and the current study tested whether such learned associations can be combined across sequences of scale degrees to predict melodic expectations, including after controlling for n -gram context or Gestalt-like rules based on intervallic information.

All three models of melodic expectation explained statistically significant unique variance in listeners' responses on the melodic cloze task. This result raises the possibility that stereotypical pitch sequences, intervallic expectations, and generalized scale degree associations all contribute to listeners' melodic expectations. None of these three types of information about pitch sequences are obviously reducible to either of the others, and it appears plausible that all three play a role in shaping listeners' predictions of melodic continuations. Still, the current design cannot establish whether all three models of melodic expectation capture distinct underlying cognitive processes – some of the models could, for example, be noisy reflections of a single underlying process.

Of the three models, expectation networks with partial repetition discounting had the highest estimated regression coefficients, indicating that the expectation networks' predictions explained the most unique variance in the human data. It should be noted that condition-specific parameters for interpolating the expectation networks with and without repetition discounting were learned from the human data, which was potentially advantageous to the expectation networks in the model comparisons. However, the results held when parameters learned from the AC+NC condition were used to interpolate predictions for the AC condition alone. The analysis was also robust to four other changes in design, including dropping the Temperley model as a parameter to allow for a direct comparison between IDyOM and the expectation networks.

Thus, it appears that listeners may be sensitive to temporal scale degree associations as captured by a computationally simple algorithm based on activation and decay. These scale degree associations can be inferred from musical corpora by increasing expectation weights between scale degrees that occur close together in time, and the resulting "expectation networks" may explain more unique variance in human melodic predictions than more complex variable-order Markov models.

Previous work suggests that generalized scale degree associations learned in this fashion better predict (asymmetrical) ratings of pitch similarity between pairs of scale degrees than transition probabilities between scale degrees (Verosky, 2019). The current work adds to that finding by showing that such learned associations between pairs of scale degrees can also be combined to predict melodic continuations. In the context of melodic prediction, model expectations are formed by adding associations of recently heard scale degrees, with expectations decaying as new melodic notes are heard such that more recent notes exert a stronger effect.

Interestingly, scale degree associations for recent notes appear to influence listeners' melodic predictions irrespective of the n -gram context those notes are embedded in, as can be seen from the fact that expectation networks explained unique variance beyond IDyOM. In the linguistic domain, Kolk, Chwilla, Herten, and Oor (2003) show that some semantic reversal anomalies (e.g., *the cat that fled*

from the mice) fail to elicit the N400 effect typically associated with semantic anomalies, indicating that individual words' semantic associations may influence semantic expectations independent of syntactic context. Meanwhile, Van Petten (2014) shows that corpus-based measures of word co-occurrence predict N400 amplitude (above and beyond other measures including free association strength). It may be that musical sequences are subject to similar cognitive processes, whereby co-occurrence information between scale degrees influences melodic expectations independent of n -gram context.

The current study did not explore cognitive mechanisms that might allow recently heard notes' scale degree associations to supersede n -grams in listeners' predictions of melodic continuations, but interference is one process that might plausibly contribute to such an effect. Verosky (2019) points out that decay as modeled in expectation networks can more precisely be thought of as representing interference than temporal decay per se (Berman, Jonides, & Lewis, 2009; Lewandowsky, Oberauer, & Brown, 2009; Oberauer & Lewandowsky, 2008; Spyra, Stodolak, & Woolhouse, 2019). Given that inserting additional musical material between pairs of pitches separated by a fixed duration interferes with recall (Pechmann & Mohr, 1992), an influence of recent notes' scale degree associations on melodic expectations irrespective of broader n -gram or intervallic context could be understood in light of interference. In particular, for n -gram context alone but not temporal scale degree associations to influence expectations, associations of recently heard individual scale degrees not congruent with n -gram probabilities would need to be suppressed, a process that could be subject to working memory constraints. Such a situation could arise, for example, if an n -gram containing an instance of the leading tone were to have a low overall probability of a tonic completion but the presence of the leading tone were to heighten the listener's expectation of a tonic resolution without regard for the broader n -gram context.

Our results suggest not only that expectation networks predict significant unique variance in human melodic expectations after taking into account IDyOM's predictions, but that expectation

networks may explain more unique variance than IDyOM. Expectation networks have two potentially advantageous properties that may contribute to this effect. One is that expectation networks are simple relative to variable-order Markov models, in terms of both time complexity and number of model parameters. IDyOM models may in fact be overfit to the training data, particularly given the relatively small corpora available for melodic expectation research (compared to, for example, natural language processing research). Similarly, it has been suggested that more compact representations of stimuli tend to be more cognitively salient (Chater & Vitányi, 2003). The other reason is that expectation networks learn from both adjacent and non-adjacent notes, using activation and decay to infer general temporal relationships between scale degrees, whereas Markov models learn only from transitions between successive notes. Endress (2010) shows that listeners track dependencies between non-adjacent notes in tonal melodies and suggests that models based on Hebbian learning might allow for tracking of non-adjacent associations between events without the explosion in computational complexity that using Markov models would entail. Expectation networks provide an example of a computationally simple, associative, non-adjacent model in this vein. Because the type of information captured by expectation networks is simple to track and does not require restricting attention to relationships between successive notes, it may be especially cognitively salient. Non-adjacent relationships may also be important in the domain of music specifically insofar as some notes might be perceived as resolving to notes occurring in the near but not immediate future. As an example, consider an AC melody from the melodic cloze task, where the leading tone F# could be heard as resolving to a tonic G provided by the listener even with the D at the end of the melodic stem separating the two (Fig. 1).

Figure 1

Figure 1: Melody AC1 from Fogel et al. (2015)/Morgan et al. (2019) with tonic resolution

Qualitatively, examining the model predictions and human data suggests that expectation networks may explain more unique variance in human expectations than the other two models partly because they distinguish well between the cadential and non-cadential melodic contexts. For expectation networks with partial repetition discounting, the tonic was the most expected melodic continuation for 38/45 AC melodies and 0/45 NC melodies, averaged across the Pearce-Wiggins and Essen training corpora (e.g., 41/45 for Pearce-Wiggins and 35/45 for Essen in the AC Condition). Those averages were 26.5/45 and 4/45 respectively for IDyOM, and 24.5/45 and 20/45 for the Temperley model. Expectation networks' ability to distinguish cadential and non-cadential contexts is consistent with the results of Analysis 4, where fitting a model constrained to have a single coefficient estimate for IDyOM and expectation networks revealed that the difference in coefficient estimates between IDyOM and expectation networks under the full multinomial logit model seems to be more meaningful for the AC+NC condition than the NC-only condition.

Additionally, expectation networks appear to give elevated predictions for all three notes of the tonic triad at cadence points, providing a better fit in some (e.g., Fig. 2, upper-left) but not all (Fig 2., upper-right) cases. As Fig 3. shows, both expectation networks and human listeners produce probability distributions with local maxima at tonic triad scale degrees when predictions are averaged across all cadential melodies. In non-cadential melodies, there were fewer obvious trends owing to the greater variety, but one qualitative pattern that should be noted is the tendency for IDyOM to sometimes be overly certain in its predictions relative to both expectation networks and human listeners (Fig. 2, lower-left).

Figure 2

Figure 2: Listener response and model probabilities for melody AC12 (upper-left), AC 1 (upper-right), NC27 (lower-left), and NC26 (lower-right). Expectation networks give elevated probabilities for all three notes of the tonic triad at cadence points, providing a good fit to human data in this case.

Figure 3

Figure 3: Models' and human listeners' predictions averaged across all authentic cadence (left column) and non-cadential (right column) melodies, with models trained on Pearce-Wiggins (top row) and Essen (bottom row) corpora

Despite the expectation networks' success in assigning the tonic the highest rank-order expectation specific to the AC condition, expectation networks tended to consistently underestimate the probability of listeners choosing the tonic at cadence points. The same pattern was previously noted for

predictions by the IDyOM and Temperley models in Morgan et al. (2019), and expectation networks do not appear to fully address this shortcoming. More generally, despite expectation networks' good fit to the human data relative to the other two models considered, much of the variance in listeners' melodic predictions remains unexplained, as is evident, for example, in Fig. 2, lower-right. Since all three models considered here explained significant unique variance in listener expectations, the expectation networks' predictions can also be interpolated with predictions from the IDyOM and Temperley models, weighting each model's predictions according to the model's coefficient estimate under multinomial logistic regression. In this sense, the discrete-choice logit models can be seen as incorporating expectation networks, n -grams, and intervallic information into a single model that best fits the melodic cloze data, and the predictions of that interpolated model averaged across all melodies are shown in Fig. 4. Future work might explore more parsimonious ways of combining these three types of information into a single model of melodic expectation. Still, although all three types of information appear to explain distinct aspects of listeners' melodic expectations, the interpolation of the three models that best fits the melodic cloze data continues to underestimate the overall probability of the tonic in the AC condition.

Figure 4

Figure 4: Predictions averaged across authentic cadence (left column) and non-cadential (right column) melodies for human listeners and all three models interpolated according to the results of multinomial logistic regression, with models trained on Pearce-Wiggins (top row) and Essen (bottom row) corpora

This limitation is perhaps unsurprising insofar as there are a variety of non-local types of melodic relationships, such as those demonstrated in Kuhn and Dienes (2005) and Krumhansl (1996), that potentially contribute to melodic expectations but appear outside the scope of models predicting

melody continuations from n -gram context (as in IDyOM) or recent scale degree associations (as in expectation networks). None of the models considered here, for example, take into account forms of melodic parallelism created through transpositions or inversions of entire groups of notes or through rhythmic patterns. Thus, while expectation networks do learn scale degree associations informed by non-adjacent notes, there is a need for future work to consider the much wider universe of possible non-adjacent relationships that may influence melodic expectations. Expectation networks may capture one relatively local type of information derived from non-adjacent melodic relationships – specifically, information about generalized temporal associations between scale degrees.

Overall, the results of the present study are consistent with the idea that scale degree associations are learned outside the bounds of immediately successive notes and that, perhaps due to interference, generalized scale degree expectations combined across recently heard notes factor into melodic predictions above and beyond n -gram context. Such a process of learning and predicting temporal associations between scale degrees can apparently be modeled with a computationally simple algorithm, which should make it possible for future corpus-based and experimental work to more firmly establish which aspects of the information captured by expectation networks are most musically important and cognitively salient.

Open Practices

Data from Fogel et al. (2015) and Morgan et al. (2019) are available at

<http://dx.doi.org/10.17605/OSF.IO/WGZ9T>. Code for all analyses in the current study is available at

<http://dx.doi.org/10.17605/OSF.IO/HU8TA>, along with supplementary tables and graphs.

References

- Agresti, A. (2013). *Categorical data analysis*. Hoboken, NJ: John Wiley & Sons, Inc.
- Bach, J. S. (1892). *Bach-Gesellschaft Ausgabe, Band 39*. Leipzig: Breitkopf & Hrtel. Retrieved from <https://kern.humdrum.org/cgi-bin/ksbrowse?type=collection&l=/musedata/bach/chorales>
- Berman, M. G., Jonides, J., & Lewis, R. L. (2009). In search of decay in verbal short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 317-333.
- Bharucha, J. J., & Todd, P. M. (1989). Modeling the perception of tonal structure with neural nets. *Computer Music Journal*, 13(4), 44-53.
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7, 19-22.
- Collins, T., Tillman, B., Barrett, F. S., Delbé, C., and Janata, P. (2014). A combined model of sensory and cognitive representations underlying tonal expectations in music: From audio signals to behavior. *Psychological Review*, 120, 33-65.
- Creighton, H. (1932). Songs and ballads from Nova Scotia. In C. Sapp (Ed.), *Songs and ballads from Nova Scotia*. Menlo Park, CA: Center for Computer Assisted Research in the Humanities. Retrieved from <https://kern.humdrum.org/cgi-bin/ksbrowse?type=collection&l=/users/craig/songs/creighton/nova>
- Croissant, Y. (2019). mlogit: Multinomial logit models (R package version 0.4-2). Available from <https://CRAN.R-project.org/package=mlogit>
- Crowder, R. G. (1989). Imagery for musical timbre. *Journal of Experimental Psychology: Human perception and performance*, 15, 472-478.
- Endress, A. D. (2010). Learning melodies from non-adjacent tones. *Acta Psychologica*, 135, 182-190.
- Endress, A. D., & Bonatti, L. L. (2007). Rapid learning of syllable classes from a perceptually continuous speech stream. *Cognition*, 105, 247-299.

- Endress, A. D., & Johnson, S. P. (2021). When forgetting fosters learning: A neural network model for statistical learning. *Cognition*.
- Fogel, A. R., Rosenberg, J. C., Lehman, F. M., Kuperberg, G. R., & Patel, A. D. (2015). Studying musical and linguistic prediction in comparable ways: The melodic cloze probability method. *Frontiers in Psychology*, 6.
- Gingras, B., Pearce, M. T., Goodchild, M., Dean, R. T., Wiggins, G., & McAdams, S. (2016). Linking melodic expectation to expressive performance timing and perceived musical tension. *Journal of Experimental Psychology: Human Perception and Performance*, 42, 594-609.
- Hannon, E. E., & Trehub, S. E. (2005). Tuning in to musical rhythms: Infants learn more readily than adults. *Proceedings of the National Academy of Sciences*, 102, 12639-12643.
- Hansen, N. C., & Pearce, M. T. (2014). Predictive uncertainty in auditory sequence processing. *Frontiers in Psychology*, 5.
- Huron, D. (2006). *Sweet anticipation: Music and the psychology of expectation*. Cambridge, MA: MIT Press.
- Keller, P. E., & Schubert, E. (2001). Cognitive and affective judgements of syncopated musical themes. *Advances in Cognitive Psychology*, 7, 142-156.
- KernScores. (2001). Retrieved from <http://kern.ccarh.org/>
- Kolk, H. H. J., Chwilla, D. J., Van Herten, M., & Oor, P. J. W. (2003). Structure and limited capacity in verbal working memory: A study with event-related potentials. *Brain and Language*, 85, 1-36.
- Krumhansl, C. L. (1979). The psychological representation of musical pitch in a tonal context. *Cognitive Psychology*, 11, 346-374
- Krumhansl, C. L. (1990). *Cognitive foundations of musical pitch*. New York, NY: Oxford University Press.

- Krumhansl, C. L. (1996). Effects of perceptual organization and musical form on melodic expectancies. In M. Leman (Ed.), *Music, gestalt, and computing: Studies in cognitive and systematic musicology* (pp. 294-320). Berlin: Springer.
- Krumhansl, C. L., & Kessler, E. J. (1982). Tracing the dynamic changes in perceived tonal organization in a spatial representation of musical keys. *Psychological Review*, 89, 334-368.
- Kuhn, G., & Dienes, Z. (2005). Implicit learning of nonlocal musical rules: Implicitly learning more than chunks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 1417-1432.
- Lewandowsky, S., Oberauer, K., & Brown, G. D. A. (2009). No temporal decay in verbal short-term memory. *Trends in Cognitive Sciences*, 13, 120-126.
- Loui, P., & Wessel, D. (2007). Harmonic expectation and affect in Western music: Effects of attention and training. *Attention, Perception & Psychophysics*, 69, 1084-1092.
- Meyer, L. B. (1956). *Emotion and meaning in music*. Chicago, IL: University of Chicago Press.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. 2013. Distributed representations of words and phrases and their compositionality. *Proceedings of the 26th International Conference on Neural Information Processing Systems* (pp. 3111-3119).
- Morgan, E., Fogel, A., Nair, A., & Patel, A. D. (2019). Statistical learning and Gestalt-like principles predict melodic expectations. *Cognition*, 189, 23-34.
- Narmour, E. (1990). *The analysis and cognition of basic melodic structures: The Implication-Realization model*. Chicago, IL: University of Chicago Press.
- Oberauer, K., & Lewandowsky, S. (2008). Forgetting in immediate serial recall: Delay, temporal distinctiveness, or interference? *Psychological Review*, 115, 544-576.
- Pearce, M. T. (2005). *The construction and evaluation of statistical models of melodic structure in music perception and composition* (Doctoral dissertation).

- Pearce, M. T., Müllensiefen, D., & Wiggins, G. A. (2010). The role of expectation and probabilistic learning in auditory boundary perception: A model comparison. *Perception*, 39, 1367-1391.
- Pearce, M. T., Ruiz, M. H., Kapasi, S., Wiggins, G. A., & Bhattacharya, J. (2010). Unsupervised learning underpins computational, behavioural, and neural manifestations of musical expectation. *NeuroImage*, 50, 302-313.
- Pearce, M. T., & Wiggins, G. A. (2004). Improved methods for statistical modelling of monophonic music. *Journal of New Music Research*, 33, 367-385.
- Pearce, M. T., & Wiggins, G. A. (2006). Expectation in melody: The influence of context and learning. *Music Perception*, 23, 377-405.
- Pechmann, T., & Mohr, G. (1992). Interference in memory for tonal pitch: implications for a working-memory model. *Memory & Cognition*, 20, 314-320.
- Rohrmeier, M. A., & Koelsch, S. (2012). Predictive information processing in music cognition: A critical review. *International Journal of Psychophysiology*, 83, 164-175.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926-1928.
- Sauvé, S., & Pearce, M. T. (2019). Information-theoretic modeling of perceived musical complexity. *Music Perception*, 37, 165-178.
- Schaffrath, H. (1995). The Essen folksong collection. In D. Huron (Ed.), *The Essen folksong collection in **kern format*. Menlo Park, CA: Center for Computer Assisted Research in the Humanities. Retrieved from <http://kern.humdrum.org/cgi-bin/browse?l=/essen>
- Schellenberg, E. G. (1997). Simplifying the Implication-Realization model of melodic expectancy. *Music Perception*, 14, 295-318.
- Schmuckler, M. A. (1989). Expectation in music: Investigation of melodic and harmonic processes. *Music Perception*, 7, 109-149.

Stevens, C. J. (2012). Music perception and cognition: A review of recent cross-cultural research.

Topics in Cognitive Science, 4, 653-667.

Temperley, D. (2008). A probabilistic model of melody perception. *Cognitive Science, 32*, 418-444.

Temperley, D. (2014). Probabilistic models of melodic interval. *Music Perception, 32*, 85-99.

Tillmann, B., Bharucha, J. J., & Bigand, E. (2000). Implicit learning of tonality: A self-organizing approach. *Psychological Review, 107*, 885-913.

Tillman, B., Poulin-Charronnat, B., & Bigand, E. (2013). The role of expectation in music: From the score to emotions and the brain. Wiley Interdisciplinary Reviews: *Cognitive Science, 5*, 105-113.

Toivianen, P., & Krumhansl, C. L. (2003). Measuring and modeling real-time responses to music: The dynamics of tonality induction. *Perception, 32*, 741-766.

Van Petten, C. (2014). Examining the N400 semantic context effect item-by-item: Relationship to corpus-based measures of word co-occurrence. *International Journal of Psychophysiology, 94*, 407-419.

Verosky, N. J. (2019). Corpus-based learning of tonal expectations with expectation networks. *Journal of New Music Research, 45*, 145-158.

Wurm, L. H., & Fisicaro, S. A. (2014). What residualizing predictors in regression analysis does (and what it does not do). *Journal of Memory and Language, 72*, 37-48