Shannon entropy predicts perceptual uncertainty in the generation of melodic pitch expectations

Niels Chr. Hansen,[*#1] Marcus T. Pearce[*#2]

[*] *School of Electronic Engineering and Computer Science, Queen Mary, University of London, United Kingdom*
[#] *Department of Computing, Goldsmiths College, University of London, United Kingdom*
[1] niels.hansen@eecs.qmul.ac.uk, nch@musikkons.dk [2] marcus.pearce@eecs.qmul.ac.uk

ABSTRACT

Background

Recent theories propose that predictive mechanisms in the human brain enhance survival (Bar, 2007). The *predictive coding theory* explains how mental representations are continuously optimised through monitoring of prediction errors (Friston, 2005). Statistical learning associated with this process has been thoroughly demonstrated in the visual (Fiser & Aslin, 2002), somatosensory (Conway & Christensen, 2005) and auditory domains; including sequences of syllables (Saffran, 2003), timbres (Tillmann & McAdams, 2004) and pitches (Loui, Wessel, & Hudson Kam, 2010). This has led researchers to consider implicit, statistical learning a domain-general mechanism (Perruchet & Pacton, 2006).

In musical contexts, short-term acquisition of statistical regularities is enhanced by the extent of exposure (Jonaitis & Saffran, 2009) and deteriorated by grammatical complexity (Rohrmeier & Cross, 2009). Moreover, melodic pitch expectations have been shown to reflect probabilities of the general tonal repertoire internalised through long-term exposure (Pearce, Ruiz, Kapasi, Wiggins, & Bhattacharya, 2010). The latter study used a computational model, based on *n*-gram methods, acquiring knowledge through unsupervised learning of a large corpus of music (Pearce, 2005).

Previous studies have solely examined the expectedness of a single continuation at a time (often retrospectively, after it has occurred). The cognitive processes of generating expectations before a forthcoming event–and the predictive uncertainty that this entails–thus remain unexplored. Probe-tone studies collecting data from multiple continuations (e.g. Cuddy & Lunney) have aimed to test theories of melodic continuation (e.g. Narmour, 1990) and have thus not addressed predictive uncertainty about the next event before it arrives.

The information-theoretic concept of entropy quantifies the uncertainty involved in predicting values of a random variable, *X*, by taking the average information content across the discrete set of possible outcomes weighted according to probabilities (Shannon, 1948):

$$H(X) = -\sum_{i=1}^{n} p(x_i)\log_2 p(x_i)$$

Hitherto, musical applications of information theory have primarily used entropy as a way of distinguishing musical styles (Margulis & Beatty, 2008) or for information-retrieval purposes like melody identification (Duane, 2010), thus ignoring its psychological potential.

Aims

The present study aimed to assess Shannon entropy as a model of perceptual uncertainty in melodic pitch expectation. Four hypotheses were tested:

1) High-entropy contexts will elicit greater uncertainty than low-entropy contexts.
2) Musicians will generally show less uncertainty than non-musicians due to a more optimised internal cognitive model of probabilistic structure in melodies.
3) An entropy-by-expertise interaction for unexpectedness ratings will result from the fact that domain-specific training is more advantageous in low-entropy contexts.
4) Earlier findings of positive correlation between perceived unexpectedness and information content (Pearce et al., 2010) will be replicated.

Method

Seventeen musicians and 17 non-musicians listened to 24 melodic contexts, presented in random order, and provided ratings (9-point scale) on perceived uncertainty (*explicit uncertainty*). Each melody was subsequently presented a further nine times each time followed by one of nine chromatically distributed probe tones centred on the median pitch. All probe-tone stimuli were presented in random order, and expectedness ratings (9-point scale) were collected (cf. Krumhansl & Shepard, 1979). Entropy computed from normalised expectedness distributions resulted in another measure of uncertainty (*implicit uncertainty*).

Complex and *simple* stimuli were selected from Schubert songs and isochronous English hymns, respectively, and assigned to low- and high-entropy conditions according to predictions of an unsupervised, variable-order Markov Model (Pearce, 2005). A modelling viewpoint linking pitch interval and scale degree was used. Musical key was estimated using Temperley's (1999) optimised key-finding algorithm.

Results

Implicit uncertainty (Fig. 1), explicit uncertainty (Fig. 2) and unexpectedness (Fig. 3) data were subjected to 2x2x2 ANOVAs with *complexity* and *entropy* as within- and *expertise* as between-participant factors.

High-entropy contexts produced significantly greater implicit uncertainty for both complex and simple stimuli, $F(1, 22) = 42.52$, $p < .01$, and greater explicit uncertainty for simple stimuli, $F(1, 30) = 6.23$, $p = .02$. Averaged across participants, implicit uncertainty correlated with entropy, $r_s = .49$, $p = .02$, but explicit uncertainty did not, $r_s = .20$, $p = .35$.

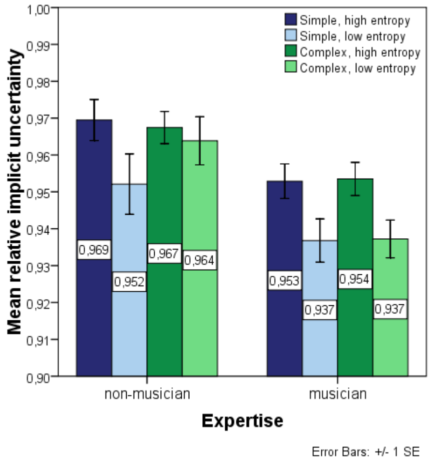

Figure 1. Mean implicit uncertainty.

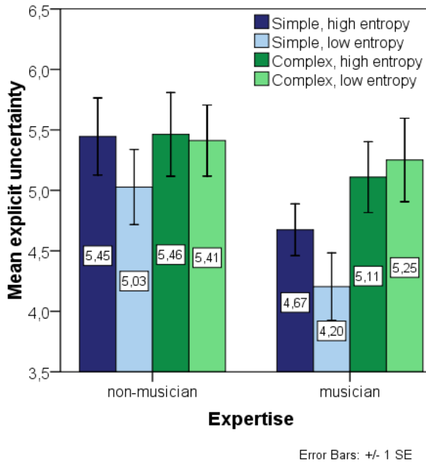

Figure 2. Mean explicit uncertainty.

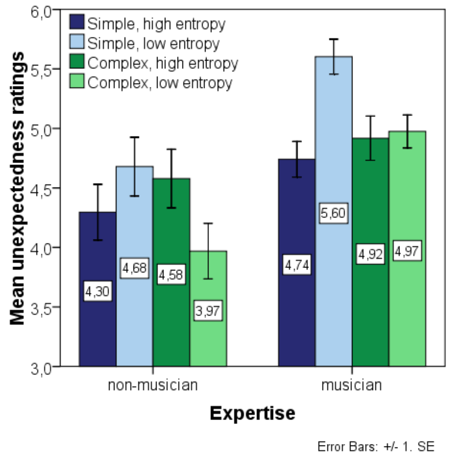

Figure 3. Mean unexpectedness.

Musicians experienced less implicit uncertainty for both complexity levels, $F(1, 22) = 5.49$, $p = .03$, and less explicit uncertainty for simple stimuli, $F(1, 30) = 4.00$, $p = .05$. Significant entropy-by-expertise, $F(1, 22) = 4.76$, $p = .04$, and complexity-by-entropy, $F(1, 22) = 4.95$, $p = .04$, interactions were found for implicit uncertainty. Moreover, complex stimuli produced higher explicit uncertainty, $F(1, 30) = 11.44$, $p < .01$, and an expertise-by-complexity interaction was found, $F(1, 30) = 4.09$, $p = .04$.

Unexpectedness ratings increased with information content, $r_s = .70$, $p < .01$. This effect was strongest in musicians, $t(215) = 3.47$, $p < .01$, and increased with musical training, $R^2 = .54$, $R^2_{adj.} = .52$, $F(1, 32) = 37.06$, $p < .01$. Additionally, the hypothesised entropy-by-expertise interaction was found for unexpectedness data, $F(1, 28) = 15.82$, $p < .01$.

Conclusions

Our results demonstrate for the first time that perceptual uncertainty reflects the Shannon entropy of an underlying process of probabilistic prediction. The strength of this effect is enhanced by: (a) simplicity in sensory input, (b) domain-relevant training, and (c) implicitness of uncertainty assessment.

Moreover, domain-relevant training leads to an increasingly accurate cognitive model of probabilistic structure. Without training, the default internal model appears to make predictions with relatively high entropy. Musical training, therefore, has the highest impact when sensory input generates probability distributions with relatively low entropy.

These findings are consistent with a statistical learning account of auditory cognition (Cristià, McGuire, Seidl, & Francis, 2011) as well as with the predictive coding theory (Friston, 2005).

Keywords

Statistical learning; music; melody; expectation; information theory; entropy; auditory cognition.

REFERENCES

Bar, M. (2007). The proactive brain: using analogies and associations to generate predictions. *Trends in Cognitive Sciences, 11*(7), 280-289. doi: 10.1016/j.tics.2007.05.005

Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. J Exp Psychol Learn Mem Cogn, 31(1), 24-39.

Cristià, A., McGuire, G. L., Seidl, A., & Francis, A. L. (2011). Effects of the distribution of acoustic cues on infants' perception of sibilants. J Phon, 39(3), 388-402. doi: 10.1016/j.wocn.2011.02.004

Cuddy, L. L., & Lunney, C. A. (1995). Expectancies generated by melodic intervals: Perceptual judgments of melodic continuity. Atten Percept Psychophys, 57(4), 451.

Duane, B. (2010). Information content in melodic and non-melodic lines. Paper presented at the 11th International Conference on Music Perception and Cognition (ICMPC), Seattle, WA, USA.

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. J Exp Psychol Learn Mem Cogn, 28(3), 458-467.

Friston, K. (2005). A theory of cortical responses. Philosophical Transactions of the Royal Society B: Biological Sciences, 360(1456), 815-836.

Krumhansl, C. L., & Shepard, R. N. (1979). Quantification of the hierarchy of tonal functions within a diatonic context. J Exp Psychol Hum Percept Perform, 5(4), 579-594. doi: 10.1037/0096-1523.5.4.579

Loui, P., Wessel, D. L., & Hudson Kam, C. L. (2010). Humans Rapidly Learn Grammatical Structure in a New Musical Scale. Music Perception, 27(5), 377-388.

Margulis, E. H., & Beatty, A. P. (2008). Musical style, psychoaesthetics, and prospects for entropy as an analytic tool. Computer Music Journal, 32(4), 64.

Narmour, E. (1990). The Analysis and Cognition of Basic Melodic Structures: The Implication-Realization Model. Chicago: University of Chicago Press.

Pearce, M., Ruiz, M. H., Kapasi, S., Wiggins, G., & Bhattacharya, J. (2010). Unsupervised statistical learning underpins computational, behavioural and neural manifestations of musical expectations. Neuroimage, 50, 302-313.

Pearce, M. T. (2005). The construction and evaluation of statistical models of melodic structure in music perception and composition. (Doctor of Philosophy), City University, London, UK.

Perruchet, P., & Pacton, S. (2006). Implicit learning and statistical learning: one phenomenon, two approaches. Trends in Cognitive Sciences, 10(5), 233-238. doi: 10.1016/j.tics.2006.03.006

Rohrmeier, M. & Cross, I. (2009). Tacit tonality: Implicit learning of context-free harmonic structure. In Louhivuori, J., Eerola, T., Saarikallio, S., Himberg, T., and Eerola, P.S. (Eds.), Proceedings of the 7th Triennial Conference of European Society for the Cognitive Sciences of Music (pp. 443-452). Jyväskylä, Finland.

Saffran, J. R. (2003). Statistical Language Learning. Curr Dir Psychol Sci, 12(4), 110-114.

Shannon, C. E. (1948). A mathematical theory of communication. Bell Systems Technical Journal, 27, 379-423, 623-656.

Temperley, D. (1999). What's key for key? The Krumhansl-Schmuckler key-finding algorithm reconsidered. Music Perception, 17(1), 65.

Tillmann.B, & McAdams.S. (2004). Implicit Learning of Musical Timbre Sequences: Statistical Regularities Confronted With Acoustical (Dis)Similarities. J Exp Psychol Learn Mem Cogn, 30(5), 1131–1142.