

Self-redundancy in Music

Chris Thornton
COGS/Informatics
University of Sussex
Brighton
BN1 9QH
UK
c.thornton@sussex.ac.uk

September 11, 2009

Abstract

Where a structural analysis can be produced for a musical artefact, variants of the artefact can often be obtained by ‘inverting’ the analysis, in much the same way we produce novel sentences from a grammar. The paper describes use of information theory for purposes of deriving structural analyses of sequences, and shows how the method can be used with musical data, for purposes of generating novel musical patterns.

Keywords: information theory, music, representation

1 Introduction

The paper introduces use of information theory (Shannon, 1948; Shannon and Weaver, 1949) for purposes of obtaining analyses in the form of self-redundancy structures. It shows how these can be used for generative purposes. Where music can be put into a purely sequential representation, the method can be used to generate variants of existing compositions. Some sample outputs are examined. A number of other pieces produced using the same method can be played from the web page at www.chrisThornton.eu/replex-music.html.

The paper divides into four sections. Section 2 contains relatively technical material relating to the method used. Sections 3 and 4 examine the results obtained. Section 5 offers some concluding comments. It is possible to skip Section 2 if necessary.

2 Informational uncertainty

The method uses the uncertainty measure from information theory (Shannon, 1948; Shannon and Weaver, 1949). This quantifies ‘degree of choice’. Given

some set of outcomes, uncertainty is defined to be the entropy of the probabilities applied:

$$\text{entropy} = - \sum_i P_i \log P_i$$

where P_i is the probability of the i 'th outcome. Entropy increases with both the number of outcomes and with the flatness (i.e., indifference) of the probabilities, in accordance with intuition about how uncertainty behaves (Shannon, 1948; Shannon and Weaver, 1949).

Regarding the possible outcomes of sun or rain, for instance, if the probabilities attributed are $P(\text{sun}) = 0.7$ and $P(\text{rain}) = 0.3$ the level of uncertainty is

$$0.7 \log 0.7 + 0.3 \log 0.3 = 0.88$$

With logs taken to base 2 (as they are throughout the paper), this value expresses the level of uncertainty in 'bits'.

Increasing the flatness of the distribution, we might have $P(\text{sun}) = 0.6$ and $P(\text{rain}) = 0.4$. This has the effect of increasing uncertainty. The new entropy value is then correspondingly higher, at

$$0.6 \log 0.6 + 0.4 \log 0.4 = 0.97$$

Where something serves to eliminate some element of uncertainty, the information content of that thing is defined to be the amount of uncertainty eliminated, regardless of what the 'thing' actually is. Should evidence of window drips serve to change probabilities to $P(\text{rain}) = 0.8$ and $P(\text{sun}) = 0.2$, uncertainty drops to 0.72. The information content of the evidence (in this context) is then $0.88 - 0.72 = 0.16$ bits.

2.1 Self-redundancy of sequences

Key to the proposed method is the concept of *self-redundancy*. For present purposes, the self-redundancy of a sequence is defined to be the maximum enhancement of mean symbol content that can be achieved through addition of uncertainty to the sequence. In general, addition of uncertainty has the effect of reducing symbol content. However, there is the possibility of symbols being eliminated altogether, which may produce an increase. Addition of uncertainty may cause mean symbol content to rise or fall then.

Consider the letter sequence 'a b c d b c a'. Imagine that we replace the third element with a choice between 'c' or 'a'. The uncertainty for an equiprobable, two-way choice is $\log 2 = 1$ bit. If each of the original elements of the sequence contained 3 bits of information, we had $7 \times 3 = 21$ bits to start with. Addition of the uncertainty relating to the third element — *uncertainization* as it will be termed — produces a loss of 1.0 bit of information. Mean symbol information (MSI) then falls to $(21-1)/7 = 2.9$.

Uncertainty can also be added by creating a choice between subsequences. Consider the choice between the 2-element subsequence starting at position two and the 2-element subsequence starting at position five. Both of these are ‘b c’ so the choice has only one outcome: there is zero uncertainty. Say we eliminate the third and sixth elements, and make the second and fourth elements be this ‘choice’. Counting those involved in the choice, we have a total of seven symbols — the original figure. The MSI then stays constant at $21/7 = 3.0$ bits.

In general, there are many ways of adding uncertainty to a sequence, including cumulative regimes where one set of choices is superimposed over another. Figure 1 illustrates a range of four cases (labeled A-D) using the sequence from above. Figure 2 provides a further four cases (labeled E-H). The general idea in these schematics is that uncertainizations of the original sequence are displayed as reformulations on a bottom-to-top, hierarchical basis, with arcs indicating the choices introduced.

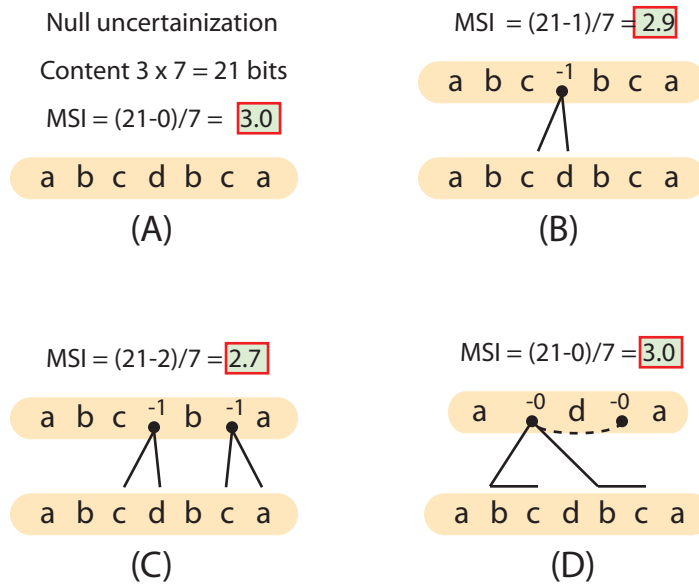


Figure 1: Illustrative uncertainizations of ‘a b c d b c a’.

Case (A) from Figure 1 represents the ‘null’ uncertainization: no change is made to the original sequence. On the assumption that each element is a choice between eight letters, the original (baseline) MSI is $\log 8 = 3$ bits. Case (B) illustrates a minimal uncertainization in which the fourth element becomes a choice between ‘c’ and ‘d’. This produces a loss of 1 bit of information, causing MSI to fall to 2.9 bits. Case (C) is similar but involving uncertainization of two elements of the sequence, with the loss of 2 bits of information.

Case (D) is a more atypical. Here, the second and fourth elements become a choice between the 2-element subsequence commencing at position two and

the corresponding subsequence commencing at position five. Because these are identical, we have zero uncertainty and therefore zero information loss. We also have elimination of some symbols but, overall, the symbol count remains static. Counting in the elements of the subsequence, we have seven symbols in total, producing a MSI of 3.0 bits (the baseline figure).

Turning now to Figure 2, case (E) illustrates cumulative uncertainization. Initially, we replace the elements at positions four and seven with a choice between ‘d’ and ‘a’. We then create a choice between the 3-element subsequence starting at position two, and the 3-element subsequence starting at position five, introducing this as the second and third element of the subsequence. The choice between ‘d’ and ‘a’ produces an information loss of 1.0 bit. But this is incurred in two places. The overall loss is then 2.0 bits. But we now have only six symbols in use so MSI rises to 3.2. Here we have *enhancement* of symbol content, providing evidence of self-redundancy.

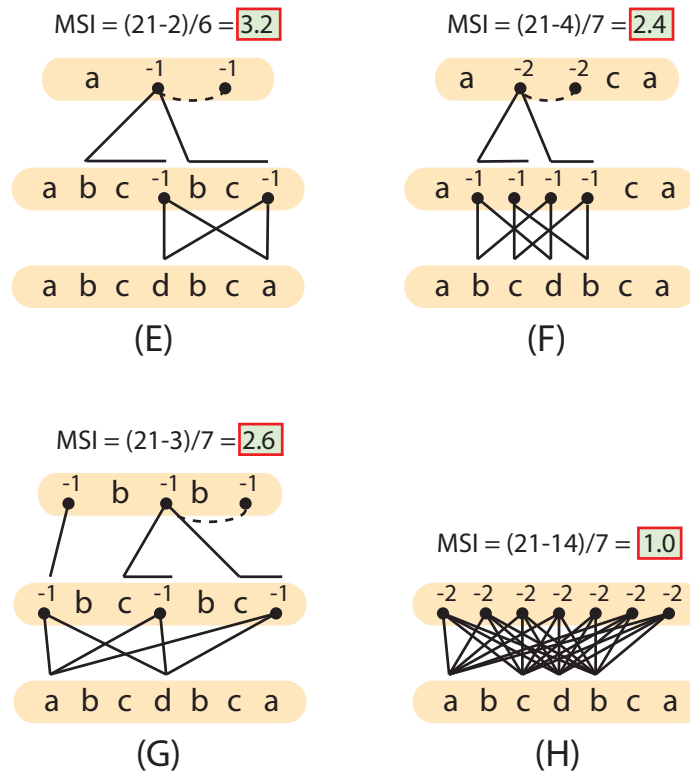


Figure 2: More uncertainizations of ‘a b c d b c a’.

Cases (F) and (G) illustrate the way in which cumulative uncertainization may have the effect of reducing MSI below the baseline. In case (F), four elements of the sequence are replaced with two-way choices, enabling a choice

between 2-element subsequences, effecting a loss of 2.0 bits of information. The general effect is loss of 4.0 bits of information but without there being any change in the number of symbols. The MSI falls to 2.4 bits. Case (G) is similar, although in this case only 3.0 bits of information are lost, producing a mean content of 2.6 bits.

Case (H) represents a rather extreme case of information loss. All seven elements of the sequence are replaced with four-way choices, producing a combined information loss of 14.0 bits. With seven symbols in use, mean content falls well below the baseline, to just 1.0 bit.

Self-redundancy has been defined to be the maximum enhancement of mean symbol content that can be achieved through uncertainization. In effect, it measures how much informational capacity is saved by moving to an optimal representation, and it is in this sense that it measures ‘redundancy’. More intuitively, it can be seen as measuring the savings that can be made by introducing ‘generalizations’ (in the form of uncertainizations) for subsequences showing identical or nearly identical repeats. This interpretation helps to make sense of case (E) particularly. Here, the second level implicitly defines a generalization covering patterns of the form ‘b c ?’, where ‘?’ can be either ‘d’ or ‘a’. The enhancement of MSI that is then achieved can be viewed as resulting from the way in which this generalization captures the relevant pattern of repeats.

2.2 Concentrative uncertainization

The general principle governing enhancement of mean content (and thus existence of self-redundancy) may now be apparent. Where the sequence exhibits




Figure 3: CID curve showing point of maximum information concentration.

repeated or nearly repeated patterns, there is the possibility for uncertainization to exploit the regularity through creation of uncertainized elements that generalize over the relevant selections. The more effectively this is done, the greater the mean content. Exploitation of a particular pattern may require a particular degree of information loss, however. Thus, we expect that the maximum achievable MSI will vary depending on the amount of loss which is introduced.

This leads to the concept of *concentrative uncertainization*, which is uncertainization carried out so as to always maintain maximum MSI subject to loss. This regime takes maximum advantage of any potential generalizations. Combinations of loss and MSI produced for particular data can be plotted out, as in Figure 3. The globally optimal value of MSI can then be identified. This identifies the uncertainization providing maximum concentration of information within symbols.

Referring back to the uncertainizations of Figure 1, we can now ask which of these are genuinely *concentrative*? In most cases, the issue cannot be decided without enumerating *all* the possibilities. However, case (D) is self-evidently optimal. Positive information concentration with zero loss can only be obtained by exploiting explicit repeats in the original sequence. Given that there is only one such repeat, we know that the uncertainization in question must yield optimal MSI. It is also possible to make a comparative judgement relating to cases (C) and (E). These show the same level of loss (2.0 bits). Since the former has higher MSI, we can infer that it must be closer to optimality (for this level of loss) than the latter.

2.3 Concentrative recursive uncertainization (CRU)

The general goal of the present approach is use of information theory for derivation of structural analyses of sequences. This can be achieved, now, by deploying




Figure 4: Self-redundancy structure produced by recursive uncertainization.

concentrative uncertainization in a *recursive* regime. Consider the case where the process is used to discover the point of maximum information concentration (i.e., globally optimal MSI). What is then produced is a version of the original sequence in which some elements are replaced with choices. By definition, the arrangement of these must optimally capture the self-redundancies (i.e., repeat effects) in the sequence. Treating the choices as symbols, we then obtain a symbolic *refinement* of the original sequence that abstracts away the informationally dominant repeat effects.

The uncertainization method can then be applied recursively to the refinement that it, itself generates. By reprocessing refinements recursively, we then ‘peel off’ self-redundancy at successive levels of description. Application of the process to the original sequence captures structural aspects relating to basic repeats. Application to the refinement captures repeats involving constructs with sub-structure. Each subsequent application captures repeats at a higher level of organization. The overall effect is then to capture the self-redundancy *structure* of the sequence. The general effect is illustrated in Figure 4.

3 Self-redundancy structures in musical sequences

It is now possible to return to the theme of music. Figure 5 shows the self-redundancy structure that concentrative recursive uncertainization (CRU) produces for the initial sixteen notes in the melody of the British national anthem. Sequence elements are shown as ovals, here, with the layers corresponding to generated refinements. Strings such as ‘\$0’, ‘\$1’ etc. are symbols standing for particular choices, with definitions being provided in the leftmost instance. Arcs are used to indicate the constitution of choices in terms of individual elements or subsequences. Negative numbers appearing in the oval shapes are information losses associated with corresponding choices. Values of MSI are marginal values, calculated on the basis of the immediately prior encoding. All symbols are assumed to cost 4.0 bits of information.

At the bottom of the figure, the sixteen notes of the melody form the primitive nodes of the hierarchy. At the next level, the informationally optimal uncertainization of those data is represented first as a set of nodes and then — to the right — as a linear string. Adjacent to this, we see calculations of aggregate content and MSI. At subsequent levels of the hierarchy, the same general arrangement is used.

Note the extensive use of the pattern ‘c d b/c/e’ in the first encoding level. This results from the fact that the relevant pattern is repeated at three locations in the original sequence (positions 2, 5 and 12). Using the symbol ‘\$0’ to represent the corresponding choice, four symbols are eliminated for 1.58×3 bits of information loss. MSI rises to 4.93 bits, in excess of the baseline figure of 4 bits. A similar effect involves ‘\$0 \$0/\$3’ in the second level refinement, and ‘\$5/c \$4’ in the third level encoding.

For comparison, Figure 6 shows a detail from the self-redundancy structure for ‘La Marseillaise’, the French national anthem. In this case, the original

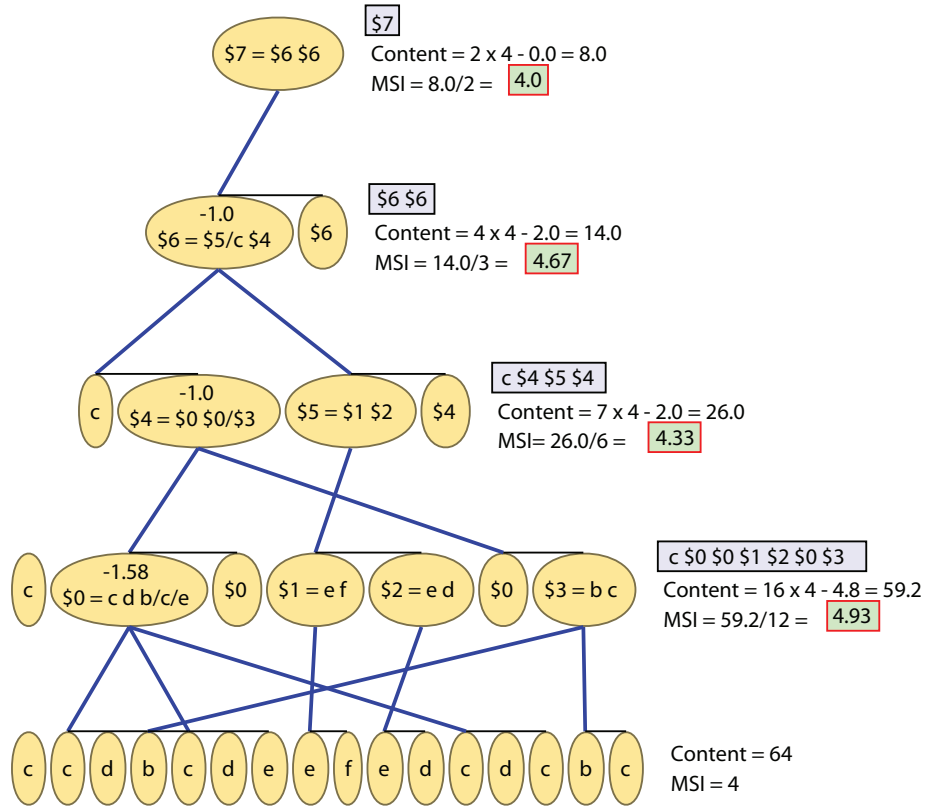


Figure 5: Self-redundancy structure of the first 16 notes from ‘God Save the Queen’ melody.

sequence is formed of strings providing an offset representation for individual notes and the hierarchy is substantially truncated. The note representation scheme used here has the form o:p:d:v, where o is the beat offset of the note from the previous note, p is the pitch offset (represented as a letter), d is a numeric duration (preceded by ‘-’ where the value is in MIDI ticks) and v is an intensity value (preceded by ‘v’). This is also the representation used in the generative experiments described below.

4 Generating musical replexes

Using concentrative uncertainization, it is possible to obtain structures capturing self-redundancies in a sequence over multiple levels. Such structures are symbol hierarchies and, like phrase-structure grammars, can be used for generative purposes. The essence of the procedure is recursive symbol-expansion.

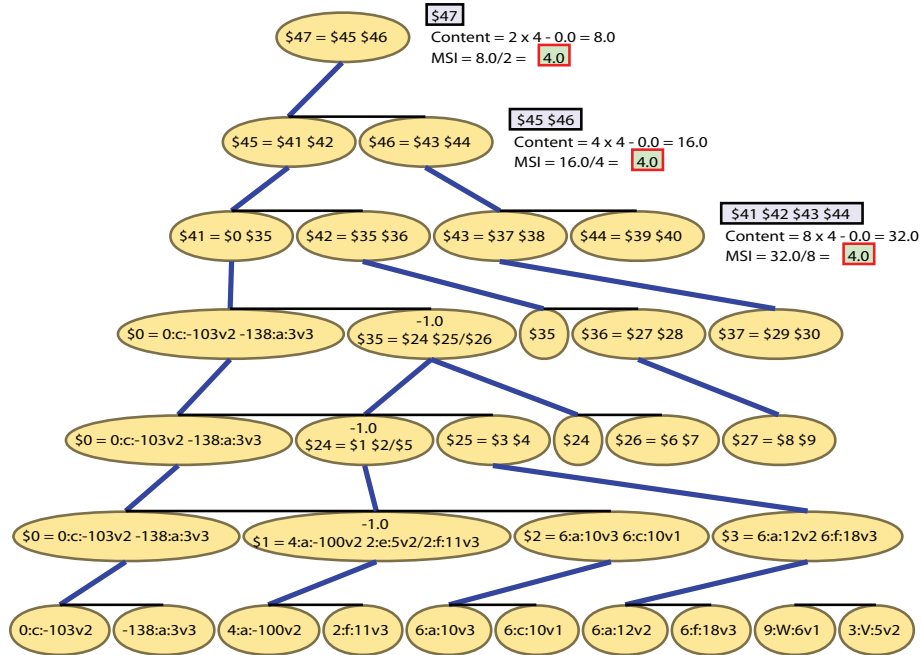


Figure 6: Upper self-redundancy structure of La Marseillaise melody.

Starting with the root symbol, symbols are expanded recursively, with random choices being made at any disjunctive branch. When no further symbol expansion is possible, the result must be a new sequence with the same self-redundancy structure as the original. Sequences constructed in this way are termed ‘replexes’ — literally ‘refoldings’.

As an illustration, four replxes of the sequence from Figure 4 are

```

c c d b c d e e f e d c d e c d e
e f e d c d e c d e e f e d c d c b c
e f e d c d e c d c e f e d c d e c d e
c c d b b c e f e d c d c c d c

```

Further examples can be generated using the applet at www.chrisThornton.eu/replex-music.html. (Access the ‘queen1’ input set from the applet’s main menu and then press ‘Analyze’ followed by ‘Make’.)

Experiments have also been carried out involving other musical sources. These have used data from MIDI files put into the offset representation described above. A score representation for part of a replx of Mozart’s Sonata in C Minor appears in Figure 7.



Figure 7: Extract from a replex of Mozart's Sonata in C Minor, 1st Movement.

4.1 Hyper-replexes

In the simplest case, obtaining generative music from uncertainization involves using self-redundancies from a single composition. Any entity generated from this can then be precisely defined: it is one of the original composition's replexes. However, it is also possible to work with structures amalgamated from more than one composition. The resulting artefact then combines (or 'crosses') informational properties from different sources. Such entities are termed 'hyper-replexes'.

One way to use this idea involves crossing the self-redundancy structures from different compositions. For example, Figure 8 shows a score representation for an extract of a replex crossing Bach's prelude No 3 (from WTC book 1) with Philip Glass's Koyaanisqatsi Suite. (Part of this piece can be played from the website.) It is also possible to amalgamate structures taken from different modalities, however. For example, we can amalgamate structures taken from textual sequences with ones derived from musical sequences. This opens up the possibility of generating musical/textual hybrids. Some examples can also be played from the website.



Figure 8: Extract from a hyper-replex crossing Bach's Prelude No. 3 with Philip Glass's Koyaanisqatsi Suite.

5 Summary

The paper has introduced the concept of self-redundancy and shown how concentrative uncertainization can be used to obtain refinement structures from sequences. Using recursive symbol expansion, such structures can be the means of generating novel sequences that reproduce the self-redundancies of an original artefact. Where it is possible to put a musical composition into a purely sequential representation, the method can be used to generate a replex of that composition, i.e., a variant with the same structure of self-redundancy. There is also the possibility of using self-redundancy structures taken from more than one source, even where those sources are in different modalities. Reproductions obtained from such amalgamated structures are ‘hyper-replexes’, with the potential to hybridize music with text.

References

- Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27 (pp. 379-423 and 623-656).
- Shannon, C. and Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, Illinois: University of Illinois Press.