Pierre Alexandre Tremblay

What is noise (music) to you?

If we take the musical act as expression, therefore communication, then signal theory forces us to accept every sound as music as long as it is intended to be. What is not intended is noise, even in the form of a Bach partita played next door when I want to listen to Bernhard Günter. Music as art is a statement of intent, and everything (and only what) the artist wants to produce is the work. There is only noise in non-intent.

That definition of music versus noise, in line with Varèse's "organised sound," has a contradicting sibling in my practice: I like to embrace uncertainty, the unexpected, and instability fascinates me as it reminds me of living things. Only human thoughts and models can dream up abstract perfection, and I prefer the filthy, gritty, humid, sticky unevenness of nature. So music, as human construct, contrasts with noise as the undesired artefact of reality. In that respect, we could take the harmonicity of a signal to declare it pure, and inharmonicity to be noisy … quite a boring scale, but quite useful, and definitely in phase with the ideals of consistency in instrument making and performance practice of the Western world. This scale has its limits though, as there are perceptual limits to saturation. One is *Gestalt* grouping, when accumulation becomes too dense; another is that we perceive in contrasts and get used to anything that is constant, even information overload. Therefore I like to think of this axis as being curved, with the maximum noisiness at the peak of perceptual saturation. In this respect, the context of listening becomes essential in the definition of noise itself.

I think we are in a post-noise period. Post-noise, as in post-glitch. In the latter, what was a mistake of digital sound devices was embraced as art,

10.5920/noise.05

first with its conceptual random value, then the sound of glitch became an object itself. Comparably, we have now heard every sound, loud and soft, saturated and pure, dense and light, with classic noise artists mostly embracing the former in each of these dualities. This soundworld has now permeated all music practices. Our soundscapes are lower-fi than ever, the sound design in multimedia is getting more daring by the day, chart-topping pop music distorts like never before, and the loudness war has reached its theoretical limits. Moreover, as individuals we are saturated with information as never before and with sheer violence at levels that are unattainable in art—radical art is *passé*: it looks quite futile compared to the daily news. More interestingly, the grey zone between these extremes is now fully assumed, full of rich crossovers and hybrids: in-between-ness reigns. This post-noise era is rich with the full breadth of the sonic world.

Maybe the desire of artists to use the dirty part of their soundworld has been there forever; they were just surrounded by cleaner sounds. With a louder world, with omnipresent background music, with piercing sirens on the street, the threshold of what is acceptable as music might have just been pushed. Here is an interesting hint that we might have reached a limit: after the re-appropriation of the soundspace by the individual, with the headphones/Walkman revolution of the 1980s, the trend of active silencing headphones points to something quite clear: noise is in the ear of

Why do you make it?

Am I making noise? Not to my ears, but I certainly embrace all levels of volume, contrast, purity, transience, density, and instability in the listening experience. This might be noise to some, but to me it is music.

I like art to be in phase with the world, and to embrace the rich multitude of different human experiences from an embodied perspective, even if this can be a little overwhelming in its contrasts. To embrace my humanity, be it simple pleasure or existential turmoil, music is the only way to express it: art is for my gut and my soul.

 10.5920/noise.05

Noise Music Information Retrieval

Nick Collins

Noise music foregrounds processes that other music might seek to minimize or avoid, from high levels of distortion to problematic and polemical subject matter.[1] The focus here is on the typical "noisy" manifestation of extreme sensory dissonance in the audio signal itself, and tracking this through automatic analysis techniques from the field of Music Information Retrieval (MIR).[2] Such MIR methods have strong applications in computational musicology when working with one or more audio files. For example, Collins[3] applies MIR tools to a study of a corpus of synth pop; Klein et al.[4] consider using MIR methods for the analysis of acousmatic electroacoustic music; Tsatsishvili[5] attempts to differentiate the overlapping sub-genres of heavy metal; and Mital and Grierson[6] explore an archive of works of Daphne Oram through a visualization method.

1 Paul Hegarty, *Noise/Music: A History* (New York: Continuum, 2007); Thomas Bey William Bailey, *Microbionic: Radical Electronic Music and Sound Art in the 21st Century* (London: Creation Books, 2009); Nick Collins, Margaret Schedel, and Scott Wilson, *Electronic Music* (Cambridge: Cambridge University Press, 2013).

2 Michael A. Casey, Remco Veltkamp, Masataka Goto, Marc Leman, Christophe Rhodes, and Malcolm Slaney, "Content-based music information retrieval: Current directions and future challenges," *Proceedings of the IEEE* 96 no.4 (2008): 668–96.

3 Nick Collins, "Computational Analysis of Musical Influence: A Musicological Case Study Using MIR Tools," *Proceedings of the International Society for Music Information Retrieval Conference, Utrecht* (2010): 177–182.

4 Volkmar Klien, Thomas Grill, and Arthur Flexer, "On Automated Annotation of Acousmatic Music," *Journal of New Music Research* 41 no.2 (2012): 153–73.

5 Valeri Tsatsishvili, "Automatic Subgenre classification of heavy metal music" (Master's Thesis, University Of Jyväskylä, 2011).

6 Parag K. Mital and Michael Grierson, "Mining Unlabeled Electronic Music Databases through 3D Interactive Visualization of Latent Component Relationships," paper presented at New Interfaces for Musical Expression (NIME), Daejeon, South Korea, 2013.

10.5920/noise.05

In this project, MIR-informed analysis is explicitly applied to noise music to assess the structure of individual recordings and to compare multiple recordings, providing new insights into the sonic content of noise. The specific targets of this study include two Merzbow albums, *Oersted* (1996) and *Space Metalizer* (1997), for which individual tracks are analyzed, and the pieces across the two albums compared. I also investigate the application of MIR analysis to a corpus of historic noise music, including Whitehouse, Masonna, and Xenakis, placing the Merzbow works in a wider context. Although there is a broad potential to this technology, which may extend beyond musicology to new compositional directions, there are also challenges. Questions remain of how to validate the results of automated analysis against human reaction, and a critical view of MIR should be maintained as we proceed. Nonetheless, the study is an essential step to approaching and evaluating MIR applications.

Noise MIR analysis techniques

MIR typically operates with respect to a corpus of audio files, though it may also examine properties of a single file in isolation. Rather than working with raw sample data, a system will extract derived features, such as energy in chroma (typically following pitch classes representing the semitonal twelve-note equal temperament typical of Western music) or Mel-Frequency Cepstral Coefficients (a measure of spectral energy distribution with good correspondence to timbre). As well as having a lower sampling rate than audio samples, making the amount of data more manageable, these features are chosen because they ideally represent more salient auditory and musical attributes of the work under study. Often in MIR, a time-varying feature is reduced further, perhaps to an average across an entire piece; this can still help to plot locations of different pieces with respect to another. Nevertheless, retaining the dynamical progress (the time series) of features within given pieces can capture musical behavior in more detail and is intuitively closer to a human-like perception over time. Having obtained features—essentially,

10.5920/noise.05

numerical summaries—for some set of files, machine learning algorithms can then be applied over the corpus to examine such questions as how all the files (music) cluster together, how to discriminate different "types" of file (music), and so forth.

This project favors such features as spectral entropy, sensory dissonance, perceptual loudness, "transientness," and spectral centroid as timbral aspects of high relevance to the perception of noise music. Some of these features, such as sensory dissonance and perceptual loudness, depend on models of human auditory perception; others are related to studies of instrumental timbre (correlating for the spectral centroid to "brightness" of tone), or are information theoretic signal processing constructions (spectral entropy is a measure of the information gained in momentary spectral change). I will use both summary features such as a mean (an average) across a whole piece, and time-varying features that describe the course of a piece, as explained further below.

Once features have been extracted, similarity measurements can compare within-piece or between-piece relations, examining the internal structure of a work, or the proximity of different audio files representing different opuses from the same or diverse composers. Similarity matrices can help assess within-piece formal relationships, including the detection of change points through a derived novelty curve. Models can be formed for individual pieces with respect to their time series, for instance, via k-means clustering from the feature vector space to cluster labels, followed by variable order Markov modeling on the symbolic sequences created. Feature statistics, or time series models, can then form the basis of comparison between pieces in corpus analysis.[7] All calculations in this study use the open source SCMIR library for SuperCollider.[8]

7 It is beyond the scope of this chapter to cover MIR techniques in general. For an introduction to content analysis, see Casey et al. "Content-based music information retrieval."

8 Nick Collins, "SCMIR: A SuperCollider Music Information Retrieval Library", *Proceedings of the International Computer Music Conference, Huddersfield* (2011): 499–502.

As an example of feature extraction, Merzbow's "cover" of *Silent Night* from a noise music Christmas compilation[9] was subjected to analysis. "Cover" is apt, as Merzbow gradually covers a strained rendition of the carol in layers of noise until the source is completely obscured. With this increase in obscuration under a wall of noise, the track provides a useful test for whether computational feature extraction discovers the progression into brutal noise clearly audible by a human listener.

Figure 1 plots how feature values vary as the initial carol submerges beneath the distortion. In part, this plot provides a validation of the types of features being extracted; they all generally show an increase that follows the highly perceptible increase in distortion as the piece progresses, though moment to moment variation is also evident. (The plotted values are maximal values for each second of audio; the FFTCrest feature, which looks at spectral "peakiness," is inverted so that the flatter spectrum of the noisier portions is evident.) There are short silences at the beginning and end of the track that disrupt feature collection a little (the SpectralEntropy in particular is maximal in the face of silence, explaining its later relatively low values, though the upwards trend is still evident).

It is also plausible to detect abrupt transitions in the time varying feature values to seek out potential sectional boundaries within a work. This may be done with respect to a single feature, or over multiple features simultaneously, for example via structural detection methods such as convolution with a checkerboard kernel along the diagonal of a similarity matrix. A 'findSections' command in SCMIR pointed to one main transition point, at 92 seconds in; this turned out to correspond to a sudden increase in the distortion and brightness of the guitar. While the routine did not reveal any additional points of change, the other layers tend to come in more slowly, or to reflect further levels of distortion that the feature detection may not fully track.

9 *Various–The Christmas Album* (Sony Records, SRCL 3723, 1996), accessed July 29, 2013, http://www.discogs.com/Various-The-Christmas-Album/release/1331008.

 10.5920/noise.05

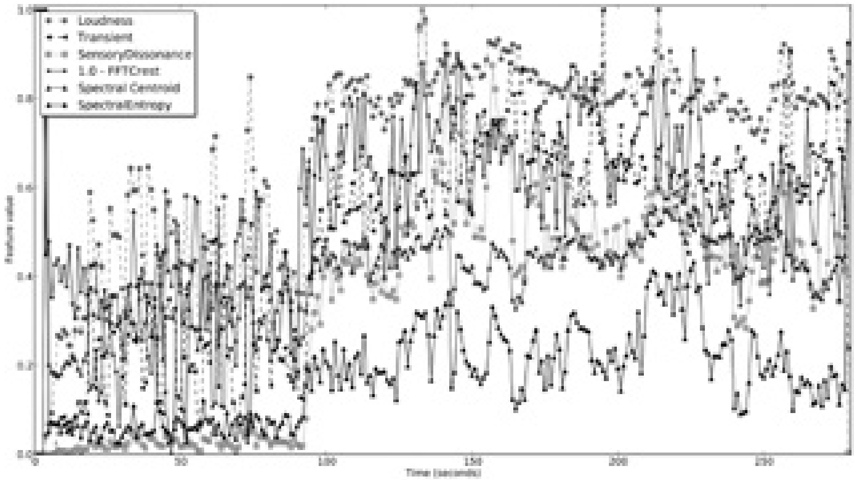

Figure 1: Feature trails over the four and a half minutes of Merzbow's *Silent Night* (1996)

The change point is also evident from the visual tipping point of the feature curves in Figure 1. I do not pursue change points further in this article, but the technique illustrates the further options for structural analysis opened up by MIR.

Merzbow vs. Merzbow

To illustrate application across multiple files, this section explores the similarity relations among the tracks over two Merzbow albums, *Oersted* (1996, 4 tracks titled by their durations) and *Space Metalizer* (1997, five tracks). Products of a period in Merzbow's 1990s output associated with harsh noise, the sound sources include various classic analog synthesizers, such as the EMS Synthi A, and metal percussion, running through effects units such as guitar pedals and filters alongside live tape manipulation. There is a feeling of spontaneous performance across the tracks, with mercurial shifts in timbre as attention wanders across the equipment at the artist's disposal. *Oersted* is perhaps a little starker, yet both albums have free time moments as well as some tight repeating loops evident at points. Both push hard, with hot mastering.

Machine analysis initially extracted six features: perceptual loudness, transientness (a measure of sudden signal change, for example caused by percussive onsets), sensory dissonance, a spectral crest measure (the "peakiness" of the spectrum), spectral centroid, and spectral entropy. All features were minimum–maximum normalized to the range 0.0 to 1.0 with respect to a globally derived maximum and minimum over all files. Means and maximums were taken across complete files as summary features, as well as one-second window means and maximums as time-varying features.

As a sanity check on the feature extraction, it is clear from the feature data that *Oersted* is louder and has higher sensory dissonance (according to the computational models) than *Space Metalizer*, somewhat confirmed by listening (all the tracks are mastered hot and general loudness is persistent, though with moment to moment fluctuation). The greatest maximum sensory dissonance of the tracks was for *7:53* on *Oersted*, the most unrelenting in its wideband noise; the lowest was "Son of Zechen" from *Space Metalizer*, which is slightly gentler (relatively!) on listening. It is clear from examination of the feature trails (time-varying one-second means and maximums) that features are often (though not inevitably) correlated; for example, loudness, transientness, and sensory dissonance all react to the sheer "energy" (in a non-technical sense) of the music.

The tracks were compared in two ways. The first simply looked at the proximity of the six mean feature values, a cruder measure. The second used time series methods, to be described below, to form a model of each piece; once a model was formed, it could be used to predict how surprising other pieces appeared with respect to that prior knowledge, so as to measure similarity.

The first mean-based method created the similarity matrix in Figure 2. Distances between the six-dimensional feature vectors (the means over a given piece for each of six features) were calculated via the Euclidean metric, and were then normalized across all distances to the range 0–1 for ease of

 10.5920/noise.05

reading (numbers are taken to 3 decimal places). This matrix is symmetric, with zeros on the diagonal, as a piece is always no distance from itself. The rows and columns allow a measure of similarity to be read between any two tracks across the two albums. (The two albums are denoted by prefixes OE and SM.)

	OE1	OE2	OE3	OE4	SM1	SM2	SM3	SM4	SM5
OE1	0	0.383	0.245	0.047	0.131	0.212	0.102	0.191	0.101
OE2	0.383	0	0.021	0.284	0.478	1	0.321	0.09	0.247
OE3	0.245	0.021	0	0.185	0.36	0.812	0.216	0.042	0.161
OE4	0.047	0.284	0.185	0	0.045	0.266	0.028	0.103	0.055
SM1	0.131	0.478	0.36	0.045	0	0.266	0.026	0.195	0.096
SM2	0.212	1	0.812	0.266	0.266	0	0.379	0.694	0.432
SM3	0.102	0.321	0.216	0.028	0.026	0.379	0	0.086	0.04
SM4	0.191	0.09	0.042	0.103	0.195	0.694	0.086	0	0.08
SM5	0.101	0.247	0.161	0.055	0.096	0.432	0.04	0.08	0

Figure 2: Similarity matrix over two Merzbow albums from mean feature vectors

There are overlaps between the two albums, given their close gestation in time within Merzbow's career. SM2, "Son of Zechen," is the most atypical track here, seen as particularly different to the central tracks on *Oersted*, though of equal distance to OE4 and SM1. It is closest to the first track of *Oersted*, which may be related to the use of a low, throbbing, bassy figuration in both tracks. The matrix may point to greater variety on *Space Metalizer* than on *Oersted*, in the sense of the degree to which tracks within the album are dissimilar to one another. Yet if *Oersted* is a little more homogenous, this seems predominantly due to SM2 on the other album. With "Son of Zechen" excluded, *Oersted* may be the more heterogeneous; for instance, the bottom-right 3 by 3 submatrix shows the close proximity of SM3 through SM5 within *Space Metalizer*, and the low scores on the SM1 row compliment this picture.

In the second method of assessing similarity, machine analysis proceeded as follows:

1. Form one-second windows of features, taking the mean in each window.
2. Vector quantize; k-means clustering is applied with 20 cluster centers over all mean feature vectors (from all seconds of all pieces) in the 6-dimensional feature space. All (continuous) mean feature vector sequences are mapped to (discrete) integer sequences.
3. Train prediction by partial-match variable-order Markov models for each audio file in the corpus based on the integer sequences.[10]
4. For each pair of files, calculate similarity based on the symmetric cross-likelihood.[11] The model from piece A is used to predict the unexpectedness of piece B, and from B to predict A; the result is a combined sense of how well A and B predict each other and, thus, their similarity.

The result shown in Figure 3 is a similarity matrix between pieces that can potentially illuminate the proximity of the musical thinking of different tracks, as it respects the time variation within those tracks far better than a gross average.

In general outline, the two matrices constructed via two different methods show similar inter-relations between tracks, which may give some confidence to both methods' applicability to noise music. In the second matrix, SM1 and OE3 are the most dissimilar, though SM2 and OE2 (the furthest apart in the

---

10 Nick Collins, "Influence in Early Electronic Dance Music: An Audio Content Analysis Investigation", *Proceedings of the International Society for Music Information Retrieval Conference, Porto* (2012): 1–6 ; Marcus Pearce and Geraint Wiggins, "Improved methods for statistical modelling of monophonic music," *Journal of New Music Research 33*, no.4 (2004): 367–85.

11 Tuomas Virtanen and Marko Helén, "Probabilistic model based similarity measures for audio query-by-example," *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics,* New York (2007): 82–85.

10.5920/noise.05

	OE1	OE2	OE3	OE4	SM1	SM2	SM3	SM4	SM5
OE1	0	0.758	0.748	0.703	0.843	0.615	0.724	0.744	0.721
OE2	0.758	0	0.596	0.615	0.939	0.835	0.76	0.671	0.692
OE3	0.748	0.596	0	0.744	1	0.766	0.812	0.682	0.747
OE4	0.703	0.615	0.744	0	0.667	0.589	0.694	0.685	0.659
SM1	0.843	0.939	1	0.667	0	0.696	0.567	0.861	0.612
SM2	0.615	0.835	0.766	0.589	0.696	0	0.539	0.769	0.697
SM3	0.724	0.76	0.812	0.694	0.567	0.539	0	0.718	0.651
SM4	0.744	0.671	0.682	0.685	0.861	0.769	0.718	0	0.584
SM5	0.721	0.692	0.747	0.659	0.612	0.697	0.651	0.584	0

Figure 3: Similarity matrix from predictive models over two Merzbow albums

first matrix) are not conversely claimed as similar. Potentially problematic in the second matrix is the low similarity between the tracks "Space Metalizer Pt. 1" and "Space MetalizerPart 2" (*sic*; this is how the track names are written in the CD liner notes). However, on listening, the difference between Parts 1 and 2 is clear as the tracks proceed; they seem to start at a common point, then diverge into their own noise worlds.

Taking a cue from some of Masami Akita's own preoccupations, one interesting facet of this study in the context of noise music is a measure of "dominance," by which a predictive model trained on one work explains another work better than the other way around. In order to calculate this, instead of the symmetric measure we can use the direct prediction scores (expressed as average logloss, where low numbers denote better predictions). The difference between B predicting A and A predicting B is a measure of the degree to which A dominates B and vice versa. An anti-symmetric matrix is presented here where 1 means that the left row track dominates the column track, -1 is the inverse, and 0 indicates that there is no dominance relation (for the self-model predictions on the diagonal).

	OE1	OE2	OE3	OE4	SM1	SM2	SM3	SM4	SM5
OE1	0	1	1	-1	1	1	1	1	1
OE2	-1	0	-1	-1	1	1	-1	1	1
OE3	-1	1	0	-1	1	1	-1	1	1
OE4	1	1	1	0	1	1	1	1	1
SM1	-1	-1	-1	-1	0	-1	-1	-1	1
SM2	-1	-1	-1	-1	1	0	-1	1	1
SM3	-1	1	1	-1	1	1	0	1	1
SM4	-1	-1	-1	-1	1	-1	-1	0	1
SM5	-1	-1	-1	-1	-1	-1	-1	-1	0

Figure 4: Dominance matrix over two Merzbow albums

According to this measure, track OE4 ("18:49") is the most predictive of other tracks; we might view it as somehow encapsulating a kernel of techniques used throughout the two albums. It is probably the most timbrally varied track on *Oersted*, and it also has an aural relation to many moments on *Space Metalizer*. SM5 ("Mirage") is the most derivative by this measure (all -1s on its row). And Part 2 of "Space Metalizer" dominates Part 1. However, these results should be set in the context of the simplification, whereby actual numerical differences have been reduced to one of three options, and are critically dependent on the aural validity of the modeling in the first place.

A small historical corpus of noise music

I now place these two albums by the same artist in the context of music by other noise musicians. The survey I present is by no means exhaustive; Merzbow's career from 1979 is itself replete with many more recordings than are examined here, and the reader will no doubt think of many other examples of noise musicians that she or he might be interested to explore.

One justification for carrying out this study, even with some reservations on the acuity of features raised in the preceding sections, is that it is

 10.5920/noise.05

impossible for an analyst to keep in mind all of the audio material across a large corpus. Automated methods at least objectify the process of hunting for interrelations on a level playing field, rather than the most recently consulted track, or the bias of a musicologist having listened more to particular material. Nonetheless, the choices made in order to establish a corpus and the decisions taken in writing the computer program are themselves a potential source of bias, if more explicitly stated.

Figure 5 lists the various sources in the corpus gathered for this part of the study. These range from electronic music by Iannis Xenakis, through the classic Lou Reed statement of feedback *Metal Machine Music*, to postpunk experimental acts such as Whitehouse who started to foreground blasts of noise, the Japanese noise artist Masonna, and a control case of the Beach Boys. The total audio content is around 10 hours of material over 96 individual files.

I applied the same processing as above, extracting the same six features with respect to the same normalization factors derived from the Merzbow

Group ID	Artist	Works, with Dates	Total Duration (minutes)
0	Merzbow	*Oersted* (1996), *Space Metalizer* (1997)	135.8
1	Xenakis	*Bohor* (1962), *Taurhiphanie* (1987), *Gendy3* (1991), *S.709* (1994)	58.5
2	Lou Reed	*Metal Machine Music* (1975)	64.2
3	Whitehouse	*Birthdeath Experience* (1980), *Great White Death* (1985)	73.2
4	Nurse with Wound	*Chance Meeting on a Dissecting Table of a Sewing Machine and an Umbrella* (1979)	47.9
5	Einstürzende Neubauten	*Strategies Against Architecture, Vol. 1* (1983)	41.3
6	Non	*Easy Listening for Iron Youth: The Best of Non* (1989)	67.5
7	Masonna	*Shock Rock – Track 2* (2002), *Mademoiselle Anne Sanglante Ou Notre Nymphomanie Auréolé* (1993), *Shinsen Na Clitoris Part 1* (1990)	55.4
8	Beach Boys	*Pet Sounds* (1966)	36.4

Figure 5: Noise music corpus

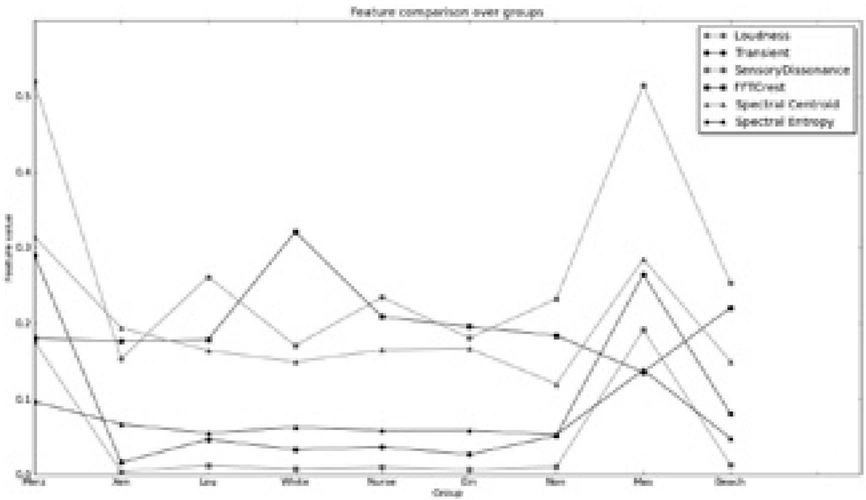

Figure 6: Mean feature values by artist group for the corpus in Figure 5

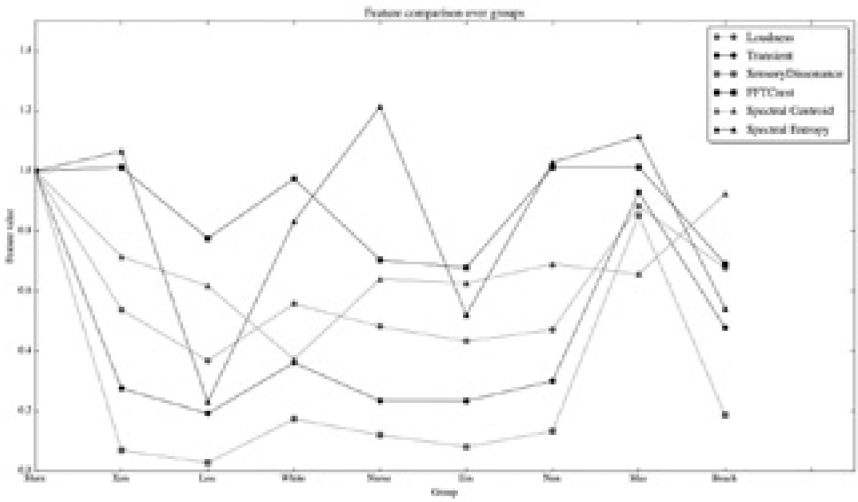

Figure 7: Maximal feature values by artist group for the corpus in Figure 5

10.5920/noise.05

albums; feature values are thus relative to the Merzbow group as the reference. All tracks from a particular artist form a group of feature data, giving nine distinct musical groups and their associated data.

Figures 6 and 7 plot the mean and the maximal feature values for each artist group over the six extracted features. Note that group 0, the Merzbow reference, has every maximum at 1, as it provided the normalization information for the feature extraction. These diagrams clearly show the relationship between the Merzbow and Masonna groups of tracks. Perhaps surprising is that Metal Machine Music does not enter so similar a space; in actual fact, Lou Reed's opus is much more based around resonant ringing pitches of feedback, and has a less broadband noise characteristic and less fierce mastering. Because of the percussion transients, the Beach Boys album takes a closer aspect (with respect to the extracted features) to other noise works than might be expected. The two Japanoise artists are clearly revealed as more aggressive in their soundworld than some earlier noise precedents.

The interrelationship of pieces was assessed via the feature-based similarity measurements described above, first using means and maximums over each group, and then with time series models trained over all the pieces within a given group. While this might also be investigated on a piece-by-piece basis, here it is restricted to artist groupings.

Using only group feature means and a Euclidean metric:

	Merz	Xen	Lou	White	Nurse	Ein	Non	Mas	Beach
Merz	0	0.962	0.669	1	0.748	0.895	0.781	**0.021**	0.652
Xen	0.962	0	0.052	0.088	0.034	**0.008**	0.05	0.916	0.07
Lou	0.669	0.052	0	0.11	**0.007**	0.028	0.011	0.631	0.012
White	1	0.088	0.11	0	0.064	**0.061**	0.09	0.999	0.074
Nurse	0.748	0.034	**0.007**	0.064	0	0.012	0.011	0.716	0.01
Ein	0.895	**0.008**	0.028	0.061	0.012	0	0.021	0.855	0.035
Non	0.781	0.05	**0.011**	0.09	**0.011**	0.021	0	0.736	0.013
Mas	**0.021**	0.916	0.631	0.999	0.716	0.855	0.736	0	0.633
Beach	0.652	0.07	0.012	0.074	**0.01**	0.035	0.013	0.633	0

Figure 8: Similarity matrix over corpus from mean feature vectors

Using models trained on the work of a particular artist, and used to predict other artists:

	Merz	Xen	Lou	White	Nurse	Ein	Non	Mas	Beach
Merz	0	0.927	0.86	0.92	0.925	1	0.787	**0.401**	0.838
Xen	0.927	0	0.323	0.294	0.322	**0.266**	0.303	0.706	0.542
Lou	0.86	0.323	0	0.407	0.372	0.384	**0.213**	0.47	0.376
White	0.92	0.294	0.407	0	0.404	0.486	**0.214**	0.614	0.497
Nurse	0.925	0.322	0.372	0.404	0	**0.215**	0.233	0.627	0.408
Ein	1	0.266	0.384	0.486	**0.215**	0	0.292	0.651	0.502
Non	0.787	0.303	**0.213**	0.214	0.233	0.292	0	0.528	0.349
Mas	**0.401**	0.706	0.47	0.614	0.627	0.651	0.528	0	0.733
Beach	0.838	0.542	0.376	0.497	0.408	0.502	**0.349**	0.733	0

Figure 9: Similarity matrix over corpus from predictive models

10.5920/noise.05

The closest proximity of artists in each row is indicated in bold type. Merzbow and Masonna are paired on both measures, while they are dissimilar to other artists in the corpus. Interestingly, Lou Reed appears much closer to Masonna than Merzbow, although Masonna is still the second most distant from him. Some overlap among experimental acts of the later 1970s and early 1980s is apparent, for instance, in the close relation of Nurse with Wound and Einstürzende Neubauten.

	Merz	Xen	Lou	White	Nurse	Ein	Non	Mas	Beach
Merz	0	1	-1	1	-1	-1	1	1	1
Xen	-1	0	1	-1	-1	-1	-1	-1	-1
Lou	1	-1	0	-1	-1	-1	-1	-1	-1
White	-1	1	1	0	-1	1	1	-1	1
Nurse	1	1	1	1	0	1	1	-1	-1
Ein	1	1	1	-1	-1	0	1	-1	-1
Non	-1	1	1	-1	-1	-1	0	-1	-1
Mas	-1	1	1	1	1	1	1	0	1
Beach	-1	1	1	-1	1	1	1	-1	0

Figure 10: Dominance matrix over corpus

The dominance matrix technique can also be applied over the artist groups and their models, as shown in Figure 10. This dominance table should not be taken too seriously and is included here as much as a provocation to consider the side effects of accepting technological musicology uncritically. Nonetheless, Merzbow dominates Masonna, and Lou Reed dominates Merzbow (but no other artist; the dominance relation is not transitive, as Lou Reed is subservient to Masonna). The Beach Boys dominate many other artists; this is likely an artifact of their (possibly) broader spectral palette and time variation with respect to the more delimited timbral world of some noise music artists.

In summary, the aggressive sound of Japanese noise music from Merzbow and Masonna is clearly indicated. However, the surprising relation of the Beach Boys to some of the other artists represented here gives pause; the combination of features extracted is probably not sufficiently representative. It certainly doesn't differentiate the common practice harmony and rhythm that appears in the Beach Boys' work.

Conclusions

Some analytical applications of MIR techniques, such as feature extraction and similarity measures, have been investigated with respect to noise music. While such procedures may reveal new formal details and new interrelations of pieces within a systematic framework, there remains a need to validate machine listening against human listening. Such analysis will remain a companion and compliment to the human analyst, but its use should be further investigated as MIR audio analysis techniques continue to develop.

I have said less about compositional applications, which are themselves at the mercy of the quality of algorithmic listening. Ideally, audio analytical methods can form the basis for critical systems for algorithmic composition and interactive music systems. For a noise music generating system, analytical techniques can provide the grounding, the listening experience, for a self-critical computer agent, whether founded on a composer's own work, or a database derived from a historical or contemporary corpus. Given the "otherness" of noise music, an exact simulacrum of human listening may not necessarily be an aesthetic requirement, and partially effective auditory modeling may allow for many noisy possibilities.

The challenges for future work are many. There is the problematic nature of many noise musicians' release catalogues, which embrace prodigious

 10.5920/noise.05

release rates over many media as a sort of cultural noise;[12] serious Merzbow or noise music history scholars may be forced to pool resources and form collaborative online databases, with associated issues of copyright. Although a musicologist's own listening is a natural centerpiece of analysis projects, the computer acts as a proxy listener, necessarily as the corpus size increases. Engaging with the vast extent of noise music releases may require computational assistance, and indeed, certain aspects of noise music structure on the boundary of human discernment may benefit from untiring neutral signal processing analysis. However, the extent to which computer extracted features reflect the encultured and embodied human auditory system remains an open question; I have provided some validation here when the computer has suggested particular conclusions, but much validation remains to be done across many musical arenas. The relationship of the Beach Boys to other works in the corpus study is in part problematic because further features working for general popular music may be required, if only to discount numerically their relevance to much noise music practice.

Noise music can reveal novel issues with feature extraction itself; for example, in this project, where normal sensory dissonance levels in popular music had not previously overloaded the detector, the threshold had to be reset to cope with Merzbow without registering a constant clipped value! The choices of thresholds of detection in the face of a wide variation of signal to noise ratios illustrates again some disconnection from a truly human-like listening experience. There may be earplugs and the stapedius reflex, or a degree of cultural learning about the possibilities of noise music, but the human auditory system does not require literal rewiring to cope with it; instead, noise music can exploit certain timbral limits of a pre-existing auditory biology. Should a computational system learn to appreciate greater

12 Matthew Blackwell, "In Left Field: Merzbow's Discography–Noise Music and the Taxonomic Drive" (Dec 7, 2011), accessed July 29, 2013, http://www.prefixmag.com/features/john-wiese-merzbow/merzbows-discography/59054/.

depths in noise music with repeated exposure, in order to model the route of human listeners?

Future computational musicology of noise music might use far larger databases that better reflect the diversity and mass release schedules of noise music across various currents of experimental and counter-culture work (consider the relationships with extreme metal, electronic body music and industrial acts, breakcore, 1960s experimental feedback pieces by Robert Ashley and Gordon Mumma, among others). An effort to form such a corpus must be accompanied by efforts to annotate properly at least a representative subset, so that a musicological ground truth is established for guidance. Finally, there are many more computational techniques to explore, such as navigational tools through corpora, or the use of complexity analysis to measure the ability of data compression algorithms to treat noise music.

 10.5920/noise.05