

INFLUENCE IN EARLY ELECTRONIC DANCE MUSIC: AN AUDIO CONTENT ANALYSIS INVESTIGATION

Nick Collins

University of Sussex

N.Collins@sussex.ac.uk

ABSTRACT

Audio content analysis can assist investigation of musical influence, given a corpus of date-annotated works. We study a number of techniques which illuminate musicological questions on genre and creative influence. By applying machine learning tests and statistical analysis to a database of early EDM tracks, we examine how distinct putatively different musical genres really are, the retrospectively labelled Detroit techno and Chicago house being the core case study. Further, by building predictive models based on works from earlier years, both by a priori assumed genre groups and by individual tracks, we examine questions of influence, and whether Detroit techno really is a sort of electronic future funk, and Chicago house an electronic extension of disco. We discuss the implications and prospects for modeling musical influence.

1. INTRODUCTION

Genre is a contentious area at the best of times [1], but an especial minefield in electronic dance music, where producers, journalists and consumers are always eager to promote new micro-genres [12]. As Brewster and Broughton have written of one highly strained genre term ‘if you name a genre of music after a club which was open for ten whole years and which was known for its eclecticism, you’re going to run into problems of definition pretty quickly. The word ‘garage’ is by far the most mangled term in the whole history of music’ [4, p. 307].¹

Electronic dance music’s origins range across African-American music and European synth pop, against a backdrop of increasingly affordable synthesis and sampling technology [6, 11, 15, 17, 18]. The important role through the 1980s of the US cities Chicago and Detroit as crucibles of new club music is unassailable, though they were not the only centres of activity (New York’s frenetic hip hop developments into electro, or the post-Moroder italo-disco movement in Europe are also worth mentioning, as indeed are general trends to danceable and synthesizer-laden pop

¹ ‘hardcore’ is another example of a heavily overloaded genre term.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

© 2012 International Society for Music Information Retrieval.

throughout the 1980s mainstream). The cities are enshrined in the genre names Chicago house and Detroit techno as two foundational pillars of later electronic dance music: they form the core of the study in this paper, though we do not assume without investigation that they are really as distinct as their names imply.

Even before audio content analysis investigation, there are good reasons for a musicologist to be wary of treating house and techno too individually in their 1980s growth. Chicago is around a five hour drive from Detroit, and Detroit artists often went to Chicago to sell their records in the larger market there; Derrick May sold Frankie Knuckles his TR-909 drum machine! The term ‘techno’ has many precedents, including track titles from Buggles, Yellow Magic Orchestra and Kraftwerk, although most famously used in *Techno City*, a 1985 Cybotron track co-produced by Juan Atkins (the elder of Techno’s ‘Belleville Three’). The genre term was finally applied as a differentiating stamp in 1988 for the *Techno! The New Dance Sound of Detroit* compilation curated by Neil Rushton, at Juan Atkins’ insistence on techno over Derrick May’s ‘Detroit house.’ Nevertheless, the compilation itself includes a ‘megamix’ at its close called *Detroit is Jacking* (jacking being a standard Chicago dance term) and another track called *Share this house* by Members of the House!

Detroit artists have themselves attempted to characterise musical differences with Chicago. In the liner notes to the *Techno!* compilation Derrick May writes ‘House still has its heart in 70s disco; we dont have any of that respect for the past, its strictly future music. We have a much greater aptitude for experiment’ (sic) [9], and most famously, that ‘It’s like George Clinton and Kraftwerk are stuck in an elevator with only a sequencer to keep them company’ [9]. The hypothesis of Chicago house as an extension of disco, and Detroit techno as a combination of electronic funk and synth pop, will be examined herein.

Two previous studies of musical influence [5, 8] published in ISMIR, on synth pop and sampling, have indicated the benefits of data-annotated corpora in new musicological investigations. Through musical similarity measures, this paper examines the use of automatic audio content analysis to establish the strength of links between historic tracks and putative genre groupings. Where Bryan and Wang [5] worked on an existing database of annotations over sample-based music concerning ‘WhoSampled who’ (<http://www.whosampled.com/>), Collins [8] examined audio similarity between date-annotated synth pop as

Genre	Dates	Num Tracks	Duration (mins)	Notes
Chicago House	1986-1989	31	197.7	Sourced in particular from <i>Chicago Trax</i> and <i>The Original Chicago House Classics</i> as well as compilations including <i>Warp10+1: Influences</i>
Detroit Techno	1986-1989	31	186.5	Including Derrick May, early Model 500, and the <i>Techno!</i> compilation. No second wave, nor Cybotron
1980s Pop	1985-1989	31	127.7	Including Michael Jackson, Madonna, Prince
Funk	1965-1978	31	118.6	Selected tracks from Parliament's <i>Mothership Connection</i> , <i>The Godfather</i> , <i>James Brown</i> , <i>The Very Best Of...</i> and <i>Funk Soul Classics</i>
Disco	1973-1980	31	112.5	Selected tracks from <i>Anthems Disco</i> and <i>Disco Fever</i>
Synth Pop	1977-1981	31	145.6	Including Kraftwerk, Human League, Gary Numan, Ultravox, Depeche Mode
Electro and Hip Hop	1980-1984	31	180.6	Some early rap, with an emphasis on the transition into electro. Includes Grandmaster Flash and the Furious Five <i>The Message</i> (1982) and excerpts from <i>The Tommy Boy Story</i>
Punk/Post-Punk	1977-1979	31	89.7	UK artists including Sex Pistols, UK Subs, Wire, The Cure, Gang of Four

Table 1. Overview of music corpus

part of the process of identifying influence. The latter might be justified as the more general case and is followed here: it is of particular import when scaling up to larger databases of audio where annotations are impractical for musicologists. Network techniques introduced in [5] are still valuable in providing applicable metrics for later analysis once similarity scores are established. However, this paper will look at direct first generation influence rather than longer-term networks spanning chains of multiple nodes.

The paper proceeds through section 2 detailing the set of 248 source audio files split over eight genre groups, and section 3 which discusses the technicalities of the predictive models used. Section 4 explores the separability of genre groups suggested, using machine learning algorithms, and the Anderson Darling statistical test to look for any rejection of the null hypothesis that they are drawn from the same distribution. In section 5 we apply the predictive models to examine questions of the strength of influence of precursor work on Detroit techno, Chicago house, and a late 80s pop control group. As well as working with models trained on whole groups of tracks, we also run tests for some famous individual tracks, such as Donna Summer's *I Feel Love* (1977). Finally, in section 6 we explore the implications of the experimental findings, and broach larger questions for studies of musical influence using MIR techniques.

2. SOURCES

Table 1 is an overview of the materials used in this study. Eight genre groups are presented, with 31 tracks per group. Five of these are precursor genres, movements in popular music from the mid 1960s to the early 1980s. The three top groupings are Chicago house, Detroit techno, and a control group of mid- to late-1980s pop including Madonna and Michael Jackson, coincident with an explosion in popularity of electronic dance music in the UK. The earlier genre groups include four important to the origins of electronic dance music: funk, disco, synth pop and hip-hop (particularly in its electro form). Some UK punk and post-punk records are included as a further control.²

²Joy Division were specifically excluded, since their New Order manifestation intersects with electro circa 1983.

Although the total duration of the genre groups differ, the critical thing is the equal number of examples in each, since tests are based on equal length excerpts from individual tracks, or otherwise involve a normalization for duration, such as the average log loss of a predictive model. 1980s examples of electronic dance music tend to involve longer tracks, where many 1960s and 1970s singles are much closer to the three minute (or less) pop song (short songs were also revived with punk's throwaway numbers); creating groups of an equal number of tracks all balanced in duration and the number of years associated is an unsolvable dynamic programming challenge.

There are many overlaps between these ostensibly separate groups, such as the shift to disco via Philly soul, the use of synthesizers by new wave acts as well as more explicitly by synth pop groups, or the appropriation of funk and disco backings in early rap records.³ The a priori use of genre groups is justified on the grounds of the musicians' statements themselves, such as Derrick May cited above, who treat 'funk' and 'disco' as known areas of musical endeavour. The groups have been constructed from well known examples of the genres in question, and one confound in particular avoided in construction; synthesized electronic instrumentation in disco is not represented in the disco group, but a few examples from the Moroder camp are included under synth pop. Part of our analysis shall be to consider the well-definition of the groups, in terms of their internal consistency; as well as considering genre based influence, we shall also take a look at influential individual tracks later in the paper, to avoid any claims of resting too heavily on genre constructions.

No categorisation can be perfect, and there are some missing early 1980s genre groups, such as European electronic body music and industrial (e.g. *Liaisons Dangereuses*, Front 242) and italo-house (e.g. Klein and MBO, Alexander Robotnick), and mid 1980s New York production developments (freestyle, Mantronix etc.). Manageability of the overall study, and the greater overlap with the formation dates in Chicago and Detroit, made these categories out of the scope of the current investigation; however, again,

³The term 'house' itself floats around in 1983, for example on *Rock the House* by Pressure Drop, a 1983 release on Tommy Boy.

we return to a few individual tracks rather than whole genre groups below. Although there are some earlier prototype Chicago house tracks, such as Jesse Saunders' *On and On* (1984), we have avoided these for some separation in date from the precursor genres in this study; most commercially available international Trax releases, for example, tend to be available from 1986 at the earliest.⁴

Complete track lists can be made available on researcher request; all music was purchased.

3. PREDICTIVE MODELING

Bag of features assumptions [7] are avoided in favour of using time series modeling to construct predictive models; in particular, Prediction by Partial-Match (PPM) variable order Markov models [2, 13]. The strength of prediction of one piece or group of pieces by another is measured by average log loss in information theoretic terms [2], as further detailed below.

Various musical attributes of the pieces under consideration are modelled, such as timbral, rhythmic and harmonic change components. The final model combines three core elements:

1. A model τ of timbre based on 11 features, with feature vectors accumulated by beats, vector quantised by a k-Means classifier into symbols, and used to train a PPM model
2. A model ι of inter-onset intervals after onset detection on polyphonic audio, using a classifier for IOI sizes into symbols, and subsequent PPM
3. A model η of harmonic change, based on extracting beat-wise 12TET chroma, forming the sum of differences between beats, a classification into symbols, and PPM

For τ , a more general set of timbral descriptors was selected than MFCCs, to try to reflect the character of analyzed audio, without high dimensionality (which would impact on the k-Means step). The timbral features were perceptual loudness, sensory dissonance (using a Sethares model [16]), two transient detection measures using the wavelet method of [10], spectral centroid, spectral percentile at 0.8% and 0.95% energy, zero crossing rate, spectral crest measure, spectral slope, and a raw onset detection function (the preprocessed signal for an onset detector). The onset detector for the raw detection function, and for the isolation of onsets for IOI detection in ι , is based on work by Stowell [19], and is applied to polyphonic audio tracks; the rectified complex deviation (RCD) onset detection function used here has proven reasonable for such applications. All features were subject to normalization with respect to corpus derived minimum and maximum values, and were gathered in beats by averaging feature vectors. Chroma for η were also accumulated in beats, the difference between

beatwise chroma vectors taken, and summed over the difference vector. This created a one-dimensional measure of harmonic change, where the summation process avoided issues with different absolute pitch centres in the music.

Feature vectors, IOIs and delta chroma sums were subject to vector quantisation into 20 tokens before PPM modeling. In order to symbolize the multidimensional timbral feature vectors in τ , vectors extracted from the training corpus were clustered with the unsupervised k-Means algorithm, with $k=20$. As one dimensional quantities, the IOIs in ι and harmonic change sums in η were classified into twenty bands by histogram equalisation [3, p.188]. A histogram for categorisation was constructed by sorting the values into order, splitting them by twenty equal size bands, and taking histogram bin positions by the maximum in each band. Twenty bands was a good compromise for a reasonable alphabet size for the PPM, without invoking too high a dimensionality. PPM models were then trained on the sequences in the 20 token alphabet, using consecutive subsequences of five values at a time. The particular model variant used here is what Pearce and Wiggins [13] denote the PPM-AX variant.

Scoring for a given PPM model γ on novel data set X was then calculated by

$$\text{average log loss}_{\gamma}(X) = \frac{-\sum_{x \in X} \log_2 P(x|\gamma)}{|X|} \quad (1)$$

where the x are all the sequence contexts of the data to be tested [2] Minimal scores correspond to high probability sequences, that is, highly expected with respect to the model. The log is critical to avoid floating point underflow on multiplication of small probabilities. This average log loss measure is robust to different durations of sequences considered; in any case, we use equal length excerpts from pieces.

In applying this to a corpus, a predictive model is trained on a subset of pieces. The model can then be used to examine one or more target pieces, via equation 1. In this work, the models are applied to equal size groups of pieces, summing the model predictions within the group to get a total score for the predictability of that group with respect to the probabilistic model.⁵ The final scores are actually the sum of those from the three models τ for timbre, ι for rhythm and η for harmonic change; these three components are individually normalized before the final sum. Whilst a given model's predictions are internally comparable, care must be taken in comparing the absolute value of scores between different models; the normalization reflects that only relative degree and order is comparable.

All implementations used the open source SuperCollider Music Information Retrieval library by the author,⁶ which includes an example with the three component model presented here. Specific client source code for the work is available on researcher request.

⁴ One example of a famous and influential track which was recorded earlier but released much later is Phuture's *Acid Tracks*, recorded late 1985, released 1987.

⁵ For a common group size, we can divide by the group size without compromising comparability, rather than taking an average over a different number of contributing members.

⁶ <http://www.sussex.ac.uk/Users/nc81/code.html>

Model	Probability of rejecting null hypothesis							
	Chicago	Detroit	Pop	Funk	Disco	Synth Pop	Electro	Punk
Chicago	0.4297	0.0033	0.3879	0.2131	0.0776	0.0832	0.3986	0.2532
Detroit	0.0033	0.0023	0.0013	0.0005	0.0001	0.0004	0.0074	0.0003
Pop	0.3879	0.0013	0.4402	0.1113	0.2218	0.0283	0.2739	0.2109
Funk	0.2131	0.0005	0.1113	0.3554	0.0040	0.1224	0.1629	0.2456
Disco	0.0776	0.0001	0.2218	0.0040	0.3057	0.0003	0.0625	0.0741
Synth Pop	0.0832	0.0004	0.0283	0.1224	0.0003	0.0982	0.0671	0.0474
Electro	0.3986	0.0074	0.2739	0.1629	0.0625	0.0671	0.3366	0.1992
Punk	0.2532	0.0003	0.2109	0.2456	0.0741	0.0474	0.1992	0.4389

Table 2. Application of Anderson-Darling tests within and between genre groups. Each cell entry is the probability of rejecting the null hypothesis that the tracks being tested together are a homogenous entity. Significant table entries are in bold, with respect to a Bonferroni significance level for multiple comparisons.

4. MACHINE LEARNING TESTS AND STATISTICAL TESTS OF SEPARABILITY FOR REPRESENTATIVE FEATURE VECTORS

An initial examination was carried out on the genre groups themselves, to see how “separable” the genre groups were from one another with respect to the ability of machine learning to differentiate them, and in terms of statistical tests for their internal and paired consistency. All tests were repeated twice, first for the timbral feature vector detailed in section 3 for model τ , and secondly for a vector of 11 MFCCs. Single vectors summarised single tracks; 45 seconds of feature vectors were extracted from a point 25% of the way into a given sound file, and averaged (normalization factors had already been calculated across the entire corpus of 248 tracks in an earlier sweep). ARFF files were exported for tests in Weka, and arrays of data into MATLAB for statistical tests.

For machine learning, we tested the discriminatory power of supervised classifiers to learn the training sets (given the genre labels 0-7), and of unsupervised clustering algorithms to match these labels (the ‘classes to clusters’ evaluation setting in Weka). Over 8 genres, the best scores came from the 11 different features rather than the MFCCs,⁷ but were still of low classification success. The best results were for a k-Means clusterer (correctly classified 69 of 248 instances, 27.8%) and naive Bayes (68 of 248, 27.4%); a range of other algorithms were investigated including neural nets and J48 decision trees. The best MFCC results were for k-Means (correct 47, 18.95%), and naive Bayes (correct 53, 21.4%). Examination of 2-dimensional subsets of features revealed a lot of overlap between genres. This result motivated using a more sophisticated time series modeling approach rather than average feature vectors, and using the mixed feature vector for timbre rather than MFCCs. With just the house and techno groups, and the heterogenous feature vector, classification accuracy was around 50% (at chance given two groups), with best performance from k-Means (correctly classified 37 of 62, 60%) and naive Bayes (34 of 62, 54.8%). For the 11 MFCCs, nothing better than 33 out of 62 (53.2%) accuracy was observed, with most algorithms performing worse than chance.

Statistical tests were also applied to the model τ feature vector data, to look for overlap between genre groups,

⁷ Vectors of 40 MFCCs were also tested without any improvement in classification scores.

and internal consistency. An Anderson-Darling test was utilized [20], which tests the null hypothesis that all feature data arose from the same distribution (without assuming normality of that distribution); a significant p-value indicates that the data is from multiple distributions. Table 2 presents the symmetric matrix of values for all pairwise (62 tracks at a time) and within-group (31 tracks) tests. The Detroit techno group is seen as more heterogenous, and the null hypothesis would be rejected if the threshold was set at 0.05% p-value. Because there are 28 pairs and 8 individual genres = 36 tests, under the Bonferroni correction the p-value of $0.05/36 = 0.0013889$ has the statistical power to cover everything up to 0.05 and may give a better sense of whether the Detroit techno result is aberrant; we may keep the null hypothesis of Detroit techno as consistent at this level. To the extent that the probabilities point to degree of homogeneity, the Detroit corpus is more heterogenous. The make-up of the Detroit corpus unbalances things when paired with other groups, whilst there seems to be a strong overlap of Chicago house and late 80s pop (given UK No.1s by Chicago house producers, this may not be so unexpected) as well as with electro and funk. Nonetheless, the average feature vector approach is quite coarse, and the predictive models are now deployed for a finer-grained examination.

5. INVESTIGATION OF INFLUENCE THROUGH PREDICTIVE MODELS

In this section, results are reported for the predictive scores given to particular genre groups, and to individual tracks, from models constructed from first genre groups, and then some interesting precursor tracks. Section 3 describes the model construction and algorithm for prediction scores. Scores are normalized for a particular run of a particular model to fall from 0 to 1, where 0 would be totally predicted by the model at probability 1, and 1 is the least expected observed situation. Models can be constructed in two directions; we favour constructing a model from an earlier historical genre or track, to predict later tracks. There is an argument that construction in the opposite direction would also indicate the degree of derivation of the later work from earlier, and we report such constructions symmetrically for the genre groups. We do not form the fully symmetric score from a matrix plus its

Model	Prediction Score							
	Chicago	Detroit	Pop	Funk	Disco	Synth Pop	Electro	Punk
Chicago	<i>*0.04013*</i>	0.81935	0.83104	0.81382	0.82633	0.7646	0.87111	0.80286
Detroit	0.7951	<i>*0.04889*</i>	0.8146	0.77913	0.79449	0.76644	0.861	0.75883
Pop	0.8961	0.86525	<i>*0.04548*</i>	0.8945	0.83403	0.85331	0.90055	0.89854
Funk	0.85148	0.85413	0.90196	<i>*0.02236*</i>	0.91059	0.84123	0.89624	0.91292
Disco	0.83158	0.76417	0.82118	0.8417	<i>*0.03575*</i>	0.76346	0.89092	0.813
Synth Pop	0.8846	0.88163	0.88508	0.87967	0.8056	<i>*0.07829*</i>	0.88813	0.86481
Electro	0.85575	0.85762	0.89833	0.87927	0.91642	0.82893	<i>*0.02486*</i>	0.87419
Punk	0.76547	0.70549	0.89158	0.91145	0.86657	0.81012	0.86613	<i>*0.03303*</i>

Table 3. Prediction scores of genres from models constructed for each genre. Starred italics on the diagonal correspond to the prediction of the data used to construct a model by that model; one bold entry in each row indicates the closest other genre to the model genre.

transpose, since model construction itself is not guaranteed to produce scores in exactly the same ranges, and the post calculation normalization reported here, whilst helpful for seeing links, is not uniform in comparison (even unnormalized, differences in model construction would question comparability; results are relative to a given model).

Table 3 presents the asymmetric matrix of results over the predictive models. The diagonal is italicized; all models find their own source database most highly predicted, as we'd expect for any sensible probabilistic modeling. The description of Chicago house as disco with a drum machine, and Detroit techno as future electro funk, is only partially borne out in these figures. One aberrant factor is the close link of Detroit techno with late 1970s punk and post punk guitar tracks; Chicago house tracks are also seen as closer to punk than Detroit techno on this view. Of the three factors in the scores, the onset detection driven IOI model is the point of similarity here; a related density of events has an impact, as does timbre to an extent, perhaps through some degree of sonic exploration in house and techno instrumental tracks. Examining relative values within rows, the links to synth pop are clear; Detroit is closer to funk than disco, but Chicago also that way round. From the funk model, Chicago is very marginally ahead of Detroit, if within the third decimal place. Electro is closest to synth pop, which is musicologically sensible; the synth pop model finds Chicago, Detroit and late 80s pop of a muchness in terms of potential influence.

The relative degree of influence of a seminal piece can be investigated by creating a predictive model from it. Table 4 compares 22 interesting tracks from the 1970s to the early 1980s; these precursor tracks were selected from mentions in sources on EDM history such as [11, 18]. Complete individual tracks are used to form predictive models, which are then deployed to predict the Chicago house, Detroit techno and late 1980s pop corpuses.

Given these mainly synthesizer-flavored precedents, Detroit makes the most whole-hearted embrace of the technologized future, and shows the greater link to James Brown to boot (though not Parliament, which links more to pop, perhaps through the inclusion of Prince in that corpus in particular). The sanity checks show some consistency, with two versions of *Planet Rock* both leading to similar results, and two runs on the same Kano track also coming out with a similar ordering. The tests were repeated over

the whole set of songs, using a version of the predictive model with 10 states rather than 20 per vector quantiser, without any great divergence from the results presented here, excepting *Mothership Connection*, *Numbers* and *Clear* being assigned to techno, *Magic Fly* to Chicago, and *Problèmes D'Amour* to pop. The greater vocal content in Kraftwerk's *The Model* may be an explanation of its stronger link to certain elements of the Chicago house corpus, or the link of the female vocal of *I Feel Love* through to pop.

Individual tracks across the corpus of 93 can be examined, to find the most predicted and the most divergent from a model. For instance, for Kano's *It's A War*, the three closest were Derrick May's *Spaced Out* and *Nude Photo*, and Blake Baxter's *Ride Em Boy*, all three from the Detroit corpus (as we might hope for this track's reception history, though there are also aural links to Prince), and the furthest away, in pop, Madonna's *Live to Tell* and the Bangles' *Eternal Flame* and *Hazy Shade of Winter*.

6. DISCUSSION

Musical influence is a complex mechanism; the assumption that strength of prediction is related to degree of influence seems reasonable, but may hide other factors, such as indirect influence, common equipment and teaching tools (such as music technology magazines), social currents, and even independent co-creation of the same idea.

The audio content analysis used here cannot be claimed to be on a par with the musicologist's ear. On the other hand, computer tools can point to useful currents of inquiry, and provide an alternative stimulus to musical historical and analytical investigation. Furthermore, it is really worth exploring the musicological applications at an early stage, to clarify the potential impact of such tools, and feedback their effectiveness.

The genre groups used in this study make categorical assumptions which can hide musical continuity. Whilst their construction was to answer some questions of influence and overlap, the most interesting results relate more to the scope of individual tracks. Future work may drop genre assumptions entirely, creating a predictive model from every individual track, to assess every other; given pairwise similarity measures and chronological distance, multidimensional scaling may give insight into structure. It may also be productive to consider rates of change per year, by

Model	Chicago	Detroit	Pop
Giorgio Moroder <i>From Here To Eternity</i> (1977)	0.6551	0.5998	0.6562
Donna Summer <i>I Feel Love</i> (1977)	0.6461	0.6824	0.6346
Kraftwerk <i>The Model</i> (1978)	0.6659	0.6929	0.7564
Kraftwerk <i>Numbers</i> (1981)	0.6146	0.5761	0.5723
Kraftwerk <i>Trans-Europe Express</i> (1977)	0.5419	0.4784	0.5648
Cerrone <i>Supernature</i> (1977)	0.9124	0.8361	0.8573
Dee D. Jackson <i>Automatic Lover</i> (1978)	0.6994	0.6878	0.7077
Space <i>Magic Fly</i> (1977)	0.6855	0.6993	0.6753
Sylvester <i>You Make Me Feel (Mighty Real)</i> (1978)	0.8448	0.7551	0.8348
Lipps Inc. <i>Funkytown</i> (1979)	0.7209	0.6336	0.6734
Kano <i>It's A War</i> (1980)	0.5009	0.4494	0.5877
Kano <i>It's A War</i> (1980)	0.5974	0.5235	0.6465
Soft Cell <i>Tainted Love</i> (1981)	0.8372	0.8085	0.8496
Depeche Mode <i>Get The Balance Right</i> (1983)	0.7934	0.727	0.7852
Afrika Bambaataa et al. <i>Planet Rock (12" Vocal Version)</i> (1982)	0.8073	0.7501	0.7697
Afrika Bambaataa et al. <i>Planet Rock</i> (1982)	0.7959	0.7562	0.7701
James Brown <i>Funky Drummer Pts. 1 and 2</i> (1970)	0.6383	0.5262	0.6318
Parliament <i>Mothership Connection (Star Child)</i> (1975)	0.8166	0.798	0.7881
Alexander Robotnick <i>Problèmes D'Amour</i> (1983)	0.6395	0.6128	0.6496
Cybotron <i>Enter</i> (1983)	0.707	0.6833	0.6852
Cybotron <i>Clear</i> (1983)	0.5503	0.5552	0.5601
Cybotron <i>Cosmic Cars</i> (1983)	0.65	0.6387	0.657

Table 4. Prediction of genres using models constructed from individual tracks. The closest genre from each model is indicated in bold.

constructing a model on one year (or on other windows of time) and testing how predictable the next is.

A human study of similarity on this corpus would be a useful follow-up, though one confound is the factor of recognition; expert musicologists of EDM would recognise many of the Chicago and Detroit tracks immediately, and the corpus used here involves many famous works. Its historical importance, however, makes it a very interesting corpus to work on, of great musicological relevance.

In future work we may consider extending out to a larger-scale investigation of the history of electronic music. Alternative time series models, such as Hidden Markov Models, could be employed, avoiding vector quantisation simplifications, and possibly using symmetrised distance measures such as the cross-likelihood discussed in [21]. More sophisticated statistical models of causality may also help to stretch the machinery for modeling influence [14].

7. CONCLUSIONS

This paper presented a study of applying MIR techniques to probe the borderline of Chicago house and Detroit techno. More generally, we related later 1980s works to 1960s to early 1980s precursors through a number of methods. We saw that the house and techno genres overlap, and are not necessarily tightly defined. Nonetheless, there was some corroboration of Derrick May's characterisation of Detroit techno as futuristic in its sound world, though less support for its separate funkiness; the disco and synth

pop heritage is a strong link to the two styles. The study presents a template for future work over the same corpus, as refined sound analysis models become available, and for more general future audio-content driven examination of musical influence.

8. REFERENCES

- [1] J. J. Aucouturier and F. Pachet. Representing musical genre: A state of the art. *Journal of New Music Research*, 32(1):83–93, 2003.
- [2] R. Begleiter, R. El-Yaniv, and G. Yona. On prediction using variable order Markov models. *Journal of Artificial Intelligence Research*, 22:385–421, 2004.
- [3] G. Bradski and A. Kaehler. *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media, Sebastopol, CA, 2008.
- [4] Bill Brewster and Frank Broughton. *Last Night a DJ Saved My Life*. Headline Book Publishing, London, 2006.
- [5] Nicholas J. Bryan and Ge Wang. Musical influence network analysis and rank of sample-based music. In *Proceedings of the International Symposium on Music Information Retrieval*, Miami, October 2011.
- [6] Mark J. Butler. *Unlocking the Groove*. Indiana University Press, Indiana, 2006.
- [7] Michael 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(4):668–696, April 2008.
- [8] Nick Collins. Computational analysis of musical influence: A musicological case study using MIR tools. In *Proceedings of the International Symposium on Music Information Retrieval*, Utrecht, August 2010.
- [9] Stuart Cosgrove. Techno! the new dance sound of Detroit (liner notes), 1988.
- [10] Laurent Daudet. Transients modelling by pruned wavelet trees. In *Proceedings of the International Computer Music Conference (ICMC)*, Havana, Cuba, 2001.
- [11] Peter Kirn, editor. *Keyboard Presents the Evolution of Electronic Dance Music*. Backbeat Books, Montclair, NJ, 2011.
- [12] Kembrew McLeod. Genres, subgenres, sub-subgenres and more: Musical and social differentiation within electronic/dance music communities. *Journal of Popular Music Studies*, 13:59–75, 2001.
- [13] Marcus Pearce and Geraint Wiggins. Improved methods for statistical modelling of monophonic music. *Journal of New Music Research*, 33(4):367–385, 2004.
- [14] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.
- [15] Simon Reynolds. *Generation Ecstasy: Into the World of Techno and Rave Culture*. Routledge, New York, 1999.
- [16] William A. Sethares. *Tuning Timbre Spectrum Scale (2nd Edition)*. Springer Verlag, Berlin, Germany, 2005.
- [17] Peter Shapiro. *Turn the Beat Around: The Secret History of Disco*. Faber and Faber Limited, London, 2005.
- [18] Dan Sisko. *Techno Rebels: The Renegades of Electronic Funk (2nd Ed.)*. Wayne State University Press, Detroit, MI, 2010.
- [19] D. Stowell and Plumbley. M. D. Adaptive whitening for improved real-time audio onset detection. In *Proceedings of the International Computer Music Conference (ICMC)*, Copenhagen, 2007.
- [20] A. Trujillo-Ortiz, R. Hernandez-Walls, K. Barba-Rojo, L. Cupul-Magana, and R.C. Zavala-Garcia. Andarkstest: anderson-darling k-sample procedure to test the hypothesis that the populations of the drawn groups are identical., 2007.
- [21] Tuomas Virtanen and Marko Helén. Probabilistic model based similarity measures for audio query-by-example. In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New York, 2007.