Audio Engineering Society

Convention Paper

Presented at the 128th Convention
2010 May 22–25 London, UK

The papers at this Convention have been selected on the basis of a submitted abstract and extended precis that have been peer reviewed by at least two qualified anonymous reviewers. This convention paper has been reproduced from the author's advance manuscript, without editing, corrections, or consideration by the Review Board. The AES takes no responsibility for the contents. Additional papers may be obtained by sending request and remittance to Audio Engineering Society, 60 East 42nd Street, New York, New York 10165-2520, USA; also see www.aes.org. All rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

Perceptually-motivated audio morphing: warmth

Duncan Williams, Tim Brookes[1]

[1] Institute of Sound Recording, University of Surrey, Guildford, GU2 7XH. United Kingdom
duncan.williams@surrey.ac.uk, t.brookes@surrey.ac.uk

ABSTRACT

A system for morphing the warmth of a sound independently from its other timbral attributes was coded, building on previous work morphing brightness only (1), and morphing brightness and softness (2). The new warmth-softness-brightness morpher was perceptually validated using a series of listening tests. A Multidimensional Scaling analysis of listener responses to paired-comparisons showed perceptually orthogonal movement in two dimensions within a warmth-morphed and everything-else-morphed stimulus set. A verbal elicitation experiment showed that listeners' descriptive labeling of these dimensions was as intended. A further 'quality control' experiment provided evidence that no 'hidden' timbral attributes were altered in parallel with the intended ones. A complete timbre morpher can now be considered for further work, and evaluated using the tri-stage procedure documented here.

1. INTRODUCTION

Audio morphing is a technique for creating a new sound with some of the characteristics of each of a pair of existing sounds (3). Many systems focus on interpolating spectral, or spectral and temporal, features from a source sound and a target sound, to create a new feature set for additive synthesis of a hybrid sound (4).

This project combines a psychoacoustic focus with the Digital Signal Processing (DSP) approach found in existing audio morphers. Our long-term goal is the implementation of a perceptually configured interface, creating a timbre morpher capable of processing individual timbral attributes independently from one another. Outside novel sound manipulation by timbre morphing, alternative future applications could also include timbral metering or synthesis.

A previous study developed a single-attribute timbre morpher and established that a morph of brightness, a single perceptual dimension with spectral correlates, was possible. This was followed by a study which established that the timbre morpher could also control softness, an additional perceptual dimension with a temporal acoustic correlate. This iteration of the timbre morpher's development was designed to establish the its adaptability to control of attributes with some degree of acoustic overlap amongst their correlates.

1.1. Selecting a third timbral attribute

Warmth was selected as the third timbral attribute to be implemented in the morpher, following a review of published research into timbral attributes which suggested that warmth had some degree of acoustic overlap with brightness (5) (6). Brightness has been shown to be well correlated with spectral centroid (7), and warmth has been correlated with spectral slope, spectral centroid, and the energy of the first three harmonics relative to the remainder of the signal (8).

It is highly likely that any processing increasing warmth will generate some decrease in perceived brightness, as the spectral centroid of the hybrid sound will ultimately become lower. Similarly, stimuli generated by the brightness-only morph may result in a decrease in warmth.

A series of perceptual tests was designed and employed to evaluate the success of the addition of this third attribute to the morpher. Firstly, a pairwise dissimilarity experiment was designed to establish whether or not the manipulation of these correlates would be perceived as orthogonal to that of other timbral attributes, including brightness and softness. Secondly, a verbal elicitation experiment was specified to provide appropriate labels for the timbral movements perceived. A third 'quality control' experiment was then designed, to confirm that no additional or unwatned timbral changes were created by the morpher, and to evaluate the morpher with 'real-world' stimuli.

2. SPECIFYING THE WARMTH MORPHING ROUTINE

The timbre morpher was developed with the open source Spectral Modeling Synthesis (SMS) platform in MATLAB. The previous iteration of the morpher was designed to morph softness and brightness independently of one another. Existing timbral literature suggested that the warmth morphing algorithm should be designed to manipulate spectral centroid and spectral slope via the amplitude of the first three harmonics relative to the remainder of the input signal [8].

2.1. Extracting and interpolating warmth

In order to be adaptable to a wide range of signals, including those that are not harmonic, the warmth morphing algorithm was designed to manipulate the energy within the spectral area from the frequency of the signal's fundamental, up to 3.5 times this frequency (the energy in this area is denoted warmth region, or WR). In a harmonic signal this area would comprise the first three harmonics; in a signal containing inharmonic partials and/or noise elements the energy in this area was intended to perform a similar timbral role to that of the first three harmonics. The energy across the full spectrum (frame magnitude, or FM) minus WR gives the energy outside the warmth region (remaining magnitude, or RM). The signal's warmth is then calculated as WR/RM. This system for extracting the correlates of warmth, combined with the previous systems for extracting softness and brightness, is illustrated in Figure 1. The system for interpolating warmth towards from that of a source toward that of a target sound is illustrated in Figure 2.

It is acknowledged that this approach may not deal appropriately with a signal containing no clear fundamental, nor with one containing significant noise-like energy at a lower frequency than its fundamental. Developing a suitable approach to these special-case signal types is reserved for further work.

2.2. Compensating acoustic overlap

In order to compensate for unwanted perceptual change, the algorithm was designed to neutralize any change in brightness caused by morphing warmth, by applying a post-warmth-morph reverse-brightness morph. This will mean that the hybrid feature set will have been morphed in warmth towards the target, and then in brightness back towards the source. This system is illustrated in Figure 3.

The knock-on effect of this compensation system could be that a 100% warmth morph, compensated in brightness with a slight spectral tilt, might then only equate to a 90% warmth morph (due to the brightness adjustments moving the warmth region amplitude a small part of the way back to where it started). Perceptual evaluation was necessary in order to determine whether there were any negative effects of this compensation routine.

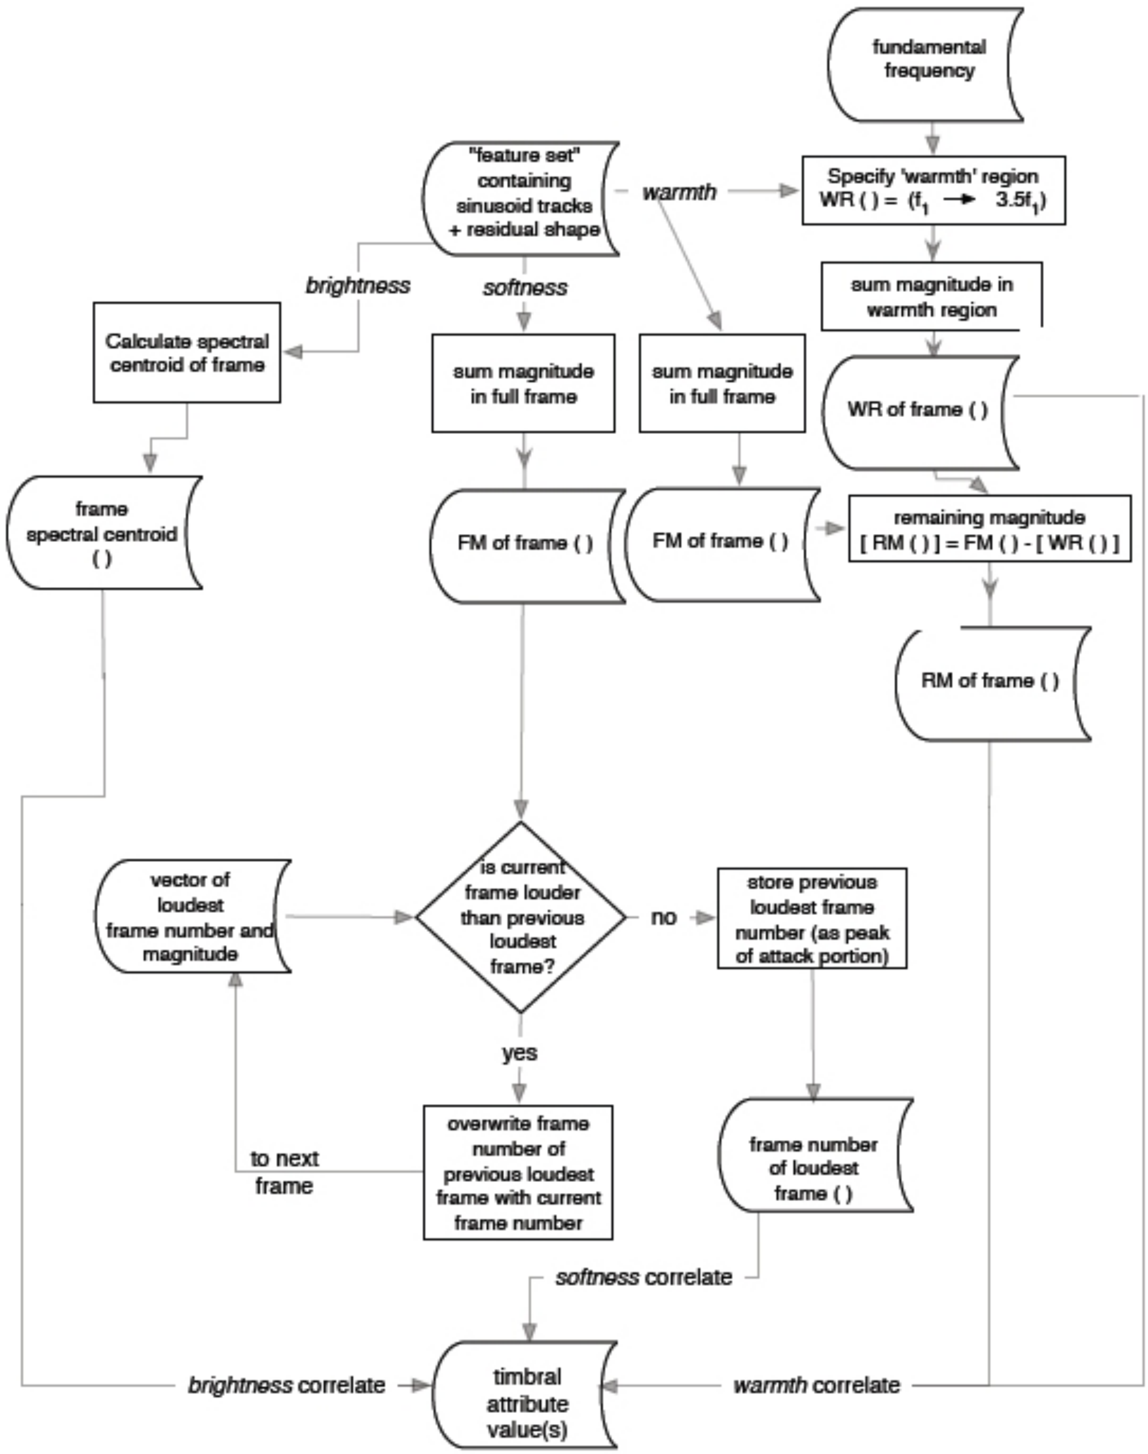

Figure 1 Signal flow of subroutine extracting acoustic correlates for all three included timbral attributes. 'Sinusoid tracks' represent discernible partials; 'residual' represents the remainder of the signal.

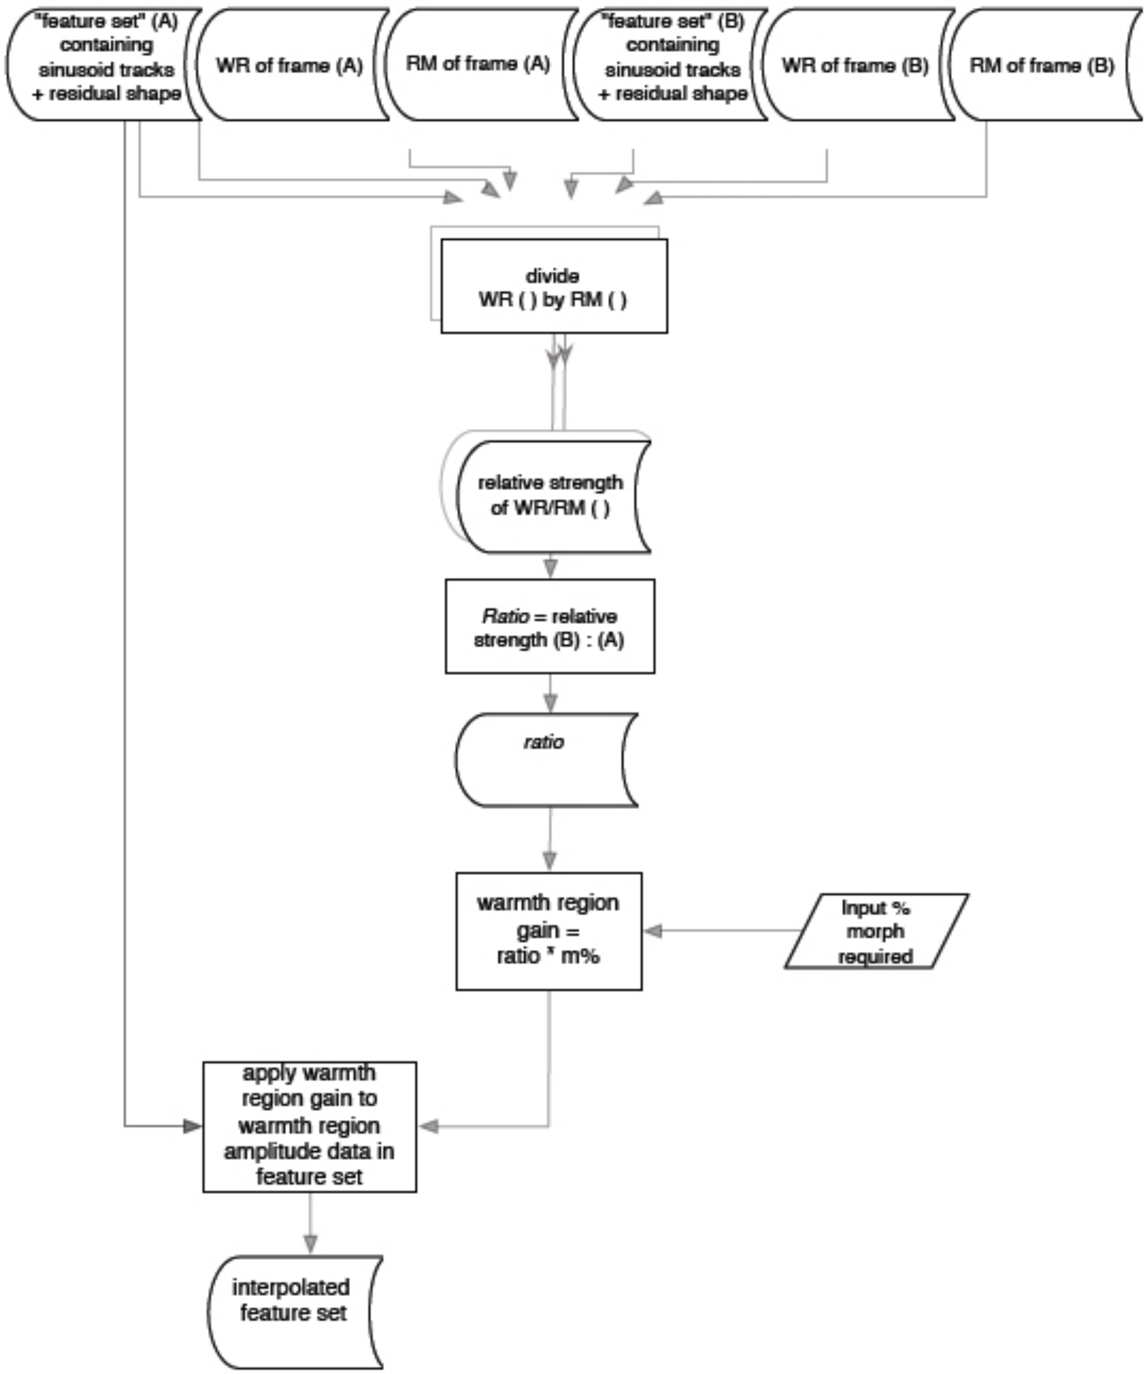

Figure 2 Signal flow of sub-routine interpolating warmth towards that of a target value

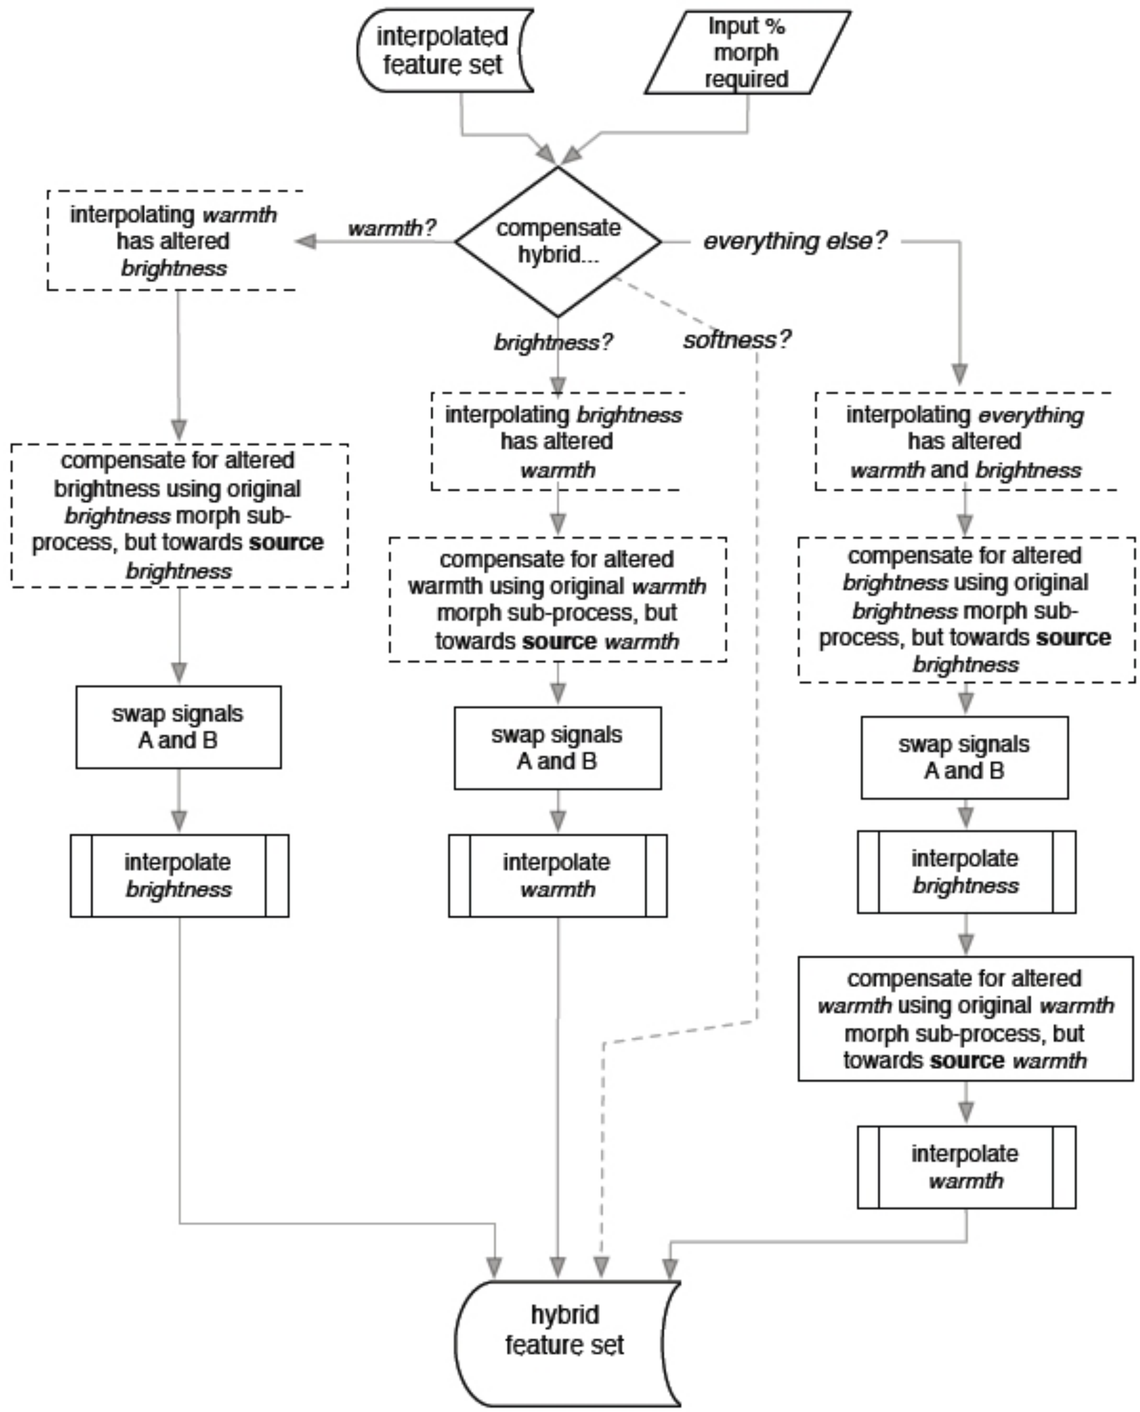

Figure 3 Signal flow of sub routine compensating for unwanted timbral changes

2.3. Discrete morph types

The addition of these algorithms to the brightness and softness morphing system should allow for the following types of audio morph:

1. Morph everything
2. Morph brightness only
3. Morph softness only
4. Morph warmth only
5. Morph brightness and softness
6. Morph brightness and warmth
7. Morph softness and warmth
8. Morph everything else apart from softness
9. Morph everything else apart from brightness
10. Morph everything else apart from warmth
11. Morph everything else apart from brightness and softness
12. Morph everything else apart from warmth and brightness
13. Morph everything else apart from warmth and softness
14. Morph everything else apart from warmth or softness or brightness

There are four orthogonal dimensions in this list of possible morph types: morph brightness (type 2), morph softness (type 3), morph warmth (type 4), and morph everything-except-warmth-softness-and-brightness (type 14). Various combinations of these orthogonal morph types can produce all the remaining types of morph listed.

3. PERCEPTUAL TESTING

A series of perceptual tests was designed and employed to evaluate the success of the addition of the third attribute to the morpher.

3.1. Pairwise dissimilarity experiment

A pairwise dissimilarity experiment was designed to test the perceived dimensionality of morphs related to the new attribute, using three of the morph types listed in section 2.3: morph warmth, morph everything, and morph everything else apart from warmth. Stimuli were generated using various percentages of morphing between one source and one target sound (synthesised, so that amplitude, fundamental frequency, envelope, and duration could remain constant). The source was a sawtooth with an attack time of 5ms; the target was a triangle wave with a 250ms attack time. These particular waves were chosen to provide clear and simply-quantifiable differences in attack time, spectral centroid and harmonic structure, allowing a wide range of morphed stimuli. Stimuli throughout the experiment were loudness equalised according to the experimenter's ear.

Allowing four stimuli per dimension meant a minimum of 16 stimuli were required in order to potentially reveal movement in up to four dimensions in subsequent MDS analysis. 19 stimuli were used, in order to have an equal number of steps across each type of morph, as shown in Table 1. For each stimulus, the percentage of morphing towards the target was chosen to provide (to the experimenter's ears) perceptually even steps across each intended dimension. The stimuli were stored in PCM Wave format, at 16 Bit, 44.1Khz, in Mono.

Label	**Type of morph**	***% softness***	***% brightness***	***% everything else***	***% warmth***
Source	*Everything (Source)*	0	0	0	0
E15	*Everything*	15	15	15	15
E30	*Everything*	30	30	30	30
E50	*Everything*	50	50	50	50
E70	*Everything*	70	70	70	70
E85	*Everything*	85	85	85	85
E100	*Everything (Target)*	100	100	100	100
W15	*Warmth*	0	0	0	15
W30	*Warmth*	0	0	0	30
W50	*Warmth*	0	0	0	50
W70	*Warmth*	0	0	0	70
W85	*Warmth*	0	0	0	85
W100	*Warmth*	0	0	0	100
EEW15	*Everything except warmth*	0	0	15	0
EEW30	*Everything except warmth*	0	0	30	0
EEW50	*Everything except warmth*	0	0	50	0
EEW70	*Everything except warmth*	0	0	70	0
EEW85	*Everything except warmth*	0	0	85	0
EEW100	*Everything except warmth*	0	0	100	0

Table 1 Stimulus set used in the pairwise dissimilarity comparison. Labeling corresponds to that used in the subsequent MDS analysis, and also intended type and amount morph, where Bx indicates an intended brightness morph of x% towards the target sound, W an intended warmth morph, and EEW an intended everything-else-except-warmth morph.

The experiment was conducted from a Macintosh Macbook Pro laptop with Sennheiser HD650 circumaural headphones, in a small room with a dry acoustic and reasonable acoustic isolation. 20 listeners, each with some experience of critical listening and sound recording, took part in the test. Before undertaking the main experiment, listeners were presented with a familiarisation exercise, in which they could listen to the full range of stimuli in a non-linear fashion for as long as they felt necessary before beginning each step. Each listener compared each element in the stimulus set to each other (and to itself), through 190 randomly ordered pairs, split over two listening tests of approximately 25 minutes in duration. Listeners were asked to grade pairs of stimuli using a hidden 100 point continuous scale with end-points labelled 'Not at all similar' and 'The Same', as shown in Figure 4.

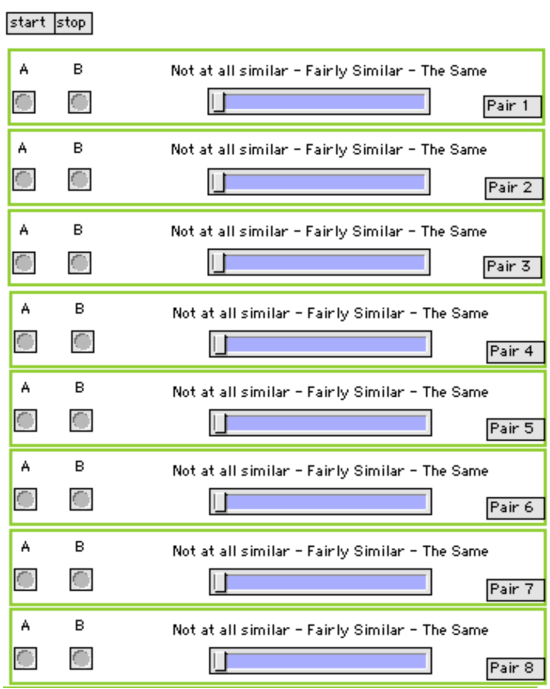

Figure 4 A section of the Max/MSP listener interface, showing two stimuli per slider. Listeners were asked to rate how similar A sounded to B using the sliders

3.1.1. Pairwise dissimilarity experiment results

The listener responses were used to create a dissimilarity matrix, and analysed by Multidimensional Scaling (MDS) INDSCAL (9) analysis, to reveal the number of dimensions of timbral variation within the stimulus set. The 'measures-of-fit' calculated by the MDS analysis, as shown in Table 2 were then examined.

Dimensionality	**RSQ**	**RSQ improvement to next dimension up**	**s-stress**
1-D	.91689	0.05110	.11481
2-D	**.96799**	**0.03019**	**.02373**
3-D	**.99818**	0.00177	.02067
4-D	**.99995**	0.00005 (max)	.00336

Table 2 Dimensionality against RSQ and 's-stress' values in MDS solutions of the pairwise comparison data. Bold dimensions/criteria indicate a statistical quality criterion has been met. Maximum possible RSQ improvement at 3-D is given by 1-(3-D RSQ).

It has been previously shown that if RSQ, or squared correlation factor, does not improve by more than ~0.05 per dimension, this is a good indicator that the optimal fit for the data has already been found, and that an RSQ of more than 0.95 indicates a confident solution (10). S-stress of under 0.1 is considered to indicate a good solution (11)

The 1-D solution shown in Table 2 satisfies none of the statistical quality criteria, with high s-stress, low RSQ, and significant improvement between 1-D and 2-D. Both 3-D and 4-D satisfy RSQ and s-stress statistical quality thresholds (>0.95 and <0.1 respectively). However, 2-D also satisfies both of these criteria as well as satisfying the third critierion - an RSQ improvement of less than 0.05 - indicating that any further improvement in the 3-D and 4-D solutions is likely to be noise. The 2-D solution (confident, good, and optimal), is therefore the best fit to the data. The corresponding 2-D timbre space is shown in Figure 5.

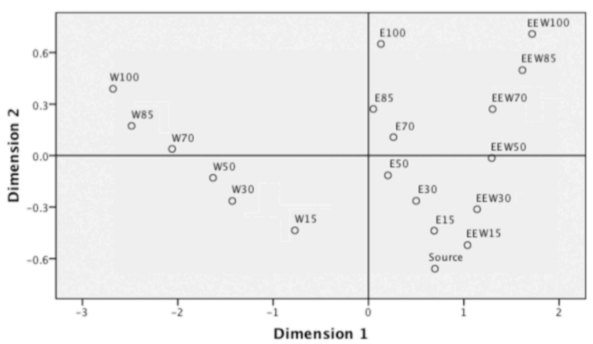

Figure 5 2-D timbre space showing perceived movement of warmth against everything and everything else except warmth hybrids.

Warmth and everything else stimuli appear in broadly orthogonal directions, indicating that listeners perceived the morphs in these directions to be independent. The everything stimuli appear in between warmth and everything except warmth. All directions and spacings are broadly as intended, evidencing the perceptual unidimensionality of the new morph types.

3.2. Verbal elicitation experiment

As with all MDS analysis, the names of the dimensions were not revealed, and a subsidiary verbal elicitation experiment was undertaken to establish whether listeners perceived the morpher to be manipulating the intended timbral attributes. A verbal elicitation experiment, and subsequent verbal protocol analysis (VPA) was undertaken in order to label the changes created by the new morphing algorithm. The listeners were asked to listen to pairs of stimuli again and to fully describe the perceived differences between the two stimuli in each pair using as many adjectives as they felt necessary. All three timbral attributes were included in the stimulus set, so that listeners would have an opportunity to remark on the complete algorithm, including brightness morphed stimuli. This was particularly important as brightness and warmth might be expected to exhibit a degree of overlap but listeners did not compare warmth with brightness in the pairwise dissimilarity experiment. The stimulus set is shown in Table 3. Stimuli A1-A3 differed only in terms of the intended warmth dimension, stimuli B1-B3 only in terms of intended brightness, stimuli C1-C3 in everything except warmth, brightness and softness, and D1-D3 in intended softness.

Label	**%** ***softness***	**%** ***brightness***	***% everything except warmth, brightness and softness***	**%** ***warmth***	**MDS label**
A1	0	0	0	0	Source
A2	0	0	0	50	W50
A3	0	0	0	100	W100
B1	0	0	0	0	Source
B2	0	50	0	0	B50
B3	0	100	0	0	B100
C1	0	0	0	0	Source
C2	0	0	50	0	EE50
C3	0	0	100	0	EE100
D1	0	0	0	0	Source
D2	50	0	0	0	S50
D3	100	0	0	0	S100

Table 3 Stimulus set for verbal elicitation experiment. 'MDS label' corresponds to the stimuli used in 2.1. Dimensions are denoted by Ax, Bx, Cx, and Dx

Participants were again presented with a Max/MSP interface on a Macintosh laptop, with the same circumaural Sennheiser HD650 headphones in the venue used in the pairwise dissimilarity experiment. 10 listeners, 2nd year undergraduate students with a similar background in audio engineering and critical listening to the pairwise dissimilarity listening panel, undertook the test.

3.2.1. Verbal elicitation experiment results

Listener responses were grouped together based on their meaning (for example, synonyms antonyms and commonalities), by an independent academic, who had no prior knowledge of the experiment's aims. The number of listener responses in each group was summed and an overall prominence, indicating the perceptual importance of each group, was calculated by dividing the number in each group by the total number of responses.

The groupings, number of occurrences, and overall prominence for each group are shown in Table 4, Table 5, Table 6, and Table 7.

Groups of similar adjectives used to describe dimension A	**Number of occurrences**	**Overall prominence (occurrences / total)**
Warm, bassy, wide, fat, thick, rounded	11	92%
Punchy, harsh, aggressive, loud	1	0.08

Table 4 Verbal elicitation groupings for dimension A

Groups of similar adjectives used to describe dimension B	Number of occurrences	Overall prominence (occurrences / total)
Bright, clear, sharp	9	76%
Warm, bassy, wide, fat, thick, rounded	1	8%
Filtered, thin, hollow, dry	1	8%
Punchy, harsh, aggressive, loud	1	8%

Table 5 Verbal elicitation groupings for dimension B

Groups of similar adjectives used to describe dimension C	Number of occurrences	Overall prominence (occurrences / total)
Hollow, filtered, thin, dry	4	50%
Warm, bassy, wide, fat, thick, rounded	3	37%
Punchy, harsh, aggressive, loud	1	13%

Table 6 Verbal elicitation groupings for dimension C

Groups of similar adjectives used to describe dimension D	Number of occurrences	Overall prominence (occurrences / total)
Soft, smooth, gentle, slow, mild, quiet	10	0.91
Punchy, harsh, aggressive, loud	1	0.09

Table 7 Verbal elicitation groupings for dimension D

The most prominent dimensions, based on the overall prominence values, were;

- Dimension A – **Warm**, bassy, wide, fat, thick, rounded
- Dimension B – **Bright**, clear, sharp
- Dimension C – **Hollow**, Filtered, thin, dry
- Dimension D – **Soft**, smooth, gentle, slow, mild, quiet

In dimension A, the overall prominence value for the intended attribute, *warmth*, is the largest of the groups by a good margin. However, this leading group also has the largest number of terms within it of any of the groups given by the independent academic.

In dimension B, the intended attribute, *brightness*, is again the most prominent of the groups, though by a smaller margin than that seen in the responses to dimension A. Also in the most prominent group in this dimension are found sharpness and clearness, both of which have been reported to be correlated to spectral centroid. The second most prominent group in this dimension features *warmth* amongst its descriptors, but there were no occurrences of *warmth* itself amongst the listener responses to Dimension B

Dimension C was not intended to have a specific label, but rather to be *everything except warmth, brightness* and *softness*. Nevertheless, the reason for selecting source and target sounds with specific spectral characteristics was such that timbral changes in *everything except* would be clearly perceivable as a change in harmonic ratio, correlated with *hollowness*. *Hollowness* is the most prominent group of this dimension, but is only more prominent by the smallest margin of any of the leading groups of descriptors in dimensions A, B, C and D. As with Dimension B, the second most prominent group in this dimension also features *warmth* amongst its descriptors, but there were no occurrences of *warmth* itself amongst the listener responses to Dimension C.

Finally, dimension D, similarly to dimension A, has a very large margin between the leading group and the next most prominent group. This leading group includes the intended attribute, *softness*, and also has a large number of other terms. It is worth noting that the academic included 'quiet' within this group – this could indicate that either the listener might have chosen quiet as a synonym for *softness*, or perhaps that the current loudness equalization employed was not appropriate for all listeners.

The prominence of *warmth* in dimension A, *brightness* in dimension B, *hollowness* in dimension C and *softness* in dimension D indicated that listeners heard the corresponding stimulus sets as changing across the intended dimensions. However, it was possible that other timbral attributes, perhaps noted in less perceptually-important groups, could have been changed in parallel with the intended ones. Parallel changes would not have shown up as additional dimensions in the MDS analysis. A quality control experiment was therefore devised.

3.3. Quality Control Experiment

This experiment required listeners to note any perceived inter-stimulus differences that were not adequately described by the intended descriptors. The stimuli for this experiment included both the morphed sawtooth and triangle wave sounds of the pairwise and verbal experiments, and real-world morphed sounds – an overdriven guitar, electric bass, and organ – to confirm the morpher's ecological validity.

The full stimulus set, including synthesised source and target sounds, is shown in Table 8.

Stimulus	Source Sound	Target Sound	Morph Type
1	EG	Tr	Softness
2	EG	BG	Warmth
3	EG	Or	Hollowness
4	BG	EG	Brightness
5	BG	Or	Hollowness
6	BG	Tr	Softness
7	Tr	EG	Brightness
8	Tr	Or	Hollowness
9	Tr	Tr	Warmer
10	Or	Tr	Softness
11	Or	EG	Brightness
12	Or	BG	Warmth
13	Saw	Tri250	Brightness
14	Saw	Tri250	Warmth
15	Saw	Tri250	Hollowness
16	Saw	Tri250	Softness
17	EG	N/a	N/a
18	BG	N/a	N/a
19	Tr	N/a	N/a
20	Or	N/a	N/a
21	Saw	N/a	N/a
22	Tri250	N/a	N/a

Table 8 Complete stimulus set for Quality Control experiment. EG represents the electric guitar sample, Tr the trumpet, Or the organ, and BG the bass guitar. Saw and Tri250 are the sawtooth source signal and triangle with soft attack envelope from previous experiments. 17 – 22 are unmorphed input signals

Listeners were presented with a paired comparison, and asked whether they agreed with the intended timbral descriptor. They were instructed as follows: "Compare the two sounds in each numbered pair. If you feel that the timbral label assigned to describe the change from the left to right stimulus is incorrect or incomplete, please enter as many alternative or additional descriptors as you wish."

Listeners who had undertaken previous experiments were not invited to take part, to avoid bias or pre-suggestion which might be caused by previous experience with the stimulus set. 10 listeners, 1st year audio engineering undergraduate students with some experience of critical listening undertook the test. Each listener was invited to repeat the test after a short break, to provide a reference for the evaluation of listener consistency, if required.

3.3.1. Quality Control experiment results

There were 6 additional responses amongst a total of 166 responses. This could be interpreted as a ~4% disagreement with the allotted timbral descriptors. Of that 4%, none of the original descriptors was removed. All of the 6 additional descriptors were directly related to loudness (i.e. "louder" or "quieter").

Amongst the repeated experiment responses, there were 7 additional (of 167 total) responses, again a ~4% disagreement with the allotted timbral descriptors. Again, none of the original descriptors was removed, and the additional responses comprised only loudness-related descriptors.

On comparing each of the listeners that had given additional responses, listener 5 and listener 7 in the initial results, and listener 3 and listener 5 in the repeat results, it became clear that these listeners were not consistent in their additional responses, suggesting that the loudness discrepancies were very small. Nonetheless an improved loudness equalisation procedure would be useful further work. However, since loudness, by definition, is not a timbral attribute, these results indicate that no additional timbral changes were perceived. Together, these results confirm that no unintended timbral changes were perceived in the stimulus set, and that the large number of terms found in the leading groups in the verbal elicitation experiment can be considered synonyms for the intended attributes, as reliable listeners introduced no additional attributes. The experiment also confirmed that the algorithm is ecologically valid and functions on real-world sounds.

4. CONCLUSIONS

A routine for selectively manipulating *warmth* was proposed, such that a blanket spectral region between the fundamental frequency in a sound and a multiple of 3.5 thereof could be selectively amplified or attenuated in order to increase or decrease perceived warmth. This amplification/attenuation would result in changes to the spectral centroid, the spectral slope, and the amplitude of the first three harmonics (or, in an inharmonic tone, a spectral area performing a similar timbral role).

The new warmth routine was combined with the previously-developed brightness and softness morphing processes to create a complete system intended to compensate for any unwanted movement in brightness that might be caused by morphing warmth. It was anticipated (and later proven) that warmth morphing would not cause any unwanted change in softness, due to its temporal, rather than spectral, acoustic correlation. Fourteen possible morph types could be generated by the full warmth-brightness-softness morphing system.

A pairwise dissimilarity experiment was undertaken, such that warmth-morphed stimuli were compared with everything-morphed and everything-except-warmth-morphed stimuli. MDS analysis of listener responses to this experiment indicated that there was variation in 2-D. Warmth was plotted orthogonally from everything-except-warmth in the 2-D timbre space, evidencing the unidimensionality of the two morph types.

A verbal elicitation experiment then showed that listeners perceived the orthogonal dimensions to be broadly as intended – but they may have heard changes in other dimensions as well, as many participants used additional descriptors.

The subsequent quality control experiment showed that, apart from loudness related issues, there were no additional changes perceived when morphing through the range of timbral attributes. The additional descriptors elicited in the verbal experiment could therefore be considered to be synonymous with the intended ones. The quality control experiment also confirmed that the morphing algorithm works with real-world sounds, demonstrating its ecological validity. The potential for bias caused by the prescribed descriptors in the quality control experiment means that the prior use of the MDS-VPA based experiments is still essential.

Together, this combination of experiments demonstrated that independent morphing of a third timbral attribute, with an acoustic correlate overlapping that of an attribute already included in the morpher, was possible and that this approach to timbre morphing, including the method for dealing with such overlap, is viable.

In the future, it should be possible to follow this revised combination of experiments to evaluate the inclusion of new attributes within the timbre morpher. However, it would be prudent to limit the number of attributes being varied independently in the VPA and QC stages if many attributes are included in the morpher. This would still form a robust combined validation process, with MDS confirming the dimensionality, VPA confirming the appropriate labeling, and the QC stage highlighting any unintended overlap with other attributes.

5. REFERENCES

1. *Perceptually-motivated audio morphing: Brightness.* **Williams, D and Brookes, T.** Vienna : Audio Engineering Society, 2007. Audio Engineering Society Convention 122. p. 7035.
2. *Perceptually-motivated audio morphing: softness.* **Williams, D and Brookes, T.** Munich : Audio Engineering Society, 2009. Audio Engineering Society Convention 126. p. 7778.
3. *Automatic Audio Morphing.* **Slaney, M, Covell, M and Lassiter, B.** Atlanta, GA : IEEE International Conference on Acoustics, Speech and Signal Processing, 1996.
4. *Toward the Perfect Audio Morph? Singing Voice Synthesis and Processing* . **Cook, P.** Proceedings of the Digital Audio Effects Workshop DAFX, 1998.
5. *A Spatial Processor for Musicians and Sound Engineers.* **Marc-Jot, J and Warufsel, O.** CIARM – International Conference on Acoustics and Musical Research, Ferrara, Italy, 1995.
6. *A subjective rating scale for timbre.* **Pratt, RL and Doak, PE.** 3, Journal of Sound and Vibration, 1976, Vol. 45.
7. *Timbre of steady sounds: A factorial investigation of its verbal attributes.* **Von Bismarck, G.** Acta Acustica, 1974, Vol. 30.
8. *SeaWave: A system for musical timbre description.* **Ethington, R and Punch, B.** 1, Computer Music Journal, 1994, Vol. 18.
9. *Maximum likelihood estimation in multidimensional scaling.* **Ramsay, JO.** Psychometrika, 1977, Vol. 42.
10. *Humanised Statistics: An investigation of social values in a senior secondary school milieu* **Astill, BR.** Australian Association for Research in Education (AARE) Conference. University of Adelaide, 1994.
11. *Multivariate Data Analysis.* **Anderson, R, et al.** New Jersey : Prentice Hall, Upper Saddle River, 1998.