



University of HUDDERSFIELD

University of Huddersfield Repository

Rice, Thomas

An Investigation into Non-linear Distortion Pedals with a Focus on Establishing a Lexicon to Accurately Describe Them.

Original Citation

Rice, Thomas (2020) An Investigation into Non-linear Distortion Pedals with a Focus on Establishing a Lexicon to Accurately Describe Them. Masters thesis, University of Huddersfield.

This version is available at <https://eprints.hud.ac.uk/id/eprint/35393/>

The University Repository is a digital collection of the research output of the University, available on Open Access. Copyright and Moral Rights for the items on this site are retained by the individual author and/or other copyright owners. Users may access full items free of charge; copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational or not-for-profit purposes without prior permission or charge, provided:

- The authors, title and full bibliographic details is credited in any copy;
- A hyperlink and/or URL is included for the original metadata page; and
- The content is not changed in any way.

For more information, including our policy and submission procedure, please contact the Repository Team at: E.mailbox@hud.ac.uk.

<http://eprints.hud.ac.uk/>

An Investigation into Non-linear Distortion Pedals with
a Focus on Establishing a Lexicon to Accurately
Describe Them

Thomas Rice

A thesis submitted to the University of Huddersfield in
partial fulfilment of the requirements for the degree of
MA by Research

The University of Huddersfield

January 2020

i. The author of this thesis (including any appendices and/ or schedules to this thesis) owns any copyright in it (the “Copyright”) and s/he has given The University of Huddersfield the right to use such Copyright for any administrative, promotional, educational and/or teaching purposes.

ii. Copies of this thesis, either in full or in extracts, may be made only in accordance with the regulations of the University Library. Details of these regulations may be obtained from the Librarian. Details of these regulations may be obtained from the Librarian. This page must form part of any such copies made.

iii. The ownership of any patents, designs, trademarks and any and all other intellectual property rights except for the Copyright (the “Intellectual Property Rights”) and any reproductions of copyright works, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property Rights and Reproductions cannot and must not be made available for use without permission of the owner(s) of the relevant Intellectual Property Rights and/or Reproductions.

Contents Page

Abstract	4
Chapter 1: Introduction	4
Chapter 2: Literature Review	5
Chapter 3: Methodology	11
Chapter 4: Text Mining	16
Chapter 5: Descriptive Survey	21
Chapter 6: Definitions Panel	27
Chapter 7: Similarity Matrix	31
Chapter 8: Dendrogram and Distortion Wheel	34
Chapter 9: Listening Tests	38
Chapter 10: Spectrograms	44
Chapter 11: Total Harmonic Distortion	51
Chapter 12: Brightness	55
Chapter 13: Roughness	58
Chapter 14: Spectral Flux & Centroids	62
Conclusion	67

List of Figures

● Fig.2.1.1- Ibanez TS-9 with Schematic	7
● Fig.2.1.2- Boss DS-1 with Schematic	7
● Fig.2.1.3- ProCo RAT with Schematic	8
● Fig.2.1.4- Electro Harmonix Big Muff with Schematic	8
● Fig.4.1.1- Word Frequency Bar Chart- Total	16
● Fig.4.2.1- Word Frequency Bar Chart- TS9	17
● Fig.4.3.1- Word Frequency Bar Chart- DS-1	19
● Fig.4.3.2- Word Frequency Bar Chart- RAT	19
● Fig.4.3.3- Word Frequency Bar Chart- Big Muff	20
● Fig.5.1.1- Word Frequency Bar Chart- Survey	22
● Fig.5.1.2- Word Cloud- Survey	23
● Fig.8.1.1- Dendrogram	34
● Fig.8.1.2- Distortion Wheel	36
● Fig.9.1.1- Boxplot- 'Thick'	39
● Fig.9.1.2- Boxplot- 'Fat'	40
● Fig.9.1.3- Boxplot- 'Thin'	41
● Fig.9.1.4- Boxplots- 'Crispy' & 'Sludgy'	41
● Fig.9.1.5- Boxplot- 'Dirty'	41
● Fig.9.1.6- Boxplots- All Remaining	42
● Fig.10.1.1- Spectrogram- Low Gain	45
● Fig.10.1.2- Spectrogram- Medium Gain	47
● Fig.10.1.3- Spectrogram- High Gain	49
● Fig.11.1.1- THD- Clean Amplifier	51
● Fig.11.1.2- THD- TS9	52
● Fig.11.1.3- THD- DS-1	52
● Fig.11.1.4- THD- RAT	53
● Fig.11.1.5- THD- Big Muff	54
● Fig.12.1.1- Brightness- TS9	55
● Fig.12.1.2- Brightness- DS-1	55

- Fig.12.1.3- Brightness- RAT 56
- Fig.12.1.4- Brightness- Big Muff 56
- Fig.13.1.1- Roughness- Clean Amplifier 58
- Fig.13.1.2- Roughness- TS9 59
- Fig.13.1.3- Roughness- DS-1 59
- Fig.13.1.4- Roughness- RAT 60
- Fig.13.1.5- Roughness- Big Muff 61
- Fig.14.1.1- Spectral Flux- Clean Amplifier 62
- Fig.14.1.2- Spectral Flux- TS9 63
- Fig.14.1.3- Spectral Flux- DS-1 64
- Fig.14.1.4- Spectral Flux- RAT 65
- Fig.14.1.5- Spectral Flux- Big Muff 66

List of Tables

- Table 6.1- Table of definitions 29-30
- Table 7.1- Similarity Matrix 31
- Table 11.1- Table of THD values 51

Abstract

This thesis attempts to identify and investigate the auditory impact distortion pedals impart on the source signal and consequently, how best to define, discuss and classify their individual, and collective, sonic signatures. This includes establishing a specialist lexicon with recontextualised descriptions for the specific adjectives of focus through means of both qualitative and quantitative experimentation. This information is then visually represented by a 'distortion wheel' based in principle off the SCAA's Coffee Tasting Wheel and how that allowed for the accessible retrieval of specialist terms within the contextual field. This is achieved through a series of etymological and audio-based analysis and experimentation. Qualitative experimentation was used to discover the initial descriptor list, critical sonic variables, and to subsequently define the words; quantitative to match respective audio and signal analysis to the linked adjectives. The results showed 'crunchy' to be the most commonly used distortion descriptor. Through further analysis, it can be concluded that the Ibanez TS-9 is the crunchiest distortion pedal as its sonic features match closest to the defining traits of 'crunchy'. Distortion pedals are subsequently successfully classified and descriptions for each adjective created to give the completed distortion lexicon.

Ch.1- Introduction

When Grady Martin's broken pre-amplifier led to his guitar tone taking on a, now ubiquitous, distorted tone at a studio session for Marty Robbins in 1961, he had accidentally stumbled on an effect that later, millions would attempt to emulate (Kosser, 2006). Merely a year later, the session engineer, Glenn Snoddy, would help sell the idea of this effect in pedal format to the Gibson Guitar Corporation, actively contributing to the production of the first widespread 'fuzzbox', the Gibson Maestro FZ-1 Fuzz Tone. Snoddy and Hobbs' germanium transistor circuit, powered by two 1.5v batteries became somewhat pioneering as an entire generation of guitarists clamoured to obtain this radical new auditory effect. With Keith Richards utilising the FZ-1 in the Rolling Stones 1965 hit 'I Can't Get No (Satisfaction)' (1965, track 7), its influence continued to grow exponentially.

Now, nearly sixty years after Martin's accidental initial foray into distortion, it is an effect that is commercially saturated and widely available to musicians of any level. In the timescale since its inception, four specific pedals have grown in stature, to the point of iconicism, and established themselves as benchmarks of the effect: Ibanez Tube Screamer, Boss DS-1, ProCo RAT and Electro Harmonix Big Muff.

Ch.2- Literature Review

According to the Cambridge Dictionary, distortion in an audio context can be defined as: ‘a change in or loss of sound quality, due to changes in the shape of the sound wave’ (“Distortion”, 2020). These “changes” are nonlinear in nature and will therefore often be referred to as nonlinearities. The debasing connotations of “loss of” are indicative that for many applications, distortion is viewed in a negative light; an unwanted discordance. This is countered by the statement that: “subtle settings can be useful for gentle enhancements and more trashy settings for more drastic results” (Izhaki, 2007). While this details the spectrum of audible changes that can be achieved through the use of distortion, the casual vernacular of ‘trashy’ doesn’t stand up to academic scrutiny as it lacks quantifiable merit; assuming it implies higher setting choices/more gain, this has many applications in a wide array of stylistic boundaries and shouldn’t be marred by negative, subjective semantic choices. The inference that perceived proximity, as a result of harmonic distortion imbuing high frequency energy, can be created through the use of distortion builds upon this notion. Zargoski-Thomas’ suggestion that “intense high frequency content can be used to make something seem closer than the loudspeaker it emanated from” (Zargoski-Thomas et al., 2012) shows distortion has many practical uses as a creative tool.

As these uses became increasingly realised among musicians, many turned to distortion pedals to help shape their instrumental timbres in search of more extreme stylistic boundaries. As rock evolved into metal, distortion became intrinsically linked with the perceptual attributes associated with the genres. The assertion that heaviness is inherently tied to the sound of the distorted rhythm guitar (Berger, 1999) demonstrates the sheer extent of distortion’s influence on genre definition. Theory and Analysis of Classic Heavy Metal Harmony (Lilja, 2009) validates these notions, as well as expanding upon them. Of special interest was the assertion that power chords were intrinsically defined by the presence of distortion. This was pertinent due to demonstrating that distortion has permeated so deeply into the fabric of popular music that basic musical notions are now viewed as incomplete without the addition of nonlinear distortion. Lilja also touches upon the additive harmonic overtones generated by nonlinear distortion and how this creates musical complexity- both specific to power chords and within general usage. This study overall provides insight into practical usage of distortion and its contemporary place within popular music theory, specifically metal and its subgenres.

Many of Jan Herbst’s papers fall into the same thematic domain as my own research. Distortion and Rock Guitar Harmony: The Influence of Distortion Level and Structural Complexity on Acoustic Features and Perceived Pleasantness of Guitar Chords (2019) directly relates to Lilja’s aforementioned research into power chords and the effect of distortion on their makeup. The expansion that establishes that due to the additional harmonic content generated by nonlinear distortion, the majority of listeners respond to the sound quality more than the harmonic structures creates an interesting tangent. This also sets up the argument that the auditory perception of distorted chords is conflicting.

Herbst's indication that familiarity will influence perception was astute and worthwhile analysis. While the research in *Heaviness and the Electric Guitar: Considering the Interaction Between Distortion and Harmonic Structures* (Herbst, 2018) provides analysis on the perception of distortion, it seems an oversight to regard distortion as a single entity rather than as a hugely changeable effect, with each iteration being slightly different from the last. This is especially surprising given the focus on metal and its subsequent subgenres whose musicians were some of the first to extensively utilise distortion pedals, often 'gain stacking' (Total Guitar, 2020) with multiple devices in attempts to achieve increasingly heavy, dense sounds. Metal pioneers such as Dave Murray of Iron Maiden were utilising devices such as the MXR Distortion + to achieve more intense gain levels than their peers could produce simply through amplifier distortion (WoodyTone, 2010). If distortion is intrinsically linked to the perception of heaviness, then a positive correlation should exist between increased distortion and increased heaviness; however, no sense of scale is conveyed in this particular study.

The first study on electronic design I found to be particularly pertinent to this research was 'Analog Musical Distortion Circuits for Electric Guitars' (Sunnerberg, 2019). This thesis provides a comprehensive overview of analog distortion circuit topologies; the relevancy extends from focusing solely on analog distortion devices, a specificity shared by this thesis also. One weakness one could identify would be the lack of validation to justify the chosen sonic descriptors- 'Fizzy' particularly seems very loosely defined and perhaps lacks academic merit to describe audible traits of distortion, without previous academic validation. However, the apparent lack of material attempting to quantify these terms was the rationale for the research in this thesis. Overall, this paper provides thorough analysis of various distortion circuits, including differentiating and comparing germanium diode limiters and class B push/pull nonlinearities. This research provided a solid electronic foundation upon which to base my own semantic research and justified the rationale to attempt to quantify the audio descriptors so commonly used to describe auditory distortion.

Building upon these ideas is 'Design and Construction of Arduino-Hacked Variable Gating Distortion Pedal' (Murthy et al., 2014). This research is more specific in its approach but provides analysis of general distortion circuits, nonetheless. This includes an articulate breakdown of the components comprising a 3-stage distortion circuit. The intricate explanation of component purpose in analog distortion circuits helped inform the electronic knowledge for the research conducted in this thesis.

The four featured pedals in this thesis (Ibanez TS-9- Fig.2.1.1, Boss DS-1- Fig.2.1.2, ProCo RAT- Fig. 2.1.3 and Electro Harmonix Big Muff- Fig. 2.1.4) can be seen here with their respective circuit diagrams alongside:



ProCo Rat Distortion

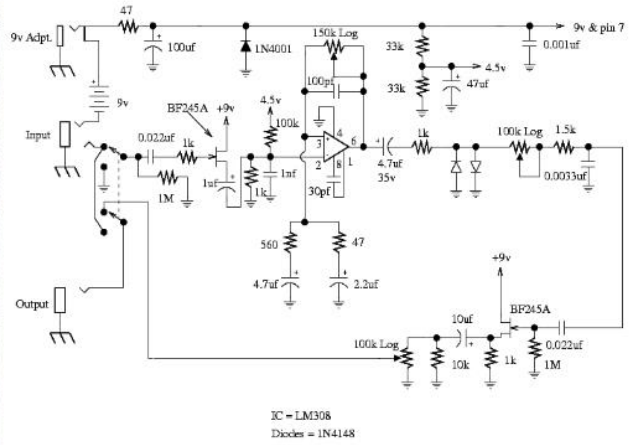


Fig.2.1.3- ProCo RAT with Schematic

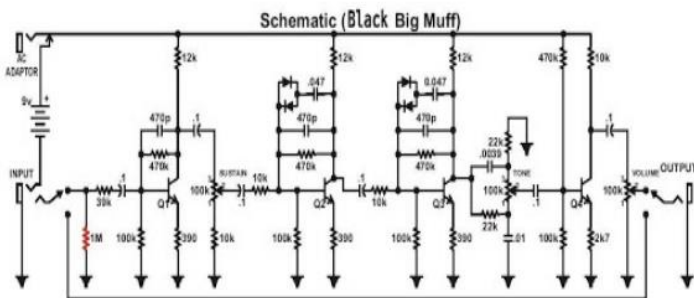


Fig.2.1.4- Electro Harmonix Big Muff with Schematic

'Analysis, Synthesis and Classification of Nonlinear Systems using Synchronized Swept-Sine Method for Audio Effects' (Novak et al., 2010) provided more of a mathematical foundation to the approach, informing both methodology and general background knowledge on nonlinearities. As the title suggests, this research demonstrates one method of data retrieval this thesis attempted to emulate, although much of the research was more mathematically advanced than was required for this thesis. The research's "capacity to distinguish both kinds of nonlinear systems through its ability to synthesize the output signals from any given input signal" (Novak et al., 2010) was the specific draw, in relation to the planned experimentation for this thesis. Coincidentally, the ST-9- one of the featured pedals in this paper, is a 4-knob variant of the Ibanez TS-9 (the earlier iteration), prominently featured in this thesis.

Moving from circuit design and topology into audio analysis, The Effect of Harmonic Overtones in Relation to 'Sharpness' for Perception of Brightness of Distorted Guitar Timbre (Tsumoto et al., 2017) is arguably the closest related research field and study to the one explored in this thesis. The final line of the study validates the research conducted here: "Future research should include the categorizing of distortion". The methodology provided a framework for the audio analysis conducted in this study and the conclusion that distortion can often induce a 'suppression effect' that acts like a dampener for the brightness of the signal, created expectations of what the results in this research may produce.

The paper that naturally succeeds Tsumoto's (2017) was also of specific relevance; Timbre of Nonlinear Distortion Effects: Perceptual Attributes Beyond Sharpness (Marui & Martens, 2005). The identification of thin-thick, sharp-dull, and dark-bright as perceptual antonyms is fairly rudimentary but a good indicator for the validity of this thesis. Analysis of the control settings was welcome alongside explanation of how these match to distortion descriptors. Discussion on the lack of understanding in general users of how settings relate to circuitry changes was also of interest, for example that 'drive' controls alter the level of electrical current fed into the central circuit.

Structured Models for Semantic Analysis of Audio Content (Chaudhuri, 2013) provided an excellent foundation for topics involving semantic analysis. The chapters involving 'Beyond Acoustic Unit Descriptors' provided useful, especially its discussion on polysemy and how semantic choice may not directly match to their acoustic counterparts being described.

Following on from this, the slightly more contextual The Semantic Space of Sounds (Pedersen, 2008) focuses on establishing lexicons and classification which proved particularly pertinent. The quantitative classification, although slightly different from what is attempted in this study, proved useful from a methodology standpoint and in its explicit similarity to the goals of the research conducted within this thesis. This was a far bigger lexicon than was attempted in this research with some lexis seeming barely relevant but a comprehensive list existing was positive to find.

The final relevant source is An Investigation into the Sound Quality Lexicon of Analogue Compression using Category Analysis (Ronan et al., 2015) which also tied closely

to the research being attempted in this thesis. These 3 audio semantic studies together formed the foundation of my understanding which enabled the research conducted regarding a distortion lexicon. The multidisciplinary approach was shared as it produced succinct quantitative and qualitative results.

Ch.3- Methodology

Many of the techniques utilised by the researchers in the aforementioned academia above, formed the foundation for the methodology used in this research. This was with the understanding that research of this nature has no predefined universal procedure by which to gather data. The approach, therefore, was inherently multidisciplinary. This included using both quantitative (objective) means of testing, alongside qualitative (subjective) to gauge both general opinion and then attempt to substantiate that with numerical data to draw comprehensive conclusions. These methodologies include electronic testing, content analysis and signal analysis as means of gathering quantitative data; alongside grounded theory, timbre studies and musicological elements to ascertain quantitative opinion and data.

The use of grounded theory would be the academic underpinning for this entire thesis. The intent was to devise a set of terms, redefined to be contextually accurate, that would form a lexicon for auditory distortion. There were no formal preconceived notions that altered the approach to any of the experiments and so the value of grounded theory became apparent early in the research (Chun Tie et.al., 2019). The fact it is an inductive method, leading to conclusions not being drawn until the end of the study and based on retrospective observations made throughout the entire process, linked closely with the intent of the entire research and subsequent thesis. The adaptability of grounded theory also made it an attractive proposition for the research; the methodology had room to grow and change to best fit where the research was heading, given there were no preconceived notions of intended direction or outcomes. Theoretical sampling allows for conclusions to be constantly revised throughout the research process, informed by the changing nature of the data gathered. This fits perfectly with the investigatory nature and lack of an open hypothesis within the entire thesis and therefore was deemed the most suitable theoretical methodology. This method will be embellished by the inclusion of Empirical Discourse Analysis (EDA), the purpose of which is to critically analyse the function of language in social or genre-specific settings (Phillips et al., 2002). This process heavily informs the etymological facets of the research.

Establishing the key communicative language used by musicians to describe the different sonic properties of distortion was the first step of the methodology. Content analysis was deemed as the most academic way of collecting this data due to its ability to systematically summarise the written communication in a quantitative manner (Hsieh, 2005). This was achieved through text mining, the purpose of this was to gather data from as large a number of sources as possible; a comprehensive list of every example of language used to describe distortion pedals online, regardless of relevancy. To begin, a data retrieval method called 'web scraping' (Persson, 2019) was used to obtain large amounts of semantic content which would subsequently enable later analysis. This involves deploying a simple piece of software that copies and stores any text present in any given body of material. The chosen material was then gathered from a multitude of sources ranging from peer-reviewed academia to informal online reviews of each of the four chosen pedals. This was deliberate for the simple reason that: desired was as expansive a vocabulary as could realistically be

obtained, as it would give a far greater indication of the populist colloquialisms used, as well as the academic equivalents; greater semantic variety gives greater credence to the results of the frequency analysis. This was designed to be a relatively fast process due to the lack of necessity for direct human contact. Many of the other qualitative methods of data collection used in this thesis required individual sessions collecting or testing for data with other musicians. This is time consuming and so to establish the fundamental lexical body and provide an academic foundation for the remaining research, the content analysis through text mining was designed to retrieve data in a timely manner. The program of choice for analysing the collective semantic body was RStudio. Adjectives (descriptors) were the only words of value to the research, unfortunately however, no software (within RStudio or externally) could be found that could lexically analyse and separate the data based purely on word-type. Therefore, the only remaining solution was to manually sift through the entire lexical catalogue and siphon out the relevant descriptors. The data for the selected descriptors was then amalgamated into a new dataset to be re-analysed which allowed me to calculate specific word frequency, to then be visually represented in bar graph and word cloud formatting. The data gathered during this part of the research was of critical importance as it informed every other stage of the methodology. By the end of this process, the goal was to have obtained the initial list of adjectives that would eventually be analysed and converted into the first iteration of the auditory distortion lexicon.

To further expand upon the lexical body of descriptors established in the 'text mining' procedure, a cross-sectional descriptive survey was devised. This was also to cross-reference and clarify any results gained previously. A different type of testing should incur different responses, if any results were repeated it validated their existence as a potential part of the upcoming lexicon and any new descriptors unveiled were of specific interest as either anomalous or contextual to the nature of the test. Each subject would provide 3 distinct descriptors for 12 independent pieces of audio- each of the four pedals recorded with low, medium and high gain settings (all other controls levelled). The subjects were also given a choice between hearing the pedal audio through single-coil or humbucker pickups, depending which they were more aurally familiar with. Alongside reaffirming the popularity of certain descriptors uncovered in the 'text mining' analysis, this survey aimed to help identify certain idiosyncratic, anomalous adjectives that emerged in reference to specific audio sources also, influenced by specific settings on a specific pedal perhaps. Its results would also aid in providing context for descriptors discovered in the previous stage. The data from this survey would be analysed individually, as well as amalgamated into a collective body of descriptors with the 'text mining' results to be analysed as a whole definitive lexicon. This deliberate separation (and consequential combination) was to preserve the integrity of the context of the data; language used when communicating generally online is not necessarily the same language used when prompted with a specific audio cue. The total results create a better dataset to be analysed, as there would simply be more results whereas individual allows the context by which they were retrieved to be factored into any results or conclusions drawn.

The next stages involved calculating word similarity. This cannot be accurately ascertained qualitatively and so a means of gathering numerical data was devised. The most

frequently occurring descriptors were entered into a matrix which allowed participants to score (out of five) each descriptor against the most prevalent variables: soft and hard clipping, low and high gain, increased and decreased sustain, increased and decreased dynamic consistency & increased and decreased bass, mid and treble responses. From these scores, averages could be calculated, and this subsequently provided numerical data, turning previously qualitative data into quantitative. The variables had been whittled down from an initially larger list and were selected based on applicability and perceived size of sonic influence. Variables such as volume, nature of clipping (symmetrical or asymmetrical) and level of harmonic content were all decided to be too diffuse to be relevant at this stage of the research. The remaining fourteen are all critical variables that directly affect output and timbre of their respective device. Trends within the results were predicted at this stage due to some of the variables being inherently linked, gain and clipping for example and so correlation between scores for different variables was expected. The average scores provided by the matrix offered insight into which sonic variable(s) was most influential on choice of descriptor, which was subsequently used to generate the dendrogram. Average scores of 4.0 or above would indicate defining traits, a variable that gives a pedal its individual sonic character. In the same sense, scores of 1.0 or below can be viewed as variables that have little to no influence on the sonic characteristics of their respective pedal. This is the first stage that previously subjective opinions, even generally, can be verified in an objective sense. The data gathered during this process is massively important to the remaining experiments, the data is visualised by the dendrogram immediately after, but also provides an objective measure to compare with the results of the quantitative signal analysis. If the most applicable variable's traits are replicated in the objective findings from analysing the pedals, it adds credence to both findings, despite being obtained by radically different means of experimentation.

The dendrogram provided an easy way of visually representing these similarities. The average scores for each field were calculated manually and then the newly generated matrix of averages was inputted into the statistical analysis software RStudio. With this dataset in place, the software calculated the similarity between descriptors and rendered the dendrogram as a visual representation of this. This graphical representation of lexical similarity provides an accessible format to view how the descriptors are grouped, based on assigning shared importance to similar sonic variables. The dendrogram adds another quantitative layer onto material originally gained through entirely qualitative means, giving its results increased traction. Its accuracy however will be defined by the results of the similarity matrix; results that generally agree to any decent extent will see an accurate dendrogram be generated which accurately groups and classifies the descriptors together. However, if the results intrinsically show signs of disagreement between participants, the dendrogram may not accurately represent the links between descriptors based on their definitive variables. This highlights the one negative of using mean averages which is that anomalous results can skew otherwise exceedingly accurate data, and a potential flaw that must be considered, no more accurate means of testing have been uncovered however and so that risk will be accepted. Visualising the data will aid accessibility the thesis offers as well, to the non-academic discerning reader, a graph depicting word similarity will be far more palpable than the matrix beforehand. To further expand upon the ease of access to

nuanced information offered by the dendrogram, a wheel chart, similar to the SCAA's 'Coffee Tasters Flavor Wheel' (1995), was created to easily map each descriptor to its representative variable. Anticipating that certain variables would have more extreme numerical scores, regardless of descriptor, than others, the final sorting procedure was a manual one. This allowed a final check that generally, each descriptor was well represented by the classification it found itself placed into and gave a chance to edit and change these if necessary. Once each of the top twenty most frequent descriptors had been matched with one of the fourteen sonic variables, where needed, the remaining spots would be filled with contextual synonyms, discussed as part of the final panel session. Once completed, this diagram would then depict what each descriptor's defining variable was, in a visually pleasing, accessible format.

The information provided by the dendrogram, from the results of the similarity matrix, also helped identify key variables, from which new, recontextualised lexical definitions could be drawn. To attempt to reduce bias and subjectivity being present within these definitions, a group of eight musicians were gathered to form a panel whose sole aim was to collectively redefine the descriptors. Initially provided were traditional definitions, etymological origins and relevant literary information. This was to provide context for each individual descriptor, so each panel participant felt comfortable with the existing definition and context for each lexi, before helping recontextualise and redefine them. Through perpetual revision over three separate sessions, and using these pre-existing definitions alongside the contextual material, a new set of definitions would be created. Interestingly, despite their colloquial prevalence in the context of audio properties, none of the descriptors had previous definitions outside of traditional contexts which added to the importance of this phase of the research. These three sessions took place bi-weekly over a six week period, deliberately planned so different pieces of research from the different experiments would be ready with each passing session, culminating in the matrix scores being finished for the final meeting which allowed for validation of existing definitory ideas, re-evaluation if an important trait was missed in discussion and correction if a false assertion had been agreed upon previously. This research method was arguably the most subjective of all methodologies explored in this research; there is no possible means of validating any definition created hence having multiple sessions and constant revision to try and collectively reach as close to the general consensus when defining each of the descriptors. These definitions alongside the distortion wheel form the etymological finale of this thesis; The most common contextual descriptors, classified by respective definitory sonic traits, presented with recontextualised definitions.

The final element of subjective, qualitative experimentation was in the form of a listening test. This involved subjects independently ranking the individual pedals against the chosen descriptors-complete with newly re-appropriated definitions. Presented with four audio stimuli for each descriptor, using a sliding scale, participants ranked each stimulus, out of ten, relating to descriptor applicability. The numerical averages drawn from this study, then allow cross-comparison with the objective analysis later, to compare which audio features are represented by which lexical choice. The data gathered from these audio tests allows for more specific descriptive analysis of each pedal; the information should inform

further, based on more specific audio cues, which sonic traits prompt specific lexical choice. This is yet another attempt to convert subjective opinion into objective data; the retrieval method is still highly opinionated, however gaining numerical scores for each pedal helps transpose this into quantitative data. With the gain becoming its own external variable in this experiment, the differences in gain for each pedal and the subsequent changes to the way the lexicon is used will only be revealed in detail during this stage. The addition of numerical data on the subject also will prove useful as qualitative data can then be transformed to quantitative with rankings able to be drawn, again adding nuance to the newly developed lexicon. As the final phase of testing before the signal analysis, the results of the listening tests and dendrogram are the culmination of every experiment beforehand. Established the most common descriptors, recontextualised and defined them and then classified them. The listening test adds the individual specificity which brings the focus back onto the featured pedals whereby numerical data is being gathered, using the results from the etymological stages of the research. This provides data that can be compared with the results of the signal analysis hence the importance of the results of the listening test.

Extensive objective analysis was then undertaken to capture as many analytical facets of the audio as possible, obtained from each gain setting of each pedal was: waveforms, frequency spectrums, spectrograms, spectral flux, spectral centroid, RMS & Low Energy, brightness and roughness graphs. These provided the objective counterpoint to the existing subjective data, allowing observation and measurements on how accurately the descriptors matched the audio features they attempt to describe. This objective experimentation provides the substantiation for many claims made beforehand. The tangible measure of the audio properties of the pedals allows validation and perspective on the etymological side of the thesis; it presents the data to cross reference every claim or assertion that emerged in any previous experiment. The data being gathered however, is far more extensive and detailed than is needed simply to complete that task, it also allows for close examination of the changes to the source signal that the pedals impart, which in turn, becomes its own strand of research in identifying these key sonic alterations, individual to each pedal.

Ch.4.1- Experiment 1- Text Mining

The first task was to discover the most frequently used adjectives that players used to describe the sonic properties of distortion pedals. The best way to ascertain this was by sifting through large amounts of relevant text and analysing which adjectives appeared most frequently. Over 500 examples of contextual material were analysed, the sources ranging from peer reviewed academia, through to informal online reviews. The language and semantics deployed in, and between, those 2 extremes of literature would obviously be massively varied and for the analysis to be as definitive as possible, as many bodies of text as possible were included. This material had to be online, in some format, to fit the research methods. Manually extracting and regurgitating written text into a digital format would have been too laborious and time-consuming a process to be considered worthwhile.

RStudio computed all the amalgamated text data and then also performed the frequency analysis which helped quantify many of the terms, in respect to applicability, based on how often they were used. When grouped together, these adjectives became the lexicon, and the individual words- 'descriptors'. Any word featured even once was included in the amalgamated lexicon as these might present interesting anomalous data or relevant synonyms that could later be utilised. Once the definitive lexicon had been established, the semantic body was subdivided by contextual source, first by pedal. These new subsets were then re-analysed to discover pedal-specific descriptors. Once the total and specific word frequency analysis had been completed, there was no computation program available that would distinguish type of adjective and so manual adjustment was required. The only adjectives of interest were describing the sonic properties, and nothing more. Lexis such as 'heavy', 'huge', 'vintage', 'classic' and 'versatile' were used frequently, however, these are not exemplifying nuanced language, given the topic, as they are all highly subjective and not indicative of any specific sonic trait.

This initial list of the twenty most frequent descriptors became a useful asset throughout the remaining research, as well as serving as a general guide to the most common language used when discussing distortion pedals. This list is displayed as a bar chart in Fig. 4.1.1. 'Crunchy' became the focus of much speculation as it emerged as the most frequent descriptor generally, and specifically, and the rationale behind its selection provided some

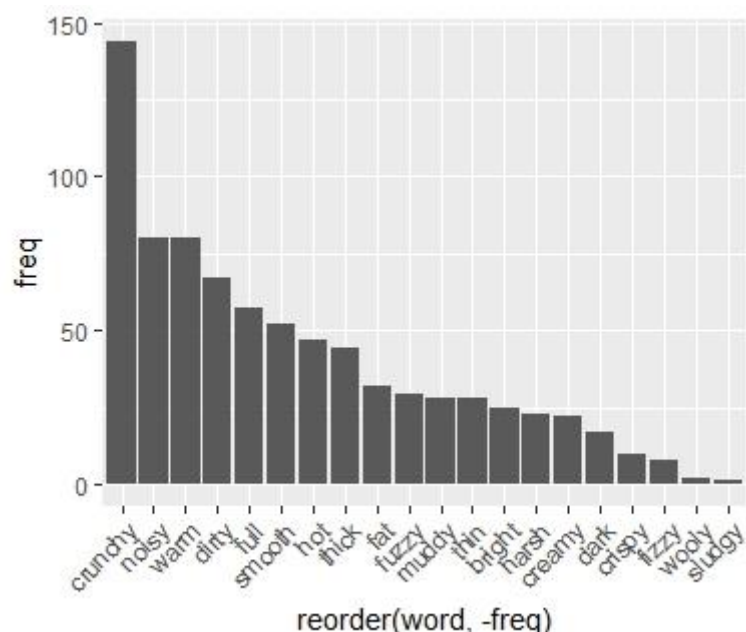


Fig.4.1.1- Word Frequency Bar Chart- Total

mystery. The popularity of ‘noisy’ also proved to be of interest, although it was suspected that its positioning amongst the most frequent descriptors was linked to focused usage in specific contexts. The suspicions lay that match of the usage of ‘noisy’ would be directed towards specific pedals, or specific traits only present in select devices. This is confirmed by the results of the combined word frequency (including survey data) which shows ‘noisy’ emerging as the most frequent descriptor used to describe the sonic properties of the Boss DS-1. The general word frequency results were pleasing as they corroborated various predictions made prior to research beginning. None of the twenty descriptors were surprising or inappropriate results and have all been heard, colloquially, in topical discussion.

Ch.4.2- The ‘Crunchy’ Paradox

By far and away the most popular descriptor overall, topping the frequency charts for both the Ibanez Tube Screamer (shown in Fig. 4.2.1) and (after extrapolating word endings to improve interpretability) the DS-1 and RAT also, was ‘crunchy’.

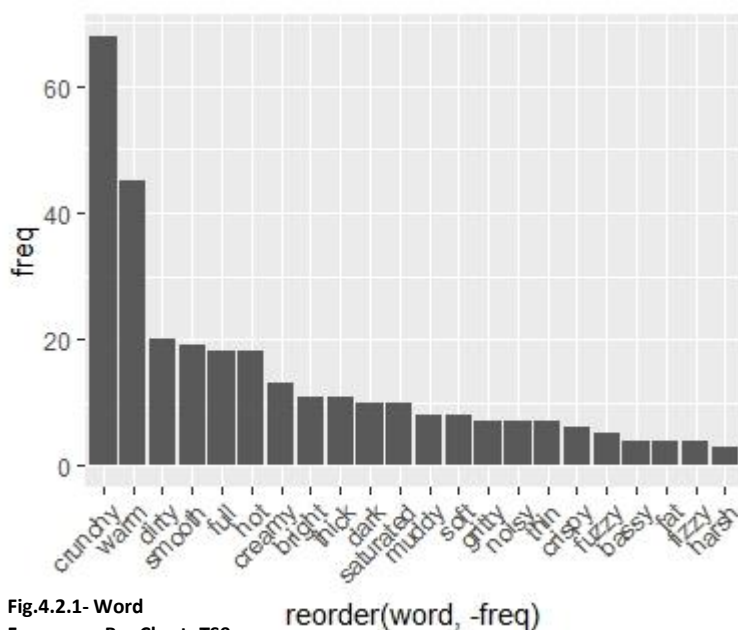


Fig.4.2.1- Word Frequency Bar Chart- TS9

This is a word that will be widely recognised by guitarists of every ability level and one commonly used when referencing guitar tone(s). It is also one of the most highly subjective in definition and overtly difficult to quantify. This, therefore, is the paradoxical element: why is a descriptor without contextual definition and so individually subjective, used so frequently, as opposed to more specific terms? Similar terms, such as ‘dirty’, ‘creamy’ and ‘sludgy’

emerged semi-frequently but nowhere near the frequency that ‘crunchy’ was used. Perhaps pedal manufacturers also spotted this colloquial adjective’s marketability potential, which could be evidenced by the increase of the word ‘crunch(y)’ in the way pedals are marketed. Joyo, a Chinese Pedal manufacturer even included the term in the name of one of its flagship distortion units- Joyo Crunch Distortion. Usage of the word in these contexts will only further validate the usage of the word in many eyes while still not providing any clarification on contextual definition.

The term is a great example of onomatopoeia which potentially goes a long way to

explaining the populist preference over more technical terminology. The 'ch' digraph in 'crunch' has obvious phonetic ties to higher levels of clipping and gain, sonic variables synonymous with overdrive and distortion; The phonetic pronunciation of the voiceless postalveolar affricate vocally mimics the clipping and compression of a distorted signal. Other examples of onomatopoeic terms that emerged in the text mining phase of the research include: 'raspy', 'fizzy', 'boomy' and 'growly'. These alternative results substantiate the idea that colloquial, easy to comprehend, onomatopoeic terms are often used, in lieu of an established terminology or contextually specific lexicon. It seems very human that when faced with a dilemma of having never been presented with formal, descriptive terms to describe the auditory features of distortion, we resort to primitive, colloquial language to increase levels of comprehension and understanding when attempting to communicate. Alongside the literature review, which revealed a surprising lack of material focusing on distortion pedals, even in a less academic sphere, there is little mention or clarification on contextual definitions for these frequently used lexical choices. This heavily implies that many people using these terms are doing so in such a highly subjective sense that its usage is constantly open to interpretation; there are currently no universal defining traits that, in regards to shared definition, link individual usage; every subjective interpretation could conceivably be entirely unique, even if only minutely. In the case of the Big Muff, 'crunchy' was only deployed three times. The Big Muff, of all the pedals, features the most extreme levels of clipping, almost pushing the signal to a square waves at times; the negative correlation, therefore, with the frequency of 'crunchy' implies there is a threshold, within clipping level, to the word's applicability.

Ch.4.3- Text Mining Cont.

Aside from 'crunchy', the most frequent descriptors for each respective pedal do a remarkably accurate job of describing the very-general sonic properties of each pedal, for the: TS9- 'Warm', DS-1- 'Noisy', RAT- 'Dirty' and Big Muff- 'Fuzzy'.

Generally, the Tube Screamer is seen as the least aggressive of the collective and so the choice of 'warm' from both a phonetic and descriptive perspective seems fitting. Equally popular descriptive terms for the Tube Screamer included: 'smooth', 'full' and 'creamy'. 'Smooth' corroborates the notion of perceived calmness (relative to the other three units), somewhat supported by the high ranking of 'full' also, which pertains more to sonic depth and consistency than any gain or clipping traits. Slightly less popular yet still notable are the inclusion of 'muddy', 'thin' and 'dark' which seemingly imply the lower gain levels present in the TS9 may obfuscate the tone. Also noteworthy was the mention, albeit uncommon (<10), of intrinsically negative descriptors such as 'noisy', 'fizzy' and 'harsh'. Occurrences of these adjectives were all documented under ten times and they could be seen as anomalous due to no tonal control consistency or gain level separation being observed during this stage of the research.

The Boss DS-1 has always notoriously had elements of transient noise present, especially at louder volumes or with the gain control set high, which perhaps explains the popularity of ‘noisy’ to emerge as the most frequent descriptor used in relation to its sound. Fig. 4.3.1 visualises the DS-1’s results in full. Many adjectives emerged that seemed to hint at the presence of negative sonic additions. It is often difficult to pinpoint transient noise and subsequently, how it is semantically represented is also often a

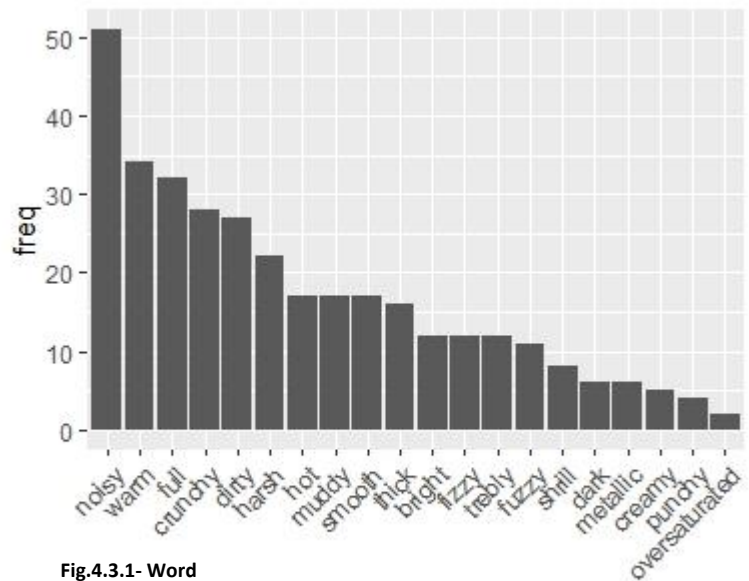


Fig.4.3.1- Word Frequency Bar Chart- DS-1

reorder(word, -freq)

challenge. Outside of sonic idiosyncrasies, the popularity of ‘harsh’, ‘hot’ and ‘muddy’ further reinforces the impression that this pedal holds certain traits that many find subjectively untenable. A slightly inappropriate result, however one that also holds relative importance, especially pertaining to general consensus and the confirmation of collective distaste, was the alarming inclusion, in 13 individual instances, of the adjective ‘crappy’ (plus another 12 instances of ‘crap’) in specific relation to Boss’ flagship pedal. ‘Warm’ was the second most popular descriptor for the DS-1 but suspiciously emerged far more regularly than any other similar terms or synonyms; ‘dark’ for example was used less than 5 times as

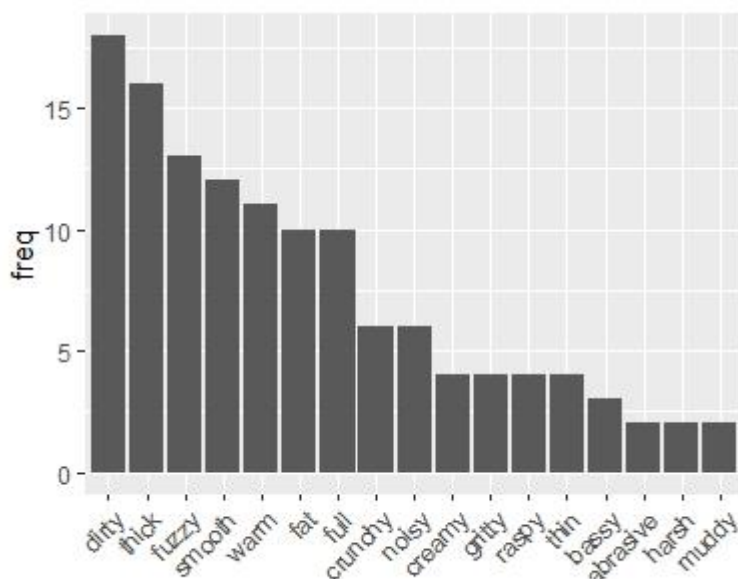


Fig.4.3.2-Word Frequency Bar Chart- ProCo RAT

reorder(word, -freq)

often. This irregularity, coupled with the emergence of the term ‘oversaturated’, a stark contrast to notions of warmth, lends credence to the idea that ‘warm’ could have been used out of context or alternatively, is simply anomalous.

One of the other previously discussed onomatopoeic terms, ‘dirty’, was the most common descriptor for the ProCo RAT, fitting its simplistic reputation as one of the more aggressive distortion units. ‘Dirty’

naturally has obvious connotations of impurity and when placed into the contextual realm of distortion, this heavily implies that either high levels of gain or clipping are present, and influential on the tone of the RAT. Fig. 4.3.2 displays the emergence of ‘dirty’ along with the remaining list. Interestingly, in more vague contextual spheres, ‘dirty’ is an intrinsically

negative word, whereas when referring to a distorted sound, its usage is generally used to indicate a positive aspect of the timbre. This extends (and fits) with the basic idea of distortion as a controlled nonlinearity; one is deliberately impurifying the signal to create a radically different, subjectively positive, change in sound. However, that act of intentionally muddying the source signal matches the lexical choice and subsequent definitory traits of 'dirty'.

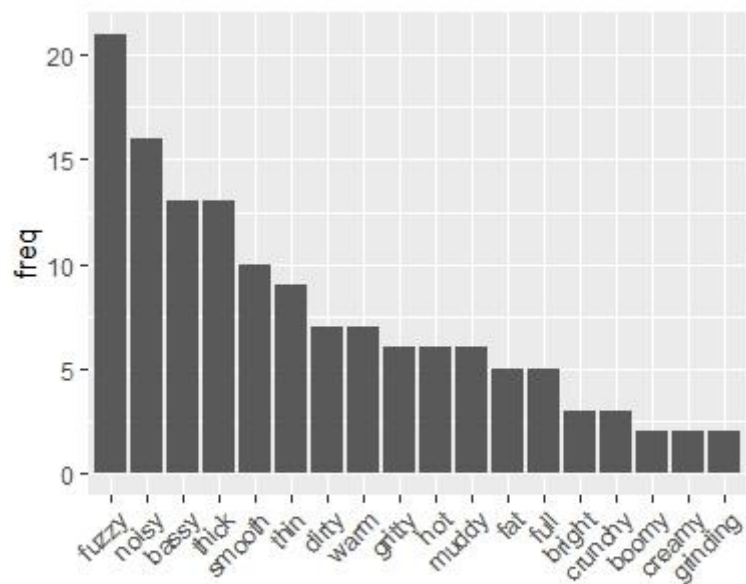


Fig.4.3.3- Word Frequency
Bar Chart- EHX Big Muff `reorder(word, -freq)`

Emerging as the most popular

lexical choice therefore validates notions that the RAT's effect impurifies the signal to a greater extent than some of the other featured distortion units. Closely trailing, 'thick', 'smooth', 'warm' and 'fat' however, all infer that ProCo's seminal pedal also regulates the levels of gain behind a filter combination that creates an objectively pleasant distortion palette. On top of this, the regularity of 'thick' and 'fuzzy' being used descriptively heavily infers higher levels of clipping to be present when using the RAT. Conversely, the lack of support for 'abrasive' and 'harsh' imply that while subjectively grating, sonically, for 2 individuals, this isn't a popular view supported by their peers. Compared to the DS-1 however, any remotely negative descriptors were relegated to single figure usage and so are treated as less representative, especially considering the top 5 descriptors for the RAT are all inherently positive sonic traits.

Predictably, the most frequent descriptor for the Big Muff was 'fuzzy', leaning into the misnomer that Electro Harmonix's creation is a fuzz pedal, as opposed to the Distortion/Sustainer it was created, and labelled, as. The inclusion of the descriptor 'grinding' (twice) was of passing interest as this also links with the much higher levels of clipping present in the Big Muff, than can be heard in any of the alternative three. It demonstrates that many individuals will extend their vocabulary outside of popular terms in attempts to accurately describe specific sonic traits, demonstrated in Fig. 4.3.3. The same could be applied for 'boomy' and similarly, 'bassy' emerged as the third most popular descriptor for the Big Muff. This result is of special interest when considering the omission of any similar terms; no 'dark' and 'warm' substantially lower. The choice of a colloquialism in itself is interesting but when represented so highly, in regard to frequency, it becomes of special interest. Also, of particular note is the difference in frequency between the previous 3 pedals and the Big Muff, for the term 'crunchy'. The fact that word frequency is hinting there are discernible sonic differences within the big miff that separate it from the others is explored in later elements of the research.

Ch.5- Extended Descriptor Survey

After amalgamating the results from the text mining phase of the research into a single body of descriptors; the second phase involved creating a survey which would ideally validate the existing results and present new data to analyse, also. The critical difference this time, however, was this decision-making process was to be prompted by audio samples of the pedals, focusing on aural perception to classify rather than unprompted, potentially irrelevant general feelings. Twenty-six samples were recorded- low, medium and high gain settings for each pedal on single coil and humbucker pickups respectively as well as the clean amp tones for both. These samples consisted of DI'd recordings of the electric guitars (two- one for each pickup type) playing a simple ostinato centred around open chords, designed to be instantly familiar in regards to voicings, which were then re-amped through the respective pedal at the specified gain setting. Sonic familiarity was of a fair importance as stimulation was predicted to yield far more interesting lexical results than if the subjects were unfamiliar with the sounds they were being exposed to. Shure Sm57s were the microphone of choice therefore, as their heavy usage within popular music production since their introduction makes the sound of their recordings intrinsically familiar to musicians. With a similar rationale, a Fender Stratocaster and Gibson Les Paul were the guitars of choice by way of their iconicism and therefore recognisability. This was all an effort to leave the pedal, and subsequent settings, as the focal point as opposed to any other sonic distraction. Participants then provided 3 separate descriptors that they thought aptly described each respective audio sample. There was an initial disclaimer that specified the descriptors given should be in the sonic domain and not referencing irrelevant factors such as cost, aesthetic, genre-usage. This was to keep the results as apt as possible. It should be noted however, that not every result fit these criteria. Certain responses included multiple word phrases and even profanity. It was decided the two instances of phrases would be omitted from the results, however due to the singular instance of profanity matching results explored in the text mining stage of the research, that would be allowed to remain.

The survey was then shared online and completed by twenty-six participants internationally. Twenty-four were guitarists and only two were not. The first three preliminary questions allowed participants to specify this and if they answered 'yes', how long they had played. These pre-screening questions were necessary given the otherwise anonymous nature of the survey. The software Qualtrics allowed for easy sharing on social media which facilitated the international responses.

The results proved what had been expected; they validated certain words that emerged frequently in text mining, the likes of 'crunchy', 'smooth' and 'warm'; however it also unearthed totally new descriptors that were added to a definitive list of every adjective found pertaining to auditory distortion. These newly found descriptors were often more lexically complex and generally less generic, these include adjectives such as: 'throaty',

'brassy' and 'growly'. These terms were often much harder to initially quantify within the context of auditory distortion but colloquially understood fairly easily. 'Growly' for example is obviously denoting a moderately aggressive clipping and/or gain level but within the context of the other variables that affect the tone of a distortion pedal, it is something of an unknown quantity. This was the case for many similar descriptors that emerged during this stage of the research. It should be noted that this descriptive survey yielded some inappropriate results which were filtered out immediately after the experiment had ended and the results were being filtered manually. These included whole phrases and terms that did not accurately fit the brief- 'expensive' for example when cost had been explicitly listed as not being a factor when choosing semantics. These inappropriate results were not overtly common however and when viewed in the context of the entire body of results, could certainly be seen as anomalous.

Similarly, to the text mining results, I staged the data as both word frequency graphs (Fig. 5.1.1) and then in word cloud format (Fig.5.1.2). These were rendered for both the individual survey results and as part of the combined, total descriptor body.

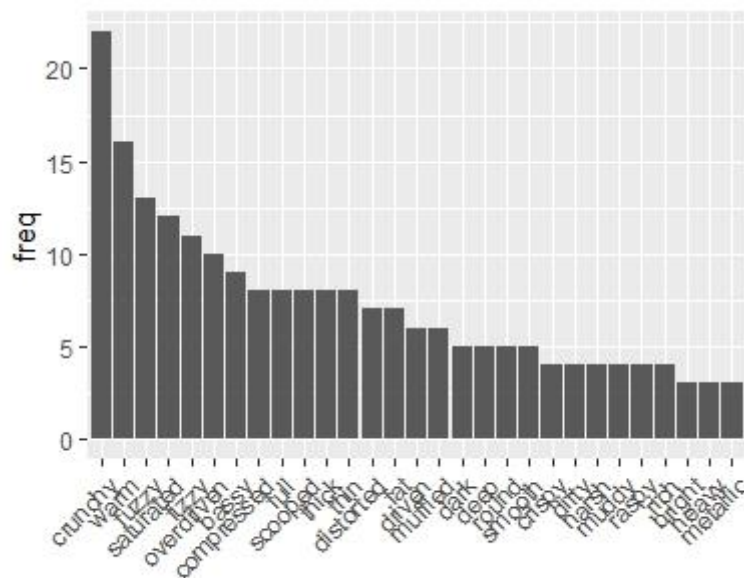


Fig.5.1.1- Word Frequency
Bar Chart- Descriptor Survey
reorder(word, -freq)

The high frequency of 'crunchy', 'warm' and 'fuzzy', given their mutually emergent popularity from the text mining, was welcomed; it corroborated the notion that these were terms used very often to describe varying sonic properties of the pedals. This confirmation between the first two experiments was certainly encouraging in its explicit hints towards the validity of the data they presented. The introduction of lexical choices such as 'scooped' proved to be interesting results; in this specific instance its usage was limited to medium and high gain settings. Within this contextual sphere, the idea of scooping is most closely related to the mids equalisation band. Mid scooping was popularised in the 1980s and 90s mainly within the metal scenes, the guitarists of which also almost exclusively favoured higher gain settings. This is the act of heavily reducing, or entirely nullifying, the mid response of your amplifier so as to produce a shriller tone with only the subsequently

pronounced bass and treble frequencies being audible. This was a control setting preferred by many prominent guitarists such as Dimebag Darrell from Pantera (Hodgson, 2016).

‘Scooped’ being used within this survey either demonstrates a sonic link between the guitar



Fig.5.1.2 (above)- Word Cloud- Descriptor Survey

tone(s) heard within this experiment and their similarity to sounds used within metal during the period where mid scooping was most apparent, or, a description of the equalisation that occurs within the filter combination of each of the respective pedals, which may denounce the mid response. The specificity of which cannot be determined by the data gathered from this descriptive survey. Only the objective, quantitative analysis can determine the outcome of which sonic traits are informing descriptor choice categorically.

Due to these results being directly prompted and not manually sorted, the results may seem more implicit or cryptic. The word cloud allows for better comparison between the 2 extremes as the results are more closely grouped, literally, subsequently the outliers exist in a physical space next to the popular descriptors. New adjectives that emerged such as: ‘articulated’ and ‘indistinct’ and ‘bloated’ seemed more concise and indicative of a higher level of descriptive specificity. ‘Articulated’ and ‘indistinct’ seem to hint towards elements of clarity, or in the latter’s case the lack thereof. This implies that some of the unit’s effects obfuscate the signal to the point of being murky, sonically, whereas alternate units do little to obscure the fundamental lucidity of the tone. ‘Bloated’ however isn’t quite as obvious in the definitory sense. According to the Cambridge Dictionary, its traditional definition is: “swollen and rounded because of containing too much air, liquid, or food” (“bloated”, 2013). This is contextually irrelevant for the most part, however. Ideas of excess are raised in ‘too much’ alongside ideas of saturation in ‘swollen and rounded’. Therefore, ‘bloated’ could be considered as adjectively describing sonic oversaturation.

Beyond this, there were certain terms that emerged anomalously that were semantically interesting enough to be worth exploring. Firstly, ‘brassy’ and ‘tubular’; both of these lexes have obvious etymological ties to higher level of volume. ‘Brassy’ being imitative of brass musical instruments, known for their projection capabilities and often harsh, piercing timbres. ‘Tubular’ is less obvious as there are potentially contextual links to tube amplifiers which would be highlighting the sonic differences between solid state amplifiers and tube amplifiers. Considering, however, that every audio sample was recorded through a tube amplifier including the non-affected clean amplifier signals that were just as reference points for every stage of the testing, technically every sample, regardless of pedal effect, could be described as tubular. Therefore, both these terms could be interpreted as being

descriptive of the more rotund timbres produced by distorted tones through tube amplifiers, as opposed to their 'tinnier' solid state counterparts. In the general contextual sphere of distortion, it is hard to imagine the deployment of these terms outside of traits directly influenced by amplifier type. Consequently, they could be considered secondary descriptors as they do not directly pertain to the sonic signature of the pedals themselves with their output being the exclusive variable in lexical choice.

Another set of seemingly linked anomalous descriptors that emerged from the word frequency graphs were 'precise' and 'unrefined'. The latter has a traditional definition of: "not processed to remove impurities or unwanted elements" ("unrefined", 2010). In stark contrast, "'precise" can be defined as: "marked by exactness and accuracy of expression or detail" ("precise", 2010). While linked by descriptive domain, the usage of these two adjectives will be radically different. 'Precise' brings up ideas of controlled clipping and gain where the internal filter combinations, likely without excessive gain or clipping amounts, can tame the overall effect to the point of sonic customisability. Precision then can be obtained with the ability to exactly obtain the desired distortion timbre. Juxtapositionally, the implication of using 'unrefined' suggests elements of the sound lack clarity in the sought-after areas. This descriptor could be seen as something of a misnomer due to the fact that none of the featured samples could be described as unprocessed, with the obvious exception of the clean amplifier reference tones. Within the specific contextual field of auditory distortion however, the idea of refinement could easily be linked with a decrease in the level of impurities; an idea also explored by 'precise' hence the noteworthy link. This is problematic because the nonlinearity in itself could be considered an impurity, however, retrospectively viewing the results from the text mining elements of the research, its results give indication that impurities exist beyond the distortion effect. The Boss DS-1's most popular descriptor was 'noisy', yet ranked much lower for each of the respective, alternate pedals. This indicates that aside from the extra harmonic content generated by the effect of each pedal, there are noticeable sonic irregularities that impact on how each pedal is described. The idea of refinement, and subsequently 'unrefined' could therefore be considered as a description of how impactful these impurities are on the overall signal of each given device. 'Unrefined' is therefore more valid, the more impurities are audible when the distortion effect is engaged; the more unwanted noise, the more unrefined the output is.

The idea of power is explored within elements of the objective analysis, but it is also a subject that emerged within semantic choice in the survey. As the word cloud shows, both 'underpowered' and 'overloaded' were mentioned as relevant contextual adjectives to describe the sonic properties of different distortion pedals. Power isn't a variable that is overly common within this domain, especially in colloquial conversation. This is perhaps due to the fact that powering these devices is normally so easy as to be overlooked. The majority of distortion units are powered by a standard 9V power supply. This can be supplied either directly from the mains or from an auxiliary power supply, usually used to power multiple pedals. Since most power units are now fully isolated (alongside the 240V mains supply, naturally) the issue of pedals being under or over-powered isn't usually a common occurrence, nor is it a discernible trait for the pedals themselves to have as it is caused by additional, external appliances rather than the unit itself. Powering multiple

pedals from the same source is a common reasoning for an individual unit to be underpowered and, arguably, the most common circumstance in which a pedal becomes overpowered is by accidentally using an incorrect adapter which can also inadvertently ruin the internal circuitry of the pedal, or in the case of clever designs, blow a capacitor designed to act as a circuit-breaker saving the rest of the components from the surge. Anecdotally, some also claim that deliberately overpowering certain pedals can give their timbre a substantial positive edge over the recommended voltage. The determining factor in which way a pedal will respond to excessive charge is the presence of an electrolytic capacitor, and its respective voltage. The best quote to summarise this idea comes from an anonymous online commenter: "There are pedals that people claim sound better at 18v, and there are pedals that will explode" (reddit, 2015). These changes in power supply do affect the sonic properties of the pedals alongside the physical. According to Sweetwater (2019) "Changing from 9 to 18 volts can give you a little more headroom and may also change the tone, but this is by design". Therefore, 'overloaded' could be interpreted as being descriptive of the additive headroom created by using an excessive voltage than specified for the respective unit. On the other hand, leaving your voltage with too little power, either through an incorrect supply or dying battery, often leaves distortion units with a 'spluttering' effect that some find subjectively desirable. This starves the transistor's bias creating the effect, which is seemingly limited to distortion or fuzz units. Consequently, 'underpowered' carries the connotation of a pedal having this unintentional, 'spluttering' timbre. This unique effect is sought-after enough, that some manufacturers have started including inbuilt options to enable this effect deliberately, rather than requiring faulty equipment. This is called 'sag' or a 'dead battery' effect. The Voodoo Lab PP2, for example, has a rotational control that allows the user to set the voltage from 4V up to 9V, directly setting the level of sag present in the signal.

Another interesting outlier was the descriptor 'digital'. Since the instructions for the survey explicitly specified that any adjectives given should be descriptive of purely the sonic qualities, this semantic choice will only be considered within this domain, aside from circuitry type in which its usage is far more common. This is referencing an argument familiar to most pedal enthusiasts which is the age-old analogue versus digital debate. The earliest pedals were entirely analog before the 1980s introduced more complex digital circuitry. The most famous of these early digital circuits was the BOSS DD-2, released in 1983, one of the first delay units to move away from the iconic bucket brigade chips that made the early pedals so retrospectively valuable. Digital circuitry allowed for fast processing of more complex audio mechanisms, using the aforementioned DD-2 for example, the change to digital circuitry allowed for much longer delay times than Boss' previous effort the analog DM-2. As time progressed, the advancements in digital audio processing grew exponentially and the pedals within which the technology was implemented, only expanded in sonic capability. Distortion however, remained largely untouched by this new wave of digital technology; this is due to digital distortion imitations largely paling in comparison to the genuine, analogue product. Digital distortion became most prevalent in multi-effects units and in 'modelling' technologies. One of the few examples that exist as a standalone distortion unit is the Digitech DF-7. This is also an example of modelling technology which is a digital circuit replicating the sound of another

(usually analog) pedal as accurately as possible. These digital replications are infamous for never quite achieving the desired goal however, and as such, the usage of 'digital' as a sonic descriptor in this context is predictably intrinsically negative. The imitations usually lack the sonic depth, especially in the lower mid or bass frequencies, of their analog counterparts and so 'digital' could be suggesting a shallowness to the timbre of the distortion pedal in question.

The final two anomalous descriptors worth mentioning are distinctly more subjective than any previous examples: 'balanced' and 'controlled'. The traditional definitions provide some insight, for 'balanced' - "a situation in which different elements are equal or in the correct proportions" ("balanced", 2018), and 'controlled' - "having been limited in intensity or level; kept in check" ("controlled", 2010). These can be grouped together by shared definitory similarity. It's worth noting initially that the settings of each of the respective pedals, along with pedal choice itself, directly affects the applicability of both of these descriptors. Cranking the gain settings, for example, on any of the pedals drastically reduces the chances that these adjectives will be suitable. However, within the individual character of each device lie sonic traits that make ideas of balance and control more accurate than others. Looking at the results of the text mining for the DS-1, the prevalence of 'noisy' suggests that neither of these two descriptors would necessarily fit the sonic signature of Boss' pedal. The more subdued Tube Screamer, on the other hand, with high rankings for 'warm' and 'crunchy' would likely see both these descriptors used far more liberally to describe its sonic properties. The likelihood is that 'balanced' refers to consistency within the equalisation spectrum, no single band is protruding above or below the others and consequently a rounded timbre is achieved. The inherent link lies in that to achieve this balance, the signal and settings must be controlled. The definitory ideas of being 'kept in check' are entirely necessary for a 'balanced' tone to be observed.

Ch.6- Definitions Panel

Once the lexicon had been established, it was quickly realised that the general descriptions of many of these words did not remotely apply to the auditory distortion context they were being used in. For example, 'crunchy' has a traditional definition of: '(especially of food) firm and crisp and making a sharp sound when you bite or crush it' ("crunchy", 2013); Most modern definitions even go so far as to specify its exclusive usage in the contextual domain of food. This lexical body, with its mismatch of inaccurate definitions, was useless without clarification on exactly what was meant when they were used in specific regard to auditory distortion. The decision was rapidly made therefore to hold, over 3 independent sessions, a discussion-based panel that would collectively determine recontextualised definitions for the top twenty most frequent semantic choices in the lexicon. The panel consisted of myself and seven other musicians (in total: 6 guitarists, 2 non-guitarists).

The first session heavily centred around existing definitions and whether anything of substance could be drawn from these that would aid us in the process of recontextualization. The process of free multiple sorting had been outlined in the planning phase of the experiment, as possibly being effective for classification; this was seen as an efficient method of initially grouping various lexis. The idea of exhausting all sorting possibilities and then providing descriptions for each was heavily informed by a paper more relevant to visual representation of later data; for their experiment Ares & Varella decided upon a 'rapid sensory descriptive method' (2012) that involved making subjects sort samples into as many clusters of groups as they see fit, and repeat the process until they can provide definitions for each sample. This is a method first used in 1967 and one found to be, despite individual data being lost/not recorded, "superior in representing all possible dimensions of categorization of the data" (Rosenberg & Kim, 1975) by 1975. The groups/clusters themselves were also discussed as understanding which perceptual attributes were being used for classification was important not only to this phase but throughout the remaining research and experiments. The most popular groups in this first session proved to be ostensibly obvious (in hindsight): level of gain, soft or hard clipping and dynamic consistency (often related to transient noise) emerged as the most popular early in the proceedings.

In this session, onomatopoeia was also heavily discussed as a large majority of the featured descriptors were examples of this literary effect. It was identified which descriptors were onomatopoeic and subsequently, individually and collectively, why onomatopoeia had emerged so frequently within the semantic choices. The onomatopoeic descriptors in question are: 'crunchy', 'fizzy', 'crispy', 'fuzzy' along with other ideophones such as 'sludgy', 'muddy' and 'woolly'. The commonality of lexical choices such as this, as previously, suggested is potentially linked to quite primitive semantic ideals, ease of understanding and interpretation namely.

'Crunchy' obviously needed lengthy specific discussion to become to gain an understanding of how to best define this particular descriptor. It's inherent links to clipping

were agreed upon immediately, as was its ties to levels of mid response. The mid response wasn't overpowering or exceedingly present above bass and treble response, just a clear emphasis placed upon the mid-range. The importance of this (and the mid response in general) is well surmised by Rob Stewart (2017): "Definitions of what the midrange of frequencies is will differ slightly depending upon who you ask. I define the midrange as the range between 200Hz and 5kHz which covers the entire critical range of the human voice (300Hz to 3.4kHz) plus a bit more. The midrange is a powerful zone, because our hearing has evolved to be most sensitive to the midrange, particularly the upper mids. Too much midrange energy can make it sound too hard, too boxy, too loud or too edgy. Too little can make it sound dull, scooped or soft. Getting the mids right is critical ". It was also suggested that the connotations of its usage were generally positive, ideas of 'too crunchy' were not ones any member of the panel had heard used in conversation, or any other context. It was suggested that excessive levels of clipping would use alternative adjectives with more negative inferences, and overwhelming mid response would again search for more a more negative semantic choice, 'muddy' for example. Therefore, 'crunchy' exists in a strangely positive definitory context due to alternative adjectives better representing its defining features far more accurately. In sharp contrast, the aforementioned 'muddy' seemed to have an intrinsically negative definitory sphere. In a case of role-reversal, 'muddy' seems only to be used to describe oversaturation of an audio property, to the point of sonic deformation and incoherence. The inherent link with overbearing mid frequencies did not become overly apparent until the third session where the results of the similarity matrix were used as an indicator for definitory sonic properties for each descriptor. However, the idea of oversaturation and subsequent incoherence as a result were the first notions that emerged in the panel's collective discussion. Further discourse led to us linking the usage of 'muddy' and its primarily negative connotations with other similar descriptors included in the lexicon; these were: 'woolly', 'fizzy', 'harsh', 'noisy' and slightly aside, 'thin' due to its association with absence instead of constructive sonic properties. The idea of congestion also emerged through conversation; 'muddy' directly connoting a lack of clarity transferred to the sonic realm whereby if recontextualised, it's new definitory sphere must refer to notions of obtrusive signal alterations that impurify the tone to the point of incoherence. This idea of predetermined positive/negative bias to the connotations of some descriptors, even within an entirely new context, only emerged as a trend during the panel discussion stage of the research.

The second and third sessions consisted mainly of clarification and refinement of ideas and definitions previously discussed in the initial session. Many of the definitions were written, re-written and then re-written again to account for different individual sorting methods or differences in opinion, on which audio features or traits defined particular semantic choices. Within the space of these three 20-40-minute sessions however, a full set of definitions was agreed upon for the top twenty most frequently occurring terms in the lexicon (text mining and survey data combined). Many required two to three revisions as we aimed to constantly critique and troubleshoot any etymological issues with the definitions we had created. We were aided also by guidance from the similarity matrix during the final session, whose results had been calculated to give a numerical approximation as to which variables of distortion affected specific semantic choice. For example, this gave an added

perspective to descriptors proving difficult to quantify such as ‘crispy’; the numerical data very clearly hinted towards a close link between decreased sustain and increased treble when ‘crispy’ was used and we subsequently adapted the existing iteration of the definition to clearly mention these sonic traits. Using this data to revise our descriptor definitions where needed, following the final session, the compiled list appeared to be cohesive and contextually relevant throughout. This list of definitions is displayed in Table 6.1.1.

<u>Descriptor</u>	<u>Definition</u>
Smooth	A sonic quality of having increased dynamic consistency often created with soft clipping, low gain and increased sustain.
Warm	Greater emphasis on the bass frequencies, often with a softer edge due to lower levels of gain.
Full	The quality of having all frequencies present, especially bass frequencies equating to increased dynamic consistency
Creamy	Lower levels of gain and clipping with increased dynamic consistency creating a rounded, rich tone.
Dark	Mellow, excessively rich sound, characterized by decreased treble frequencies, often with a heightened bass response.
Bright	Greater emphasis on the treble frequencies with a hard, crisp edge and attenuation of the bass response.
Thin	Lack of prominence of any frequency band, notable lack of treble frequencies, creating an unfulfilling sonic presence

Table 6.1.1 (left)- Definitions

Muddy	Overbearing mid frequencies, coupled with soft clipping causing sonic incoherence
Fizzy	Decreased bass response and sharp, oversaturated treble frequencies
Woolly	Lacking clarity or sharpness, often due to loose, ill-defined bass frequencies and a greatly reduced treble response
Dirty	Hard clipping with high gain creating a raspy tone
Crispy	Brittle in texture, characterized by decreased sustain and peaks in the mid-treble region
Crunchy	Hard clipping and high gain with peaks in the mid regions creating the onomatopoeic qualities
Thick	Sodden bass with pronounced mid and treble responses producing a complete, rounded tone
Hot	High levels of gain and hard clipping creating loud, unstable sounds
Fuzzy	Extremely hard clipping with high gain creating a muffled timbre
Fat	Exaggeration of the mid and upper bass ranges creating a turbid tone
Sludgy	A viscous sound characterized by muted treble frequencies and high levels of gain
Harsh	Unpleasantly rough or jarring to the ear denoted by shrill treble and excessive gain or clipping
Noisy	Loud, with extreme levels of gain and/or clipping, often featuring unwanted, transient noise

Ch.7- Similarity Matrix

The similarity matrix was designed to allow for quantitative measure and data to be gathered, in a subjective thematic domain. Semantic choice is inherently unique and theoretically, even when using the same lexis, individuals could have entirely different connotations and definitions in mind which affects their usage. The chances of this are heightened when using lexis that are distinctly hard to quantify, that exist in a qualitative realm almost, with fluid definitions. The top twenty most frequent descriptors, which in terms of analytical material by this point have become their own subset-lexicon, was comprised heavily of lexical choices that fit this brief: easy to comprehend colloquially but intrinsically difficult to quantify or universally define. In an attempt to apply scientific principle and numerically quantify these terms, against the most applicable sonic variables, a matrix was devised by which participants could rate each descriptor in regard to each variable. After being completed by twenty individual participants, the averages of the ratings were calculated which provided the numerical data to analyse, displayed in Fig. 7.1.1. The preliminary results were used to inform the final stages of the panel discussion on definitions which allowed the definitions to reflect the quantitative viewpoint in some way, otherwise they would be totally independently exclusive of each other, which in the context of the thesis didn't make much sense. The final results were the academic grounding and provided the data for the dendrogram to be constructed, which in turn, allowed the wheel to be rendered, creating a satisfying visual representation of the results from the similarity matrix.

	Crunchy	Warm	Smooth	Hot	Fuzzy	Noisy	Dirty	Muddy	Thick	Thin	Wooly	Creamy	Fizzy	Harsh	Fat	Bright	Dark	Crispy	Sludgy	Full
Soft Clipping	1.3	3.7	3.2	1.3	1	1.8	1.7	2.8	3.5	1.3	2.3	3	0.8	0.8	3	2.3	2.5	1.5	1.2	2.3
Hard Clipping	3.5	0.8	1.2	4	4.5	4	4.2	2.5	2.7	1.8	1.2	1.5	4	4.7	3.8	2.5	1.5	2.8	2.8	2.3
Low Gain	1	3	3.7	1	1	1.5	1.2	2.7	2.7	2	3	2.5	1	0	1.5	2.3	2.8	1.2	0.7	1.8
High Gain	3.7	1.8	1.7	4.5	4.8	4.2	4	2.8	3.5	2.7	1.7	2.3	4.3	4.7	3.7	2.7	1.7	3.2	3.5	2.3
Quiet	1.2	2.3	3.7	1	1	0.3	1	2.2	2	2	2	2.3	1.3	0	1.7	0.8	1.7	1.2	0.8	1.2
Loud	3.2	1.5	2.2	3.8	4	4.8	2.8	2.5	2.8	2	1.8	2.3	3.7	4.7	2.5	1.5	1	3.5	3	2.3
Decrease Sustain	3.3	1.2	0.7	2.7	1.7	2.5	3	2.8	2	3	2.5	0.5	2.8	2.7	1.7	1	0.7	4	1.7	1.8
Increase Sustain	0.8	2.5	4.2	1.7	2.7	1.5	2.3	2.3	2.5	1	1.7	3	1.3	1.2	3.8	0.8	2.2	0.5	4.2	2.8
Decrease Dynamic Consistency	1.8	1	0.7	3.2	2.3	3.3	2.7	3.3	1.7	3.2	3.2	0.8	3.2	3	1.3	2.7	1.8	2.7	1.7	1
Increase Dynamic Consistency	2.8	1.7	4.5	2.5	2.3	1.2	2.3	1.5	3.7	1.8	2.3	3.3	1.7	0.7	4.8	1.3	2.3	2	4.8	4.3
Decrease Bass Response	2	0.5	1.8	2.8	3	3.2	2.7	2.8	0.2	4	2.7	1.3	4.5	4	0.7	3.3	0.7	2.7	0.8	0.3
Increase Bass Response	2.5	4.5	2.3	2.5	2.2	0.8	2.5	3.2	4	0.3	2.3	3	0.3	1	4.5	0.3	3.3	1	4.3	3.2
Decrease Mid Response	1.2	1.7	2	1.3	1.3	1.2	2	2.8	0.8	3.3	3	0.5	1.8	1.5	1.7	2.3	2.2	1.8	1.2	0.8
Increase Mid Response	4	2.7	2.3	3.3	3.7	2.7	3.7	3	3.5	1	2	3.5	3.2	3.8	3	2	2.8	2.5	2.5	2.8
Decrease Treble Response	1.3	2.2	2.5	2.2	2.5	2.3	4	3.8	2	3.3	4.3	1.8	2.2	1.2	1.5	0.3	4.5	1.8	4	0.8
Increase Treble Response	2.8	1	1.5	3.3	2.5	2.8	2.7	2	2.2	3	1	2.3	4.2	4.7	1.8	5	0.3	3	0.5	2

Fig.7.1.1- Similarity Matrix (Mean Average Scores)

The results from the similarity matrix proved to be integral to the entire line of research, despite only seeming simple compared to other lines of analysis being undertaken simultaneously. It proved estimated hypotheses such as the choice of 'fuzzy' was very closely linked to high levels of clipping, along with 'crispy' being inherently tied to lower or decreased sustain being present in the given audio sample. While being satisfactorily predictable, in that from the average scores, there was little to no anomalous data recorded, the analytical verification of notions established prior to research beginning was welcomed. 'Warm' and 'Bright' having close links to increased bass response and increased treble response, respectively, is an excellent example of this. For the descriptors that were proving difficult to quantify in any definitive sense, the results of the matrix helped provide insight as to what musicians mean to communicate when using these terms; 'woolly' being most closely linked with a decreased treble response is a prime example of the nuance that this matrix offered. In the instance of 'woolly' with the panel, everyone understood the reference to a lack of clarity but specifying this further proved challenging. To obtain clear, numerical data which so strongly visualised its defining trait was decreased treble response aided the definition and general research greatly. Similarly, 'dirty' being tied with hard clipping and increased mid response was interesting, usually hard clipping and high gain have similar correlations but, in this case, it was interesting to note the specificity of increased levels of hard clipping without necessarily needing the high gain. I highly doubt this separation would have been made as in discussion, the two did not seem distinguishable in this context, the matrix therefore gave great definitory insight not previously available. The results for 'fizzy' seemed fairly comprehensive; high gain, hard clipping, decreased bass response and increased treble response all averaged 4.0 or more. This was expected as the onomatopoeic descriptors in their etymological make-up alone, hint at which sonic traits they are supposed to phonetically resemble. In this instance, the harsh fricative voicing of the 'z's imitate the shrill overtones of an oversaturated treble response. The results corroborating this colloquial estimation was very pleasing. Another descriptor that, even in discussion, proved difficult to collectively quantify in any way was 'sludgy'. Traditional definitions and etymology refer to viscosity and thick consistency but transposing those definitory features into a sonic domain proved hard to articulate. The matrix's results showing clear links to increased sustain, increased dynamic consistency and increased bass response seemed to, retrospectively, match closely the muculent qualities connoted in traditional definitions. Distinctly low average scores for low gain and soft clipping perhaps also indicate that 'sludgy' is often referring to more aggressive forms of auditory distortion, where the effect, with added sustain, creates a thick, viscous distorted tone. 'Thick', in that context refers to its own matrix scores which showed heavy bias towards increased bass response and increased dynamic consistency, mirroring 'sludgy' demonstrating the two descriptors are inherently similar, with minor differences. In the case of 'sludgy', the strongest link was to increased sustain, implying length of note mimicked the traditional viscous definition. This was not a trait it had in common with 'thick', whose results were negligible enough to not indicate a preference towards increased or decreased sustain. Likewise, 'smooth' seemed to have categorical results which showed close connections with increased sustain and increased dynamic consistency but even numerically, it showed a clear, definitive focus on increased dynamic consistency, which was

expected. The dendrogram was designed to explore semantic similarity in much greater detail but it was interesting to observe shared traits that linked words in an intermediate fashion, this early in the process. The descriptor with perhaps the clearest numerical indication of sonic values was 'harsh'; hard clipping, high gain and increased treble response were all rated 4.7/5.0. Every one of the additive or more aggressive form of the traits had the overwhelming numerical weighting for 'harsh' except the score of 4.0 for decreased bass response. This seems to heavily indicate the negativity of this descriptors context in the sphere of auditory distortion, hinting lack of balance and oversaturation of higher frequencies.

Ch.8- Dendrogram and Distortion Wheel

One of the final actions of the qualitative process, was finding an optimal way of visually representing the data. Using the numerical average scores from the similarity matrix, a dendrogram was decided upon as an effective means of measuring semantic similarity- calculating how similar the descriptors were to each other by analysing and comparing the definitory strength of the variables they were measured and rated against. This is displayed in Fig. 8.1.1. The results from this dendrogram were predictable for the most part, with a couple of surprise results that perhaps stemmed from the linearity of the similarity matrix results. A quick glance at the results of the similarity matrix would allow colloquial comparisons to be drawn between descriptors based on high scores in similar fields, the dendrogram visualises these links and allows closer analysis of the less obvious similarity links.

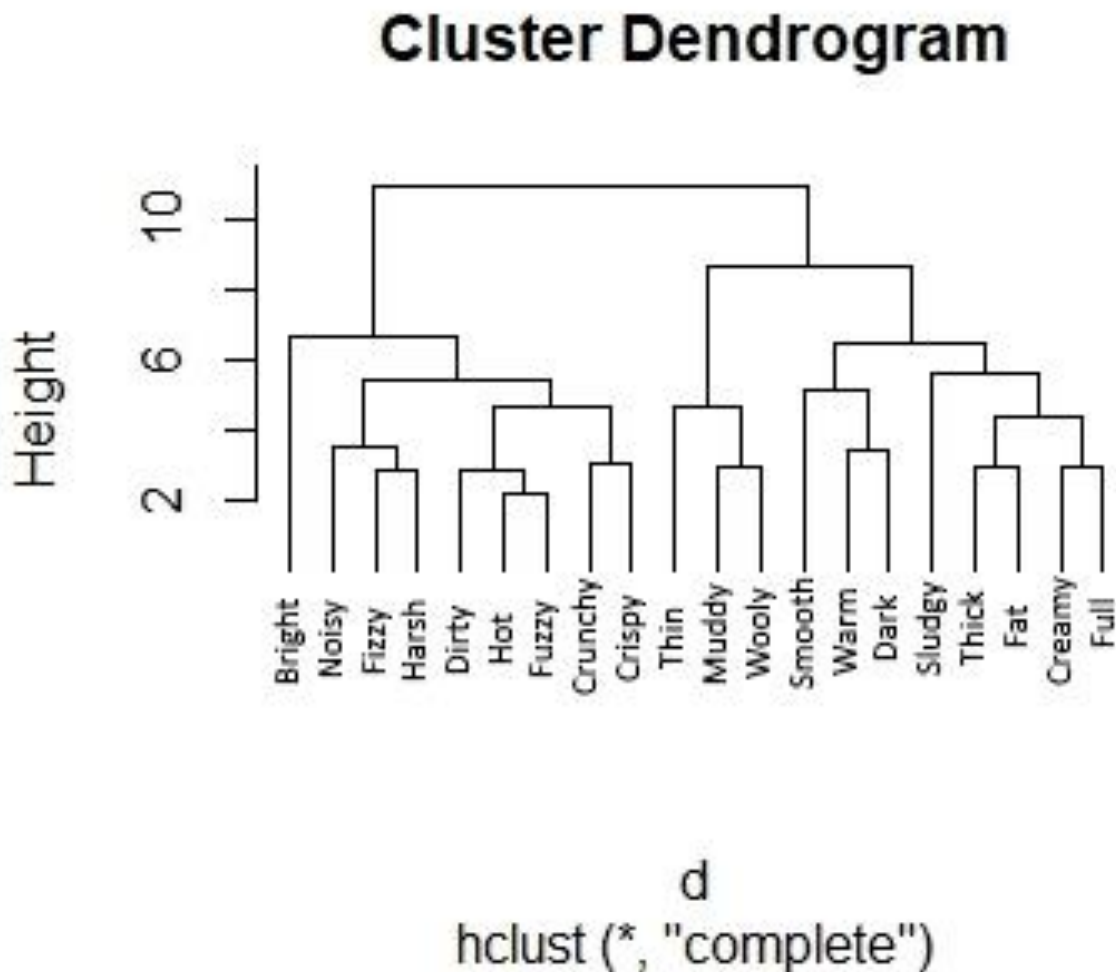


Fig.8.1.1 (above)- Dendrogram

Some of the closest observable links obviously are between descriptors that share similar definitions. This was to be expected, as part of the defining process was to analyse and implement each descriptor's strongest sonic traits in the definitions themselves. For example, 'warm' and 'dark' were proved to be intrinsically linked in regard to similarity. Their shared properties of typically lower treble response and higher than average bass response facilitated this match. Similar tendencies leaning towards soft clipping, low gain, increased sustain and distinct lack of importance given to dynamic consistency will also have contributed greatly to their close proximity in regards to similarity. Even outside the parameters of the matrix definition, these two descriptors are inherently similar within the contexts of which they are often used. Another very close similarity match was between 'thick' and 'fat' which was again, somewhat expected given their definitions and close links with similar influencing variable. 'Muddy' and 'woolly' I had always personally linked in regards to their attempts to communicate sonic incoherence, to see the dendrogram validate this through the variables was an excellent result. Their shared trait of decreased bass response seemed the main analytical indicator of their similarity, with decreased dynamic consistency being the only other variable where consistent correlation was shown between the two descriptors. Interestingly, 'thin' was rated as very closely related to the immediate pairing of 'muddy' and 'woolly'. 'Thin' also attempts to describe absence instead of additive sonic traits making it similar, in purpose, to these two negatively tinged descriptors.

The grouping of 'noisy', 'fizzy' and 'harsh' into an inherently negative subset provided an interesting piece of information. 'Thin', 'muddy' and 'woolly' are similarly negative descriptors however the negativity surrounding them stems from subtractive or omissive sonic traits, the lack of treble presence or pronunciation being the easiest observable link. With 'noisy', 'fizzy' and 'harsh', the opposite seems to be true with many of the critical definitory features being additive. 'Noisy' fits this description best, with its usage being solely linked to over-exaggeration of particular frequencies, or the jarring inclusion of unwanted transient noise. 'Harsh' seems specifically matched with over-pronunciation of certain sonic elements to the point of auditory annoyance. 'Fizzy' is more specific but still references a sonic exaggeration, pushed to the point of impacting negatively on the overall timbre. These two subsets therefore demonstrate that the negative descriptors usually relate to additive or subtractive sonic properties, as opposed to focusing on more negative aspects of the general timbres. Using subsets to identify common traits is useful alongside individual analysis. Directing attention towards the grouping of 'dirty', 'hot' and 'fuzzy', without delving into shared sonic characteristics, there is an immediate implication that gain and clipping are the variables that have brought these three descriptors together. The results validate this, with the average scores explicitly show hard clipping and high gain as the two highest rated numerical values, confirming their importance.

Most of the descriptors were grouped into a subset by the second tier, if not the first. Only two descriptors were not grouped early into the process, later amalgamated into a much wider subset by the third or fourth generation of matching- 'bright' and 'sludgy'. 'Bright' has numerical focus on decreased bass response and increased treble response, specific traits not shared by any other descriptor on the list perhaps explaining its isolation. 'Sludgy' has clearly definable traits but in such a specific sense which could justify the lack of

close links initially. High gain, increased dynamic consistency, increased sustain, increased bass response and decreased treble response all attained average scores of 4.0 or more which was not uncommon in the individual categories, but when amalgamated, proved to be fairly unique, hence the separation of 'sludgy'. Another likely pairing that the dendrogram analytically validated was 'crunchy' and 'crispy'. The obvious shared variable is decreased sustain as both descriptors phonetically imitate quite brittle, sharp sounds. 'Crunchy' is a lot more closely associated with increased mid response but every other variable shows correlation between the two descriptors.

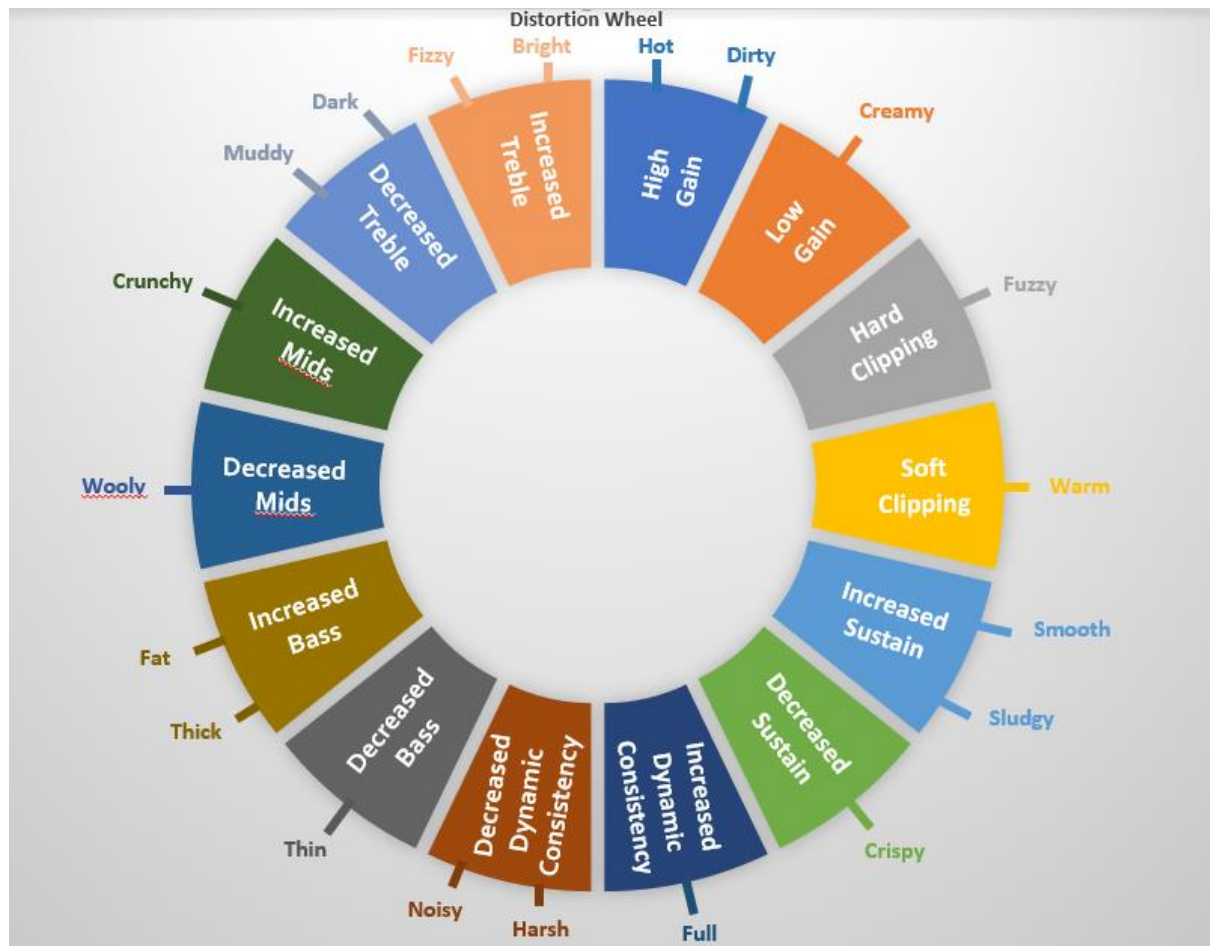


Fig.8.1.2- Distortion Wheel

The wheel is the product of every previous stage of testing and research. It provides a simple visualisation (Fig. 8.1.2) that links the most prominent descriptors with their respective definitory variables. The established top twenty most frequent descriptors were assigned to the variable of best fit, using the data from the similarity matrix averages.

Ch.9- Listening Tests

The fundamental aim of the listening tests was to add quantitative specificity to elements of the data already gathered. By this point the lexicon had been created, refined and then defined so the vocabulary was in place to be utilised within ongoing experiments. The listening tests were a perfect opportunity to use them, along with the objective analysis and experimentation. With each descriptor now attached with a relevant, contextually accurate definition, the testing went through a role reversal whereby now the objective was to ascertain which sonic properties from each pedal prompt specific contextual responses, rather than vice-versa. With the twelve pieces of audio already recorded, these became the stimuli for the listening test- low, medium and high gain settings for each pedal. Using the data from the similarity matrix, for each of the top twenty most frequent descriptors, the most appropriate set of stimuli was chosen. For example, for 'crunchy', the strong association with increased mid response warranted the humbucker pickups with medium gain, for each pedal, as the set of stimuli. However, for 'bright', single coil pickups with low gain was far more fitting for the descriptor's definition. Each participant had to assign a numerical value between 0-100 to each stimulus based on perceived cohesion. For each test, the ordering of these stimuli was randomised in an attempt to remove subconscious bias that may emerge as a result of having a consistent order which allows predetermined notions to filter into responses. This also consequently meant that it was anticipated the results of this experiment may be less correlated and show distinctly less cohesion than previous tests. This is due to the inherent subjectivity present at multiple stages of this particular experiment. The way each individual perceives each audio sample, relative to the given descriptor, which is then internally transposed into a numerical value provides many stages where personal understanding and perception will likely separate individuals apart entirely, over the course of the four different stimuli for a single pedal. When this is repeated numerous times, the results could be massively varied. Once calculated however, the average scores should still provide a decent platform to draw conclusions surrounding how relevant each descriptor is to each respective pedal. Suitability of semantic choice is the main thematic domain with which this particular experiment hopes to inform.

The results, therefore, provide insight specifically into the pedals themselves, showing a quantitative specificity towards them not previously seen in any of the previous experiments. Hultigen provided the perfect means of enabling this testing, with their MUSHRA interface proving to be ideal. Once the stimuli had been entered and the twenty tests finalised with matching descriptors, the test was completed by twenty musicians. As with the similarity matrix, once the participants had successfully completed the listening test, average scores were calculated for each stimulus, rounded to the nearest integer.

The boxplots provide visual representations that allow deeper analytical analysis on the discernible qualities of each pedal, in specific relation to the comparative descriptor. This is best evidenced with the boxplots of 'thick' and 'fat' and the conclusions we can draw from a comparative look at the two. For 'thick', the boxplot (Fig. 9.1.1) displays that the RAT

emerged as the closest matching pedal to the descriptive qualities of 'thick'. While the DS-1 is ranked the lowest of the four devices, the comparatively low ranking for the Big Muff is also notable, given that its sonic qualities seemed to lend itself to antonyms of 'thin'. The closest antonym is 'fat', with the main etymological difference being that 'fat' has further connotations of saturation, than is present, definitively, for 'thick'.

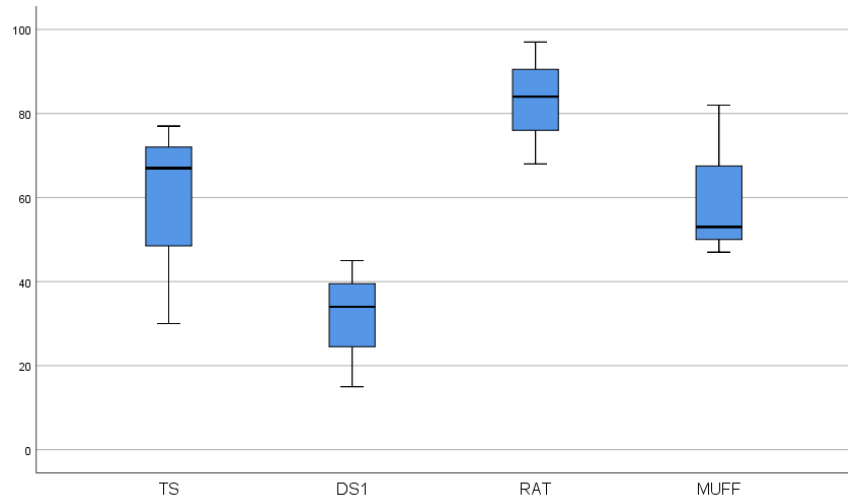


Fig.9.1.1- Boxplot- 'Thick'

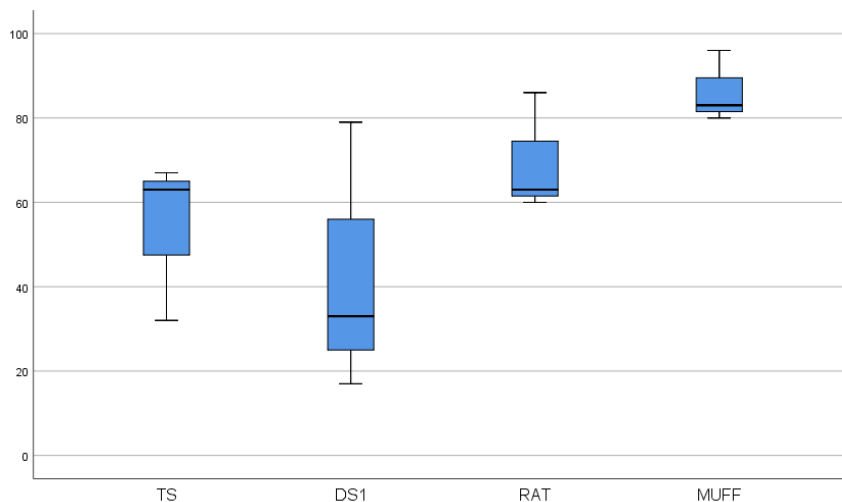


Fig.9.1.2- Boxplot- 'Fat'

Comparison with the results for 'fat' (Fig. 9.1.2) then should reveal the critical perceptive differences between the two similar descriptors. The DS-1 scores the lowest again, confirming the notion that Boss' product does not fulfil any of the descriptive criteria for descriptors

denoting pronounced bass or mid frequencies. The emergence of the Big Muff as the closest applicable descriptor adds credibility to the idea that the definitive descriptive difference between 'thick' and 'fat' are oversaturation. The fuzz-like sonic properties of the Big Muff, most clearly defined by excessive levels of clipping and gain, lend themselves well to the descriptive properties of 'fat'. The RAT finishing behind only the Big Muff shows the propinquity of these two particular descriptors.

Conversely, analysing the results for the antonym 'thin' provide a different angle upon which to draw conclusions. 'Thin' is the direct etymological opposite of 'thick' and so it

could be expected that the results will mimic this diametric opposition. While not exactly matching, the results for 'thin' do add further insight and credibility to the results and subsequent notions and consequences drawn from the previous two boxplots.

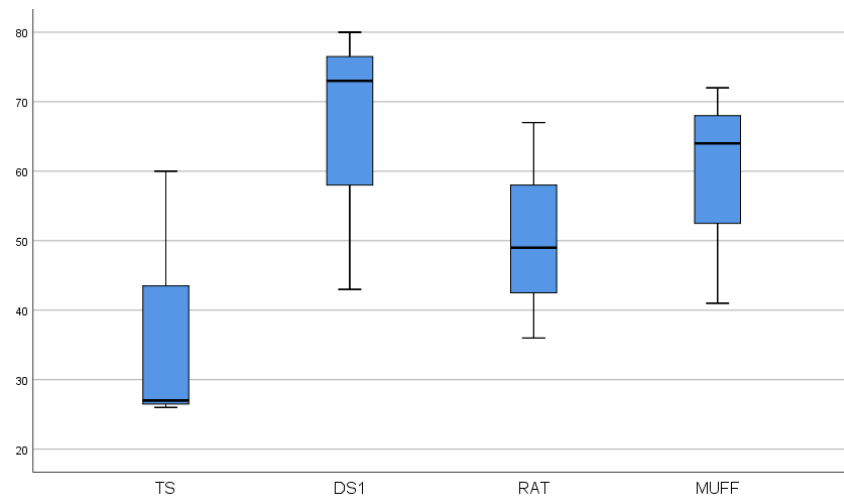


Fig.9.1.3- Boxplot- 'Thin'

The DS-1 emerges as the most applicable descriptor to accurately match the sonic implications of 'thin' (Fig. 9.1.3). With the upper limit of the range recorded at 80, this is indicative of just how fitting participants thought that 'thin' was for the sound of the DS-1. The least applicable pedal was determined to be the Tube Screamer. This could be seen as unsurprising since the TS9 features vastly reduced levels of gain or clipping, compared to the other devices. This means the outputted signal retains many sonic characteristics of the clean tone, meaning participants were less likely to liken its sound to the subtractive definitive qualities of semantic choice 'thin'. The RAT scoring below 50 adds credibility to the results as it reaffirms its ranking for 'thick' and 'fat'. The Big Muff's results in this instance could be seen as negligible due to the large range and the aforementioned variance in ranking between 'thick' and 'fat'.

In similar fashion, side-by-side comparison of the boxplots for contextual antonyms 'crispy' and 'sludgy' allows closer examination and analysis on the subtle nuances of each pedal's outputted signal, measured subjectively against opposing descriptors (Fig. 9.1.4).

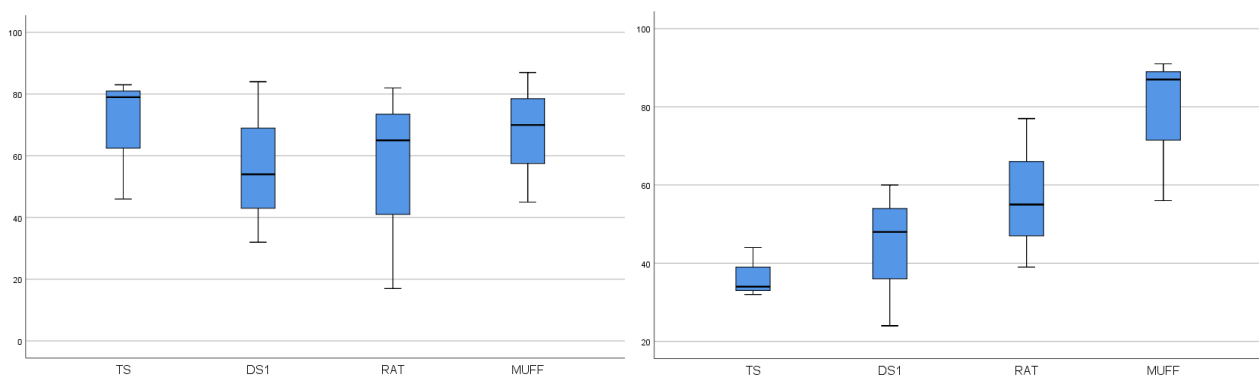


Fig.9.1.4- Boxplots- 'Crispy' & 'Sludgy'

The results for 'crispy' show that participants believed the Tube Screamer endowed the signal with the shortest amount of sustain, while for 'sludgy', the results show conclusively that the Big Muff was the pedal to feature the highest perceptible levels of sustain, gain and muted or imperceptible treble frequencies. While the ranges for each pedal under 'crispy' seem large, the results obtained are fairly marginal; in stark contrast, 'sludgy' has far more

clearly defined results but with slightly smaller ranges. This potentially indicates that the participants found 'crispy' to be far more subjective when definitively ranking each pedal, while for 'sludgy' there was more of a consensus unknowingly reached between all participants, reflected in the smaller ranges and more discernible rankings.

While some results are best analysed as part of a group, others need closer individual analysis to draw conclusive notions from. An example of this is for 'dirty', one of the more intrinsically subjective descriptors from the list. The rankings from least to most applicable of: Muff, TS9, DS-1, RAT; demonstrates this aptly (Fig. 9.1.5). For the Big Muff to be ranked last, this disputes one element of the definition which places emphasis on notions of hard clipping. This particular result refutes the defining qualities established by the panel discussion, therefore. The Tube Screamer ranks slightly above however, potentially indicating that excessive clipping- as creates the ubiquitous Muff-effect, is too sonically

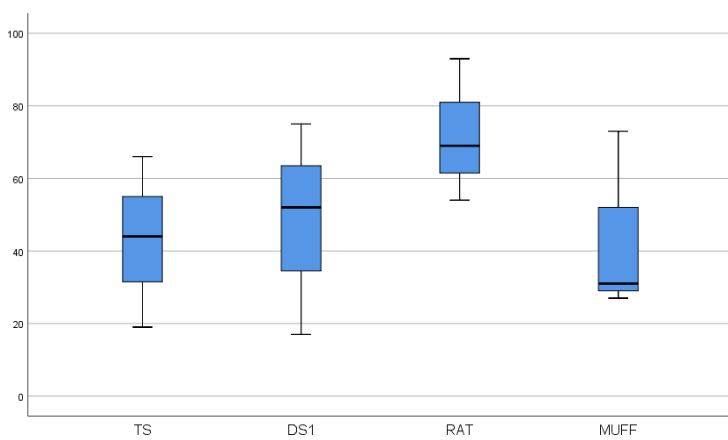


Fig.9.1.5- Boxplots- 'Dirty'

extreme to fit 'dirty' but high gain could still be an influencing variable. The DS-1 ranked second, finishing only below the RAT. These pedals share a focus on high gain, without needing as extreme levels of clipping as the Big Muff, strongly indicating that perception and semantic application of 'dirty' is heavily informed by increased levels of gain.

The remaining boxplots have been compiled together for ease of viewing and comparison. Brief analysis of these results is discussed below:

**Fig.9.1.6
(right)-
Boxplots
(labelled)**

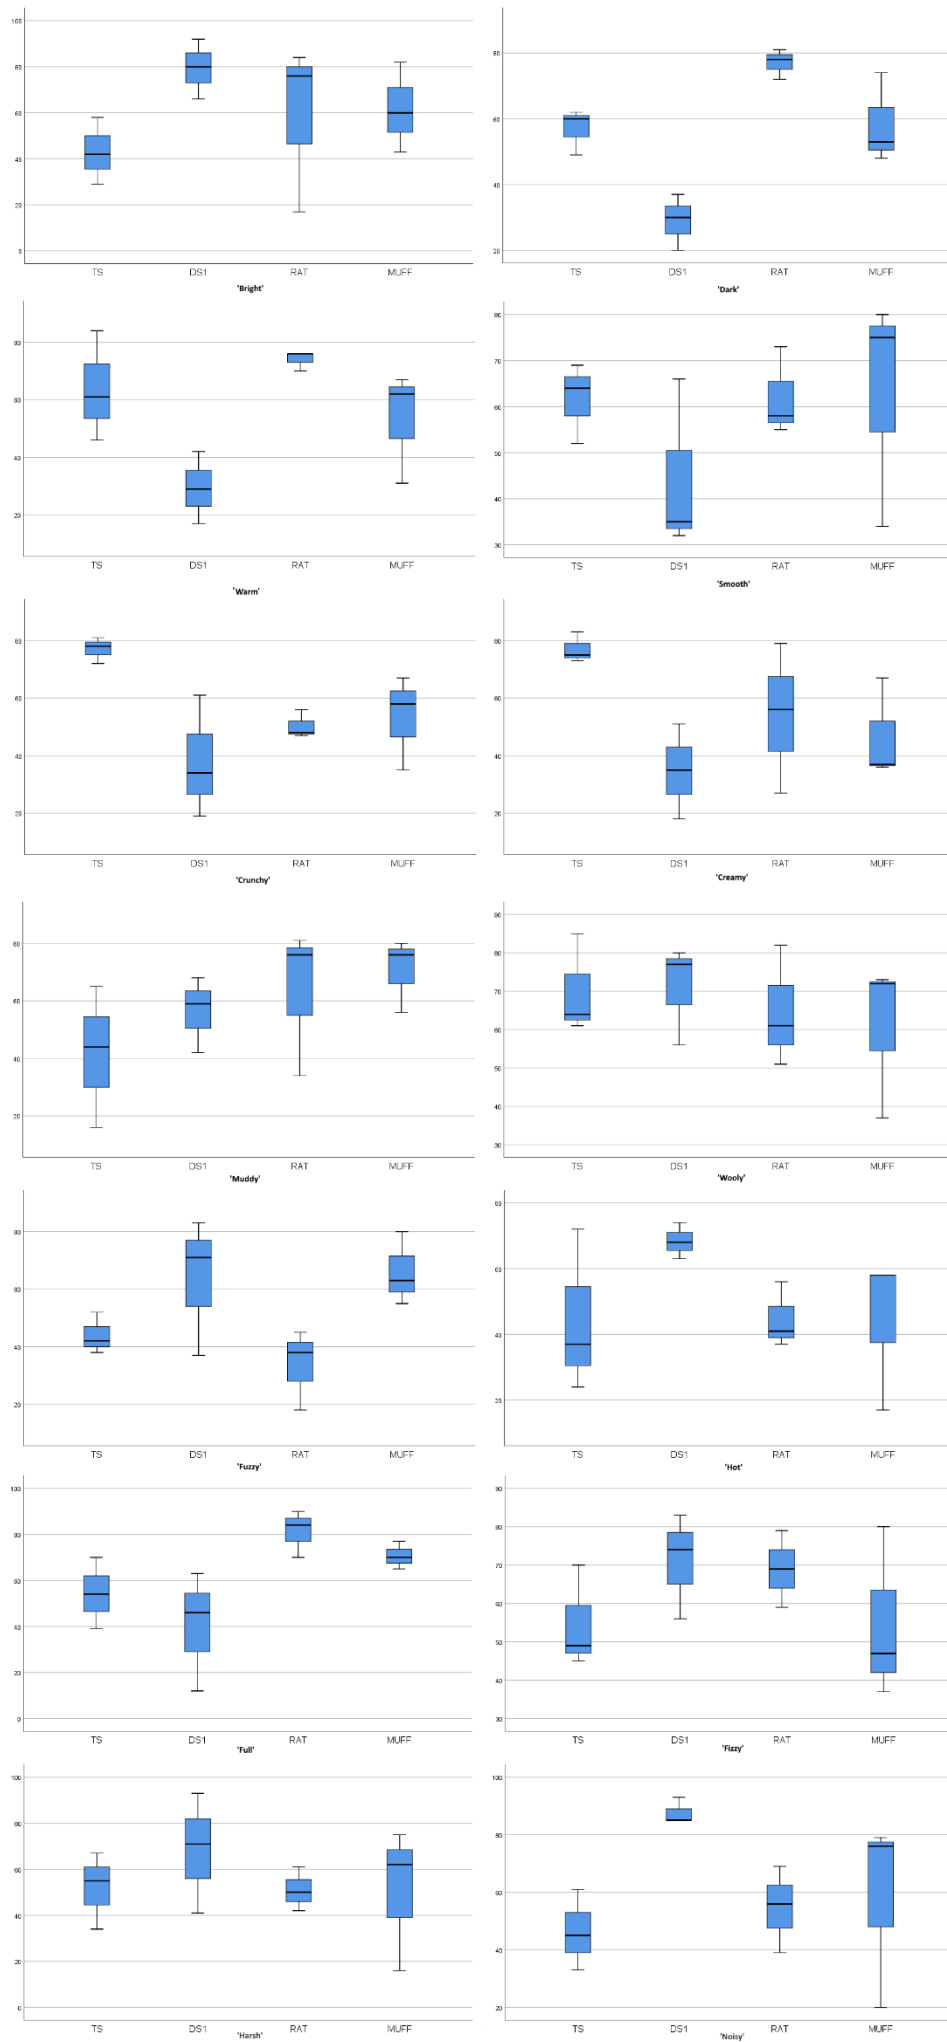


Fig. 9.1.6 displays the boxplots for the remaining descriptors. The results for 'bright' and 'dark' display consistency of result, even while diametrically opposed in regard to semantic implication. The DS-1 is highlighted as the brightest and the least dark device. The most applicable pedal for 'warm' was deemed to be the RAT, with a very small range given meaning unanimous agreement. This narrow range makes the result an outlier, however. On the opposite end, the DS-1 was found to be the definitively least 'warm' pedal out of the four featured. This is an example of the the shared definitory properties between 'dark' and 'warm', most notably the lack of prominence of the treble frequencies. The most 'smooth' outputted sound came from the Big Muff demonstrating that lower levels of gain can potentially negate the harsher qualities of hard clipping. The TS9 which finished just below better embodies the soft clipping, low gain measures of 'smooth', hence the slightly surprising emergence of the Big Muff as the most applicable. The Tube Screamer was ranked as the most applicable pedal for descriptors 'crunchy' and 'creamy'. This implies that depending on setting, the TS9 is capable of effectively characterising low gain, soft clipping tones reflective of the usage of 'creamy', but also the hard clipping, high gain 'crunchy' sound. The result for 'muddy' saw a tied favourite, the participants decided that both the RAT and the Big Muff equally summarised its definitory qualities, primarily overbearing mid frequencies. Whereas for 'woolly' whose differentiation is achieved by specifying reduced treble response and ill-defined bass frequencies, the closest matching tone came from the DS-1. For 'fuzzy', the DS-1 also emerged as the most applicable, however the margin between the DS-1 and the Big Muff is single figures, inferring that both could be considered to produce a 'fuzzy' sound. In contrast, the margin was far more extreme for 'hot', when it determined the DS-1 to produce a signal with sonic instability due to high levels of gain and hard clipping. Furthermore, in a predictable result, the Tube Screamer was ranked the least 'hot' pedal out of the featured four. With slightly more positive connotations relating to heightened dynamic consistency, the most apt 'full' tone was decided to be the RAT. Inherently linked to 'bright', the results for 'fizzy', also denoting excessive treble frequencies to the point of being abrasive, determined that the DS-1 audibly demonstrated this quality to the greatest degree. Finally, the 2 inherently negative adjectives- 'noisy' and 'harsh'. The DS-1 also emerged as the most applicable pedal for both descriptors though the previous discussion on the transient noise present must be considered as an influencing variable for these two. The excessive levels of clipping and gain available to Boss' flagship distortion unit is another strong rationale as to why this emerged as the undisputed favourite for both 'noisy' and 'harsh'.

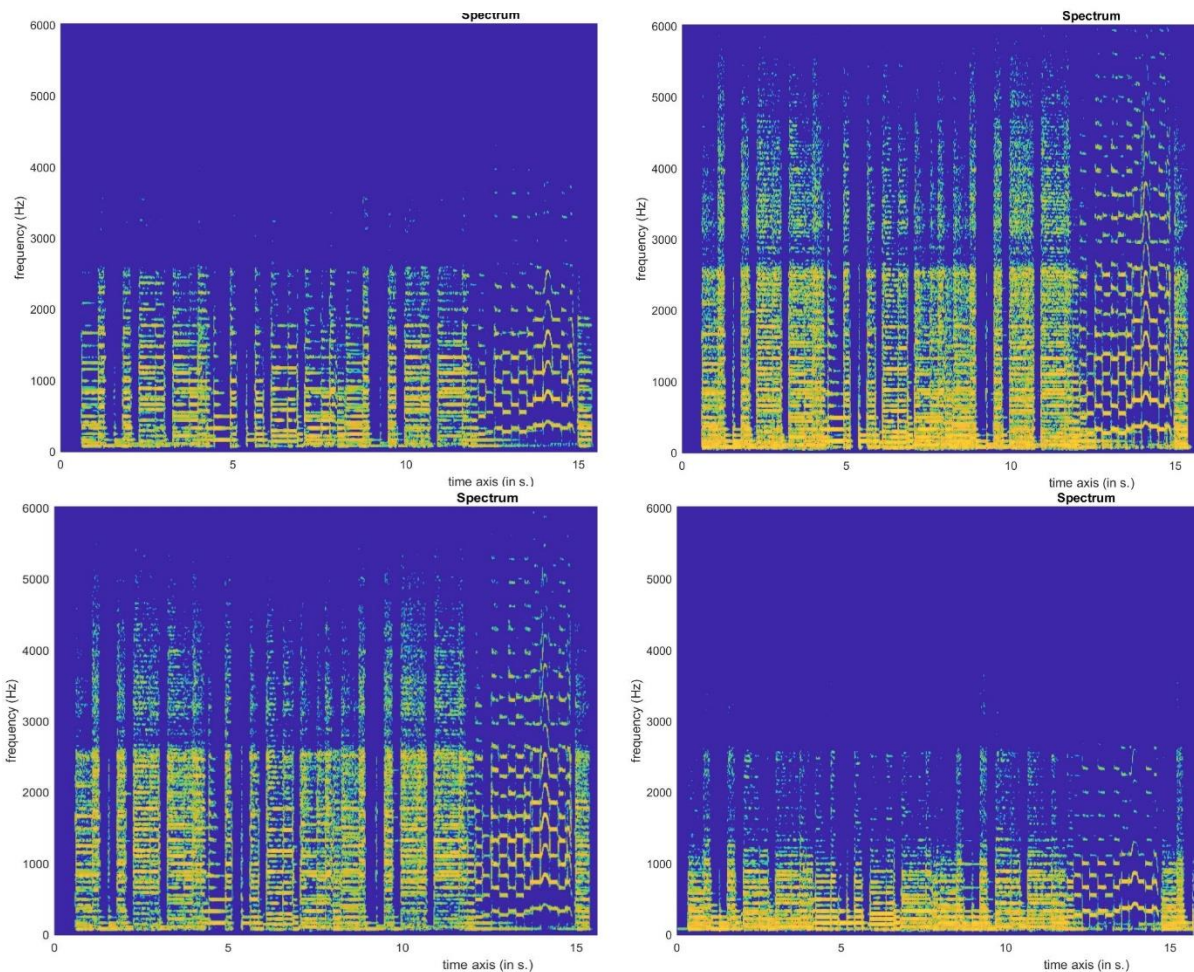
Objective Analysis

The objective means of testing was achieved through multiple different methods, in an attempt to gain conclusive results from which accurate conclusions could be drawn, ideally substantiating etymological hypotheses simultaneously. The waveforms and frequency spectrums were generated to graphically inform and identify the individual frequency components of the signals. The spectral centroid allows close examination of where the central emphasis of the spectrum lies, which together with spectral flux, provides a fairly comprehensive viewpoint of the fluctuation and timbre of any given signal. The Total Harmonic Distortion (THD) measurements then compound this information with clear, empirical evidencing of the level of auditory distortion present. The remaining calculations and graphical representations of brightness and roughness both provide context and extra pertinent information. Brightness, very generally, allows one to ascertain the 'sharpness' of the audio, whereas roughness provides juxtaposition to this measure by informing data more relating to timbre.

Ch.10- Spectrograms

The spectrogram provides excellent insight into the specifics of how each device imparts upon the source signal, sonically. By providing a visualisation of the strength across the sonic spectrum, judgements can be made regarding the sonic signatures of each pedal respectively. Pictured below (Fig 10.1.1) is a collation of the spectrograms for each device, using single coil pickups, with the gain controls set to low. (Left to right, top to bottom: TS9, DS-1, RAT, Big Muff)

Fig.10.1.1 (below)- Spectrogram- Low Gain

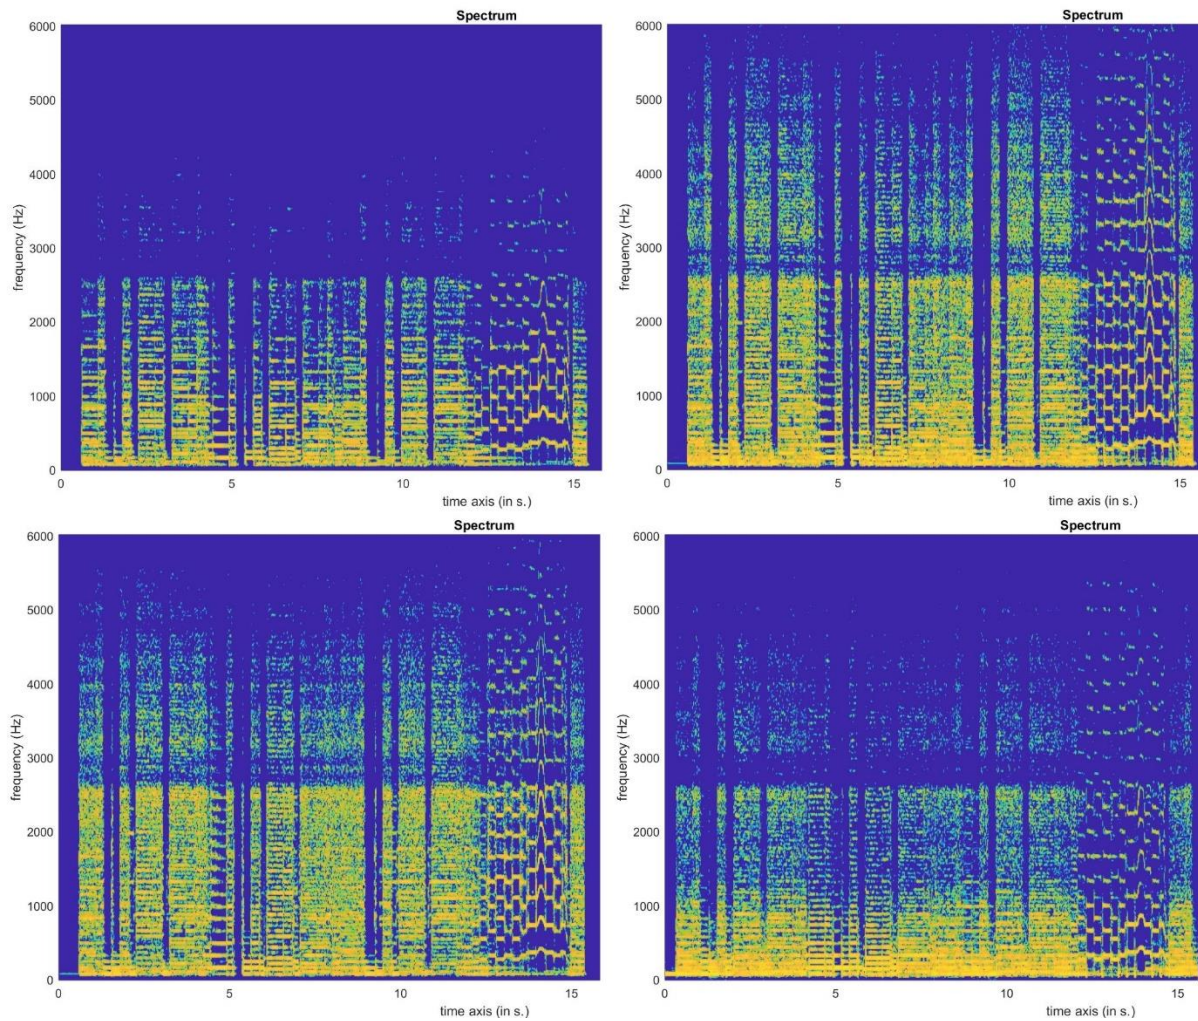


These initial results are very interesting as they actively go against some of the preconceptions laid out earlier in the thesis. The most notable example of this is the spectrogram for the Big Muff. Electro Harmonix's pedal is colloquially seen as the most aggressive and featuring one of the largest footprints of any of the featured pedals. To then see its spectrogram displaying moderate strength below 1kHz but little to no strength at any frequency above that, was not a result that was not predicted. The readings of intensity at 1000Hz or below was only matched by the DS-1 however, showing the Big Muff's obvious focus on the bass frequencies of the signal. These results indicate the Big Muff places obvious emphasis on the lower frequencies making its sound 'bassy' and 'growly' as opposed to the shriller properties that would be taken up if it placed more emphasis on the higher frequencies. The only other device to share the sheer sonic strength placed upon the low end (0-100Hz) is the DS-1, however, given its attempts to seemingly apply sonic emphasis across the spectrum, rather than focusing upon one frequency band, this gives Boss's product a different sonic edge, hence why the two pedals sound audibly different. Based on these results therefore, the Big Muff is most comparable with the Tube Screamer. The most noticeable difference between the two graphs is the obvious pronunciation of signal strength below 1kHz for the Big Muff whereas the Tube Screamer's readings are more gradual in their decreasing intensity until approximately 2600Hz. For the Big Muff, 100Hz seems like a threshold from which there is an obvious, instant reduction in the intensity of

frequencies above that whereas the Tube Screamer has no visible moment where intensity levels discernibly drop. The Tube Screamer also has superficial intensity markings between 3000-4000Hz with little to no activity in the 500Hz or so below this, above the last activity at approximately 2.6kHz. This shows that the Tube Screamer has moderate overall intensity up to 2.5kHz but the absence of major activity above this demonstrates there are few shrill overtones present in the signal at low gain.

The Boss DS-1 has the most active graph with it affecting frequencies up to 6000Hz, more than double the highest threshold the Big Muff reached at the same gain level. Similarly to the other three graphs, there is a noticeable absence of activity roughly between 2600-3000Hz proving this is not anomalous or exclusive to any singular device. Arguably displaying the most intensity of any pedal, across the spectrum, the Boss DS-1's graph only starts to subside in frequency strength just prior to 5000Hz. Above this are still tangible markers of intensity however, explicitly showing the sheer breadth of sonic content the DS-1 produces. The spectrogram for the RAT was also slightly surprising in its obvious similarities to the DS-1's graph. The RAT also shows an extremely expansive sonic footprint. Close observation can establish the intensity across the spectrum is not quite as strong as the DS-1, however it shows consistent readings up to 5000Hz and even beyond, although not as strong beyond this threshold. From 0 to 2500Hz, the intensity of the readings is very consistent. The graphs aside from showing close, observable similarities to the results for the DS-1, also show a dissimilarity at these gain settings from the RAT to the Big Muff. The RAT's reading above the shared blip are noticeably less intense than the DS-1 implying that ProCo's device likely does not share the same 'shrill' properties than Boss' pedal does. Whereas the Big Muff places obvious strength of emphasis on the very bottom end (0-1000Hz), the RAT shows no obvious points of emphasis, instead placing equal weighting to the entirety of its affectable sonic spectrum.

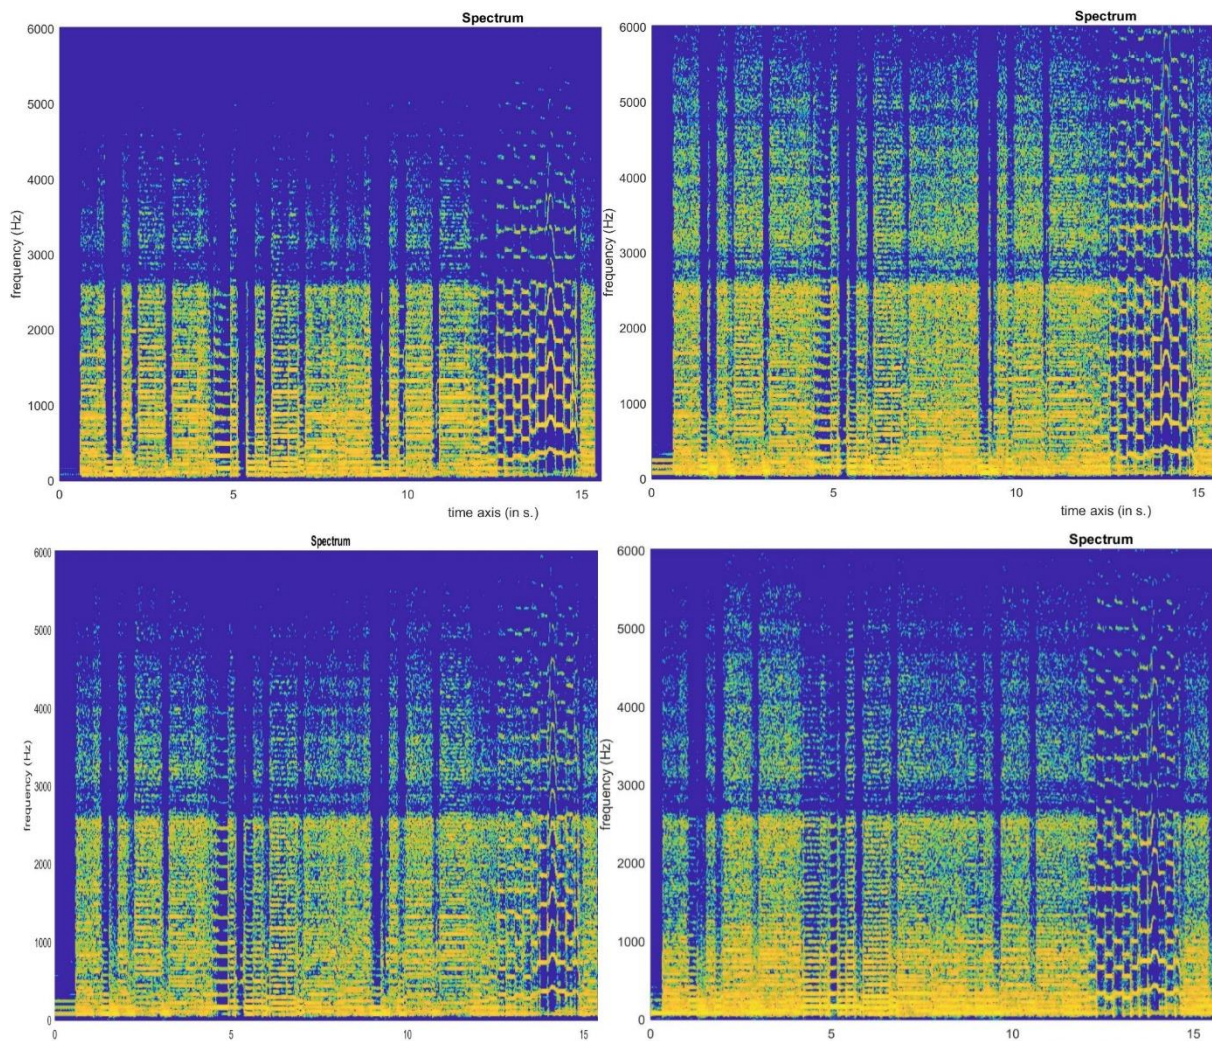
Fig.10.1.2 (below)- Spectrogram- Medium Gain



The panels above (Fig. 10.1.2) show the spectrograms generated with each of the distortion pedals set to medium gain. Generally, these are more pronounced iterations of the same plotting structure observable in the low gain set of spectrogram results. The most demonstrable difference lays with the Big Muff; its initially small area of effect (0-1000Hz), with a small increase in gain, has expanded greatly. The frequency strength of the lowest frequencies is maintained but with added emphasis beyond the 2.5kHz threshold observed previously (excluding anomalous outliers). Also, newly observable is the shared absence of any major frequency strength between approximately 2500-3000Hz, this is thanks to new results showing visible markings beyond 3kHz up to just above 5kHz where there were no markings at all previously. To gain confirmation of this frequency gap, there had to exist markings strong enough above to highlight the absence. This demonstrates an extreme rise from the Big Muff, its peak at low gain was around 2500Hz whereas at medium gain, it stretches up to over 5000Hz. Generally, the biggest graphical changes seem to exist above the shared 2500-3000Hz gap, irrespective of which pedal is focused upon. With low gain, the Tube Screamer only had outliers dotted above this gap whereas with medium gain, there are observable intensity markings consistently along the recording. This shows that with increased gain, increased levels of clipping are also present, creating more harmonic overtones which create the increased graphical readings at higher frequencies. The DS-1's

spectrogram at medium gain shares many similarities with its low gain spectrogram; the general shape is near identical with the readings simply consistently stronger, with slightly higher peak frequencies unanimously. While only faint markings reached near the 6kHz domain at low gain, scatterings of intensity markings now touch this threshold. The colour spectrum becomes most useful at this stage. With RStudio, the default colour scheme is simply based around different hues of yellow; by expanding this to include more visibly disparate green palette, it allows for more distinct separation of intensity. This allows for easier reading of the results of each spectrogram as the differences in the blue-green spectrum provide a more accurate analytical foundation for low intensity markings, while the green-yellow spectrum provides the same for stronger intensity readings. This gives the results far more nuance than the aforementioned rudimentary blue-yellow spectrum with no middle ground. Therefore, for the DS-1, the low gain chart shows most notably green results above the frequency gap whereas at medium gain, these markings become more discernibly yellow in colour, subsequently indicating these frequencies to be a fair amount stronger. These results are generally shared by most pedals, increased gain leads to increased intensity. The RAT's medium gain spectrogram follows the same progression pattern as the DS-1; marginal increases in intensity consistently across the timeline. Below 2500Hz, the spectrograms for Boss' and ProCo's pedals are remarkably similar. The one discernible difference is the higher levels of intensity visible below 500Hz for the DS-1. Consistently higher observable levels of intensity compared to the low gain spectrogram are the only discernible changes for the RAT.

Fig.10.1.3 (below)- Spectrogram- High Gain



The spectrogram results for high gain settings (Fig 10.1.3) follow the observable trends seen between the previous sets of graphs. All pedals feature increased intensity unselectively applied consistently along their affected frequency spectrums. The Tube Screamer's peak frequency has expanded to beyond 5kHz, with similarly marginal increases in intensity more visible with decreasing frequency levels. This is to the point that its intensity levels <500Hz are at a similar intensity to its counterparts the DS-1 and the Big Muff. The Tube Screamer proves itself to be the most sonically concise of the four, with maintained clarity even at the highest gain level. This is juxtaposed by the remaining three spectrograms which at this gain level have all started to show signs of graphical blurring, indicative of extreme intensities. The Big Muff arguably is most demonstrative of this, however both the DS-1 and RAT feature noticeable 'fuzziness' also. Below 500Hz both the DS-1 and Big Muff have been blurred to the point of being illegible, showing their extreme, respective intensities at lower frequencies. The main discernible difference between the spectrograms of the DS-1 and the RAT is the blurred nature of the lowest frequencies, on top of the slightly higher peak threshold frequency the DS-1 registers. The DS-1's high gain spectrogram proves it to be the most consistent, when faced with increasing gain. At high gain levels, the graph shows regular markings at 6kHz, a threshold only previously

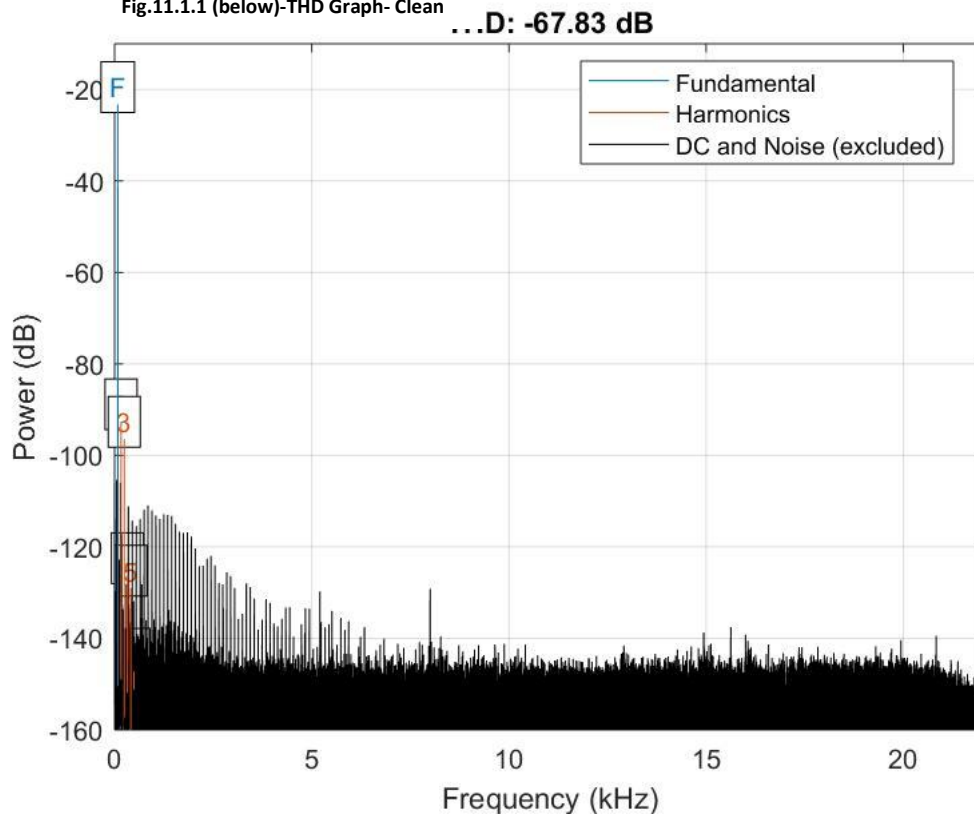
approached, rather than reached, by the results at lower gain levels. Of any pedal, the DS-1 spectrograms show it to be the least affected by changing gain levels as a variable. This includes the recurring feature of the shrill, single note section of the audio sample registering intensities at far higher frequencies than any of its pedal counterparts, irregardless of gain level. Barring marginal increases in intensity, most present above the shared absence band (>3kHz), the RAT shows very little change in its spectrograms generated for medium gain and high gain. If anything, the graph seems to indicate decreased intensity between 1000-2000Hz, the first time a result such as this has been observed when the gain has been raised. With little to no changes visible to the human eye, it's easily established that ProCo's product matches Boss' in their spectrogram consistency against a changing variable of gain. In stark contrast, the final spectrogram for the Big Muff shows explicitly that Electro Harmonix's device as changing the most, sonically, with different levels of gain. With the flecked outliers now touching the 6kHz threshold, the change from the initial, low gain graph is conspicuous. At low gain, every spectrogram was visibly discernible from each other and had obvious identifying features; at high gain, however, below approximately 2600Hz, the spectrograms are remarkably similar. All feature high levels of intensity consistently from 0Hz up to this threshold.

Ch.11- Total Harmonic Distortion

	<u>THD</u>
Clean Amp	-67.83dB
Tube Screamer- Low Gain	-21.10dB
Tube Screamer- Medium Gain	-16.68dB
Tube Screamer- High Gain	-13.19dB
DS-1- Low Gain	-14.46dB
DS-1- Medium Gain	-12.55dB
DS-1- High Gain	-11.00dB
RAT- Low Gain	-11.68dB
RAT- Medium Gain	-7.31dB
RAT- High Gain	-7.34dB
Big Muff- Low Gain	-16.55dB
Big Muff- Medium Gain	-16.00dB
Big Muff- High Gain	-22.32dB

Table 11.1 (above)- THD

Fig.11.1.1 (below)-THD Graph- Clean



Listed above are the numerical results from the total harmonic distortion calculations for each relevant audio sample (Table 11.1), along with the matching graphical results for the clean amplifier signal (Fig 11.1.1).

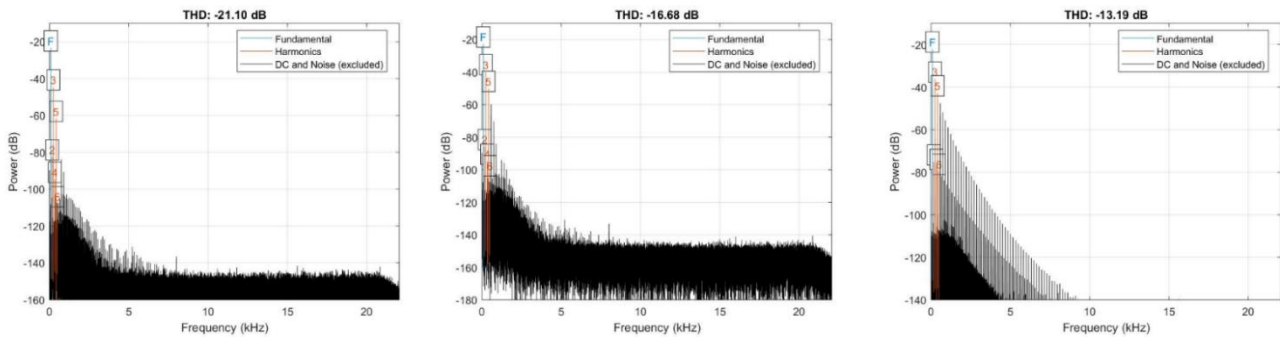


Fig.11.1.2- THD- TS9

The Tube Screamer's THD results (Fig. 11.1.2) provide a good analytical foundation to remark on the auditory distortion properties of the device. The low gain graph shows minute increases in distortion strength from the clean signal, although not remarkable amounts. The main graphical body displays consistent results of just below -150dB, with the initial exception of the early peak extending to approximately -105dB. This is indicative of a non-aggressive distortion palette, whereby only moderate levels of clipping and gain are present, hence why the signal is not massively altered in any way. The medium gain graph shows distinct continuation from its low gain counterpart. General structure and shape are retained, with general increases in distortion strength observable throughout. This sees the main body of the graph now sit at close to -145dB showing slight alteration but not seismic change. This medium gain graph also enables the observation that the signal begins to wane below approximately -175dB showing its main affectable THD-band lies between -120dB and -175dB. The high gain graph, surprisingly, sees unrecognisable change from its previous iterations. The affected frequency spectrum is greatly reduced, but visualising stark increases to distortion level in this shorter range. At the inception, beginning at 0dB, the graph now rises to -40dB. From this point, the graph steeply declines in consistent fashion. From these results, it is a fair conclusion to draw that the TS9 is fairly mild in the auditory distortion that it applies to a given signal.

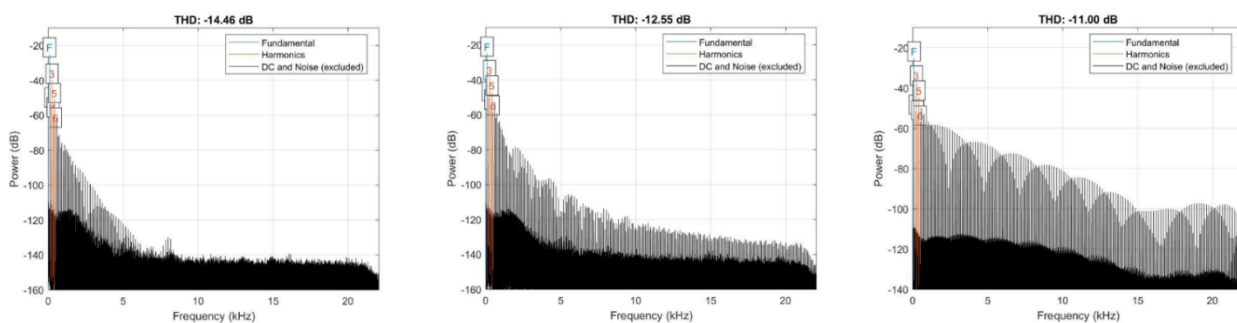


Fig.11.1.3- THD- TS9

Compared to the Tube Screamer, the THD graphs for the DS-1 (Fig. 11.1.3) show more lively readings, indicating a more effervescent response to increasing levels of gain. The low gain iteration shows remarkable similarity to the Tube Screamer, similar overall shape and intensity of readings. The DS-1 readings are initially higher than the Tube Screamer, however, extending over 20dB higher. From 5kHz onwards, there is very little difference between the Tube Screamer at low gain, and the DS-1 at low gain. The

differences become more immediately obvious at medium gain; the most intense readings below -140dB remain consistent, the secondary markings extending consistently above this threshold distinguish this graph from its low gain counterpart. While distinctly higher and more pronounced, graphically, than previously, medium gain still did not fully demonstrate the distortion properties linked with the DS-1 prior to this phase of testing. The graphical representation of the DS-1 at high gain does fulfil some of these auditory distortion properties however, in the results it displays. The lower, most-intense body of the graph sits approximately 20dB higher than previously, now resting at -120dB until just after 12.5kHz where it then dips to around -130dB. The secondary markings (indicative of additive transient noise) are far more extreme on this graph, illustrating the more unhinged properties linked with the pedal throughout the subjective, qualitative stages of experimentation. The majority of these markings register above -100dB showing just how sonically impactful these additions would be. These results almost conclusively demonstrate that the DS-1 fulfils its reputation as being more aggressive and untamed with its auditory distortion effect, than the Tube Screamer, against direct comparison.

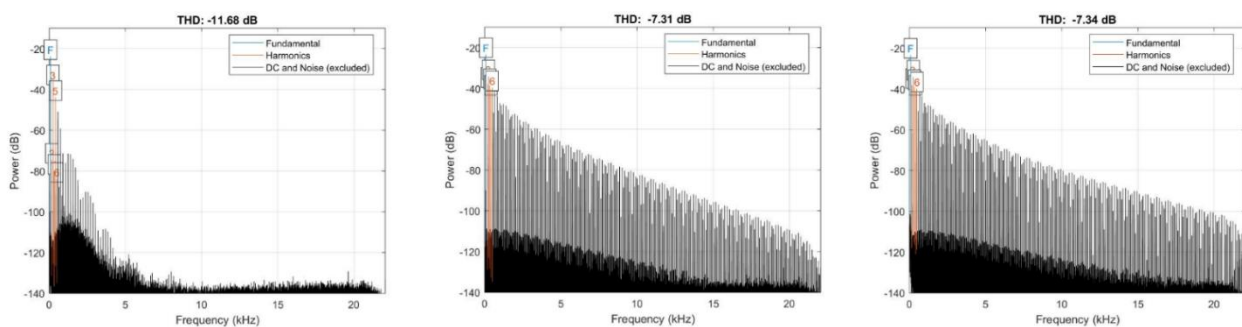


Fig.11.1.4- THD- RAT

The RAT again throws up interesting, and seemingly anomalous results at the THD stage of testing, also (Fig. 11.1.4). The graph representing the signal at low gain shows remarkably sparse readings. The initial peak rises to just below -100dB but then quickly recedes to approximately -135dB where it plateaus and remains well beyond 20kHz. These results match the initial, low gain results for both the Tube Screamer and the DS-1 also. At medium gain, the graph rapidly transforms and sees the consistency of the highest intensity readings alter to a slope, beginning at -110dB and ending just above -140dB after passing the 22.5kHz marker. The readings for transient noise also present far more intense results than were expected. Mimicking the curve presented by the raw THD data, the noise exists within the sonic territory up to 40dB higher than the main affected body, at any given point. This seems demonstrative of the fact that the transition between low and medium gain has a massive impact on the qualities (and level) of distortion present in the outputted signal. Most surprising is the change, or lack thereof, when transitioning from medium to high gain. The first half of the graph, 0-12.5kHz, displays only very marginal increases from its medium gain counterpart, with the main shape of plotting and structure remaining intrinsically similar. The remaining second half, 12.5-22.5kHz, still only displays very small increases however they are slightly more discernible than previously. Fundamentally, these results prove that the largest change in distortion occurs when transitioning between low and

medium gain, with the change between medium and high resulting in only very minute changes to the sonic qualities imparted by the pedal itself. The quality of the distortion seems to exist in a similar sphere to the DS-1 but with a more noticeable decline from low end to high, indicative of an emphasis on bass response with slightly less shrill treble qualities.

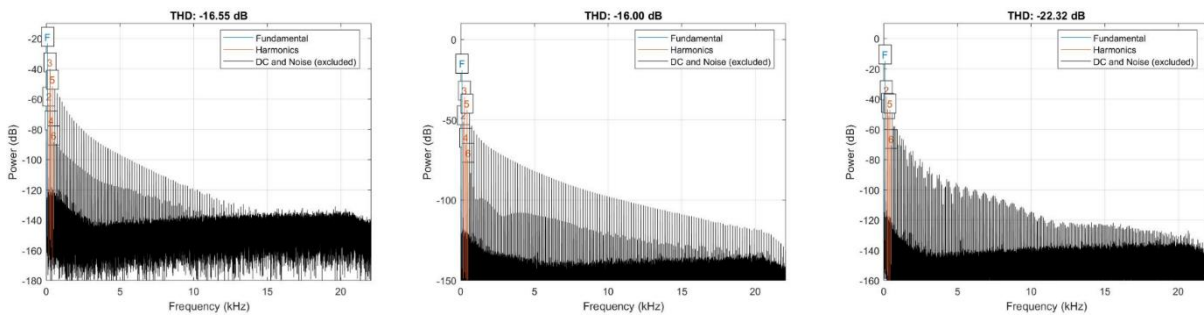


Fig.11.1.5- THD- Big Muff

The THD results for the Big Muff (Fig. 11.1.5) provide a return to consistency. At low gain, the main graphical body sits at a resting point of -140dB with slight fluctuation and a marginal rise from 5kHz onwards. Initially, up until the 12.5kHz marker, there is a noticeable level of transient noise present, however this rapidly decreases until 15kHz from which point, it is entirely negligible. Also noteworthy is the absence of any markings below -180dB aptly demonstrating the lower sonic limits of the Big Muff. The transition to medium gain sees the graphical structure generally retained. The transient noise has increased so that it is now visible across the entire graph; until the 10kHz marker, this transient noise constantly exists above -100dB. The main intensity plotting for the general THD remains consistent with the previous iteration, however. This trend is continued into high gain, with the main readings remaining remarkably consistent, hovering around -140dB generally. The external, transient noise also continued its increase also; at high gain, this noise exists at a higher intensity, but at lower frequencies. This is denoted by the lack of separation into two distinct bands of plotting, and therefore the darker colour implying greater intensity. These results show that the greatest alternating factor in the Big Muff's signal with changing gain, is the level of transient noise present.

Ch.12- Brightness

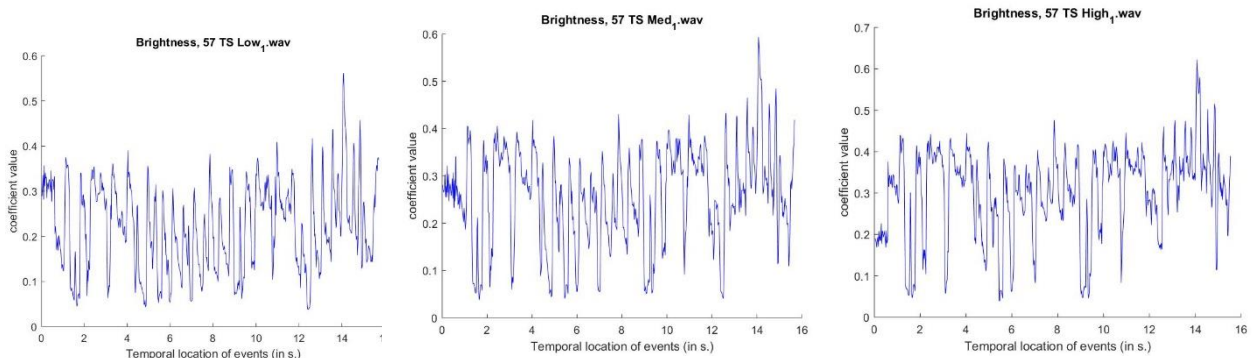


Fig.12.1.1- Brightness Graphs- TS9

The brightness of each audio sample provides valuable timbral insight. Starting with the least volatile of the four pedals, the Tube Screamer, the graphs pictured (Fig. 12.1.1) show the differences in brightness between low, medium, and high gain settings (using single coil pickups, favoured here for their brightness compared to their humbucker counterparts). The low gain graph shows consistent fluctuation between 0.05 and 0.4 for the first 10 seconds of the sample. The peak at 14 seconds, measuring approximately 0.55 is interesting as this is the section of the sample where the guitar part strays from block chords to a single tone focused ostinato. This result implies that the Ibanez Tube Screamer lends a higher level of timbral brightness to the more shrill, higher pitches than to the lower, chordal sounds. Barring the early troughs, in a surprising set of results, the high gain setting for the Tube Screamer was most consistent in specific regard to brightness.

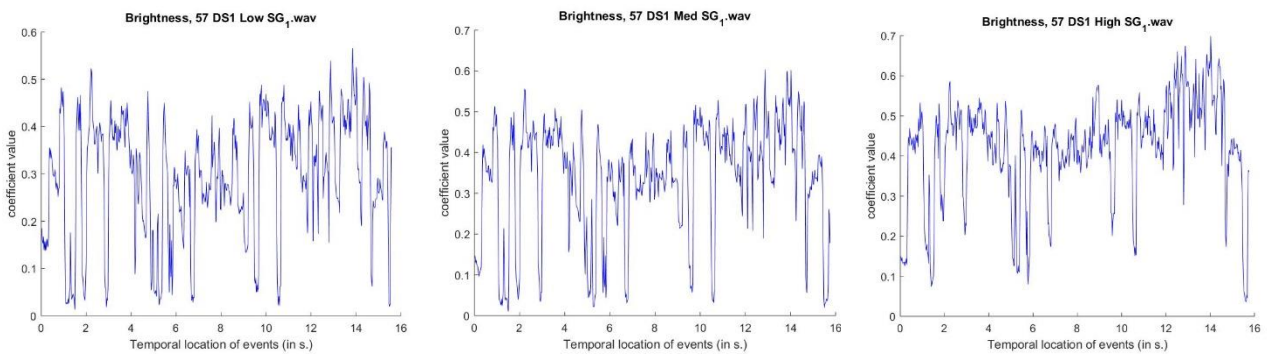


Fig.12.1.2- Brightness Graphs- DS-1

The Boss DS-1 is generally far brighter than the Tube Screamer. The results explicitly show this (Fig. 12.1.2), even when using the audio recorded through humbucker pickups which are duller compared to single coil pickups. While the main body of the graph mainly sits below 0.4 for the TS9, the DS-1 far exceeds this and is often peaking towards 0.5. That being said, the plotting for the DS-1 is much more volatile than the Tube Screamer and fluctuates between extremes far more readily. The low gain settings show the most fluctuation, while again, the high gain provides the most consistency. The DS-1's peak measurement for brightness was at 0.7 after 14 seconds, an entire 0.1 higher than the maximum brightness achieved by the Tube Screamer. Unusually, the low gain and medium gain graphs show very little difference in their results with the medium gain only being marginally brighter, generally. Similar to the TS9 however, the high gain settings

demonstrably showed the most consistency in regard to sonic brightness. The results for both these pedals match the expectation set out before the results were gathered; the DS-1 is noticeably brighter and fluctuates more often and more violently than the milder Tube Screamer.

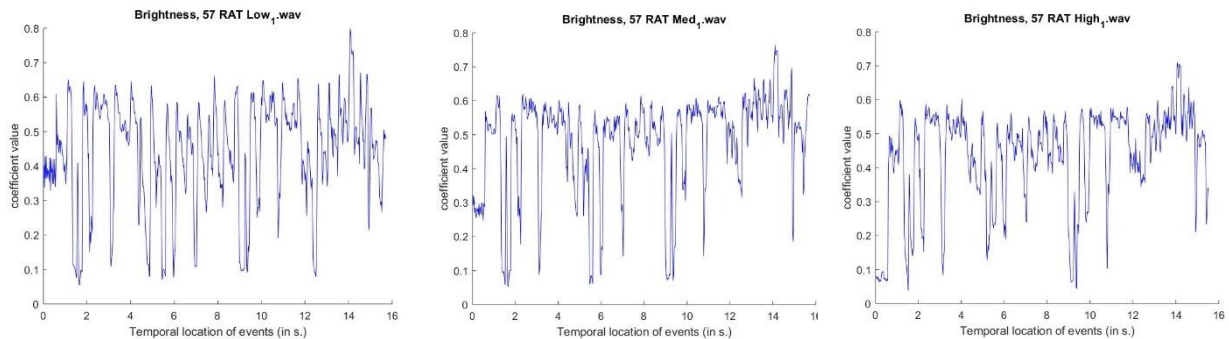


Fig.12.1.3- Brightness- RAT

The ProCo RAT’s brightness appears to be the most consistent of any of the pedals, irregardless of gain setting (Fig. 12.1.3). In contrast to the Tube Screamer and the DS-1 however, as the gain increases on the RAT, the timbral brightness decreases. The troughs visible in the RAT’s graphs as well indicate as the gain increases, the majority of the troughs get less severe. This explicitly shows therefore that the brightness gets narrower, giving the outputted audio more sonic focus. This, in many ways, is the opposite of what the other pedals seem to aim for which is more expansive brightness with increasing amounts of gain, perhaps to compensate for the harder levels of clipping also. This all considered, at its peak, the RAT also shows the highest level of brightness recorded for any of the four devices. On the lowest gain setting, the final peak approaches 0.8 which is far brighter than any setting the DS-1 or Tube Screamer can muster. It is worth noting that the chosen graphs for the RAT were taken from recordings using single coil pickups. The added brightness of these compared to the humbucker audio used to measure the DS-1 should be considered when making direct comparisons. The results for the RAT, however, do fit the observable pattern of the high gain settings providing the most consistent results when measuring the brightness.

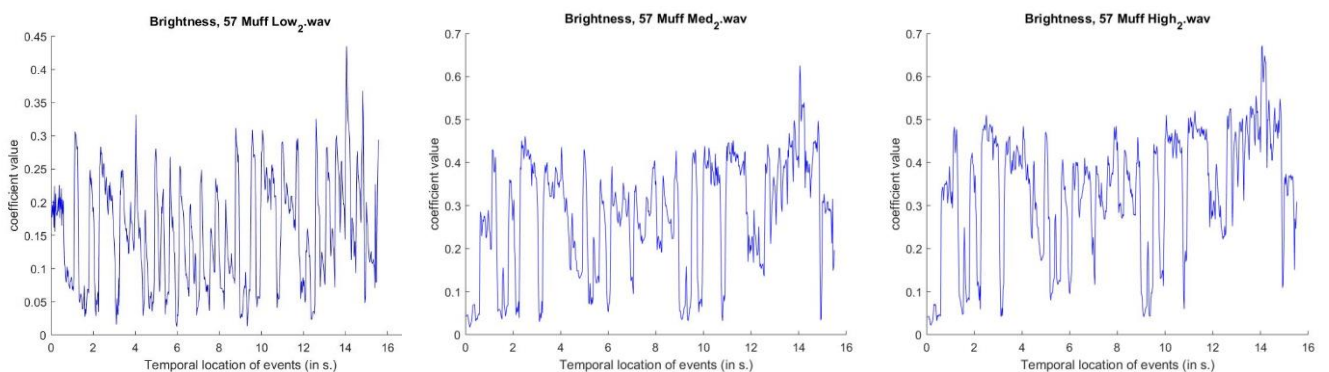


Fig.12.1.4- Brightness- Big Muff

The Big Muff's brightness results were easily the most consistent (Fig. 12.1.4), though this is perhaps due to its control settings not having a pot to exclusively set the gain. For comparison's sake, the results from the humbucker set of audio recordings was used, due to them being unusually far brighter than the single coil results. The big muff has the potential to be the least bright of the four pedals, as the deep troughs actively demonstrate but also one of the brightest, as the peaks approach 0.7. This will likely be thanks to the excessive extra levels of harmonic content that are generated by the extreme levels of clipping and gain added to the signal. Electro Harmonix's flagship unit therefore has the most extreme results at either end of the brightness spectrum. This could be partially negated by the fact that it has far more fuzz-related properties than any of the other three featured units, however it is fair to call this the most 'full' sounding pedal in specific regard to brightness. The most fluctuation can be observed in the plotting of the graph for the lowest gain settings.

Ch.13- Roughness

Roughness is a quantitative measure of rapid amplitude modulation. The unit used to measure this effect is the asper. Creating graphical representations of these measurements allows examination of how coarse a signal or sound, is. For example, the graph for the clean amplifier signal shows the peak roughness measurement to be approximately 160 aspers.

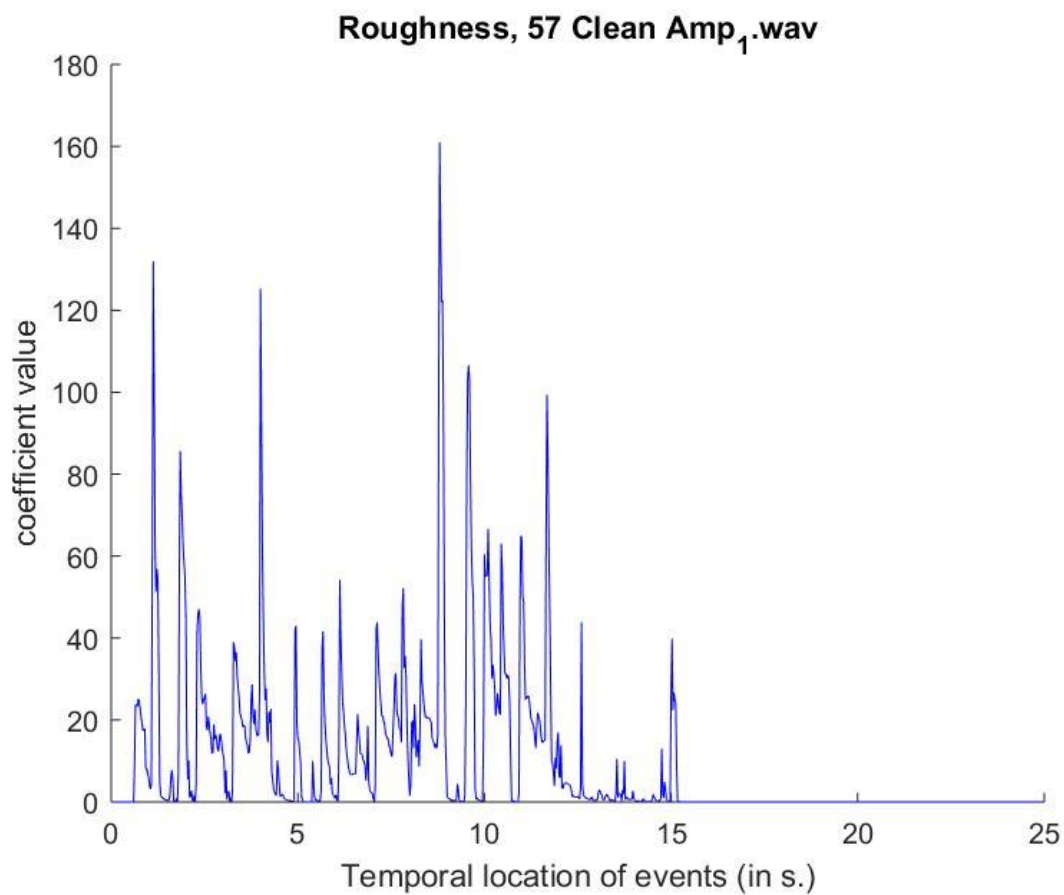


Fig.13.1.1- Roughness- Clean

As a point of reference for every other graph, this initial graphical representation (Fig. 13.1.1) shows the changing levels of roughness of the audio sample, along the approximately fifteen second timeline. As this is essentially a clean signal, any subsequent fluctuation to the roughness level can be considered as solely the responsibility of the given pedal.

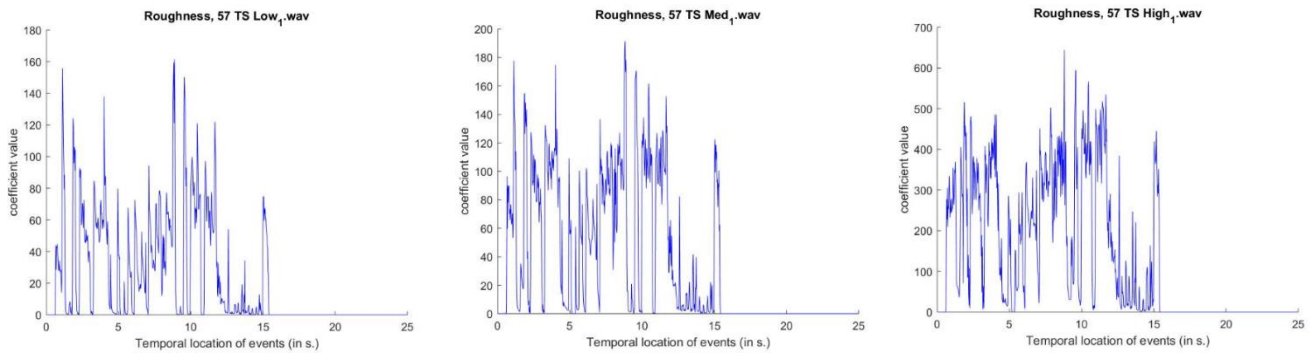


Fig.13.1.2- Roughness- TS9

The Tube Screamer at low gain (Fig 13.1.2) shows marginal increases from the roughness graph for the clean amplifier. General increases of approximately 20 aspers can be observed throughout the recorded peaks, with the notable exception of the largest peak which remains largely untouched by the device at low gain. The change to medium gain follows a similar trend, general rises of approximately 20-40 aspers. By this point, any shared resemblance between the Tube Screamer’s graphical qualities and the original, clean amplifier have been all but lost. The highest recorded peak increases to 180 aspers at medium gain. In stark contrast to the small, marginal increases observable from low to medium gain, and even clean amplifier through to medium gain, the roughness graph for the Tube Screamer on high gain sees sharp, drastic increases from its counterparts. The general, median weighting of the results seems to have doubled, now sitting close to approximately 300 aspers, with peaks exceeding the 500 aspers threshold easily. These stark increases are indicative that auditory roughness stemming from the distortion only takes real effect in the final third of the Tube Screamer’s gain control, when compared to the clean amplifier sample.

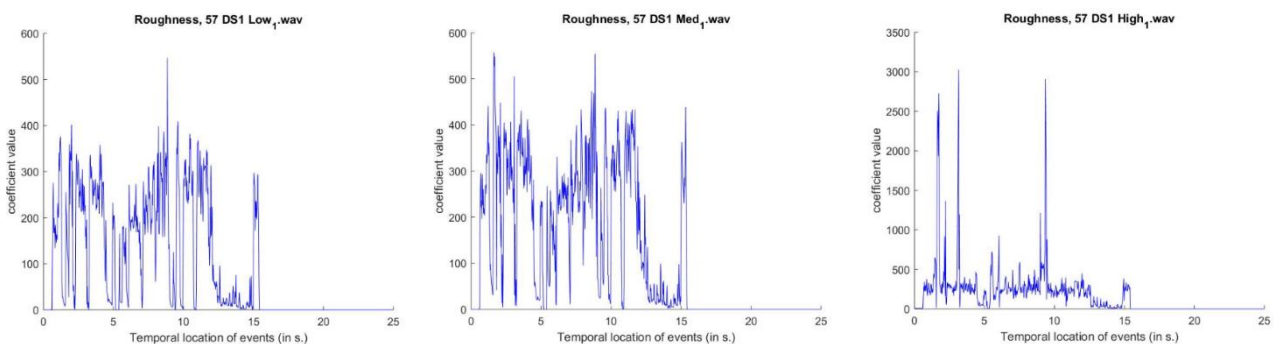


Fig.13.1.3- Roughness- DS-1

The line of graphical results for the Boss DS-1 (Fig. 13.1.3) are expectedly far more turbulent than the results gathered for the Tube Screamer. The first graphical iteration at low gain, already far supersedes the comparatively nominal results for the clean amplifier sample; recorded values beyond 500 aspers, more than triple the value of the highest peak with the clean amplifier, confirm this. It is more than fair to say that, even at low gain, the DS-1 creates an auditory distortion of strength more than double that of its original, clean counterpart. The final peak, just sitting shy of 40 aspers originally, measures at just under

300 aspers with the DS-1 on its lowest gain setting. At medium gain, there is a contextually gradual increase, not as large as the jump from clean amplifier to DS-1 on low gain, however still significant. Results slightly beyond 550 aspers are recorded at medium gain, again greatly beyond the upper threshold reached with no pedal interaction. The true degree of auditory roughness for the DS-1 is determined with the settings on high gain, in this instance. Closely mimicking the stark contrast and upward fluctuation observed in the Tube screamer between medium and high gain, the DS-1 records far higher roughness results at high gain. With the highest peak breaking the 3000 aspers threshold, many of the results recorded under these conditions are the highest found in this experiment. These results are very interesting as they are the first to radically move away from the general plotting pattern observed in every previous graphical representation. While the peaks usually extend 100-200 aspers beyond the main body of the graph, at high gain, the DS-1's peaks stand apart from the rest of the graphical body in an extreme way, often more than 2500 aspers above where the peak begins and returns to. This plotting implies the median roughness does not vary massively, in spite of the peaks fluctuating in size hugely between gain variants.

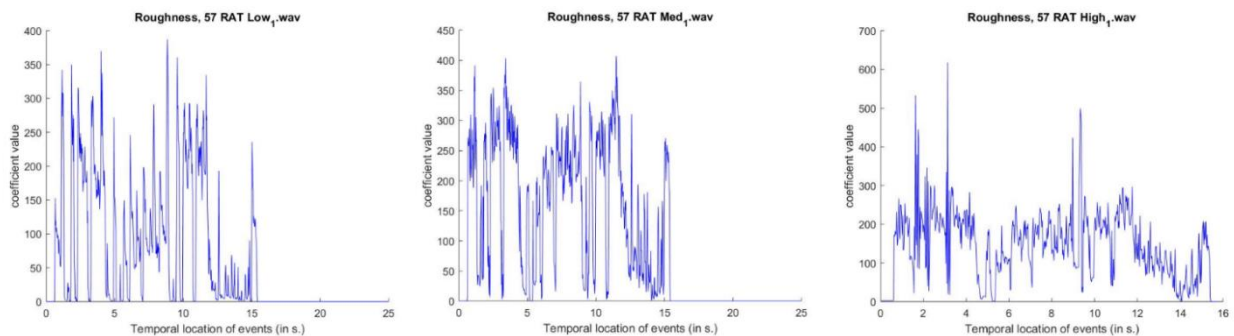


Fig.13.1.4- Roughness- RAT

The graphical results for the RAT (Fig 13.1.4) present a contextual return to normalcy, after the extreme results of the Boss DS-1 on high gain. At low gain, the RAT's roughness results show a distinct raise from the levels recorded for the clean amplifier and the Tube Screamer. The highest peak at low gain extends to approximately 380 aspers, while the median roughness seems to exist around the 150 aspers mark. The RAT then follows the established trend of rising incrementally as the gain increases, unlike its previous counterparts however, this extends along the full range of the gain scale. For medium gain, therefore, the highest peak encroaches just beyond the 400 aspers threshold, with the median existing at roughly 225 aspers. The graphical representation of the results at medium gain seems far more compact than low gain, with the plotting between 200-350 aspers seeming more pronounced. The high gain results see their top peak extend beyond the 600 aspers threshold, similarly to the Tube Screamer, although reached far more incrementally. The median exists at approximately 200 aspers, seeing an apparent decrease from medium gain. The RAT presents arguably the most consistent set of auditory roughness results; steady increases in the levels of roughness correlating directly with the incremental increases in gain.

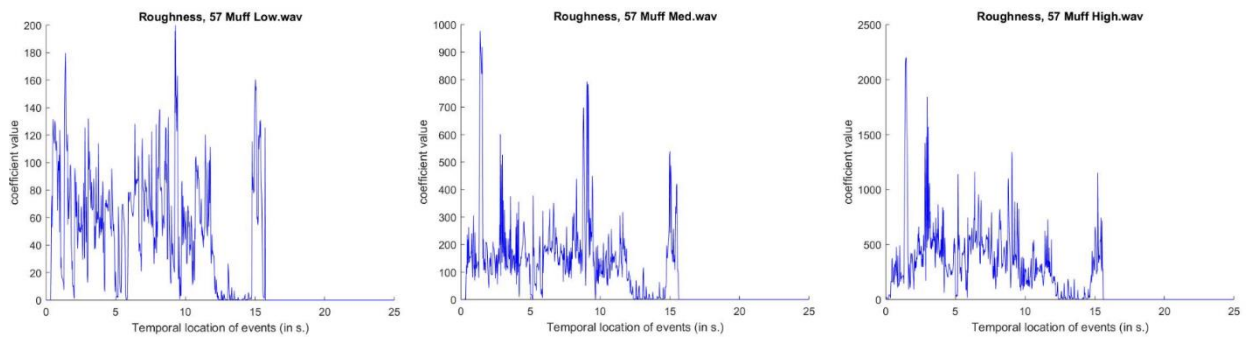


Fig.13.1.5- Roughness- Big Muff

The Big Muff's roughness results (Fig. 13.1.5) follow a similar upward trend but more aggressive in its increases. At low gain, a very active graph sees the highest peak exceed values of 200 aspers, however the median lays significantly lower at approximately 70 aspers. This is followed by a steep rise at medium gain with the highest value falling just short of 1000 aspers, again with a significantly lower median of around 150 aspers, in a somewhat less active graph. The graphical representation of auditory roughness at high gain sees the tallest peak reach values of 2200 aspers. The median for high gain works out to be an estimated 400 aspers and this graph is arguably the least active of all graphical iterations for the Big Muff.

Ch.14- Spectral Centroid & Flux

While the spectrograms provide great general insight into the effect of each pedal on the source signal's frequency range and the varying intensities within their affectable ranges, and THD allows closer analysis of the levels of distortion present in the given sample, calculating the spectral centroid gives accurate indications of where the median frequency lays for each audio sample provided. For added transparency, this is provided both as a numerical value accompanied with full graphical representation. As a reference guide, initially displayed is a table comprised of every spectral centroid numerical value for every audio sample, both single coil and humbucker pickup variants included. This allows for easy comparison of these values alongside each other. Spectral Flux can provide another good graphical means of analysing the timbre of a given signal. It allows this by showing the rate of change of the spectrum. Lower values signify slow spectral changes, whereas higher values indicate far quicker changes to the spectrum.

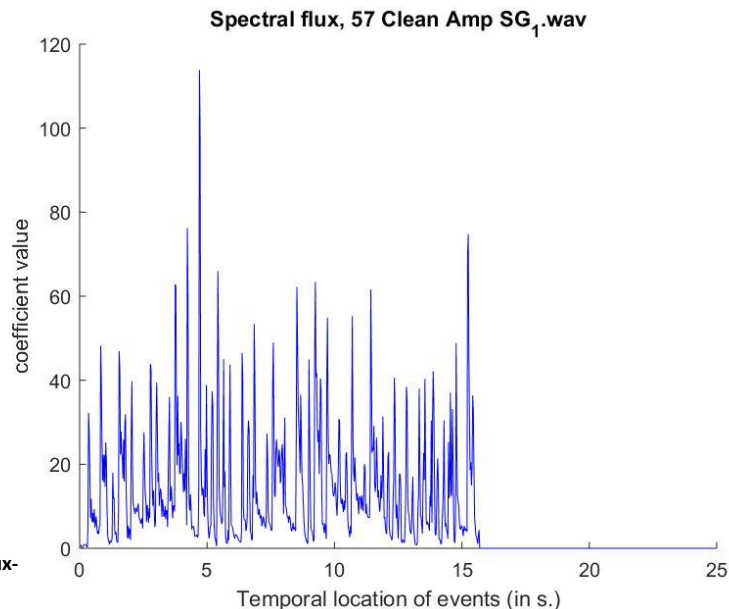


Fig.14.1.1- Spectral Flux-Clean

The clean amplifier signal, through humbucker pickups, shows a relatively low spectral flux level (Fig 14.1.1). The main body of the graphical content exists below 40dB, with only a singular peak exceeding the 100dB threshold. With only a singular, notable exception, each device observes the same trend of the spectral centroid rising as the gain is increased. This exception was the results for the RAT at high gain which for both pickup types saw a significant decrease (approximately 150Hz lower) from its results at medium gain. Generally, the single coil audio samples returned higher spectral centroids than their humbucker counterparts; the Big Muff was the only pedal not to display this, with the humbucker results registering over 100Hz higher than the single coils when set to high gain.

Examination of each pedal's spectral centroid graphs respectively, gives clearer indication as to why these exceptions exist, along with closer frequency analysis for each unit.

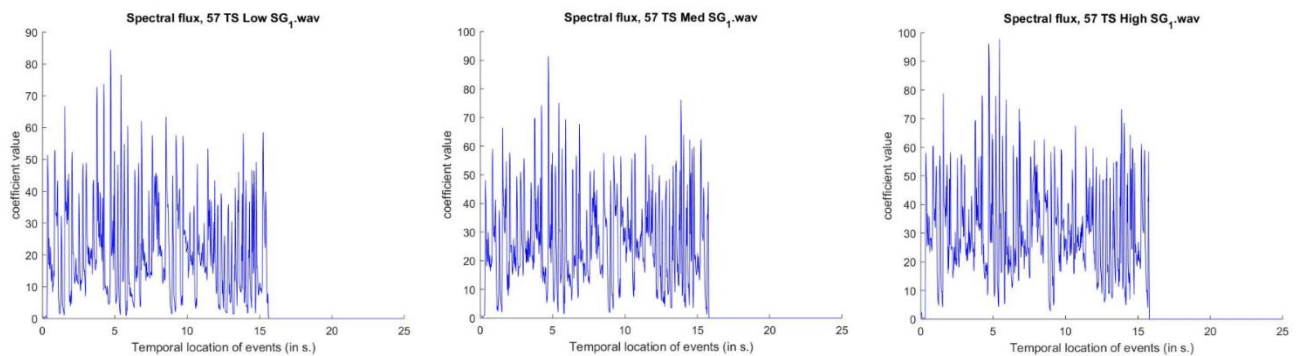


Fig.14.1.2- Spectral Flux- TS9

The flux graphs over the 3 gain iterations provide a great foundation for timbral analysis. As Fig. 14.1.2 shows, for the Tube Screamer, the low gain graph indicates that the TS9 provides certain blunting sonic properties. This is due to the peaks existing at lower levels than the clean amplifier signal. Without any pedals present, the clean amplifier graph's tallest peak exceeds 110 whereas the Tube Screamer at low gain doesn't register a peak exceeding 90, a significant decrease. The majority of graphical plotting content exists in the 0-50 range with the average laying at approximately 28, indicating a general decrease in flux levels when compared to the unaltered clean amplifier signal. At medium gain, the general structuring of the graph bears a striking resemblance to its low gain counterpart, with observable marginal increases. The most prominent peak now reaches beyond 90, with the second-most prominent seeing a steeper increase of around 20. Likewise, the average of the flux content exists at approximately 35, confirming an observable trend of general increases in flux content as the gain level increases. This trend, predictably, continues as the gain is increased to its highest setting. Interestingly, the average at high gain is around 40dB, approximately identical to the average perceived through simply the unaffected, clean amplifier signal. From these results, it can be deduced that the Tube Screamer's 'warm' properties relate to a loss of treble presence due to decreased fluctuation which translates to lower peaks and generally a bassier spectral output.

The Tube Screamer has the closest spectral centroid to the clean amplifier results, with both pickup-type results at the lowest gain settings existing at $\pm 250\text{Hz}$ of the clean amplifier's initial centroids. The single coil results exist only 72.54Hz apart from each other demonstrating how close the spectral centroid for the TS9 exists to the unaltered clean signal's centroid. As the gain rises, the Tube Screamer's spectral centroid does not rise unpredictably, instead increasing by consistent margins, showing the level of gain to steadily alter and raise the median frequency, rather than radically altering the affectable frequency band. The spectral centroid for the TS9 at low gain is 1124Hz, at medium gain 1260Hz and at high gain, 1320Hz. These numerical results prove that the Tube Screamer's sonic presence impacts the least on the spectral centroid of the signal.

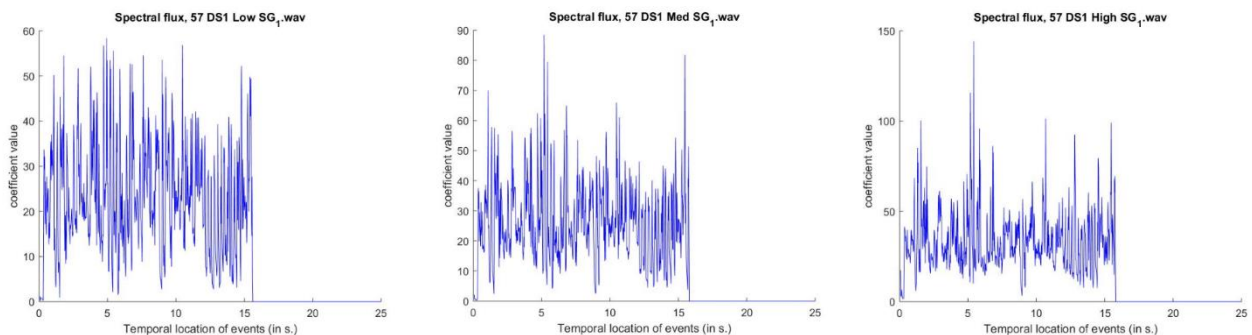


Fig.14.1.3- Spectral Flux- DS-1

The DS-1 set of spectral flux graphs (Fig.14.1.3) provide very different, yet equally valuable timbral results. As expected, the pedal at low gain presents the least fluctuation of any gain setting. The largest peak does not even reach the 60 threshold, indicating a similarly subdued fluctuation pattern, when compared with the clean amplifier graph. The main body of the fluctuation occurs between 10-40 with the average existing, in consistent fashion, at 25. The compact structure of the plotting for the low gain graph is lost with the progression to medium gain, with much higher levels of fluctuation immediately visible. The peaks now extend beyond 85, with these peaks standing more noticeably above the main body of plotting. One particularly interesting result at medium gain, is the result for the final chord of the recording: the flux content according to the graph extends from 0-80 demonstrating extreme levels of rapid fluctuation. The transition to high gain sees drastic change to the results for spectral flux. The peaks at high gain reach beyond the highest recorded value with the clean amplifier, for the first time. The highest peak sits just below 150, indicating very high levels of sonic fluctuation present just after the 5 second marker. The majority of the plotting still exists far lower however, at around 40 regaining its similarity to previous iterations at lower gain levels.

As the numerical results confirm, the spectral centroids for the Boss DS-1 sit far higher than the Tube Screamer and subsequently, significantly raised from the centroids recorded for the clean amplifiers. At its highest recorded point (single coil pickups with the gain settings on high), the DS-1's spectral centroid exists just shy of 1kHz (954.05Hz) above the equivalent centroid through the clean amplifier. This notable upward shift in median frequency values is best exemplified on the graphical representation of the aforementioned high gain-single coils output, where the plotting takes a sharp rise after 15 seconds reaching frequencies approaching 4500Hz. This is the highest recorded frequency value gathered from this section of experimentation and signifies the DS-1 validates the contextual usage of words such as 'shrill' to describe its sound. While increases are explicitly visible between low and medium gain also, the sudden rise when at high gain is distinctly more sudden and intense than any trend noted prior. The spectral centroid for the DS-1 at low gain is 1566Hz, at medium gain it's 1719Hz, while at high gain it exists at 2006Hz.

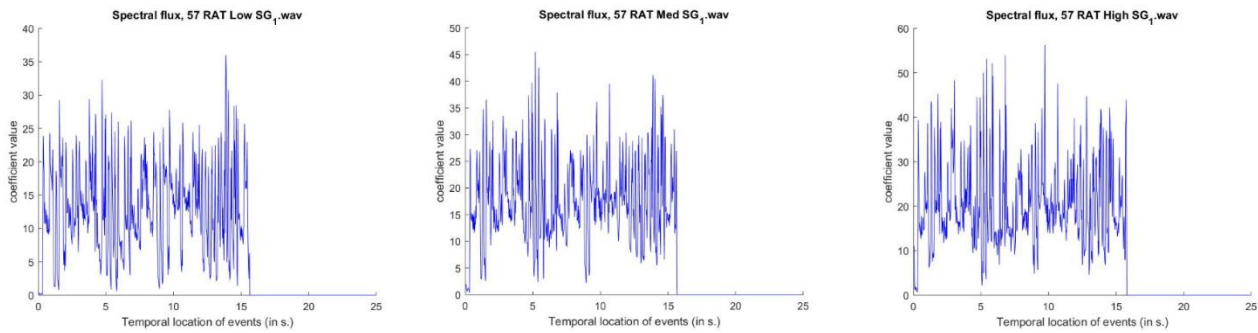


Fig.14.1.4- Spectral Flux- RAT

The spectral flux results for the RAT (Fig 14.1.4) present a more subdued response than expected. At low gain, the highest peak barely exceeds the 35 threshold. The majority of the content exists beneath 25 with the average level of fluctuation existing at approximately 15. This is indicative of the RAT not sonically fluctuating very much when set to lower gain levels. It could be stated, therefore, that the RAT is spectrally consistent at low gain. The rise to medium gain sees the results follow a similar pattern as observed in the Tube Screamer and the DS-1. General, consistent increases see the peaks reach highs of 45. The average level of fluctuation has also increased to approximately 20 and the general shape of plotting has been retained from the low gain graphical representation. Unlike the DS-1, however, the final iteration at high gain does not see a radically steep inclination to the flux results. Instead this follows the steadier trajectory observed also with the Tube Screamer. This sees the graph's highest value at around 56 while the average has also steadily rose to approximately 25. These results indicate that the RAT does not fluctuate greatly sonically, regardless of gain level. This could be a valid rationale as to why it is perceived to be similar to the DS-1 but slightly tamer.

The RAT displays the least consistent results of any pedal tested. This is best demonstrated by the spectral centroid results between medium and high gain, using humbucker pickups. The majority of spectral centroids results follow an observable pattern of steadily increasing with each iteration of the graph, as gain increases. The RAT does not follow this trend in any way. The numerical results of the spectral centroid at low gain being 1822Hz, at medium gain existing at 1936Hz and high gain being 1801Hz explicitly show the spectral inconsistency that the RAT presents.

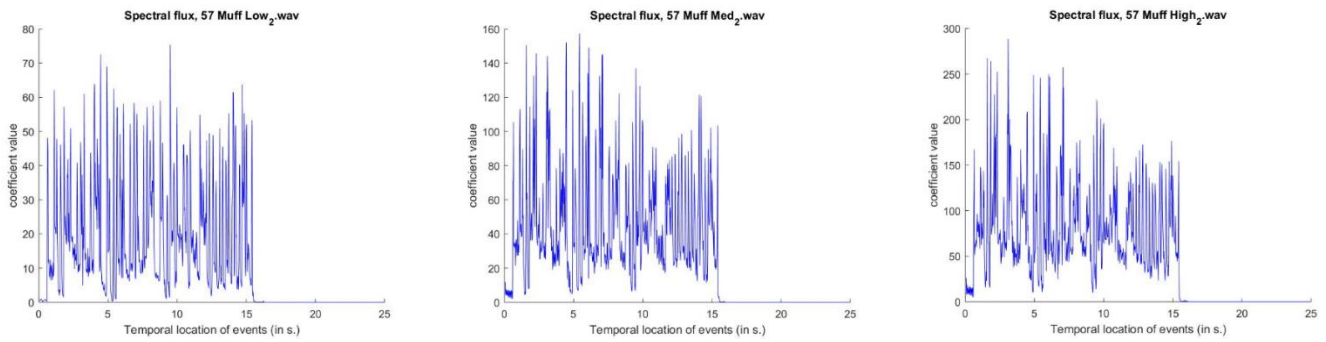


Fig.14.1.5- Spectral Flux- Big Muff

The spectral flux results for the Big Muff (Fig. 14.1.5) were the most extreme out of any tested, both in terms of progression and general plotting. The low gain graph shows very active results, with closely grouped markings fluctuating greatly. The average of the readings lays at approximately 30, likely the highest of any of the pedals. The peaks at low gain extend beyond 70, roughly 75 demonstrating that even at low gain, the Big Muff features high levels of sonic fluctuation. Progressing onwards to medium gain, the transition shows bigger changes than any previous pedal recorded, between these two gain settings. The average starkly spikes to around 70 and the same peak now extends to beyond 150 showing a dramatic increase in the highest level of fluctuation present, shortly after the five second mark. This radical upwards shift is observable again in the transition from medium to high gain. The tallest peak now reaches approximately 280, the highest value recorded in this experiment; the average of the results also exceeds the previous highest, now existing well beyond 100. These results prove the Big Muff is the most volatile pedal tested, in regards to sonic fluctuation. The highest values recorded emerged from the Big Muff tests and this shows its flux levels to be distinctly higher than the other devices tested.

The spectral centroid results for the Big Muff show a return to the same trend as observed in the Tube Screamer and DS-1, in contrast to the RAT: as gain increases, the spectral centroid increases. This trend is not as strong as the TS9 or the DS-1, however. With single coil pickups, transitioning from low to medium gain shows explicit graphical increases. The same rise exists between medium and high gain, although this increase is not as significant. Oddly enough then, the Big Muff's spectral centroid is the second closest, numerically, to the clean amplifier centroid, after the Tube Screamer. At low gain, the centroid is 1083Hz, at medium gain it is 1554Hz, while at high gain the centroid is 1756Hz. Given its fuzz-like properties, both colloquially and as can be seen in the total harmonic distortion experiment results, the result of its centroid sitting closer to the clean amplifier centroid value than the DS-1 or the RAT was somewhat unexpected.

Conclusion

To summarise, through text mining, the most frequent adjectives used to describe the sonic properties of distortion pedals were discovered to be: Crunchy, Warm, Smooth, Hot, Fuzzy, Noisy, Dirty, Muddy, Thick, Thin, Woolly, Creamy, Fizzy, Harsh, Fat, Bright, Dark, Crispy, Sludgy, Full. While 'crunchy' emerged as the most popular choice of descriptor overall, the favourite for each individual pedal was: Tube Screamer- 'warm', Boss DS-1- 'noisy', ProCo RAT- 'dirty' and Big Muff- 'fuzzy'. The next stage was to conduct a descriptive survey with the aim of validating the descriptors unearthed in the initial text mining experiment, but also to uncover new adjectives when prompted aurally. 'Crunchy' emerged with the highest usage again, confirming its importance and strong contextual usage. The new descriptors that emerged tended to feature onomatopoeia more commonly, this can be linked to the listeners attempting to phonetically mimic the distorted tones they were hearing. These descriptors were adjectives such as 'raspy' and 'saturated'. With the lexicon beginning to take shape and the descriptors most popularly used having been established, it was essential to redefine these adjectives to better fit the contextual sphere of distortion in which they were being specifically deployed. Creating these new definitions was achieved by means of panel discussions, with constant revisions over three sessions allowing a consensus to be achieved for each descriptor. This concluded the etymological element of the research; the relevant descriptors had been uncovered and redefined for the specific distortion context.

Following on, the aim was to research which sonic traits prompt specific descriptor choice. The first element of this involved devising a similarity matrix which would allow users to numerically score each of the twenty most popular descriptors against influential variables such as: gain, clipping and sustain. The results of this were used in the final panel session to be included in the descriptions. The numerical results were also used to create the dendrogram. The dendrogram allowed for visual representation of descriptor similarity, grouped and calculated against the variables specified in the similarity matrix. The dendrogram combined with the results of the matrix were also critical in the creation of the 'Distortion Wheel' which presented a more casual visual representation of each descriptor and the most prominent sonic variable in influencing their usage. The final stage of the subjective, qualitative experimentation was the listening tests. This involved getting participants to listen to four audio samples (one for each featured pedal) and then numerically score them against each of the twenty most popular descriptors. This provided numerical data that could be interpreted into boxplots which visualised the ranking of each pedal. These charts then consequently give an idea as to which pedal best embodies the definitory traits of each descriptor.

The second half of the experimentation was based on ascertaining quantative data, so that alongside the qualitative, opinion-based data, thorough conclusions could be drawn. This involved gathering data for each pedal that provided insight into the prominent sonic measures. For each device, these were: Spectrograms, Total Harmonic Distortion plots, brightness, roughness, flux alongside the spectral centroids. Using the results of these

calculations, precise observations can be made about the sonic signatures of each device. These succinct observations are as follows:

The Tube Screamer's pronounced mid frequencies create a crunchy timbre, while the soft clipping properties give the tone a creamy edge, with the effect of decreased sustain meaning its sound is often also described as crispy.

The DS-1 is characterized by its oversaturated treble response, creating shrill tones that are described as bright, fizzy and harsh. This sits alongside an underwhelming bass response that validates the use of thin, while the intrinsically high gain nature of the pedal means its outputted sound is often hot and fizzy. Transient noise that emanates from the high gain settings lead to the descriptor noisy being levelled against the DS-1.

The RAT is well balanced, if lacking high end treble response, hence being labelled both dark and warm. This balance, coupled with high levels of gain create a thick, full sonic texture while the preference towards hard clipping helps lend the distinctive dirty sound.

The Big Muff's heightened low-end bass response and increased levels of sustain help define the pedal as sludgy, yet smooth. The extremely hard clipping creates a formidable timbre often simply described as fat.

Reference List

- Balanced. (2018). In Collins English Dictionary. Retrieved from <https://www.collinsdictionary.com/dictionary/english/balance>
- Bloated. (2013). In Cambridge Dictionary. Retrieved from <https://dictionary.cambridge.org/dictionary/english/bloated>
- Controlled. (2010). In Oxford English Dictionary. Retrieved from <https://www.lexico.com/en/definition/controlled>
- Distortion. (2020). In Cambridge Dictionary. Retrieved from <https://dictionary.cambridge.org/dictionary/english/distortion>
- Herbst, J-P. (2018). Heaviness and the Electric Guitar: Considering the Interaction between Distortion and Harmonic Structures. *Metal Music Studies*, 4(1), 95-113. https://doi.org/10.1386/mms.4.1.95_1
- Hodgson, P. (2016). The Sound and Style of Dimebag Darrell. *Mixdown Mag*. <http://www.mixdownmag.com.au/sounds-and-style-dimebag-darrell>
- Kossler, M. (2006). *How Nashville Became Music City U.S.A.: 50 Years of Music Row* (1st ed.). Hal Leonard Corporation .
- Murthy, A. A., Rao, N., Beemaiah, Y. R., Shandilya, S. D., & Siddegowda, R. B. (2014). Design and Construction of Arduino-Hacked Variable Gating Distortion Pedal. *IEEE Access*, 2, 1409 - 1417. doi: 10.1109/ACCESS.2014.2374195
- Novak, A., Simon, L., & Lotton, P. (2010). Analysis, Synthesis, and Classification of Nonlinear Systems Using Synchronized Swept-Sine Method for Audio Effects. *EURASIP Journal on Advances in Signal Processing*, doi: <https://doi.org/10.1155/2010/793816>
- Precise. (2010). In Oxford English Dictionary. Retrieved from <https://www.lexico.com/en/definition/precise>
- reddit. (2015). Overpowering Distortion. Retrieved from https://www.reddit.com/r/guitarpedals/comments/29ng2p/overpowering_distortion/
- Stewart, R (2017). Working the Mids. Retrieved from <https://www.justmastering.com/article-working-the-mids.php>
- Total Guitar. (2020). *Pedal gain stacking: everything you need to know*. Retrieved from <https://www.musicradar.com/how-to/pedal-gain-stacking-order-and-more-explained#:~:text=What%20is%20gain%20stacking%3F,to%20mention%20your%20amp%20itself!>
- Unrefined. (2010). In Oxford English Dictionary. Retrieved from <https://www.lexico.com/en/definition/unrefined>
- Yeh, D. T., Abel, J. S., & Smith, J. O. (2007). SIMPLIFIED, PHYSICALLY-INFORMED MODELS OF DISTORTION AND OVERDRIVE GUITAR EFFECTS

PEDALS. 10th Int. Conference on Digital Audio Effects (DAFx-07), Bordeaux, France, Retrieved from <https://pdfs.semanticscholar.org/9903/44e66aa3851e90fc7a825f7481a33be21ecf.pdf>

- Yeh, D. T. (2009). DIGITAL IMPLEMENTATION OF MUSICAL DISTORTION CIRCUITS BY ANALYSIS AND SIMULATION (PhD Thesis). Retrieved from <https://ccrma.stanford.edu/~dtyeh/papers/DavidYehThesissinglesided.pdf>

Bibliography

- Avila, F. R., Carvalho, H. T., & Biscainho, L. W. P. (2016). Bayesian Blind Identification of Nonlinear Distortion with Memory for Audio Applications. *IEEE Signal Processing Letters*, 23(4), 414-418. doi: 10.1109/LSP.2016.2525005
- Berger, H. M. (1999). *Metal, Rock, and Jazz: Perception and the Phenomenology of Musical Experience*. Middletown, CT: Wesleyan University Press.
- Bowers, J., & Yaremchuk, V. (2007). The Priority of the Component, or in Praise of Capricious Circuitry. *Leonardo Music Journal*, 17, 39. doi: 10.1162/lmj.2007.17.39
- Burt, W. (1999). Matthew Burtner: Portals of Distortion. *Computer Music Journal*, 23(4), 112-113. doi: 10.1162/comj.1999.23.4.112
- Chaudhuri, S. (2013). *Structured Models for Semantic Analysis of Audio Content* (Doctoral Thesis). Retrieved from https://www.lti.cs.cmu.edu/sites/default/files/sourish_chaudhuri_structured_model_s_for_semantic_analysis_of_audio_content.pdf
- Chun Tie, Y., Birks, M., & Francis, K. (2019). Grounded theory research: A design framework for novice researchers. *SAGE Open Medicine*, 7, <https://doi.org/10.1177/2050312118822927>
- Defraene, B., van Waterschoot, T., Diehl, M., & Moonen, M. (2016). Subjective audio quality evaluation of embedded-optimization-based distortion precompensation algorithms. *The Journal of the Acoustical Society of America*, 140(1), EL101-EL106. doi: 10.1121/1.4955025
- Doets, P. J. O., & Lagendijk, R. L. (2008). Distortion Estimation in Compressed Music Using Only Audio Fingerprints. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(2), 302-317. doi: 10.1109/TASL.2007.911716
- Gay, L. C. (1998). Acting up, Talking Tech: New York Rock Musicians and Their Metaphors of Technology. *Ethnomusicology*, 42(1), 81. doi: 10.2307/852827
- Grey, J. M., & Gordon, J. W. (1978). Perceptual effects of spectral modifications on musical timbres. *The Journal of the Acoustical Society of America*, 63(5), 1493-1500. doi: 10.1121/1.381843
- Herbst, J-P. (2019). Distortion and Rock Guitar Harmony: The Influence of Distortion Level and Structural Complexity on Acoustic Features and Perceived Pleasantness of Guitar Chords. *Music Perception*, 36(4), 335-352. <https://doi.org/10.1525/mp.2019.36.4.335>
- Herbst, J-P. (2018). Heaviness and the Electric Guitar: Considering the Interaction between Distortion and Harmonic Structures. *Metal Music Studies*, 4(1), 95-113. https://doi.org/10.1386/mms.4.1.95_1
-
- Inui, M., Takemoto, K., & Hamasaki, T. (2018). Harmonics and Intermodulation Distortion Analysis of the Even-Order Nonlinearity Controlled Effector Pedal. *Audio*

Engineering Society Convention 144, Retrieved from <http://www.aes.org/e-lib/browse.cfm?elib=19500>

- Izhaki, R. (2007). *Mixing Audio: Concepts, Practices and Tools*. Waltham, MA: Focal Press.
- Juchniewicz, J., & Silverman, M. J. (2013). The influences of progression type and distortion on the perception of terminal power chords. *Psychology of Music*, 41(1), 119-130. doi: 10.1177/0305735611422506
- Kendall, G. S., Haworth, C., & Cádiz, R. F. (2014). Sound Synthesis with Auditory Distortion Products. *Computer Music Journal*, 38(4), 5-23. doi: 10.1162/COMJ_a_00265
- Kim, C. D., Jee, S. K., & Masato, A. (1995). Measurement of the distortion level of a loudspeaker using an adaptive filter algorithm in a reverberation and noisy environment. *The Journal of the Acoustical Society of America*, 97(5), 3253. doi: 10.1121/1.413032
- Lazzarini, V., & Timoney, J. (2013). Synthesis of Resonance by Nonlinear Distortion Methods. *Computer Music Journal*, 37(1), 35-43. doi: 10.1162/COMJ_a_00160
- Marui, A., & Martens, W. L. (2005). Timbre of Nonlinear Distortion Effects: Perceptual Attributes Beyond Sharpness. *Conference on Interdisciplinary Musicology*, Montreal (Quebec), 5, Retrieved from https://www.researchgate.net/publication/243463251_Timbre_of_nonlinear_distortion_effects_Perceptual_attributes_beyond_sharpness
- Moore, A. (2017). *An Investigation into Non-Linear Sonic Signatures with a Focus on Dynamic Range Compression and the 1176 Fet Compressor (PhD Thesis)*. Retrieved from <http://eprints.hud.ac.uk/id/eprint/34118/>
- Murthy, A. A., Rao, N., Beemaiah, Y. R., Shandilya, S. D., & Siddegowda, R. B. (2014). Design and Construction of Arduino-Hacked Variable Gating Distortion Pedal. *IEEE*, 2(1), 1409-1417. doi: 10.1109/ACCESS.2014.2374195
- Mynett, M. (2017). *Metal Music Manual: Producing, Engineering, Mixing, and Mastering Contemporary Heavy Music*. New York ; London: Routledge.
- Novak, A., Simon, L., & Lotton, P. (2010). Analysis, Synthesis, and Classification of Nonlinear Systems Using Synchronized Swept-Sine Method for Audio Effects. *EURASIP Journal on Advances in Signal Processing*, , doi: 10.1155/2010/793816
- Olofsson, A., & Hansen, M. (2006). Objectively measured and subjectively perceived distortion in nonlinear systems. *The Journal of the Acoustical Society of America*, 120(6), 3757-3769. doi: 10.1121/1.2372591
- Pakarinen, J., & Yeh, D. T. (2009). A Review of Digital Techniques for Modeling Vacuum-Tube Guitar Amplifiers. *Computer Music Journal*, 33(2), 85-100. doi: 10.1162/comj.2009.33.2.85
- Pakarinen, J., Välimäki, V., Fontana, F., & Abel, J. S. (2011). Recent Advances in Real-Time Musical Effects, Synthesis, and Virtual Analog Models. *EURASIP Journal on Advances in Signal Processing*, 2011(1), doi: 10.1155/2011/940784

- Partti, H. (2014). Cosmopolitan musicianship under construction: Digital musicians illuminating emerging values in music education. *International Journal of Music Education*, 32(1), 3-18. doi: 10.1177/0255761411433727
- Pedersen, T. H. (2008). *The Semantic Space of Sounds* (2nd ed.). IL: DELTA.
- Persson, E. (2019). Evaluating tools and techniques for web scraping [Masters dissertation, KTH Royal Institute of Technology- School of Electrical Engineering and Computer Science]. <http://www.diva-portal.org/smash/get/diva2:1415998/FULLTEXT01.pdf>
- Phillips, N., & Hardy, C. (2002). *Discourse Analysis*. SAGE University . /10.4135/9781412983921
- Ronan, M., Ward, N., & Sazdov, R. (2015). An investigation into the sound quality lexicon of analogue compressors using category analysis. *Audio Engineering Society Convention 138*, Warsaw, Retrieved from https://www.researchgate.net/publication/303678004_An_investigation_into_the_sound_quality_lexicon_of_analogue_compressors_using_category_analysis
- SCAA. (1995). *The Coffee Taster's Flavor Wheel* [Infographic]. <https://sca.coffee/research/coffee-tasters-flavor-wheel>
- Schneiderman, M., & Sarisky, M. (2009). Analysis of a Modified Boss DS-1 Distortion Pedal. *Audio Engineering Society Convention 126*, Retrieved from <http://www.aes.org/e-lib/browse.cfm?elib=14987>
- Sunnerberg, T. D. (2019). *Analog Musical Distortion Circuits for Electric Guitars (Masters)*. Retrieved from <https://scholarworks.rit.edu/cgi/viewcontent.cgi?article=11227&context=theses>
- Toop, D. (2007). Rush Pep Box. *Leonardo Music Journal*, 17, 21-23. doi: 10.1162/lmj.2007.17.21
- Tsumoto, K., Marui, A., & Kamekawa, T. (2017). The Effect of Harmonic Overtones in Relation to 'Sharpness' for Perception of Brightness of Distorted Guitar Timbre. *Proceedings of Meetings on Acoustics*, 29(1), doi: 10.1121/2.0000380
- Williams, D. (2014). Tracking timbral changes in metal productions from 1990 to 2013. *Metal Music Studies*, 1(1), 39-68. doi: 10.1386/mms.1.1.39_1
- WoodyTone. (2010). *Early Maiden: Murray, Smith Gear*. Retrieved from <https://www.woodytone.com/2010/12/14/early-maiden-murray-smith-gear/>
- Yeh, D. T., Abel, J. S., Vladimirescu, A., & Smith, J. O. (2008). Numerical Methods for Simulation of Guitar Distortion Circuits. *Computer Music Journal*, 32(2), 23-42. doi: 10.1162/comj.2008.32.2.23
- Zargoski-Thomas, S., Isakoff, K., Lacasse, S., & Stevance, S. (2019). *The Art of Record Production: Creative Practice in the Studio* (2nd ed.). UK: Routledge.