

STEALING GUITAR EFFECTS

J. Max Morris

University of Rochester
jmorr32@ur.rochester.edu

Alex Kim

University of Rochester
akim65@ur.rochester.edu

ABSTRACT

Guitarists often look to other guitarists for new tones and sounds to use. We propose a method to identify various guitar effects and replicate them via inputting a recording with effects and applying that to a dry recording to produce that effect without knowing the parameters. This is done using features extracted from a signal with effect with non-negative matrix factorization to compare to a known dataset.

1. INTRODUCTION

Guitarists often hope to replicate other guitarist’s tones. Whether this is through using the same pedals, same amps, same guitar, or same plugins, these are usually quite expensive and do not necessarily include what settings they use. We propose a method to identify and reverse engineer guitar effects from a single source to apply to other guitar recordings. This includes time-based effects as well as timbre-based effects.

There have been several different methods that attempt to do this individually, like dynamic convolution using impulse responses [1], neural networks [2], or virtual analog representations [3], but this paper proposes various unique and novel methods to accomplish this task. These are all based on the features retrieved with non-negative matrix factorization (NMF) [4]. This paper proposes ways to identify and replicate delay, tremolo, and reverb effects given a wet signal with one effect applied. We also reviewed methods for timbre-based effects for further work.

2. BACKGROUND

NMF is a method used to decompose a V matrix into two matrices, W and H [4]. In terms of audio, after applying a short time Fourier transform (STFT), we can extract various characteristics of the sound. The H matrix for instance is the activation matrix which shows note onsets. The W matrix has the spectral characteristics of said activations. To retrieve these W and H matrices, they can either be constantly changed or saved. In this case, for time-based effects, we save the H matrix of an initial training set. This set was of various clean (dry) guitar recordings. Useful information can also be extracted with the NMF of signals with effects on them (wet) for identification. These were recorded with direct input using a Focusrite Scarlet audio interface. For time-based effects, this was based on taking the Fast Fourier Transform (FFT) of the activation matrix for both identification and in some case, finding the properties of the effects we hope to retrieve. We also hoped to accomplish a method to retrieve a distortion and/or overall timbre from a different recording and apply it to a dry signal. The program asks for a wet signal that one would like to identify, a clean signal to apply it to, datasets for dry and

wet effects. There are choices for pretrained data, or new data if one would like to use their own.

3. DATASET

The dataset we used consisted of two main parts: the baseline dataset, and the effects dataset. The effects dataset consists of guitar audio with various effects applied, in various amounts including, reverb, delay, tremolo, and some with no effects, all of which were monophonic music. This dataset was used for training the identification weights as well as for initial testing. For any training for the NMF, we used 100 iterations and r value of 25.

On the other side, the baseline dataset can be broken down into two components, the frequency and time focused audio. For the frequency focused audio, we used recordings of a clean guitar playing every possible note on the guitar. This consisted of 49 notes ranging from an E2 to an E6, corresponding to a 24-fret guitar’s range in standard tuning. This audio was run through an NMF algorithm and used to create a Frequency Dictionary matrix (W) which represents the frequency content of a typical clean guitar. To aid with accuracy, each column of the initial W matrix consists of 0’s below the 49 fundamental frequencies in a guitar’s range, which can be seen in Figure 1.

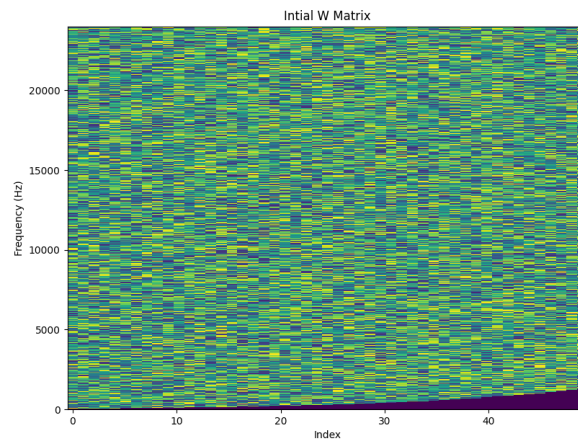
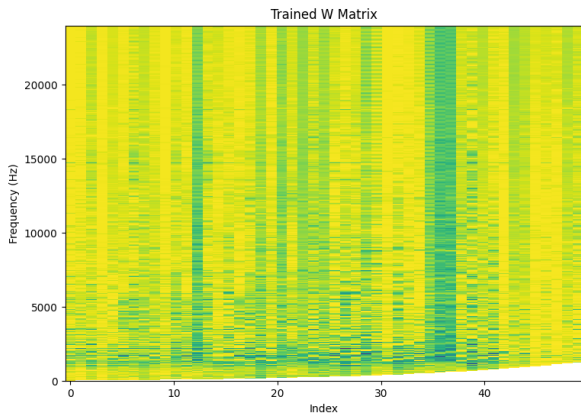



Figure 1. Initial W matrix shaped to follow pre-existing knowledge.

The W matrix output is then fed back into the next NMF algorithm, in the hopes of fine tuning the W matrix with each audio file. The resulting W matrix can be seen in figure 2.

The time focused audio consisted of recordings of a guitarist playing a song on the guitar. This audio was also run through an NMF algorithm to create an Activation Dictionary Matrix (H) which could be used to compare to temporal based effects. An FFT is then applied to the combination of all the rows, providing a baseline curve that can be used as a comparison for identification. An example of this can be seen in Figure 3.



89 **Figure 2.** Trained Baseline W matrix.



91 **Figure 3.** FFT of Baseline H matrices.

93 4. METHOD

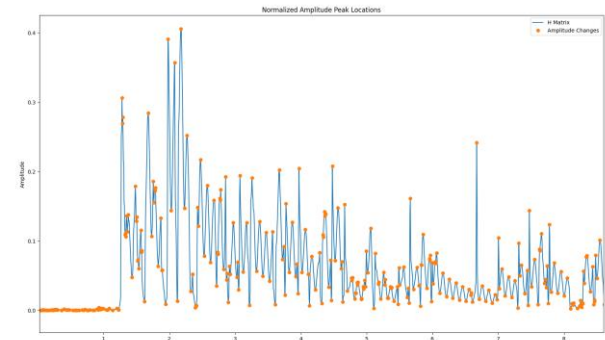
94 4.1 Time Based Effects

95 The three time-based effects we looked at were tremolo,
 96 delay, and reverb. For each of these effects, the first step
 97 of each algorithm is to run the input signal through an
 98 NMF algorithm to extract the activation matrix (H) with
 99 size $[r,n]$. In some cases, we combined all the rows of the
 100 H matrix to produce a master H matrix, with size $[n]$, for
 101 the audio. This ends up approximating the amplitude per
 102 frame for the audio. Either the amplitude per frame or the
 103 full H matrix is useful when attempting to derive its char-
 104 acteristics for replication.

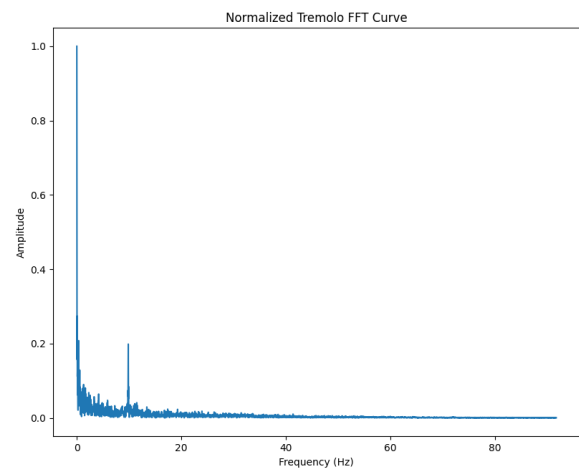
105 4.1.1 Tremolo

106 The two controls that the tremolo algorithm aims to find
 107 are the frequency and depth of the amplitude modulation.
 108 To find the depth, we first take the time derivative of the
 109 master H matrix and then find all the zero crossing indexes.
 110 This provided us with a map to show when the amplitude
 111 changed directions. An example of this can be seen in Fig-
 112 ure 4.

113 Using this we then looked to find the average percent
 114 change in amplitude between two consecutive peaks. On
 115 top of this, we only considered peaks where the following
 116 peak had a lower amplitude. The ratios between all these
 117 peaks are then averaged to find the predicted depth value.

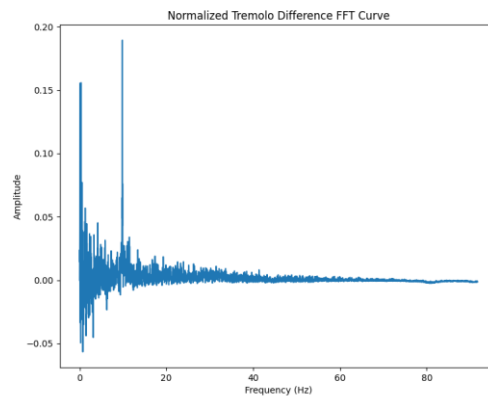


118 **Figure 4.** Zoomed in graph of peak amplitudes on top of
 119 the Master H matrix for a tremolo signal.
 120

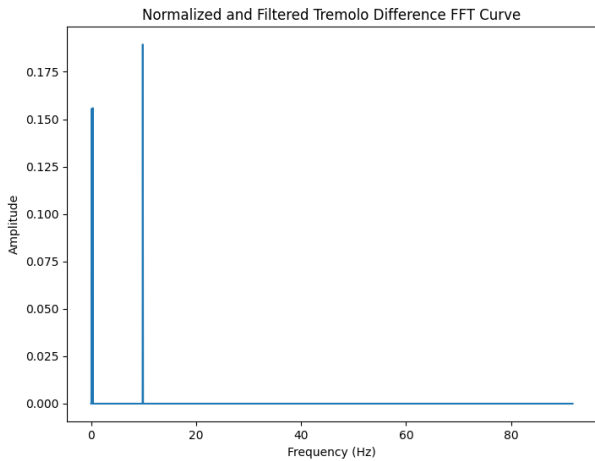


121 **Figure 5.** FFT of the master H matrix for a signal with a
 122 10Hz tremolo effect applied.
 123

124 For the frequency control, we applied an FFT to the
 125 master H matrix, which can be seen in Figure 5. For com-
 126 parison, we also used our baseline dataset to find an ex-
 127 pected FFT curve. We then subtract the baseline FFT from
 128 the tremolo FFT. This results in a frequency response with
 129 a strong peak at the frequency of the tremolo effect. An
 130 example of this can be seen in Figure 6. We then filter out
 131 amplitudes less than or equal to 0.1 to isolate this peak. An
 132 example of this can be seen in Figure 7.



133 **Figure 6.** Difference of FFTs of the master H matrix for a
 134 10Hz tremolo signal and baseline signal.
 135



136

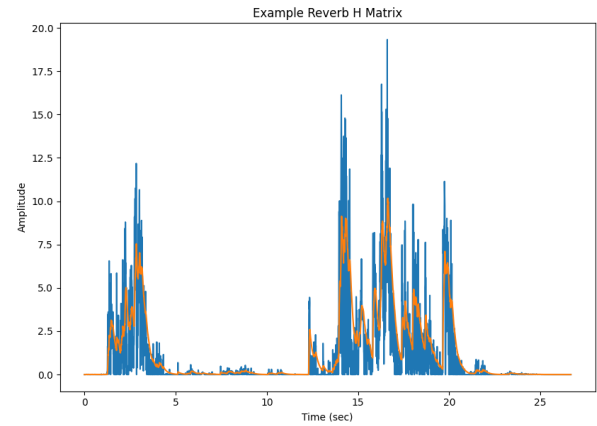
137 **Figure 7.** Filtered difference of FFTs of the master H ma-
 138 trix for a 10Hz tremolo signal and baseline signal.

139 We identify this peak and then use parabolic interpola-
 140 tion to determine the exact frequency identified. Finally,
 141 using this information we can generate a low-frequency
 142 oscillator (LFO) to control the amplitude of the provided
 143 clean audio. Once this is applied, the modified signal is
 144 then output to a rendered file.

145 4.1.2 Reverb

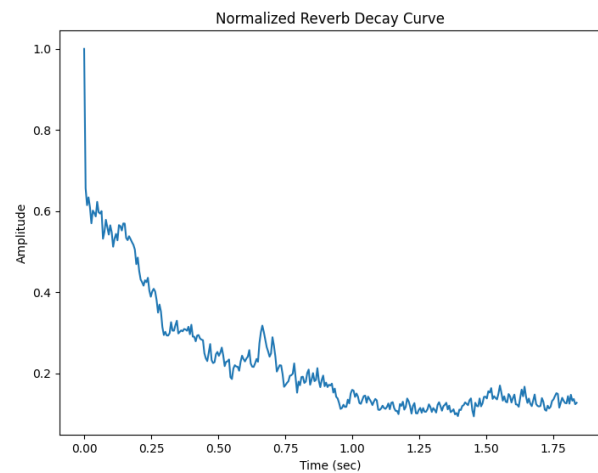
146 For this effect, the algorithm takes a slightly different ap-
 147 proach compared to the previous one. This algorithm first
 148 finds the decay time, and then uses it to create an impulse
 149 response (IR) that can be convolved with the dry signal.
 150 For the decay time, it looks to find the average number of
 151 frames it takes the envelope of each row in the H matrix to
 152 go from a peak amplitude to below a specified threshold.
 153 For our tests, we used a threshold of 0.05. An example of
 154 one instance of an H matrix can be seen in Figure 8. Once
 155 this is known, it looks to find the average amplitude curve
 156 over that number of frames. Again, this is done for every
 157 row of the H matrix, and the results are averaged. An ex-
 158 ample of this can be seen in Figure 9.

159 Once the decay time and decay curve are found, we then
 160 look at the STFT of the wet audio. It then finds the decay
 161 curve, using the same method mentioned for the amplitude
 162 curve, for each frequency bin. Once each frequency bin's
 163 decay is found, it is multiplied by the amplitude curve, to
 164 ensure the IR follows the amplitude decay. An example
 165 STFT of a calculated IR can be seen in Figure 10. An in-
 166 verse STFT function is then applied to the calculated IR to
 167 convert it into the time domain. Once this is done, it is con-
 168 volved with the dry audio, and the algorithm outputs the
 169 modified audio and impulse response.



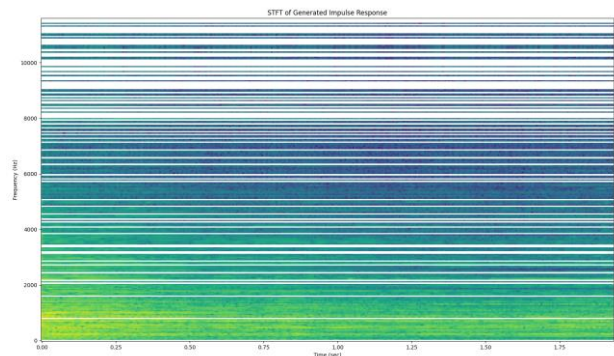
170

171 **Figure 8.** Example of H matrix row, with envelope and
 172 peak identification.



173

174 **Figure 9.** Example of calculated decay curve.




175

176 **Figure 10.** Example of calculated impulse response.

177 4.1.3 Delay

178 For a delay effect, there are two controls that our proposed
 179 algorithm aims to find, the delay time and the echo ampli-
 180 tude curve. The first control that the algorithm looks for is
 181 the delay time. To find this we used the same method used
 182 to find the frequency for the tremolo effect. Unlike the FFT
 183 results gotten from the tremolo effect, this FFT results in a
 184 frequency spectrum with small evenly spaced spikes

185 which correlate to the delay “frequency.” An example of
 186 this can be seen in Figure 11. To determine this “fre-
 187 quency” we first subtract the baseline FFT from the delay
 188 FFT, giving us the graph in Figure 12. We then use peak
 189 detection to find the location of each onset.




190
 191 **Figure 11.** FFT of the H matrix of a signal with a 500ms
 192 delay applied.

193 We then use an algorithm to find the most common
 194 spacing between peaks that is also equal to a peak value.
 195 This corresponds to the “fundamental frequency.” Invert-
 196 ing this frequency yields the delay time, which we can eas-
 197 ily convert from frames to seconds to get a usable value.

198 The next control parameter the algorithm looks for is the
 199 echo amplitude curve. This curve represents how loud each
 200 delayed signal is compared to the clean signal. To find this
 201 value, we use a similar method to find the decay curve for
 202 the reverb effect. However, in this case we use the already
 203 identified peaks and compare those amplitude peaks to the
 204 amplitude peaks that are integer multiples of the delay time
 205 away.

206 Essentially, we are using the fact that we know when to
 207 expect another peak, based on the delay time, to help de-
 208 termine the amplitude change. This algorithm is run on
 209 each row in the activation matrix to get an average curve
 210 for each row. When we combine these curves, we also ig-
 211 nore any large rising amplitudes (≥ 0.2) as we assume
 212 that is a new note, rather than an echo. An example of this
 213 curve can be seen in Figure 16.

214 Once this is performed the results are averaged to find
 215 the predicted delay mix. Using the delay time and echo
 216 curve, we add the clean audio to itself, delayed by the
 217 amount we found, and multiplied by the averaged curve of
 218 the amplitude decays to output the signal. We then output
 219 the clean file with delays.

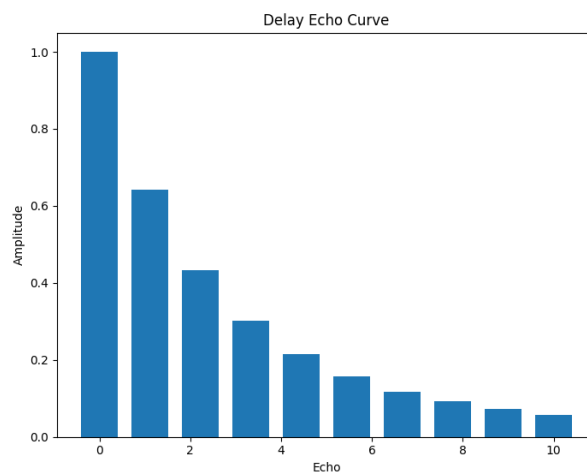


220
 221 **Figure 12.** Difference of FFTs of the H matrix for a de-
 222 layed signal and baseline signal.


223 4.2 Timbral Based Effects

224 The goal for the timbral based effects is to be able to trans-
 225 fer the timbral information from a wet signal to a dry sig-
 226 nal. The focus was on distortion effects, but the algorithm
 227 could be expanded to any timbral based effect, like equal-
 228 ization and amplifier/cabinet modeling.

229 Three main methods were tested for these effects. These
 230 were all generally based on the principal of cross synthesis
 231 [5]. The first was using the average envelope shape to ad-
 232 just the frequency information of a clean signal to match
 233 the envelope of the wet signal. For this, both audio files
 234 were run through an NMF algorithm to extract the W ma-
 235 trix. From this, the envelope of each instance (column in
 236 the W matrix) is calculated before they are averaged to-
 237 gether. For the clean signal, the envelope is also averaged
 238 with the baseline dataset’s averaged envelope. The wet and
 239 dry averaged envelopes are then compared to find a trans-
 240 form function that is then applied to each instance in the
 241 clean signal’s W matrix. From this, the modified W matrix
 242 and clean signal’s H matrix are recombined with the orig-
 243 inal phase, and an inverse-STFT is performed to generate
 244 the modified audio.




245
 246 **Figure 16.** Example of echo amplitude curve.

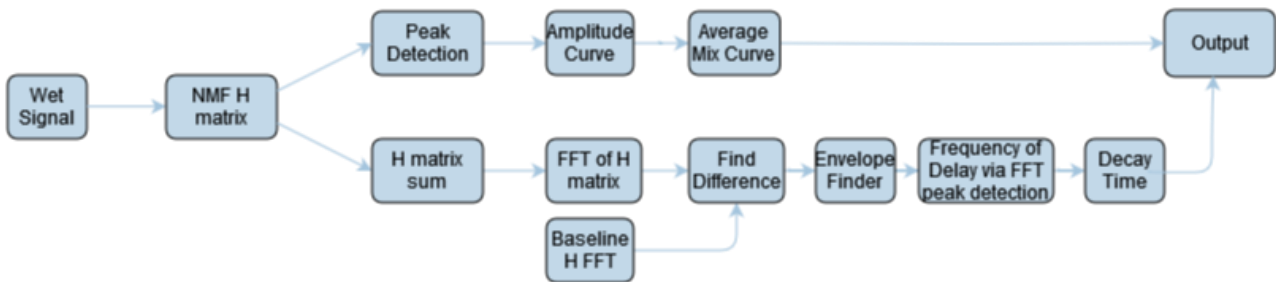


247 **Figure 13.** Reverb algorithm block diagram.

281
282



248 **Figure 14.** Tremolo algorithm block diagram.




249 **Figure 15.** Delay algorithm block diagram.

250 The second method improves upon the first by looking
251 at the average harmonic series amplitudes, rather than the
252 average envelope. Similarly, to the first method, the clean
253 signal's W matrix is joined with the baseline W matrix.
254 From this, the peaks in each instance of both the wet and
255 clean W matrices are identified. Then using the same algo-
256 rithm used for the delay effect, find the most common
257 spacing between all the peaks. This allows us to identify
258 the fundamental frequency and the peaks that correspond
259 to the harmonic series.

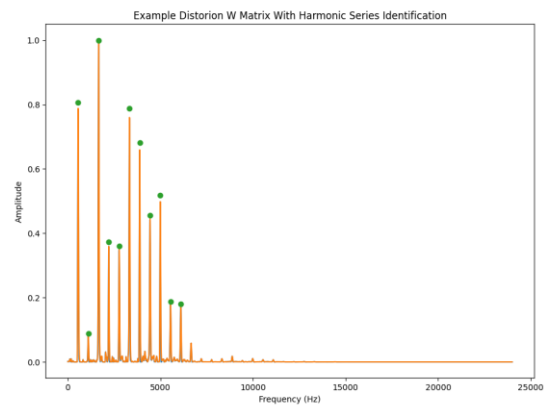
260 Examples of this can be seen in Figure 17 and Figure
261 18. We can average the amplitude of the first 16 overtones,
262 which can be seen in Figure 19. We then calculate a trans-
263 fer function based on the ratio between the wet and clean
264 signal's harmonic series amplitudes. This transform func-
265 tion is then applied to the provided clean signal, which
266 goes through the same algorithm to identify the harmonic
267 series, before the transform function is applied.

268 The final method we looked at used the tanh function to
269 approximate distortion specifically [6]. As a result, we
270 tried to implement a machine learning based method that
271 applies a tanh function to the inputted clean signal, and
272 then compares it to the provided wet signal, to update the
273 approximation. The planned method to find the difference
274 between the two signals was to compare Mel Frequency
275 Cepstrum Coefficients (MFCC) of both signals. However,
276 we were unable to fund an effective equation to handle
277 this.

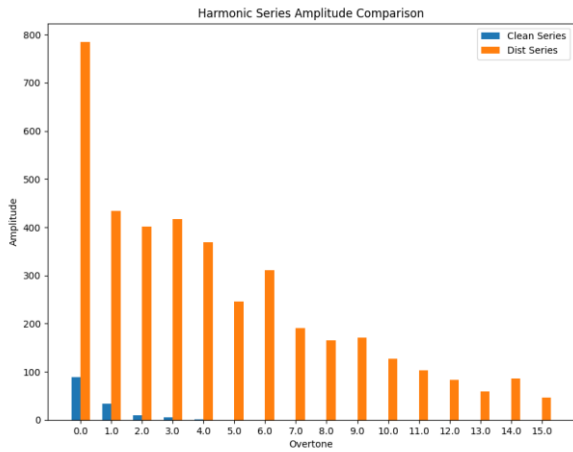
278
279
280



283 **Figure 17.** Predicted overtone series of one column of W
284 matrix of clean signal.
285



286 **Figure 18.** Predicted overtone series of one column of W
287 matrix of distorted signal.
288

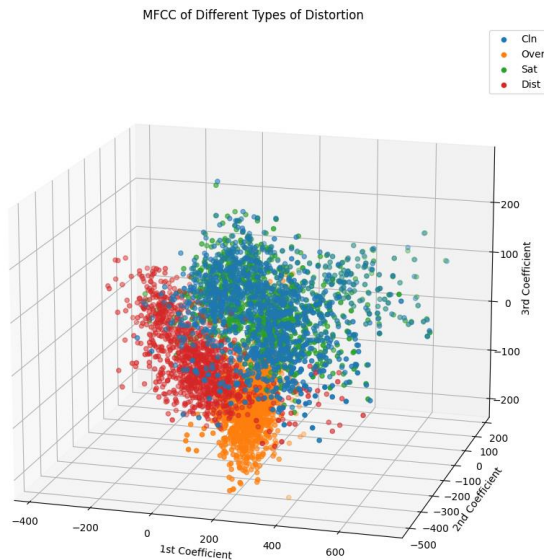


289 **Figure 19.** Comparison of overtone series for clean and
 290 distorted signal
 291

292 4.3 Identification

293 The final aspect of this algorithm to consider is its ability
 294 to identify what effect is applied to the provided signal. For
 295 timbral based effects, this method has not been imple-
 296 mented. As a result, this will be a discussion of a proposed
 297 method. The proposed algorithm will calculate the MFCC
 298 of the provided audio. An example of these can be seen in
 299 Figure 20.

300 It will then compare this to a pretrained library of
 301 MFCCs. This pretrained library will be used to find some
 302 equation that is able to represent where most values of the
 303 MFCCs are expected to be. Then these equations will be
 304 applied to the MFCC of the provided audio, and the per-
 305 centage of values that fall within the ranges of this equa-
 306 tion will be recorded. The equation that produces the best
 307 match will then be selected as the predicted effect.



308 **Figure 20.** MFCC Coefficients 1, 2, and 3 for different
 309 types of distortion
 310

311 For time-based effects, it will use the FFT of the H matrices for comparison. From our testing, we have identified
 312 that the FFTs of the H matrices vary significantly enough
 313 between a tremolo, delay, and reverb effect to allow for

316 identification. To differentiate between them, we look at
 317 the number of peaks (Peak) above 0.01 and the number of
 318 large peaks (LPeak), which are the peaks above 0.1. Since
 319 the FFT of a tremolo effect results in one large spike and
 320 significant low amplitude noise, it has many total peaks
 321 and the largest peaks. The many small amplitude peaks in
 322 the FFT of a delay effect, results in a moderate number of
 323 total peaks, and typically has no large peaks. The FFT of a
 324 reverb effect is typically very similar to the clean baseline,
 325 and as a result it has the least number of total peaks, and
 326 few large peaks. To help with identification e also looked
 327 at the difference between a clean guitar FFT curve and the
 328 baseline. We found that the FFT of a clean signal is very
 329 similar to the baseline, but has a lot of low amplitude noise,
 330 like the tremolo effect. As a result, it has a similar number
 331 of total peaks as the tremolo effect, but with less large
 332 peaks, allowing the two to be differentiated. To determine
 333 these values, we trained the algorithm on our effects da-
 334 taset, and the resulting values can be seen in Table 1.

	Clean	Tremolo	Delay	Reverb
Peak	82.77	86.33	29.75	13.5
LPeak	1.69	3.166	0	1.5

336 **Table 1.** Trained identification weights

337 To calculate the match percentage, we use the following
 338 equations. In all cases, if the value calculated for A from
 339 Eqn (1) is less than 0, we use 0 in place of a negative value.
 340 If the expected LPeak (Exp LPeak) is above zero, we use
 341 Eqn (2), and in this case, if B is below 0, we use 0 in place
 342 of a negative value. If B is 0 and Exp LPeak is 0, then we
 343 assume B/(Exp LPeak) is 1. In all other cases we assume
 344 the second term is 1/B, in this case we take the absolute
 345 value of B.

$$346 \quad A = \text{Exp Peak} - |\text{Exp Peak} - \text{Peak}|$$

$$347 \quad B = \text{Exp LPeak} - |\text{Exp LPeak} - \text{LPeak}|$$

348 **Equation 1.** Expected minus calculated distance.

$$349 \quad \left(\frac{A}{\text{Exp Peak}} + \frac{B}{\text{Exp LPeak}} \right) * 0.5 = \text{Closeness}$$

350 **Equation 2.** Closeness percentage for when Exp LPeak
 351 does not equal 0.

352 5. RESULTS

353 5.1 Time Based Effects Results

354 Our results for replicating the delay are as follows. For
 355 identifying the delay time, the algorithm can correctly
 356 identify the delay time within 14% error. However, due to
 357 the nature of using NMF algorithms, which are random,
 358 the algorithm can occasionally result in guesses within
 359 22% error. This is likely due to some of the peaks in the
 360 FFT being less prominent, making it harder for the algo-
 361 rithm to separate the peaks from noise. For the mix percent
 362 value, the algorithm can correctly identify the value within
 363 44% error. This is most likely due to the calculated echo
 364 curve, still being susceptible to other note onsets adjusting
 365 the amplitude values. The specific results for the delay ef-
 366 fect can be seen in Table 2.

	Run 1	Run 2	Run 3	Expected	Lowest Error	Avg Error
Delay Time	585ms	519ms	611ms	500ms	4%	14%
Mix	71%	73%	72%	50%	42%	44%

369 **Table 2.** Delay effect results for delayed signal with
370 500ms delay time.

371

372 Next, our results for replicating the tremolo are as fol-
373 lows. For the frequency, the algorithm can correctly iden-
374 tify the value within 2% error. On the other hand, for the
375 depth, the algorithm can correctly identify the value within
376 31% error. This is likely due to noise in the activation ma-
377 trix, and the peak identification, that causes inaccuracies in
378 the amplitudes. The specific results for the tremolo effect
379 can be seen in Table 3.

380

	Run 1	Run 2	Run 3	Expected	Lowest Error	Avg Error
Freq	9.84Hz	9.84Hz	9.84Hz	10.055Hz	2%	2%
Depth	41%	42%	42%	60%	30%	31%

381 **Table 3.** Tremolo effect results for tremolo signal with
382 10Hz modulation frequency.

383

384 Finally, our results for replicating the reverb are as fol-
385 lows. For the decay time, the algorithm can correctly iden-
386 tify the value within 8% error. The specific results for the
387 reverb effect can be seen in Table 4.

388

	Run 1	Run 2	Run 3	Expected	Lowest Error	Avg Error
Decay Time	1.838 sec	1.835 sec	1.932 sec	2 sec	3%	7%

389 **Table 4.** Reverb Effect results for reverb signal with 2 sec-
390 ond decay time.

391

392 Listening to the outputted audio, the algorithm can ap-
393 ply a similar sounding reverb to the audio, however, there
394 is still some artifacts and lack of clarity.

395 For the time-based identification, we found that our al-
396 gorithm was able to consistent correctly identify the effect.
397 A confusion matrix showing the percent match the algo-
398 rithm found for each inputted signal can be seen in Table
399 5. For clarity, the highest match has been highlighted in
400 green, and any other match percents 50% or above have
401 been highlighted in orange.

402

Guess	Input Signal			
	Clean	Tremolo	Delay	Reverb
Clean	69.83%	25.58%	50%	37.8%
Tremolo	65.82%	65.6%	35.37%	23.66%
Delay	25%	16.67%	92.86%	72.97%
Reverb	33.33%	0%	33.33%	81.48%

403 **Table 5.** Confusion matrix between inputted signal and the
404 algorithm’s identification guess.

405

406 Looking at this table, while the algorithm correctly iden-
407 tified each effect, some had higher percentages with other
408 effects. This is most likely due to similarities with the iden-
409 tification weights. For example, the clean and tremolo
410 weights are almost identical, and the inputted clean signal
411 had a large match percent for being a tremolo. As a result,
412 it would be expected that these two have similar match per-
413 cents, as the difference between the number of large peaks
414 is the only value that is useful for differentiating the two
415 effects. For the delay/clean and reverb/delay confusion, it
416 is likely due to the closeness of the number of large peaks.

417 Since they are so close, it is very likely that the match per-
418 cent is inflated as a result.

419 5.2 Timbre Based Effects Results

420 Our results for replicating timbre-based effects are as fol-
421 lows. All three methods discussed are unable to accurately
422 recreate the provided effect. The first two methods, involv-
423 ing the envelope and harmonic series amplitudes, result in
424 audio which consists of the clean audio with significant ar-
425 tifacts and unintentional distortion. On the other hand, the
426 third method, tanh approximation, can apply a more “typ-
427 ical” distortion sound but is unable to match the exact tone
428 from the provided wet signal. There is also the issue of re-
429 turning the original phase, which there are some solutions,
430 like the Griffin-Lim algorithm. This is important for when
431 we return the inverse-STFT to reproduce the sound with
432 minimal artifacts.

433

433 6. DISCUSSION

434 6.1 Summary

435 This work can apply multiple effects to a provided clean
436 guitar audio file, given a guitar recording with a single ef-
437 fect. This was tested using monophonic guitar recordings.
438 The baseline dataset did include polyphonic guitar record-
439 ings (strummed with chords), but the effects data set did
440 not include any. We found success in retrieving tremolo,
441 delay, and reverb effects with usable results, while timbre
442 was not useable. The timbral results we did obtain were
443 interesting as a unique effect, but not the replication we
444 intended.

445 6.2 Future Work

446 The large scope of this project requires various techniques
447 to process each individual effect. While this may be more
448 intuitive to figure out and process than a neural network, it
449 requires a lot more study into methods per effect. We hope
450 that improvements on timbre replication as well as a better
451 approximation of delays can be done.

452 While these techniques are not exclusive to guitar, we
453 focused on typical effects that guitarists use. We have not
454 attempted to see how our program reacts with non-guitar
455 effects, which may produce interesting results. We also
456 have only tested with supplying our program with input
457 signals of singular notes. Our methods should be robust to
458 chords for time domain effects. We have not tested iden-
459 tifying signals that have more than one effect. This project
460 set out to obtain guitar effects from a supplied recording.
461 With some training already applied, we can take a single
462 effect and apply it to a signal we want. This proves that at
463 a minimum, this is a viable concept, but there are various
464 aspects to improve. This include, on a broader range of
465 topics, various effects and being robust to multiple effects.
466 This would mean adding more common effects as well as
467 being able to identify them. On a smaller scale, we believe
468 that each method we use in this: tremolo, reverb, and delay,
469 could be more robust to minute changes, as well as more
470 consistent in identification and replication. We also
471 acknowledge the small dataset we used, which were self-

472 supplied recordings. Another consideration, given more
473 time, would be to use a much larger dataset to develop a
474 more accurate reference and see if and or how the results
475 change.

476

7. REFERENCES

- 477 [1] M. Kemp, "Analysis and Simulation of Non-Linear
478 Audio Processes using Finite Impulse Responses De-
479 rived at Multiple Impulse Amplitudes," *AES 106th*
480 *Convention*, Munich, Germany, 1999,
- 481 [2] J. Engel, L. Hantrakul, C. Gu, A. Roberts, "DDSP:
482 Differentiable Digital Signal Processing," *Interna-*
483 *tional Conference on Learning Representations*,
484 Online 2020.
- 485 [3] D. T. Yeh and J. S. Abel, Automated Physical Mod-
486 eling of Nonlinear Audio Circuits for Real-Time Au-
487 dio Effect – Part 1: Theoretical Development, "*IEEE*
488 *Transactions on Speech and Audio Procession*," vol.
489 18, no. 3, March 2010.
- 490 [4] D. D. Lee and H. S. Seung, "Learning the parts of
491 objects by non-negative matrix factorization," *Na-*
492 *ture Journal*, vol. 401, pp 788-791, October 1999.
- 493 [5] G. Roma, O. Green, P. Tremblay, "Audio Morphing
494 Using Matrix Decomposition and Optimal
495 Transport," in *International Conference on Digital*
496 *Audio Effects*, Online, 2020.
- 497 [6] J. T. Colonel, M. Comunita, J. Reiss, "Reverse Engi-
498 neering Memoryless Distortion Effects with Differ-
499 entiable Waveshapers," *AES 153rd Convention*, New
500 York, NY, 2020.