

Cover Song Synthesis by Analogy

Christopher J. Tralie
Duke University Department of Mathematics
ctralie@alumni.princeton.edu

Abstract

In this work, we pose and address the following “cover song analogies” problem: given a song A by artist 1 and a cover song A' of this song by artist 2, and given a different song B by artist 1, synthesize a song B' which is a cover of B in the style of artist 2. Normally, such a polyphonic style transfer problem would be quite challenging, but we show how the cover songs example constrains the problem, making it easier to solve. First, we extract the longest common beat-synchronous subsequence between A and A' , and we time stretch the corresponding beat intervals in A' so that they align with A . We then derive a version of joint 2D convolutional NMF, which we apply to the constant-Q spectrograms of the synchronized segments to learn a translation dictionary of sound templates from A to A' . Finally, we apply the learned templates as filters to the song B , and we mash up the translated filtered components into the synthesized song B' using audio mosaicing. We showcase our algorithm on several examples, including a synthesized cover version of Michael Jackson’s “Bad” by Alien Ant Farm, learned from the latter’s “Smooth Criminal” cover.

1 Introduction




Figure 1: A demonstration of the “shape analogies” technique in [22]. In the language of our work, the statue head with a neutral expression (A') is a “cover” of the low resolution mesh A with a neutral expression, and this is used to synthesize the surprised statue “cover” face B' from a low resolution surprised face mesh B .

The rock group Alien Ant Farm has a famous cover of Michael Jackson’s “Smooth Criminal” which is faithful to but stylistically unique from the original song. However, to our knowledge, they never released a cover of any other Michael Jackson songs. What if we instead wanted to know how they would have covered Michael Jackson’s “Bad”? That is, we seek a song which is identifiable as MJ’s “Bad,” but which also sounds as if it’s in Alien Ant Farm’s *style*, including timbral characteristics, relative tempo, and instrument types.

In general, multimedia style transfer is a challenging task in computer aided creativity applications. When an example of the stylistic transformation is available, as in the “Smooth Criminal” example above, this problem can be phrased in the language of analogies; given an object A and a differently stylized version of this object A' , and given an object B in the style of A , synthesize an object B' which has the properties of B but the style of A' . One of the earliest works using this vocabulary is the “image analogies” work [12], which showed it was possible to transfer both linear filters (e.g. blurring, embossing) and nonlinear filters

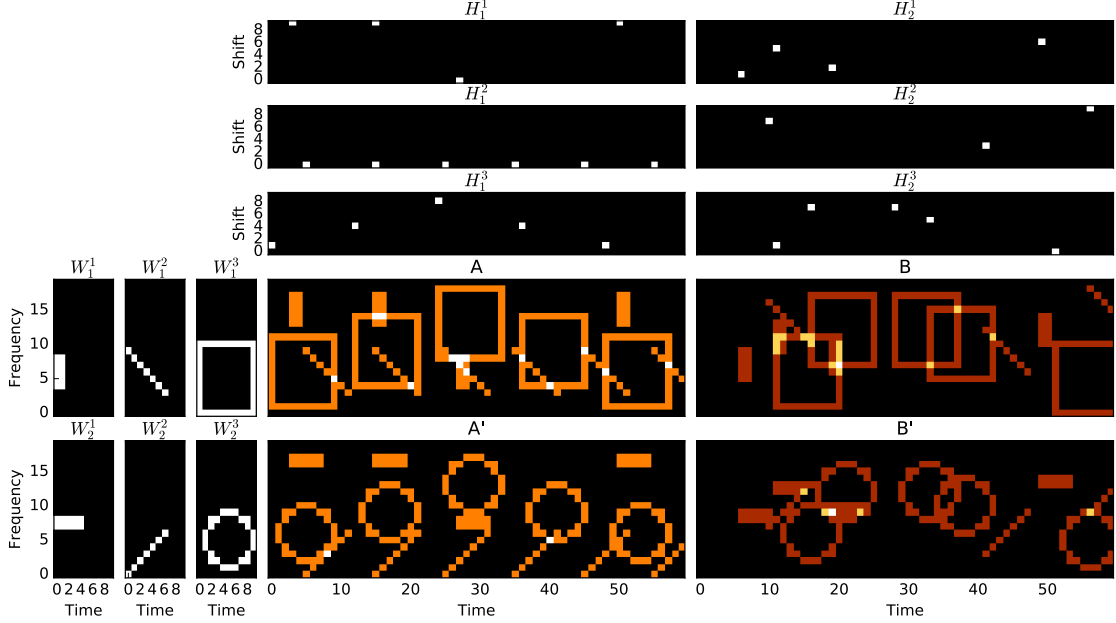


Figure 2: An ideal cartoon example of our joint 2DNMF and filtering process, where $M = 20$, $N = 60$, $T = 10$, $F = 10$, and $K = 3$. In this example, vertical time-frequency blocks are “covered” by horizontal blocks, diagonal lines with negative slopes are covered by diagonal lines with positive slopes, and squares are covered by circles. When presented with a new song B , our goal is to synthesize a song B' whose CQT is shown in the lower right green box.

(e.g. watercolors) in the same simple framework. More recent work with convolutional networks has shown even better results for images [10]. There has also been some work on “shape analogies” for 3D meshes [22], in which nonrigid deformations between triangle meshes A and A' are used to induce a corresponding deformation B' from an object B , which can be used for motion transfer (Figure 1).

In the audio realm, most style transfer works are based on mashing up sounds from examples in a target style using “audio mosaicing,” usually after manually specifying some desired path through a space of sound grains [16]. A more automated audio mosaicing technique, known as “audio analogies” [19], uses correspondences between a MIDI score and audio to drive concatenated synthesis, which leads to impressive results on monophonic audio, such as stylized synthesis of jazz recordings of a trumpet. More recently, this has evolved into the audio to musical audio setting with audio “musaicing,” in which the timbre of an audio source is transferred onto a target by means of a modified NMF algorithm [7], such as bees buzzing The Beatles’ “Let It Be.” A slightly closer step to the polyphonic (multi source) musical audio to musical audio case has been shown to work for drum audio cross-synthesis with the aid of a musical score [5], and some very recent initial work has extended this to the general musical audio to musical audio case [9] using 2D nonnegative matrix factorization, though this still remains open. Finally, there is some recent work on converting polyphonic audio of guitar songs to musical scores of varying difficulties so users can play their own covers [1].

In this work, we constrain the polyphonic musical audio to musical audio style transfer problem by using *cover song* pairs, or songs which are the same but in a different style, and which act as our ground truth A and A' examples in the analogies framework. Since small scale automatic cover song identification has matured in recent years [17, 18, 4, 23], we can accurately synchronize A and A' (Section 2.1), even if they are in very different styles. Once they are synchronized, the problem becomes more straightforward, as we can blindly factorize A and A' into different instruments which are in correspondence, turning the problem into a series of monophonic style transfer problems. To do this, we perform NMF2D factorization of A and A' (Section 2.2). We then filter B by the learned NMF templates and mash up audio grains to create B' , using the aforementioned “musaicing” techniques [7] (Section 2.3). We demonstrate our techniques on snippets of A , A' , and B which are about 20 seconds long, and we show qualitatively how the instruments of A' transfer

onto the music of B in the final result B' (Section 3).

2 Algorithm Details

In this section, we will describe the steps of our algorithm in more detail.¹

2.1 Cover Song Alignment And Synchronization




Figure 3: An example feature fused cross-similarity matrix D for the first 80 beats of Michael Jackson’s “Smooth Criminal,” compared to the cover version by Alien Ant Farm. We threshold the matrix and extract 20 seconds of the longest common subsequence, as measured by Smith Waterman. The alignment path is shown in red.

As with the original image analogies algorithm [12], we find it helpful if A and A' are in direct correspondence at the sample level. Since cover song pairs generally differ in tempo, we need to align them first. To accomplish this, we draw upon the state of the art “early fusion” cover song alignment technique presented by the authors of [23]. Briefly, we extract beat onsets for A and A' using either a simple dynamic programming beat tracker [8] or slower but more accurate RNN + Bayesian beat trackers [13], depending on the complexity of the audio. We then compute beat-synchronous sliding window HPCP and MFCC features, and we fuse them using similarity network fusion [25, 26]. The result is a $M \times N$ cross-similarity matrix D , where M is the number of beats in A and N is the number of beats in A' , and D_{ij} is directly proportional to the similarity between beat i of A and beat j in A' . Please refer to [23] for more details.

Once we have the matrix D , we can then extract an alignment between A and A' by performing Smith Waterman [20] on a binary thresholded version of D , as in [23]. We make one crucial modification, however. To allow for more permissive alignments with missing beats for identification purposes, the original cover songs algorithm creates a binary thresholded version of D using 10% mutual binary nearest neighbors. On the other hand, in this application, we seek shorter snippets from each song which are as well aligned as possible. Therefore, we create a stricter binary thresholded version B , where $B_{ij} = 1$ only if it is in the top $3\sqrt{MN}$ distances over all MN distances in D . This means that many rows of B_{ij} will be all zeros, but we will hone in on the best matching segments. Figure 3 shows such a thresholding of the cross-similarity matrix for two versions of the song “Smooth Criminal,” which is an example we will use throughout this section. Once B has been computed, we compute a X -length alignment path P by back-tracing through the Smith Waterman alignment matrix, as shown in Figure 3.

Let the beat onset times for A in the path P be t_1, t_2, \dots, t_X and the beat times for A' be s_1, s_2, \dots, s_X . We use the rubberband library [3] to time stretch A' beat by beat, so that interval $[s_i, s_i + 1]$ is stretched by

¹Note that all audio is mono and sampled at 22050hz.




Figure 4: Joint 2DNMF on the magnitude CQTs of “Smooth Criminal” by Michael Jackson and Alien Ant Farm, with $F = 14$ (7 halfsteps at 24 bins per octave), $T = 20$ (130ms in time), and $K = 2$ components. In this case, W_1^1 and W_2^1 hold percussive components, and W_1^2 and W_2^2 hold guitar components.

a factor $(t_{i+1} - t_i)/(s_{i+1} - s_i)$. The result is a snippet of A' which is the same length as the corresponding snippet in A . Henceforth, we will abuse notation and refer to these snippets as A and A' . We also extract a smaller snippet from B of the same length for reasons of memory efficiency, which we will henceforth refer to as B .

2.2 NMF2D for Joint Blind Factorization / Filtering

Once we have synchronized the snippets A and A' , we blindly factorize and filter them into K corresponding tracks A_1, A_2, \dots, A_K and A'_1, A'_2, \dots, A'_K . The goal is for each track to contain a different instrument. For instance, A_1 may contain an acoustic guitar, which is covered by an electric guitar contained in A'_1 , and A_2 may contain drums which are covered by a drum machine in A'_2 . This will make it possible to reduce the synthesis problem to K independent monophonic analogy problems in Section 2.3.

To accomplish this, the main tool we use is 2D convolutional nonnegative matrix factorization (2DNMF) [15]. We apply this algorithm to the magnitude of the constant-Q transforms (CQTs) [2] \mathbf{C}_A , $\mathbf{C}_{A'}$ and \mathbf{C}_B of A , A' , and B , respectively. Intuitively, we seek K time-frequency templates in A and A' which represent small, characteristic snippets of the K instruments we believe to be present in the audio. We then approximate the magnitude constant-Q transform of A and A' by convolving these templates in both time and frequency. The constant-Q transform is necessary in this framework, since a pitch shift can be approximated by a *linear* shift of all CQT bins by the same amount. Frequency shifts can then be represented as convolutions in the vertical direction, which is not possible with the ordinary STFT. Though it is more complicated, 2DNMF is a more compact representation than 1D convolutional NMF (in time only), in which it is necessary to store a different template for each pitch shift of each instrument. We note that pitch shifts in real instruments are more complicated than shifting all frequency bins by the same perceptual amount [14], but the basic version of 2DNMF is fine for our purposes.

More concretely, define a K -component, F -frequency, T -time lag 2D convolutional NMF decomposition for a matrix $\mathbf{X} \in \mathbb{R}^{M \times N}$ as follows

$$\mathbf{X} \approx \mathbf{\Lambda}_{\mathbf{W}, \mathbf{H}} = \sum_{\tau=1}^T \sum_{\phi=1}^F \mathbf{W}^{\tau} \mathbf{H}^{\phi} \quad (1)$$

where $\mathbf{W}^{\tau} \in \mathbb{R}^{M \times K}$ and $\mathbf{H}^{\phi} \in \mathbb{R}^{K \times N}$ store a τ -shifted time template and ϕ -frequency shifted coefficients, respectively. By $\downarrow_{\phi} A$, we mean down shift the rows of A by ϕ , so that row i of $\downarrow_{\phi} A$ is row $i - \phi$ of A , and the first ϕ rows of $\downarrow_{\phi} A$ are all zeros. And $\leftarrow_{\tau} A$ means left-shift A , so that column j of $\leftarrow_{\tau} A$ is column $j - \tau$ of A , and the first τ columns of $\leftarrow_{\tau} A$ are all zeros.

In our problem, we define an extension of 2DNMF to *jointly* factorize the 2 songs, A and A' , which each have M CQT coefficients. In particular, given matrices $\mathbf{C}_A, \mathbf{C}_{A'} \in \mathbb{C}^{M \times N_1}$ representing the complex CQT frames in each song over time, we seek $\mathbf{W}_1^\tau, \mathbf{W}_2^\tau \in \mathbb{R}^{M \times K}$ and $\mathbf{H}_1^\phi \in \mathbb{R}^{K \times N_1}$ that minimize the sum of the Kullback-Leibler divergences between the magnitude CQT coefficients and the convolutions:

$$D(|\mathbf{C}_A| \parallel \Lambda_{\mathbf{W}_1, \mathbf{H}_1}) + D(|\mathbf{C}_{A'}| \parallel \Lambda_{\mathbf{W}_2, \mathbf{H}_1}) \quad (2)$$

where the Kullback-Leibler divergence $D(X \parallel Y)$ is defined as

$$D(\mathbf{X} \parallel \mathbf{Y}) = \sum_i \sum_j \mathbf{X}_{i,j} \log \frac{\mathbf{X}_{i,j}}{\mathbf{Y}_{i,j}} - \mathbf{X}_{i,j} + \mathbf{Y}_{i,j} \quad (3)$$

That is, we *share* \mathbf{H}_1 between the factorizations of $|\mathbf{C}_A|$ and $|\mathbf{C}_{A'}|$ so that we can discover shared structure between the covers. Following similar computations to those of ordinary 2DNMF [15], it can be shown that Equation 2 is non-decreasing under the alternating update rules:

$$\mathbf{W}_1^\tau \leftarrow \mathbf{W}_1^\tau \odot \frac{\sum_{\phi=1}^F \left(\frac{\uparrow \phi}{\Lambda_{\mathbf{W}_1, \mathbf{H}_1}} \right) \mathbf{H}_1^\phi}{\sum_{\phi=1}^F \mathbf{1} \cdot \mathbf{H}_1^\phi} \quad (4)$$

$$\mathbf{W}_2^\tau \leftarrow \mathbf{W}_2^\tau \odot \frac{\sum_{\phi=1}^F \left(\frac{\uparrow \phi}{\Lambda_{\mathbf{W}_2, \mathbf{H}_1}} \right) \mathbf{H}_1^\phi}{\sum_{\phi=1}^F \mathbf{1} \cdot \mathbf{H}_1^\phi} \quad (5)$$

$$\mathbf{H}_1^\phi \leftarrow \mathbf{H}_1^\phi \odot \left(\frac{\sum_{\tau=1}^T \mathbf{W}_1^{\downarrow \phi \tau} \left(\frac{\leftarrow \tau}{\Lambda_{\mathbf{W}_1, \mathbf{H}_1}} \right) + \mathbf{W}_2^{\downarrow \phi \tau} \left(\frac{\leftarrow \tau}{\Lambda_{\mathbf{W}_2, \mathbf{H}_1}} \right)}{\sum_{\tau=1}^T \mathbf{W}_1^{\downarrow \phi \tau} \mathbf{1}^{\leftarrow \tau} + \mathbf{W}_2^{\downarrow \phi \tau} \mathbf{1}^{\leftarrow \tau}} \right) \quad (6)$$

where $\mathbf{1}$ is a column vector of all 1s of appropriate dimension. We need an invertible CQT to go back to audio templates, so we use the non-stationary Gabor Transform (NSGT) implementation of the CQT [24] to compute $\mathbf{C}_A, \mathbf{C}_{A'}$, and \mathbf{C}_B . We use 24 bins per octave between 50hz and 11.7kHz, for a total of 189 CQT bins. We also use $F = 14$ in most of our examples, allowing 7 halfstep shifts, and we use $T = 20$ on temporally downsampled CQTs to cover a timespan of 130 milliseconds. Finally, we iterate through Equations 4, 5, and 6 in sequence 300 times.

Note that a naive implementation of the above equations can be very computationally intensive. To ameliorate this, we implemented GPU versions of Equations 1, 4, 5, and 6. Equation 1 in particular is well-suited for a parallel implementation, as the shifted convolutional blocks overlap each other heavily and can be carefully offloaded into shared memory to exploit this². In practice, we witnessed a 30x speedup of our GPU implementation over our CPU implementation for 20 second audio clips for A , A' , and B .

Figure 2 shows a synthetic example with an exact solution, and Figure 4 shows a local min which is the result of running Equations 4, 5, and 6 on real audio from the “Smooth Criminal” example. It is evident from \mathbf{H}_1 that the first component is percussive (activations at regular intervals in \mathbf{H}_1^1 , and no pitch shifts), while the second component corresponds to the guitar melody (\mathbf{H}_1^2 appears like a “musical score” of sorts). Furthermore, \mathbf{W}_1^1 and \mathbf{W}_2^1 have broadband frequency content consistent with percussive events, while \mathbf{W}_1^2 and \mathbf{W}_2^2 have visible harmonics consistent with vibrating strings. Note that we generally use more than $K = 2$, which allows finer granularity than harmonic/percussive, but even for $K = 2$ in this example, we observe qualitatively better separation than off-the-shelf harmonic/percussive separation algorithms [6].

Once we have $\mathbf{W}_1, \mathbf{W}_2$, and \mathbf{H}_1 , we can recover the audio templates A_1, A_2, \dots, A_K and A'_1, A'_2, \dots, A'_K by using the components of \mathbf{W}_1 and \mathbf{W}_2 as filters. First, define $\Lambda_{\mathbf{W}, \mathbf{H}, \mathbf{k}}$ as

²In the interest of space, we omit more details of our GPU-based NMFD in this paper, but a documented implementation can be found at <https://github.com/ctrlalieu/CoverSongSynthesis/>, and we plan to release more details in a companion paper later.

$$\Lambda_{\mathbf{W}, \mathbf{H}, \mathbf{k}} = \sum_{\tau=1}^T \sum_{\phi=1}^F \mathbf{W}^{\mathbf{k}^\tau} \mathbf{H}_{\mathbf{k}}^{\phi} \quad (7)$$

where $\mathbf{W}^{\mathbf{k}}$ is the k^{th} column of \mathbf{W} and $\mathbf{H}_{\mathbf{k}}$ is the k^{th} row of \mathbf{H} . Now, define the filtered CQTs by using soft masks derived from \mathbf{W}_1 and \mathbf{H}_1 :

$$\mathbf{C}_{\mathbf{A}_k} = \mathbf{C}_{\mathbf{A}} \odot \left(\frac{\Lambda_{\mathbf{W}_1, \mathbf{H}_1, \mathbf{k}}^p}{\sum_{m=1}^K \Lambda_{\mathbf{W}_1, \mathbf{H}_1, \mathbf{m}}^p} \right) \quad (8)$$

$$\mathbf{C}_{\mathbf{A}'_k} = \mathbf{C}_{\mathbf{A}'} \odot \left(\frac{\Lambda_{\mathbf{W}_2, \mathbf{H}_1, \mathbf{k}}^p}{\sum_{m=1}^K \Lambda_{\mathbf{W}_2, \mathbf{H}_1, \mathbf{m}}^p} \right) \quad (9)$$

where p is some positive integer applied element-wise (we choose $p = 2$), and the above multiplications and divisions are also applied element-wise. It is now possible to invert the CQTs to uncover the audio templates, using the inverse NSGT [24]. Thus, if the separation was good, we are left with K independent monophonic pairs of sounds between A and A' .

In addition to inverting the sounds after these masks are applied, we can also listen to the components of \mathbf{W}_1 and \mathbf{W}_2 themselves to gain insight into how they are behaving as filters. Since \mathbf{W}_1 and \mathbf{W}_2 are magnitude only, we apply the Griffin Lim algorithm [11] to perform phase retrieval, and then we invert them as before to obtain 130 millisecond sounds for each k .

2.3 Musaicing And Mixing

We now describe how to use the corresponding audio templates we learned in Section 2.2 to perform style transfer on a new piece of audio, B .

2.3.1 Separating Tracks in B

First, we compute the CQT of B , $\mathbf{C}_{\mathbf{B}} \in \mathbb{C}^{M \times N_2}$. We then represent its magnitude using \mathbf{W}_1 as a basis, so that we filter B into the same set of instruments into which A was separated. That is, we solve for \mathbf{H}_2 so that $|\mathbf{C}_{\mathbf{B}}| \approx \Lambda_{\mathbf{W}_1, \mathbf{H}_2}$. This can be performed with ordinary 2DNMF, holding \mathbf{W}_1 fixed; that is, repeating the following update until convergence

$$\mathbf{H}_2^\phi \leftarrow \mathbf{H}_2^\phi \odot \left(\frac{\sum_{\tau=1}^T \mathbf{W}_1^{\tau} \left(\frac{|\mathbf{C}_{\mathbf{B}}|}{\Lambda_{\mathbf{W}_1, \mathbf{H}_2}} \right)^{\leftarrow \tau}}{\sum_{\tau=1}^T \mathbf{W}_1^{\tau} \mathbf{1}^{\leftarrow \tau}} \right) \quad (10)$$

As with A and A' , we can now filter B into a set of audio tracks B_1, B_2, \dots, B_K by first computing filtered CQTs as follows

$$\mathbf{C}_{\mathbf{B}_k} = \mathbf{C}_{\mathbf{B}} \odot \left(\frac{\Lambda_{\mathbf{W}_1, \mathbf{H}_2, \mathbf{k}}^p}{\sum_{m=1}^K \Lambda_{\mathbf{W}_1, \mathbf{H}_2, \mathbf{m}}^p} \right) \quad (11)$$

and then inverting them.

2.3.2 Constructing B' Track by Track

At this point, we could use \mathbf{H}_2 and let our cover song CQT magnitudes $|\mathbf{C}_{\mathbf{B}'}| = \Lambda_{\mathbf{W}_2, \mathbf{H}_2}$, followed by Griffin Lim to recover the phase. However, we have found that the resulting sounds are too “blurry,” as they lack all but re-arranged low rank detail from A' . Instead, we choose to draw sound grains from the inverted, filtered tracks from A' , which contain all of the detail of the original song. For this, we first reconstruct each track of B using audio grains from the corresponding tracks in A , and then we replace each track with the corresponding track in B . To accomplish this, we apply the audio musaicing technique of Driedger [7] to each track in B , using source audio A .

For computational and memory reasons, we now abandon the CQT and switch to the STFT with hop size $h = 256$ and window size $w = 2048$. More specifically, let N_1 be the number of STFT frames in A and A' and let N_2 be the number of STFT frames in B . For each A_k , we create an STFT matrix $\mathbf{S}_{\mathbf{A}_k}$ which holds the STFT of A_k concatenated to pitch shifted versions of A_k , so that pitches beyond what were in A can be represented, but by the same instruments. We use ± 6 halfsteps, so $\mathbf{S}_{\mathbf{A}_k} \in \mathbb{C}^{w \times 13N_1}$. We do the same for A'_k to create $\mathbf{S}_{\mathbf{A}'_k} \in \mathbb{C}^{w \times 13N_1}$. Finally, we compute the STFT of B_k without any shifts: $\mathbf{S}_{\mathbf{B}_k} \in \mathbb{C}^{w \times N_2}$.

Now, we apply Driedger’s technique, using $|\mathbf{S}_{\mathbf{A}_k}|$ as a spectral dictionary to reconstruct $|\mathbf{S}_{\mathbf{B}_k}|$ ($\mathbf{S}_{\mathbf{A}_k}$ is analogous to the buzzing bees in [7]). That is, we seek an $\mathbf{H}_k \in \mathbb{R}^{13N_1 \times N_2}$ so that

$$|\mathbf{S}_{\mathbf{B}_k}| \approx |\mathbf{S}_{\mathbf{A}_k}| \mathbf{H} \quad (12)$$

For completeness, we briefly re-state the iterations in Driedger’s technique [7]. For L iterations total, at the ℓ^{th} iteration, compute the following 4 updates in order. First, we restrict the number of repeated activations by filtering in a maximum horizontal neighborhood in time

$$\mathbf{R}_{\mathbf{km}}^{(\ell)} = \begin{cases} \mathbf{H}_{\mathbf{km}}^{(\ell)} & \text{if } \mathbf{H}_{\mathbf{km}}^{(\ell)} = \mu_{\mathbf{km}}^{\mathbf{r},(\ell)} \\ \mathbf{H}_{\mathbf{km}}^{(\ell)} (1 - \frac{n+1}{N}) & \text{otherwise} \end{cases} \quad (13)$$

where $\mu_{\mathbf{km}}^{\mathbf{r},(\ell)}$ holds the maximum in a neighborhood $\mathbf{H}_{\mathbf{k}, \mathbf{m}-\mathbf{r}, \mathbf{m}+\mathbf{r}}$ for some parameter r (we choose $r = 3$ in our examples). Next, we restrict the number of simultaneous activations by shrinking all of the values in each column that are less than the top p values in that column:

$$\mathbf{P}_{\mathbf{km}}^{(\ell)} = \begin{cases} \mathbf{R}_{\mathbf{km}}^{(\ell)} & \text{if } \mathbf{R}_{\mathbf{km}}^{(\ell)} \geq \mathbf{MP}^{(\mathbf{n})} \\ \mathbf{R}_{\mathbf{km}}^{(\ell)} (1 - \frac{n+1}{N}) & \text{otherwise} \end{cases} \quad (14)$$

where $\mathbf{MP}^{(\ell)}$ is a row vector holding the p^{th} largest value of each column of $\mathbf{R}^{(\ell)}$ (we choose $p = 10$ in our examples). After this, we promote time-continuous structures by convolving along all of the diagonals of \mathbf{H} :

$$\mathbf{C}_{\mathbf{km}}^{(\ell)} = \sum_{i=-c}^c \mathbf{P}_{(\mathbf{k}+\mathbf{i}),(\mathbf{m}+\mathbf{i})}^{(\ell)} \quad (15)$$

We choose $c = 3$ in our examples. Finally, we perform the ordinary KL-based NMF update:

$$\mathbf{H}^{(\ell+1)} \leftarrow \mathbf{H}^{(\ell)} \odot \frac{|\mathbf{S}_{\mathbf{A}_k}|^T \frac{|\mathbf{S}_{\mathbf{B}_k}|}{|\mathbf{S}_{\mathbf{A}_k}| \mathbf{C}^{(\ell)}}}{|\mathbf{S}_{\mathbf{A}_k}| \cdot \mathbf{1}} \quad (16)$$

We perform 100 such iterations ($L = 100$). Once we have our final \mathbf{H} , we can use this to create B'_k as follows:

$$\mathbf{S}_{\mathbf{B}'_k} = \mathbf{S}_{\mathbf{A}'_k} \mathbf{H} \quad (17)$$

In other words, we use the learned activations to create $\mathbf{S}_{\mathbf{B}_k}$ using $\mathbf{S}_{\mathbf{A}_k}$, but we instead use these activations with the dictionary from $\mathbf{S}_{\mathbf{A}'_k}$. *This is the key step in the style transfer*, and it is done for the same reason that \mathbf{H}_1 is shared between A and A' . Figure 5 shows an example for the guitar track on Michael Jackson’s “Bad,” translating to Alien Ant Farm using A and A' templates from Michael Jackson’s and Alien Ant Farms’ versions of “Smooth Criminal,” respectively. It is visually apparent that $|\mathbf{S}_{\mathbf{A}_k}| \mathbf{H} \approx \mathbf{S}_{\mathbf{B}_k}$, it is also apparent that $\mathbf{S}_{\mathbf{B}'_k} = \mathbf{S}_{\mathbf{A}'_k} \mathbf{H}$ is similar to $\mathbf{S}_{\mathbf{B}_k}$, except it has more energy in the higher frequencies. This is consistent with the fact that Alien Ant Farm uses a more distorted electric guitar, which has more broadband energy.

To create our final B' , we simply add all of the STFTs $\mathbf{S}_{\mathbf{B}'_k}$ together for each k , and we perform and inverse STFT to go back to audio.




Figure 5: Driedger’s technique [7] using audio grains from pitch shifted versions of A_k (the k^{th} track in Michael Jackson’s “Smooth Criminal”) to create B_k (the k^{th} track in Michael Jackson’s “Bad”), and using the activations to create B'_k (the k^{th} synthesized track in Alien Ant Farm’s “Bad”).

2.4 A Note About Tempos

The algorithm we have described so far assumes that the tempos $t_A, t_{A'}$, and t_B of A, A' , and B , respectively are similar. This is certainly not true in more interesting covers. Section 2.1 took care of the disparity between A and A' during the synchronization. However, we also need to perform a tempo scaling on B by a factor of t_A/t_B before running our algorithm. Once we have computed B' , whose tempo is initially t_A , we scale its tempo back by $(t_B/t_A) \cdot (t_{A'}/t_A)$. For instance, suppose that $t_A = 60$, $t_{A'} = 80$, and $t_B = 120$. Then the final tempo of B' will be $60 \times (120/60) \times (80/60) = 160$ bpm.

3 Experimental Examples

We now qualitatively explore our technique on several examples³. In all of our examples, we use $K = 3$ sources. This is a hyperparameter that should be chosen with care in general, since we would like each component to correspond to a single instrument or group of related instruments. However, our initial examples are simple enough that we expect basically a harmonic source, a percussive source, and one other source or sub-separation between either harmonic and percussive.

First, we follow through on the example we have been exploring and we synthesize Alien Ant Farm’s “Bad” from Michael Jackson’s “Bad” (B), using “Smooth Criminal” as an example. The translation of the guitar from synth to electric is clearly audible in the final result. Furthermore, a track which was exclusively drums in A included some extra screams in A' that Alien Ant Farm performed as embellishments. These

³Please listen to our results at <http://www.covers1000.net/analogies.html>

embellishments transferred over to B' in the “Bad” example, further reinforcing Alien Ant Farm’s style. Note that these screams would not have been preserved had we simply used inverted the CQT $\mathbf{A}_{\mathbf{W}_2, \mathbf{H}_2}$, but they are present in one of the filtered tracks and available as audio grains during mosaicing for that track.

In addition to the “Bad” example, we also synthesize Alien Ant Farm’s version of “Wanna Be Startin Something,” using “Smooth Criminal” as an example for A and A' once again. In this example, Alien Ant Farm’s screams occur consistently with the fast drum beat every measure.

Finally, we explore an example with a more extreme tempo shift between A and A' ($t_{A'} < t_A$). We use Marilyn Manson’s cover of “Sweet Dreams” by the Eurythmics to synthesize a Marilyn Manson cover of “Who’s That Girl” by the Eurythmics. We found that in this particular example, we obtained better results when we performed 2DNMF on $|\mathbf{C}_A|$ by itself first, and then we performed the optimizations in Equation 5 and Equation 6, holding \mathbf{W}_1 fixed. This is a minor tweak that can be left to the discretion of the user at runtime.

Overall, our technique works well for instrument translation in these three examples. However, we did notice that the vocals did not carry over at all in any of them. This is to be expected, since singing voice separation often assumes that the instruments are low rank and the voice is high rank [21], and our filters and final mosaicing are both derived from low rank NMF models.

4 Discussion / Future Directions

In this work, we demonstrated a proof of concept, fully automated end to end system which can synthesize a cover song snippet of a song B given an example cover A and A' , where A is by the same artist as B , and B' should sound like B but in the style of A' . We showed some promising initial results on a few examples, which is, to our knowledge, one of the first steps in the challenging direction of automatic polyphonic audio mosaicing.

Our technique does have some limitations, however. Since we use \mathbf{W}_1 for both A and B , we are limited to examples in which A and B have similar instruments. It would be interesting to explore how far one could push this technique with different instruments between A and B , which happens quite often even within a corpus by the same artist.

We have also noticed that in addition to singing voice, other “high rank” instruments, such as the fiddle, cannot be properly translated. We believe that that translating such instruments and voices would be an interesting and challenging future direction of research, and it would likely need a completely different approach to the one we presented here.

Finally, out of the three main steps of our pipeline, synchronization (Section 2.1), blind joint factorization/source separation (Section 2.2), and filtering/mosaicing (Section 2.3), the weakest step by far is the blind source separation. The single channel source separation problem is still far from solved in general even without the complication of cover songs, so that will likely remain the weakest step for some time. If one has access to the unmixed studio tracks for A , A' , and B , though, that step can be entirely circumvented; the algorithm would remain the same, and one would expect higher quality results. Unfortunately, such tracks are difficult to find in general for those who do not work in a music studio, which is why blind source separation also remains an important problem in its own right.

5 Acknowledgements

Christopher Tralie was partially supported by an NSF big data grant DKA-1447491 and an NSF Research Training Grant NSF-DMS 1045133. We also thank Brian McFee for helpful discussions about invertible CQTs.

References

- [1] Shunya Ariga, Satoru Fukayama, and Masataka Goto. Song2guitar: A difficulty-aware arrangement system for generating guitar solo covers from polyphonic audio of popular music. In *18th International Society for Music Information Retrieval (ISMIR)*, 2017.

- [2] Judith C Brown. Calculation of a constant q spectral transform. *The Journal of the Acoustical Society of America*, 89(1):425–434, 1991.
- [3] C Cannam. Rubber band library. *Software released under GNU General Public License (version 1.8.1)*, 2012.
- [4] Ning Chen, Wei Li, and Haidong Xiao. Fusing similarity functions for cover song identification. *Multimedia Tools and Applications*, pages 1–24, 2017.
- [5] Christian Dittmar and Meinard Müller. Reverse Engineering the Amen Break – Score-informed Separation and Restoration applied to Drum Recordings. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(9):1531–1543, 2016.
- [6] Jonathan Driedger, Meinard Müller, and Sascha Disch. Extending harmonic-percussive separation of audio signals. In *ISMIR*, pages 611–616, 2014.
- [7] Jonathan Driedger, Thomas Prätzlich, and Meinard Müller. Let it bee-towards nmf-inspired audio mosaicing. In *ISMIR*, pages 350–356, 2015.
- [8] Daniel PW Ellis. Beat tracking by dynamic programming. *Journal of New Music Research*, 36(1):51–60, 2007.
- [9] Hadrien Foughmand and Geoffroy Peeters. Multi-source musaicing using non-negative matrix factor 2-d deconvolution. In *18th International Society for Music Information Retrieval (ISMIR), Late Breaking Session*, 2017.
- [10] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*, 2015.
- [11] Daniel Griffin and Jae Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(2):236–243, 1984.
- [12] Aaron Hertzmann, Charles E Jacobs, Nuria Oliver, Brian Curless, and David H Salesin. Image analogies. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 327–340. ACM, 2001.
- [13] Florian Krebs, Sebastian Böck, and Gerhard Widmer. An efficient state-space model for joint tempo and meter tracking. In *ISMIR*, pages 72–78, 2015.
- [14] Tomohiko Nakamura and Hirokazu Kameoka. Shifted and convolutive source-filter non-negative matrix factorization for monaural audio source separation. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 489–493. IEEE, 2016.
- [15] Mikkel N Schmidt and Morten Mørup. Nonnegative matrix factor 2-d deconvolution for blind single channel source separation. In *International Conference on Independent Component Analysis and Signal Separation*, pages 700–707. Springer, 2006.
- [16] Diemo Schwarz, Roland Cahen, and Sam Britton. Principles and applications of interactive corpus-based concatenative synthesis. In *Journées d’Informatique Musicale (JIM)*, pages 1–1, 2008.
- [17] Joan Serra, Xavier Serra, and Ralph G Andrzejak. Cross recurrence quantification for cover song identification. *New Journal of Physics*, 11(9):093017, 2009.
- [18] Diego F Silva, Chin-Chin M Yeh, Gustavo Enrique de Almeida Prado Alves Batista, Eamonn Keogh, et al. Simple: assessing music similarity using subsequences joins. In *International Society for Music Information Retrieval Conference, XVII*. International Society for Music Information Retrieval-ISMIR, 2016.
- [19] Ian Simon, Sumit Basu, David Salesin, and Maneesh Agrawala. Audio analogies: Creating new music from an existing performance by concatenative synthesis. In *ICMC*. Citeseer, 2005.

- [20] Temple F Smith and Michael S Waterman. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.
- [21] Pablo Sprechmann, Alexander M Bronstein, and Guillermo Sapiro. Real-time online singing voice separation from monaural recordings using robust low-rank modeling. In *ISMIR*, pages 67–72, 2012.
- [22] Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. In *ACM Transactions on Graphics (TOG)*, volume 23, pages 399–405. ACM, 2004.
- [23] Christopher J Tralie. Mfcc and hpcp fusion for robust cover song identification. In *18th International Society for Music Information Retrieval (ISMIR)*, 2017.
- [24] Gino Angelo Velasco, Nicki Holighaus, Monika Dörfler, and Thomas Grill. Constructing an invertible constant-q transform with non-stationary gabor frames. *Proceedings of DAFX11, Paris*, pages 93–99, 2011.
- [25] Bo Wang, Jiayan Jiang, Wei Wang, Zhi-Hua Zhou, and Zhuowen Tu. Unsupervised metric fusion by cross diffusion. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2997–3004. IEEE, 2012.
- [26] Bo Wang, Aziz M Mezlini, Feyyaz Demir, Marc Fiume, Zhuowen Tu, Michael Brudno, Benjamin Haibe-Kains, and Anna Goldenberg. Similarity network fusion for aggregating data types on a genomic scale. *Nature methods*, 11(3):333–337, 2014.