

Adaptive synchrosqueezing based on a quilted short-time Fourier transform

Alexander Berrian* and Naoki Saito†

Department of Mathematics, University of California, One Shields Avenue, Davis, CA 95616, USA

ABSTRACT

In recent years, the *synchrosqueezing transform* (SST) has gained popularity as a method for the analysis of signals that can be broken down into multiple components determined by instantaneous amplitudes and phases. One such version of SST, based on the short-time Fourier transform (STFT), enables the sharpening of instantaneous frequency (IF) information derived from the STFT, as well as the separation of amplitude-phase components corresponding to distinct IF curves. However, this SST is limited by the time-frequency resolution of the underlying window function, and may not resolve signals exhibiting diverse time-frequency behaviors with sufficient accuracy. In this work, we develop a framework for an SST based on a “quilted” short-time Fourier transform (SST-QSTFT), which allows adaptation to signal behavior in separate time-frequency regions through the use of multiple windows. This motivates us to introduce a discrete reassignment frequency formula based on a finite difference of the phase spectrum, ensuring computational accuracy for a wider variety of windows. We develop a theoretical framework for the SST-QSTFT in both the continuous and the discrete settings, and describe an algorithm for the automatic selection of optimal windows depending on the region of interest. Using synthetic data, we demonstrate the superior numerical performance of SST-QSTFT relative to other SST methods in a noisy context. Finally, we apply SST-QSTFT to audio recordings of animal calls to demonstrate the potential of our method for the analysis of real bioacoustic signals.

Keywords: synchrosqueezing, instantaneous frequency, short-time Fourier transform, adaptive time-frequency representations, time-frequency analysis, audio signal processing, chirped windows, reassigned spectrogram

1. INTRODUCTION

In the field of signal processing, one frequently seeks to model a signal with time-varying oscillatory properties as a sum of distinct *amplitude-phase components* containing information on instantaneous amplitudes (IAs) and instantaneous frequencies (IFs). In particular, one might characterize a signal $f : \mathbb{R} \rightarrow \mathbb{C}$ using an *amplitude-phase decomposition* given by $f(t) = \sum_{m=1}^M f_m(t)$; $f_m(t) := A_m(t)e^{2\pi i\phi_m(t)}$, where the $\{A_m\}$ represent *instantaneous amplitudes* (IAs), the $\{\phi_m\}$ represent *instantaneous phases* (IPs), and the $\{\phi'_m\}$ represent *instantaneous frequencies* (IFs). Then, the problem is to retrieve the IFs ϕ'_m and amplitude-phase components (*modes*) f_m , given that only f is known.

One may use a time-frequency representation such as the short-time Fourier transform (STFT) or continuous wavelet transform (CWT) to analyze the signal. However, these transforms provide blurry amplitude and frequency information, thereby complicating the task of accurately determining each amplitude-phase component. Hence, often a post-processing method is used to sharpen the blurry signal information. The tool of interest to us is the *synchrosqueezing transform* (SST). Originally introduced by Daubechies and Maes¹ in the context of audio signal processing, the SST is a time-frequency representation that provides a sharpened picture of IAs and IFs and enables the separation and reconstruction of each separate amplitude-phase component for a certain class of signals.²⁻⁴

The SST technique has been applied to problems in many different disciplines, and researchers have further developed the SST idea to work in an increasingly large collection of signal processing contexts. In 2011, Daubechies, Lu, and Wu² gave theoretical proofs of the effectiveness of the CWT-based version of the SST

* Email: aberrian@math.ucdavis.edu – Website: <http://www.math.ucdavis.edu/~aberrian>

† Email: saito@math.ucdavis.edu – Website: <http://www.math.ucdavis.edu/~saito>

and demonstrated its applicability to sets of medical data. The same year, Thakur and Wu³ generalized the SST to a version based on STFT. Since then, there has been an uptick of research on SST, and the technique has been successfully adapted to a number of physical problems, including speaker identification from an audio signal,¹ fault diagnosis in planetary gearboxes for wind turbines,⁵ quantifying the effect of solar radiation on a paleoclimate change on Earth,⁶ and extracting heart-rate variability (HRV) from an ECG signal.²

However, researchers continue to develop techniques for cases where the original SST techniques are insufficient. Often, this insufficiency is due to the resolution constraints of the underlying time-frequency representation (STFT or CWT). One area of recent work can be categorized under the umbrella of *non-stationary Gabor transforms* (NSGT),⁷ where one analyzes the signal using a kernel function with changing time-frequency resolution over the time-frequency plane. This allows for generalizations of the STFT and CWT that allow the user to adapt the time-frequency resolution to the signal content. Balazs et al.⁷ developed the idea of NSGT in a frame theory context, and Dörfler further generalized this idea to the notion of *quilted Gabor frames*.⁸ Her work refers to the time-frequency plane as a “quilt” with different “patches” corresponding to regions where a signal exhibits different time-frequency behavior.

In the context of SST, recent developments under the umbrella of NSGT include generalizations of SST where different windows may be used to adapt to the signal at different times,⁹ or at different frequencies.¹⁰ In a previous work,¹¹ we further generalized this idea by introducing an SST based on a *quilted STFT* (QSTFT), where the window is allowed to change depending on the time-frequency region of interest. More recently, Sheu et al. derived an algorithm for SST with the automatic selection of optimal windows in different time-frequency regions,¹² but did not provide a method for signal reconstruction in the context of joint time-frequency window variation. In this work, we provide theoretical results for the effectiveness of the SST-QSTFT in both the continuous and the discrete context, and we derive a slightly different algorithm for automatic window selection. Furthermore, we explore the usage of chirped windows in depth and provide a methodology for signal reconstruction when the selected windows vary in both time and frequency.

We give an outline of this paper as follows. In Sec. 2, we review the relevant background material on the STFT and the corresponding synchrosqueezing transform. We then introduce the QSTFT and SST-QSTFT in the continuous setting (Sec. 3) and in the discrete setting (Sec. 4), and we provide theorems that demonstrate the effectiveness of the SST-QSTFT for accurate IF detection and mode reconstruction. Numerical results and applications follow in Sec. 5. We conclude with a summary of our results in Sec. 6.

2. BACKGROUND

2.1 Admissible function class

In this analysis, we restrict our study to a class of signals satisfying certain theoretical properties.

DEFINITION 2.1. A signal f is said to be in the *weakly modulated IA and IF signal class* $\mathcal{B}_{\epsilon,d}$ ^{2-4,6} if for some $M \in \mathbb{N}$ we can write

$$f(t) = \sum_{m=1}^M f_m(t); \quad f_m(t) := A_m(t)e^{2\pi i\phi_m(t)}; \quad (1)$$

and if there exist $\epsilon, d > 0$ such that for each $m \in \{1, \dots, M\}$,

- A_m and ϕ'_m are *bounded and sufficiently smooth*: $A_m \in C^1 \cap L^\infty$, $\phi_m \in C^2$, $\phi'_m \in L^\infty$, $\inf_{t \in \mathbb{R}} A_m(t) > 0$ and $\inf_{t \in \mathbb{R}} \phi'_m(t) > 0$;
- the IA A_m and IF ϕ'_m are *slowly-varying*: $\|A'_m\|_\infty \leq \epsilon$ and $\|\phi''_m\|_\infty \leq \epsilon$;
- ϕ'_m is *well-separated* from the other IFs: $\phi'_m(t) - \phi'_{m-1}(t) > d$ for each $t \in \mathbb{R}$, provided that $2 \leq m \leq M$.

The assumptions above are necessary to ensure the accurate isolation of the IFs $\{\phi'_m\}$ and reconstruction of the modes $\{f_m\}$ in the theory that follows.

REMARK 2.2. One models real-valued signals by the real part in (1), yielding $f_{\text{real}}(t) = \sum_{m=1}^M A_m(t) \cos(2\pi\phi_m(t))$.

2.2 Fourier transforms

The standard building block for determining the spectral information of a signal is the *Fourier transform*, which comes in different forms and can be defined in many different ways. In this paper, we define the *continuous Fourier transform* of a continuous-time signal $f \in L^2(\mathbb{R}) \cap L^1(\mathbb{R})$ by $\hat{f}(\xi) := \int_{\mathbb{R}} f(x)e^{-2\pi i x \xi} dx$ for each $\xi \in \mathbb{R}$.

Also, for the purposes of this paper, we define the *semi-discrete Fourier transform* of the periodic sequence g by $\hat{g}(u) := \sum_{\ell=0}^{L-1} g[\ell]e^{-2\pi i u \ell}$ for the continuous argument $u \in \mathbb{R}$.

REMARK 2.3. It is a slight abuse of notation to use the same hat symbol for the continuous Fourier transform \hat{f} of a continuously-time function f and for the semi-discrete Fourier transform \hat{g} of a (discrete) periodic sequence g . However, the separate meaning of these two transforms is clear based on whether the transform input is continuously or discretely valued. Moreover, in the rest of this paper, we will only consider the Fourier transform of continuously-time window functions and discrete periodic window sequences as defined in Sec. 2.3, and we will not use the hat symbol in ambiguous cases where (for instance) a discrete periodic window sequence is explicitly said to be derived from discretizing a continuous-time window function.

2.3 Short-time Fourier transform (STFT)

In order to analyze the local frequency content at different times of the signal f , we first consider the *continuous short-time Fourier transform* of f with respect to a window function $g \in L^2(\mathbb{R})$ centered at 0, which we define¹³ by

$$\mathcal{V}_g f(t, \xi) := \int_{\mathbb{R}} f(x) \overline{g(x-t)} e^{-2\pi i \xi(x-t)} dx.$$

Here, the effect of the window g is to essentially truncate the signal f around the time of interest t in a smooth fashion, in order to enable the user to examine the local spectral information of the signal near the time t . If the window g is compactly supported, then we say g is *time-limited*. If \hat{g} is compactly supported, then we say g is *band-limited*. Note that a function h is said to be *compactly supported* if the set $\text{supp}(h) := \{x \in \mathbb{R} : h(x) \neq 0\}$ satisfies $\text{supp}(h) \subseteq [-R, R]$ for some finite $R > 0$.

In the context of discrete implementation, we assume that f is sampled at sampling rate f_s , denoting $\Delta t := 1/f_s$. Here and in the rest of this paper, we abuse notation slightly by denoting $f[\ell] := f(\ell\Delta t)$, $f_m[\ell] := f_m(\ell\Delta t)$, $A_m[\ell] := A_m(\ell\Delta t)$, $A'_m[\ell] := A'_m(\ell\Delta t)$, and similarly for ϕ_m and all its derivatives.[‡] We also further assume that f is strictly zero outside a compact region having left boundary at index $\ell = 0$, as it is practical to assume for many real-world applications where we are only interested in a finite-length signal.

We may now define the *discrete short-time Fourier transform* of f for a discrete, time-limited window function g sampled at L points, with $g[\ell] := g_\ell$ for $\ell = 0, \dots, L-1$, and with *hop size* $H \in \mathbb{N}$, via¹³

$$V_g f[n, k] := \sum_{\ell=0}^{L-1} f[\ell + nH] \overline{g[\ell]} e^{-2\pi i k \ell / L}, \quad (2)$$

for *frequency bins* $k \in \{0, \dots, L-1\}$, and *frames* $n \in \{0, \dots, N-1\}$ where $N \in \mathbb{N}$ is large enough so that $f[\ell] = 0$ for all $\ell \geq (L-1) + (N-1)H$. For signals with many samples, it is often practical to set the hop size H larger than 1, yielding a subsampling of the STFT. In such a context, an *overlap-add formula* can provide a perfect reconstruction of the original signal, provided that the window satisfies a certain formula for the given hop size^{14§}.

[‡]In this paper, we adhere to the convention of using square brackets for functions with discrete inputs, and parentheses for functions with continuous inputs.

[§]See http://ccrma.stanford.edu/~jos/sasp/Choice_Hop_Size.html for the specific formula, and which windows satisfy it.




Figure 1. **Top left:** Glockenspiel signal f , sampled at $f_s := 44100$ Hz. **Top right:** Window functions used to analyze f . *Short window (in blue):* Blackman window g^1 of length 2000 samples. *Tall window (in red):* Blackman window g^2 of length 500 samples. g^1 and g^2 are normalized to have unit ℓ^1 -norm and zero-padded to 2048 and 512 samples respectively. **Bottom left:** $\log_2(1 + |V_{g^1} f[n, k]|^2)$, for times 2.7 to 4.1 seconds and frequencies 2000 to 7500 Hz. **Bottom right:** $\log_2(1 + |V_{g^2} f[n, k]|^2)$, over the same frequency range. $V_{g^1} f$ and $V_{g^2} f$ are both computed with hop size $H = 250$ samples. In the STFT plots, the vertical content corresponds to transient events (note onsets) and the horizontal content corresponds to tonal events (enduring notes). With g^1 , the frequency resolution is better, while the time resolution is worse, leading to better visualization of note pitches but worse visualization of note onset times. With g^2 , the time resolution is better, leading to sharp onset resolution, but the frequency resolution is worse, leading to obscured pitch information.

The STFT allows one to visualize the IAs and IFs of the signal f , as well as the presence of transient and noise content spanning over many frequencies. Due to the Fourier uncertainty principle^{15, Sec. 7.3}, there is a tradeoff between the time precision and the frequency precision of the STFT. That is, a window g that enables the STFT to render more precise frequency information will yield less precise time information, and vice-versa. This point is illustrated using STFTs with respect to two different windows in Figure 1. Thus, the Fourier uncertainty principle limits the robustness of the STFT, and so various methods have been designed to sharpen the STFT information and approximately retrieve the IAs and IFs of a signal. In the next subsection, we describe the *STFT-based synchrosqueezing transform*, which enables the sharpening of the frequency information derived from the STFT.

2.4 Synchrosqueezing transform based on STFT (SST-STFT)

2.4.1 Continuous theory

The idea of the continuous STFT-based synchrosqueezing transform³ is to calculate a *reassignment frequency* $\Xi_g f$ for each STFT coefficient with magnitude above a certain tolerance. This reassignment frequency is an estimate of the actual frequency location of the nearest IF. Then, for each time-frequency point (t, ξ) , the synchrosqueezing transform finds all other frequencies η where the STFT coefficient $\mathcal{V}_g f(t, \eta)$ at this time has reassignment frequency $\Xi_g f(t, \eta) = \xi$, and approximately sums up all such STFT coefficients to yield the SST-STFT coefficient $\mathcal{S}_{g, \gamma}^\beta f(t, \xi)$. (We will explain the meaning of “approximately” summing shortly.) The effect is to “squeeze” the blurred-out STFT visualization in the frequency direction, leading to a better-concentrated time-frequency representation along the frequency axis. The SST-STFT is defined as follows.

DEFINITION 2.4. The *continuous STFT-based synchrosqueezing transform* (SST-STFT)³ with tolerance $\gamma \geq 0$

and limiting parameter $\beta > 0$ is given by

$$\mathcal{S}_{g,\gamma}^\beta f(t, \xi) := \int_{\mathcal{A}_{g,\gamma} f(t)} \mathcal{V}_g f(t, \eta) \frac{1}{\beta} b\left(\frac{\xi - \Xi_g f(t, \eta)}{\beta}\right) d\eta,$$

where $b \in C_c^\infty(\mathbb{R})$ is a ‘‘bump function’’ satisfying $\hat{b}(0) = 1$, $\mathcal{A}_{g,\gamma} f(t) := \{\eta \in \mathbb{R}^+ : |\mathcal{V}_g f(t, \eta)| > \gamma\}$, and $\Xi_g f(t, \eta) := \frac{\partial_t \mathcal{V}_g f(t, \eta)}{2\pi i \mathcal{V}_g f(t, \eta)}$ is the *STFT-based reassignment frequency*.

REMARK 2.5. As $\beta \searrow 0$, the term $\frac{1}{\beta} b\left(\frac{\xi - \Xi_g f(t, \eta)}{\beta}\right)$ converges in the distributional sense to $\delta(\xi - \Xi_g f(t, \eta))$, where δ denotes the Dirac delta. This is the meaning of the ‘‘approximate summation’’ described earlier.

REMARK 2.6. The choice of $\Xi_g f$ is motivated by the fact that for a constant chirp $f(t) := Ae^{2\pi i c t}$, with $A, c > 0$, we have exactly $\Xi_g f(t, \eta) = c$ for all (t, η) .^{1,2}

REMARK 2.7. Here, $C_c^\infty(\mathbb{R})$ denotes the class of functions where derivatives of all orders exist and are continuous (a property which we call *C[∞]-smoothness*) and where each function is compactly supported.

The following theoretical result ensures the accuracy of the SST-STFT:

THEOREM 2.8. (Thakur & Wu;³ Oberlin, Meignen & Perrier⁴) Let $\epsilon > 0$, $\nu \in (0, 1/2)$, $\tilde{\epsilon} := \epsilon^\nu$, $d > 0$. Suppose that $f = \sum_{m=1}^M f_m \in \mathcal{B}_{\epsilon, d}$. Let $g \in S(\mathbb{R})$, where $S(\mathbb{R})$ denotes the Schwartz class of *C[∞]-smooth, rapidly-decaying functions*. Assume that g is real-valued and satisfies $\text{supp}(\hat{g}) \subseteq [-d/2, d/2]$. Then, if ϵ is sufficiently small:

- **(Concentration of STFT around IF curves)** $|\mathcal{V}_g f(t, \xi)| > \tilde{\epsilon}$ only when there is an $m \in \{1, \dots, M\}$ such that $(t, \xi) \in \mathcal{Z}_m := \{(t, \xi) \in \mathbb{R} \times \mathbb{R}^+ : |\phi'_m(t) - \xi| < d/2\}$.
- **(Closeness of reassignment frequency $\Xi_g f$ to nearby IF)** For all $m \in \{1, \dots, M\}$ and all $(t, \xi) \in \mathcal{Z}_m$ such that $|\mathcal{V}_g f(t, \xi)| > \tilde{\epsilon}$, we have $|\Xi_g f(t, \xi) - \phi'_m(t)| \leq \tilde{\epsilon}$.
- **(Accuracy of reconstruction)** For each $m \in \{1, \dots, M\}$ there is a constant $C_m > 0$ such that for all $t \in \mathbb{R}$,

$$\left| \lim_{\beta \rightarrow 0^+} \left(\int_{\{\xi : |\xi - \phi'_m(t)| < \tilde{\epsilon}\}} \frac{1}{g(0)} \cdot \mathcal{S}_{g,\tilde{\epsilon}}^\beta f(t, \xi) d\xi \right) - f_m(t) \right| \leq C_m \tilde{\epsilon}. \quad (3)$$

REMARK 2.9. The assumption that g is real-valued in Theorem 2.8 is not essential. In Sec. 3.5, we will state a generalization of Theorem 2.8 (Theorem 3.6), from which the statement of Theorem 2.8 will follow for complex-valued g , provided that one normalizes by $\overline{g(0)}$ instead of $g(0)$ in the reconstruction formula given in (3).

2.4.2 Discrete theory

The discrete theory of the SST-STFT was partially addressed in the work of Thakur and Wu,³ resulting in a theorem concerning the concentration of the STFT around the IF curves and the closeness of the reassignment frequency to the nearby IF. The authors additionally showed robustness to noise and accuracy in the setting of nonuniform sampling. However, the theory prescribes a reassignment frequency formula that may lead to aliasing, for the following reason. Thakur and Wu, in computing the analogue of $\partial_t \mathcal{V}_g f$ for the discrete case, passed the derivative ∂_t through the integral and computed

$$\partial_t \mathcal{V}_g f(t, \xi) = \int_{\mathbb{R}} f(x) (2\pi i \xi g(x-t) - g'(x-t)) e^{-2\pi i \xi(x-t)} dx,$$

which enables one to derive an analogous discrete reassignment frequency formula to replace $\partial_t \mathcal{V}_g f$. However, if g' is nonzero at the boundaries of the support of g (i.e., if g' does not tail off to zero when g tails off to zero), then an aliased discrete Fourier transform (DFT) will result, since the signal is assumed to be periodic. Moreover, the explicit form of g' may not be available. Indeed, one may wish to use a discrete-time window sequence that is not simply the discretized version of a continuous-time window function, in which case it does not make sense to talk about the continuous-time derivative g' . Examples of such window sequences include the minimal-latency windows used by Su and Wu for real-time SST.¹⁶ These considerations lead us to suggest a reassignment frequency formula based on finite differencing of the phase spectrum in Sec. 4.3. We also address the issue of reconstruction in Sec. 4.7.

3. SST-QSTFT: CONTINUOUS SETTING

Since the STFT and SST-STFT only permit a single window choice, these transforms are limited in their capability to adapt to signals whose behavior changes depending on the time-frequency region. Due to the Fourier uncertainty principle, the window cannot have both good time resolution and good frequency resolution. Moreover, while the SST-STFT is designed to sharpen frequency resolution, it does not improve time resolution, and the reassignment frequency is generally less accurate for signals with fast-varying instantaneous frequencies. Hence we consider the notion of *adaptive* time-frequency transforms defined from a family of window functions, where different windows are used for time-frequency regions containing different phenomena. We call such a family a *quilted window family*, borrowing the notion of a time-frequency “quilt” from Dörfler, who coined the term in the context of her work on *quilted Gabor frames*.⁸ Here, different “patches” of the quilt refer to time-frequency regions with different signal behavior. We define a modified version of the STFT for the case of a quilted window family. Next, we define the SST based on this “quilted” STFT. We then define the notion of an *adaptive quilted window family* for a signal f of the class $\mathcal{B}_{\epsilon,d}$. Then we state a theorem for the theoretical accuracy of the reassignment frequency and mode reconstruction for the SST based on quilted STFT, in the case when an adaptive quilted window family is used. Due to the page limitation, we leave the proof of this theorem for our future work.

3.1 Continuous quilted window families

DEFINITION 3.1. We define a *continuous quilted window family* to be a collection $\{h_{t,\xi}\}_{(t,\xi) \in \mathbb{R} \times \mathbb{R}^+}$ where for each $(t,\xi) \in \mathbb{R} \times \mathbb{R}^+$, $h_{t,\xi} \in L^2(\mathbb{R})$ is a window function. Hence, we associate a window function to each time-frequency point.

As a simple example of a continuous quilted window family, we consider first a collection consisting of two Blackman windows $\{g^1, g^2\}$, where g^1 is a wide Blackman window and g^2 is a narrow Blackman window. The wider window enables sharp resolution along the frequency axis, while the narrower window provides sharp resolution along the time axis. In Figure 1, we showed an example of such a collection, together with the different magnitude spectra that would be produced when computing the STFT with respect to each window.

Strictly speaking, the collection $\{g^1, g^2\}$ by itself does not define a continuous quilted window family, since we have not yet associated a window to each time-frequency point of $\mathbb{R} \times \mathbb{R}^+$. In general, one should associate time-frequency points to window functions in an automatic fashion based on the signal content, and we describe algorithms in Sec. 5.1 for this purpose. Heuristically, one may associate a narrow window to time-frequency regions containing either transient events resembling delta spikes or onsets of amplitude-phase components, in order to better resolve the onset times of these events. Similarly, one may associate a wide window to regions containing amplitude-phase components enduring over a long time period, in order to better resolve the instantaneous frequency information. Figure 1 demonstrated the capacity of the collection $\{g^1, g^2\}$ to sharply resolve these different categories of signal content for a glockenspiel signal.

The usage of a more general, larger window family consisting of several dilations of a single Gaussian window to compute SST was explored in depth by Sheu et al.¹² The dilation of the Gaussian window corresponds to varying its effective bandwidth, enabling for sharp representation of a greater variety of time-frequency events. As another example, we consider a *chirped window family* of the form $\{g^\sigma\}_{\sigma \in \mathcal{S}}$, where

$$g^\sigma(t) := g(t)e^{2\pi i \sigma t^2/2}$$

for each σ in a finite set of real numbers \mathfrak{S} , with g a fixed window function. Here, the parameter σ is called the *chirp parameter* or *chirp rate*. The STFT V_{g^σ} will then sharply concentrate around instantaneous frequency curve segments that can be closely approximated by $\phi'(t) \approx \sigma t + c$ for some $c \in \mathbb{R}$. Families of chirped windows provide an alternative to varying the window width. Moreover, by determining which parameter $\sigma^* \in \mathfrak{S}$ yields the STFT $V_{g^{\sigma^*}}$ that best concentrates the signal content in a small region, one may directly infer that the signal content in that region is better linearized by a component of the form $\sigma^* t + c$ than $\sigma t + c$ for any other $\sigma \in \mathfrak{S}$. Thus, chirped windows allow us an immediate estimate of the signal information that does not follow directly from a simple window dilation. We will further explore the use of chirped window families in Sec. 5.

Using the definition of continuous quilted window family, one may consider a tiling of the time-frequency plane where each tile $\mathcal{T} \subseteq \mathbb{R}^2$ has a window $h^\mathcal{T}$ associated to it; i.e., $h_{t,\xi} = h^\mathcal{T}$ for each $(t,\xi) \in \mathcal{T}$. Thus, it is possible to consider an optimization problem where the window $h^\mathcal{T}$ is chosen to adapt to the signal behavior in \mathcal{T} . Indeed, such adaptivity will become necessary when considering the accuracy of an SST in the quilted window context, as we discuss in Secs. 3.4 and 3.5.

3.2 Continuous quilted short-time Fourier transform (QSTFT)

DEFINITION 3.2. Suppose $\{h_{t,\xi}\}_{(t,\xi) \in \mathbb{R} \times \mathbb{R}^+}$ is a continuous quilted window family, and define the function h by $h(x,t,\xi) := h_{t,\xi}(x)$ for each $t, x \in \mathbb{R}$ and $\xi \in \mathbb{R}^+$. Then we define the *quilted short-time Fourier transform*¹¹ of the signal f with respect to h by

$$\mathcal{V}_h^Q f(t, \xi) := \int_{\mathbb{R}} f(x) \overline{h_{t,\xi}(x-t)} e^{-2\pi i \xi(x-t)} dx. \quad (4)$$

3.3 Continuous QSTFT-based SST (SST-QSTFT)

DEFINITION 3.3. We define the *continuous QSTFT-based synchrosqueezing transform*¹¹ (continuous SST-QSTFT) of a signal f , with respect to the function h defining the continuous quilted window family $\{h_{t,\xi}\}_{(t,\xi) \in \mathbb{R} \times \mathbb{R}^+}$, and with tolerance $\gamma \geq 0$ and limiting parameter $\beta > 0$, as follows:

$$\mathcal{S}_{h,\gamma}^{Q,\beta} f(t, \xi) := \int_{\mathcal{A}_{h,\gamma}^Q f(t)} \mathcal{V}_h^Q f(t, \eta) \frac{1}{\beta} b \left(\frac{\xi - \Xi_h^Q f(t, \eta)}{\beta} \right) d\eta,$$

with b as before, $\mathcal{A}_{h,\gamma}^Q f(t) := \{\eta \in \mathbb{R}^+ : |\mathcal{V}_h^Q f(t, \eta)| > \gamma\}$, and where $\Xi_h^Q f(t, \xi) := \frac{\partial_t \mathcal{V}_h^Q f(t, \xi)}{2\pi i \mathcal{V}_h^Q f(t, \xi)}$ is the *QSTFT-based reassignment frequency*.

3.4 Adaptive continuous quilted window function families

In order to ensure the accuracy of reassignment frequency and mode reconstruction for the SST-QSTFT, one must make some assumptions on the quilted window family. In particular, one must assume that the family is *adaptive* to the signal in the following manner.

DEFINITION 3.4. Suppose that $f \in \mathcal{B}_{\epsilon,d}$. We say that the continuous quilted window family $\{h_{t,\xi}\}_{(t,\xi) \in \mathbb{R} \times \mathbb{R}^+}$ is of class $\mathcal{W}_{d,\epsilon,f}^Q$ if the following conditions hold:

- **Smoothness, non-triviality, and band-limitation:** For each $(t, \xi) \in \mathbb{R} \times \mathbb{R}^+$ we have $h_{t,\xi} \in S(\mathbb{R})$, $h_{t,\xi}(0) \neq 0$, and $\text{supp} \left(\widehat{h_{t,\xi}} \right) \subset [-d/2, d/2]$.
- **Window choice remains constant in the frequency band around an IF value:** For each $t \in \mathbb{R}$ and $m \in \{1, \dots, M\}$, there exists a single window function $g_{t,m}$ such that $h_{t,\xi} \equiv g_{t,m}$ for all ξ in the frequency band $\{\xi : |\phi'_m(t) - \xi| < d/2\}$.




Figure 2. SST-QSTFT of glockenspiel signal with respect to $\{\tilde{g}^1, \tilde{g}^2\}$ and with tolerance $\gamma = 0$, using Algorithm 2 for automatic adaptive window selection (see Sec. 5.1), with SST applied only where \tilde{g}^1 is used, and frequency range between 2000 and 11000 Hz. \tilde{g}^1 and \tilde{g}^2 are normalized to have unit ℓ^1 -norm. Colors denote the window chosen in each time-frequency region (blue for \tilde{g}^1 and red for \tilde{g}^2). For improved energy concentration, instead of reassigning the original QSTFT coefficients, we reassign their squared magnitudes (See Secs. 4.4 and 4.7). We also apply a threshold to the SST-QSTFT, to give a cleaner visualization in low-energy regions where the automatic window choice algorithm becomes unreliable. The hop size is $H = 128$ samples, and the parameters used for Algorithm 2 are $A^* = B^* = 5$, $\alpha = 0.25$, $T_{\text{shift}} = Y_{\text{shift}} = 3$, $T_{\text{step}} = Y_{\text{step}} = 1$, with $P_T = P_Y = \{-1, 1\}$.

- **Integration bounds:** For each $p \in \{0, 1, 2\}$, there exists $I_p \in \mathbb{R}^+$ such that

$$\sup_{(t, \xi) \in \mathbb{R} \times \mathbb{R}^+} \int_{\mathbb{R}} |u|^p |h_{t, \xi}(u)| \, du \leq I_p. \quad (5a)$$

- **The window family does not change too quickly over time:** Defining $h(x; t, \xi) := h_{t, \xi}(x)$, we have

$$\sup_{(t, \xi) \in \mathbb{R} \times \mathbb{R}^+} \int_{\mathbb{R}} |\partial_t h(u; t, \xi)| \, du \leq \epsilon \text{ and } \sup_{(t, \xi) \in \mathbb{R} \times \mathbb{R}^+} \int_{\mathbb{R}} |u| |\partial_t h(u; t, \xi)| \, du \leq J_1 \text{ for some } J_1 \in \mathbb{R}^+. \quad (5b)$$

Moreover, we call $\mathcal{W}_{d, \epsilon, f}^Q$ the *class of f -adaptive continuous quilted window families*.

REMARK 3.5. Of course, one does not generally know the information ϕ'_m in advance. Hence, an algorithm must be prescribed to automatically select the windows in a signal-adaptive manner. We describe algorithms for this purpose in Sec. 5.1. Figure 2 demonstrates the usage of an automatic adaptive window selection algorithm to compute the SST-QSTFT, using a collection of Blackman windows $\{\tilde{g}^1, \tilde{g}^2\}$, where \tilde{g}^1 is of length 2048 samples and \tilde{g}^2 is of length 512 samples. In this figure, transient note actions are well-concentrated, owing to the better time resolution of \tilde{g}^2 in transient note regions.

3.5 Theorem

We state a theoretical result analogous to that of Theorem 2.8 for the continuous SST-QSTFT here:

THEOREM 3.6. Let $\epsilon > 0$, $\nu \in (0, 1/2)$, $\tilde{\epsilon} := \epsilon^\nu$, $d > 0$. Suppose that $f = \sum_{m=1}^M f_m \in \mathcal{B}_{\epsilon, d}$. Assume that $\{h_{t, \xi}\}_{(t, \xi) \in \mathbb{R} \times \mathbb{R}^+}$ is of the class $\mathcal{W}_{d, \epsilon, f}^Q$. Then, if ϵ is sufficiently small we have:

- **(Concentration of QSTFT around IF curves)** $|\mathcal{V}_h^Q f(t, \xi)| > \tilde{\epsilon}$ only when there is an $m \in \{1, \dots, M\}$ such that $(t, \xi) \in \mathcal{Z}_m := \{(t, \xi) \in \mathbb{R} \times \mathbb{R}^+ : |\phi'_m(t) - \xi| < d/2\}$.
- **(Closeness of reassignment frequency $\Xi_h^Q f$ to nearby IF)** For all $m \in \{1, \dots, M\}$ and all $(t, \xi) \in \mathcal{Z}_m$ such that $|\mathcal{V}_h^Q f(t, \xi)| > \tilde{\epsilon}$, we have $|\Xi_h^Q f(t, \xi) - \phi'_m(t)| \leq \tilde{\epsilon}$.
- **(Accuracy of reconstruction)** For each $m \in \{1, \dots, M\}$ there is a constant $C_m > 0$ where for any $t \in \mathbb{R}$,

$$\left| \lim_{\beta \rightarrow 0^+} \left(\int_{\{\xi : |\xi - \phi'_m(t)| < \tilde{\epsilon}\}} \frac{1}{h_{t, \xi}(0)} \cdot \mathcal{S}_{h, \tilde{\epsilon}}^{Q, \beta} f(t, \xi) \, d\xi \right) - f_m(t) \right| \leq C_m \tilde{\epsilon}.$$

We will give a proof of this theorem in future work.

4. SST-QSTFT: DISCRETE SETTING

The discrete theory and implementation of SST-QSTFT differs depending on whether one uses a window that is band-limited, time-limited, or neither band-limited nor time-limited. While the continuous theory prescribes the use of a band-limited window, it has nonetheless been proven¹⁷ that the SST is still accurate for windows that are almost band-limited (i.e., having Fourier transform almost zero for all frequencies outside a certain passband). Moreover, for many signal processing applications, it is more common to use time-limited windows such as the well-known Hanning, Hamming, Blackman, and Kaiser windows. Indeed, it is generally not practical to use band-limited windows for real-time applications, since the implementation requires computing the DFT of the entire signal (which is not available). Still, the usage of band-limited windows may enable sharper frequency resolution than is possible in the case of time-limited windows, which all exhibit some sort of frequency roll-off. Moreover, certain windows may be neither band-limited nor time-limited, but are still of practical interest, such as the Gaussian window. In practice, one truncates such windows and implements them in either the same way as a band-limited window or as a time-limited window.

For the sake of brevity, we will focus on the most common case of time-limited windows, which enables extensions to the recently developed real-time SST theory.^{16,18,19}

4.1 Discrete quilted windows

NOTATION. For $L \in \mathbb{N}$, we define $\mathbb{Z}_L := \{0, \dots, L-1\}$. This notation will help to simplify the following definitions.

DEFINITION 4.1. We define a *discrete quilted window family* for a given hop-size $H \in \mathbb{N}$, discrete-time signal of interest f , and maximal window length $L \in \mathbb{N}$, to be a tuple $\{(h_{n,k}, L_{n,k})\}_{[n,k] \in \mathbb{Z}_N \times \mathbb{Z}_L}$, where for each $n \in \mathbb{Z}_N$ and $k \in \mathbb{Z}_L$,

- $h_{n,k}$ is a discrete sequence of length $L_{n,k} \in \mathbb{N}$ defining a discrete window function,
- the $L_{n,k}$ are chosen such that $L = \max_{[n,k] \in \mathbb{Z}_N \times \mathbb{Z}_L} L_{n,k}$, and
- $N \in \mathbb{N}$ is large enough so that $f[(L-1) + (N-1)H] = 0$.

Hence, we associate to each time-frequency point $[n, k] \in \mathbb{Z}_N \times \mathbb{Z}_L$ a window function $h_{n,k}$ of length $L_{n,k}$.

REMARK 4.2. It is well known that windows of different lengths have different time-frequency resolution. On the other hand, one often extends windows by zeros along their boundary to a length of 2^P for $P \in \mathbb{N}$ (a practice commonly known as *zero-padding*) for the purpose of fast computation using the Fast Fourier Transform (FFT). For our purposes, one should interpret the parameter $L_{n,k}$ that determines the length of each window in the discrete quilted window family as *the window length after zero-padding*.

REMARK 4.3. As in the continuous case, we do not generally choose an entirely different window for each time-frequency point, opting instead to divide the discrete time-frequency plane into tiles $T \subset \mathbb{Z}_N \times \mathbb{Z}_L$, with a window h^T associated to each T ; that is, where $h_{n,k} = h^T$ for each $[n, k] \in T$. The constructions described in Sec. 5.1 proceed in such a manner.

4.2 Discrete quilted short-time Fourier transform (QSTFT)

DEFINITION 4.4. In analogy with the manner in which we defined the discrete STFT in Eqn. (2), we may now define the *discrete quilted short-time Fourier transform* of f with hop size $H \in \mathbb{N}$, for a discrete quilted window family $\{(h_{n,k}, L_{n,k})\}_{[n,k] \in \mathbb{Z}_N \times \mathbb{Z}_L}$, via

$$V_h^Q f[n, k] := \sum_{\ell=0}^{L_{n,k}-1} f[\ell + nH] \overline{h_{n,k}[\ell]} e^{-2\pi i k \ell / L_{n,k}}, \quad (6)$$

for each frame $n \in \mathbb{Z}_N$, and each frequency bin $k \in \mathbb{Z}_L$, where we define the discrete function h by $h[\ell; n, k] := h_{n,k}[\ell]$ for each $n \in \mathbb{Z}_N$ and $\ell, k \in \mathbb{Z}_L$.

4.3 Discrete reassignment frequency formula

In the continuous context, the reassignment frequency for SST-QSTFT requires the computation of the continuous-time derivative of the QSTFT $\mathcal{V}_h^Q f(t, \xi)$. Since one cannot directly compute $\partial_t \mathcal{V}_h^Q f(t, \xi)$ in the discrete context, one option is to observe that $\partial_t \mathcal{V}_h^Q f(t, \xi)$ can be expressed as the QSTFT of f with respect to the window function $(2\pi i \xi h_{t,\xi}(\cdot - t) - \partial_t(h_{t,\xi})(\cdot - t))$. This is the method used by Thakur and Wu for the theory of the discrete SST-STFT.³ However, as observed in Sec. 2.4.2, $\partial_t(h_{t,\xi})$ may be nonzero at the boundary of the support of $h_{t,\xi}$, which violates the periodicity requirement for taking the DFT and may cause aliasing. Alternatively, we may wish to use $h_{n,k}$ that does not arise from any continuous-time function h , in which case such a quantity $\partial_t(h_{t,\xi})$ is unavailable.

Instead, one may consider for instance a first-order finite difference approximation $D_{\Delta t}$ to the derivative ∂_t . However, using this approximation together with the continuous reassignment frequency formula will yield an inaccurate result in general, which we show in the following. First we put forth the following notation.

NOTATION. Suppose $s(\cdot)$ is a continuous-time signal defined on a set $\mathcal{R} \subset \mathbb{R}$, and we define the discretized version of s by $s[n] := s(n\Delta t)$ for all $n \in \mathbb{Z} \cap \{n : \exists r \in \mathcal{R}, n = r/\Delta t\}$. Then we define the *backward-shifted discrete signal* s^+ by $s^+[n] := s((n+1)\Delta t)$.

The notation above becomes important because in general we may deal with a QSTFT having hopsize $H > 1$, and hence more coarsely sampled than the original signal. Now consider the approximation of the continuous reassignment frequency using the forward difference $(V_h^Q f^+ - V_h^Q f)/(\Delta t)$, for a complex-valued constant chirp of the form $f(t) = Ae^{2\pi i c t}$, where $A, c \in \mathbb{R}^+$:

$$\begin{aligned} \frac{\partial_t \mathcal{V}_h^Q f(nH\Delta t, k/L_{n,k})}{2\pi i \mathcal{V}_h^Q f(nH\Delta t, k/L_{n,k})} &\approx \frac{V_h^Q f^+[n, k] - V_h^Q f[n, k]}{2\pi i \Delta t V_h^Q f[n, k]} \\ &= \frac{\sum_{\ell=0}^{(L_{n,k}-1)} A (e^{2\pi i c(\ell+nH+1)\Delta t} - e^{2\pi i c(\ell+nH)\Delta t}) \overline{h_{n,k}[\ell]} e^{-2\pi i k\ell/L_{n,k}}}{2\pi i \Delta t \sum_{\ell=0}^{(L_{n,k}-1)} A e^{2\pi i c(\ell+nH)\Delta t} \overline{h_{n,k}[\ell]} e^{-2\pi i k\ell/L_{n,k}}} \\ &= \frac{e^{2\pi i c \Delta t} - 1}{2\pi i \Delta t} = c \cdot e^{\pi i c \Delta t} \text{sinc}(\pi c \Delta t). \end{aligned}$$

Hence, this forward difference does not isolate the IF c in this simple case. Instead, one may consider using the alternative continuous reassignment frequency proposed by Oberlin et al.⁴ where one instead computes the time derivative of the *phase spectrum* $\arg(\mathcal{V}_h^Q f(t, \xi))$. Then, we see that the forward difference approximation exactly retrieves the IF c for the complex-valued constant chirp f , provided that c is below the Nyquist frequency $1/(2\Delta t)$:

$$\begin{aligned} \frac{1}{2\pi} \partial_t \arg(\mathcal{V}_h^Q f(nH\Delta t, k/L_{n,k})) &\approx \frac{\arg(V_h^Q f^+[n, k]) - \arg(V_h^Q f[n, k])}{2\pi \Delta t} \\ &= \frac{1}{2\pi \Delta t} \arg \left(\frac{\sum_{\ell=0}^{(L_{n,k}-1)} A e^{2\pi i c(\ell+nH+1)\Delta t} \overline{h_{n,k}[\ell]} e^{-2\pi i k\ell/L_{n,k}}}{\sum_{\ell=0}^{(L_{n,k}-1)} A e^{2\pi i c(\ell+nH)\Delta t} \overline{h_{n,k}[\ell]} e^{-2\pi i k\ell/L_{n,k}}} \right) = \frac{\arg(e^{2\pi i c \Delta t})}{2\pi \Delta t} = c. \end{aligned}$$

Therefore, defining the *discrete SST-QSTFT reassignment frequency*

$$\tilde{\Xi}_h^Q f[n, k] := \frac{1}{2\pi \Delta t} \arg \left(\frac{V_h^Q f^+[n, k]}{V_h^Q f[n, k]} \right) \quad (7)$$

for each frame $n \in \mathbb{Z}_N$ and frequency bin $k \in \mathbb{Z}_L$ where $V_h^Q f[n, k] \neq 0$, we see that $\tilde{\Xi}_h^Q f[n, k]$ is *exact* for the complex-valued constant chirp of frequency $c \in (0, \frac{1}{2\Delta t})$ considered. Moreover, the derivation of the continuous formula for the reassignment frequency is motivated by the fact that it is exact for a constant chirp.^{1,2} Hence, we use the discrete reassignment formula given by (7) in the following definition of the discrete version of SST-QSTFT.

4.4 Discrete QSTFT-based SST (SST-QSTFT)

DEFINITION 4.5. We define the *discrete QSTFT-based synchrosqueezing transform* (discrete SST-QSTFT) of a discrete-time signal f , with respect to the function h defining the discrete quilted window sequences $\{h_{n,\ell}\}_{[n,\ell] \in \mathbb{Z}_N \times \mathbb{Z}_L}$, and with tolerance $\gamma \geq 0$, as follows:

$$S_{h,\gamma}^Q f[n,k] := \sum_{\ell \in A_{h,\gamma}^Q f[n]} V_h^Q f[n,\ell] \mathbf{1}_{B_{h,k}^Q f[n]}[\ell],$$

for each frame $n \in \mathbb{Z}_N$ and for each *reassignment frequency bin* $k \in \mathbb{Z}_K$ with $K \geq L$, where $A_{h,\gamma}^Q f[n] := \{\ell \in \mathbb{Z}_L : |V_h^Q f[n,\ell]| > \gamma\}$, $B_{h,k}^Q f[n] := \{\ell \in \mathbb{Z}_L : -\frac{1}{2} \leq K\Delta t \tilde{\Xi}_h^Q f[n,\ell] - k < \frac{1}{2}\}$ is the set of QSTFT frequency bins ℓ at the frame n where the corresponding reassignment frequency is closer to k than any other reassignment frequency bin, $\mathbf{1}_X$ is the characteristic function on the set X (i.e., $\mathbf{1}_X(x) = 1$ if $x \in X$ and $\mathbf{1}_X(x) = 0$ otherwise), and $\tilde{\Xi}_h^Q f$ as defined in (7) is the *discrete QSTFT-based reassignment frequency*.

REMARK 4.6. Here, we do not use any function b to do an approximate summation, as in the case of the continuous SST-QSTFT. The use of b in the continuous case is for the sake of the proof of reconstruction, and in the following we do not provide a direct reconstruction formula from the discrete SST-QSTFT coefficients, passing back instead to the original QSTFT coefficients. See Sec. 4.7 for more details.

REMARK 4.7. Note that the SST-QSTFT may be computed over more frequency bins K than the QSTFT has (L), enabling the possibility of an even more precise estimation of IF. Indeed, since the discrete reassignment frequency formula is exact in the case of a constant chirp *regardless of whether the IF coincides with a frequency bin*, it follows that the use of $K > L$ may well yield a more precise concentration around the true IF.

DEFINITION 4.8. Alternatively, we may also calculate the *discrete QSTFT-based magnitude synchrosqueezing transform* (discrete magnitude SST-QSTFT), where we reassign the *magnitude-squared* of the QSTFT coefficients, by

$$S_{h,\gamma}^{|Q|^2} f[n,k] := \sum_{\ell \in A_{h,\gamma}^Q f[n]} \left| V_h^Q f[n,\ell] \right|^2 \mathbf{1}_{B_{h,k}^Q f[n]}[\ell].$$

REMARK 4.9. Calculating $S_{h,\gamma}^{|Q|^2} f$ enables us to preserve the total energy of the QSTFT coefficients, and hence yields higher energy peaks along IF curves than the original QSTFT. Since the original QSTFT coefficients are summed together for $S_{h,\gamma}^Q f$, their total energy might decrease after summation. Both SST-QSTFT representations will be more concentrated than QSTFT, due to the squeezing procedure summing together high-energy coefficients closer to the IF curves, but $S_{h,\gamma}^{|Q|^2} f$ will generally contain more energy along the IF curves than $S_{h,\gamma}^Q f$.

REMARK 4.10. One may set $\gamma > 0$ to ensure an accurate reassignment frequency for all reassigned coefficients. In fact, this is necessary in the continuous case to guarantee reconstruction accuracy, as per Theorem 3.6. However, as stated in Remark 4.6, we do not perform direct reconstruction from the discrete SST-QSTFT coefficients. Moreover, we use $S_{h,\gamma}^{|Q|^2} f$ in all our numerical experiments (for reasons given in Sec. 4.7), and the issue of inaccuracy for reassignment frequency only concerns low-energy coefficients that do not contribute much to $S_{h,\gamma}^{|Q|^2} f$. Hence, we always set $\gamma = 0$ in our experiments.

4.5 Adaptive discrete quilted window function families

For the case of discrete quilted window families, we make the following definition of an adaptive family:

DEFINITION 4.11. Suppose that $f \in \mathcal{B}_{\epsilon,d}$. We say that the quilted window family $\{(h_{n,k}, L_{n,k})\}_{[n,k] \in \mathbb{Z}_N \times \mathbb{Z}_L}$ is of class $W_{d,\epsilon,f}^Q$ if the following conditions hold:

- **Fourier-side decay (almost band-limitation):** There exists a constant $C_h \in \mathbb{R}^+$ such that

$$\sup_{\substack{u \in \mathbb{R} \setminus [-d/2, d/2] \\ [n, k] \in \mathbb{Z}_N \times \mathbb{Z}_L}} \left| \widehat{h_{n,k}}(u) \right| = C_h \epsilon, \quad (8)$$

so that the semi-discrete Fourier transform of $h_{n,k}$ is bounded outside the frequency band $[-d/2, d/2]$.

- **Window choice remains constant in the frequency band around an IF value:** For all $n \in \mathbb{Z}_N$ and $m \in \{1, \dots, M\}$, there exists a single discrete window function $g_{n,m}$ such that $h_{n,k} \equiv g_{n,m}$ for all $k \in \{k : |k/(L_{n,k}\Delta t) - \phi'_m[nH]| < d/2\}$.

Moreover, we call $W_{d,\epsilon,f}^Q$ the *class of f -adaptive discrete quilted window families*.

4.6 Theorem

Now, we state the following theorem for the discrete SST-QSTFT:

THEOREM 4.12. *Let $\epsilon > 0$, $\nu \in (0, 1/2)$, $\tilde{\epsilon} := \epsilon^\nu$, $d > 0$. Suppose that $f = \sum_{m=1}^M f_m \in \mathcal{B}_{\epsilon,d}$, with $\phi'_m[nH] \leq 1/(2\Delta t)$ for all $m \in \{1, \dots, M\}$ and $n \in \mathbb{Z}_N$. Assume that $\{(h_{n,k}, L_{n,k})\}_{[n,k] \in \mathbb{Z}_N \times \mathbb{Z}_L}$ is of the class $W_{d,\epsilon,f}^Q$. Then, if ϵ is sufficiently small we have:*

- **(Concentration of QSTFT around IF curves)** $|V_h^Q f[n, k]| > \tilde{\epsilon}$ only when there is an $m \in \{1, \dots, M\}$ such that $[n, k] \in Z_m := \{[n, k] \in \mathbb{Z}_N \times \mathbb{Z}_L : |\phi'_m[nH] - k/(L_{n,k}\Delta t)| < d/2\}$.
- **(Closeness of reassignment frequency $\tilde{\Xi}_h^Q f$ to nearby IF)** For all $m \in \{1, \dots, M\}$ and all $[n, k] \in Z_m$ such that $|V_h^Q f[n, k]| > \tilde{\epsilon}$, we have $|\tilde{\Xi}_h^Q f[n, k] - \phi'_m[nH]| \leq C_{\tilde{\epsilon}}$, where $\lim_{\tilde{\epsilon} \rightarrow 0^+} C_{\tilde{\epsilon}} = 0$.

The proof of this theorem is left for our future work, due to the page limitation.

4.7 Discrete reconstruction

We do not provide a reconstruction theorem for the discrete case, because the technique of summing over SST-QSTFT reassignment frequency bins for each frame is not accurate in the case of time-limited windows used with hop size $H > 1$. Instead, we refer to the technique of Holighaus et al.,¹⁰ who suggested to store the *inverse reassignment map* defined by

$$(\tilde{\Xi}_h^Q f)^{-1}[n, k] := \{\ell \in \{0, \dots, L-1\} : \tilde{\Xi}_h^Q f[n, \ell] = k\}$$

for $n \in \mathbb{Z}_N$ and $k \in \mathbb{Z}_K$. $(\tilde{\Xi}_h^Q f)^{-1}[n, k]$ is the set of all frequencies $\ell \in \mathbb{Z}_L$ whose reassignment frequency is k . Then, rather than doing reconstruction along the narrow SST IF ridges, the reconstruction can be done over the thicker QSTFT ridges via frame synthesis, using an *overlap-add formula*¹⁴ and the construction of dual windows using canonical tight frames [7, Sec. 3.3, Theorem 1]. Note that for this type of reconstruction, the condition of maintaining the same quilted window over all frequencies in a given IF band (for a fixed frame) is essential.

We note that the usual technique for reconstruction using SST is to: (i) analyze the magnitude-squared of the SST representation for peaks that form continuous ridges; (ii) use an algorithm to extract IF curves following these ridges,^{20–22} and then (iii) employ a discrete version of a reconstruction formula such as the one in Theorem 3.6 to reconstruct from the original SST coefficients.^{2,4} By contrast, the technique above enables us to replace step (iii) with reconstruction from QSTFT coefficients, and we only need the SST for extracting the IF curves. Hence, rather than using $S_{h,\gamma}^Q f$, it makes more sense to use the generally more strongly concentrated $S_{h,\gamma}^{|Q|^2} f$ for analysis (see Remark 4.9). In the following, we will refer to both $S_{h,\gamma}^Q f$ and $S_{h,\gamma}^{|Q|^2} f$ as *SST-QSTFT*, and use only $S_{h,\gamma}^{|Q|^2} f$ in our numerical experiments.

5. NUMERICAL IMPLEMENTATION

In the following section, we describe the numerical implementation of SST-QSTFT. We have developed a Python suite for the adaptive time-frequency transforms described in this work, including SST-QSTFT, available to researchers upon request.[¶]

5.1 Automatic adaptive window selection

The question remains how to adapt the window selection to the signal. In 2007, Jaillet & Torr sani introduced the idea of the *time-frequency jigsaw puzzle*, where different windows could be adaptively associated to different regions (representing “jigsaw pieces”) in the time-frequency plane, based on the time-frequency content.²³ As a starting point, we may incorporate only a single iteration of their algorithm in order to choose an optimal window for each time-frequency region, from a given window collection. Our first algorithm for automatic adaptive window selection proceeds as follows:

ALGORITHM 1.

1. Choose windows h^w , $w = 1, \dots, W$ for $W \in \mathbb{N}$, and corresponding sampling lattices, the latter determined by the hop size H (which must be the same for all windows, by Sec. 4.1) and FFT size L^w .
2. For each h^w , compute the discrete STFT $V_{h^w} f$ with hop size H and FFT size L^w (chosen in Step 1).
3. Fix parameters $A, B \in \mathbb{N}$ to decompose the time-frequency plane \mathbb{R}^2 into “supertiles” $\square_{r,s}$, where $\square_{r,s} := [r\tilde{A}, (r+1)\tilde{A}) \times [s\tilde{B}, (s+1)\tilde{B})$, $\tilde{A} := AH\Delta t$, $\tilde{B} := B/(F_{\min}\Delta t)$, and $F_{\min} := \min_w \{F^w\}$. We note that A (resp. B) represents the number of points on the coarsest grid along the time (resp. frequency) axis contained in each supertile.
4. For each $\square_{r,s}$ and h^w , calculate the total energy $E_{r,s}^w := \sum_{(\ell,k) \in \square_{r,s}} |V_{h^w}^Q f[\ell, k]|^2$.
5. Fix $\alpha \in (0, 1)$. Then for each $\square_{r,s}$ and h^w , calculate the *sampled R nyi entropy*^{23,24} $\mathcal{R}_{r,s}^{\alpha,w}$ of the STFT coefficients whose corresponding lattice points are within $\square_{r,s}$:

$$\mathcal{R}_{r,s}^{\alpha,w} := \frac{1}{1-\alpha} \log \left[\sum_{(\ell,k) \in \square_{r,s}} \left(\left(\frac{H}{L^w} \right)^{1-\alpha} \frac{|V_{h^w}^Q f[\ell, k]|^2}{E_{r,s}^w} \right)^\alpha \right]$$

6. To each supertile $\square_{r,s}$, associate the window h^{w^*} with the smallest sampled R nyi entropy $\mathcal{R}_{r,s}^{\alpha,w^*}$. Thus we adapt to signal content in $\square_{r,s}$.

REMARK 5.1. The R nyi entropy can be seen as a generalization of the ℓ^2 -normalized ℓ^1 -norm, and can hence be seen as a time-frequency sparsity measure. Indeed, with $\alpha = 0.5$, one recovers the ℓ^2 -normalized ℓ^1 -norm. Furthermore, minimization of the ℓ^2 -normalized ℓ^p -norm with $p = 2\alpha$ is equivalent to minimizing the R nyi entropy.²⁵

The left plot in Figure 3 demonstrates the application of Algorithm 1 to compute the SST-QSTFT of a noise-contaminated signal $f(t) = \sum_{m=1}^4 f_m(t) + f_{\text{noise}}(t)$ containing crossing linear chirp components of the form $f_m(t) = \cos(2\pi\phi_m(t))$ with $\phi'_m(t) = c_m + \sigma_m t$, where $(c_m)_{m=1}^4 = (5000, 8000, 12000, 15000)$ and $(\sigma_m)_{m=1}^4 = (2000, 800, -800, -2000)$, and with zero-mean white noise component $f_{\text{noise}}(t)$ such that f has SNR equal to 4.0 dB. Here, SNR for a generic signal $f(t) = y(t) + f_{\text{noise}}(t)$, where y is noiseless, is defined by²

$$\text{SNR}(f) := 10 \log_{10} \left(\frac{\text{var}(y)}{\text{var}(f_{\text{noise}})} \right).$$

[¶]Please email Alex Berrian at aberrian@math.ucdavis.edu for a copy of the Python suite.




Figure 3. **Left:** SST-QSTFT of the crossing chirps signal f using Algorithm 1. **Right:** SST-QSTFT of the same signal using Algorithm 2. In both plots, colors correspond to the window chosen in each supertile for the calculation of the QSTFT. Blue: h^1 . Green: h^2 . Purple: h^3 . Red: h^4 . Frequency range is restricted from 4500 Hz to 15500 Hz. The perturbed supertiles algorithm leads to fewer instances of QSTFT coefficients near IFs from windows of very dissimilar chirp rate.

To compute the SST-QSTFT, we use the chirped window collection $\{h^w\}$ where $h^w(t) := h^0(t)e^{2\pi i\sigma^w t^2/2}$. Here h^0 is a Hanning window of width 4000 samples, zero-padded to $L = 2^{12} = 4096$ samples, and $\{\sigma^w\}_{w=1}^4 = \{1900, 900, -900, -1900\}\text{Hz}^2$. We set $H = 250$, $K = 4L = 16384$, $\gamma = 0$, $\alpha = 0.5$, and $A = B = 24$.

REMARK 5.2. We use this particular set of chirp rates $\{\sigma^w\}$ only because it demonstrates that a reasonable result may be obtained using windows with chirp rates that are close, but not quite equal, to the actual local chirp rates $\{\sigma_m\}$. Often, $\{\sigma^w\}$ may be chosen by visual inspection of a single STFT of the signal. However, in general, the local chirp rates are unknown, and one should have some systematic method of deciding the collection $\{\sigma^w\}$. This topic is out of the scope of this paper, and we leave it for our future work.

In Figure 3, we see that supertiles in the neighborhood of the chirps may be assigned a window with chirp parameter that does not correspond to the slope of the nearby chirp component. This may happen in the case that the supertile contains very little signal content, in which case sharply-concentrated content may be mistaken for noise. There is also the possibility that an IF curve does not pass through a supertile at all, but is analyzed by a window that yields a very blurry representation of the IF, which ends up “leaking” into the supertile. In both of these cases, Algorithm 1 may fail to pick the window that yields the sparsest signal representation in the supertile.

We note that Algorithm 1 is essentially equivalent to a simplified version of the algorithm of Sheu et al.,¹² but with the Rényi entropy written out explicitly for the discrete setting. Sheu et al. allow for overlapping supertiles, and update the optimal entropy choice as they move forward in time and frequency. To improve upon the deficiencies of Algorithm 1, we go with a more general approach based on *perturbed supertiles*. In Figure 4, we plot a single supertile $\square_{r,s}$ together with eight perturbations of $\square_{r,s}$ along the time and frequency directions. The basic idea of Algorithm 2 is to calculate the Rényi entropy in the original supertile $\square_{r,s}$ as well as in several perturbations $\square_{r,s}^{\tilde{t},\tilde{y}}$ of $\square_{r,s}$, and then to compute the average $\tilde{\mathcal{R}}_{r,s}^{\alpha,w}$ of all the calculated entropies, for all windows (indexed by w). Then, the window that minimizes this averaged entropy $\tilde{\mathcal{R}}_{r,s}^{\alpha,w}$ is the optimal window for the supertile $\square_{r,s}$. We give a rigorous description of our algorithm as follows:

ALGORITHM 2.

1–3. Do Steps 1-3 of Algorithm 1.




Figure 4. All perturbations of the supertile $\square_{r,s}$, for Algorithm 2 parameters given as follows: $P_T = P_Y = \{-1, 1\}$ (backward and forward directional perturbation in both time and frequency), $T_{\text{shift}} = Y_{\text{shift}} = 2$ (two perturbations per direction in both time and frequency), $\tau = T_{\text{step}} \cdot (H\Delta t)$, and $W = Y_{\text{step}} / (F_{\text{min}}\Delta t)$ for arbitrary Δt , H , F_{min} , T_{step} , and Y_{step} . The label $\square_{r,s}^{\tilde{t},\tilde{y}}$ of each perturbed supertile is shown in the upper-left corner of that supertile.

4. Fix permutation step sizes T_{step} and Y_{step} , the total amounts of permutations T_{shift} and Y_{shift} , and the sets P_T and P_Y of permutation directions satisfying $P_T, P_Y \subseteq \{-1, 1\}$, for the time (T) and frequency (Y) axes respectively. We explain these parameters as follows:

- T_{step} (resp. Y_{step}) is given in units of points on the coarsest grid in the time (resp. frequency) direction. For example, suppose one sets $Y_{\text{step}} = 3$. Then, since the distance between two points on the coarsest frequency grid is $1/(F_{\text{min}}\Delta t)$ Hz, one permutes the supertile with a step size of $W := 3/(F_{\text{min}}\Delta t)$ Hz. If one sets $Y_{\text{shift}} = 2$, then the supertile is permuted two separate times in the frequency direction, each time by W Hz (see Figure 4).
- The possible forms of P_T and P_Y are $\{-1\}$, $\{1\}$, and $\{-1, 1\}$. The first option means that the permutations of the supertile grid will be in the backward direction. The second option refers to the forward direction, and the last option refers to permutation in both directions.

5. Define $\tilde{T} := \{p_T \cdot t \cdot T_{\text{step}} \cdot (H\Delta t) : p_T \in P_T; t = 1, \dots, T_{\text{shift}}\}$, and $\tilde{Y} := \{p_Y \cdot y \cdot Y_{\text{step}} / (F_{\text{min}}\Delta t) : p_Y \in P_Y; y = 1, \dots, Y_{\text{shift}}\}$. Then:

5a. For each $\tilde{t} \in \tilde{T}$, define $\square_{r,s}^{\tilde{t},0} := [r\tilde{A} + \tilde{t}, (r+1)\tilde{A} + \tilde{t}] \times [s\tilde{B}, (s+1)\tilde{B}]$.

5b. For each $\tilde{y} \in \tilde{Y}$, define $\square_{r,s}^{0,\tilde{y}} := [r\tilde{A}, (r+1)\tilde{A}] \times [s\tilde{B} + \tilde{y}, (s+1)\tilde{B} + \tilde{y}]$.

6. Now, fix the Rényi entropy parameter $\alpha \in (0, 1]$ and do the following for each pair (r, s) :

6a. For each $\square \in \mathfrak{B}_{r,s} := \{\square_{r,s}, \square_{r,s}^{\tilde{t},0}, \square_{r,s}^{0,\tilde{y}} : \tilde{t} \in \tilde{T}; \tilde{y} \in \tilde{Y}\}$ and h^w , calculate the total energy

$$E_{r,s}^{w,\square} := \sum_{(\ell,k) \in \square} |V_{h^w}^Q f[\ell, k]|^2.$$

6b. For each $\square \in \mathfrak{B}_{r,s}$ and h^w , calculate the sampled Rényi entropy $\mathcal{R}_{r,s}^{\alpha,w,\square}$ of the STFT coefficients

	1st-order SST-STFT			2nd-order SST-STFT			RM			SST-QSTFT		
	6 dB	0 dB	-6 dB	6 dB	0 dB	-6 dB	6 dB	0 dB	-6 dB	6 dB	0 dB	-6 dB
ϕ_1^l	3.3%	2.6%	1.5%	9.7%	5.4%	2.1%	15.4%	10.8%	4.6%	16.9%	13.0%	6.3%
ϕ_2^l	10.5%	7.9%	4.0%	19.8%	13.7%	6.0%	21.8%	16.1%	7.6%	20.6%	15.7%	8.1%
ϕ_3^l	10.6%	8.4%	4.4%	20.2%	14.2%	6.5%	22.2%	16.7%	8.2%	21.0%	16.2%	8.9%
ϕ_4^l	3.3%	2.6%	1.4%	9.9%	5.3%	2.0%	15.3%	10.6%	4.5%	17.1%	12.8%	6.1%

Table 1. Comparison of percent total energy along IF ridges relative to the total energy between 5 and 15 kHz over all times, for time-frequency representations 1st-order SST, 2nd-order SST, RM, and SST-QSTFT, and SNRs of 6 dB, 0 dB, and -6 dB. Specifically, we compute $E_C(R_m^q(C, f))/E_C(\mathbb{Z}_N \times \mathfrak{W})$ for each time-frequency representation C , where $E_C(R)$ denotes $\sum_{[n,k] \in R} C[n, k]$ over the given region R , f is the crossing chirps signal, and \mathfrak{W} denotes bins between 5 and 15 kHz.

whose corresponding lattice points are within \square :

$$\mathcal{R}_{r,s}^{\alpha,w,\square} := \frac{1}{1-\alpha} \log \left[\sum_{(\ell,k) \in \square} \left(\left(\frac{H}{L^w} \right)^{1-\alpha} \frac{|V_{h^w}^Q f[\ell, k]|^2}{E_{r,s}^{w,\square}} \right)^\alpha \right]$$

6c. For each h^w , calculate the averaged entropy measure

$$\tilde{\mathcal{R}}_{r,s}^{\alpha,w} := \frac{1}{1 + |\tilde{T}| + |\tilde{Y}|} \sum_{\square \in \mathfrak{B}_{r,s}} \mathcal{R}_{r,s}^{\alpha,w,\square}.$$

7. To each supertile $\square_{r,s}$, associate the window h^{w^*} with the smallest averaged entropy $\tilde{\mathcal{R}}_{r,s}^{\alpha,w^*}$. Thus we adapt to signal content in $\square_{r,s}$.

REMARK 5.3. Algorithm 2 is significantly different from the overlapping supertiles algorithm of Sheu et al., since their algorithm only proceeds forward in time and frequency, and does not incorporate an averaging procedure.

REMARK 5.4. Algorithm 2 is similar to a method used by Coifman and Donoho in the context of wavelet-based signal denoising, called “cycle-spinning,” that involved averaging a certain quantity over grid perturbations.²⁶

The intended effect of averaging the entropy over perturbations of the given supertile is to avoid the errors that result from supertiles containing minimal signal content and from blurred signal representation leaking into surrounding supertiles. In the right plot of Figure 3, we used Algorithm 2 with $\alpha = 0.5$, $A = B = 24$, $T_{\text{step}} = Y_{\text{step}} = T_{\text{shift}} = Y_{\text{shift}} = 4$, and $P_T = P_Y = \{-1, 1\}$. We see that Algorithm 2 improves the performance of SST-QSTFT, leading to far fewer blurry patches along the IF curves. This sharper representation results from the more frequent selection of windows with chirp parameter closest to the nearest IF component’s chirp rate.

REMARK 5.5. We note that the selection of the parameters T_{step} , Y_{step} , T_{shift} , and Y_{shift} for Figure 3 (and later on for Figures 5 and 6) is somewhat arbitrary. In general, some tuning of these parameters is necessary to achieve a desired result. We do not address optimal strategies for tuning these parameters in this paper, leaving this topic for our future work. However, we generally recommend using $P_T = P_Y = \{-1, 1\}$, so that the entropy averaging calculation is not biased either forward or backward in time or frequency.

5.2 Numerical results

5.2.1 Comparison with SST-STFT, second-order SST, and reassignment method

One way to measure the sharpness of a time-frequency representation for a given signal is to calculate the amount of energy along the *ridges* of the IF components. Given a signal $f = \sum_{m=1}^M f_m \in \mathcal{B}_{\epsilon,d}$ and a discrete time-frequency representation $C(f)$ (for instance, STFT or SST-STFT) operating on f with N frames, K frequency bins, and hop size H , we define the m^{th} ridge ($m = 1, \dots, M$) of bandwidth q bins to be $R_m^q(C, f) :=$

$\{C(f)[n, k] : -\frac{1}{2} - q \leq K\Delta t\phi'_m[nH] - k < \frac{1}{2} + q\}$; i.e., all the coefficients $C(f)[n, k]$ within q frequency bins from the closest frequency bin to $\phi'_m[nH]$.

Using the crossing chirps signal, we compare the ridge concentration performance of SST-QSTFT with SST-STFT, *second-order SST-STFT*,²⁷ and the *reassignment method* (RM).^{28,29} The central idea of second-order SST-STFT is to use an improved, “second-order accurate” reassignment frequency formula that exactly retrieves the IF in the case of a noiseless single-component signal with IF of the form $\phi'(t) = \sigma t + c$ for real-valued constants σ and c . Second-order SST-STFT requires the side computation of a *reassignment time* quantity, but still only reassigns the frequency locations of the STFT coefficients. By contrast, RM uses reassignment times together with reassignment frequencies to reassign *both* the time and frequency locations of the STFT coefficients.

In order to compute second-order SST-STFT and RM, we introduce the *discrete reassignment time formula*

$$T_g f[n, k] := nH\Delta t - \frac{1}{2\pi} \arg \left(\frac{V_g^+ f[n, k]}{V_g f[n, k]} \right), \quad (9)$$

where $V_g^+ f[n, k] := \sum_{\ell=0}^{L-1} f[\ell + nH] \overline{g[\ell]} e^{-2\pi i \ell (k/L + \Delta t)}$ is the STFT of f with normalized frequency shifted ahead by Δt . It is easy to show that for a single spike $f(t) = A \cdot \mathbf{1}_{\{t=\ell_0\Delta t\}}(t)$ with constant amplitude $A > 0$ located at $t = \ell_0\Delta t$, $T_g f[n, k] \equiv \ell_0\Delta t$ whenever $nH - \ell_0 \in \mathbb{Z}_L$. This formula is a discretization of the continuous reassignment time formula given by $\tau_g f(t, \xi) := t - \frac{1}{2\pi} \partial_\xi \arg(\mathcal{V}_g f(t, \xi))$.²⁷ For second-order SST-STFT specifically, we introduce the *discrete second-order reassignment frequency formula* given by

$$\tilde{\Xi}_g^{(2)} f[n, k] := \tilde{\Xi}_g f[n, k] + \frac{D_t \tilde{\Xi}_g f[n, k]}{D_t T_g f[n, k]} ((nH + [0.5L])\Delta t - T_g f[n, k]), \quad (10)$$

where $D_t \tilde{\Xi}_g f[n, k] := \frac{1}{2\Delta t} (\tilde{\Xi}_g f^+[n, k] - \tilde{\Xi}_g f^-[n, k])$, $D_t T_g f[n, k] := \frac{1}{2\Delta t} (T_g f^+[n, k] - T_g f^-[n, k])$, and $f^-[l] := f((l-1)(\Delta t))$ is the *forward-shifted discrete signal*. This is a discretization of the continuous formulation of second-order reassignment frequency given in Eqn. (32) in the seminal paper by Oberlin et al.²⁷ on second-order SST. We use this discretization instead of the alternative formulation proposed by those authors in Eqns. (13) and (31) of the aforementioned paper, because it allows us to avoid the issues mentioned in Sec. 4.3.

For this experiment, we set $q = 1$ and calculate the total energy on $R_m^q(C, f)$ for all $m = 1, 2, 3, 4$ and all representations C given by SST-STFT (which we also call *first-order SST-STFT*), second-order SST-STFT, RM, and SST-QSTFT. For SST-QSTFT we use Algorithm 2 with the parameters used for Figure 3. For all the transforms, we reassign the magnitude-squared of the STFT or QSTFT coefficients, in order to preserve the signal energy (see Remark 4.9). For first-order SST-STFT, second-order SST-STFT, and RM, we set the window $g = h^0$. For all transforms, $H = 250$, $K = 16384$, and $\gamma = 0$ as before. Table 1 demonstrates that SST-QSTFT achieves superior ridge concentration as the SNR decreases, especially for the IFs ϕ'_1 and ϕ'_4 with higher absolute slope. This underlies the point that the usage of chirped windows effectively exposes the presence of chirped signal content in heavy noise (as shown in Figure 5). By contrast, the second-order accurate reassignment frequency formula used for second-order SST-STFT is not as effective in detecting the local chirp rates $\{\sigma_m\}_{m=1}^4$ when noise is present. Similarly, the reassignment time used in RM is not enough to compensate for the noise. However, both second-order SST-STFT and RM improve over first-order SST-STFT, which is known to be ineffective when $|\sigma_m| \gg 0$.²⁷

5.2.2 Application to analysis of gibbon calls

In this section, we describe the potential of using SST-QSTFT together with a chirped window family to analyze audio recordings of animal calls. In particular, we analyze a dataset of recordings of female Bornean gibbon great calls, recorded at the Stability of Altered Forest Ecosystems site in Sabah, Malaysia.³⁰ One application of interest is distinguishing the individual gibbons from each other based on features extracted from their calls.

In what follows, we compute the SST-QSTFT of two calls from different gibbons, both with sampling rate $f_s = 44100$ Hz. To analyze these signals, we use the chirped window collection $\{h^w\}$ where $h^w(t) := h^0(t) e^{2\pi i \sigma^w t^2/2}$, h^0 and H are as before, $K = 3L = 12288$, and where we select $\{\sigma^w\}_{w=1}^5 = \{50, 100, 200, 400, 600\} \text{Hz}^2$ by visual




Figure 5. Time-frequency representations of noiseless (**top row**) and -6 dB SNR (**bottom row**) crossing chirps signal. **From left to right:** 1st-order SST, 2nd-order SST, RM, and SST-QSTFT. For visualization purposes, we focus on the time-frequency region between 1.5 and 2.5 seconds and 8.5 and 11.5 kHz. In SST-QSTFT plots, colors correspond to the window chosen in each supertile for the calculation of the QSTFT. Blue: h^1 . Green: h^2 . Purple: h^3 . Red: h^4 . The SST-QSTFT enables the IFs to be more easily visualized under noise, and concentrates the most signal energy along the IF ridges, as per Table 1. *Note:* The calculation in Table 1 is done over a larger time-frequency region than the one shown here.

inspection of a single STFT of the signal. We use Algorithm 2 for window selection, with $\alpha = 0.5$, $A = B = 24$, $T_{\text{step}} = Y_{\text{step}} = 4$, $T_{\text{shift}} = Y_{\text{shift}} = 2$, and $P_T = P_Y = \{-1, 1\}$.

We plot the QSTFT and SST-QSTFT of both signals in Figure 6. Here, we see that the SST-QSTFT is a more concentrated time-frequency representation of the signal than the QSTFT. Moreover, we may see also that the different curvature of the IF curves for the different gibbons is well-represented by the difference in optimal chirp rates. In future work, we seek to use the QSTFT and SST-QSTFT to extract features from the gibbon calls that may be combined with a machine learning algorithm to distinguish the gibbons from one another.

6. CONCLUSION

We have developed a synchrosqueezing transform in the context of a quilted Gabor framework, enabling for improved adaptivity to the signal, and coming closer to the goal of an ideal time-frequency representation. The SST-QSTFT yields an improved visualization of the IF information of the signal with adaptation to the time-frequency content, permitting for the precise isolation of diverse time-frequency events. We have furthermore implemented a new algorithm for automatically and adaptively selecting optimal windows depending on the time-frequency content. Reconstruction of each component is possible, even in the considered case of time-limited windows, by passing back to the original QSTFT coefficients. Theoretical results demonstrate the concentration of the QSTFT around each IF curve and the closeness of the reassignment frequency to each true IF in both the continuous and discrete frameworks, as well as the accuracy of modes reconstruction in the continuous case. Our numerical results show the effectiveness of the SST-QSTFT in adapting to the signal to achieve improved time-frequency concentration.

ACKNOWLEDGMENTS

The authors were supported in part by ONR grant N00014-16-1-2255, as well as NSF grants DMS-1418779 and IIS-1631329, and the GAANN fellowship grant P200A120162. We thank Dena Jane Clink and Mark Grote of the UC Davis Department of Anthropology for providing us with the gibbon call dataset. We also thank Dmytro Iatsenko for granting us permission to freely distribute the part of our Python code suite that we based on his




Figure 6. **Top row:** First gibbon. **Bottom row:** Second gibbon. **From left to right:** Original signal, QSTFT, and SST-QSTFT of the respective gibbon call. Here, we restrict the frequency range of the plots between 500 and 1500 Hz, where the fundamental frequency information is captured. Different colors correspond to the window chosen in each supertile for the calculation of the QSTFT. Blue: h^1 . Green: h^2 . Purple: h^3 . Red: h^4 . Black: h^5 .

MATLAB codes^{||} for SST.^{21,31} We used this Python code suite, including the part based on Dmytro Iatsenko’s codes, to generate the results in this paper.

REFERENCES

- [1] Daubechies, I. and Maes, S., “A nonlinear squeezing of the continuous wavelet transform based on auditory nerve models,” in [*Wavelets in Medicine and Biology*], Aldroubi, A. and Unser, M., eds., ch. 20, 527–546, CRC Press (1996).
- [2] Daubechies, I., Lu, J., and Wu, H.-T., “Synchrosqueezed wavelet transforms: an empirical mode decomposition-like tool,” *Applied and Computational Harmonic Analysis* **30**(2), 243–261 (2011).
- [3] Thakur, G. and Wu, H.-T., “Synchrosqueezing-based recovery of instantaneous frequency from nonuniform samples,” *SIAM Journal on Mathematical Analysis* **43**(5), 2078–2095 (2011).
- [4] Oberlin, T., Meignen, S., and Perrier, V., “The Fourier-based synchrosqueezing transform,” *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 315–319, IEEE (2014).
- [5] Feng, Z., Chen, X., and Liang, M., “Iterative generalized synchrosqueezing transform for fault diagnosis of wind turbine planetary gearbox under nonstationary conditions,” *Mechanical Systems and Signal Processing* **52**, 360–375 (2015).
- [6] Thakur, G., Brevdo, E., Fučkar, N. S., and Wu, H.-T., “The synchrosqueezing algorithm for time-varying spectral analysis: robustness properties and new paleoclimate applications,” *Signal Processing* **93**(5), 1079–1094 (2013).
- [7] Balazs, P., Dörfler, M., Jaillet, F., Holighaus, N., and Velasco, G., “Theory, implementation and applications of nonstationary Gabor frames,” *Journal of Computational and Applied Mathematics* **236**(6), 1481–1496 (2011).
- [8] Dörfler, M., “Quilted Gabor frames—a new concept for adaptive time-frequency representation,” *Advances in Applied Mathematics* **47**(4), 668–687 (2011).
- [9] Guo, Y., Fang, Z., and Chen, X., “A new improved synchrosqueezing transform based on adaptive short time Fourier transform,” *Proceedings of the IEEE Far East Forum on Nondestructive Evaluation/Testing (FENDT)*, 329–334, IEEE (2014).
- [10] Holighaus, N., Průša, Z., and Søndergaard, P. L., “Reassignment and synchrosqueezing for general time-frequency filter banks, subsampling and processing,” *Signal Processing* **125**, 1–8 (2016).

^{||}Dmytro Iatsenko’s codes can be found at <http://www.physics.lancs.ac.uk/research/nbmphysics/diats/tfr/>.

- [11] Berrian, A., Leung, J., and Saito, N., “シンクロ・スクィージング変換による信号の時間 - 周波数特徴の抽出とデータ・ソニフィケーションへの応用 (Time-frequency feature extraction via synchrosqueezing transform and its application to data sonification),” *Proceedings of the Annual Conference of the Japan Society for Industrial and Applied Mathematics*, 336–337 (2015).
- [12] Sheu, Y.-L., Hsu, L.-Y., Chou, P.-T., and Wu, H.-T., “Entropy-based time-varying window width selection for nonlinear-type time–frequency analysis,” *International Journal of Data Science and Analytics* **3**, 1–15 (2017).
- [13] Müller, M., [*Fundamentals of Music Processing: Audio, Analysis, Algorithms, Applications*], Springer International Publishing (2015).
- [14] Smith, J. O., [*Spectral Audio Signal Processing*], <http://ccrma.stanford.edu/~jos/sasp/> (accessed 2015-02-04). Online book, 2011 edition.
- [15] Folland, G., [*Fourier Analysis and Its Applications*], American Mathematical Society (1992).
- [16] Su, L. and Wu, H.-T., “Minimum-latency time-frequency analysis using asymmetric window functions,” *arXiv preprint arXiv:1606.09047* (2016).
- [17] Behera, R., Meignen, S., and Oberlin, T., “Theoretical analysis of the second-order synchrosqueezing transform,” *Applied and Computational Harmonic Analysis* (2016).
- [18] Fourer, D., Auger, F., and Flandrin, P., “Recursive versions of the Levenberg-Marquardt reassigned spectrogram and of the synchrosqueezed STFT,” *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4880–4884 (2016).
- [19] Chui, C. K., Lin, Y.-T., and Wu, H.-T., “Real-time dynamics acquisition from irregular samples—with application to anesthesia evaluation,” *Analysis and Applications* **14**(04), 537–590 (2016).
- [20] Auger, F., Flandrin, P., Lin, Y.-T., McLaughlin, S., Meignen, S., Oberlin, T., and Wu, H.-T., “Time-frequency reassignment and synchrosqueezing: An overview,” *IEEE Signal Processing Magazine* **30**(6), 32–41 (2013).
- [21] Iatsenko, D., McClintock, P. V., and Stefanovska, A., “Extraction of instantaneous frequencies from ridges in time–frequency representations of signals,” *Signal Processing* **125**, 290–303 (2016).
- [22] Meignen, S., Oberlin, T., and McLaughlin, S., “Fully adaptive mode decomposition from time-frequency ridges,” *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3884–3888, IEEE (2017).
- [23] Jaillet, F. and Torrèsani, B., “Time-frequency jigsaw puzzle: adaptive multiwindow and multilayered Gabor expansions,” *International Journal of Wavelets, Multiresolution and Information Processing* **5**(02), 293–315 (2007).
- [24] Liuni, M., *Automatic adaptation of sound analysis and synthesis*, PhD thesis, Université Pierre et Marie Curie – Paris VI (2013).
- [25] Ricaud, B. and Torrèsani, B., “A survey of uncertainty principles and some signal processing applications,” *Advances in Computational Mathematics* **40**(3), 629–650 (2014).
- [26] Coifman, R. R. and Donoho, D. L., “Translation-invariant de-noising,” in [*Wavelets and Statistics*], Antoniadis, A. and Oppenheim, G., eds., *Lecture Notes in Statistics* **103**, 125–150, Springer New York (1995).
- [27] Oberlin, T., Meignen, S., and Perrier, V., “Second-order synchrosqueezing transform or invertible reassignment? Towards ideal time-frequency representations.,” *IEEE Transactions on Signal Processing* **63**(5), 1335–1344 (2015).
- [28] Kodera, K., De Villedary, C., and Gendrin, R., “A new method for the numerical analysis of non-stationary signals,” *Physics of the Earth and Planetary Interiors* **12**(2-3), 142–150 (1976).
- [29] Auger, F. and Flandrin, P., “Improving the readability of time-frequency and time-scale representations by the reassignment method,” *IEEE Transactions on Signal Processing* **43**(5), 1068–1089 (1995).
- [30] Clink, D. J., Bernard, H., Crofoot, M. C., and Marshall, A. J., “Investigating individual vocal signatures and small-scale patterns of geographic variation in female Bornean gibbon (*Hylobates muelleri*) great calls,” *International Journal of Primatology*, 1–16 (2017).
- [31] Iatsenko, D., McClintock, P. V., and Stefanovska, A., “Linear and synchrosqueezed time–frequency representations revisited: Overview, standards of use, resolution, reconstruction, concentration, and algorithms,” *Digital Signal Processing* **42**, 1–26 (2015).