

Super-resolution

1 Super-resolution of point sources

The word *super-resolution* has different meanings in different disciplines. In optics it is usually used to describe data-acquisition techniques designed to overcome the diffraction limit [8]. In image-processing and computer-vision applications the term tends to refer to the problem of obtaining a high-resolution image either from several low-resolution images [10] or by upsampling a single image while preserving its edges and hallucinating textures in a reasonable way [5]. In this lecture, super-resolution denotes the inverse problem of estimating a signal from low-pass measurements. In contrast to the previous definitions, we assume that the data-acquisition mechanism is *fixed* and we aim to *recover* rather than hallucinate the lost high-resolution information.

1.1 Spatial super-resolution

As Lord Rayleigh pointed out in his seminal 1891 paper On Pin-hole Photography, *it has long been known that the resolving power of lenses, however perfect, is limited*. Diffraction imposes an inflexible limit on the resolution of any optical system. However we are often interested in information that is only apparent beyond this limit. For instance, in microscopy, astronomy or medical imaging it may be challenging to discern cellular structures, celestial bodies or incipient tumors from the available data. This is illustrated by the image on the left of Figure 1, which shows measurements of the interior of a cell obtained by fluorescence microscopy. The limited resolution of the microscope produces aliasing artifacts that completely obscure the fine-scale details of the image. The aim of super-resolution is to uncover such fine-scale structure from coarse-scale measurements.

Let x be a high-resolution representation of an object of interest. Mathematically, a reasonable model for the data in many of the applications mentioned in the previous paragraph is

$$x_{\text{LR}} := \phi * x. \quad (1)$$

The signal x is convolved with a low-pass point spread function (PSF) ϕ that depends on the sensing mechanism. This convolution smooths out the fine-scale details producing a low-resolution version of the original signal. The aim of *spatial super-resolution* is to estimate x from x_{LR} , a problem which is often also referred to as *deconvolution*. Recovering the lost fine-scale information amounts to extrapolating the spectrum of x . This becomes apparent




Figure 1: An example of super-resolution applied to cell imaging using fluorescence microscopy data. In both images a superposition of a sequence of frames is shown. The data on the left are contaminated by noise and heavy aliasing. Super-resolving the probes reveals the fine-scale details of the cell, as we can see on the right. The image is due to Veniamin Morgenshtern.

when we consider the measurement process in the frequency domain. If the cut-off frequency of the PSF is equal to f_c , the spectrum of x_{LR}

$$\widehat{x}_{\text{LR}} = \widehat{\phi} \widehat{x} = \widehat{\phi} \Pi_{[-f_c, f_c]^d} \widehat{x}, \quad (2)$$

where d is the ambient dimension, $\widehat{\cdot}$ denotes the frequency representation of a signal and $\Pi_{[-f_c, f_c]^d}$ is an indicator function that is zero out of the set $[-f_c, f_c]^d$. The high frequency information in the signal is suppressed in the data. Super-resolution aims to extrapolate this information from the low-pass measurements.

In order to super-resolve a signal it is necessary to leverage some prior knowledge about its structure. Otherwise the problem is hopelessly ill posed; the missing spectrum can be filled in arbitrarily to produce estimates that correspond to the data. In this lecture we consider signals that are well modeled as superpositions of point sources,

$$x := \sum_{t_j \in T} c_j \delta_{t_j}, \quad (3)$$

where δ_τ is a Dirac measure at τ , T is a set of locations in the unit interval and the amplitudes a_j may be complex valued. In this case, the low-resolution data are of the form

$$x_{\text{LR}}(t) := \phi * x(t) \quad (4)$$

$$= \sum_{t_j \in T} c_j \phi(t - t_j), \quad (5)$$

as illustrated at the top of Figure 2.

Point sources are used to represent celestial bodies in astronomy, neuron spikes in neuroscience or line spectra in signal processing and spectroscopy. In addition, locating pointwise fluorescent probes is a crucial step in some optical super-resolution procedures capable of handling more complicated objects. Techniques such as photoactivated localization microscopy (PALM) or stochastic optical reconstruction microscopy (STORM) are based on localizing probes that switch randomly between a fluorescent and a non-fluorescent state. To obtain an image of a certain object, multiple frames are gathered. Each frame consists of a superposition of blurred light sources that correspond to the active probes. Deblurring these sources and combining the results allows to super-resolve the object of interest. The image to the right of Figure 1 was generated in this way.

An additional assumption that we make in this lecture is that the PSF of the sensing mechanism is known. This is usually a realistic assumption in point-source super-resolution as long as the measurement process is indeed space-invariant. In such cases the PSF can be estimated by locating an isolated blurred source in the data. For general images or PSFs that are not low pass, which arise for instance due to motion blurring, it is necessary to jointly estimate the PSF and the signal; a problem known as *blind deconvolution* which we will not discuss here.

An important instance of the model given by (3) and (4) is when the signal is one dimensional and the Fourier transform of the PSF ϕ is constant over $[-f_c, f_c]$, i.e. ϕ is a periodized sinc or Dirichlet kernel. In this case, since the support of x_{LR} is restricted to the unit interval, it follows from the sampling theorem that its spectrum is completely determined by the discrete samples

$$y_k := \widehat{x}_{\text{LR}}(k) \tag{6}$$

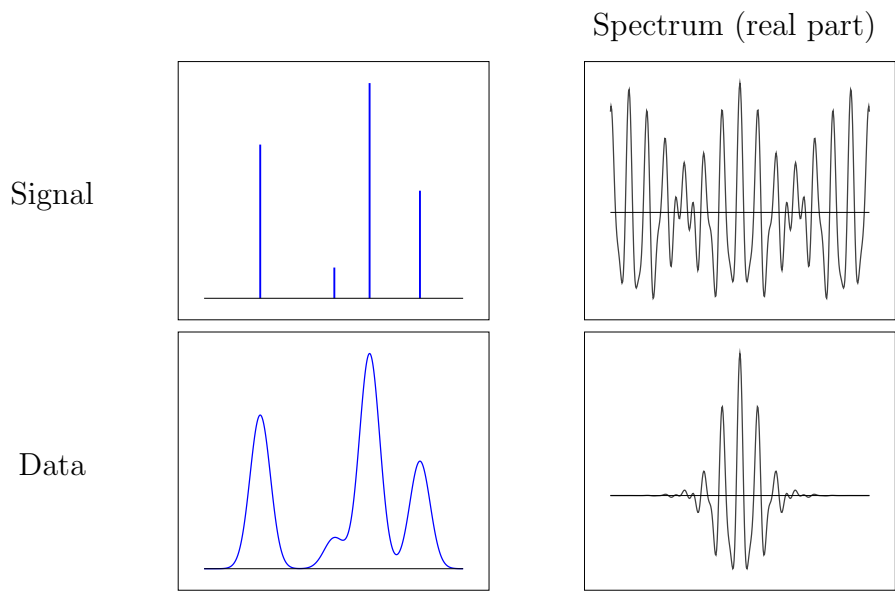
$$= \int_0^1 \exp(-i2\pi kt) x(dt) \tag{7}$$

$$= \sum_{t_j \in T} c_j \exp(-i2\pi kt_j), \quad k \in \mathbb{Z}, |k| \leq f_c, \tag{8}$$

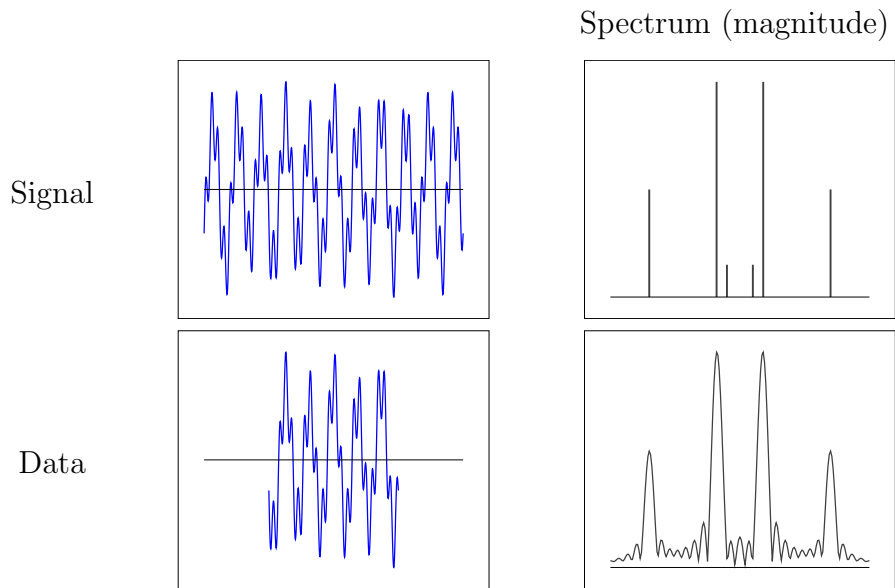
where we assume for simplicity that f_c is an integer. In a more compact form,

$$y = \mathcal{F}_n x \tag{9}$$

where $y \in \mathbb{C}^n$ and \mathcal{F}_n is the linear operator that maps a measure or function to its lowest $n := 2f_c + 1$ Fourier coefficients.



Spatial super-resolution



Spectral super-resolution

Figure 2: Schematic illustration of spatial and spectral super-resolution.

1.2 Spectral super-resolution

A basic problem in signal processing is to estimate the spectrum of a multisinusoidal signal

$$g(t) := \sum_{f_j \in T} c_j \exp(-i2\pi f_j t) \quad (10)$$

from a finite number of samples. The corresponding spectrum of the signal is of the form

$$\hat{g} = \sum_{f_j \in T} c_j \delta_{f_j}. \quad (11)$$

Spectra that are composed of a superposition of Dirac measures are called *line spectra*. Estimating line spectra is important for direction-of-arrival (DOA) problems in radar where the aim is to determine the direction from which a propagating wave arrives at an array of sensors.

Let us assume that the spectrum of g is restricted to the interval $[-1/2, 1/2]$. This means that its cut-off frequency is $f = 1/2$. Regular samples taken at the rate dictated by the Nyquist-Shannon theorem $2f = 1$

$$g(k) := \sum_{f_j \in T} c_j \exp(-i2\pi k f_j), \quad \infty \leq k \leq \infty, \quad (12)$$

suffice to recover \hat{g} . Indeed, these samples $g(k)$ are the Fourier coefficients of \hat{g} !

Unfortunately, in most practical situations we only have access to a finite number of samples. Truncating the measurements in the time domain is equivalent to convolving the spectrum of the signal with a periodized sinc. This induces aliasing in the frequency domain as depicted in the lower half of Figure 2. *Spectral super-resolution* is the problem of estimating the line spectra of g from such measurements. This is exactly equivalent to the spatial super-resolution in 1D where the PSF is a sinc. Indeed, \hat{g} is a superposition of Dirac measures just like x and the measurements are a contiguous subset of the Fourier coefficients of the measure.

1.3 Deconvolution in reflection seismography

Characterizing underground geological structure, such as the strata portrayed in Figure 3, from surface measurements is an important problem in geophysics. Reflection seismology is a technique to achieve this by sending pulses into the ground and analyzing the reflected signals. Figure 4 provides a cartoon description of the problem. The reflections at the interface between two strata are governed by a *reflection coefficient* that corresponds to the difference between the impedance of the two layers. Estimating the reflection coefficients and their positions from the reflected pulse allows to determine the composition of the different strata.




Figure 3: Geological strata in Salta, Argentina.

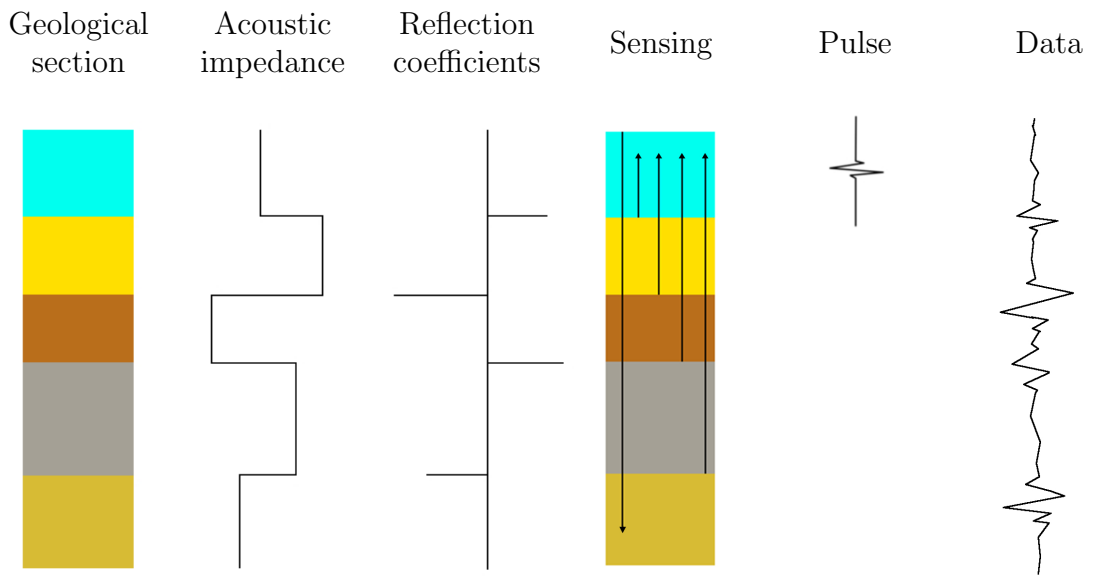


Figure 4: Cartoon description of reflection seismology.

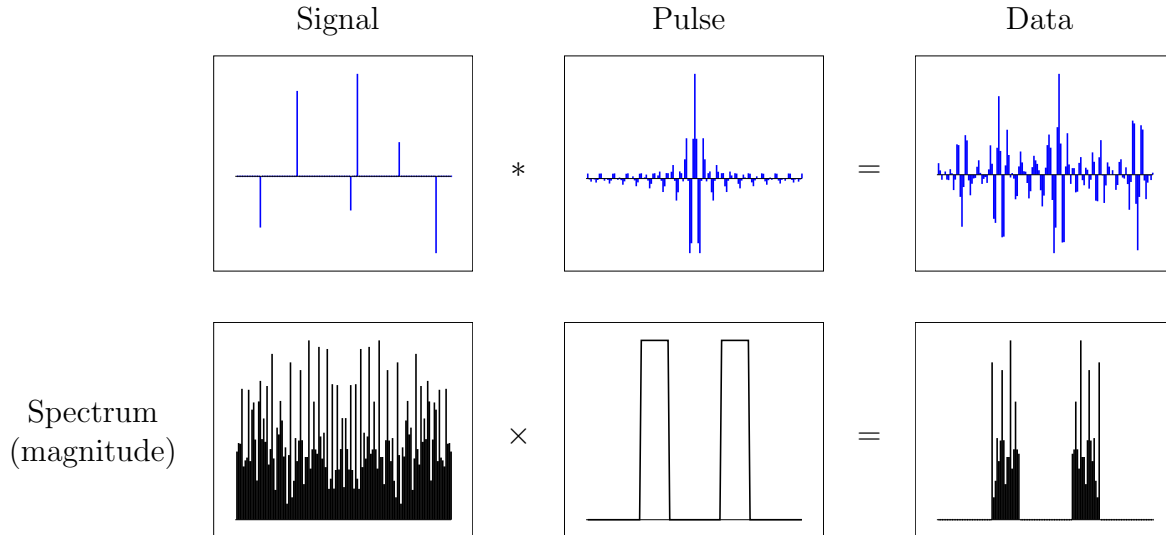


Figure 5: Simplified model for a seismography problem where the aim is to estimate reflection coefficients from bandpass measurements.

A simplified mathematical model of the measurement process is that the data are equal to the convolution of the pulse ϕ and the reflection coefficients x ,

$$y = \phi * x \quad (13)$$

where the reflection coefficients can be modeled as a superposition of Dirac measures,

$$x := \sum_{t_j \in T} c_j \delta_{t_j}. \quad (14)$$

The spectrum of pulses used in reflection seismology is often approximately bandpass; it is mostly concentrated in a certain band of frequencies. If we model the pulse as being exactly bandpass, then the spectrum of the data corresponds to a contiguous subset of Fourier coefficients of x , or two contiguous subsets if we consider real bandpass filters as illustrated in Figure 5. The problem is consequently very similar to the simplified models for spatial and spectral super-resolution introduced in the previous sections.

2 Conditioning of the super-resolution problem

In this section we analyze the conditioning of the super-resolution problem when the signal is equal to a superposition of Dirac measures and the data are equal to the first n low-pass Fourier coefficients

$$y = \mathcal{F}_n x. \quad (15)$$




Figure 6: Illustration of the minimum separation of a signal.

As we discussed in the previous lecture, a fundamental insight in the theoretical analysis of compressed sensing is that random measurements conserve the energy of sparse signals with high probability. More specifically, the randomized operator obeys the restricted-isometry property (RIP). This ensures that the inverse problem of sparse recovery from random measurements is well posed at least in principle.

In contrast, sparsity is not enough to make super-resolution a well posed problem. The crucial difference between super-resolution and compressed sensing is the correlation structure of the measurement process. The sensing operator \mathcal{F}_c defined by (19) can be interpreted as an infinite matrix with n rows and infinite columns parametrized by $t \in [0, 1]$. Each of the columns is a sinusoidal atom of the form $\exp i2\pi kt$ for $-f_c \leq k \leq f_c$. If the difference between t_1 and t_2 is small, the corresponding columns

$$(\mathcal{F}_n)_{t_1} = \begin{bmatrix} \exp -i2\pi f_c t_1 \\ \exp -i2\pi (f_c - 1) t_1 \\ \dots \\ \exp i2\pi f_c t_1 \end{bmatrix} \approx (\mathcal{F}_n)_{t_2} = \begin{bmatrix} \exp -i2\pi f_c t_2 \\ \exp -i2\pi (f_c - 1) t_2 \\ \dots \\ \exp i2\pi f_c t_2 \end{bmatrix} \quad (16)$$

are extremely correlated. As a result, the measurement process does not even satisfy the RIP with sparsity level equal to two! In this case, we cannot recover a 1-sparse signal reliably in the presence of noise. However, if the difference between t_1 and t_2 is large then the atoms are not too correlated. In fact, as we shall demonstrate numerically in a moment, submatrices corresponding to columns that are spread out are well conditioned. This suggests restricting the class of signals of interest to signals that have a support that is not too clustered. To this end, we define the *minimum separation* of the support T of a signal. The concept is illustrated in Figure 6.

Definition 2.1 (Minimum separation). *For a set of points $T \subseteq [0, 1]$, the minimum separation (or minimum distance) is defined as the closest distance between any two elements from*

T ,

$$\Delta(T) = \inf_{(t,t') \in T: t \neq t'} |t - t'|. \quad (17)$$

In order to determine how the linear operator \mathcal{F}_n acts upon signals with supports that satisfy a certain minimum separation, we compute the singular values of the submatrix formed by the corresponding columns. Figure 7 shows the results. For minimum separations above $0.5/f_c$ the submatrices are very well conditioned, but as soon as the minimum separation falls below $0.5/f_c$ some of the singular values of the operator plummet. Below $0.35/f_c$ a fourth of the singular values are extremely small. Below $0.2/f_c$ the magnitude of half of them is negligible. The effect becomes more pronounced as the cardinality s of the support of the signal increases.

Note that signals in the span of the singular vectors that correspond to very small singular values in Figure 8 are essentially mapped to zero by the low-resolution operator. These signals are consequently almost in the null space of the measurements: they are impossible to estimate under very small perturbations to the data *even if we know the support beforehand*. As a result, a fundamental limit for signal recovery in terms of the minimum separation is $1/f_c$, the inverse of the cut-off frequency (it also corresponds to the width of the main lobe of the point-spread function ϕ , when the PSF is a sinc). To see why, consider two signals x_1 and x_2 with a minimum separation just below $1/f_c$ such their difference $d := x_1 - x_2$ has a minimum separation under $0.5/f_c$. From Figure 7 we can choose d so that

$$\mathcal{F}_n d = \mathcal{F}_n x_1 - \mathcal{F}_n x_2 \approx 0. \quad (18)$$

Even under very low levels of noise it will be impossible to distinguish x_1 and s_2 from low-resolution data. Figure 8 illustrates this: the measurements corresponding to two signals with disjoint supports and a minimum distance of $0.9\lambda_c$ for $f_c = 10^3$ are indeed almost indistinguishable. The phenomenon can be characterized theoretically in an asymptotic setting using Slepian's prolate-spheroidal sequences [12] (see also Section 3.2 in [2]). More recently, [9] (see Theorem 1.3) provides a non-asymptotic analysis.

The take-home lesson is that we can only hope to achieve stable super-resolution of signals with a minimum separation above $0.5/f_c$ if the signal consists of a large number of point sources. For signals with a small number of point sources, robust recovery may be possible below that limit. However, even for very small supports, the problem is always hopelessly ill posed for a small enough value of the minimum distance.

3 Linear methods

In this section we describe some simple linear techniques to tackle the super-resolution problem.

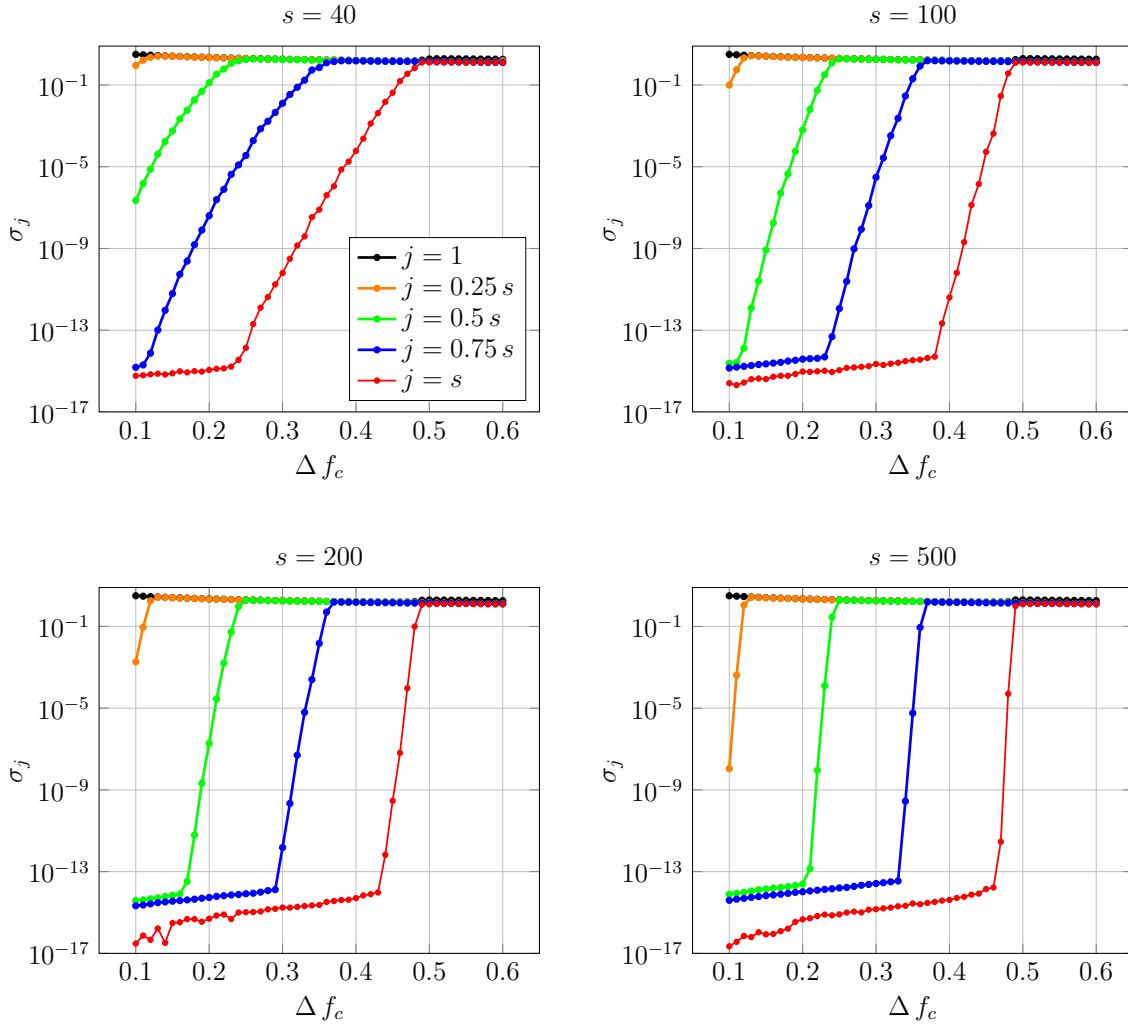


Figure 7: Singular values of the linear operator \mathcal{F}_n when restricted to act upon signals of sparsity s that have a minimum separation of Δ .

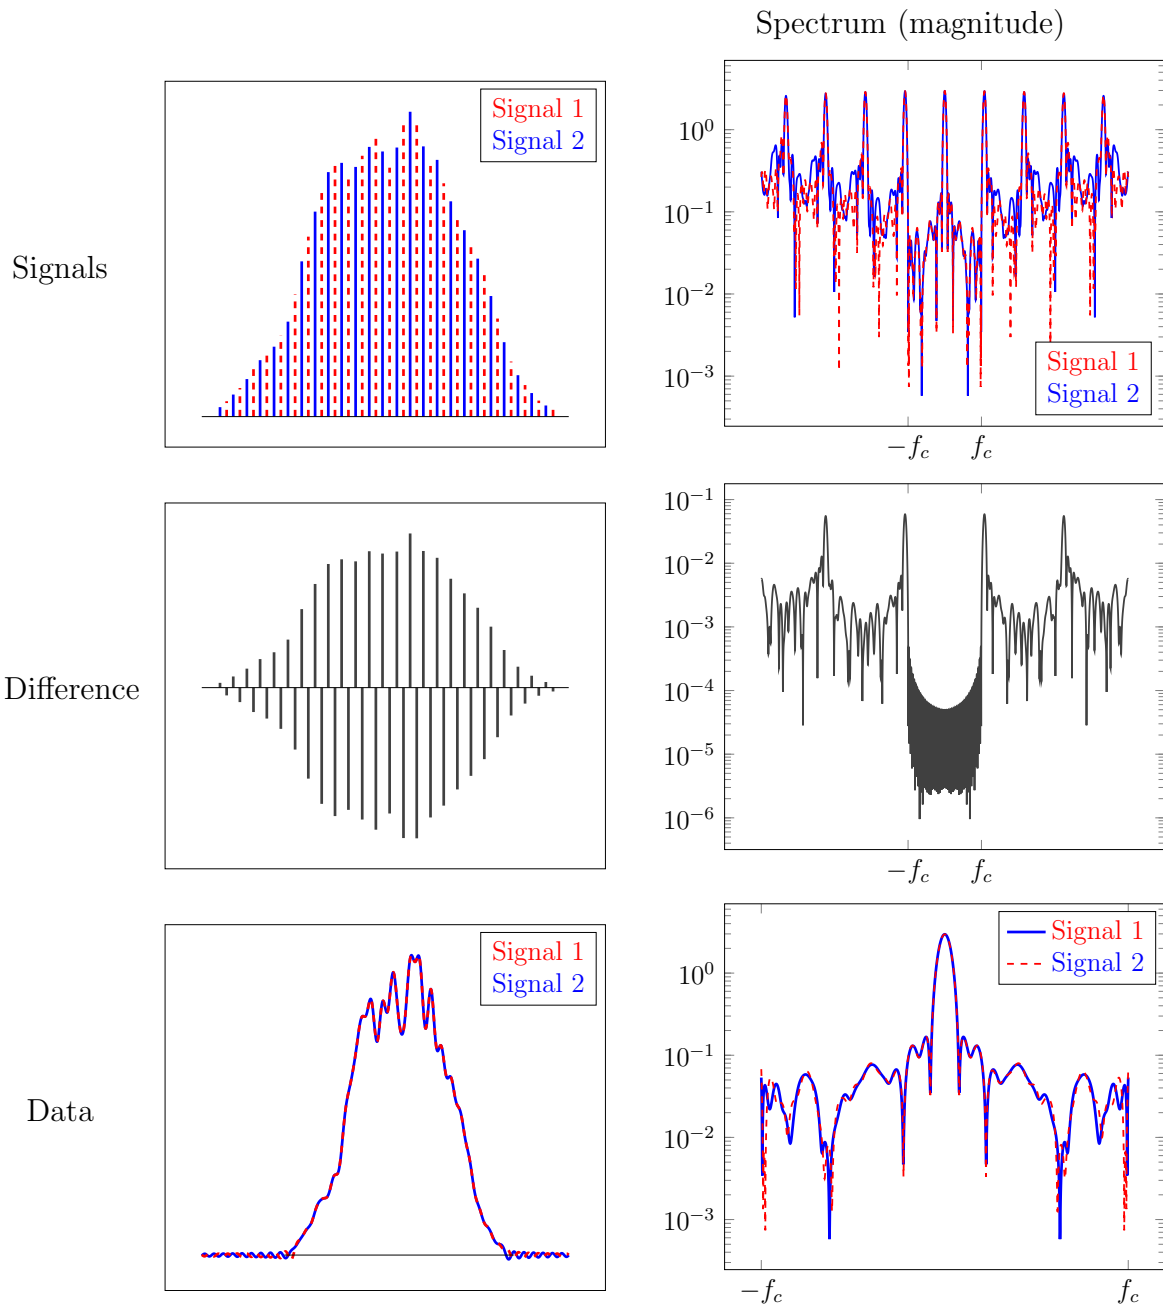


Figure 8: Two signals with disjoint supports that satisfy the minimum-separation condition for $\Delta(T) = 0.9f_c$ when $f_c = 10^3$ (top left) and their spectrum (top right). Their difference (center left) has a spectrum that is concentrated away for the low-pass band between $-f_c$ and f_c (center right). As a result, it is very difficult to distinguish the data corresponding to the two signals (bottom).

3.1 The periodogram

As explained in Section 1.2 in the spectral super-resolution problem the data (12) correspond to a contiguous subset of the Fourier coefficients of the signal

$$y = \mathcal{F}_n x. \quad (19)$$

If x corresponds to a superposition of line spectra, this is equivalent to convolving the sparse spectrum with a periodized sinc function. The periodogram recovers this representation by applying the adjoint of the measurement operator \mathcal{F}_n to the data. This is essentially equivalent to projecting the data onto the signal space.

$$P(t) = \mathcal{F}_n^* y \quad (20)$$

$$= \sum_{k=-f_c}^{f_c} y(k) \exp(i2\pi kt) \quad (21)$$

$$= \sum_{k=-f_c}^{f_c} \left(\sum_{t_j \in T} c_j \exp(i2\pi kt_j) \right) \exp(i2\pi kt) \quad (22)$$

$$= \sum_{t_j \in T} c_j \sum_{k=-f_c}^{f_c} \exp(i2\pi k(t - t_j)) \quad (23)$$

$$= \sum_{t_j \in T} c_j D_{f_c}(t - t_j), \quad (24)$$

where D_{f_c} is the periodized sinc or Dirichlet kernel

$$D_{f_c}(t) := \sum_{k=-f_c}^{f_c} \exp(i2\pi kt) = \begin{cases} 1 & \text{if } t = 0 \\ \frac{\sin((2f_c+1)\pi t)}{(2f_c+1)\sin(\pi t)} & \text{otherwise.} \end{cases} \quad (25)$$

Just to clarify, in most applications that use the periodogram t would index the frequency domain, *not* the time domain, but we keep this notation for the sake of consistency. Also, f_c no longer has a physical meaning (beyond determining the number of time-domain samples) in contrast to spatial super-resolution where it represents the cut-off frequency of the sensing mechanism.

Computing the periodogram does not solve the super-resolution problem, it just allows to visualize the aliased spectrum corresponding to the data samples available in spectral super-resolution. If the line spectra of the signal are far apart, a straightforward way to estimate their location is to locate local maxima of P . The problem with this approach is that the side lobes corresponding to large blurred spikes may mask the presence of smaller spikes. As a result, the periodogram is not very useful if the spikes are not far enough from each other

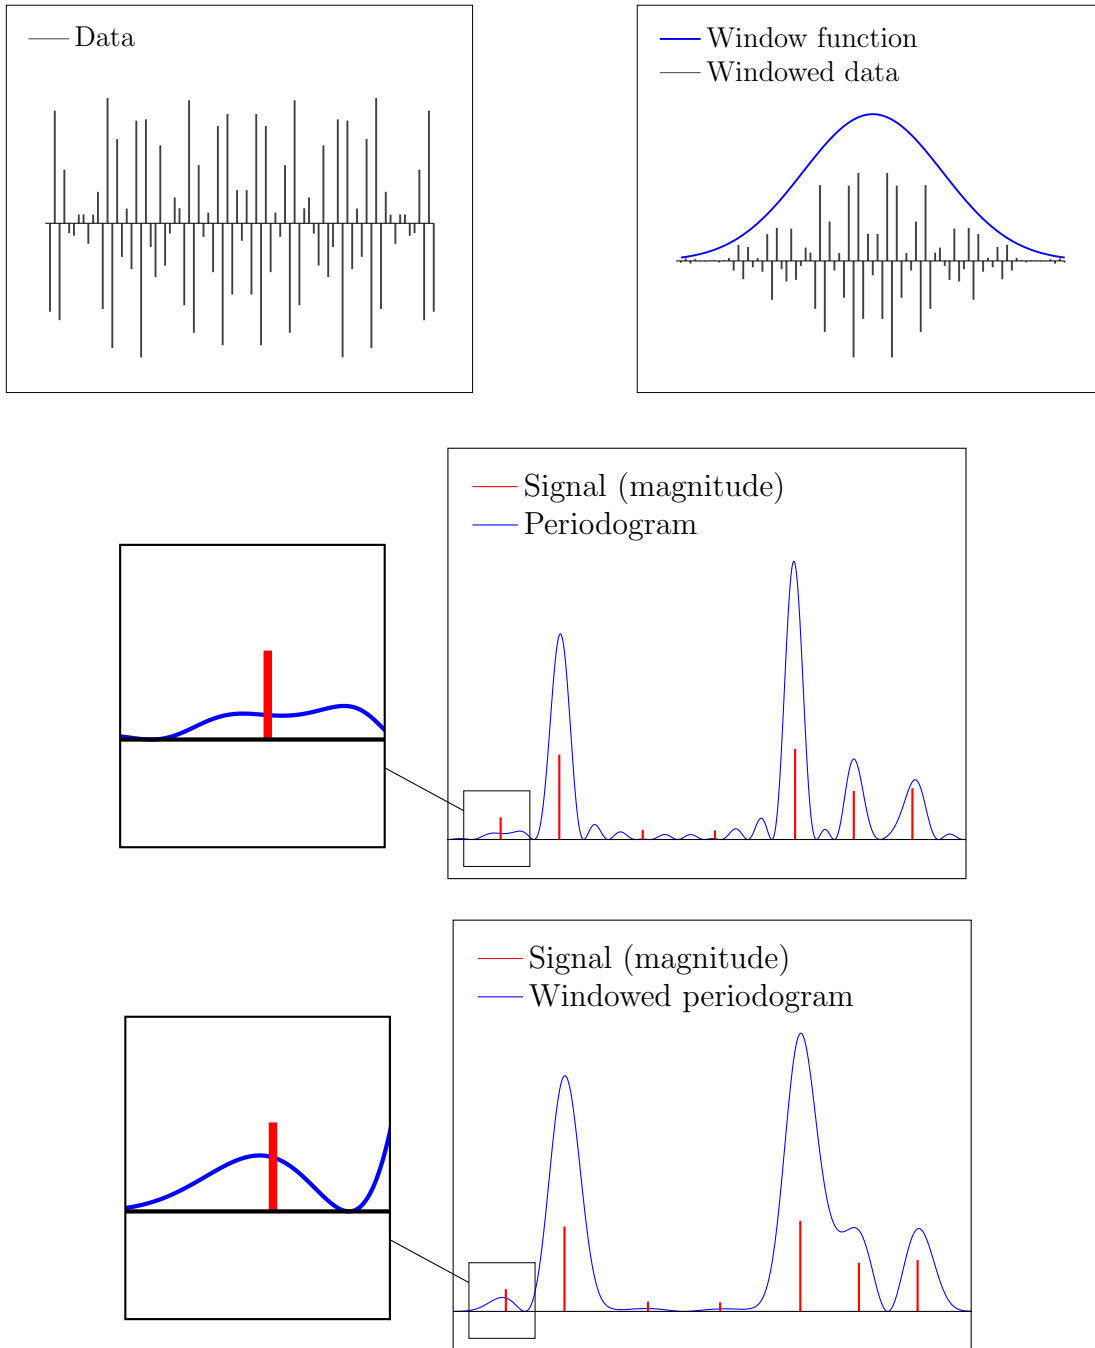


Figure 9: Example of data in spectral super-resolution generated according to (6) before (top left) and after applying a window function (top right). No noise is added to the data. Below we can see the periodogram (center) and windowed periodogram (bottom) computed from the data. The scaling of the periodograms is set so that both the large and small peaks can be seen on the plot.

or if their amplitudes differ substantially, *even if no noise is present in the data*. The image at the center of Figure 9 illustrates this: detecting some of the lower-amplitude spikes from the periodogram is impossible.

In order to alleviate this problem one can apply a window function $\hat{w} \in \mathbb{C}^n$ to the data before computing the periodogram,

$$y_{\hat{w}} = y \cdot \hat{w}, \quad (26)$$

where \cdot denotes pointwise multiplication. The windowed periodogram

$$P_{\hat{w}}(f) = \mathcal{F}_n^* y_{\hat{w}} \quad (27)$$

$$= \sum_{t_j \in T} c_j w(t - t_j), \quad (28)$$

where w denotes the inverse Fourier transform of the window function. Ideally, w should be as *spiky* as possible to make it easier to locate the support of the signal from the windowed periodogram. However, this is challenging due to the constraint that \hat{w} has finite support and hence w is a low-pass function.

In the image on the top right of Figure 9 we apply a Gaussian window to the data. To be more precise, we set \hat{w} to be a truncated Gaussian, so that w is equal to the convolution between a periodized sinc and a periodized Gaussian. The resulting periodogram, shown at the center of Figure 9, has much less *spectral leakage* from the largest signal components, due to the fact that the Gaussian window has lower side lobes than the periodized sinc. However, the latter is spikier at the origin, which allows to better distinguish neighboring spikes with similar amplitudes. In general, designing an adequate window implies finding a good tradeoff between the width of the main lobe and the height of the side lobes. We refer the reader to [6] for a detailed account of design considerations and types of window function. As discussed in Lecture Notes 4, windowing is also important when computing short-time Fourier transforms.

Figure 10 compares the result of applying the periodogram and the windowed periodogram to noisy data from two signals with different minimum separations. The signal-to-noise ratio is 20 dB (equivalently the ratio between the ℓ_2 -norm of the signal and the noise is equal to 10). When the minimum separation is large the windowed periodogram provides a relatively good estimate of the number of point sources, although their locations do not exactly coincide with the local maxima for some of the spikes. Without windowing, the periodogram provides sharper localization but suffers from spurious spikes produced by the side lobes of the sinc function. When the minimum separation is small, it is impossible to super-resolve many of the spikes, which is not surprising: the problem becomes very ill-conditioned as discussed in Section 2.

To conclude, the periodogram is a useful method to obtain a general idea of the spectral structure of the signal and can provide insight as to the number of sources and their approximate location, especially if the data is processed with a suitable window function and the

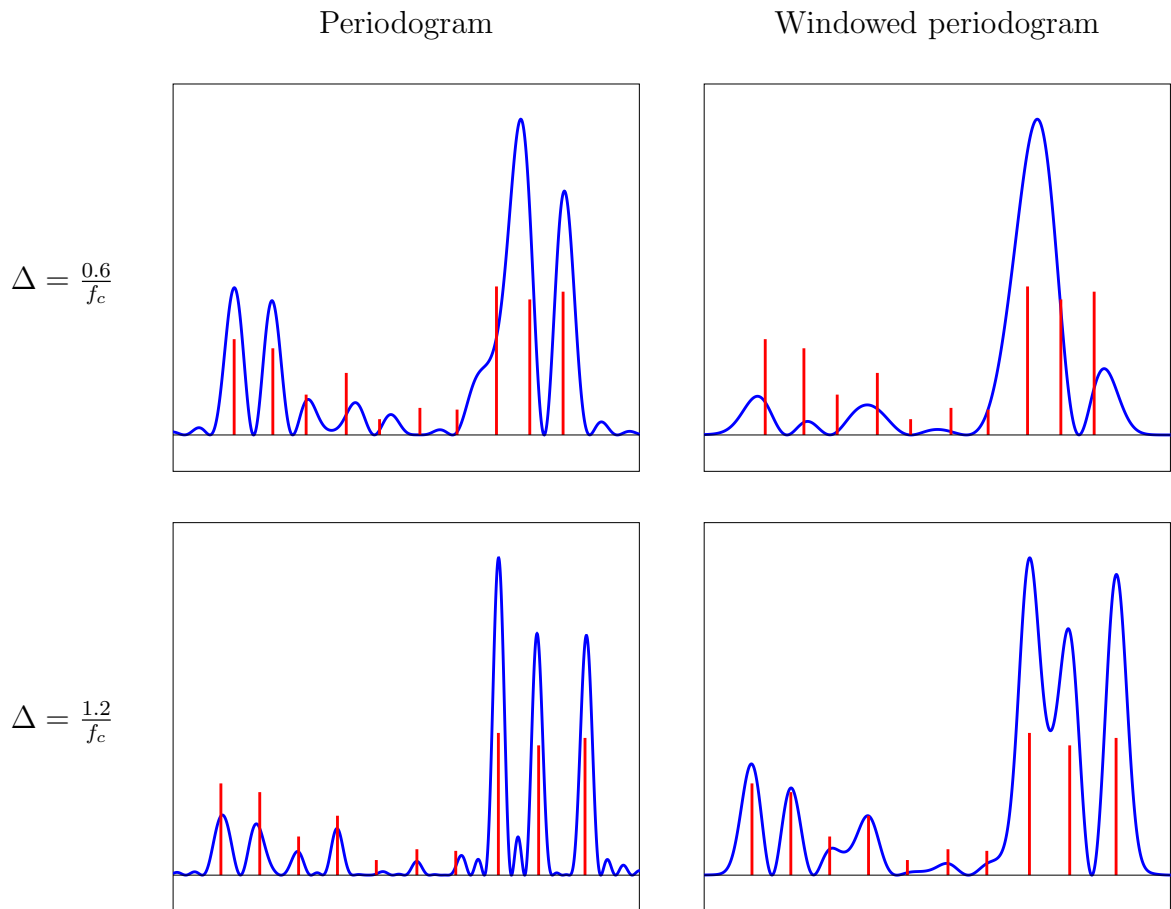


Figure 10: Periodogram (left) and windowed periodogram (right) for two signals with identical amplitudes but different minimum separations.

support of the signal is not too clustered together. However, the method does not provide a precise estimate of the support even in the absence of noise. This is the aim of the techniques described in the following sections.

3.2 Local fitting

If we have low-resolution measurements of the form (4) and the signal consists of a single blurred source then

$$x_{\text{LR}}(t) := c_1 \phi(t - t_1). \quad (29)$$

In this case, it seems reasonable to estimate the location of the source by fitting the shifted PSF to the data, for instance, by minimizing a least-squares cost function. For a single source, performing a least-squares fit amounts to finding the shift that maximizes the inner product between the data and the PSF ϕ shifted by \tilde{t} , which is denoted by $\phi_{\tilde{t}}$. Let us assume that x_{LR} and $\phi_{\tilde{t}}$ are in \mathcal{L}_2 and that $\|\phi_{\tilde{t}}\|_{\mathcal{L}_2} = 1$. For a fixed \tilde{t} the best coefficient is obtained by projecting x_{LR} onto $\phi_{\tilde{t}}$

$$\langle x_{\text{LR}}, \phi_{\tilde{t}} \rangle = \arg \min_{\tilde{t}, \alpha \in \mathbb{C}} \|x_{\text{LR}} - \alpha \phi_{\tilde{t}}\|_2. \quad (30)$$

As a result, choosing the first shift is equivalent to maximizing the inner product

$$\begin{aligned} t_{\text{est}} &= \arg \min_{\tilde{t}} \min_{\alpha \in \mathbb{C}} \|x_{\text{LR}} - \alpha \phi_{\tilde{t}}\|_2 \\ &= \arg \min_{\tilde{t}} \|x_{\text{LR}} - \langle x_{\text{LR}}, \phi_{\tilde{t}} \rangle \phi_{\tilde{t}}\|_2 \\ &= \arg \min_{\tilde{t}} -2 |\langle x_{\text{LR}}, \phi_{\tilde{t}} \rangle|^2 \\ &= \arg \max_{\tilde{t}} |\langle x_{\text{LR}}, \phi_{\tilde{t}} \rangle|. \end{aligned} \quad (31)$$

This procedure is known as *matched filtering* in the signal-processing literature.

When more than one point source is present, the overlap between neighboring blurred sources may be negligible if the sources are far enough with respect to the decay of the PSF. In this case, one can just fit the PSF locally to each of them. The approach is essentially equivalent to applying matching pursuit (Algorithm 3.2 in Lecture Notes 4) with a dictionary of shifted copies of ϕ . This method is quite popular in fluorescence microscopy. However, if there is interference between the blurred sources, then the problem becomes much more challenging. As a result, in fluorescence microscopy the data is usually measured at a rate that ensures a certain separation between the active fluorophores.

4 Parametric methods for spectral super-resolution

In this section we review parametric methods for spectral super-resolution based on Prony's method. These algorithms take the number of line spectra of the original signal as input and return an estimate of their location. As in our description of the periodogram, we use the variable t to denote frequency for the sake of consistency with other sections.

4.1 Prony's method

Prony's method solves the spectral super-resolution problem by encoding the position of s sources $T = \{t_1, \dots, t_s\}$ as the zeros of a trigonometric polynomial of order s . In the absence of noise, it turns out that we can always find such a polynomial and achieve perfect spectral super-resolution by locating its roots.

We consider a signal x

$$x := \sum_{t_j \in T} c_j \delta_{t_j}, \quad (32)$$

and data that correspond to the n first Fourier coefficients of x

$$y_k := \widehat{x}(k), \quad 0 \leq k \leq n-1. \quad (33)$$

If we set $n := 2f_c + 1$, this model is essentially equivalent to (6).

The following lemma shows that for *any* x of the form (32) there exists a nonzero polynomial of order s which is exactly zero on the support of the x .

Lemma 4.1 (Existence of Prony polynomial). *Assume that x is as in (3) and $|T| = s$, then*

$$P_{\text{prony}}(t) := \prod_{j=1}^s (1 - \exp(i2\pi(t - t_j))) \quad (34)$$

$$= 1 + \sum_{l=1}^s v_l \exp(i2\pi lt), \quad v_0 := 1, \quad (35)$$

is a nonzero polynomial of order s with roots located exactly on T .

Proof. Since $v_0 = 1$ the polynomial is nonzero. By construction $P_{\text{prony}}(t_j) = 0$ for every $t_j \in T$. \square

If we are able to compute such a polynomial, then finding its roots immediately reveals the support of x . The problem is how to compute P_{prony} from the available data y . Prony's

method consists of setting up a system of linear equations that only depends on y such that the solution is equal to the coefficients of P_{prony} .

Since P_{prony} is zero wherever x is nonzero, the inner product of x and P_{prony} is equal to zero,

$$\langle P_{\text{prony}}, x \rangle = \int_0^1 P_{\text{prony}}(t) \bar{x}(dt) \quad (36)$$

$$= \sum_{j=0}^s \bar{c}_j \exp(-i2\pi lt_j) (1 - \exp(i2\pi(t_j - t_j))) \quad (37)$$

$$= 0. \quad (38)$$

By Parseval's Theorem this inner product equals the inner product between the Fourier coefficients of x and P_{prony} which implies

$$\langle v, \hat{x} \rangle = \langle P_{\text{prony}}, x \rangle = 0. \quad (39)$$

We can use this identity to obtain an equation tying the coefficient vector v to the data.

$$0 = \langle v, \hat{x} \rangle = \sum_{k=-\infty}^{\infty} v_k \overline{\hat{x}(k)} \quad (40)$$

$$= \sum_{l=0}^s v_l \overline{\hat{x}(k)} \quad (41)$$

$$= \sum_{l=0}^s v_l \overline{y(k)} \quad \text{if } s+1 \leq n. \quad (42)$$

In order to estimate v we need at least $s-1$ other equations. Consider taking the inner product between the coefficients of P_{prony} and \hat{x} shifted by k' , which we denote by

$$\hat{x}_{k'} := \hat{x}(k+k'). \quad (43)$$

By basic Fourier identities, $\hat{x}_{k'}$ is the Fourier representation of $\exp(2\pi k't)x$. This measure has exactly the same support as x . As a result, we can repeat the same argument as above to conclude that

$$0 = \langle P_{\text{prony}}, \exp(2\pi k't)x \rangle \quad (44)$$

$$= \langle v, \hat{x}_{k'} \rangle \quad (45)$$

$$= \sum_{l=0}^s v_l \overline{y(k+k')} \quad \text{if } s+k' \leq n-1 \quad (46)$$

$$= 0. \quad (47)$$

If we select k' to be between 0 and $s-1$ we obtain a system of s equations. This system only involves the data y and the unknown coefficients v as long as $n \geq 2s$.

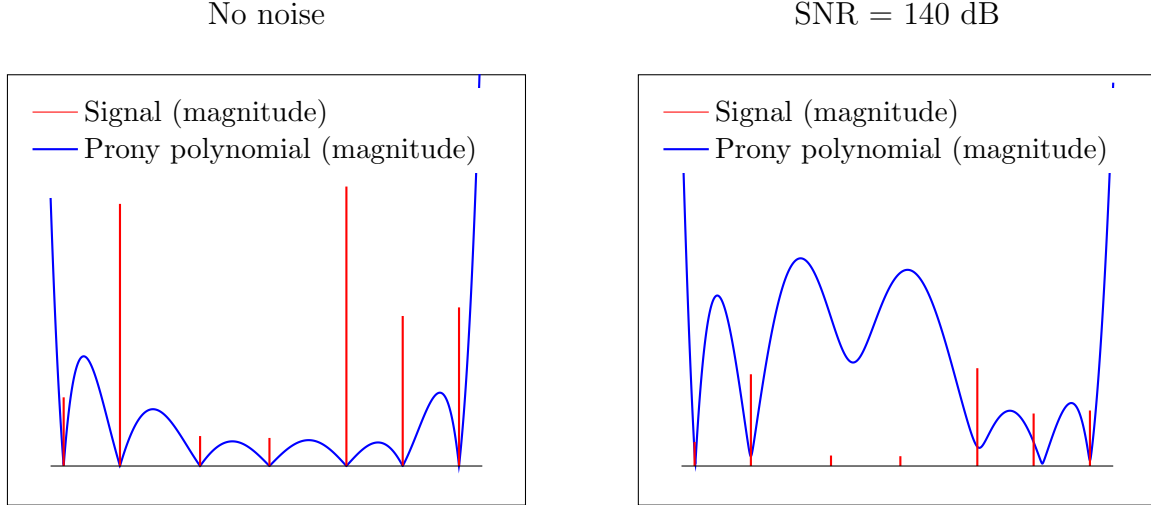


Figure 11: Prony polynomial applied on noiseless data (left). The image on the right shows the effect of adding a very small quantity of noise to the data. The roots of the polynomial no longer coincide with the support of the original signal. Note that the vertical axis is scaled differently in the two images.

Prony’s method consists of solving this system and then decoding the support of the signal by rooting the corresponding polynomial. For simplicity we assume that $n = 2s$.

Algorithm 4.2 (Prony’s method). *The input to the algorithm is the cardinality of the support s and the data y , which are assumed to be of the form (33).*

1. Form the system of equations

$$\begin{bmatrix} y_1 & y_2 & \cdots & y_s \\ y_2 & y_3 & \cdots & y_{s+1} \\ \cdots & \cdots & \cdots & \cdots \\ y_s & y_{s+1} & \cdots & y_{n-1} \end{bmatrix} \begin{bmatrix} \tilde{v}_1 \\ \tilde{v}_2 \\ \cdots \\ \tilde{v}_s \end{bmatrix} = - \begin{bmatrix} y_0 \\ y_1 \\ \cdots \\ y_{s-1} \end{bmatrix}. \quad (48)$$

2. Solve the system to obtain v_1, \dots, v_s and set $v_0 = 1$.
3. Root the polynomial corresponding to v_0, \dots, v_s to obtain its s roots z_1, \dots, z_s .
4. For every root on the unit circle $z_j = \exp(i2\pi\tau)$ include τ in the estimated support.

This procedure is guaranteed to achieve exact recovery of the original signal.

Lemma 4.3. *In the absence of noise, the output of Prony’s method is equal to the support of the original signal.*

Proof. The coefficients of the polynomial (35) are a feasible solution for the system of equations (48). In fact, they are the unique solution. To show this we compute the factorization

$$\begin{aligned}
 & \begin{bmatrix} y_1 & y_2 & \cdots & y_s \\ y_2 & y_3 & \cdots & y_{s+1} \\ \cdots & \cdots & \cdots & \cdots \\ y_s & y_{s+1} & \cdots & y_{n-1} \end{bmatrix} = & (49) \\
 & \begin{bmatrix} e^{-i2\pi t_1} & e^{-i2\pi t_2} & \cdots & e^{-i2\pi t_s} \\ e^{-i2\pi 2t_1} & e^{-i2\pi 2t_2} & \cdots & e^{-i2\pi 2t_s} \\ \cdots & \cdots & \cdots & \cdots \\ e^{-i2\pi s t_1} & e^{-i2\pi s t_2} & \cdots & e^{-i2\pi s t_s} \end{bmatrix} \begin{bmatrix} c_1 & 0 & \cdots & 0 \\ 0 & c_2 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & c_s \end{bmatrix} \begin{bmatrix} 1 & e^{-i2\pi t_1} & \cdots & e^{-i2\pi(s-1)t_1} \\ 1 & e^{-i2\pi t_2} & \cdots & e^{-i2\pi(s-1)t_2} \\ \cdots & \cdots & \cdots & \cdots \\ 1 & e^{-i2\pi t_s} & \cdots & e^{-i2\pi(s-1)t_s} \end{bmatrix}. & (50)
 \end{aligned}$$

The diagonal matrix is full rank as long as all the coefficients c_j are nonzero, whereas the two remaining matrices are full rank by the following lemma, proved in Section A.1 of the appendix (set $z_j := \exp(-i2\pi t_j)$).

Lemma 4.4 (Vandermonde matrix). *For any distinct set of s nonzero complex numbers z_1, z_2, \dots, z_s and any positive integers m_1, m_2, s such that $m_2 - m_1 + 1 \geq s$ the Vandermonde matrix*

$$\begin{bmatrix} z_1^{m_1} & z_2^{m_1} & \cdots & z_s^{m_1} \\ z_1^{m_1+1} & z_2^{m_1+1} & \cdots & z_s^{m_1+1} \\ z_1^{m_1+2} & z_2^{m_1+2} & \cdots & z_s^{m_1+2} \\ \cdots & \cdots & \cdots & \cdots \\ z_1^{m_2} & z_2^{m_2} & \cdots & z_s^{m_2} \end{bmatrix} \quad (51)$$

is full rank.

As a result, the matrix in (48) is full rank, so the system of equations has a unique solution equal to (35). This completes the proof, as rooting (35) obviously yields the support of the signal. \square

Unfortunately, Prony's method as presented above cannot be applied to real data even if the signal-to-noise ratio is exceptionally high. The image on the left of Figure 11 shows how the Prony polynomial allows to super-resolve the support to very high accuracy from noiseless

data. However, on the right we see the result of applying the method to data that have a very small amount of noise (the ratio between the ℓ_2 norm of the noise and the noiseless data is around 10^{-8} !). The roots of the Prony polynomial are perturbed away from the points of the unit circle that correspond to the support, so that it is no longer possible to accurately locate the support of the signal. It is consequently necessary to adapt the method to deal with noisy data if there is to be any hope of applying it in any realistic scenario. The following sections discuss such extensions.

4.2 Subspace methods

In order to motivate subspace-based methods to tackle noisy data, we first describe an alternative interpretation of Prony's method.

Solving the system (48) is essentially a way to find a nonzero vector in the null space of $Y(s+1)^T$, where $Y(m)$ is defined for any integer m as the Hankel matrix

$$Y(m) := \begin{bmatrix} y_0 & y_1 & \cdots & y_{n-m} \\ y_1 & y_2 & \cdots & y_{n-m+1} \\ \cdots & \cdots & \cdots & \cdots \\ y_{m-1} & y_m & \cdots & y_{n-1} \end{bmatrix}. \quad (52)$$

The vector in the null space of $Y(s+1)^T$ corresponds to the coefficients of the Prony polynomial, which we can root to find the support of the signal.

Recall that in the absence of noise the data are of the form

$$\begin{aligned} y_k &:= \widehat{x}(k) & (53) \\ &= \sum_{t_j \in T} c_j e^{-i2\pi k t_j}, & 0 \leq k \leq n-1. & (54) \end{aligned}$$

Consider the decomposition

$$\begin{aligned} Y(m) &= \begin{bmatrix} a_{0:m-1}(t_1) & a_{0:m-1}(t_2) & \cdots & a_{0:m-1}(t_s) \end{bmatrix} \begin{bmatrix} c_1 & 0 & \cdots & 0 \\ 0 & c_2 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & c_s \end{bmatrix} \begin{bmatrix} a_{0:n-m}(t_1)^T \\ a_{0:n-m}(t_2)^T \\ \cdots \\ a_{0:n-m}(t_s)^T \end{bmatrix} \\ &= A_{0:m-1}(T) C A_{0:m}(T)^T, & (55) \end{aligned}$$

where for $k > 0$ we define

$$a_{0:k}(t) := \begin{bmatrix} 1 \\ e^{-i2\pi t} \\ e^{-i2\pi 2t} \\ \dots \\ e^{-i2\pi kt} \end{bmatrix}, \quad A_{0:k}(T) := \begin{bmatrix} a_{0:k}(t_1) & a_{0:k}(t_2) & \dots & a_{0:k}(t_s) \end{bmatrix}. \quad (56)$$

This decomposition suggests an alternative way of estimating the support T from $Y(m)$: finding sinusoidal atoms $a_{0:m-1}(t)$ that are in the column space of $Y(m+1)$. Lemma 4.5 below proves that the only atoms of this form that belong to the column space of $Y(m)$ are precisely $a_{0:m-1}(t_1), a_{0:m-1}(t_2), \dots, a_{0:m-1}(t_s)$.

In order to make this procedure robust to noise, in practice we check what atoms are *close* to the column space of $Y(m)$. To quantify this we compute the orthogonal complement \mathcal{N} of the column space of $Y(m)$ and construct the *pseudospectrum*

$$P_{\mathcal{N}}(t) = \log \frac{1}{|\mathcal{P}_{\mathcal{N}}(a_{0:m-1}(t))|^2}, \quad (57)$$

where $\mathcal{P}_{\mathcal{N}}$ denotes a projection onto \mathcal{N} . If an atom is almost orthogonal to \mathcal{N} then $P_{\mathcal{N}}$ will have a very large value at that point.

In order to find the subspace \mathcal{N} , one possibility is computing the null space of the empirical covariance matrix

$$\Sigma(m) = \frac{1}{n-m+1} Y Y^* \quad (58)$$

$$= \frac{1}{n-m+1} \sum_{j=0}^{n-m} \begin{bmatrix} y_j \\ y_{j+1} \\ \dots \\ y_{j+m-1} \end{bmatrix} [\bar{y}_j \quad \bar{y}_{j+1} \quad \dots \quad \bar{y}_{j+m-1}]. \quad (59)$$

The following lemma shows that the maxima of the pseudospectrum reveal the support of the signal in the noiseless case.

Lemma 4.5. *Let \mathcal{N} be the null space of the empirical covariance matrix $\Sigma(m)$ for $m \geq s$. If the data are of the form (53),*

$$P_{\mathcal{N}}(t_j) = \infty, \quad \text{for } t_j \in T, \quad (60)$$

$$P_{\mathcal{N}}(t) < \infty, \quad \text{for } t \notin T. \quad (61)$$




Figure 12: Pseudospectrum corresponding to the data used to construct the Prony polynomials in Figure 11. The true location of the support is marked with red circles.

Proof. By (55) the atoms $a_{0:m-1}(t_1), \dots, a_{0:m-1}(t_s)$ span the column space of $Y(m)$ and $\Sigma(m)$. As a result they are orthogonal to the null space \mathcal{N} of the empirical covariance matrix, which proves (60).

We prove (61) by contradiction. The atoms $a_{0:m-1}(t_1), \dots, a_{0:m-1}(t_s)$ span the orthogonal complement to \mathcal{N} . As a result, if $a_{0:m-1}(t)$ is orthogonal to \mathcal{N} for some t then $a_{0:m-1}(t)$ is in the span of $a_{0:m-1}(t_1), \dots, a_{0:m-1}(t_s)$. This would imply that $A_{0:m-1}^T(T \cup \{t\})$ is not full rank, which can only hold if $t \in T$ by Lemma 4.4. \square

Figure 12 shows the pseudospectrum corresponding to the data used to construct the Prony polynomials in Figure 11 when $n = 2s - 1$. In the noiseless case, the pseudospectrum allows to locate the support T perfectly. Unfortunately, the locations of the local maxima are severely perturbed by even a very small amount of noise.

In order to obtain an estimate that is robust to noise, we need to use more data to estimate the support of the signal. The definition of the empirical covariance matrix $\Sigma(m)$ in (58) provides a way to do this, we can just average over more data. If we fix m and increase n , the column space of $\Sigma(m)$ remains the same, but the averaging process may cancel out the noise to some extent. This is the principle underlying the multiple-signal classification (MUSIC) [1, 11].

Algorithm 4.6 (Multiple-signal classification (MUSIC)). *The input is the cardinality of the support s , the data y , which are assumed to be generated as in (66) and the value of the parameter $m \geq s$.*

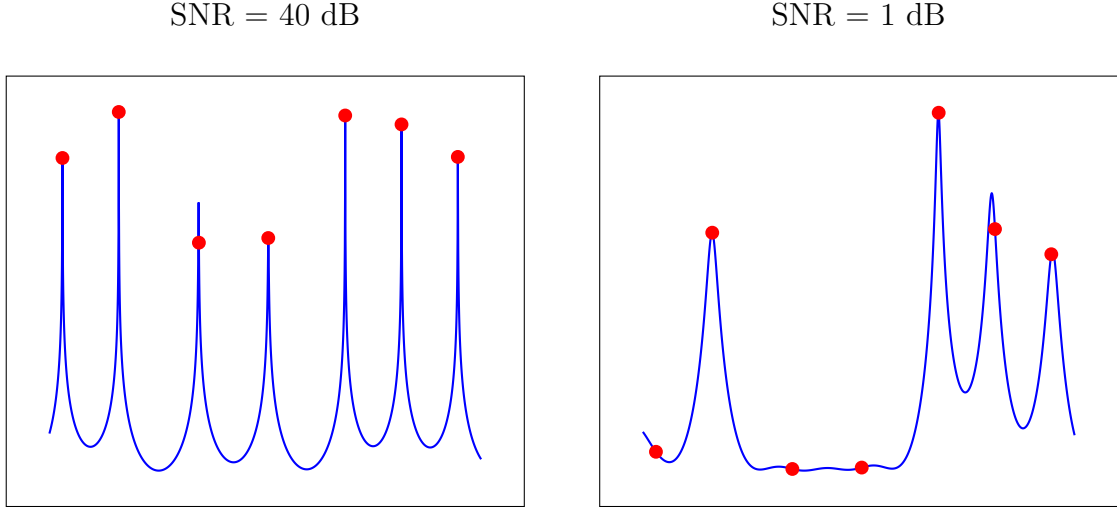


Figure 13: Pseudospectrum constructed by the MUSIC algorithm with $n = 81$ and $m = 30$ for the same signal used in Figure 11 and different noise levels. The true location of the support is marked with red circles.

1. Build the empirical covariance matrix $\Sigma(m)$ defined in (58).
2. Compute the eigendecomposition of $\Sigma(m)$ to select the subspace \mathcal{N} corresponding to the $m - s$ smallest eigenvalues.
3. Output an estimate of s estimated positions for the support by locating the s highest peaks of the pseudospectrum

$$P_{\mathcal{N}}(t) = \log \frac{1}{|\mathcal{P}_{\mathcal{N}}(a_{0:m-1})|^2}, \quad (62)$$

By Lemma 4.5, in the absence of noise MUSIC allows to estimate the support of the signal perfectly. When additive Gaussian noise is present in the data, MUSIC is much more robust than Prony's method. Figure 13 shows the result of applying MUSIC algorithm with $n = 81$ and $m = 30$ to the same data used in Figure 11 and different noise levels. The method is able to locate the support of the signal at a noise level of 40 dB. At 1 dB the pseudospectrum does not detect the smaller spikes (the true magnitudes are shown in Figure 11), but the estimate for the rest is still rather accurate.

In order to provide a theoretical justification of why MUSIC is stable we study the method in an asymptotic regime where the two following assumptions on the signal and the noise are met:

- **Assumption 1:** The original signal,

$$x = \sum_{t_j \in T} c_j \delta_{t_j} = \sum_{t_j \in T} |c_j| \exp(i\phi_j) \delta_{t_j}, \quad (63)$$

follows a probabilistic model where the phases ϕ_j are independent and uniformly distributed in the interval $[0, 2\pi]$, whereas the amplitudes are arbitrary and deterministic. Note that this implies that the expectation of x is equal to zero and that the covariance matrix equals

$$\mathbb{E}[cc^*] = D_c := \begin{bmatrix} |c_1|^2 & 0 & \cdots & 0 \\ 0 & |c_2|^2 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & |c_s|^2 \end{bmatrix}. \quad (64)$$

- **Assumption 2:** The measurements are corrupted by white Gaussian noise with zero mean and standard deviation σ , which is independent from the signal. At time k we measure

$$y_k := \int_0^1 \exp(-i2\pi kt) x(dt) + z_k \quad (65)$$

$$= \sum_{t_j \in T} c_j \exp(-i2\pi kt_j) + z_k, \quad (66)$$

where $z \sim N(0, \sigma^2 I)$ is a zero-mean Gaussian random vector with covariance matrix $\sigma^2 I$.

Under these assumptions, a sequence of m measurements y_0, y_1, \dots, y_{m-1} is a random vector of the form

$$\tilde{y} = A_{0:m-1}(T) c + z, \quad (67)$$

The following proposition, proved in Section A.2 of the appendix, characterizes the covariance matrix of \tilde{y} .

Proposition 4.7. *Let \tilde{y} be an m dimensional vector of data satisfying Assumptions 1 and 2. The eigendecomposition of the covariance matrix of \tilde{y} is equal to*

$$\mathbb{E}[\tilde{y}\tilde{y}^*] = \begin{bmatrix} U_S & U_N \end{bmatrix} \begin{bmatrix} \Lambda + \sigma^2 I_s & 0 \\ 0 & \sigma^2 I_{n-s} \end{bmatrix} \begin{bmatrix} U_S^* \\ U_N^* \end{bmatrix}. \quad (68)$$

The eigenvectors are divided into two unitary matrices.

- $U_S \in \mathbb{C}^{m \times s}$ contains an orthonormal basis of the signal subspace which corresponds to the span of $a_{0:m-1}(t_1), \dots, a_{0:m-1}(t_s)$.




Figure 14: Result of applying MUSIC to data with an SNR of 20 dB for two different minimum separations. In both cases, $f_c = 40$ (so $n = 81$) and $m = 40$.

- $U_{\mathcal{N}} \in \mathbb{C}^{m \times (m-s)}$ is a unitary matrix spanning the noise subspace, which is the orthogonal complement of the signal subspace.

In addition, all eigenvalues are positive: $\sigma^2 I_k \in \mathbb{C}^{k \times k}$ is a diagonal matrix with diagonal entries equal to σ^2 and Λ is another diagonal matrix with positive entries.

This proposition provides a rather intuitive interpretation of the MUSIC algorithm. The eigendecomposition of the covariance matrix of the data allows to estimate a signal subspace and a noise subspace. As a result, the term *subspace methods* is often used to describe MUSIC and related algorithms. Computing the pseudospectrum from these subspaces allows us to locate the support of the signal.

In practice, we approximate the covariance matrix using the empirical covariance matrix $\Sigma(m)$ defined in (58). Asymptotically, if we fix s and m and let $n \rightarrow \infty$, $\Sigma(m)$ converges to the true covariance matrix (see Section 4.9.1 in [13]). However that this does *not* necessarily imply that MUSIC will allow to find the support! To ensure that we can actually identify the noise subspace correctly, the eigenvalues in Λ must all be large with respect to the variance of the noise σ^2 . In the case of signals that have a small separation (with respect to n), some of these eigenvalues may be small due to the correlation between the atoms $a_{0:m-1}(t_1), \dots, a_{0:m-1}(t_s)$. This is not surprising because at small separations the inverse problem is ill posed, as explained in Section 2. Figure 14 shows the result of applying MUSIC to data with the same noise level for two different minimum separations. When the minimum separation is too small the method fails to detect some of the smaller spikes.

Figure 15 compares the performance of MUSIC for different noise levels and different values of the parameter m . On the left column, we see the decay of the eigenvalues of the empirical

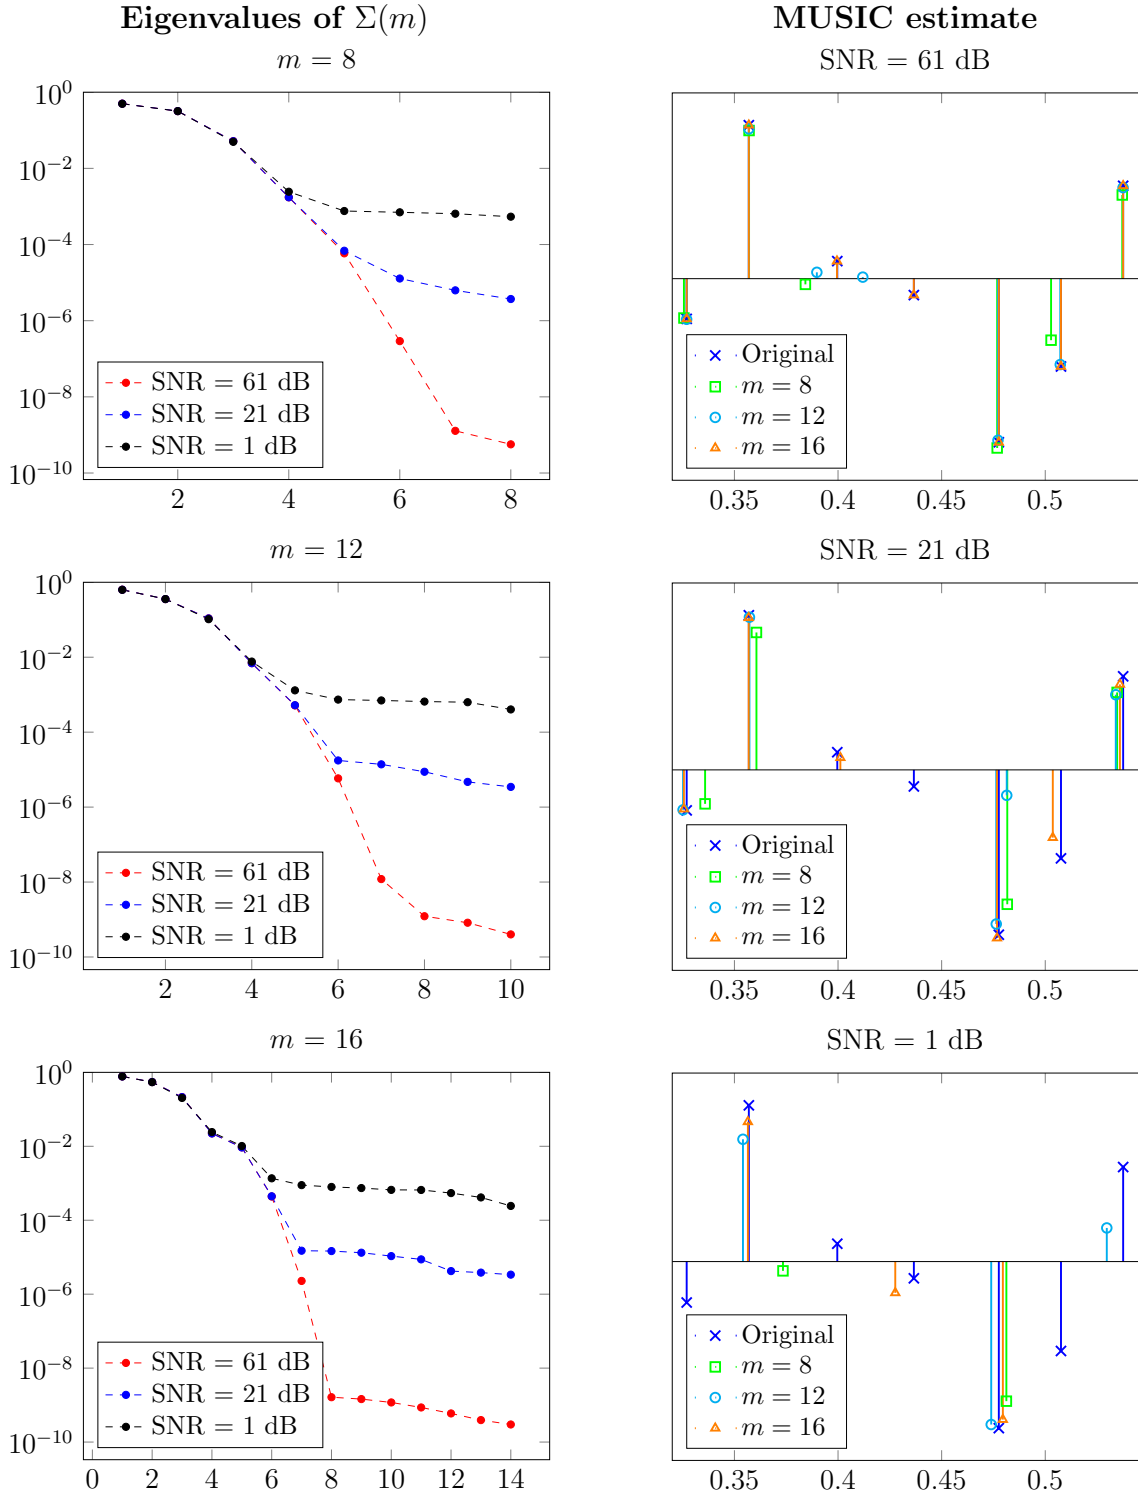


Figure 15: Eigenvalues of the empirical covariance matrix $\Sigma(m)$ used by MUSIC (left) and corresponding estimates (right) for different values of the parameter m and of the SNR. The cardinality of the true support is $s = 7$.




Figure 16: Line-spectra estimates obtained by Root MUSIC when the estimated number of sources is equal to $s - 1$ (left) and $s + 1$ (right) for the same data as in Figure 15.

covariance matrix. At high signal-to-noise ratios (SNR) there is a clear transition between the eigenvalues corresponding to the signal subspace (in this case $s = 7$) and the noise subspace, but this is no longer necessarily the case when the noise is increased. On the right column, we see the performance of the algorithm for different values of the SNR and the parameter m . At relatively high SNRs MUSIC is an effective algorithm as long as the assumptions on the signal (random phases), noise (Gaussian) and measurement model (equispaced time samples) are satisfied. In Figure 16 we show the result of running the algorithm for the wrong value of the parameter s . If the value is not too different to s and the SNR not too low, the method is still capable of approximately locating the support.

4.3 Matrix-pencil methods

In this section we describe an alternative approach to perform spectral super-resolution. In the previous section, we saw that MUSIC exploits the fact that in the absence of noise the matrix

$$Y_0 = \begin{bmatrix} y_0 & y_1 & \cdots & y_{n-m} \\ y_1 & y_2 & \cdots & y_{n-m+1} \\ \cdots & \cdots & \cdots & \cdots \\ y_{m-1} & y_m & \cdots & y_{n-1} \end{bmatrix} \quad (69)$$

has the rank- s factorization

$$Y_0 = A_{0:m-1}(T) C A_{0:n-m}(T)^T \quad (70)$$

$$= \sum_{t_j \in T} c_j a_{0:m-1}(t_j) a_{0:n-m}(t_j)^T \quad (71)$$

by computing the eigendecomposition of the empirical covariance matrix. Recall that

$$A_{m_1:m_2}(T) := [a_{m_1:m_2}(t_1) \ a_{m_1:m_2}(t_2) \ \cdots \ a_{m_1:m_2}(t_s)]. \quad (72)$$

Matrix-pencil methods make use of this factorization in a different way.

Let us begin by defining a matrix pencil and its rank-reducing values.

Definition 4.8 (Matrix pencil). *The matrix pencil of two matrices M_1 and M_2 is the matrix-valued function*

$$L_{M_1, M_2}(\mu) := M_2 - \mu M_1 \quad (73)$$

where $\mu \in \mathbb{C}$.

The set of rank-reducing values \mathcal{R} of a matrix pencil satisfy

$$\text{rank}(L_{M_1, M_2}(\mu)) = \text{rank}(L_{M_1, M_2}(\mu_j)) + 1 \quad \text{for all } \mu_j \in \mathcal{R} \text{ and } \mu \notin \mathcal{R}. \quad (74)$$

Now consider the matrix

$$Y_1 = \begin{bmatrix} y_1 & y_2 & \cdots & y_{n-m+1} \\ y_2 & y_3 & \cdots & y_{n-m+2} \\ \cdots & \cdots & \cdots & \cdots \\ y_m & y_{m+1} & \cdots & y_n \end{bmatrix} \quad (75)$$

which can be written as

$$Y_1 = A_{1:m}(T) C A_{0:n-m}(T)^T \quad (76)$$

$$= \sum_{t_j \in T} c_j a_{1:m}(t_j) a_{0:n-m}(t_j)^T. \quad (77)$$

The insight behind matrix-pencil methods for spectral super-resolution is that $\exp(i2\pi t_1)$, $\exp(i2\pi t_2)$, \dots , $\exp(i2\pi t_s)$ are all rank-reducing values of the matrix pencil L_{Y_0, Y_1} .

Lemma 4.9. *$\exp(i2\pi\tau)$ is a rank-reducing value of L_{Y_0, Y_1} if and only if $\tau \in T$.*

This lemma, proved in Section A.3 of the appendix implies that we can estimate the support of the signal by finding the rank-reducing numbers of the matrix pencil L_{Y_0, Y_1} . The following lemma, proved in Section A.4 of the appendix, shows how to compute them.

Lemma 4.10. *Let $Y_0 = U_0 \Sigma_0 V_0^*$ be the singular-value decomposition of Y_0 . The s eigenvalues of the matrix*

$$M = V_0 \Sigma_0^{-1} U_0^* Y_1 \quad (78)$$

are equal to $\exp(i2\pi t_j)$ for $1 \leq j \leq s$.

The matrix-pencil algorithm [7] estimates the rank-reducing values from the noisy data, hoping that the noise will not disrupt the eigenvalues of M too much.

Algorithm 4.11 (Matrix-pencil method). *The input is the cardinality of the support s , the data y , which are assumed to be generated as in (66) and the value of the parameter m .*

1. *Build the matrices Y_0 and Y_1 .*
2. *Compute the singular-value decomposition of $Y_0 = U_0 \Sigma_0 V_0^*$.*
3. *Compute the s largest eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_s$ of the matrix*

$$M = V_0 \Sigma_0^{-1} U_0^* Y_1. \quad (79)$$

4. *Output τ such that $\lambda_j = \exp(i2\pi\tau)$ for $1 \leq j \leq s$.*

For more information we refer the interested reader to Chapter 4 of [13], where it is reported that matrix-pencil algorithms exhibit a similar performance to subspace methods. For theoretical guarantees under a minimum-separation condition see [9].

5 Super-resolution via convex programming

In this section we describe how to perform super-resolution by penalizing a sparsity-inducing norm. This is reminiscent of compressed sensing, which consists of estimating sparse signals from randomized measurements via ℓ_1 -norm minimization. As we have explained in the previous sections, in super-resolution problem the signal is of the form

$$x := \sum_{t_j \in T} c_j \delta_{t_j}, \quad (80)$$

and the measurements are given by

$$y = \mathcal{F}_n x. \quad (81)$$

Since the signal can be supported anywhere in the unit interval, we cannot use the ℓ_1 norm to induce sparsity, as it is only defined for finite-dimensional vectors. In order to promote




Figure 17: Results of applying ℓ_2 and ℓ_1 -norm minimization to the inverse problem described in Figure 5. ℓ_1 -norm minimization achieves exact recovery.

sparsity within a continuous interval, we resort to a continuous counterpart of the ℓ_1 norm: the total-variation (TV) norm. The total-variation norm of a measure x is defined as

$$\|x\|_{\text{TV}} = \sup_{\|f\|_{\infty} \leq 1, f \in C(\mathbb{T})} \int_{\mathbb{T}} \overline{f(t)} x(dt). \quad (82)$$

For a superposition of Dirac deltas $\sum_j a_j \delta_{t_j}$, this norm is equal to the ℓ_1 norm of the coefficients, i.e. $\|x\|_{\text{TV}} = \sum_j |a_j|$. The term total-variation norm is somewhat unfortunate because *total variation* may also refer to the ℓ_1 norm of the discontinuities of a piecewise-constant function, as we have seen in previous lectures.

Super-resolution via TV-norm minimization requires solving the optimization problem

$$\text{minimize} \quad \|\tilde{x}\|_{\text{TV}} \quad (83)$$

$$\text{subject to} \quad \mathcal{F}_n \tilde{x} = y, \quad (84)$$

where the minimization is carried out over the set of all finite complex measures \tilde{x} supported on $[0, 1]$. This infinite-dimensional convex problem can be solved by recasting it into a finite-dimensional problem involving matrices (more specifically a semidefinite program), as we will see later on in the course.

If the support of the signal is restricted to lie on a grid, problem 83 is equivalent to ℓ_1 -norm minimization, exactly as in compressed sensing. Figure 17 shows the result of applying this algorithm to the reflection-seismography inverse problem described in Figure 5, where ℓ_1 -norm minimization achieves exact recovery. In fact, it was geophysicists working in this area who first suggested using ℓ_1 -norm minimization to tackle underdetermined linear inverse problems. Figure 18 shows results on real seismography data from a paper in 1979 [14]. These results spurred the interest of statisticians like David Donoho, who would eventually spearhead the use of ℓ_1 -norm minimization in other areas of signal processing and compressed sensing.




Figure 18: Result of applying ℓ_1 -norm minimization to real seismography data from a paper in 1979 [14].




Figure 19: Interpolating kernel K used in the proof of Theorem 5.1 (left), along with its asymptotic decay (center) and its low-pass spectrum (right).

5.1 Exact recovery

The following theorem establishes that TV-norm minimization succeeds for signals that have a minimum separation of $1.26/f_c$, very near the limit of $1/f_c$ at which the problem becomes ill posed.

Theorem 5.1 (Total-variation minimization). *Let T be the support of a signal x . If the minimum separation obeys*

$$\Delta(T) \geq 1.26 / f_c \tag{85}$$

then x is the unique solution to (83).

The proof of this theorem is based on the construction of a dual certificate for the TV norm. To simplify matters, while still conveying the main ideas, let us assume that the signal lies on a grid and we are minimizing the ℓ_1 norm. In that case, as we saw in the lecture on compressed sensing, a dual certificate is a vector v such that $q = \mathcal{F}_n^* v$ interpolates the sign pattern of the signal on its support and has magnitude smaller than one on its off-support. In compressed sensing q is a superposition of random vectors, whereas in super-resolution q is a superposition of low-pass trigonometric polynomials (the rows of the linear operator \mathcal{F}_n). This difference is crucial: if the support of the signal is cluttered together and the sign of its coefficients varies rapidly, it may not be possible to achieve the interpolation with a bounded polynomial.

We will sketch the proof of Theorem 5.1, which can be found in [2, 4]. The construction of dual certificates for compressed sensing typically relies on concentration arguments from probability theory. However there is nothing random about Theorem 5.1. We resort to a different approach: interpolating the sign of the signal using a low-pass kernel K to construct

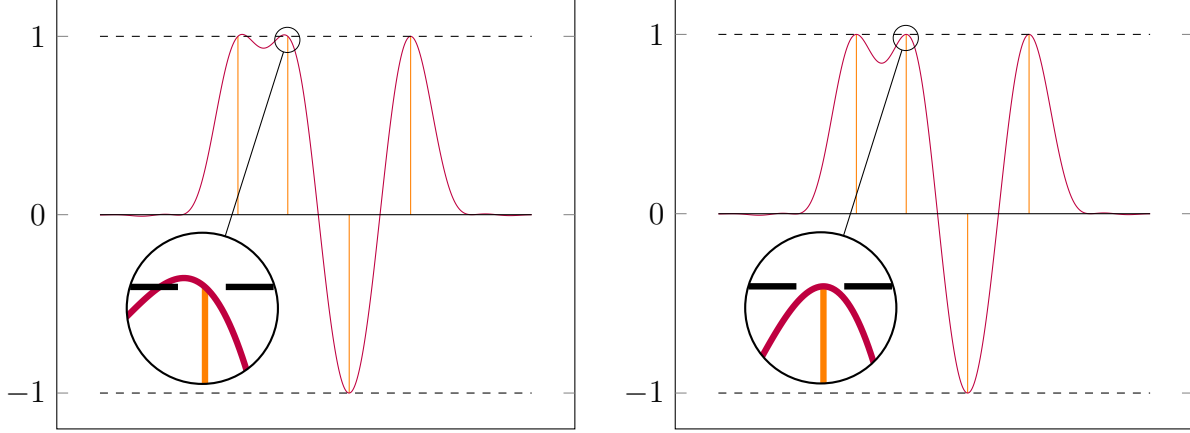


Figure 20: Dual-polynomial if we only use K to interpolate the sign pattern (left), as opposed to also incorporating its derivative into the construction and forcing the derivative of the polynomial to be zero at the interpolation points (right).

q directly. Since any shift of a low-pass kernel is low pass, this means that $q = \mathcal{F}_n^* v$ for some $v \in \mathbb{C}^n$. Figure 19 shows the kernel and its spectrum.

Interpolating the sign pattern with an adequately selected kernel is not sufficient to produce a valid construction. The reason is that the magnitude of the polynomial tends to exceed one near the elements of the support T . This can be avoided, however, by forcing the derivative of the kernel to be zero at those points. As a result, the magnitude of the polynomial has a local extremum at the interpolation point and its magnitude remains below one for a sufficiently large minimum separation, as illustrated by Figure 20. In order to enforce the extra constraint on the derivative of q , we need more degrees of freedom in the construction. For this purpose, we incorporate the derivative of the kernel, so that the q is of the form

$$q(t) = \sum_{t_j \in T} \alpha_j K_\gamma(t - t_j) + \beta_j K'_\gamma(t - t_j), \quad (86)$$

where $\alpha, \beta \in \mathbb{C}^{|T|}$ are coefficient sequences satisfying

$$q(t_k) = \sum_{t_j \in T} \alpha_j K_\gamma(t_k - t_j) + \beta_j K'_\gamma(t_k - t_j) = v_j, \quad t_k \in T, \quad (87)$$

$$q'_R(t_k) + iq'_I(t_k) = \sum_{t_j \in T} \alpha_j K'_\gamma(t_k - t_j) + \beta_j K''_\gamma(t_k - t_j) = 0, \quad t_k \in T, \quad (88)$$

and q_R and q_I are the real and imaginary parts of q . Since the derivative of a low-pass function is still low pass, q is in the row space of the linear operator \mathcal{F}_c . All is left is to show that the system of equations has a solution and that the magnitude of q is indeed bounded on the complement of T . We refer the interested reader to [2, 4] for the gory details.




Figure 21: Result of applying TV-norm minimization to super-resolve data with an SNR of 20 dB for two different minimum separations. In both cases, $f_c = 40$.

5.2 Super-resolution from noisy data

Consider the additive-noise model

$$y = \mathcal{F}_n x + z \quad (89)$$

where z is a noise term. In order to account for the perturbation, we can adapt Problem (83) by using an inequality constraint to quantify the uncertainty,

$$\min_{\tilde{x}} \|\tilde{x}\|_{\text{TV}} \quad \text{subject to} \quad \|\mathcal{F}_n \tilde{x} - y\|_2^2 \leq \delta, \quad (90)$$

where δ is an estimate of the noise level. Alternatively, we could also consider a Lagrangian formulation of the form

$$\min_{\tilde{x}} \|\tilde{x}\|_{\text{TV}} + \gamma \|\mathcal{F}_n \tilde{x} - y\|_2^2, \quad (91)$$

where the regularization parameter $\gamma > 0$ governs the tradeoff between data fidelity and the sparsity of the estimate. Figure 21 shows the result of applying TV-norm minimization to super-resolve data with an SNR of 20 dB for two different minimum separations. For a theoretical analysis of the performance of Problem 90 see [3].

References

For more information on spectral super-resolution using the periodogram, subspace methods and algorithms based on matrix pencil we refer the reader to [13] and references therein.

- [1] G. Bienvenu. Influence of the spatial coherence of the background noise on high resolution passive methods. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, volume 4, pages 306 – 309, 1979.
- [2] E. J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956.
- [3] E. J. Candès and C. Fernandez-Granda. Super-resolution from noisy data. *Journal of Fourier Analysis and Applications*, 19(6):1229–1254, 2013.
- [4] C. Fernandez-Granda. Super-resolution of point sources via convex programming. *Information and Inference*, 2016.
- [5] W. T. Freeman, T. R. Jones, and E. C. Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 22(2):56–65, 2002.
- [6] F. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1):51 – 83, 1978.
- [7] Y. Hua and T. K. Sarkar. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 38(5):814–824, 1990.
- [8] J. Lindberg. Mathematical concepts of optical superresolution. *Journal of Optics*, 14(8):083001, 2012.
- [9] A. Moitra. Super-resolution, extremal functions and the condition number of Vandermonde matrices. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, 2015.
- [10] S. C. Park, M. K. Park, and M. G. Kang. Super-resolution image reconstruction: a technical overview. *IEEE Signal Processing Magazine*, 20(3):21 – 36, 2003.
- [11] R. Schmidt. Multiple emitter location and signal parameter estimation. *IEEE Transactions on Antennas and Propagation*, 34(3):276 – 280, 1986.
- [12] D. Slepian. Prolate spheroidal wave functions, Fourier analysis, and uncertainty. V - The discrete case. *Bell System Technical Journal*, 57:1371–1430, 1978.
- [13] P. Stoica and R. L. Moses. *Spectral Analysis of Signals*. Prentice Hall, 2005.
- [14] H. L. Taylor, S. C. Banks, and J. F. McCoy. Deconvolution with the ℓ_1 norm. *Geophysics*, 44(1):39–52, 1979.

A Proofs

A.1 Proof of Lemma 4.4

Let us define

$$Z := \begin{bmatrix} 1 & 1 & \cdots & 1 \\ z_1 & z_2 & \cdots & z_s \\ z_1^2 & z_2^2 & \cdots & z_s^2 \\ \cdots & \cdots & \cdots & \cdots \\ z_1^{s-1} & z_2^{s-1} & \cdots & z_s^{s-1} \end{bmatrix}. \quad (92)$$

The determinant of the first s rows of our matrix of interest is equal to

$$|Z| \prod_{1 \leq i \leq s} z_i = \prod_{1 \leq j < k \leq s} (z_j - z_k) \prod_{1 \leq i \leq s} z_i \neq 0 \quad (93)$$

This implies that the first s rows are linearly independent and consequently that the whole matrix is full rank.

A.2 Proof of Proposition 4.7

Due to the assumptions,

$$\mathbb{E}[yy^*] = \mathbb{E}[A_{0:m-1}cc^*A_{0:m-1}^* + A_{0:m-1}cz^* + zc^*A_{0:m-1}^* + zz^*] \quad (94)$$

$$= A_{0:m-1}\mathbb{E}[cc^*]A_{0:m-1}^* + A_{0:m-1}\mathbb{E}[c]\mathbb{E}[z^*] + \mathbb{E}[z]\mathbb{E}[c^*]A_{0:m-1}^* + \mathbb{E}[zz^*] \quad (95)$$

$$= A_{0:m-1}D_cA_{0:m-1}^* + \sigma^2I. \quad (96)$$

We begin by writing the full eigendecomposition of $A_{0:m-1}D_cA_{0:m-1}^*$, which is a positive semidefinite symmetric matrix of rank s . By the spectral theorem, this matrix has a singular value decomposition of the form

$$A_{0:m-1}D_cA_{0:m-1}^* = \begin{bmatrix} U_S & U_N \end{bmatrix} \begin{bmatrix} \Lambda & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} U_S^* \\ U_N^* \end{bmatrix}, \quad (97)$$

where U_S and U_N are as defined in the statement of the proposition. Λ is a diagonal matrix with positive entries, because $A_{0:m-1}D_cA_{0:m-1}^*$ is positive semidefinite. Indeed if we define

$\sqrt{D_c}$, a diagonal matrix with $|c|$ in the diagonal, we have that for any nonzero vector u

$$u^* A_{0:m-1} D_c A_{0:m-1}^* u = u^* A_{0:m-1} \sqrt{D_c} \sqrt{D_c} A_{0:m-1}^* u \quad (98)$$

$$= \left\| \sqrt{D_c} A_{0:m-1}^* u \right\|_2^2 \quad (99)$$

$$> 0, \quad (100)$$

where the inequality is strict because $\sqrt{D_c} A_{0:m-1}^*$ is full rank.

To complete the proof, we decompose the identity matrix using U_S and U_N to obtain

$$E[yy^*] = A_{0:m-1} D_c A_{0:m-1}^* + \sigma^2 I \quad (101)$$

$$= \begin{bmatrix} U_S & U_N \end{bmatrix} \begin{bmatrix} \Lambda & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} U_S^* \\ U_N^* \end{bmatrix} + \begin{bmatrix} U_S & U_N \end{bmatrix} \begin{bmatrix} \sigma^2 I_s & 0 \\ 0 & \sigma^2 I_{n-s} \end{bmatrix} \begin{bmatrix} U_S^* \\ U_N^* \end{bmatrix}. \quad (102)$$

A.3 Proof of Lemma 4.9

For any $\tau \in T$

$$a_{1:m}(\tau) = \exp(i2\pi\tau) a_{0:m-1}(\tau) \quad (103)$$

so

$$L_{Y_0, Y_1}(\exp(i2\pi\tau)) = Y_1 - \exp(i2\pi\tau) Y_0 \quad (104)$$

$$= \sum_{t_j \in T} c_j (a_{1:m}(t_j) - \exp(i2\pi\tau) a_{0:m-1}(t_j)) a_{0:n-m}(t_j)^T \quad (105)$$

$$= \sum_{t_j \in T/\{\tau\}} c_j (a_{1:m}(t_j) - \exp(i2\pi\tau) a_{0:m-1}(t_j)) a_{0:n-m}(t_j)^T, \quad (106)$$

which has rank $s - 1$.

If $\tau \notin T$ the column space of $L_{Y_0, Y_1}(\exp(i2\pi\tau))$ is spanned by $\exp(i2\pi t_1), \dots, \exp(i2\pi t_s)$ and $\exp(i2\pi\tau)$. By Lemma 4.4 these vectors are linearly independent, so the matrix pencil has rank s .

A.4 Proof of Lemma 4.10

From (103) we have

$$A_{0:m-1}(T) = A_{0:m-1}(T) \Phi, \quad (107)$$

where

$$\Phi := \begin{bmatrix} \exp(i2\pi t_1) & 0 & \cdots & 0 \\ 0 & \exp(i2\pi t_2) & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \exp(i2\pi t_s) \end{bmatrix}. \quad (108)$$

Y_0 can be decomposed as

$$Y_0 = A_{0:m-1}(T) C A_{0:n-m}(T)^T. \quad (109)$$

Let $U \Sigma V^*$ be the singular value decomposition of $C A_{0:n-m}(T)^T$, where $U \in \mathbb{C}^{s \times s}$, $\Sigma \in \mathbb{C}^{s \times s}$ and $V \in \mathbb{C}^{n-m+1 \times s}$. If the coefficients c_1, \dots, c_s are nonzero then C is full rank. $A_{0:n-m}(T)^T$ is full rank by Lemma 4.4. This implies that $V \Sigma^{-1} U^*$ is a right inverse of $C A_{0:n-m}(T)^T$,

$$C A_{0:n-m}(T)^T V \Sigma^{-1} U^* = I. \quad (110)$$

We now use these facts to decompose Y_1 ,

$$Y_1 = A_{1:m}(T) C A_{0:n-m}(T)^T \quad (111)$$

$$= A_{0:m-1}(T) \Phi C A_{0:n-m}(T)^T \quad \text{by (107)} \quad (112)$$

$$= A_{0:m-1}(T) C A_{0:n-m}(T)^T V \Sigma^{-1} U^* \Phi C A_{0:n-m}(T)^T \quad \text{by (110)} \quad (113)$$

$$= Y_0 V \Sigma^{-1} U^* \Phi U \Sigma V^*. \quad (114)$$

The row space of Y_0 is the same as the row space of $C A_{0:n-m}(T)^T$ because $A_{0:m-1}(T)^T$ is full rank by Lemma 4.4. This implies that

$$V_0 \Sigma_0^{-1} U_0^* Y_0 V = V_0 V_0^* V \quad (115)$$

$$= V. \quad (116)$$

As a consequence

$$V_0 \Sigma_0^{-1} U_0^* Y_1 = V_0 \Sigma_0^{-1} U_0^* Y_0 V \Sigma^{-1} U^* \Phi U \Sigma V^* \quad (117)$$

$$= V \Sigma^{-1} U^* \Phi U \Sigma V^* \quad (118)$$

$$= P^{-1} \begin{bmatrix} \Phi & 0 \\ 0 & 0 \end{bmatrix} P, \quad (119)$$

where

$$P := \begin{bmatrix} U & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \Sigma & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} V^* \\ V_\perp^* \end{bmatrix} \quad (120)$$

and V_\perp is any orthogonal matrix whose column space is the orthogonal complement to the column space of V . This establishes that the nonzero eigenvalues of the matrix $V_0 \Sigma_0^{-1} U_0^* Y_1$ correspond to the diagonal entries of Φ .