

This is a repository copy of *Comparison of modelled pursuits with ESPRIT and the matrix pencil method in the modelling of medical percussion signals*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/208317/>

Version: Published Version

Article:

Brown, Kenneth Ian orcid.org/0000-0002-1840-6151 and Wells, Jez orcid.org/0000-0002-4159-7179 (2024) Comparison of modelled pursuits with ESPRIT and the matrix pencil method in the modelling of medical percussion signals. *Biomedical signal processing and control*. 105777. ISSN 1746-8094

<https://doi.org/10.1016/j.bspc.2023.105777>


Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Comparison of modelled pursuits with ESPRIT and the matrix pencil method in the modelling of medical percussion signals

Kenneth I. Brown^{*}, Jeremy J. Wells

University of York, Heslington, York YO10 5DD, UK

ARTICLE INFO

Keywords:

Medical percussion
EDS
Modelled pursuits
MoP
ESPRIT
MPM

ABSTRACT

The objective of this paper is to compare *Modelled Pursuits* (MoP), a recently developed iterative signal decomposition method, with more established matrix based subspace methods used to aid or automate medical percussion diagnoses.

Medical percussion is a technique used by clinicians to aid the diagnosis of pulmonary disease. It requires considerable expertise, so it is desirable to automate this process where possible. Previous work has examined the application of modal decomposition techniques, since medical percussion signals (MPS) can be intuitively characterised as combinations of exponentially decaying sinusoidal (EDS) vibrations. Best results have typically been reported with matrix based subspace methods such as *Estimation of Signal Parameters via Rotational Invariance Techniques* (ESPRIT) and the *Matrix Pencil Method* (MPM).

Since ESPRIT and MPM are computationally expensive, this paper investigates whether an iterative method such as MoP can produce similar results with less computation and/or memory overheads. Using randomly generated synthetic signals designed to replicate typical ‘tympanic’ and ‘resonant’ percussion signals, we compared each method: MoP, ESPRIT, and MPM, for accuracy, speed and memory usage.

We find that for low Signal to Noise Ratios (SNRs) MoP gives less accuracy than both ESPRIT and MPM, however for high SNRs (as would be typically encountered in a clinical setting) it is more accurate than MPM but less accurate than ESPRIT. We conclude that in embedded clinical applications where both operations-per-second and memory-usage are a factor, MoP is less computationally intensive than ESPRIT and thus is worth considering for use in those contexts.

1. Introduction

1.1. Focus of this paper

The specific focus of this paper is to introduce to the field of MPS modelling, a recently developed method, MoP, which was originally developed for resource-efficient audio analysis of impulse responses. We compare its effectiveness and resource requirements with two more established methods (ESPRIT and MPM) in the analysis of damped vibrational systems such as the human chest cavity with percussive input, i.e. Medical Percussion. In order to evaluate this over a significant number of runs, an approach to the parametrisation and synthesis of many thousands of synthetic medical percussion signals (MPS) in two broad clinical categories is described.

1.2. Medical percussion

For over 250 years medical percussion has been used to detect abnormalities within the human thorax and abdomen. Trained physicians are able to recognize many different kinds of percussion sounds, the main three of which are widely known as “resonant,” “tympanic,” and “dull” [1]. Changes in the presence of air or liquid affects the acoustic properties of the area being percussed, for instance over normal lung tissue, low-pitch, hollow *resonant* sounds are typically heard, whereas if the patient has, for example, pneumothorax (air outside the lungs) a more drum-like *tympanic* sound is heard.

The quality of diagnosis using manual percussion techniques remains highly variable, being dependent on the clinician’s skills, experience and subjective evaluation [2,3]. If the sound of the percussion is recorded and digitised resulting in a medical percussion signal, it should be possible to perform an automated, or assisted analysis with the aim of

^{*} Corresponding author.

E-mail address: kenneth.brown@york.ac.uk (K.I. Brown).

<https://doi.org/10.1016/j.bspc.2023.105777>

Received 12 June 2023; Received in revised form 26 October 2023; Accepted 17 November 2023

Available online 3 December 2023

1746-8094/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

reducing variability and subjectivity and improving the quality and efficiency of the diagnostic process.

1.3. Earlier work in MPS analysis

A variety of percussion signal capture and analysis techniques are described in the literature [e.g. 1,2,4]. These techniques model the captured medical percussion signals (which are the response of the body to the percussive input) as sums of exponentially decaying sinusoids (EDS), since this model is an intuitive and efficient way of describing the impulse responses (IRs) of vibrational systems.

More recently, there has been interest in the development of self-contained devices that mechanically produce their own medical percussion and can capture and analyse the resultant vibrations from the human body. In [5] a self-contained system for the analysis and classification of percussion signals in the non-invasive diagnosis of acute chest syndrome in sufferers of sickle cell disease is described. This uses a device (known as a Tabla [3]) to impart a percussion signal into the sternum. The transmitted signal is then picked up via an Eko digital stethoscope and analysed by an associated smartphone app [6]. In [7], a general system for automatic percussion, capture and analysis, *iApp* (*iApp: An Autonomous Inspection, Auscultation, Percussion, and Palpation Platform*) is proposed for telediagnosis. In [8] the effectiveness of a system based around a portable embedded computing system (the Nvidia Jetson) for classification of captured thoracic percussion response signals into categories of dull, resonant and tympanic is explored.

In general a trend can be observed in these examples and elsewhere, for deployment on processing and memory-constrained portable devices, such as smart phones, whose convenience and low-power requirements might be suitable in a much wider range of clinical scenarios, rather than fixed, specialist equipment. Because of this trend, it is important that computational efficiency with respect to processing speed and memory usage is maximised.

All the systems described above use some form of frequency, time-frequency or time-scale (wavelet) model in the pipeline between signal capture and classification. Modelling in this way converts a time-domain MPS into a smaller set of descriptive parameters that can later be used in classification to aid diagnosis.

1.4. Impulse responses and the exponentially damped sinusoidal [EDS] model

When considering medical percussion of the chest cavity, the cavity can be considered an unvarying resonator being excited by impulsive-like input. Any linear time-invariant vibrational system can be described by its impulse response. In a digital system a single non-zero sample is an impulse that contains equal energy at all frequencies up to the Nyquist limit. Medical percussion involves introducing energy into the bodily region of interest by tapping with the fingers [4]. A single tap is an impulse-like signal (although it is certainly not perfectly impulsive - it will likely be much longer in duration than a single non-zero sample and therefore will have less energy at high frequencies due to the integrating effect of that longer duration), and so is likely to excite all vibrational modes of the object, or objects, being investigated.

In a vibrational system, energy is dissipated via damping of the medium(s) through which that energy passes, and that damping effect can be frequency dependent. Where there is an impedance change then energy will be reflected or diffracted and so may pass through a single point many times. The amount of energy lost due to damping is directly proportional to the amount of energy remaining in the signal and so it is seen to decay exponentially over time (it is a recursive, i.e. feedback, process). The frequencies of vibrational modes depend on the dimensions of the vibrating object and the speed of sound through its medium(s). Therefore the impulse responses of vibrating objects can be intuitively characterised as a sum of exponentially decaying sinusoids at

different frequencies and with different amplitude and decay rates.

Such a sum $\tilde{s}(t)$ can be described for digital systems as:

$$\tilde{s}(t) = \sum_{m=1}^M e^{a_m + j\varphi_m - t(\alpha_m + j2\pi f_m)} + \varepsilon \quad (1)$$

where M is the total number of components in the model, a_m is the starting amplitude of the m th component, φ_m is the starting phase, α_m is the exponential amplitude change rate (decay-rate) in Nepers, f_m is the frequency of the component and ε is the experimental noise component. For digital signals and systems, which are the focus of this paper, time is represented by indices which are spaced by a time period of $1/F_s$, where F_s is the sample rate of the system. Throughout this paper we refer to such sampled signals as a vector (e.g. \mathbf{s} of length N , or its constituent samples s_n). The generic signal type described by equation (1) is referred to in the literature as the EDS model [9]. It can also be viewed as the IR of a single resonant filter where $e^{a+j\varphi}$ is the complex amplitude and $e^{(\alpha+j2\pi f)}$ is the complex pole of the filter.

By modelling a signal as a sum of EDS, i.e. estimating the modelled parameters from the noisy data (as described later in this paper), it is possible to transform the signal from a stream of time-domain samples to an array of parameters. The set of four parameters f , a , α and φ for a given n (representing a single exponentially decaying sinusoid) are referred to as an ‘atom’.

There exist different methods for obtaining these parameter estimates.

1.5. Summary of existing EDS parameter estimation methods

When seeking the combination of EDS that best fits an input signal, a typical approach is to use methods from linear algebra (LA), since the component atoms of the EDS sum are independent of each other, i.e. the addition or removal of one component does not affect the other components. For each atom, its EDS component signal at some time t can be completely determined if the atom’s parameter values are known. Generating a time domain signal from atom parameters is called *resynthesis*. Conversely, using LA, the signal data can be manipulated in matrix form such that it can be ‘solved’ (like a simultaneous equation) to produce its constituent atom parameters. In the presence of noise it is only possible to estimate the parameters rather than calculate them exactly. Various methods to do this have been presented, each having advantages and disadvantages in their estimation accuracy and efficiency. Matrix methods in general have computational complexity proportional to N^3 , where N is length of the signal. In this paper we discuss two LA methods: ESPRIT and MPM, which are computationally expensive but are capable of delivering excellent results in atomic decomposition of EDS [2,10]. These techniques are also applied in other fields and their development is an ongoing research activity in areas such as 6G communication [11] and radar [12].

There are methods that use a single orthogonal transform such as Fourier [2] or wavelet [8] to analyse the signal in one pass, but although these are cheap in terms of processing requirements, they typically perform poorly [13]. Matching Pursuits (MP) [4,14] is an iterative approach to signal decomposition that can use any dictionary (or dictionaries) of atoms. These dictionaries can be (and usually are) ‘over-complete’ (i.e. the number of dictionary elements is greater than the number of samples in the signal being analysed) in order to create a decomposition that is optimal in some way. The optimal decomposition is often deemed to be the one that is sparsest, since this suggests that its atoms are best matched to the signal. MP is an iterative process: the atom which has the most energetic correlation with the signal is chosen and subtracted from the signal, leaving a residual. This process is repeated on the residual and so on. Provided the signal is in the space spanned by the dictionary (guaranteed for over-complete dictionaries which contain at least one set of orthogonal elements) the algorithm converges (i.e. the residual energy approaches zero as the algorithm progresses). Although

it can yield an optimal (sparse) representation the iterative nature of the method means it is more computationally expensive than Fourier or wavelet methods.

Modelled Pursuits (MoP) [13] is an iterative method similar to Matching Pursuits which estimates the parameters of the highest magnitude component using an initial fast Fourier transform (FFT) of the zero-padded signal, synthesises the component in the time domain from those estimates, and subtracts it from the signal, leaving a residual. This residual is then analysed iteratively in the same way (requiring a new FFT to be calculated.) This iteration process is continued until a stopping criterion is reached, e.g. the number of atoms has reached some threshold, or the residual energy increases instead of decreasing. Each iteration produces an estimate of the frequency, amplitude, phase and decay rate of that iteration's component atom.

This 'maximum atoms' threshold is useful in the cases where the actual number of EDS in the original signal is known to be low. In these cases, MoP can significantly out-perform both LA based methods and dictionary-based iterative methods. This is particularly significant in the case of MPS analysis (see below.)

1.6. Contribution of this paper

It is of interest in many fields to understand the relative trade-offs in employing LA, Fourier or hybrid methods to solve signal analysis and processing problems, e.g. [11,12]. This paper examines how three EDS decomposition methods perform with MPS in the presence of different levels of noise. In [2] the MPM algorithm is compared with FFT-based algorithms for the modal analysis of MPS, and in previous work, the second author of this paper considered how MoP and ESPRIT compare for acoustic IRs [13].

In [4] the author states that 'a small number of low frequency atoms can be sufficient for classification of MPS'. In cases where it is known *a priori* that the number of EDS in the original signal is small, a modelling process can be used that only calculates the first few most significant atoms rather than the maximum possible by that modelling method. If, upon re-synthesis, such a short set of atoms yields a good approximation of the original signal (without noise), then that modelling process can be deemed to have efficiently and intuitively modelled the signal.

The specific novel contribution of the work reported here is to compare MPM, MoP and ESPRIT with the short duration, lower resolution signals which are typical of MPS. In order to do this in a robust, parametrised manner over many (nearly two million) tests, we develop a method for generating synthetic MPS which embody the statistics of real MPS described in existing literature. This enables rigorous benchmarking of a cheaper, in terms of computation and memory requirements, method against two particular examples of state of the art, but which are expensive matrix based approaches. We do this using synthetic signals which embody the qualities of clinical signals, whilst giving us sufficient experimental control. These synthetic signals are divided into two broad groups: 'tympanic' and 'percussive'.

The rest of the paper is organised as follows. In the next section established methods for EDS modelling of signals, including the recently introduced MoP method, are surveyed. Section 3 outlines the experimental procedure for comparing the methods described in the previous section. Section 4 presents the results from this procedure, in terms of both signal estimation quality, and computational cost in terms of execution time and memory requirements. In the final section the results are discussed and conclusions drawn about the potential suitability of MoP for analysis of medical percussion signals, compared with already established approaches.

2. Methods

2.1. MPM and ESPRIT: Background for exact algorithms used in this paper

MPM [15,16,17] and ESPRIT [18] have their basis in Prony's original technique for matrix decomposition of mixed EDS signals [19]. In their original references, one of the major differences was that MPM acted directly on the signal data, whereas ESPRIT first formed a covariance matrix and acted upon that. Subsequent developments of both methods have led to some degree of convergence especially in the use or not of Singular Value Decomposition (SVD), and in later papers there is no universal agreement as to exactly which algorithmic steps are used when MPM or ESPRIT are mentioned. The main difference in our MPM and ESPRIT implementations is the matrices from which the poles are estimated. For MPM, steps 5 and 6 (Algorithm 1, described in Section 2.2.1) operate on matrices derived from a Hankel matrix of the input signal. For ESPRIT, the equivalent steps (8 and 9 in Algorithm 2, described in Section 2.3.1) operate on matrices derived from the signal subspace, calculated as the singular value decomposition (SVD) of the input signal Hankel matrix (\mathbf{H}). The columns of the matrix output of this SVD are the eigenvectors of $\mathbf{H}\mathbf{H}^*$ which represents the covariance of the signal.

In this paper we have used the most recent and fullest documented algorithms for 'standard' MPM and ESPRIT methods that could be obtained following a literature review. Our algorithmic steps for each are shown below, and our Matlab source code is available [20] to allow readers to examine the detail of the exact algorithms we have used. For each algorithm we present the memory requirements for the input, output and internal structures and the computational complexity of each stage.

2.2. Matrix pencil method

The Matrix Pencil Method involves forming an $(N - L) \times (L + 1)$ Hankel matrix \mathbf{Y} from the noisy data vector length N , where L is the *pencil parameter*, described as the "free moving window length," that is "closely related to the *polynomial degree* or *polynomial prediction order*" [17].

From this matrix, the parameters of up to L atoms can be estimated by solving a generalized eigenvalue problem. To do this, two new matrices are calculated from the Hankel matrix, \mathbf{Y}_1 with the last column deleted, and \mathbf{Y}_2 with the first column deleted. A Moore-Penrose pseudo inverse is calculated from \mathbf{Y}_1 , then a complex eigenvalue matrix \mathbf{Y}^+ is calculated from this matrix multiplied by \mathbf{Y}_2 . From \mathbf{Y}^+ , the amplitude, frequency, decay rate and phase can be calculated.

For maximum accuracy in the presence of noise, L should be in the range $N/3$ to $(N/2)-1$ [17].

Algorithm 1: MPM Estimation

	Step	Memory	Complexity
1.	Input: $\mathbf{s} = [s_1, s_2, \dots, s_N]$, L	$N + 1$	
2.	$\mathbf{Y} = \text{hankel}(\mathbf{s}(1:(N - L)), \mathbf{s}(N - L:N));$	$NL - L^2 + N - L$	
3.	$\mathbf{Y}_1 = \mathbf{Y}(:, 1:L);$	$NL - L^2$	
4.	$\mathbf{Y}_2 = \mathbf{Y}(:, 2:L + 1);$	$NL - L^2$	
5.	$\mathbf{Y}^+ = \text{pinv}(\mathbf{Y}_1) \mathbf{Y}_2$	$(N - L)^2$	$O((N-L)^2L)$
6.	$\mathbf{p} = \text{eig}(\mathbf{Y}^+);$	$N - L$	$O((N-L)^3)$
7.	$\mathbf{Z} = \text{zeros}(N, L);$	NL	
8.	for $ii = 1:\text{length}(\mathbf{p})$		
9.	$\mathbf{Z}(:, ii) = \text{transpose}(\mathbf{p}(ii).^{\wedge}(0:N-1));$		$O(N + L)$
10.	end for		
11.	$\mathbf{a} = \mathbf{Z}\mathbf{s};$	L	$O(N^2L)$
12.	Output: $\mathbf{p} = [p_1, \dots, p_L]$, $\mathbf{a} = [a_1, \dots, a_L]$		

The exact algorithm used is adapted from the Matlab code provided in [15]. This method does not compute an SVD explicitly, however it does compute a pseudo inverse which requires an SVD. In order to meet the MPM optimal noise rejection criterion for L , the pencil parameter is set

to $N/3 + 1$, and the resulting atoms truncated during post-processing to the desired quantity M .

At post-processing the complex poles and amplitudes obtained from the MPM algorithm are converted into standard atom parameters (denoted f , a , α and φ), removing duplicate (but negative frequency and phase) atoms, then sorting into decreasing energy order and truncating to the required number of atoms.

Algorithm 1b: MPM Post Processing

Step	Memory	Complexity
1. Input: poles = $[p_1, p_2, \dots, p_N]$, amplitudes = $[a_1, a_2, \dots, a_N]$, $Ts, M, lenIR$	$N + 1$	
2. $\alpha = \log(\text{abs}(\mathbf{poles}))/Ts$;	N	$O(N)$
3. $\mathbf{f} = \text{atan2}(\text{imag}(\mathbf{poles}), \text{real}(\mathbf{poles}))/2\pi Ts$;	N	$O(N)$
4. $\mathbf{a} = \text{abs}(\mathbf{amplitudes})$;	N	$O(N)$
5. $\varphi = \text{atan2}(\text{imag}(\mathbf{amplitudes}), \text{real}(\mathbf{amplitudes}))$;	N	$O(N)$
6. $[\mathbf{f}, \mathbf{a}, \alpha, \varphi] = \text{sortDEORemoveNegFreqAtomsAndTrimtoM}(\mathbf{f}, \mathbf{a}, \alpha, \varphi, M)$;		$O(N \log(N))$
7. Output: $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_M]$, $\mathbf{f} = [f_1, f_2, \dots, f_M]$, $\varphi = [\varphi_1, \varphi_2, \dots, \varphi_M]$, $\mathbf{a} = [A_1, A_2, \dots, A_M]$		

The function `atan2` is the four quadrant inverse tangent. The function `sortDEORemoveNegFreqAtomsAndTrimtoM` sorts the atom arrays into descending energy order, removes those atoms which have negative frequencies, and trims the resulting arrays to length M (or the number of remaining atoms, whichever is the smaller.) It must double the amplitude of the positive frequency atom corresponding to any negative frequency atoms it deletes.

2.3. ESPRIT

ESPRIT as described in [19] is similar to MPM in that it first forms a Hankel matrix from the signal, except the pencil length is fixed at $(N/2) - 1$. It uses singular value decomposition (SVD) to remove the rows and columns of estimated noise components, before calculating the EDS parameters (as complex arrays of poles and amplitudes) from the eigenvalues of this new (smaller) matrix.

The exact algorithm used in this paper is adapted from the DESAM Toolbox [21], namely `HR_Analysis.m`, which explicitly computes an SVD.

Algorithm 2: ESPRIT Estimation

Step	Memory	Complexity
1. Input: $\mathbf{s} = [s_1, s_2, \dots, s_N]$, M	$N + 1$	
2. $N_x = \text{ceil}(N/2) + 1$;	1	
3. $N_y = \text{floor}(N/2)$;	1	
4. $\mathbf{H} = \text{hankel}(\mathbf{s}(1:N_x), \mathbf{s}(N_y:N))$;	$N^2/4 + N/2$	
5. $\mathbf{U} = \text{svd}(\mathbf{H}, 'economy')$;	$N^2/4 + N/2$	$O(N_y^2 N_x)$
6. $\mathbf{W}_1 = \mathbf{U}(1:N_x - 1, 1:M)$;	$N_x M/2$	
6. $\mathbf{W}_2 = \mathbf{U}(2:N_x, 1:M)$;	$N_x M/2$	
8. $\mathbf{W}^+ = \text{pinv}(\mathbf{W}(1: M - 1, :)) * \mathbf{W}(2: M, :)$;		$O(N_x^2 M)$
9. $\mathbf{p} = \text{eig}(\mathbf{W}^+)$;	M	$O(M^3)$
10. $\mathbf{V} = \text{zeros}(N, M)$;	NM	
11. for $m = 1: M$		$O(MN)$
12. if $\text{abs}(\mathbf{p}(m)) < 1$		
13. $\mathbf{V}(:, m) = \mathbf{p}(m). \wedge ((0: N - 1)')$;		
14. else		
15. $\mathbf{V}(:, m) = \mathbf{p}(m). \wedge ((-(N - 1) : 0)')$;		
16. end if		
17. end for		
18. $\mathbf{a} = \text{pinv}(\mathbf{V}) * \mathbf{s}$	M	$O(N^2 M)$
19. Output: $\mathbf{p} = [p_1, \dots, p_M]$, $\mathbf{a} = [a_1, \dots, a_M]$		

Post-processing is similar to MPM, however the set of atoms produced does not need to be truncated after sorting, because the algorithm allows the specification of the *order* (number of atoms to produce) as part of the algorithm. To allow for the fact that this implementation generates symmetric positive and negative frequencies, an order of $M * 2$ is specified.

2.4. Modelled Pursuits

The Modelled Pursuits algorithm used in this paper is as described in [13] except we have modified it to halt once M atoms have been found. Note that MoP may still halt before the limit of M atoms is reached, if an increase in residual energy is detected during an iteration (which would indicate a mis-detection of parameters and does not meet the require-

ment for a monotonic decrease in residual energy).

Algorithm 3: Modelled Pursuits Estimation

Step	Memory	Complexity
1. Input: $\mathbf{s} = [s_1, s_2, \dots, s_N]$, M	$N + 1$	
2. Initialize: $m = 0$, residual $\mathbf{r}_0 = \mathbf{s}$.	$N + 1$	
3. while ($m < M$ or $\ \mathbf{s} - \mathbf{r}_m\ > \epsilon$) and $\ \mathbf{r}_{m-1}\ - \ \mathbf{r}_m\ \geq 0$ do 1		$O(MN \log(N))$
4. $\mathbf{S} = \text{DFT}(\mathbf{r}_m)$	N	
5. Select $k_m: S_{k_m} _{\text{argmax}}$	N	
6. Estimate α_m, φ_m, f_m and a_m	$4M$	
6. Synthesise $\mathbf{g}_m = e^{j\alpha_m} + j\varphi_m - t^{(a_m + j2\pi f_m)}$	N	
8. $\mathbf{r}_{m+1} = \mathbf{r}_m - \mathbf{g}_m$		
9. end while		
10. Output: $\alpha = [\alpha_1, \alpha_2, \alpha_M]$, $\mathbf{f} = [f_1, f_2, \dots, f_M]$, $\varphi = [\varphi_1, \varphi_2, \dots, \varphi_M]$, $\mathbf{a} = [a_1, a_2, \dots, a_M]$		

2.5. Summary

Each method has a different performance in terms of the quality of parameter estimation and the computational time and resources required. This paper compares the SNR, complexity and memory requirements (provided in the preceding algorithm descriptions), along with execution time of each method within a common framework, Matlab [9], to give a real-world performance comparison.

3. Comparison methodology

If the number of expected EDS is small then a significant data reduction after modelling can be expected. Also the modelling process can be optimised in comparison with obtaining the maximum possible number of EDS. The categories of MPS that this paper focusses on are so-called resonant and tympanic, since these are best described in the literature. For these signals there are only a few EDS components in the signal. Hence a small fixed number of modelled atoms M_{resynth} is used for modelling in this paper, that number being chosen as the maximum of the number typically observed in these MPS categories.

This study utilises synthetic MPS. This is to enable data gathering based on a large set of individually unique signals that each broadly adhere to the observed properties of real signals reported in the literature.

The methodology is briefly as follows:

- Generate M_{test} synthetic EDS components with parameters randomly selected from ranges that match the type of MPS as observed in [2,4].




Fig. 1. Relative execution times. a) (left) Resonant template: MoP, MPM and ESPRIT, b) (right) Tympanic template: MoP, MPM and ESPRIT.

- Add white Gaussian noise (WGN) to simulate the additive noise that contaminates MPS when they are captured.
- Perform an atomic EDS decomposition of the noisy signal using each of the three methods under test and produce a resynthesis from the M_{resynth} derived model parameters.
- Calculate the signal to noise ratio (SNR) of the resynthesised signal and compare it to the original.

The strategy used to generate a synthetic EDS signal of a given category is to form a 'template' (see Algorithm 4) based on the mean and standard deviation of the atom parameters from examples in the published literature. The examples' atoms were not initially ordered in a way amenable to this, thus it was necessary to match atoms between examples using an exhaustive search of best fits of a Euclidean distance metric calculated from frequency and energy (a function of amplitude and decay), ignoring phase. From these now best-matched ordered lists of atom parameters, the mean and standard deviation (SD) are calculated. Each test (see Algorithm 5) can randomly simulate any of the EDS components in a new synthetic signal as if they had been drawn from the same population as the published signal examples.

3.1. Evaluation procedure

Algorithm 4: Template Generation

1. for signal categories $sc = \{\text{resonant, tympanic}\}$
2. **sigatoms** = load relevant examples of atom sets for this signal category
3. $M_{\text{max}} = \max(\text{number of atoms per example in sigatoms})$
4. **metrics** = calcEuclidianDistanceMetricFreqEnergy (**sigatom**)
5. **matchedatoms** = findBestMatch(**sigatoms**, **metrics**)
6. for $m = 1:M_{\text{max}}$
7. **components** = $\{f, a, \alpha\}$ of each **matchedatoms**' atom index m
8. for each component index $compind$
9. TPLMEANS($sc, m, compind$) = mean(**components**)
10. TPLSTDs($sc, m, compind$) = std(**components**)
11. end for
12. end for
13. end for

(continued on next column)

(continued)

Algorithm 4: Template Generation

11. The resulting template is a matrix of sc signal categories, each containing its relevant M_{max} atoms' component means and standard deviations.
12. Store a minimum M_{min} and maximum M_{max} to use in generation per sc (see next paragraph), and global minima and maxima per component e.g. $f > 0$; $a > 0$, $\alpha < -0.5$, etc.

The number of atoms in the published examples of resonant and tympanic signals are 14, 11 and 12, 8 respectively. Based on observations of significant EDS [2] in a typical tympanic signal, the number of atoms that can be chosen for that category is in the range one to four. For the resonant signals that range is between five and ten atoms.

Using each template, the following steps are performed to generate one synthetic signal to be used in testing:

Algorithm 5: Generate Random EDS from template

1. Choose the first M_{test} atoms from the relevant template, where M_{test} is chosen randomly between that template's minimum and maximum M .
2. Deviate each of these atoms' components $\{f, a, \alpha\}$ by $\max(N(\text{mean}, \text{SD}), 0.05)$ where N is the normal distribution.
3. The phase ϕ per atom is selected from $U(0, 2\pi)$ where U is the uniform distribution.
4. From the resulting M_{test} atoms obtained, a total signal is generated that is the sum of each atom's resynthesised EDS signal. The result is the noise-free ground-truth signal x to be compared-to after analysis.

The test parameters are based on existing published findings and synthetic signals [2]. As well as the specification of the synthetic signals above, the choice of appropriate sample rate, test signal duration and number of atoms to extract from each model is determined as follows. A sample rate of 2 kHz is employed since we are only interested in frequencies below 1 kHz [2]. Despite an assumption in a previous study that windows of 40 ms are sufficient to capture MPS [22], analysis of the slowest decay rate of the six EDS components in synthetic signals in [2] shows that the -90 dB point is only reached after approx. 1000 samples, and so therefore 0.5s is chosen as the total signal duration in this study. For each unknown test signal the maximum number of atoms M_{resynth} to request from the modeller (after post processing) is chosen to be 10, in

line with the maximum number of significant EDS components found in [22].

The test SNRs used are -20 dB to $+90$ dB in steps of 10 dB (12 different SNRs.) where:

$$\text{SNR} = 10 \cdot \log_{10} \left(\frac{\sum_{n=1}^N x_n^2}{\sum_{n=1}^N \varepsilon_n^2} \right) \quad (2)$$

x_n are the reference test signal samples, ε_n are the (white Gaussian) noise samples and N is the number of samples. At each SNR level 5000 random atom-sets are synthesised using equation [13] from each of the two signal category templates. The total number of tests (noisy signals modelled by each modeller) is thus 2 [signal categories] * 5000 [random sets of atoms] * 12 [different SNRs] * 16 [repeats of different random noise at same SNR] = 1,920,000 (N) total test signals generated.

The tests were performed on the University of York Viking high speed computing cluster (see acknowledgements), running Matlab 2019 [9]. Each of the test signals was modelled with the MoP, MPM and ESPRIT algorithms described previously, each resulting in a model containing up to 10 atoms.

The produced atoms are resynthesised, according to equation (1), to produce a signal y which is the sum of each atom generated from the estimated parameters. Then SNR_{mod} (in dB) is calculated using the following equation, where x is the original random test signal (prior to any noise being added):

$$\text{SNR}_{\text{mod}} = 10 \cdot \log_{10} \left(\frac{\sum_{n=1}^N x_n^2}{\sum_{n=1}^N (y_n - x_n)^2} \right) \quad (3)$$

4. Results

4.1. Execution times

Fig. 1 shows boxplots of the execution time in seconds taken to model each of the test signals by each modeller MoP, ESPRIT and MPM. The left hand plot shows the results for the set of resonant test signals, and the right hand plot, those from the set of tympanic test signals. In this figure, the top and bottom of each blue box represents the upper and lower quartiles of the data respectively. The black whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted




Table 1

Mean and standard deviation (SD) of execution times per test (number of tests: 1,920,000).

Modeller	Signal Category			
	Resonant		Tympanic	
	Mean	SD	Mean	SD
MoP	1.00	0.271	1.00	0.268
MPM	15.0	1.65	14.9	1.63
ESPRIT	2.54	0.259	2.51	0.256

individually using the red '+' symbol. Points are deemed outliers if they are greater than $q3 + (1.5 \times (q3 - q1))$ or less than $q1 - (1.5 \times (q3 - q1))$, where $q1$ and $q3$ are the 25th and 75th percentiles of the sample data, respectively. The mean is plotted as a green diamond symbol, and the mean plus one standard deviation and mean minus one standard deviation are plotted as green x symbols.

Fig. 1b uses the same data as Fig. 1, but normalised so as to correspond with Table 1 and with outliers clipped at a maximum of 22.5 to show more detail.

To allow for easier comparison, the mean execution time and its standard deviation (for each method and signal category) is listed below in Table 1, after first normalising each with respect to the shortest mean per signal category (that of MoP in both cases.)

4.2. Accuracies

Fig. 2 shows the mean (with error bars at plus and minus one standard deviation) of the modelled SNR in decibels on the vertical axis, against the actual SNR in decibels on the horizontal axis, for the set of resonant (left hand plot) and tympanic (right hand plot) test signals. In each plot, MoP is shown in dashed black, MPM dotted red, and ESPRIT dot-dashed blue. For reference, a plot of equal modelled SNR against actual SNR is shown as a solid grey line with x markers.

5. Discussion

MoP is clearly the cheapest of the three methods in terms of execution time for the short length signals discussed in this paper. The large




Fig. 1b. Relative execution times as Fig. 1 but normalised so as to correspond with Table 1 and with outliers clipped at a maximum of 22.5 to show detail.

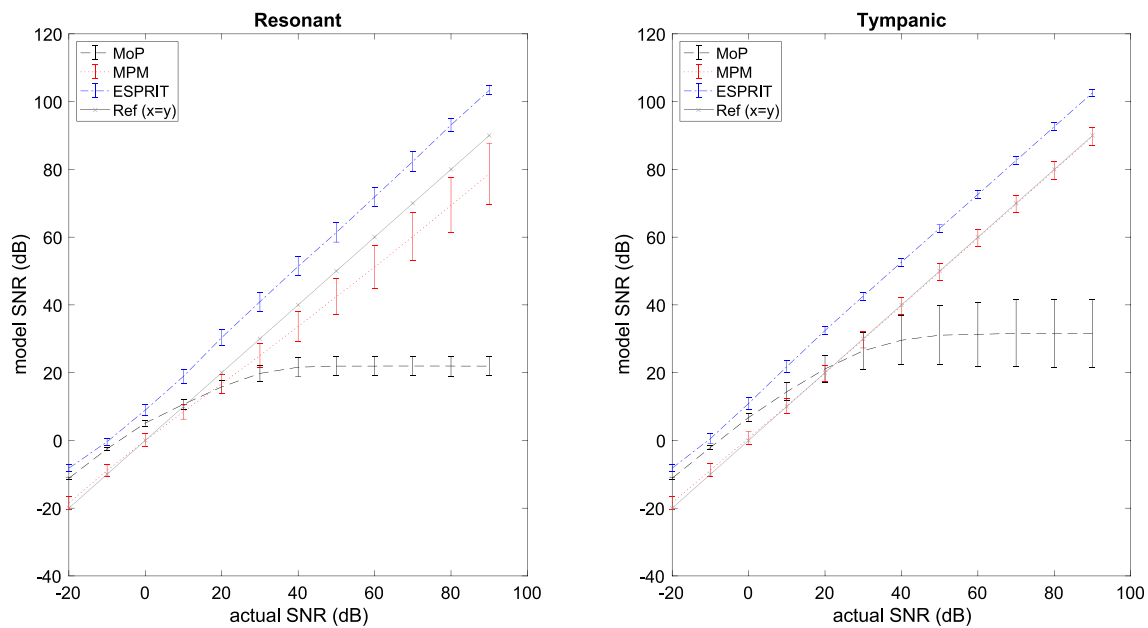


Fig. 2. Modelled SNR v original SNR for Resonant and Tympanic templates.

difference between the execution times of MPM and ESPRIT can be accounted for by the fact that ESPRIT uses singular value decomposition (SVD) as an early step in the algorithm, which significantly reduces the size of the matrices on which subsequent operations such as pseudo inverse and eigenvalue are deployed. Therefore the overall computation for ESPRIT is significantly less than it is for MPM.

It can be seen from Table 1 that there is no significant difference between the execution times of the resonant signal category and the tympanic signal category, for any of the methods under test. This is to be expected because the algorithms under test do not attempt any derivation of the underlying number of noise free EDS in the original signals, and that is the main difference between the two synthetic test signal categories. The outliers can be reasonably explained by random perturbations in the timing system used due to operating system influence.

It can be seen in Fig. 2 that ESPRIT performs the best in terms of accuracy. It is in all cases able to recover a SNR greater than the actual SNR by about 12 to 13 dB. MPM can only recover an SNR approximately equal to that of the added noise for tympanic signals, and is slightly worse for resonant signals, having a reasonably constant slope that starts at about 1.5 dB better for minus 20 dB actual SNR, and finishes at about 11 dB worse than the actual SNR at an actual SNR of plus 90 dB.

MoP can, for SNRs less than approx. 20 dB, recover the signal with a better SNR than the test signal, but this tails off rapidly, shelving at approximately 22 dB for resonant and 33 dB for tympanic signals. This upper limit of the performance of MoP for high SNRs is due to errors in the atom parameter estimation from individual Fourier spectrum components, as opposed to their derivation from a global, matrix-based, minimisation approach. The lower error bar of MoP for the resonant test signal category is roughly the same SNR as that for tympanic signals, however in the tympanic case, the upper error bar and mean are higher, indicating that MoP performs better for signals with fewer EDS components (as in the tympanic case.).

6. Conclusions

ESPRIT and MPM do better overall than MoP in terms of parameter estimation, an expected result since the first two are matrix decomposition methods. The implication is that they would be better in clinical situations where signals can be captured with a very high signal to noise ratio and where ample computation resources are available. However, in noisy conditions (lower than 20 dB SNR) MoP performs as well as

ESPRIT and MPM but requiring considerably less memory. In such situations the data that we have produced indicates that MoP offers the best compromise between computational demands and effectiveness.

CRediT authorship contribution statement

Kenneth I. Brown: Methodology, Software, Validation, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Jeremy J. Wells:** Conceptualization, Methodology, Software, Formal analysis, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of Competing Interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The project was publicly funded via the authors' institution, and was undertaken in collaboration with RPPTv Ltd who provided non-financial support in the form of consultancy.

Data availability

No data was used for the research described in the article.

Acknowledgements

This project was supported by the University of York through their institutional impact acceleration account from EPSRC (EP/R51181X/1).

The Industrial partner and collaborator on this project was RPPTv Ltd.

This project was undertaken on the Viking Cluster, which is a high performance compute facility provided by the University of York. We are grateful for computational support from the University of York High Performance Computing service, Viking and the Research Computing team.

References

- [1] A. Murray, J.M.M. Neilson, Diagnostic percussion sounds: 1, A Qualitative Analysis. *Med. Biol. Eng.* 13 (1) (1975) 19–28.

- [2] M. Bhuiyan, E.V. Malyarenko, M.A. Pantea, F.M. Seviaryn, R.G. Maev, Advantages and limitations of using matrix pencil method for the modal analysis of medical percussion signals, *IEEE Trans Biomed Eng.* 60 (2013) 417–426, <https://doi.org/10.1109/TBME.2012.2227318>.
- [3] A. Rao, J. Ruiz, C. Bao, S. Roy, Tabla: A Proof-of-Concept Auscultatory Percussion Device for Low-Cost Pneumonia Detection, *Sensors* 18 (8) (2018) 2689, <https://doi.org/10.3390/s18082689>.
- [4] J. Dech, M. Bhuiyan, R. Maev, Analysis of Clinical Percussion Signals Using Matching Pursuit, *IJCEE* 7 (4) (2015) 248–260. ISSN: 1793–8163.
- [5] B. Allen, et al., Early detection of ACS through electronic recording and analysis of auscultatory percussion, *IEEE J. Trans. Eng. Health Med.* 8 (2020).
- [6] Ekohealth, Ekohealth web page, <https://www.ekohealth.com/> (accessed 10th Jan 2023).
- [7] S. Ryu et al., iApp: An Autonomous Inspection, Auscultation, Percussion, and Palpation Platform. *Frontiers in Physiology*, February 2022.
- [8] K.P. Ayodele, et al., A medical percussion instrument using a wavelet-based method for archivable output and automatic classification, ISSN 0010-4825, *Comput. Biol. Med.* 127 (2020) 104100, <https://doi.org/10.1016/j.compbimed.2020.104100>.
- [9] MATLAB, version 9.8.0.1417392 (R2019). Natick, Massachusetts: The MathWorks Inc., 2019.
- [10] R. Badeau, G. Richard, B. David, Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials, *IEEE J_SP* 56 (2) (2008) 492–504.
- [11] W. Long, R. Chen, M. Moretti, Recursive ESPRIT Algorithm for Multi-User OAM Low-Overhead AoA Estimation, *IEEE Trans. Vehicular Technol.* 72 (2) (2023) 2672–2677.
- [12] K. Sangdong, B. Kim, Y. Jin, J. Lee, Extrapolation-RELAX Estimator Based on Spectrum Partitioning for DOA Estimation of FMCW Radar, *IEEE Access* 7 (2019) 98771–98780.
- [13] J. Wells, Modal Decompositions of Impulse Responses for Parametric Interaction, *J. Audio Eng. Soc.* 69 (7/8) (2021) 530–541.
- [14] S.G. Mallat, Z. Zhang, Matching pursuits with time-frequency dictionaries, *IEEE t. Signal Proces.* 41 (12) (1993) 3397–3415.
- [15] A. Fernández Rodríguez, et al., Coding Prony's method in MATLAB and applying it to biomedical signal filtering, *BMC Bioinf.* 26;19 (1) (2018) 451, <https://doi.org/10.1186/s12859-018-2473-y>.
- [16] T.K. Sarkar, O. Pereira, Using the matrix pencil method to estimate the parameters of a sum of complex exponentials, *IEEE Antennas Propag Mag.* 37 (1995) 48–55, <https://doi.org/10.1109/74.370583>.
- [17] Y. Hua, T. Sarkar, Matrix pencil method for estimating parameters of exponentially damped sinusoids in noise, *IEEE Trans. Acoust. Speech Sig. Proc.* 38 (5) (1990) 814–824.
- [18] R. Roy, A. Paulraj, T. Kailath, ESPRIT—A subspace rotation approach to estimation of parameters of complex sinusoids in noise, *IEEE Trans Acoust Speech Signal Process.* 34 (1986) 1340–1342, <https://doi.org/10.1109/TASSP.1986.1164935>.
- [19] D. Potts, M. Tasche, Parameter estimation for nonincreasing exponential sums by Prony-like methods, *Linear Algebra Appl.* 439 (2013) 1024–1039.
- [20] Source code available at https://www.jezwells.org/modal_analysis_medical_percussion.
- [21] M. Lagrange et al., The DESAM Toolbox: Spectral Analysis of Audio. Proceedings of the 13th Intl. Conf. on Digital Audio Effects (DAFx10, 2010, pp. 254-261.
- [22] M. Bhuiyan, Automated Classification of Medical Percussion Signals for the Diagnosis of Pulmonary Injuries, *Electronic Theses and Dissertations.* 4941 (2013). <https://scholar.uwindsor.ca/etd/4941>.