A COMPARISON OF MODAL PARAMETER EXTRACTION METHODS WHEN APPLIED TO MEASUREMENTS OF STRINGED INSTRUMENTS

Mark Rau[1*] **Julius O. Smith**[1] **Jonathan S. Abel**[1]

[1] Center for Computer Research in Music and Acoustics (CCRMA), Stanford University, USA

ABSTRACT

Modal testing is a commonly used method to measure the transfer function or frequency response function of musical instruments or their components. Various excitation and measurement tools are used, and the recorded signals are analyzed to estimate the modal frequencies, damping ratios, and mode shapes of the object. These modal parameters are used to compare musical instruments, study changes to their geometry and materials, create synthesis models, and verify finite element and other simulation models. An object is typically excited with an impact hammer or shaker, while the resulting vibrations are measured with a microphone, accelerometer, or laser Doppler vibrometer. However, musical instrument builders don't typically have access to expensive measurement equipment, so more ad-hoc methods may be used. Multiple methods exist to extract modal parameters from the measured transfer function, each with its own strengths and weaknesses. This study compares commonly used modal extraction methods when applied to measurements of stringed instruments made with a wide spectrum of excitation and measurement sensors. The methods are evaluated based on generated modal data, and then tested with measurements of musical instruments and other objects.

Keywords: *modal testing, modal fitting, stringed instruments.*

**Corresponding author*: *mrau@ccrma.stanford.edu.*

1. INTRODUCTION

When studying the acoustics and sound of musical instruments, it is useful to take vibrational and acoustic measurements. For a struck, plucked, or bowed instrument, such as a marimba bar, guitar, or violin, we can assume the vibrations are small and the response stays in the linear region. Generally, of interest are the modal frequencies, damping, and amplitudes or mode shapes of the structure. To extract these modal parameters, various algorithms for *mode or modal fitting* can be performed [1–4].

When performing modal fitting on measurements, care can be put into the measurement setup and post-processing, so the algorithms hopefully work reasonably well. Then, you must choose an algorithm, and tune the parameters of the algorithm to properly extract the modal data. This is not a trivial task, and each method seems to have certain shortcomings that are hard to determine before testing each method and parameter space. More sophisticated algorithms such as those using modal optimization can be used, but even they typically require a reasonable initial estimate of the modal frequencies, damping, and amplitudes [5, 6].

Modal fitting becomes especially difficult when working with non-ideal measurement conditions or inexpensive sensors. Situations like this are common as instrument builders wish to extract modal data from the instruments they build, but may not have access to expensive measurement equipment, a vibration, and sound-isolated space, and the tools and software necessary for the mode fitting algorithms [7]. Instrument builders are becoming more and more interested in acquiring this data, so one motivation for this study is to investigate the reliability, accuracy, and ease of use of some common modal fitting algorithms when applied to musical instruments and similar objects.

This paper outlines the initial investigation into testing four modal fitting algorithms. Modal time-domain responses are generated from modal frequencies, damping, and amplitudes to test the algorithms against ground truth results. The generated data includes clean as well as noisy data. Additionally, modal measurements were made with 10 guitars, 10 other instruments, 10 everyday objects, and 20 boards of guitar top wood. The algorithms are briefly described and evaluated on various metrics when applied to the generated and measured vibration data.

Section 2 outlines the generated and modal data. Section 3 briefly introduces the four algorithms being tested. Section 4 describes the evaluation criteria used and presents results from the mode fitting and evaluations. Section 5 discusses the evaluation, shortcomings, and areas for future study.

2. DATA

To evaluate the algorithms, vibration responses are needed. The vibration response is assumed to be linear, having the response of a set of M damped simple harmonic oscillators. The impulse response is then:

$$h(t) = \sum_{m=1}^{M} \gamma_m e^{2\pi f_m t(i - \zeta_m)}, \qquad (1)$$

where γ_m, f_m, and ζ_m are the amplitude, natural frequency, and damping ratios of modes $m = 1, 2, ..., M$ [1,2].

The measurements and generated impulse response data are available at `https://ccrma.stanford.edu/~mrau/FA2023/`. Additional information about the instruments and objects as well as photos are provided.

2.1 Generated Data

In order to evaluate each algorithm, the true values of the mode frequencies, damping ratios, and amplitudes need to be known beforehand which is not possible with measurements. Impulse responses were generated to be used as synthetic measurements. Each impulse response was formed using eq. 1, and is of length 1 second at a sample rate of 48 000 samples/second.

Impulse responses were generated with M = $[1, 2, 3, \ldots, 50]$ modes with 100 examples generated for each. The mode frequencies were chosen randomly between 30 to 10 000 Hz and logarithmically distributed. Damping ratios were randomly chosen between 0.001 and 0.1. Each impulse response had two sets of amplitudes, one with positive-real values, simulating single input/output measurement locations, and one with complex values simulating different input and output measurement locations. The real and complex parts of each amplitude value were randomly chosen to be between 0.0001 and 1.

To simulate non-ideal measurements, the generated impulse responses were augmented in three different ways. In the first case, white noise having an amplitude of 0.01 was added. In the second case, AC hum was simulated by adding 60, 120, and 180 Hz components at amplitudes of 0.1, 0.01, and 0.001. In the third case, the data was clipped by enforcing $-0.5 < h < 0.5$.

This resulted in a total of 45 000 simulated impulse responses with known mode frequencies, damping ratios, and amplitudes.

2.2 Measurements

Ultimately, to test the effectiveness of the algorithms, they need to be evaluated on physical measurements. Four categories of measurements were chosen:

- 10 guitars – 1890s parlour, four dreadnoughts, orchestra model, 00-style, two 000-style, and one classical guitar.
- 10 other instruments – five violins, viola, two cellos, double bass, and one mandolin.
- 10 everyday objects – mug, espresso cup, water pitcher, thermos, frothing pitcher, doorbell, hole saw, tamper, milk pitcher, and a bowl.
- 20 wood top plates for guitars – four German spruce, four Italian spruce, four Sitka spruce, four redwood, and four sinker redwood.

The guitars, other instruments, and everyday objects were measured with lab-quality equipment as well as custom-built low-cost equipment. All measurements were taken with the impact hammer method [2]. The lab equipment consisted of a force hammer for the impact, a laser Doppler vibrometer (LDV), an accelerometer, and a calibrated microphone for the receivers. The low-cost equipment consisted of a 3D printed piezo impact hammer, piezo accelerometer, and inexpensive measurement microphone. The wood boards were measured with a custom-built tonewood measurement device [8].

Five measurements were taken with each object and sensor configuration, resulting in 850 measurements.

3. MODE FITTING ALGORITHMS

Four modal fitting algorithms are tested. Only brief explanations of the methods are provided and the reader is directed to the relevant citations for more details on each algorithm. Each algorithm was run in Matlab 2022b on a 2023 Macbook Pro [9].

1. Peak-Picking (pp) - In the frequency domain, significant local peaks are assumed to correspond to a single mode. Each mode is assumed to be a simple harmonic oscillator, and the damping and amplitudes are solved by setting up and solving a system of equations for each mode [10].
2. Least-Squares Complex Exponential (lsce) - The impulse response is compared to a matrix of complex damped sinusoids formed using Prony's method to find the roots that give the mode frequencies and damping. The mode amplitudes are found using least-squares to solve the system of equations corresponding to the basis sinusoids, mode amplitudes, and initial impulse response [10].
3. Least-Squares Rational Function (lsrf) - The transfer function coefficient constraints are expressed in terms of orthonormal rational basis functions on the unit circle, and the mode frequencies, gains, and amplitudes are determined using an iterative scheme [11].
4. Hankel Impulse Response (Hankel) - a Hankel matrix of time-domain impulse response samples is created, and the eigenstructure is analyzed to determine the mode frequencies and damping. The mode amplitudes are found using least-squares to solve the system of equations corresponding to the basis sinusoids, mode amplitudes, and initial impulse response [12].

Multiple other mode fitting algorithms exist, but these four were chosen as the first three are reasonably common methods and are implemented in the Matlab Vibration Analysis Toolbox [9], and the author has generally had good experiences with the 4^{th}. All algorithms are tested on 1 second of vibration response recorded at 48 000 samples/second. While better results could be obtained by changing the sample rate, processing the raw data in various ways, and adjusting algorithm parameters, that was not the objective of the study, and each algorithm is evaluated on its performance when implemented in a naive manner.

4. EVALUATION

The generated modal data has a ground truth, so the modal fitting can be evaluated based on the difference between the fit and ground truth. This is not true with the measured modal data, so a qualitative evaluation is provided instead.

4.1 Evaluation Criteria

The modal fitting of the generated impulse responses was evaluated on 6 different criteria.

1. Time that is taken for the modal fitting to run in MATLAB on a 2023 Macbook Pro.
2. The number of Modes fit by the algorithm compared against the true number of modes.
3. Mean absolute error in Decibels of the frequency response when comparing the generated and fit frequency response functions.
4. Mean absolute error between the generated and fit mode frequencies.
5. Mean absolute error between the generated and fit mode damping ratios.
6. Mean absolute error between the generated and fit mode amplitudes.

4.2 Results

4.2.1 Generated Data

Figures 1 to 6 show the modal fitting evaluations for 10 measurements at each of the M = $[1, 2, 3, \ldots, 50]$ mode cases. Only the real amplitude with no added noise, AC hum, or clipping is shown for brevity. However, in general, the modal fitting is worse when additional noise, AC hum, or clipping is added. The graphs of the other cases can be found at `https://ccrma.stanford.edu/~mrau/FA2023/`.

All algorithms are quite fast, with the exception of lsrf. Only pp received a speedup when a lower number of modes was fit. Both the pp and Hankel methods fit a reasonably close number of modes to the true value, and underfit the number of modes at higher numbers, while the lsce and lsrf generally fit the same number of modes regardless of the true number of modes. Ideally, the algorithm would fit the correct number of modes, so pp and the Hankel method seem better in this case.

When observing the frequency response error, each algorithm does reasonably well, but the lsrf seems to do better in general. At low numbers of modes, the lsrf appears to fit the mode frequencies but gets worse with a higher number of modes. The other algorithms are all roughly the same, with around 1 Hz of error in the found modes.

The lsce and lsrf appear to fit the damping ratios better than pp and the Hankel method. The Hankel method appears to fit the amplitudes best at low mode numbers but gets worse when more modes are added.

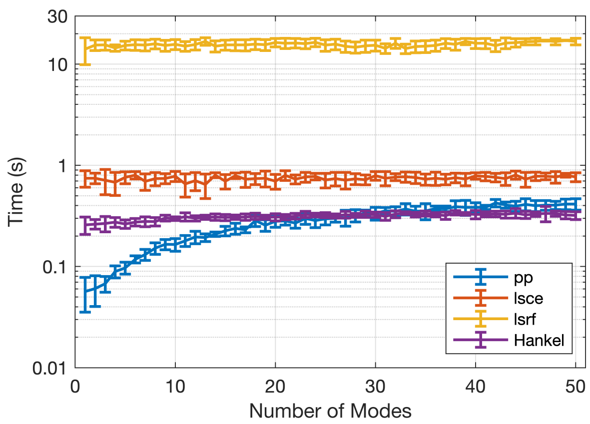

Figure 1: Generated fitting time.

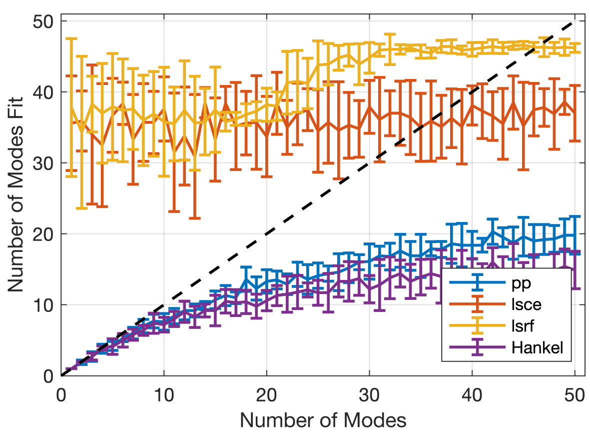

Figure 2: Generated number of modes fit.

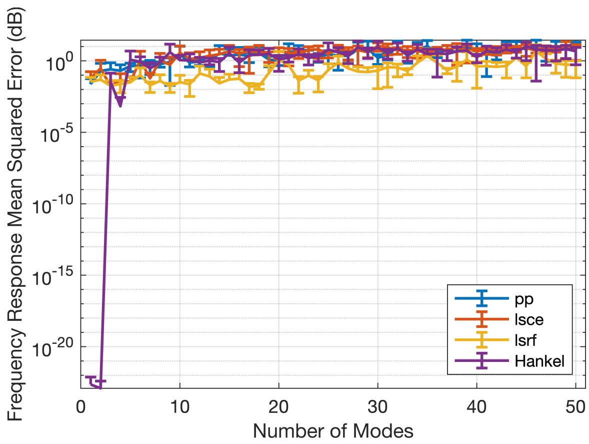

Figure 3: Generated fitting frequency response error.

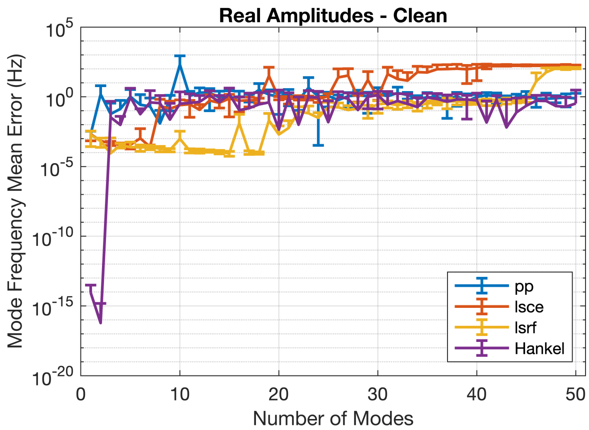

Figure 4: Generated fitting mode frequency error.

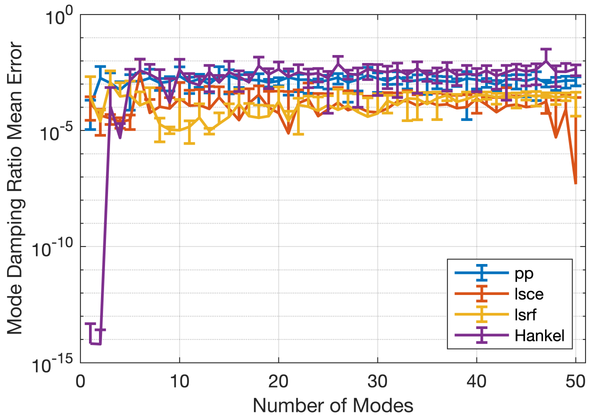

Figure 5: Generated fitting mode damping ratio error.

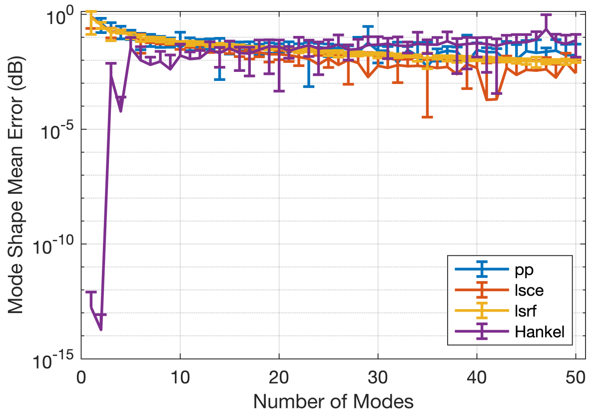

Figure 6: Generated fitting mode amplitude error.

4.2.2 Measurements

While no concrete evaluation can be performed on the measurements, visual observations can be made when comparing the measured frequency response function to that of the modal fit. Figures 7 through 11 show admittance measurements of a guitar, upright bass, cello, mug, and wood board. With the mug, having fewer modes with low damping ratios, all four of the mode fitting algorithms do a reasonable job. However, for each of the instrument and wood board examples, none of the algorithms result in what would be considered a good fit. Some are reasonable in certain frequency ranges, but none are perfect. It also appears that no single algorithm works well in each case, with some failing wildly for some instruments and working well for others. In general, it seems that the peak-picking and Hankel methods work the best, but neither is perfect in this naive implementation.

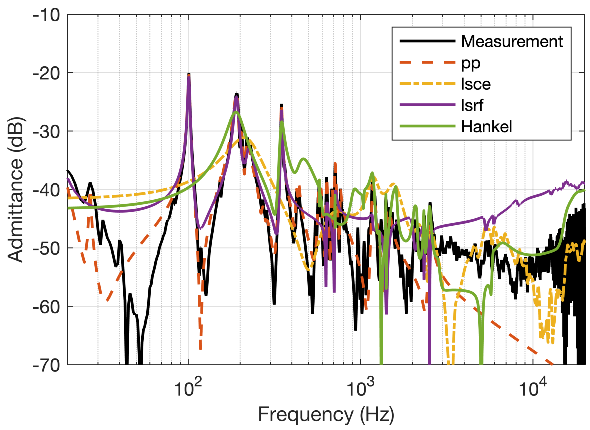

Figure 7: Guitar Admittance.

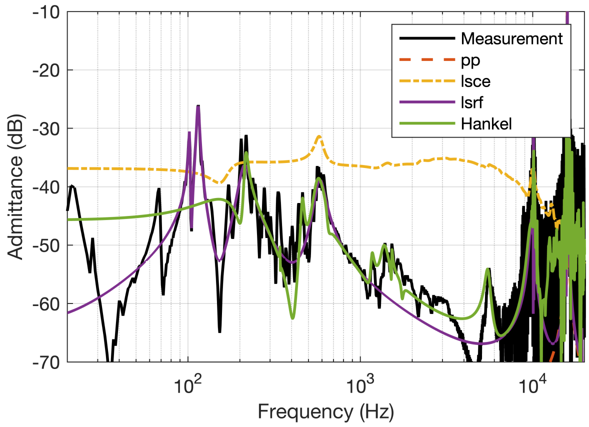

Figure 8: Bass Admittance.

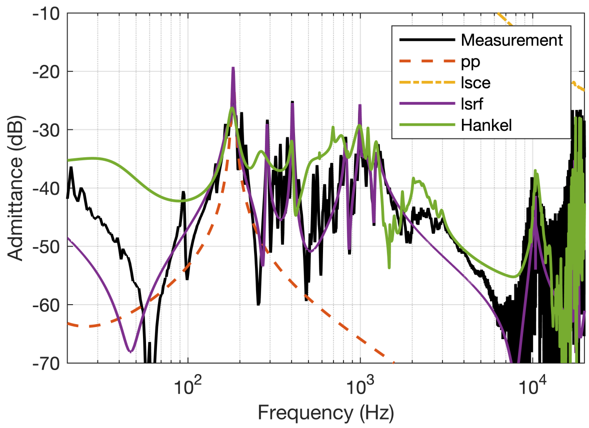

Figure 9: Cello Admittance.

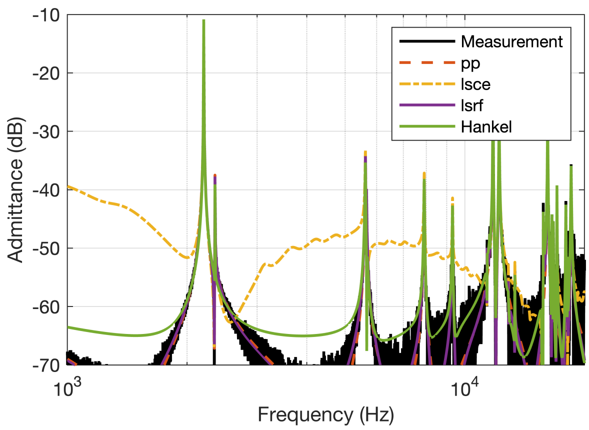

Figure 10: Mug Admittance.

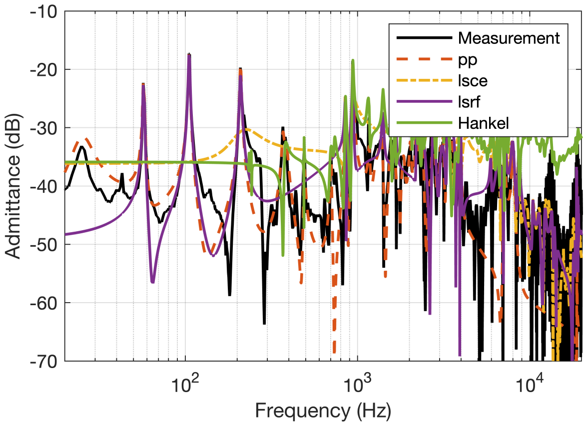

Figure 11: Redwood board Admittance.

5. DISCUSSION AND CONCLUSIONS

This study is not, and was not meant to be a conclusive investigation of all mode fitting algorithms, but rather serves as a beginning and framework for future studies. However, some conclusions and directions for future investigation can be drawn.

Generated data provides the only true way to evaluate the data, but the modal algorithms tend to perform better on generated data than they do on the measured data. This suggests, that physical measurements should always be used as a final check when performing modal fitting.

None of the methods tested perform well enough to be used and trusted in a non-ideal and unsupervised setting such as when being used in an instrument builder's workshop. More work needs to be done to test different algorithms, and methods to optimize them in a simple and repeatable manner, and verify their reliability.

One avenue for future study the authors hope to explore is a machine-learning approach to modal fitting. The generated and measured modal data collected for this study could be used as a training/test set for this machine learning approach. Hopefully, machine learning could potentially be leveraged to provide a lightweight, simple-to-use, and reliable mode-fitting approach that can be used by instrument builders and hobbyists.

6. ACKNOWLEDGMENTS

I would like to thank Samuel Clarke, Daiki Nakajima, Joseph Malinowski, the Santa Cruz Guitar Company, Herbert Myers, and the Stanford Harry R. Lange Collection of Musical Instruments and Bows for graciously lending instruments for this study.

7. REFERENCES

[1] A. Technologies, "Fundamentals of modal testing," 2000.

[2] A. Brandt, *Noise and vibration analysis: signal analysis and experimental procedures*. John Wiley & Sons, 2011.

[3] D. L. Brown, R. J. Allemang, R. Zimmerman, and M. Mergeay, "Parameter estimation techniques for modal analysis," *SAE transactions*, pp. 828–846, 1979.

[4] E. Reynders, "System identification methods for (operational) modal analysis: review and comparison," *Archives of Computational Methods in Engineering*, vol. 19, pp. 51–124, 2012.

[5] E. Maestre, G. P. Scavone, and J. O. Smith, "Joint modeling of bridge admittance and body radiativity for efficient synthesis of string instrument sound by digital waveguides," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 5, pp. 1128–1139, 2017.

[6] E. Maestre, G. P. Scavone, and J. O. Smith, "Design of recursive digital filters in parallel form by linearly constrained pole optimization," *IEEE Signal Processing Letters*, vol. 23, no. 11, pp. 1547–1550, 2016.

[7] M. Rau, R. Hoover, R. Barto, and J. O. Smith, "A self-contained and automated tonewood measurement device," *Vienna Talk on Musical Acoustics*, Vienna, Austria, September 11–14, 2022. published abstract.

[8] M. Rau, "Measurements and analysis of acoustic guitars during various stages of their construction," *in 180th Meeting of the Acoustical Society of America*, Virtual, June 8–10, 2021. published abstract.

[9] The MathWorks Inc., "Matlab version: 9.13.0 (r2022b)," 2022.

[10] The MathWorks Inc., "Modal parameters from frequency-response functions - matlab2017," 2017.

[11] A. A. Ozdemir and S. Gumussoy, "Transfer function estimation in system identification toolbox via vector fitting," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6232–6237, 2017.

[12] M. Rau, J. S. Abel, D. James, and J. O. Smith III, "Electric-to-acoustic pickup processing for string instruments: An experimental study of the guitar with a hexaphonic pickup," *The Journal of the Acoustical Society of America*, vol. 150, no. 1, pp. 385–397, 2021.