Available online at www.eccomasproceedia.org
Eccomas Proceedia UNCECOMP (2023) 82-93

ECCOMAS Proceedia

UNCECOMP 2023
5th ECCOMAS Thematic Conference on
Uncertainty Quantification in Computational Science and Engineering
M. Papadrakakis, V. Papadopoulos, G. Stefanou (eds.)
Athens, Greece, 12-14 June 2023

A NOVEL VARIATIONAL BAYESIAN APPROACH TO STOCHASTIC SUBSPACE IDENTIFICATION

Brandon J. O'Connell[1], **Max D Champneys**[1], **Elizabeth J. Cross**[1] **and Timothy J. Rogers**[1]

[1]Department of Mechanical Engineering, The University of Sheffield
Sir Frederick Mappin Building, Mappin Street, Sheffield, S1 3JD
e-mail: {bjoconnell1, max.champneys, e.j.cross, tim.rogers}@sheffield.ac.uk

---

Abstract. *Covariance-driven stochastic subspace identification (SSI) is a frequently employed modal analysis technique, often used in operational modal analysis (OMA) applications, as a reliable means of recovering the modal properties of a structural dynamic system. At its core, this method relies on a mathematical concept known as canonical correlation analysis (CCA) that seeks to find the correlation between Hankel matrices of the future and the past observations, from a set of response sensors, measuring a dynamic system. In previous work by the authors, a probabilistic formulation of SSI was presented that saw the replacement of traditional CCA with its probabilistic equivalent, using the theory of latent variable models. This change in formulation provides new insight into this well established approach. Subsequently, the probabilistic method was further extended by the authors to a so-called, fully Bayesian approach and solved using Markov Chain Monte Carlo (MCMC) sampling to recover the posterior distributions over the modal properties. The availability of the posterior uncertainty provides additional information to the user which can impact future decision making or modelling exercises. This paper presents a continuation of the Bayesian SSI formulation in the form of a novel variational Bayesian SSI approach, capable of approximating a surrogate posterior distribution over the modal properties. It is shown on a simple case study how suitable approximations to the posterior distributions over the modal properties can be recovered that show good agreement with the truth, whilst also encompassing the SSI estimate in the posterior. This is also followed by a brief discussion on its overall performance and the possible limitations and how these could be addressed.*

Keywords: Variational Inference, System Identification, Modal Analysis, Stochastic Subspace, Bayesian

---

ISSN:2623-3339 © 2023 The Authors. Published by Eccomas Proceedia.
Peer-review under responsibility of the organizing committee of UNCECOMP 2023.
doi: 10.7712/120223.10326.19796

1 INTRODUCTION

The characterisation of structural dynamic systems is a well-established area of research and continues to be a key component of modern engineering practice. Modal analysis is a common type of system identification used to determine the modal properties of linear dynamic systems from experimental data; familiar to the dynamicist as the natural frequencies, damping ratios and mode shapes. The availability of this modal information is useful to the practitioner as it provides a factual basis for informed decision making throughout the entire lifecycle of an engineering structure. Rarely used in isolation, modal analysis plays a central role in a much larger set of tools and processes such as model updating, structural health monitoring (SHM) and digital twins.

Stochastic subspace identification (SSI) [1, 2] is a prominent modal analysis technique for which the main objective is to estimate the linear state-space model of a structural dynamic system from correlated sequences of observed data. An eigen-decomposition of the recovered state-space model, obtained using traditional linear algebra techniques, yields point estimates for the desired modal properties. Frequently employed across industry and research, SSI (and its many algorithmic variants) continue to be a reliable means of experimental and operational modal analysis.

Operational modal analysis (OMA), one of the leading applications of SSI and the focus of this work, is a subset of modal analysis concerned with the recovery of modal properties in the absence of measured input information; often the case when testing in-situ (operation). Operational testing is routinely required due to scale, complexity, or an inability to accurately recreate forcing conditions in a laboratory setting. This requirement is typical of in-service, large scale structures such as civil infrastructure, aircraft and off-shore structures. The response of these systems is instead measured during ambient excitation, which is seldom measured and often assumed to be broadband white noise. Inevitably, OMA has gained increased popularity in recent years due to its high economic value and convenience in many engineering applications, especially for high-value large-scale assets [3, 4].

Most OMA techniques, including SSI, operate deterministically; recovering only point estimates of the modal properties and rarely accounting for or calculating any associated uncertainty. However, the inclusion of uncertainty, in its various forms[1], has proven to be greatly useful when tackling certain dynamics focused tasks such as; automatic selection of poles in the consistency diagram[2], modal parameter estimation [7] and robustness [8].

In recent years, Bayesian approaches to system identification have gained increasing attention in the dynamics community [9]. The successful use of Bayesian methodologies and application of posterior uncertainty in 'downstream' applications (those reliant on modal data) such as SHM and digital twins, has highlighted an explicit need to expand the scope of Bayesian frameworks to 'upstream' tasks such as system identification. The ultimate objective of which is to to recover or approximate the posterior uncertainty and in doing so, provide the experienced dynamicist with valuable additional information in the decision making process. Furthermore, viewing system identification as an inference problem can offer a more powerful and insightful perspective that explicitly addresses the problem of uncertainty.

[1]It is important to note, that there are many types of uncertainty, whether that be a posterior uncertainty in a Bayesian sense, estimates for bounds, confidence intervals, fiducial intervals, etc. In this work, a Bayesian paradigm is used.

[2]This often referred to as the stabilisation diagram, however the authors believe this choice of terminology may cause confusion with the 'stability' of poles referenced in control theory. Hence, the term consistency diagram is used. This transition has also been adopted by others in the field of dynamics [5, 6].

In previous work by the authors [8] it was shown that, through leveraging the presence of canonical correlation analysis (CCA), covariance-driven SSI can be reformulated probabilistically through direct replacement of CCA with the theory of probabilistic projections (probabilistic CCA) [10]. This was later extended, using hierarchical methods, to a novel Bayesian interpretation of SSI [11] using Bayesian CCA [12, 13, 14] with a Markov-chain Monte-Carlo (MCMC) sampling scheme to obtain the posterior uncertainty. As an alternative to the MCMC approach, this paper introduces the variational inference approach to Bayesian SSI. It will be shown how suitable approximations to the posterior distributions over the modal properties can be recovered, with discussion on the ability of the method to accurately represent the uncertainty over these properties and possible avenues to improve future predictions. A key advantage of this variational approach, is the overall reduction in computational load associated with MCMC.

2 RELATED WORK

Perspectives on Bayesian system identification have existed for a number of years and over the last decade many attempts to formulate Bayesian/Probabilistic OMA techniques have arisen in the literature taking a variety of approaches [15, 16, 9]. The most prominent in the field of OMA is the work of Au et.al., summarised in [17, 18], that develop a Bayesian Fast Fourier transform (FFT) algorithm. This algorithm has subsequently led to a body of works labelled Bayesian OMA or 'BAYOMA' [19, 20]. BAYOMA can be seen as the process of coupling the Bayesian FFT algorithm with a familiar modal analysis technique, such as Frequency Domain Decomposition [21], to recover a most-probable value and a representation of the uncertainty in the form of a coefficient of variation. This method, although Bayesian in its initial formulation, appears to reduce to a maximum likelihood approach since prior information is not employed [22]. In contrast, the method proposed by this paper will aim to recover a suitable approximation in the form of a surrogate posterior, using flexible but informative priors.

3 THEORY

3.1 Covariance-Driven SSI

Given standard derivations of covariance-driven SSI [23, 1, 24], it is widely known that for a classically defined, output-only state-space model, the extended observability ($\mathcal{O}$) and controllability ($\mathcal{C}$) matrices can be computed using the singular value decomposition (SVD)

$$\mathbf{\Sigma}_{ff}^{-1/2}\mathbf{\Sigma}_{fp}\mathbf{\Sigma}_{pp}^{-\mathsf{T}/2} \;=\; \mathbf{V}_1\mathbf{\Lambda}\mathbf{V}_2^{\mathsf{T}} \;\backsimeq\; \breve{\mathbf{V}}_1\breve{\mathbf{\Lambda}}\breve{\mathbf{V}}_2^{\mathsf{T}} \tag{1}$$

where $\mathbf{\Sigma}_{fp}$ is a block cross-covariance matrix between Hankel matrices of the future and past, and $\mathbf{\Sigma}_{ff}$ and $\mathbf{\Sigma}_{pp}$ are block auto-covariance matrices, derived using matrices of lags in the form of past and future Hankel matrices, $\mathbf{Y}_p = \mathbf{Y}_{0|2j-1}$ and $\mathbf{Y}_f = \mathbf{Y}_{j|2j-1}$. The vectors $\mathbf{V_1}$ and $\mathbf{V}_2$ correspond to the left and right singular vectors of the SVD, respectively and $\breve{\mathbf{\Lambda}}$ neglects sufficiently small singular values (equivalent to the canonical correlations) in $\mathbf{\Lambda}$ such that the resultant state vector has the dimension $d = \dim(\breve{\mathbf{\Lambda}})$. Consequently, the cross-covariance matrix, $\mathbf{\Sigma}_{fp}$, can be decomposed into a product of the extended observability and controllability matrices using $\mathbf{\Sigma}_{fp} = \mathcal{O}\mathcal{C}$, such that

$$\mathcal{O} \;=\; \mathbf{\Sigma}_{ff}^{1/2}\breve{\mathbf{V}}_1\breve{\mathbf{\Lambda}}^{1/2} \quad , \quad \mathcal{C} \;=\; \breve{\mathbf{\Lambda}}^{1/2}\breve{\mathbf{V}}_2^{\mathsf{T}}\mathbf{\Sigma}_{pp}^{\mathsf{T}/2} \tag{2}$$

where rank($\mathcal{O}$) = rank($\mathcal{C}$) = d.

3.2 Prob-SSI

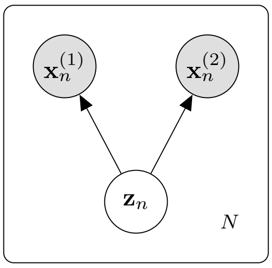

Figure 1: Graphical model of probabilistic CCA

In [8] it was shown that through simple substitution of the SVD (equivalent to CCA) in equation 1 with probabilistic CCA [10], covariance-driven SSI can be reformulated probabilistically. The latent variable model used to make this substitution is shown graphically in Figure 1 and described mathematically by equations,

$$\mathbf{z}_n \sim \mathcal{N}(0, \mathbb{I}) \tag{3}$$

$$\mathbf{x}_n^{(m)}|\mathbf{z}_n \sim \mathcal{N}(\mathbf{W}^{(m)}\mathbf{z}_n + \boldsymbol{\mu}^{(m)}, \boldsymbol{\Sigma}^{(m)}) \tag{4}$$

$$\mathbf{x}_n|\mathbf{z}_n \sim \mathcal{N}(\mathbf{W}\mathbf{z}_n + \boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{5}$$

where $\mathcal{N}(\mu, \Sigma)$ corresponds to a Gaussian distribution characterised by a mean μ and covariance Σ, $\mathbb{I}$ represents an identity matrix, $\boldsymbol{\Sigma}$ is a block-diagonal covariance matrix with $\boldsymbol{\Sigma}^{(1)}$ and $\boldsymbol{\Sigma}^{(2)}$ on its diagonal, $\boldsymbol{\mu} = [\boldsymbol{\mu}^{(1)}; \boldsymbol{\mu}^{(2)}]$ and $\mathbf{W}^{(m)}$ is a linear transformation on the latent variable $\mathbf{z}_n$ where $\mathbf{W} = [\mathbf{W}^{(1)}; \mathbf{W}^{(2)}]$. In this probabilistic form, assuming $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ are $\mathbf{Y}_f \in \mathbb{R}^{D_1 \times N}$ and $\mathbf{Y}_p \in \mathbb{R}^{D_2 \times N}$ respectively, the maximum likelihood estimates (MLE) of the weight matrices can be shown to be equivalent to the observability matrix and controllability matrix transposed.

$$\hat{\mathbf{W}}^{(1)} = \boldsymbol{\Sigma}_{ff}^{1/2}\mathbf{V}_1\mathbf{P}^{1/2}\mathbf{R} = \mathcal{O} \tag{6}$$

$$\hat{\mathbf{W}}^{(2)} = \boldsymbol{\Sigma}_{pp}^{1/2}\mathbf{V}_2\mathbf{P}^{1/2}\mathbf{R} = \mathcal{C}^{\mathsf{T}} \tag{7}$$

where $\mathbf{P}$ is a diagonal matrix of the largest d canonical correlations and $\mathbf{R}$ is an arbitrary square rotation matrix of size d. Although often unnecessary[3], the arbitrary rotation matrix $\mathbf{R}$ can be recovered through a simple post-processing step [25][4].

4 BAYESIAN STOCHASTIC SUBSPACE IDENTIFICATION

4.1 Bayesian CCA

Developing a hierarchical extension to PCCA, Wang [12] and Klami and Kaski [13] proposed Bayesian forms of CCA by introducing prior distributions over the model parameters. This extension to the model is shown graphically in Figure 2. The distribution over the likelihood and the prior distribution over the latent variable $\mathbf{z}$ are the same as in PCCA and are exactly as

[3]This arbitrary rotation is often ignored as its omission still results in data being transformed into the relevant subspace.

[4]Note the article contains an errata to the appendix with the relevant derivation.

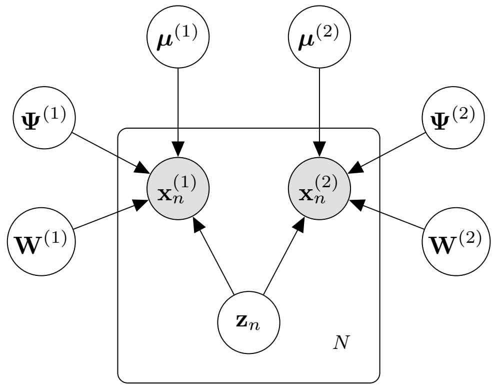

Figure 2: Graphical model of Bayesian CCA

written in equations (3 - 5), whilst the prior distributions over the model parameters are given by,

$$\begin{aligned} \mathbf{w}_i &\sim \mathcal{N}(\boldsymbol{\mu}_{w_i}, \sigma_{w_i}^2 \mathbb{I}) && (8) \\ \boldsymbol{\Psi} &\sim \mathcal{W}(\mathbf{K}_0^{-1}, \nu_0) && (9) \\ \boldsymbol{\mu} &\sim \mathcal{N}(0, \sigma_\mu^2 \mathbb{I}) && (10) \end{aligned}$$

where $\mathbf{w}_i$ denotes the ith column of the full weight matrix $\mathbf{W}$, where each column of $\mathbf{W}$ is considered independent and therefore assumed to have its own prior of the same form, $\mathcal{W}$ is shorthand notation for a Wishart distribution, and the priors over the mean $\boldsymbol{\mu}$ and the precision matrices $\boldsymbol{\Psi}^{(m)}$ are conventional conjugate priors. Note here that unlike the PCCA model, precision matrices are used as they simplify the variational model. Moreover, also note that compared to [13] and [12], the sparsity inducing (ARD) prior over the independent columns of the weight matrix has been removed.

4.2 Bayesian SSI

Akin to the derivation of Prob-SSI, CCA can be replaced with Bayesian CCA in SSI, such that the posterior estimates for the weights $\mathbf{W}^{(1)}$, $\mathbf{W}^{(2)}$ are analogous to the observability matrix and controllability matrix transposed. This results in a Bayesian interpretation of SSI and thus, an inference problem. There are several methods available to compute the posteriors, such as MCMC, Laplace approximation and variational approximation. In this paper, a variational inference scheme, similar to that employed by Wang, is used to recover analytical approximations to the posterior.

4.3 Variational Bayes (Inference)

Variational Bayes [26, 27] is often described as an extension to the expectation-maximisation (EM) algorithm, from a maximum-a-posteriori (MAP) estimation of the most-probable value to a more complete Bayesian approach capable of computing an approximate (surrogate) posterior over the parameters and latent variables. For Bayesian SSI, using a mean field approach, the surrogate posterior $q_\phi(\mathbf{z})$ was selected as having the following factorised form

$$q(\mathbf{z}, \boldsymbol{\mu}, \boldsymbol{\Psi}, \mathbf{W}) = q(\mathbf{z}) \prod_{i=1}^{d} q(\mathbf{w}_i) q(\boldsymbol{\mu}) q(\boldsymbol{\Psi}) \tag{11}$$

such that

$$q(\mathbf{z}_n) = \mathcal{N}\left(\mathbf{z}_n|\boldsymbol{\mu}_{\mathbf{z}_n}, \boldsymbol{\Sigma}_{\mathbf{z}_n}\right) \tag{12}$$

$$q(\boldsymbol{\Psi}) = \mathcal{W}\left(\boldsymbol{\Psi}|\tilde{\mathbf{K}}^{-1}, \tilde{\boldsymbol{\nu}}\right) \tag{13}$$

$$q(\mathbf{w}_i) = \mathcal{N}\left(\mathbf{w}_i|\boldsymbol{\mu}_{\mathbf{w}_i}, \boldsymbol{\Sigma}_{\mathbf{w}_i}\right) \tag{14}$$

$$q(\boldsymbol{\mu}) = \mathcal{N}\left(\boldsymbol{\mu}|\boldsymbol{\mu}_{\boldsymbol{\mu}}, \boldsymbol{\Sigma}_{\boldsymbol{\mu}}\right) \tag{15}$$

where $\boldsymbol{\phi} = \{\boldsymbol{\mu}_{\mathbf{z}_n}, \boldsymbol{\Sigma}_{\mathbf{z}_n}, \tilde{\mathbf{K}}^{-1}, \tilde{\boldsymbol{\nu}}, \boldsymbol{\mu}_{\mathbf{w}_i}, \boldsymbol{\Sigma}_{\mathbf{w}_i}, \boldsymbol{\mu}_{\boldsymbol{\mu}}, \boldsymbol{\Sigma}_{\boldsymbol{\mu}}\}$ are the moments of the surrogate distributions. Given that the intention is to obtain a surrogate posterior as close to the real posterior as possible, an optimisation problem arises seeking to minimise a dissimilarity function between the two distributions. A common approach, and the one employed here, is the minimisation of the Kullback-Leibler (KL) divergence. The KL divergence is typically chosen as it makes the minimisation process tractable. Rather than directly minimising the KL divergence, it is common practice to define the problem in terms of the evidence lower bound (ELBO), which when maximised, is equivalent to minimising the KL divergence. Obtaining the ELBO is also beneficial as it can be used to monitoring convergence. The ELBO is defined as

$$\mathcal{L}(\boldsymbol{\phi}) = \mathbb{E}_{q(\mathbf{z},\boldsymbol{\theta})}\left[\log p(\mathbf{x}, \mathbf{z}, \boldsymbol{\theta}) - \log q_{\boldsymbol{\phi}}(\mathbf{z}, \boldsymbol{\theta})\right] \tag{16}$$

where $\boldsymbol{\theta} = \{\boldsymbol{\mu}, \boldsymbol{\Psi}, \mathbf{W}\}$. Given the ELBO, and using coordinate ascent variational inference (CAVI) [27], the distributional form of each component of the surrogate posterior and the accompanying statistical moment equations, can be determined in full using

$$\hat{\phi} = \arg\min_{\boldsymbol{\phi}} \mathbb{KL}(q_{\boldsymbol{\phi}}(\mathbf{z}, \boldsymbol{\theta})||p(\mathbf{x}, \mathbf{z}, \boldsymbol{\theta})) = \arg\max_{\boldsymbol{\phi}} \mathcal{L}(\boldsymbol{\phi}) \tag{17}$$

The equations and their derivations are long and tedious but necessary, requiring numerous pages of algebra. As such, they are omitted here for space and the reader's sanity, but can be derived if needed.

4.4 Modal Parameter Recovery

Once posterior distributions to the observability and controllability matrices have been approximated using variational inference, the final phase is the recovery of the posterior distributions to the modal properties. However, propagating uncertainty in closed form through standard matrix problems and eigen-decompositions, two mathematical operations necessary for modal parameter recovery, is currently not possible, meaning there is no tractable solution to the posterior. Therefore, sampling based (Monte Carlo) methods are needed to obtain approximations to the posterior distributions of the modal properties.

5 RESULTS AND DISCUSSION

To assess the identification performance of Bayesian SSI, response data was generated using a simulated three degree of freedom (DOF) linear dynamic system with natural frequencies $\omega_n = \{4.74, 6.44, 10.65\}$ Hz and damping ratios $\zeta_n = \{0.0015, 0.0020, 0.0033\}$, given a white noise excitation. The system was simulated at a sample rate of 1000Hz, generating 8192 data points. Both SSI and Bayesian SSI were applied to the dataset and used to identify the modal properties of the system, assuming a model order of 20 (equivalent to 10 unique modes as poles exist as conjugate pairs).

The prior mean for the weights $\boldsymbol{\mu}_{w_i}$ was set equal to those obtained through SSI. This was deemed a sensible and logical approach as the truth is expected to lie in somewhat close proximity to the result from SSI. This significantly improves computational efficiency whilst also reducing the likelihood of getting trapped in a local minima, a common trait in iterative optimisation methods such as EM and variational coordinate ascent. The remaining prior parameters were set as: $\sigma^2_{w_i} = 10^{-4}$, $\mathbf{K_0} = 10^{-3}\mathbb{I}$, $\nu_0 = D + 1$ and $\sigma^2_{\mu} = 1$, where D is the dimension of the dataset. Once posterior estimates for the observability and controllability matrices were recovered using variational inference, 10 000 samples were drawn from the distributions and used to obtain estimates for the posterior distributions of the modal properties.

Following application of both methods, full posteriors over $\mathcal{O}$ and $\mathcal{C}$ were obtained and samples drawn to obtain distributions over the modal properties. Figure 3 illustrates the prior and posterior distributions over the natural frequencies, with indicators to the SSI found result and the truth, and similarly, Figure 4 shows the same but for the damping ratios. Unsurprisingly, SSI performs as expected finding estimates relatively close to the true values but with some deviation seen. Bayesian SSI appears to perform well, recovering posterior distributions with means close to the truth, whilst also having sufficient variance such that the distribution encompasses both the SSI solution and the true values, given the priors. Interestingly, the priors over the damping ratios give greatly inflated estimates, which is expected as damping is strongly influenced by the magnitude of the dynamic response and is often hard to recover. Nevertheless, the posterior estimates do appear to converge towards the true damping ratios, as expected.

Figures 5 and 6 also shows the prior and posterior distributions for the mode shapes of the system. Generally, there seems to be good agreement between the posterior and the truth, as the expectation of the distributions appearing to align well with the true modes. Another promising observation is the removal of the spurious modes, seen in the second mode shape prior.

Lastly, Figure 7 shows the plot of obtained posterior poles, with a secondary plot better highlighting the pole estimates near the SSI solution and truth. This plot provides an alternative look at the posterior and is useful as it highlights the presence of posterior poles that lie in the right-hand plane, corresponding to unrealistic estimates for the damping ratio ($\zeta < 0$). The management of this is an open research question, one discussed later in the closing remarks. Nevertheless, regardless of whether a Bayesian approach is used, nothing constrains classical SSI to enforce ($\zeta > 0$).

An important feature worthy of discussion, is the apparent distributional shape of the recovered properties. Locally, it could be said that the individual modal properties appear to follow a Gaussian distribution and could be approximated as such, however, one can reason that the expected distributions cannot be analytically Gaussian due to the inherent mathematical operations needed to recover them. Furthermore, non-Gaussian distributions over the natural frequencies and damping ratios are expected, as the two modal parameters are physically bounded $(0, \infty)$ to ensure a real and stable dynamic system. As a Gaussian has support over the whole real line $(-\infty, \infty)$, estimates to the posterior may result in unrealistic, nonphysical estimates of the modal properties. Again made clear through Figure 7, where the distinction between poles corresponding to ($\zeta > 0$) and ($\zeta < 0$) is made.

It is also worth noting there are visible differences in the variance of the prior distributions over all modal parameters. Despite the columns of the prior weight matrix being normally distributed with the same variance, the same non-linear operations needed to obtain the parameters, such as the eigen-decomposition, directly alter the mapping of this variance. This

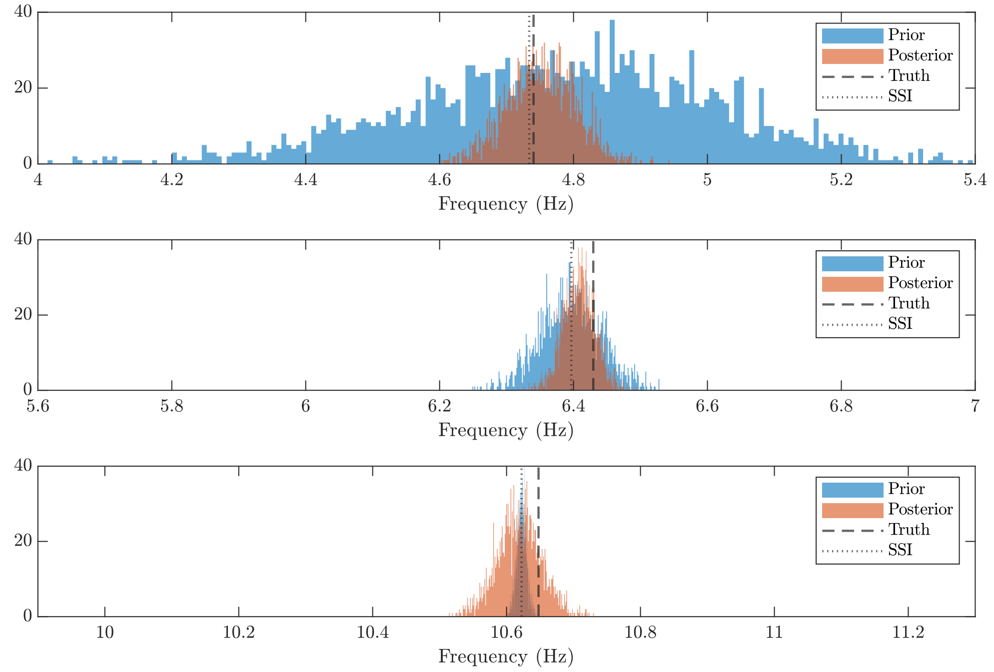

Figure 3: Histograms of the prior and posterior distributions over the three natural frequencies, constructed from the posterior samples.

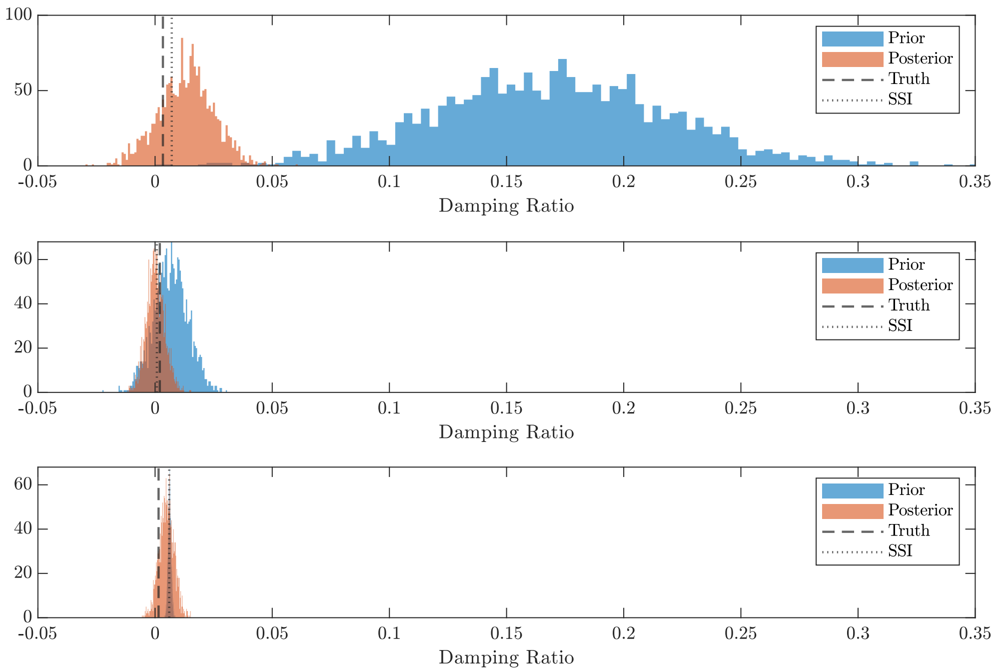

Figure 4: Histograms of the prior and posterior distributions over the three damping ratios, constructed from the posterior samples.

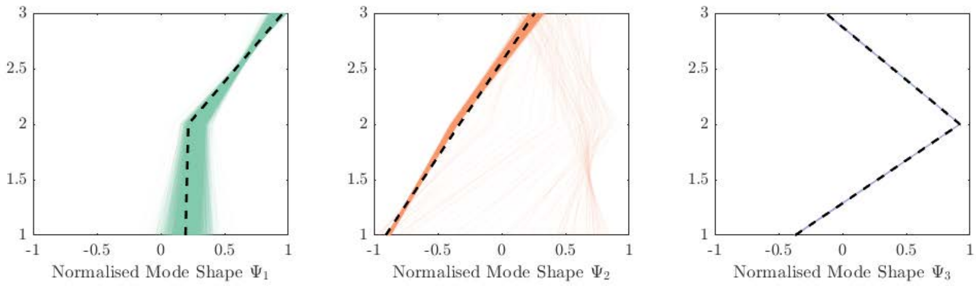

Figure 5: Plot of the prior samples of the mode shapes. The mode shapes have been normalised using $\|\cdot\|_2$

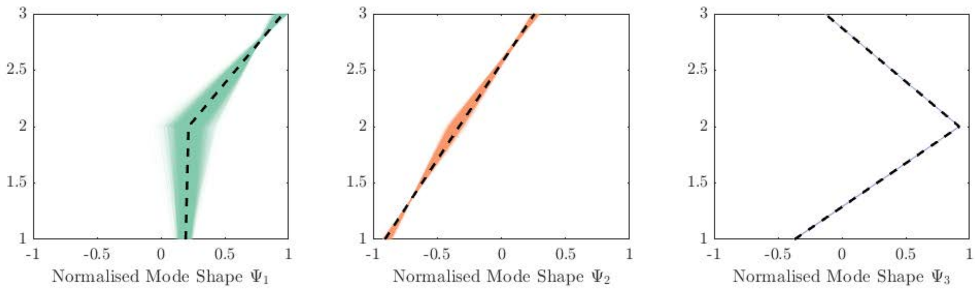

Figure 6: Plot of the posterior samples of the mode shapes. The mode shapes have been normalised using $\|\cdot\|_2$

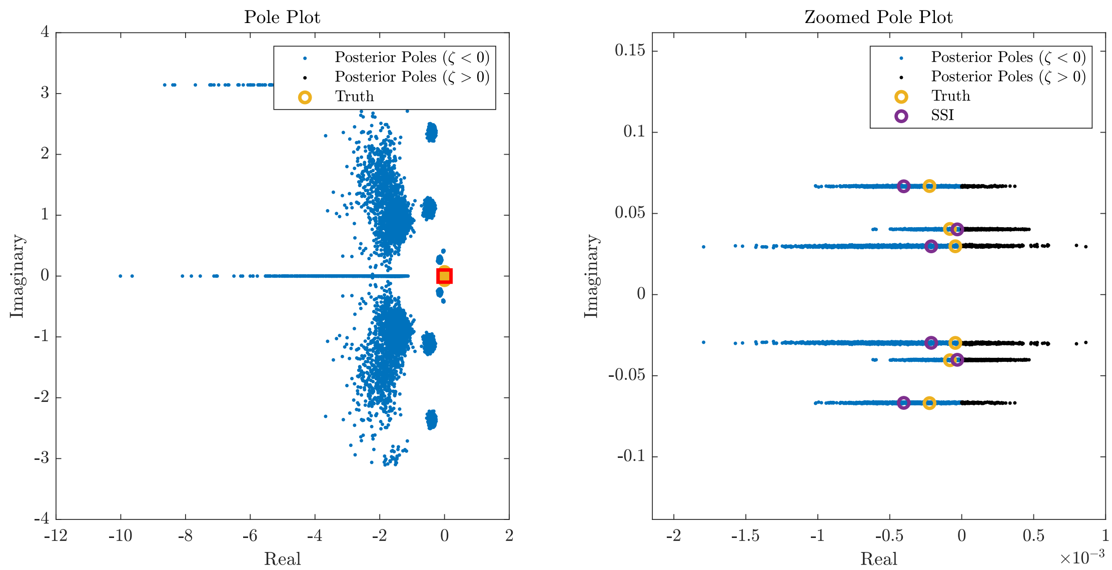

Figure 7: Posterior samples of the system poles, (left) shows the full set of recovered poles, (right) provides a clearer perspective of the poles close to the truth.

behaviour was largely expected from the eigen-decomposition, as all the available 'flexibility' is primarily used in finding the first component. This reduces the effective dimension each time a subsequent orthogonal eigenvector is added. Consequently, this results in a reduction to the variance of the prior for the increasing mode number. This does have its limitations, as can be seen from the prior and posterior distributions of the third mode. Despite the variance of the posterior increasing with the inclusion of data, there is little movement of the posterior mean toward the truth, as the tightness of the prior heavily influences the estimate.

Considering the above, two interesting questions arise. Can more suitable priors be constructed such that a real and stable dynamic system is guaranteed? Furthermore, does a closed form, analytical solution exist for a true or surrogate posterior? One could hypothesise this could be Gamma or log normal for example. Overall, the ability to recover posterior distributions over the modal properties in this way, despite the limitations highlighted, is an inherently useful feature of Bayesian SSI that provides a powerful insight into the uncertainty in both the data and the SSI algorithm.

6 CLOSING REMARKS

This paper presented a novel variational Bayesian approach to covariance-driven stochastic subspace identification (SSI) for use in operational modal analysis (OMA). It was shown that by leveraging the presence of canonical correlation analysis (CCA) in SSI, a replacement with Bayesian CCA can be made, thus reformulating SSI as an inference problem. As a result, approximations to the posterior distributions over the observability and controllability matrices can be recovered in closed form using a variational inference scheme. Moreover, by extension, approximations to the posterior distributions over the modal properties can also be recovered through sampling; necessary due to the intractability of the posterior once operated upon by non-linear mathematical operations, such as an eigen-decomposition. The method was successfully applied to response data from a simulated three degree-of-freedom linear dynamic system and unsurprisingly, demonstrated comparable identification performance to SSI, whilst also recovering suitable approximations to the posteriors that include both the truth and SSI estimates of the modal properties.

Following this body of work, future work will aim to address the issue of recovering physically meaningful posterior estimates to the modal properties. This may include the exploration and development of a novel approach to embedding physical understanding into the priors in this practical context. The authors will also investigate the use of the now available posterior uncertainty to aid model order selection, without the need for consistency diagrams. The role and inclusion of the sparsity inducing prior, used by Wang [12] but not employed here or discussed in depth, may yet play an interesting role in automatic model order selection for modal analysis.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the Engineering and Physical Sciences Research Council (EPSRC), UK through grant number EP/W002140/1. The authors also thank Ramboll Energy for their support in particular Ulf Tygesen with whom they work closely. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

REFERENCES

[1] Peter Van Overschee and Bart De Moor. *Subspace Identification for Linear Systems*. Springer, 1996.

[2] Bart Peeters and Guido De Roeck. Stochastic system identification for operational modal analysis: A Review. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 123(4):659–667, 2001.

[3] Carlo Rainieri and Giovanni Fabbrocino. *Operational Modal Analysis of Civil Engineering Structures: An Introduction and Guide for Applications*. Wiley, 2014.

[4] Rune Brincker and Carlos E. Ventura. *Introduction to Operational Modal Analysis*. John Wiley and Sons, Ltd, Chichester, UK, aug 2015.

[5] A. W. Phillips, R. J. Allemang, and D. L. Brown. Autonomous modal parameter estimation: Methodology. In *Modal Analysis Topics, Volume 3*, pages 363–384, New York, NY, 2011. Springer New York.

[6] R. J. Allemang and D. L. Brown. A unified matrix polynomial approach to modal identification. *Journal of Sound and Vibration*, 211:301–322, 1998.

[7] Edwin P.B. Reynders. Uncertainty quantification in data-driven stochastic subspace identification. *Mechanical Systems and Signal Processing*, 151:107338, 2021.

[8] Brandon J O'Connell, Elizabeth J Cross, and Timothy J Rogers. Robust probabilistic canonical correlations for stochastic subspace identification. In *9th IOMAC - International Operational Modal Analysis Conference*, 2022.

[9] Timothy J Rogers. *Towards Bayesian System Identification: With Application to SHM of Offshore Structures*. PhD thesis, University of Sheffield, 2018.

[10] Francis R Bach and Michael I. Jordan. A probabilistic interpretation of canonical correlation analysis. *Dept. Statist., Univ. California, Berkeley, CA, Tech. Rep*, 688:1–11, 2005.

[11] B J O'Connell, E J Cross, U T Tygesen, and T J Rogers. Bayesian canonical correlations for stochastic subspace identification. In *30th ISMA, 9th USD - International Symposium on Modal Analysis and International Conference on Uncertainty in Structural Dynamics*, 2022.

[12] Chong Wang. Variational Bayesian Approach to Canonical Correlation Analysis. *IEEE Transactions on Neural Networks*, 18(3):905–910, 2007.

[13] Arto Klami and Samuel Kaski. Local dependent components. In *ACM International Conference Proceeding Series*, volume 227, pages 425–432, 2007.

[14] Arto Klami, Seppo Virtanen, and Samuel Kaski. Bayesian canonical correlation analysis. *Journal of Machine Learning Research*, 14(1):965–1003, 2013.

[15] Omid Sedehi, Lambros S. Katafygiotis, and Costas Papadimitriou. Hierarchical Bayesian operational modal analysis: theory and computations. *Mechanical Systems and Signal Processing*, 140:106663, 2020.

[16] Heung Fai Lam, Jun Hu, and Jia Hua Yang. Bayesian operational modal analysis and Markov chain Monte Carlo-based model updating of a factory building. *Engineering Structures*, 132:314–336, 2017.

[17] Siu Kui Au. Fast Bayesian ambient modal identification in the frequency domain, Part I: Posterior most probable value. *Mechanical Systems and Signal Processing*, 26(1):60–75, 2012.

[18] Siu Kui Au. Fast Bayesian ambient modal identification in the frequency domain, Part II: Posterior uncertainty. *Mechanical Systems and Signal Processing*, 26(1):76–90, 2012.

[19] James Mark William Brownjohn, Alison Raby, Siu Kui Au, Zuo Zhu, Xinrui Wang, Alessandro Antonini, Athanasios Pappas, and Dina D'Ayala. Bayesian operational modal analysis of offshore rock lighthouses: Close modes, alignment, symmetry and uncertainty. *Mechanical Systems and Signal Processing*, 133:106306, 2019.

[20] Siu Kui Au, James M.W. Brownjohn, and John E. Mottershead. Quantifying and managing uncertainty in operational modal analysis. *Mechanical Systems and Signal Processing*, 102:139–157, 2018.

[21] Rune Brincker, Lingmi Zhang, and Palle Andersen. Modal identification of output-only systems using frequency domain decomposition. *Smart Materials and Structures*, 10(3):441–445, 2001.

[22] Edwin Reynders, Kristof Maes, Geert Lombaert, and Guido De Roeck. Uncertainty Quantification in Operational Modal Analysis with Stochastic Subspace Identification: Validation and Applications. *Mechanical Systems and Signal Processing*, 66-67:13–30, 2016.

[23] Tohru Katayama. *Subspace Methods for System Identification: A Realisation Approach*. Springer, 2005.

[24] B. Peeters, G. De Roeck, and P. Andersen. Stochastic system identification: uncertainty of the estimated modal parameters. In *IMAC XVII - 17th International Modal Analysis Conference*, volume 32, pages 231–237, 1999.

[25] Cédric Archambeau, Nicolas Delannay, and Michel Verleysen. Robust probabilistic projections. In *Proceedings of the 23rd International Conference on Machine Learning*, volume 148, pages 33–40, Pittsburgh, 2006.

[26] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, Cambridge, 2006.

[27] Kevin P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023.