Proceedings of ISFOG 2025
5TH INTERNATIONAL SYMPOSIUM ON
FRONTIERS IN OFFSHORE GEOTECHNICS
Nantes, France | June 9-13 2025

ISBN 978-2-85782-758-0

Comparison of machine learning models in a data-driven method for design of laterally loaded monopiles embedded in layered clay sites

I. Kamas*
University of Oxford, Oxford, Oxfordshire, United Kingdom
H.J. Burd, B.W. Byrne
University of Oxford, Oxford, Oxfordshire, United Kingdom
**ioannis.kamas@eng.ox.ac.uk*

ABSTRACT: Monopiles are widely used as foundations for offshore wind turbine support structures. The PISA project introduced a 1D model, known as the 'PISA design model', to predict the monotonic response of laterally loaded monopiles. Offshore wind turbine sites typically comprise distinct soil layers with different geotechnical properties. For layered soil configurations, the PISA design model employs soil reaction curves derived from 3D finite element analyses of homogeneous soils, assuming independent behaviour for each layer. However, this approach does not account for interactions between soil layers, which can significantly impact monopile lateral behaviour. In the current work, an extended version of the PISA design model—referred to as the 'data-driven 1D design model'—has been developed. This model uses machine learning techniques to define the soil reaction curves, allowing for direct calibration across various soil types and layering configurations. This paper presents a calibration procedure for layered clay sites, considering variations in clay strength, stiffness, and layer thickness. Four machine learning models—sparse Gaussian process regression, artificial neural network regression, support vector regression and eXtreme Gradient Boosting—are assessed to determine the most effective approach.

Keywords: Laterally loaded monopiles; Machine learning; Surrogate model; Layered clay soil profiles

1 INTRODUCTION

Monopile foundations are commonly used to support offshore wind turbine (OWT) structures. Current design methods for OWT monopile foundations typically employ simplified models. The *p-y* method, commonly used for laterally loaded piles, models the pile as a beam and represents the lateral soil response using non-linear *p-y* curves.

The PISA (Pile Soil Analysis) project introduced an alternative one-dimensional (1D) model referred as the 'PISA design model' to predict the monotonic behaviour of laterally loaded monopiles (Byrne et al., 2020; Burd et al., 2020a). This model enables rapid calculations by using 'soil reaction curves' which are derived from detailed model calibration by three-dimensional (3D) finite element analysis (FEA).

OWT monopiles are typically embedded in layered soils. The nature of the soil layering can significantly influence the performance of the foundation (e.g. Yang and Jeremic 2005). Although the PISA design model is calibrated for homogeneous soil profiles, it has also been applied to layered soils using an approach termed as the 'Independent Layer Method' (ILM), in which soil reaction curves determined from 3D FEA of homogeneous soils are employed within a layered soil configuration (Burd et al., 2020b). While ILM has been shown to be an effective approach, it does not account for interactions between adjacent soil layers. These interactions can have a significant influence on the behavior of the monopile.

In the current work, an alternative version of the PISA design model, called the 'data-driven 1D design model', has been developed (Kamas 2024). This model utilises machine learning (ML) techniques to establish soil reaction curves, enabling direct calibration (via 3D FEA) across different soil types and layering configurations. The data-driven 1D design model therefore explicitly incorporates the interactions between adjacent soil layers via the calibration analyses. This paper outlines a calibration procedure specifically for layered clay sites, accounting for variations in clay strength, stiffness, and thickness of each layer. Four ML regression models—sparse Gaussian process regression, artificial

neural network regression, support vector regression and eXtreme Gradient Boosting—are evaluated to identify the most effective method.

2 DATA-DRIVEN APPROACH

The data-driven 1D design model builds on the 1D finite element framework of the PISA design model but uses a different parametrisation for the soil reaction curves. The data-driven 1D design model employs soil reaction curves that are defined using 8-parameter Piecewise Cubic Hermite Interpolating Polynomial (PCHIP) functions. These functions closely replicate the numerical soil reaction curves (see Fig. 1). The numerical soil reaction curves employed for clay sites in the data-driven 1D design model adopt the dimensionless forms presented in Table 1. Knot point force/moment data define the soil reaction curves. ML techniques are employed to predict these knot point data for unseen scenarios. Key parameters such as soil stiffness, strength, pile aspect ratio and soil layering serve as training features for the model.

Table 1. Dimensionless forms of parameters adopted in the data-driven 1D design model for clay sites (Kamas, 2024). Here, D is the pile diameter, and s_u is the undrained shear strength corresponding to triaxial compression tests. Variables with a bar (e.g. $\bar{p}$, $\bar{v}$) represent the normalised form of the corresponding soil reaction components.

Normalised variable	Clay framework
Distributed lateral load, $\bar{p}$	$p/(s_u D)$
Lateral displacement, $\bar{v}$	v/D
Distributed moment, $\bar{m}$	$m/(s_u D^2)$
Pile cross-section rotation, $\bar{\psi}$	ψ
Base horizontal force, $\bar{H}_B$	$H_B/(s_u D^2)$
Base moment, $\bar{M}_B$	$M_B/(s_u D^3)$

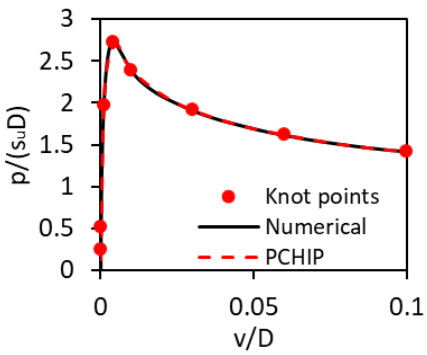

Figure 1. An example of a PCHIP spline soil reaction curve for use in the data-driven 1D design model demonstrates an excellent match with numerical soil reaction curves derived from 3D finite element analysis.

3 GENERATION OF TRAINING DATA

3.1 Layered configurations

The layered configurations considered in the current work are shown in Fig. 2. The term 'matrix' in Fig. 2 refers to the main soil type, while 'layer' describes any embedded or base layers. In Case A, the model features two layers, while Case B introduces a relatively thin embedded layer within a homogeneous matrix. Both configurations allow for variability in soil strength and stiffness properties within the matrix and layers. The normalised thickness of each soil layer penetrated by the pile (z_L), relative to the pile's embedment length (L), can vary between 0 and 1.

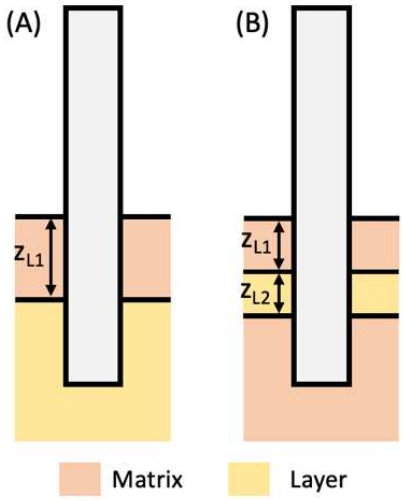

Figure 2. *Layered soil configurations considered: (a) case A, two layers; (b) case B, single embedded layer.*

Ground models for these layered clay soil profiles are constructed using data from the reference homogeneous soil models shown in Fig. 3, which depicts the profiles of small-strain shear modulus (G_0) and undrained shear strength (s_u). The properties of each clay layer are based on the PISA representative offshore glacial clay till site model described by Byrne et al. (2020), with a submerged unit weight of $\gamma' = 11.38\,\mathrm{kN/m^3}$. The reference homogeneous models are determined using an approach in which homogeneous soil is considered to have been previously consolidated solely by the weight of an ice sheet with a thickness of d_{ice}, using one-dimensional compression theory. Soil properties such as s_u and G_0, as shown in Fig. 3, are determined following the method outlined by Kamas et al. (2023). The clay soil profiles used in this study correspond d_{ice} ranging from 25m to 150m.

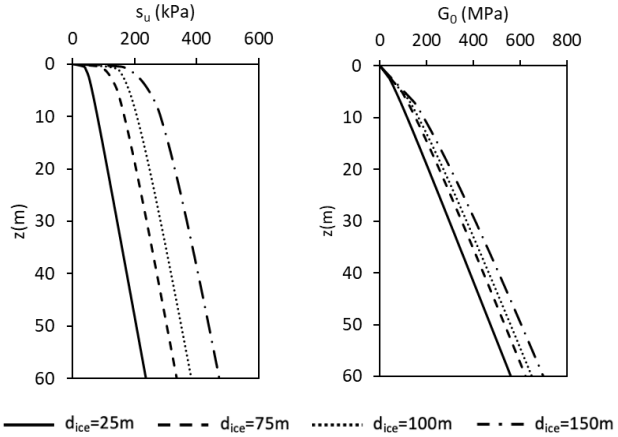

Figure 3. Profiles of s_u and G_0 used in the reference homogeneous soil models for the layered soil profile configurations shown in Fig. 2.

3.2 Pile geometry calibration space

The pile geometries defining the calibration space for the data-driven 1D design model are detailed in Table 2. These calibration piles feature geometries typical of offshore configurations, with L/D between 2 and 5 and D between 8m and 12m. The pile wall thickness is $t = D/110$; it is selected to prevent unrealistically high bending stresses developing in the pile wall.

Table 2. Pile geometry calibration space. L is the embedded length; D is the pile diameter; h is the load eccentricity (height of load application above seabed level).

Pile	$L(m)$	$D(m)$	$h(m)$
LC1	16	8	40
LC2	16	8	120
LC3	24	12	60
LC4	24	12	180
LC5	40	8	40
LC6	40	8	120
LC7	60	12	60
LC8	60	12	180

3.3 Finite element analyses procedure

3D FEA was performed using PLAXIS 3D software (Brinkgreve et al., 2018) employing the NGI-ADP (total stress) constitutive model (Grimstad et al., 2012). The NGI-ADP model parameter values were determined on the basis of the representative offshore glacial clay till site in Byrne et al. (2020) with modifications to account for the strength and stiffness profiles in Fig. 3 (see Kamas et al. 2023).

Fig. 4 illustrates a typical finite element mesh employed in the calibration analyses. The mesh incorporates a plane of symmetry to halve the problem domain. The pile is modelled using plate elements, assuming linear elastic material, having a Young's modulus of $E = 200GPa$ and a Poisson's ratio of $\nu = 0.30$. The pile is treated as 'wished in place', meaning installation effects are disregarded. Interface elements are included around the pile to represent pile-soil interaction, based on the approach detailed in Kamas et al. (2023).

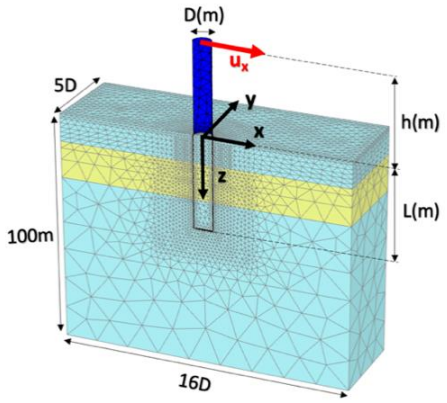

Figure 4. FE mesh, for pile LC5 embedded in three-layered soil profile Case B-IV (see Table 3).

3.4 Calibration finite element analyses

The calibration analyses must account for variations in soil properties and layer thickness across the selected pile calibration space for the Case A and Case B layered configurations in Fig. 2. A total of 152 FEA calibration analyses were conducted. The soil reaction curves extracted from the analyses were used to develop the data-driven 1D design model for the layered clay profiles of Fig. 2 and Table 3.

Table 3. Normalised thicknesses of soil layers penetrated by the pile, adopted for the calibration of layered soil profile FEA models (see Fig. 2).

Case A	z_{L1}/L	Case B	z_{L1}/L	z_{L2}/L
I	0.2	I	0.2	0.25
II	0.5	II	0.6	0.25
III	0.8	III	0.2	0.6
		IV	0.3	0.4
		V	0.2	0.1
		VI	0.6	0.1

For each profile, FEA was conducted for all pile configurations listed in Table 2. To account for variations in soil properties in the data-driven 1D design model, two calibration cases were applied to each profile. In the first case, the strength and stiffness of the soil 'matrix' were set to correspond to $d_{ice} = 25m$, while the soil 'layer' was assigned properties corresponding to $d_{ice} = 75m$, or vice versa. In the second case, the soil 'matrix' was assigned properties corresponding to $d_{ice} = 75m$, with the soil 'layers' corresponding to $d_{ice} = 150m$, or vice versa. Additionally, FEA was performed for calibration piles embedded in homogeneous soil profiles.

4 TRAINING INPUTS

Identifying the most influential inputs is key to train accurate ML models. To minimize the input set, dimensionless synthetic features (e.g., L/D, t/D, G_0/s_u, z_{L1}/L) were introduced. A heatmap-based correlation matrix and pair plots were employed to investigate relationships among key features. Spearman correlations and feature importance from an Extreme Gradient Boosting (XGB) model were computed and plotted. This combined approach ensures that the most impactful features are selected to train the ML models.

The inputs used by the ML models are as follows. The ratio (z/L) captures the depth variation of the knot points for $\bar{p}$ and $\bar{m}$, but is excluded for $\bar{H}_B$ and $\bar{M}_B$, which provide point predictions. The (L/D) value indicates pile slenderness, while (G_o/s_u) distinguishes soil reactions for the same (z/L) and (L/D). The (s_u/σ'_v) ratio, where σ'_v denotes the local vertical effective stress, differentiates clay profiles of varying strength. Layering effects are captured by (z_{Li}/L), where z_{Li} denotes the thickness of each soil layer penetrated by the pile. The number of (z_{Li}/L) inputs corresponds to the number of layers in the pile's embedded length. For example, in Case B (Fig. 2), two inputs (z_{L1}/L, z_{L2}/L) are applied.

5 REGRESSION METHODS

5.1 Machine learning models

Sparse Gaussian Process Regression (SGPR) (Snelson and Ghahramani, 2006) is an optimised version of Gaussian Process Regression (GPR), designed for large datasets. GPR models functions probabilistically, using a Gaussian distribution updated via Bayes' rule based on observed data. It provides both predictions and a built-in measure of uncertainty by quantifying the variance. The model is fully defined by a mean and covariance (kernel) function. SGPR reduces computational cost by using a subset of representative data points, called 'inducing points', while maintaining predictive accuracy.

Artificial Neural Network Regression (ANNR) predicts continuous values by adjusting weights between neurons based on input and target data. The model's architecture includes hidden layers, neurons, and their connections, while parameters like weights and biases are optimised to capture complex patterns (Goodfellow et al., 2016).

Support Vector Regression (SVR) (Smola and Schölkopf, 2004), estimates continuous values by finding the optimal hyperplane within a margin, called the 'epsilon-tube'. It minimizes errors while keeping most data within this margin. SVR handles both linear and non-linear problems using kernel functions to capture complex patterns.

Extreme Gradient Boosting (XGB), (Chen and Guestrin, 2016), uses a regression algorithm built upon decision tree concepts. Building on the original Gradient Boosting Machines (Friedman, 2001), XGB improves generalisation with L1/L2 regularisation, enhances gradient updates using second-order derivatives, and accelerates computation through parallelised tree construction.

5.2 Training procedure

A separate ML model is trained for each of the eight knot points for each soil reaction component ($\bar{p}$, $\bar{m}$, $\bar{H}_B$, $\bar{M}_B$). For $\bar{p}$ and $\bar{m}$, 6,080 data points are used to train each knot point, while 152 data points are used for $\bar{H}_B$ and $\bar{M}_B$. The dataset for each knot point is randomly split, with 80% used for training and the remaining 20% reserved for testing. Prior to training, the input features for each ML model are standardised by subtracting the mean and dividing by the standard deviation, both calculated from the training data. The procedures used to train the ML models for each knot point parameter, as described in Section 5.1, are outlined below:

Sparse Gaussian Process Regression: Due to the large dataset, training a standard GPR model resulted in prolonged computation times, exceeding 8 hours for the full set of $\bar{p}$ and $\bar{m}$ knot points. To mitigate this, an SGPR model was employed. The Variational Free Energy (VFE) method (Titsias, 2009) was implemented using the Gpflow Python library (Matthews et al., 2017). A zero mean function and the Matérn kernel ($\nu = 5/2$) were employed. Based on preliminary experiments, the model utilised 600 inducing variables, selected using the K-means algorithm (Hartigan and Wong, 1979). The hyperparameters were optimised by maximising the log marginal likelihood. For the knot points assocated with the base soil reaction curves ($\bar{H}_B$, $\bar{M}_B$), which involved significantly fewer data points (152 points per GPR model), standard GPR was applied.

Artificial Neural Network Regression: A feedforward neural network was employed for regression, implemented using the Keras library (Chollet, 2021). The rectified linear unit (ReLU) activation function was used. During model training, 20% of the training data (equivalent to 16% of the total dataset) is set aside as a validation set. In this study, the model is trained using the Mean Squared Error (MSE) as the loss function, employing the Adam optimiser. Preliminary experiments indicated that two hidden layers provided good prediction performance. The structure of the ANNR models and their

hyperparameters were optimised using a randomised search. The hyperparameters analysed included: (i) the number of neurons in each hidden layer, with options of 100 to 500; (ii) batch sizes of 16 to 512; (iii) dropout rates of 0.1 to 0.2; (iv) epoch counts of 25 to 300; and (v) learning rates of 0.01, 0.001, 0.002 and 0.005.

Support Vector Regression: SVR model was implemented using the Scikit-learn Python library (Pedregosa et al., 2011). A randomised search was conducted to determine the optimal set of hyperparameters for the model. The *C* (regularisation parameter) was varied between 10 and 1000. The RBF kernel was used, and the epsilon-tube value was explored in the range of 0.01 to 0.1.

Extreme Gradient Boosting: XGB model was implemented using the XGB library in Python (Chen and Guestrin, 2016). A randomised search was used to identify the optimal hyperparameters. The search evaluated various settings, including: (i) 'max_depth': ranging from 3 to 15, (ii) 'colsample_bytree': ranging from 0.5 to 0.9, (iii) 'learning_rate': ranging from 0.1 to 0.3, (iv) 'subsample': ranging from 0.5 to 1.0, (v) 'n_estimators': ranging from 100 to 1000, (vi) 'alpha': ranging from 0.1 to 1, (vii) 'gamma': ranging from 0 to 1, (viii) 'reg_lambda': ranging from 1 to 3, (viiii) 'reg_alpha': ranging from 0.1 to 1.0.

5.3 Assessment of the machine learning models

To evaluate the performance of the regression methods, two commonly used metrics in regression analysis were employed: the coefficient of determination (R^2) and mean squared error (MSE). Table 4 presents a comparison of the average R^2 scores for all knot points associated with the distributed lateral load, $\bar{p}$, and the distributed moment, $\bar{m}$, on both the training and test datasets. R^2 values typically range from 0 to 1, with higher values indicating a stronger correlation between the predicted and observed data. It can be seen that all the ML regression models examined provide satisfactory R^2 scores. However, the ANN, XGM and SGPR models clearly outperform the SVR model, yielding higher R^2 scores.

Table 4. Average R^2 regression performance metrics for the analysed ML regression models as applied to the knot point data for the distributed lateral load $\bar{p}$ and the distributed moment $\bar{m}$.

R^2 scores	SGPR	ANNR	XGB	SVR
$\bar{p}$ train data	0.932	0.920	0.943	0.798
$\bar{p}$ test data	0.915	0.908	0.892	0.783
$\bar{m}$ train data	0.907	0.903	0.927	0.840
$\bar{m}$ test data	0.905	0.894	0.897	0.828
Average	0.915	0.906	0.915	0.812

6 DESIGN APPLICATION

Two design cases (DLC1 and DLC2) were explored to demonstrate the predictive capability of the data-driven 1D design model, utilising the SGPR, ANNR and XGB models presented in Section 5 (that all have average R^2 scores above 0.9). Both cases consider a pile with $D = 9m$, $L = 30m$, and $h = 90m$; the L/D and h/D ratios were both previously unseen by the ML models. The G_0 and s_u profiles of the two design cases are presented in Fig. 5. DLC1 has a two-layer soil profile with unseen G_o/s_u, s_u/σ'_v, and z_{L1}/L, while DLC2 involves a three-layer profile with unseen G_o/s_u, s_u/σ'_v, z_{L1}/L, and z_{L2}/L.

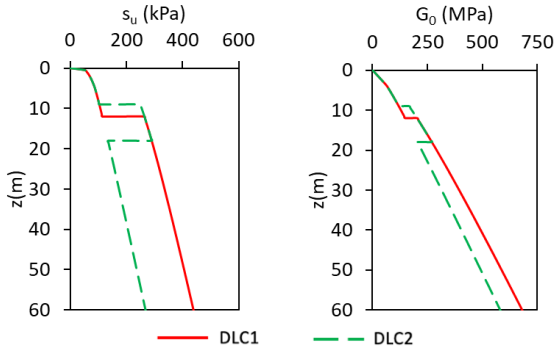

Figure 5. Profiles of s_u and G_0 used in the DLC1 and DLC2 design cases.

Table 5 presents horizontal force values at a ground-level displacement of $v_G = D/10000$, indicating the pile's response at small displacements range. Fig. 6 shows the pile's load-displacement response for ultimate reference ground-level displacements ($v_G = D/10$), calculated using the data-driven 1D design model, alongside validation data from 3D FEA simulations. The data-driven 1D design models using the SGPR, ANNR, and XGB ML models show good agreement with the 3D calculations. Differences among these three ML models are minimal, as reflected in their average R^2 scores presented in Table 4.

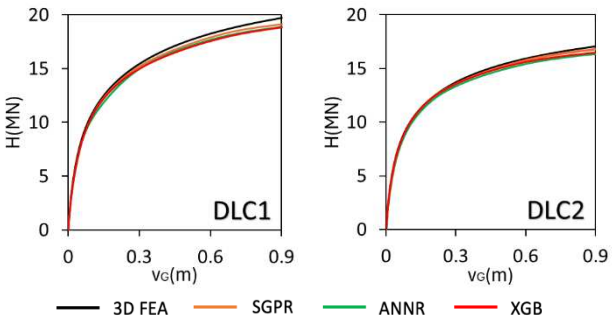

Figure 6. Comparison of monopile ultimate response between the validation 3D FEA and the data-driven 1D design models (using the ML regression models discussed in Section 5) for design cases DLC1 and DLC2.

Table 5. Values of the horizontal force at a ground level displacement of $v_G = D/10000$ determined by the validation 3D FEA and the data-driven 1D design models for design cases DLC1 and DLC2 (rounded to 1kN).

Case	H: kN			
	3D FEA	SGPR	ANNR	XGB
DLC1	408	399	393	397
DLC2	422	413	412	411

7 CONCLUSIONS

A data-driven 1D design model was developed to predict the response of monopiles subjected to monotonic lateral loading in layered clay soil profiles. The results show that the ANNR, SGPR, and XGB ML models are highly effective in training on soil reaction curve data derived directly from layered soil analyses, outperforming the SVR model. In contrast, the GPR model used by Kamas et al. (2023) is inefficient for large datasets due to its high computational demands.

Calibrating the data-driven 1D design model with data from layered soil analyses enhances its accuracy by explicitly accounting for interactions between adjacent soil layers. These models can rapidly predict the pile's lateral response under monotonic loading, significantly streamlining the design process.

AUTHOR CONTRIBUTION STATEMENT

Ioannis Kamas: Conceptualization, Methodology, Formal analysis, Software, Writing – original draft. **Harvey Burd**: Conceptualization, Supervision, Reviewing & Editing. **Byron Byrne**: Supervision, Funding acquisition, Reviewing & Editing.

ACKNOWLEDGEMENTS

The authors acknowledge the support of Ørsted via the PICASO project. Byrne is supported by the Royal Academy of Engineering under the Research Chairs and Senior Research Fellowships scheme.

REFERENCES

Brinkgreve, R.B.J., Kumarswamy, S. and Swolfs, W.M. (2018). *Plaxis 2018*. Plaxis, Delft, the Netherlands.

Burd, H.J., Taborda, D.M.G., Zdravkovic, L., Abadie, C.N., Byrne, B.W., Houlsby, G.T., Gavin, K.G., Igoe, D.J.P., Jardine, R.J., Martin, C.M., McAdam, R.A., Pedro, A.M.G. and Potts, D. (2020a). PISA design model for monopiles for offshore wind turbines: application to a marine sand. *Géotechnique* 70(11), pp. 1048-1066.

Burd, H.J., Abadie, C.N., Byrne, B.W., Houlsby, G.T., Martin, C.M., McAdam, R.A., Jardine, R.J., Pedro, A.M., Potts, D.M., Taborda, D.M. and Zdravković, L., (2020b). Application of the PISA design model to monopiles embedded in layered soils. *Géotechnique*, 70(11), pp.1067-1082.

Byrne, B.W., McAdam, R.A., Burd, H.J., Beuckelaers, W.J.A., Gavin, K.G., Houlsby, G.T., Igoe, D.J.P., Jardine, R.J., Martin, C.M., Wood, A.M., Potts, D.M., Gretlund, J.S., Taborda, D.M. G., and Zdravković L. (2020). 'Monotonic laterally loaded pile testing in a stiff glacial clay till at Cowden'. *Géotechnique* 70(11), pp. 970-985.

Chen, T., and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd International Conference on Knowledge Discovery and Data Mining*, San Francisco, pp. 785-794.

Chollet, F. (2021). *Deep learning with Python*. Simon and Schuster, New York, United States.

Friedman, J. H. (2001). Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics*, 29(5), pp. 1189–1232.

Grimstad, G., Andersen, L., Jostad, H.P. (2012). NGI-ADP: Anisotropic shear strength model for clay. *International Journal for Numerical and Analytical Methods in Geomechanics,* 36(4), pp.483-497.

Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. (2016). *Deep learning* (Vol. 1). MIT press, Cambridge, United Kingdom.

Hartigan, J.A. and Wong, M.A. (1979). Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society*, 28(1), pp.100-108.

Kamas, I., Burd, H.J., Byrne, B.W., Suryasentana, S.K. (2023). Generation of a general Cowden clay model (GCCM) using a data-driven method. *9th International SUT OSIG Conference 2023*, London, pp. 1637-1644.

Kamas (2024). *Generalised methods for the development of 1D models of monopiles for monotonic lateral loading*. DPhil Thesis, University of Oxford, UK.

Matthews, A.G.D.G., Van Der Wilk, M., Nickson, T., Fujii, K., Boukouvalas, A., Le, P., Ghahramani, Z. and Hensman, J. (2017). GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40), pp.1-6.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, pp. 2825-2830.

Smola, A.J. and Schölkopf, B. (2004). A tutorial on support vector regression. *Statistics and Computing*, 14(3), pp. 199-222.

Snelson, E., and Ghahramani, Z. (2006). Sparse Gaussian Processes using Pseudo-inputs. *In Advances in Neural Information Processing Systems*, 18, pp. 1259-1266.

Titsias, M. (2009). Variational Learning of Inducing Variables in Sparse Gaussian Processes. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, Florida, pp. 567-574.

Yang, Z. and Jeremic, B. (2005). Study of soil layering effects on lateral loading behaviour of piles. *Journal of Geotechnical and Geoenvironmental Engineering* (ASCE) 131(6), pp. 762-770.