Proceedings of the 7th International Conference on Geotechnical and Geophysical Site Characterization
Barcelona, 18 - 21 June 2024

A Case Study on Quantile and Percentile Regression to Select Design Lines for Complex, Real Work Profiles

Juan Valderrama[1#], Michael O'Neill[1], Fraser Bransby[1] and Phil Watson[1]

[1]The University of Western Australia, Department, M053, Perth WA, Australia
[#]Corresponding author: juan.valderramagiraldo@research.uwa.edu.au

ABSTRACT

Site characterisation necessarily relies on engineering judgement, usually combined with some level of statistics to define characteristic values for design purposes. A suitable method for this task is quantile regression, which allows for the definition of lower, upper, and best-estimate characteristic values. The application of quantile regression to homogeneous profiles is relatively straightforward. Although such sites are common in some areas, there is need for a more comprehensive approach to quantile regression that covers the more general scenario of heterogeneous stratified profiles. This paper takes piezocone penetrometer data from a relatively complex seabed site and demonstrates the streamlined application of quantile regression, highlighting and analysing some of the assumptions and choices behind the approach. The work shows the nuances of the method and suggests workarounds for potential scenarios where its application may be challenging.

Keywords: Quantile regression; design lines; percentile regression.

1. Introduction

The geotechnical design of foundations for offshore infrastructure often involves the definition of lower and higher estimate design lines, in order to address the sometimes conflicting requirements of assessing capacity and installability. Such profiles rely on engineering judgement supported by statistical tools and methods, such as Quantile Regression (QR). A model commonly used in econometrics (Koenker 2017), QR is a generalised version of linear regression that modifies the loss function to account for levels of exceedance, making it ideal for applications where lower and upper bounds are required. While Uzielli et al. (2019) demonstrated the use of QR to define undrained shear strength (s_u) profiles for a site offshore northern Australia, to the best understanding of the authors its application remains relatively rare in geotechnical practice and research.

A version of QR, utilising its main concepts (fitting functions associated with levels of exceedance) is referred to in this paper as Percentile Regression (PR). PR involves dividing the data into discrete intervals at arbitrarily chosen depths (i.e., 'binning' the data) and then ranking the data from each bin from lower to higher to find relevant percentiles that represent the variability of each bin. Least squares linear functions are then fitted to these percentiles.

This paper explores the use of QR and PR and the nuances of their application in both a uniform synthetic dataset and a relatively complex seabed site.

2. Design lines in geotechnical engineering

Traditionally, defining characteristic values to represent the mechanical behaviour of soils has been a subjective task, with engineers using their judgment to select design lines. However, the integration of statistical tools into geotechnical engineering practice has allowed for more objective definitions of the term 'design line'.

Phoon (2023), for example, cites Eurocode 7 (clause 2.4.5.2(2) (CEN 2004)), where the characteristic value of a geotechnical parameter is defined as a '*cautious estimate of the value affecting the occurrence of the limit state*'. This definition, as explained by Ching et al. (2020), can be approached from two perspectives; first, the physics of the problem needs to be understood to determine *the value affecting the occurrence of the limit state*, and second, providing a *cautious estimate* requires an assessment of the uncertainty and variability of the mobilised value (i.e., the statistics of the problem).

Both parts of the definition should be considered when assessing the consequences of the occurrence of the limit state. This consequence assessment, in turn, should indicate how cautious the estimate should be. A question then arises as to how to objectively measure the level of caution of an estimate.

Eurocode 7 (Clause 2.4.5.2(11)) provides a useful (yet debatable) approach to this question by suggesting: "If statistical methods are used, the characteristic value should be derived such that the calculated probability of a worse value governing the occurrence of the limit state under consideration is not greater than 5%".

In addition to the practical implications of Eurocode's definition (and aside from discussions around the appropriateness of a single quantile (5^{th}) for different problems), linking design lines to levels of exceedance is particularly useful and opens the door to models such as QR and PR, where this is explicitly addressed.

This was illustrated by Uzielli et al. (2019), who applied QR for the analysis of undrained shear strength (s_u) data derived from offshore Piezocone Penetration Tests (PCPT, also called CPTU). They provided a lower limit for bearing capacity (i.e., a low quantile) and an upper limit for installation (i.e., a high quantile).

Two things distinguish the current study from what was presented by Uzielli et al. (2019). Firstly, the analyses presented in this paper were performed on net cone pressure (q_{net}) profiles (O'Neill et al. 2022b) – this choice was made to avoid introducing additional uncertainty beyond the paper's scope, by avoiding the need for correlations between q_{net} and s_u (i.e., N_{kt}). Secondly, the PCPT dataset used in this work was acquired on a site with greater variability, which introduces additional considerations regarding layering.

The q_{net} profiles used in this study were derived from offshore PCPT data through the following formulation: $q_{net} = q_c + u_2(1-\alpha) - \sigma_v$, where q_c and u_2 are the measured cone resistance and pore pressure, respectively, α is an area ratio that represents the shape of the cone penetrometer and σ_v is the total in situ vertical stress.

2.1. Linear models

Linear models, such as Eq. (1) below, are commonly used to represent the trend of soil resistance (or strength) increasing with depth. It is worth noting that the soundness of such models depends on the variability and genesis of the soils being represented, but they are appropriate for most cases in a single soil layer.

As mentioned above, the analyses presented in this paper were conducted on q_{net} profiles, meaning that the function for representing soil resistance will be as follows:

$$q_{net} = a + b \cdot z \tag{1}$$

where a and b are fitting coefficients, and z is the depth below the mudline (seabed).

The following subsections will provide a brief introduction to QR and PR and demonstrate their application to synthetic data. The objective is to illustrate how to generate linear functions, like the one presented in Eq. (1), that are associated with arbitrary levels of exceedance (i.e., arbitrary quantiles).

2.2. Quantile regression

QR, as formalised by Koenker and Hallock (2001), involves minimising a generalised objective function $\rho_\tau(u)$, defined as follows:

$$\rho_\tau(u) = u\big(\tau - I(u<0)\big) \tag{2}$$

Here, τ is the quantile of interest, I is an indicator function, and u is the residual error, which is calculated using the formula $u = y_i - X_i^T\beta$, where y_i is the i^{th} observation of the response variable, X_i the vector of predictors for the i^{th} observation, and β is the vector of coefficients to be estimated.

To determine, for example, a regression of the 10th quantile of the data τ is set to 0.1. Then, the equations representing a model such as the one presented in Eq. (1) are as follows:

$$X_i^T = [1 \quad z_i] \tag{3}$$

$$\beta = \begin{bmatrix} a_{QR10} \\ b_{QR10} \end{bmatrix} \tag{4}$$

$$u_i = q_{net_i} - [1 \quad z_i]\begin{bmatrix} a_{QR10} \\ b_{QR10} \end{bmatrix} \tag{5}$$

$$\rho_{\tau_i} = u_i\big(0.1 - I(u_i<0)\big) \tag{6}$$

$$\min \textstyle\sum_{i=1}^{n}(u_i)\,\big(0.1 - I(u_i<0)\big) \tag{7}$$

Finding the values of a_{QR10} and b_{QR10} that minimise the sum shown in Eq. (7) gives the regression for the 10th quantile of the data. The notation for such a fit is shown below.

$$(q_{net})_{QR10} = a_{QR10} + b_{QR10} \cdot z \tag{8}$$

2.3. Percentile regression

Some geotechnical practitioners use a variant on the concept of QR, which in this paper is referred to as Percentile Regression (PR), and which also aims to find regressions conditioned on levels of exceedance.

PR is based on bin statistics, where the data are divided depth-wise, into bins of constant size, as shown in Fig.1, where the bins are 1 m deep. The data within each bin are sorted to determine the required percentiles. Once the percentiles have been determined, their values are assigned to a depth in their respective bin's average depth. Subsequently, traditional linear regression analyses are carried out to determine the PR fits.

When determining the percentiles, care should be taken to select algorithms that provide exclusive percentiles (i.e., the formulation assumes that the data are only a sample of the population and more extreme values are possible).

It is also important to note that the bin-statistical analysis required for this approach should be performed on the residuals (detrended data) rather than on the raw measurements. Failure to use detrended data results in an artificial increase in the variability of the data.

To explain the concepts of trend and residuals, we can divide the data into two components (Phoon and Kulhawy 1999): a mean process (or trend) and a deviation process. The residuals are the values that the deviation process takes. In this framework, the spatial variation of stationary soil properties can be modelled as follows:

$$q_{net}(z) = trend(z) + deviation(z) \tag{9}$$

where z is the depth. The trend function of geotechnical properties is typically modelled as a deterministic linear function such as the ones shown in Eq. (1); and the random deviation process is typically modelled as a Gaussian process with zero mean and an autocorrelation structure (Cai et al. 2019).

In the context of QR and PR, and for Gaussian stationary data, the trend component would be given by the 50th quantile (i.e., the median), with the deviation process fluctuating around it.

Note that a rigorous implementation of the PR method requires the modelled soil property to follow a Gaussian distribution (which is not always found in natural deposits) and prior knowledge of the trend, which creates a circular problem. A workaround for this issue (when sufficient data is available) is to use relatively thin bins (e.g., 20 cm) to minimise the effect of the trend in the variation of the property.

3. Synthetic data

A simple synthetic stationary dataset was generated for the demonstration of QR and PR. The rationale behind doing this, instead of using real data, lies in the fact that the trend and variability are known in the synthetic dataset, which allows for the discussion to focus on the performance of the regression methods rather than on the data properties.

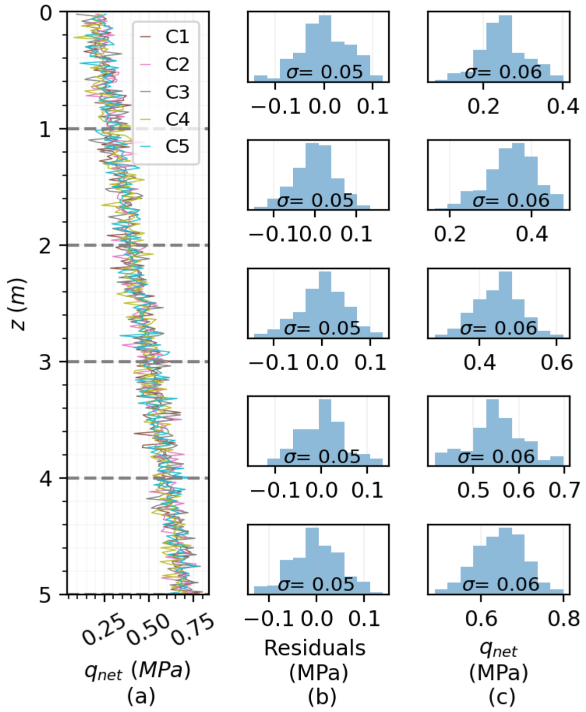

Figure 1. (a) Synthetic q_{net} profiles. (b) bin-wise distribution of residuals. (c) bin-wise distribution of the raw data.

Fig. 1 shows five synthetic stationary q_{net} profiles. These profiles were generated by adding a trend function, μ, and a random deviation process, with no autocorrelation function, $N(0,\sigma^2)$. The values for the constants are provided below.

$$q_{net} = (0.2\ + z \cdot 0.1) + N(0,(0.05)^2)\ [MPa] \qquad (10)$$

It is acknowledged that Eq. (10) could also be expressed in logarithmic values, which would allow for a more practical modelling approach by avoiding negative values. The need for this is avoided in the (synthetic) case by using an offset of 0.2 MPa at mudline. The values shown in Eq. (10) were chosen subjectively to avoid negative values and to provide what would be a homogeneous site in geotechnical practice.

The term 'homogenous' is used here to describe sediments with a single geological origin and relatively low variability. These sediments typically produce q_{net} profiles similar to those shown in Fig 1 (a).

Table 1 presents the constants a and b for six regression analyses, three PR and three QR, performed to obtain regressions for the 10th, 50th and 90th percentiles.

The analyses were conducted according to the formulations presented in the previous subsections. It is important to note that the PR was performed using residuals that were obtained by subtracting an assumed trend (as would be done in a real-world application), which was calculated as the average of the five profiles.

Regarding the regression results (shown in Fig. 2), both PR and QR perform well in capturing the properties of the trend process (first term in Eq. (10)). These results are shown in Table 1 in the rows corresponding to $(q_{net})_{PR50}$ and $(q_{net})_{QR50}$. Additionally, it is noteworthy that the PR fit is marginally closer to the known/real values.

Table 1. QR and PR fits for the synthetic dataset.

Case	a	b
$(q_{net})_{PR10}$	0.1314	0.1007
$(q_{net})_{QR10}$	0.1300	0.1016
$(q_{net})_{PR50}$	0.1995	0.1005
$(q_{net})_{QR50}$	0.1991	0.1007
$(q_{net})_{PR90}$	0.2700	0.0981
$(q_{net})_{QR90}$	0.2670	0.0986

When comparing these regression approaches, given their almost identical results, it is important to consider the practicality of their implementation.

On one hand, QR is a straightforward approach that does not need assumptions to be made or conditions to be met; it is necessary, nevertheless, to program it (in languages such as Python or R or using spreadsheets).

On the other hand, PR requires a set of steps (illustrated in Fig. 2 (a)) that are relatively easy to follow using built-in spreadsheet tools. Firstly, a bandwidth for the binning of the data must be set, this decision should be made case by case with consideration of the characteristics of the studied dataset. Subsequently, it is necessary to assume a trend process, which needs to be checked for consistency. After this, the percentiles are determined and assigned to a representative depth. Only after all these steps are the actual regression analyses performed.

Of all the steps mentioned above, the most critical is the assumption of a trend process. It must be ensured that the assumed trend effectively represents the whole dataset. This is typically done by checking that the resulting residuals follow a normal distribution with approximately zero mean. The dependence of PR on trend selection makes it a potentially subjective approach, especially in sites with great variability.

This potential subjectivity can be avoided by practitioners by determining the percentiles from the raw data instead of selecting a trend. This, however, needs to be done carefully. Fig. 3 (a), below, shows (in dotted lines) the plots of PR fits resulting from raw data; for context, QR fits are shown as well.

It is evident that the 10th and 90th PR fits tend to more extreme values. This is due to the influence of the trend in the distribution – its value is small for the first values of each bin and large for the last, artificially increasing the variability of the data. This is illustrated in Fig. 1 (b) and (c), which display the distribution and standard deviation of each bin. Visually, the histograms for the raw data (Fig. 2 (c)) appear to have heavier tails, which is confirmed by the larger standard deviation values (0.06).

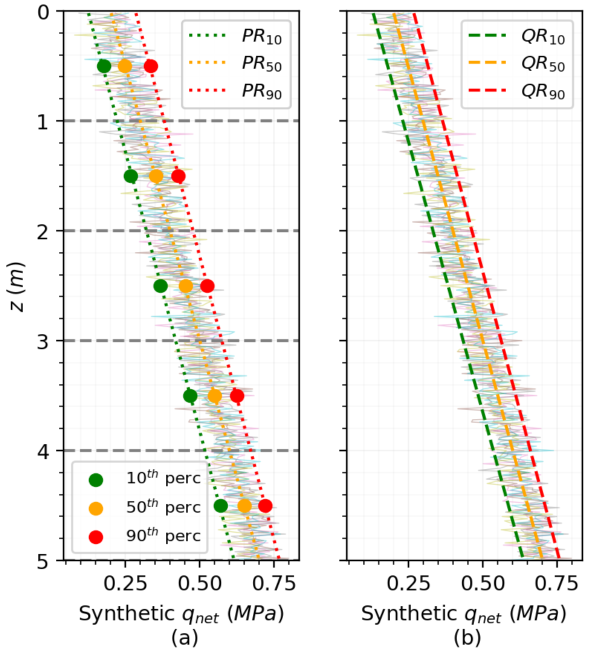

Figure 2. (a) PR fits. (b) QR fits.

The variability is said to be artificially increased as the real value (the one with which the data was generated) is 0.05. To address this issue, a practical solution is to decrease the bin width. This reduces the influence of the trend on the data distribution at each bin. Fig. 3 (b) illustrates how the PR fits move towards the QR when the bin width is reduced to 0.2 m.

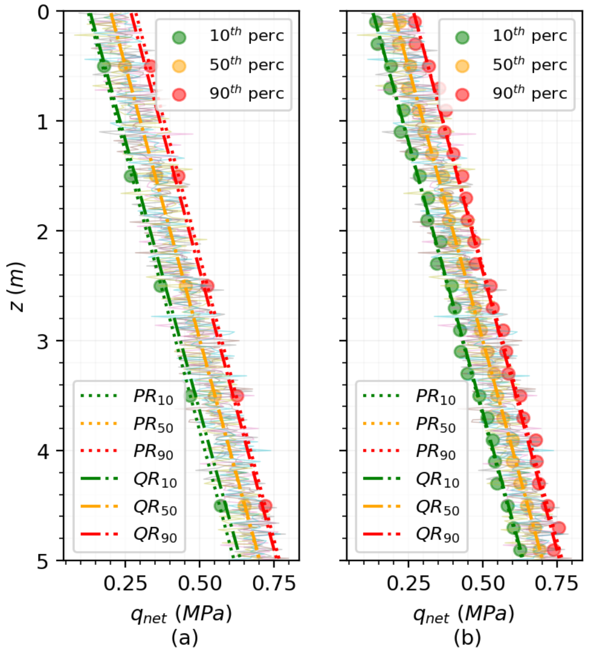

Figure 3. Comparison of QR and PR fits done with raw data. (a) Analyses with 1.0 m wide bins. (b) Analyses with 0.2 m wide bins.

The above discussion introduced the key concepts behind QR and PR, and demonstrated their implementation in statistically homogenous (i.e., stationary) soil profiles. In the case adopted, both QR and PR produce consistent regressions associated with levels of exceedance, demonstrating the methods' potential to be used as tools for design line selection. A comparison of their performance showed that QR stands out for its relative simplicity and the authors believe it should be preferred for practical applications.

4. Regression analyses on layered profiles

Although uniform sites, like the one simulated in the previous section, do occur in some areas, they are likely to be an exception rather than the rule – especially for the current focus on offshore wind, with such developments tending to be in shallower waters and more variable ground conditions. Therefore, methods for site characterisation, particularly those that aim to define design lines, must be able to accommodate and represent layering.

This section presents the implementation of QR and PR on a site with relatively high variability.

4.1. Selected real-world dataset

Soil deposits are typically made of layers of materials with varying origins that have undergone different geological processes, which creates marked differences in mechanical behaviour. The dataset used in this part of the work is an example of that. It consists of 13 PCPTs shown in Fig. 5, covering an area (sketched in Fig. 4) of approximately 105 m x 105 m.

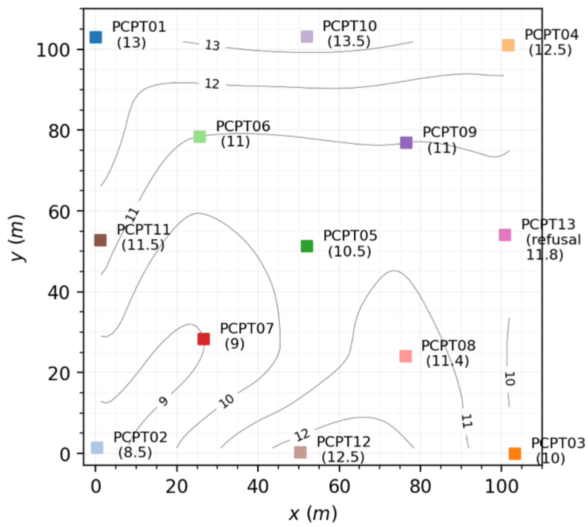

Figure 4. Plan view of the location of the PCPTs.

All tests but PCPT13 reached a depth of 20 m, for which cone refusal was reached at 12.1 m. The data were obtained at a deep-water site offshore north west Australia. The seabed in this region is composed mainly of carbonate sediments (Watson et al. 2019).

The q_{net} profiles were derived from q_c and u_2 measurements taken at 0.02 m depth intervals using 15 cm^2 projected area cones with 60° tips, with an area ratio $\alpha = 0.59$. Subsequently, the net resistance was calculated using the effective unit weights shown in Table 2 to determine the total in situ vertical stress. As the data were originally acquired in a slightly irregular manner (i.e., the readings were taken at depths that were not multiples of 0.02 m), for the purposes of this study the q_{net} datapoints were linearly interpolated to 0.02 m multiples – in order to simplify processing for PR.

Table 2. Effective unit weights for data processing.

Depth (m)	$\gamma'(kN/m^3)$
0 - 10	5.2
10-16	6.4
16-20	6.3

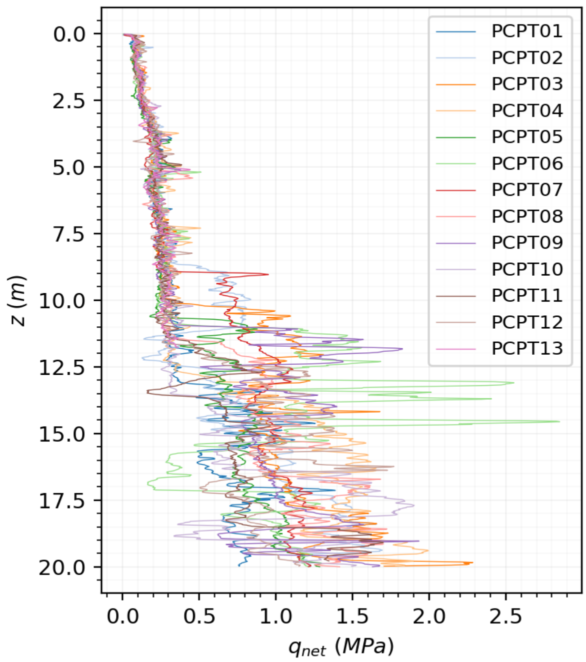

Figure 5. Profiles of net cone resistance (q_{net}).

In terms of mechanical behaviour and statistical properties, the stratigraphy of this site can be divided into two distinct layers:

1. The first is from the mudline to a depth that varies between 8.5 m and 13 m, is a relatively soft layer that is described (in geotechnical reports for the site) as comprising interbedded turbidites (very soft carbonate sandy mud) and pelagic sediments (very soft carbonate mud). Fig. 6 details the first 5 m of the profile and it is interesting to note how, besides some occasional peaks, this pelagic layer is strikingly similar to the synthetic dataset used in the previous section. This similarity suggests stationarity in the statistical properties of this layer – and such uniformity would allow direct implementation of QR or PR in this depth range.
2. The second layer covers the remainder of the profile to 20 m, and has highly variable conditions and (overall) higher resistance. This layer has been described as comprising debris flow material consisting of carbonate clasts, sands and clays. This mixture of materials produces a nonstationary layer – and the non-stationarity of the statistical properties manifests in the non-normally distributed histograms and wide ranges of values shown in Fig. 7.

The difference in magnitude and distribution between the first and second layers implies that it is virtually impossible to find single trend and deviation processes to accurately represent the whole site.

Aside from the cone resistance values in each layer, an additional feature worth analysing is the depth at which the transition from the first to the second soil unit (shown in brackets below the name tags in Fig. 4 and in profile form in Fig. 7). It is interesting to note that this depth is spatially dependent – and such dependence should be considered when assessing the appropriateness models such as QR and PR, as regression models produce functions that are completely independent of the plan location (x-y-wise) of the data. This is not to say that the models should not be used, but rather than in such sites it is important to recognise that their outputs are more a tool for characterising a site as a whole site, rather than a tool for predicting profiles at specific locations.

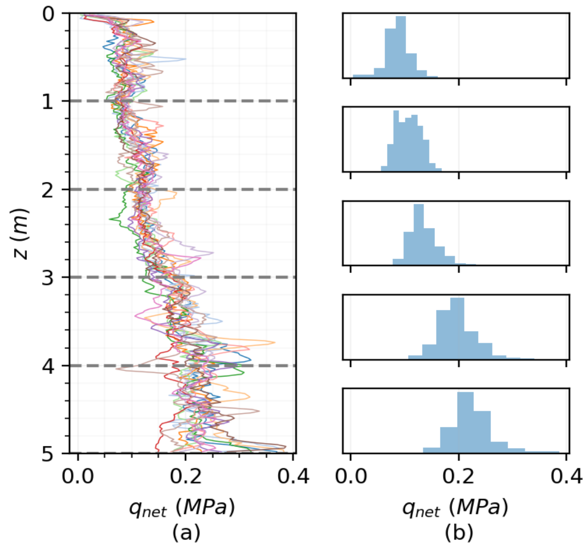

Figure 6. q_{net} data for the first five metres. (a) profiles. (b) data distribution.

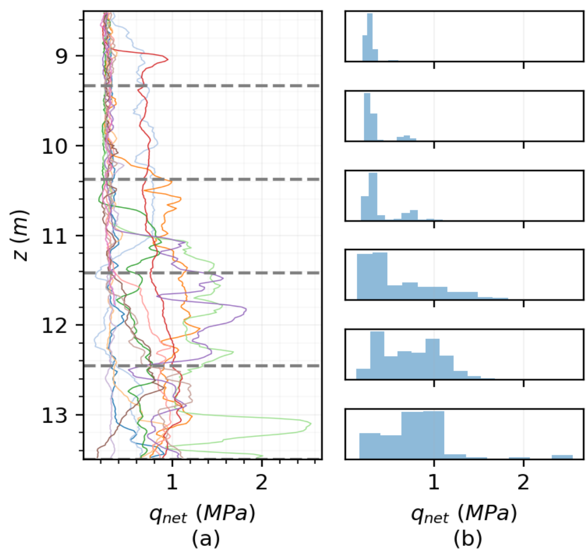

Figure 7. q_{net} data from 8.5 to 13.5 m depth. (a) profiles. (b) data distribution.

4.2. PR and QR analyses

This subsection details the authors' approach to implementing QR and PR for design line selection based on the 20 m profiles shown in Fig. 5.

It is crucial to note the subjective nature of this subsection – and it is accepted that some decisions are subject to debate. The exercise is intended to explore the key questions and challenges practitioners might face when implementing these models to real-world problems.

The initial question is: how to determine the appropriate quantiles? Addressing this starts with an exploration of the applicable codes and standards. However, interpreting and interrogating the codes requires a nuanced process that considers the intended uses of the analysis outputs and the potential consequences of design failure – for instance, distinctions between the repercussions of an offshore wind turbine foundation failure and those of an onshore bridge pier underscore the need for tailored quantile selection.

Beyond the normative considerations, quantile selection must also recognise that the fits from these regression analyses essentially provide a form of ground profile – and therefore, their value should be calibrated within the framework of Burland's geotechnical triangle (1987). According to this framework, a comprehensive ground profile should integrate findings from ground exploration, description, and testing, and incorporate an understanding of the geological processes responsible for sediment formation. Such a profile must also seamlessly interact with the conceptual idealisation of the soil as part of an engineered system.

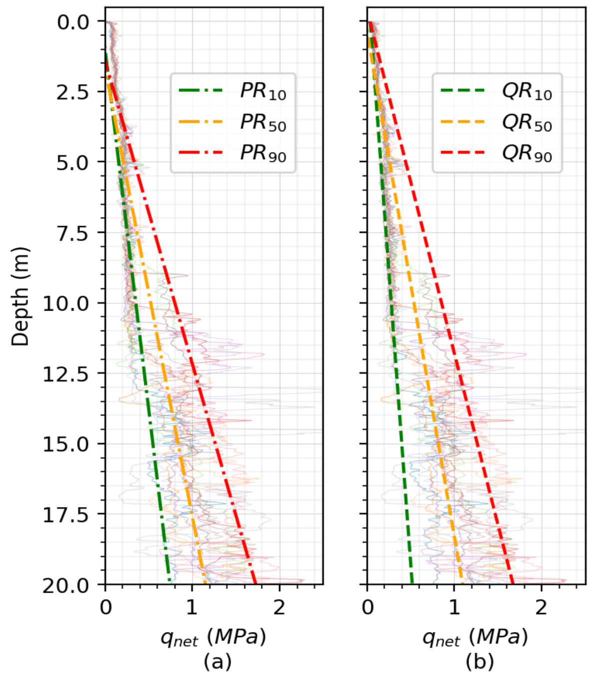

Figure 8. Fits to the q_{net} dataset without consideration to layering. (a) PR. (b) QR.

This conceptualisation necessitates a clear definition of the specific use for each output of the regression analyses. For instance, in foundation design for offshore infrastructure, two design lines are often defined: a low estimate for bearing capacity computation, and a high estimate for installation requirements definition. In this paper three quantiles – 10th, 50th and 90th – were selected, which does not rule out the existence of extremes. For example, opting for a 90th quantile introduces the possibility of overlooking localised hard layers, while using the 10th (rather than the 5th, as Eurocode recommends) can yield a rather optimistic estimate of the profiles' resistance.

Following the quantile selection, QR and PR regression analyses can be conducted. The fits illustrated in Fig. 8 show what would be obtained in analyses undertaken without considering the layering (i.e. using the full profile). They are clearly not optimal, with the following points standing out:

- The PRs show negative intercepts at the mudline, attributed to the stronger sediments of the second layer tilting the regression lines and resulting in negative values for the initial depths, as demonstrated in Fig. 9.
- Although the QRs do not exhibit negative values, there is still some tilting in the fits. Notably, the 90th percentile fit significantly exceeds the values of the data for the first layer of the profile.

The results of this first exercise pave the way to the second question a practitioner would face: how to address the challenge of layering in practice?

4.2.1. Layering

As outlined above, the profile of this site can be simplified as a two-layer system. A first attempt to address this would be to designate a single depth for the transition from the first to the second layer. Employing 11 m as the transitional depth, the QR and PR fits shown in Fig. 9 were derived.

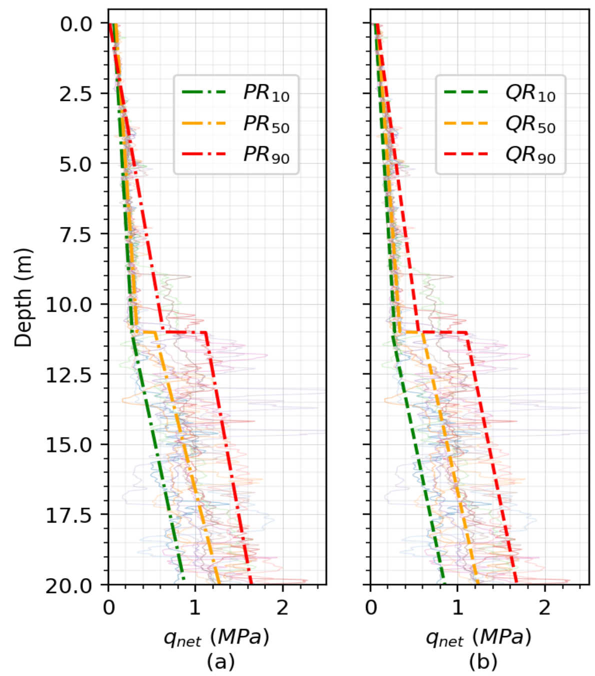

Figure 9. Fits to the q_{net} dataset using a two-layer system with a single transitional depth of 10 m. (a) PR. (b) QR.

While this improves the characterisation of the first layer (by preventing negative values in the PR fits), it overlooks the significant spatial dependence of this transitional depth on the site – thereby introducing risk in design. Notably, within the 90th quantile fits, a layer of relatively high resistances between 9 and 11 m is overlooked, and the 'tilting' effect described above causes the PR_{90} fit to cross the other two regression lines, as Fig. 9 (a) shows.

To address this, a more nuanced strategy was adopted involving a 0.2 m bin-statistics analysis. Here, the strategy involved plotting the bin-wise q_{net} values for the selected quantiles, as illustrated in Fig. 10. Examination of this figure led to the definition of a system with separate transitional depths (O'Neill et al. 2022a):

- For the 10th and 50th quantile fits, a single transitional depth, at 13.6 m and 12.0 m, was respectively adopted.
- Two transitional depths for the 90th quantile fits, at 9.0 m and 11 m.

This refined approach addressed not only the spatial nuances in the transitional depth, but also provided a more accurate representation of the site's profile – overcoming the limitations observed in the initial single-depth designation. This improvement in the model's ability to represent the site comes, nevertheless, at the cost of some statistical rigour.

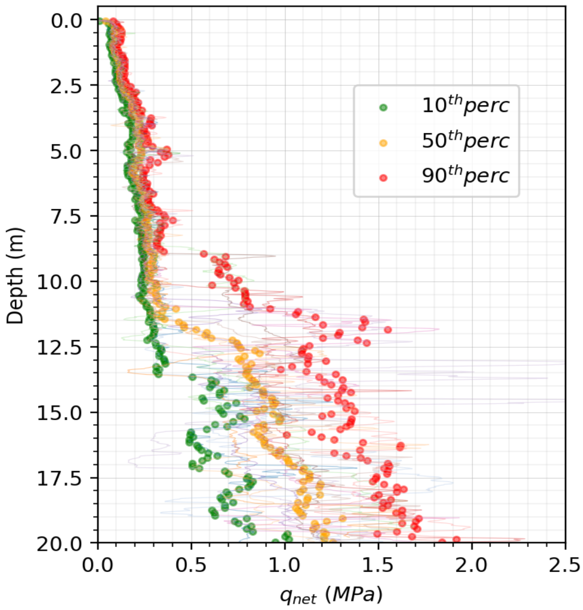

Figure 10. 10th, 50th and 90th of the 0.2 m binned q_{net} dataset.

4.2.2. Final regression analyses

Fig. 11 presents the outcomes of the analyses, featuring the application of QR and PR models to the layered data. The comparison in Fig. 12 displays the results generated by both models, utilizing 0.2 m-bin quantiles as a frame of reference. From this figure it is worth noting how, under certain conditions, PR and QR can produce virtually identical results.

The approach to regression analyses for design line definition shown in this paper is part a practical guide for the implementation of such models in real-world problems and is intended to prompt further discussion, as it shows how these models struggle to represent such variable sites and require significant input from the engineer – which increases the degree of subjectivity in the analyses.

5. Conclusions

This paper was motivated by the need to highlight the utility of models that can generate design lines associated with levels of exceedance. Two such models (PR and QR) were employed to analyse both synthetic and real datasets. The results demonstrate that under specific conditions both approaches yield practically identical outcomes. However, it was observed that these conditions are site-specific and susceptible to subjective handling of the data. Notably, PR appears to produce less reliable results compared to QR.

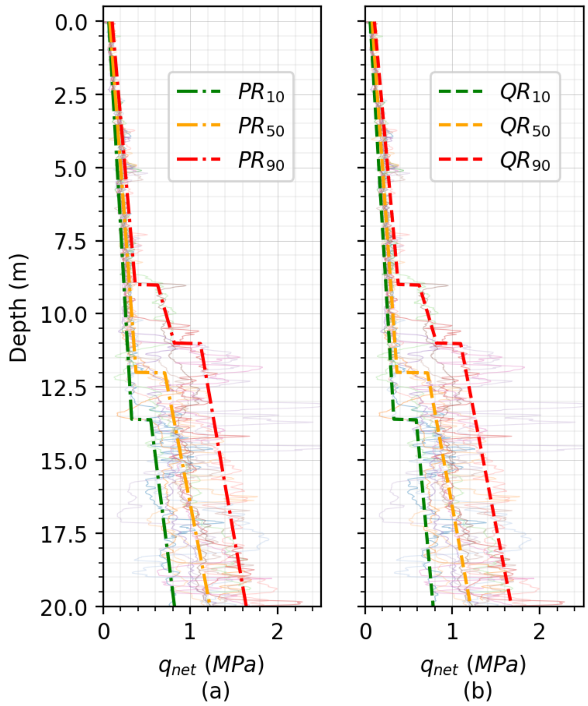

Figure 11. Fits to the q_{net} dataset. (a) PR. (b) QR.

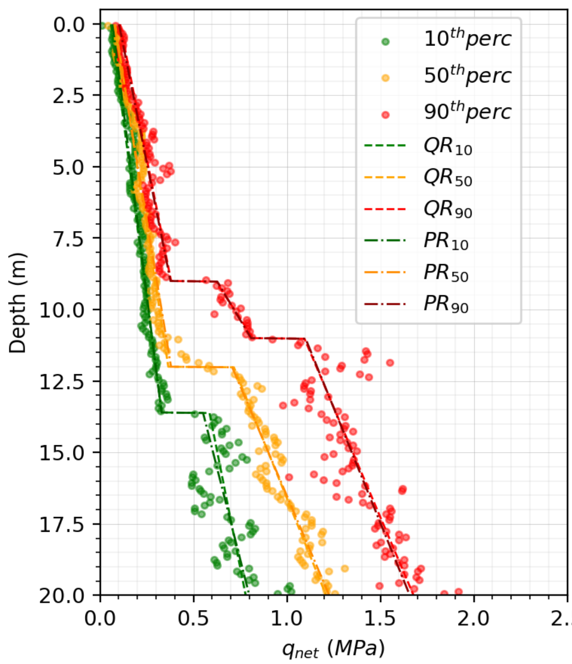

Figure 12. Comparison of the PR and QR fits to the q_{net} dataset.

The observed disparities in performance can be attributed (in part) to the fact that PR is a simplified approach to generating such design lines, whereas QR stems from a more rigorous mathematical foundation. Consequently, for rigorous analysis, the authors believe QR should be prioritized over PR.

One of the main advantages of QR is its potential to reduce the subjectivity in design line definition. However, the approach still requires a fair degree of expertise from the engineer, particularly in the selection of quantiles for the regression analyses and the definition of an appropriate layering system.

Despite the usefulness and flexibility of QR, it has a crucial limitation – namely its independence from the location of the input data. This independence translates into an inability to address spatial variability, meaning that it produces a single output regression for the whole dataset and cannot be used to predict location-specific profiles. An alternative use of the model to overcome this limitation would involve running location-specific regressions using relevant subsets of data – for instance, to get an interpolated q_{net} profile for a point located at x = 30 m and y = 50 m (referring to Fig. 4), QR fits can be obtained using PCPTs 5, 6, 7, and 11. However, these accommodations require judgement about which parts of the data set are relevant to a specific location (and infer statistical stationarity in that region) – this implies that part of the data goes unused, which is not an ideal scenario as it could lead to the overlooking of important features of the site. Approaches to consider spatial variability explicitly have been developed (Uzielli 2022) and a companion paper at this conference (Valderrama et al. 2024) shows the implementation of such a model – GeoWarp (Bertolacci et al. 2024), for the analysis of the same real-world dataset used for this work.

In summary, regression models are valuable tools in providing an overall understanding of the features and variability of a site, but they should be used carefully in geologically complex sites.

Acknowledgements

This research is supported by the ARC ITRH for Transforming Energy Infrastructure through Digital Engineering (TIDE, http://TIDE.edu.au), which is led by The University of Western Australia (UWA), delivered with The University of Wollongong and a number of Australian and international research partners, and funded by the Australian Research Council, INPEX Operations Australia, Shell Australia, Woodside Energy, Fugro Australia Marine, Wood Group Kenny Australia, RPS Group, Bureau Veritas, and Lloyd's Register Global Technology (grant No. IH200100009). Fraser Bransby holds the Fugro Chair in Geotechnics at UWA, whose support is gratefully acknowledged. Phil Watson leads the Shell Chair in Offshore Engineering research team at UWA, which is supported by Shell Australia.

References

Bertolacci, M., A. Zammit-Mangion, J. Valderrama Giraldo, M. O'Neill, F. Bransby, and P. Watson. 2024. "GeoWarp: Warped spatial processes for inferring subsea sediment properties." In preparation.

Burland, J. B. 1987. "Kevin Nash Lecture. The teaching of soil mechanics – a personal view." *Proceedings of the 9th European Conference on Soil Mechanics and Foundation Engineering, Dublin.* Vol. 3, 1427–1447.

Cai, Y., J. Li, X. Li, D. Li, and L. Zhang. 2019. "Estimating soil resistance at unsampled locations based on limited CPT data." *Bulletin of Engineering Geology and the Environment* 1435-9537. https://doi.org/10.1007/s10064-018-1318-2.

CEN. 2004. *Eurocode 7 Geotechnical Design—Part 1: General Rules. EN 1997-1.* European Committee for Standardization.

Ching, J., K. K. Phoon, K. F. Chen, T. L.L. Orr, and H. R. Schneider. 2020. "Statistical determination of multivariate characteristic values for Eurocode 7." *Structural Safety* 101893. 10.1016/j.strusafe.2019.101893.

Fenton, G. A. 1999. "Random Field Modeling of CPT Data." *Journal of Geotechnical and Geoenvironmental Engineering* 125 (6): 486-498. 10.1061/(ASCE)1090-0241(1999)125:6(486).

Koenker, R. 2017. "Quantile regression: 40 years on." *Annual review of economics* (9): 155-176. https://doi.org/10.1146/annurev-economics-063016-103651.

Koenker, R., and K. F. Hallock. 2001. "Quantile Regression." *Journal of Economic Perspectives* 15: 143-156. www.aeaweb.org/articles?id=10.1257/jep.15.4.14.

Lyu, B., Y. Hu, and Y. Wang. 2023. "Data-driven development of three-dimensional subsurface models from sparse measurements using Bayesian compressive sampling: A benchmarking study." *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering* 9 (2): 04023010. https://doi.org/10.1061/AJRUA6.RUENG-935.

O'Neill, M., F. Bransby, and P. Watson. 2022a. "The effectiveness of Spatial Interpolation of Sparse PCPT data to optimise Offshore Design." *8th International Symposium for Geotechnical Safety & Risk (ISGSR 2022).* Newcastle, NSW, Australia.

O'Neill, M., M. Bransby, J. Doherty, and P. Watson. 2022b. "Spatial interpolation of sparse PCPT data to optimise infrastructure design." *Cone Penetration Testing 2022.* 1023-1028. 1023-1028. https://doi.org/10.1201/9781003308829-154.

Phoon, K. K. 2023. "What Geotechnical Engineers Want to Know about Reliability." *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering* 9 (2): 03123001. ascelibrary.org/doi/abs/10.1061/AJRUA6.RUENG-1002.

Phoon, K. K., and F. H. Kulhawy. 1999. "Characterization of geotechnical variability." *Canadian Geotechnical Journal* 36 (4): 612-624. https://doi.org/10.1139/t99-038.

Uzielli, M. 2022. "Non-deterministic interpretation and applications of CPT testing data." Edited by G Gottardi and L Tonni. *Cone Penetration Testing 2022.* Bologna, Italy: CRC Press. 81-93. https://www.taylorfrancis.com/chapters/oa-edit/10.1201/9781003308829-6/non-deterministic-interpretation-applications-cpt-testing-data-uzielli.

Uzielli, M., M. Zei, and M. J. Cassidy. 2019. "Probabilistic assignment of design undrained shear strength using quantile regression." Edited by Jianye Ching, Dian-Qing Li and Jie Zhang. *Proceedings of the 7th International Symposium on Geotechnical Safety and Risk (ISGSR).* Taipei: Research Publishing. 188-193. ISBN: 978-981-11-2725-0; doi:10.3850/978-981-11-2725-0 IS15-8-cd.

Valderrama, J., N. O'Neill, F. Bransby, P. Watson, M. Bertolacci, and A. Zammit-Mangion. 2024. "Can seabed spatial uncertainty be quantified using advanced statistical approaches?" *Proceedings of the 7th International Conference on Geotechnical and Geophysical Site Characterization.* Barcelona.

Watson, P. G., M. F. Bransby, Z. L. Delimi, C. T. Erbrich, I. Finnie, H. Krisdani, C. Meecham, et al. 2019. "Foundation design in offshore carbonate sediments-building on knowledge to address future challenges." Edited by N.P. et al. López-Acosta. *XVI Pan-American Conference on Soil Mechanics and Geotechnical Engineering (XVI PCSMGE).* Cancun, Mexico: ISSMGE. 240-274.