Optimising Foundation Layouts over Existing Tunnels

Arthur Chai, supervised by Professor Charles Augarde

Abstract — An automated foundation layout optimiser was developed and utilises the principles of genetic algorithms to minimise hoop stresses within underlying tunnels. The achieved accuracy of maximum hoop stress values was within ±0.1kPa of the theoretical minimum for a single-tunnel problem, with a corresponding precision of ±0.3m in foundation locations, limited by the coarse tunnel lining mesh used to reduce computation times. This research aims to demonstrate the potential of open-source tools to streamline engineering workflows, fostering accessible geotechnical design methods.

Index Terms— Geotechnical Finite Element Analysis (FEA), Geotechnical Optimisation, Foundation Layout Optimisation, Automated Design Tools

I. Introduction

A. *Motivation*

Over the past century, geotechnical theories have expanded into non-linear soil mechanics, which are prevalent in real-world conditions and deviate significantly from the idealised linear scenarios found in earlier geotechnical theories. Soils are inherently non-linear in their behaviour and have historically posed challenges to accurate analysis, such as the heterogeneity of soils in Mexico City leading to unexpectedly excessive differential settlements of its Cathedral in 1989 [1]. Solutions have been developed for such non-linear problems in the past, but their intensive computational costs rendered such calculations infeasible.

Advancements in computational power over recent decades have transformed the field of geotechnical engineering, allowing engineers to tackle these intricate problems with remarkable precision. These advancements have led to the development of robust and comprehensive geotechnical software tools such as Plaxis [2], which leverages finite element analysis (FEA) and other sophisticated numerical techniques and have become central to modern geotechnical design. However, despite their accuracy and reliability, such software often comes with significant drawbacks: high costs and steep learning curves. For practitioners and researchers seeking quick, efficient solutions, these tools can often pose significant barriers.

More recently, the advent of open-source platforms and high-level programming languages has opened new possibilities for geotechnical engineers to develop tailored tools that address specific problems. Python, with its extensive libraries and user-friendly design, has emerged as a leading choice for developing custom computational solutions. This has paved the way for the development of simpler, cost-effective tools that, although less precise than their commercial counterparts, offer significant advantages in terms of accessibility and usability. With growing demands for cost-effective and sustainable infrastructure, the optimisation of infrastructure designs is both a critical challenge and therefore an opportunity to effectively and efficiently address these increasingly important demands.

As such, this project aims to leverage these advancements in technology to develop an optimiser for foundation layouts for geotechnical problems involving existing tunnels in the geometry. By prioritising simplicity and computational efficiency, the tool is designed to offer a practical solution for engineers and researchers who require a fast yet effective approach to foundation design, especially where available resources are more limited than those available to larger scale engineering firms. The emphasis of user accessibility and computational simplicity positions this tool as an approachable alternative for real-world applications, enabling a broader range of engineers to design optimised foundations, as well as providing a strong and flexible foundation for future modifications. By extension, it will facilitate broader accessibility to efficient and sustainable solutions, while advancing the adoption of modern design methodologies.

B. *Literature Review*

1) *Historical context of geotechnical analyses*

The comprehensive geotechnical FEA tools available today lead back to the numerous findings from early soil mechanics. These include the fundamental theories developed by pioneering civil engineers such as Karl Terzaghi and Ralph Peck [3], which served as the basis for a wide range of subsequent geotechnical engineering studies and laid the groundwork for advancements in non-linear soil mechanics. These include theories such as Terzaghi's Principle of Effective Stress [3], allowing higher levels of accuracy in geotechnical analysis to be achieved.

Alongside non-linear soil mechanics, FEA has also been developed over 70 years and has been combined with geotechnical engineering fundamentals to solve geotechnical problems, which can vary significantly compared to the structural engineering problems that FEA was initially developed for.

Prior to the incorporation of FEA in geotechnical engineering, traditional geotechnical analyses relied heavily on empirical methods and analytical solutions, which were often limited in their ability to handle or compute complex soil-structure interactions and deal with heterogeneous conditions, often resulting in the use of overconservative designs.

In the early stages of geotechnical FEA development, advancements focused on linear elastic models, but as computational power gradually increased, it became possible to perform calculations of greater complexity. Incorporation of non-linear behaviour and corresponding constitutive soil models such as Critical State models and visco-plastic models were gradually implemented into geotechnical FEA tools, allowing better representation of real-world geotechnical problems. Long-term settlements, slope stability, and tunnelling were able to be accurately captured by geotechnical FEA, which were previously infeasible due to the otherwise vast amount of hand calculations that would have been required. This has therefore led to the predominance of geotechnical FEA use in industry, due to its practicality and ability to accurately analyse complex geotechnical problems.

Geotechnics and geotechnical FEA differ in methodology, accuracy, and consequently, their applications. Traditional geotechnical methods rely on empirical techniques and simplified analytical models, making them more suited for routine and quick design tasks. However, complex problems often involve non-linear soil behaviour, anisotropy and irregular geometries, rendering hand calculations impractical, especially when compared with the numerical capabilities of modern software. In contrast, geotechnical FEA utilises advanced numerical techniques, enabling more accurate and precise solutions to complex problems. Challenging analyses involving phenomena such as plasticity, creep, and complex loading paths can be handled with relative ease, offering a significant advantage over traditional methods [4].

Despite its capabilities, geotechnical FEA presents several issues which must be addressed to achieve accurate and reliable results, which [4] and [5] also describe. A key issue is the need for accurate soil modelling to capture non-linearity and anisotropy. The inherent heterogeneity of soil can lead to significant variability in properties such as stiffness, permeability, and strength. Neglecting these variabilities can lead to uncertainties in the analysis or provide inaccurate approximations of soil responses if they are not accounted for when setting up the FEA. Additionally, defining realistic boundary conditions are another issue that, if improperly addressed, result in inaccurate stress and deformation predictions.

Furthermore, numerical methods are often employed in most geotechnical FEA, which can introduce problems inherent to numerical methods. One such challenge is convergence difficulty, which is frequently encountered in large-scale problems where complex geometries or highly non-linear material properties can lead to increased numerical instability. These challenges highlight the need for efficient iterative algorithms to enhance the practicality and reliability of geotechnical FEA in engineering design.

Results are also affected by the quality and refinement of meshes used in the FEA, and the set of assumptions made in the FEA may be too tailored towards the specific model analysed and reduce the generalisability of the FEA tool itself. For example, using small-scale lab test data to calibrate parameters in the model will become inaccurate when the model is then used to assess much larger geometries.

Ultimately, geotechnical FEA, while powerful and effective, remains an approximation relying on numerical discretisation. The closed-form analytical solutions of traditional geotechnical methods provide precise solutions under simplified assumptions, despite their more limited applications. For geotechnical FEA, the reliance on the assumptions used for the chosen constitutive model and mesh quality means that while useful insights can be obtained for a wider range of problems, it cannot fully replicate the exact behaviour of soils in the real world.

2) Numerical Optimisation in Geotechnical Engineering

With geotechnical FEA capable of solving a wide range of problems, the integration of numerical optimisation schemes has significantly enhanced its ability to efficiently explore complex design spaces and identify optimal solutions.

One of the most common optimisation techniques is topology optimisation, which optimises the material layout within a defined space to meet predefined stress conditions, constraints and loading scenarios. In geotechnical engineering, this is frequently applied to foundation geometries and similar components used in geotechnical structures such as tunnel linings and are often used to reduce material costs in designs.

Parametric optimisation is another widely used optimisation technique in geotechnical engineering, in which design parameters are varied to explore potential configurations. It is useful for multi-variable problems common to geotechnical engineering, such as the design of a shallow foundation to maximise factors of safety against multiple failure mechanisms.

As different optimisation schemes are better suited to specific problem types, it is essential to define a clear objective function before selecting the appropriate optimisation method.

A key distinction between different optimisation schemes lies in the types of algorithms employed, which can determine the efficiency and accuracy of the optimisation scheme when applied to complex geotechnical problems, as documented in [6]. Furthermore, these algorithms can be combined to optimise different stages of the process, allowing specific optimisation techniques to be applied to stages where they are most effective. For example, reliability-based algorithms be applied to address stages containing uncertainties in design variables such as material properties and loading conditions, addressing sensitivities of parameters such as slope stability. Methods such as the Monte Carlo method [7] or the Normal Distribution are used, where probabilistic models are utilised to achieve a desired level of reliability.

For particularly complex problems, such optimisation algorithms may become excessively expensive in terms of computational cost, such that re-running the optimisation can lead to significant delays in project timelines.

To overcome these computational costs, heuristic algorithms can be utilised in the optimisation scheme to generate effective solutions to complex, non-linear problems with a trade-off between accuracy and efficiency. Global solutions are reached more quickly, which can prove useful for projects where time or budgets are limited, thus making heuristic algorithms advantageous for complex geotechnical problems.

Genetic algorithms (GAs) are a widely used heuristic algorithm based on the principles of natural selection and genetics and have been successfully implemented in geotechnical FEA, commonly for slope and foundation designs [8], [9]. A population of potential solutions are iteratively refined through selection and crossovers, with the introduction of mutations to avoid convergence at local optima, iteratively refining solutions towards an optimum through a mix of exploration and exploitation (improvement of previously explored solutions). Randomness is utilised for exploration of the solution space, for which the process is commonly guided by adaptive mechanisms to balance exploration and exploitation. Stochastic trial-and-error is combined with feedback with each cycle of exploration, to improve the accuracy and convergence of generated solutions. It is also known as an evolutionary algorithm (as a subset of heuristic algorithms), and is effective in handling non-linear and multi-modal problems, such as complex geotechnical problems. Similarly used evolutionary algorithms in geotechnical engineering can include simulated annealing and ant-colony optimisation [10], to efficiently compute optimised solutions for the objective function.

Such heuristic algorithms can be easily implemented via Python scripts – for example, GAs are widely available as Python packages, making them accessible and easier to incorporate into optimisation workflows. Indeed, heuristic methods have been used to iteratively solve geotechnical problems before, as seen in [11], but are often employed to the design of the structures themselves, as opposed to their layout. Furthermore, foundation layout optimisation on soils containing existing tunnels have not yet been explored, likely due to the challenges associated with integrating geotechnical FEA software and optimisation algorithms into a singular and efficient workflow. Foundation optimisation has been covered multiple times in the past but often use highly specialised hardware/software for the optimisation process, which are very costly, and focus on refining individual foundations and/or superstructures, as opposed to the layout of such foundations, narrowing the range of objective functions that are selected in geotechnical optimisations.

Moreover, engineers often need to make incremental changes to designs as projects progress, especially for complex, multi-stage projects which may require quick performance checks to evaluate these adjustments. This renders computationally costly optimisation methods impractical, as they are often too time-consuming and inflexible for multiple iterative, real-time changes made in relatively quick succession.

As such, this underscores the importance of incorporating appropriate optimisers alongside a user-friendly software with an intuitive input interface, to overcome the limitations of traditional design processes and expand the range of objective functions considered in geotechnical optimisations.

3) Current Trends in Geotechnical FEA Optimisation Schemes

As geotechnical optimisation techniques continue to evolve, several emerging techniques are poised to enhance the design and analysis process, with advancements mainly focused on optimisation algorithms employed. One notable development is the increasing use of artificial intelligence (AI) driven optimisation and machine learning (ML) with geotechnical FEA [12] to refine the algorithms used. AI can identify patterns and modify algorithm parameters during optimisations to generate design solutions that may not be immediately obvious through traditional methods, while ML leverages previous data from iterative cycles to increase the rate of convergence towards optimal solutions and significantly accelerating the solution process. More recently, real-time analyses for geotechnical structure monitoring have been coupled with ML [13], enabling more accurate and continuous results from geotechnical analyses, alongside earlier anomaly detection and prevention.

Another key development is the integration of geotechnical FEA with cloud-based software and Building Information Modelling (BIM) [14], which allows for better collaboration and modelling between engineers and stakeholders. Continuous feedback loops between FEA analyses and BIM models can significantly improve the accuracy of foundation designs and streamline the process, whilst advances in parallel computing and cloud-based platforms are enabling faster and more scalable geotechnical FEA analyses.

One such example is Gofer [15], a cloud-based geotechnical FEA software developed by Oasys, the software arm of ARUP, and utilised as the main FEA software in this project. Its cloud-based nature enables it to efficiently handle geotechnical analyses without being limited by local hardware constraints, whilst its API allows users to integrate the setup and analysis of geotechnical FEA models into automated workflows, facilitating direct communication between Gofer and external applications without relying on Gofer's manual user interface. Most commercial geotechnical software often operate as standalone desktop applications with limited interoperability, and whilst manually updating models to reflect changes in design configurations within such software is possible, it is often time-consuming and labour-intensive. Tasks such as re-meshing and re-calibration of material models can lead to substantial computational and operational burdens, such that most optimisations are undertaken for a limited range of objective functions which do not require as much re-calibration in the software used.

This therefore serves to highlight the significance of Gofer, which can be easily combined with various optimisation schemes within a Python-based integrated development environment (IDE), greatly increasing the feasibility of iteratively testing multiple design configurations for a wider range of problems. As computational power and optimisation algorithms continue to improve, the introduction of cloud-based geotechnical FEA software will therefore enable more efficient and scalable analyses alongside integration with other tools such as optimisers, significantly improving the feasibility of optimising a larger spectrum of geotechnical problems.

As such, this project aims to demonstrate the potential and feasibility of such software, through the automated optimisation of foundation layouts over existing tunnels.

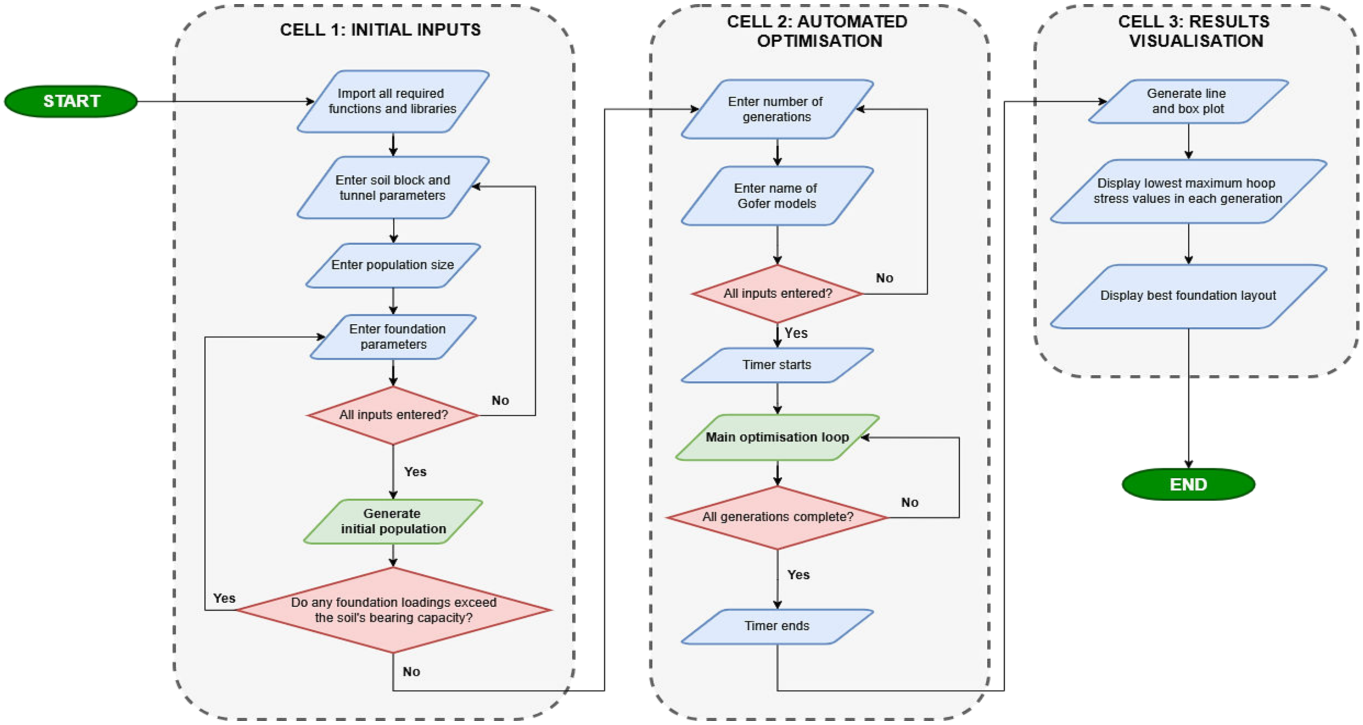

Figure 1: Simplified flowchart illustrating the structure of the developed genetic algorithm-based optimisation tool. The processes are grouped according to their corresponding Jupyter Notebook cells, reflecting the modular execution framework of the optimiser.

II. Implementation and Deployment of the Tool

A. Overview and objectives

This project develops an optimiser tool incorporating the cloud-based geotechnical FEA software Gofer to automate foundation layout designs over shallow tunnels. Its novelty lies in combining automated geotechnical FEA with GA-based optimisation for foundation layouts, as opposed to the more common optimisation of individual foundation topologies.

The primary objective is to evaluate the optimiser's performance on a benchmark case consisting of a single circular tunnel beneath two square pad foundations in uniform soil. Key GA parameters such as population size and mutation rates are individually varied to assess their effect on convergence rates and maximum hoop stress reductions. The accuracy and precision of the maximum hoop stresses of optimised foundation layouts will therefore be assessed, alongside the convergence rates achieved during their optimisations. Successful convergence in the single-tunnel setup will validate the optimiser's functionality and informs its application to more complex geometries, which are also tested later in the study.

Due to lack of real-world data to validate the computed tunnel hoop stresses, results are benchmarked against minima graphs, which are line graphs that plot maximum hoop stresses across systematically varied foundation layouts. Performance is further evaluated using box plots showing generational stress reduction and convergence patterns. Gofer's calculations are assumed to be reliable, supported by prior hand-calculations for very simplified geotechnical problems, although these checks are inherently limited as they cannot be performed for more complex geometries.

B. Main structure of the optimiser tool

The optimiser is implemented as a Jupyter Notebook script, with its structure outlined in Figure 1, split into three cells, which are individual code blocks that can be run independently within a Notebook script. The first cell allows the user to enter most of the required inputs, and the second cell consisting of the remaining inputs and main optimisation loop, while the third cell visualises the results. The script optimises the foundation locations using their centre coordinates, with corresponding results displayed in a line and box plot.

1) Initial user inputs and constraints

In the first cell, the user is required to define the problem geometry and GA input parameters, such as the population size. The user is required to input the relative distances between adjacent foundations and their corresponding loadings. Using this information, the script calculates adjusted loadings for each randomly generated foundation layout in the initial population after the inputs have been entered. As the relative distances between the foundations are analogous to the span lengths between columns in the superstructure, varying the relative distances between foundations would alter the loads transferred to each foundation. Therefore, this approach accounts for these changes to the foundation loadings depending on their relative distances, for more accurate models.

A minimum spacing of 1 metre is imposed between adjacent foundations, whilst the foundation location boundaries (x-coordinates) are defined relative to tunnel geometry parameters to keep foundations within a relevant range, equal to 2.5 times the sum of the tunnel radius and lining thickness, extending in both directions from the centre of the tunnel.

Non-geometric constraints are implemented through a function that checks whether any adjusted foundation loadings exceed the bearing capacity of the soil, in which a simplified version of the Brinch-Hansen bearing capacity equation [16] is used, for undrained soils with square pad foundations:

$$q_f = 5.14 \times c_u \times 1.2 + p'_{ov} \quad (1)$$

where q_f is the ultimate bearing capacity of the soil (kPa), c_u is the undrained shear strength of the soil (kPa), p'_{ov} is the overburden pressure on the soil (mainly due to foundation loading), and 1.2 is a correction factor accounting for the foundation geometry. This check prevents FEA failures in Gofer, which would stop the optimisation process altogether and delay the workflow.

The soil is modelled as a linear-elastic-perfectly-plastic London clay under undrained conditions, with a unit weight of 20kN/m^3 and Young's Modulus of 70,000kPa. The selected failure criterion of the soil is the Mohr-Coulomb criterion, as it captures both elastic and plastic failure, including yielding and shear failure. However, the selected failure criterion of the tunnel lining is linear elastic, as plastic behaviour in concrete typically indicates catastrophic structural failure.

2) *Main optimisation loop*

This is the main calculation stage of the optimiser, found in the second cell. The user is prompted to enter the number of generations and assign names to the models uploaded to Gofer, allowing for easier tracking across iterations. Calculations are then performed within a nested loop, encapsulated in a core function which consists of 10 structured steps ("Main optimisation loop" in Figure 1). Steps 1-5 define the Gofer model – setting up the geometry, material properties, and analysis stages – compiled into a JSON-like structure. The required format of the JSON-like structure was obtained from downloading complete models (with analysis stages fully defined) from the Gofer GUI and inspecting their metadata. Steps 6-8 use HTTP requests (utilising API endpoints documented in Gofer's Swagger UI page [17]) to upload the model, run the analyses, and retrieve results. Step 9 extracts stress data for the tunnel lining from the returned results, before Step 10 computes the magnitude of the maximum hoop stress across all tunnels in the model. These are then appended to an array that contains all of the maximum hoop stresses across all iterations, for plotting after all iterations are completed. To calculate hoop stresses within the tunnel, the following equation is used (derived from 3D stress equations in [18]):

$$\sigma_{\theta\theta} = \sigma_{xx} \sin^2\theta + \sigma_{yy} \cos^2\theta - \tau_{xy} sin2\theta \quad (2)$$

where $\sigma_{\theta\theta}$ is the hoop stress (kPa), σ_{xx} and σ_{yy} are the horizontal and vertical effective stresses (kPa), and τ_{xy} is the shear stress (kPa). The angle θ represents the anticlockwise angle between the vertical line projected upwards from the centre of the tunnel and the line connecting the centre to the sampled elements. For each model, this calculation is performed on every element in the outer layer of the tunnel lining. This begins in Step 3, where sample points are defined within the tunnel. In Step 9, the script then searches for the nearest elements to these points within the tunnel lining. Therefore, the number of sample points needs to be large enough to capture all of these elements. For the tunnels used in the subsequent tests, this was 110.

3) *Genetic algorithm-styled optimisation stage and results visualisation*

These processes constitute the remaining stages of the optimiser tool. The maximum hoop stresses achieved are sorted from smallest to largest, and this order is then used to reorder the initial population array, which is then truncated and passed to a mutation function, which randomly perturbs the foundation locations. This helps to introduce variation whilst retaining the best layouts for the next generation. A smaller proportion of the best layouts of the previous generation are also retained in the next generation. Once the retained and mutated foundations have been collected into the next generation, they are input back into the nested loop that obtains the maximum hoop stresses. This then repeats until the specific number of generations is reached. The best foundation layouts of each generation and total computation time are then displayed to the user, alongside a line graph of the lowest maximum hoop stress values across all tunnels and a corresponding box plot of all the stresses in each generation.

III. Execution of the Script

The minima graph displaying the achievable range of maximum hoop stress values is generated first. Then, plots showing the effect of the fixed boundaries are shown. Following this, 3 GA input parameters are varied, alongside the incorporation of a modification to the optimisation process which consists of a decreasing perturbation range. Finally, examples of optimisations performed on more complex geometries are shown.

A. *Single-tunnel, two-foundation problem*

The problem to be optimised in the following tests is shown below, with the problem geometry shown in Figure 2.

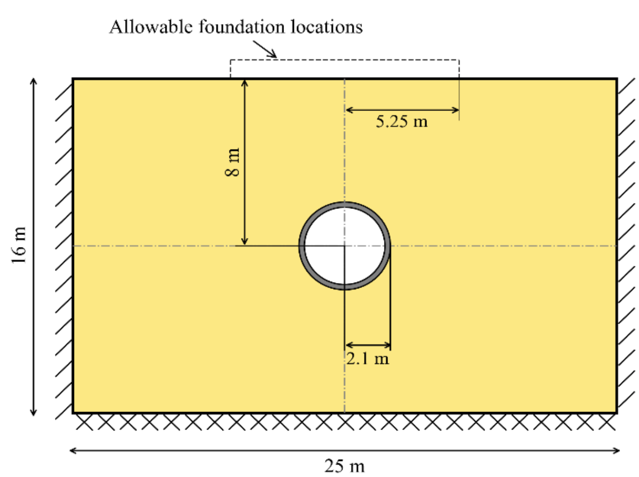

Figure 2: The default geometry of the single-tunnel problem, in which hoop stresses within the tunnel lining at the centre of the rectangular soil block is to be minimised.

Foundation boundaries are also overlaid on this diagram. The two foundations are of width 2m, and at least 1m apart. Fixed boundaries are also shown, for which the vertical edge boundaries are restrained horizontally, and the base boundary restrained in both the horizontal and vertical directions.

The tunnel lining thickness is 0.25m with a mesh size of 0.15m, which, while relatively coarse (resulting in a tunnel lining only two elements thick), was selected to balance accuracy with simulation times. Simulations with 100 iterations at this mesh size averaged 53 minutes (roughly 30 seconds per iteration), whilst reducing the mesh size to 0.08m (increasing tunnel lining to three elements thick) would double simulation times but only slightly increase the reduction and consistency of maximum hoop stress values post-optimisation.

1) Symmetric minima graph

The symmetric minima graph generated for this setup is shown by Figure 3 below:

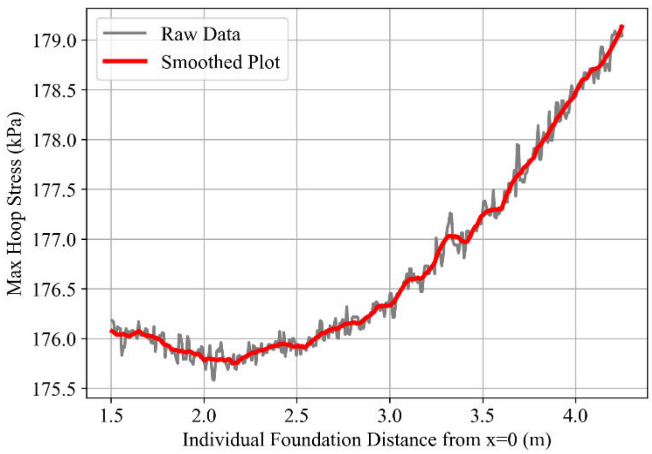

Figure 3: A graph showing the maximum hoop stress in the tunnel lining for a symmetrical foundation layout centred around the tunnel. A smoothed line is shown to account for the noise from the raw data.

To maintain precision with the optimiser (to 0.01m), the foundations were symmetrically arranged about the tunnel, each starting 0.25m from the tunnel centre and moved away in 0.01m intervals until reaching the outer foundation boundaries (4.25m). A smoothed line, using a Savitzky-Golay filter (red), more clearly shows the impact of foundation layout on maximum hoop stresses and its relationship with foundation distance from the tunnel. The minimum point of this smoothed line is found where foundations are approximately 2.15m away from the tunnel centre, with a corresponding maximum hoop stress value in the tunnel of around 175.7kPa, although slightly lower values would be reached by the optimiser due to the noise. In the raw data, the overall minimum is where foundations are slightly closer, just under 2.1m.

It is assumed that the symmetric foundation layout with this minimum stress is the best possible layout for simplicity, as the graph shows that for similar layouts, the hoop stress would be very similar. The range of achievable hoop stresses is in the range of 175.7-179kPa, indicating that improvements are limited. This prompted the use of 'bad' foundation pairs for the initial population of the following tests to better highlight stress reductions during optimisation. Note that for smaller mesh sizes, which were not used due to the high computation cost relative to the resulting noise reduction, hoop stress values would be shifted higher by a few kPa due to the closer proximity of the outer tunnel elements to the outer tunnel boundary, where hoop stresses are higher.

2) Effect of fixed boundaries

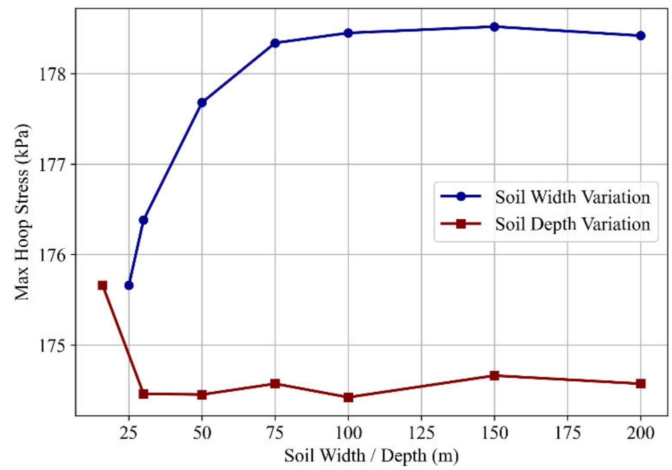

Figure 4: A graph showing the separate effects of the fixed edge and base boundaries on the maximum hoop stresses within the tunnel.

The effects of the boundaries at the soil block edges and the base edge on maximum hoop stresses in the tunnel were examined separately due to their differing fixity conditions. This was selected for a block with foundations at locations of ±2.15m, corresponding to a maximum hoop stress of 175.66kPa for the standard soil block size of 25 by 16m.

From Figure 4, as soil width increased from the base 25m, the corresponding rise in maximum hoop stress diminished, reaching a maximum increase of around 6kPa at 100m. The opposite occurred when soil depth was decreased, reducing by around 1.5kPa by a soil depth of 30m, from the base 16m. The observed fluctuations for both plots are likely caused by the coarse tunnel lining mesh. Due to its thin geometry relative to the soil block, the lining is highly sensitive to stress redistribution, amplifying variations in hoop stress. The reduction in maximum hoop stress with increased soil depth is attributed to a reduction in artificial confinement below the tunnel, allowing the additional stresses to redistribute more broadly and reducing stress concentrations on the base of the tunnel lining.

In contrast, increasing soil width led to greater hoop stresses in the tunnel lining. Observing the displacement vectors from the Gofer GUI in Figure 5 (next page) shows that arching occurs following the tunnel insertion, which is subsequently amplified by the foundation loadings. Horizontal confinement at the fixed edge boundaries lessens the extent of stress redistribution towards the bottom half of the tunnel, preventing the full development of the arching effect. As such, this affects the stress redistributions, and results in a less realistic calculation of hoop stresses within the tunnel due to the incomplete arching effect.

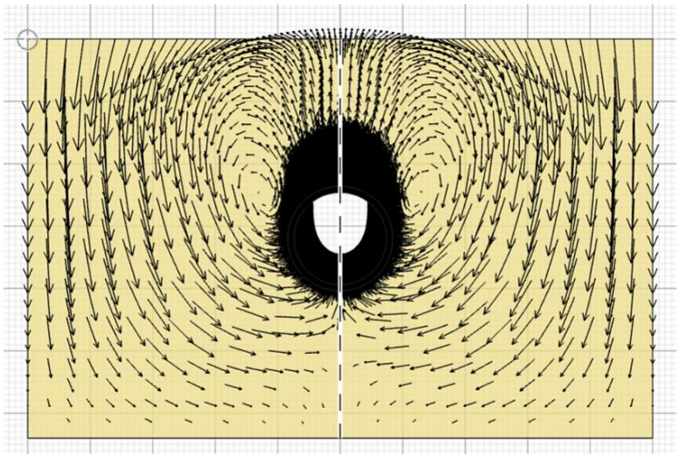

Figure 5: Screencap of a Gofer model post-analysis, displaying the resultant displacement vector plot around the tunnel, where arching is shown around the tunnel. Fixed horizontal constraints at the edge boundaries are also shown, limiting the extent of arching in the model.

Arching causes more soil displacement towards the base of the tunnel, which therefore leads to greater compressive stresses acting on the bottom half of the tunnel. Figure 6 below displays the hoop stress distribution in the tunnel lining, in which it can be seen that peak hoop stress values are found at the base of the tunnel.

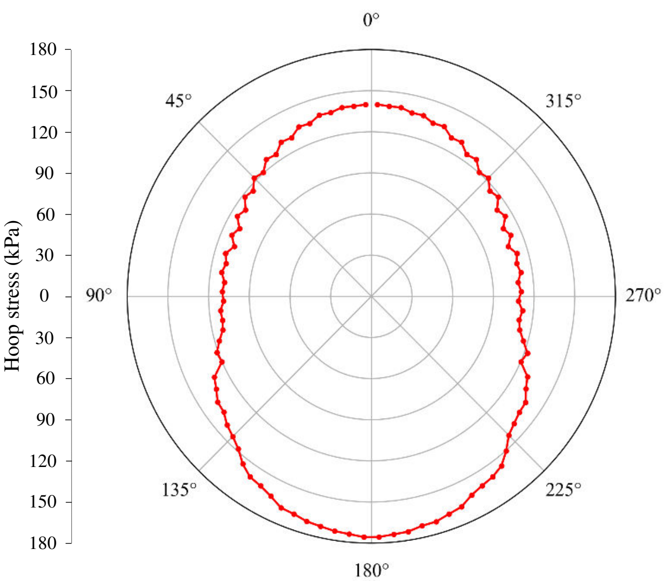

Figure 6: Polar plot of the hoop stress distribution around a circular tunnel lining, for which hoop stresses are compressive.

As such, increased arching would amplify these peak compressive hoop stresses at the tunnel base, leading to the increased maximum hoop stresses observed when the soil block width was increased.

The tunnel's zone of influence on the surrounding soil redistribution is restricted in the narrower soil blocks due to proximity to the fixed edge boundaries, limiting arching and resulting in reduced stress redistribution towards the tunnel base. Therefore, as the edge boundaries are moved further apart, the arching is less restricted and stresses are more concentrated at the bottom half of the tunnel, leading to increased maximum hoop stresses. Given that real-world conditions have effectively infinite boundaries, the confined soil block used in the optimisations will require correction to account for inaccuracies due to the constrained boundaries. As such, it is recommended that computed hoop stresses be validated against real tunnel data to verify findings and establish correction factors for boundary-induced inaccuracies. Furthermore, given that the increase in stresses due to increased soil block width was greater than the decrease for increased soil block depth, it is expected that the combined boundary effects are to reduce maximum hoop stresses from their actual values, due to the lack of arching and its stress redistributions.

3) Varying population size

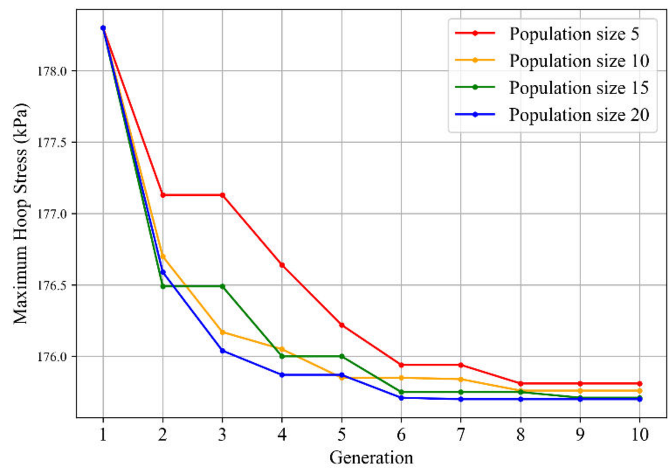

Figure 7: Graph showing the effects of using an increasing population size in the optimisation process on maximum hoop stresses within a tunnel across the 10 generations of the optimisation.

The number of generations was kept at 10, and the retention rate (of best solutions between generations) was kept at 20% and the mutation cutoff rate kept at 50%.

Larger population sizes resulted in lower final maximum hoop stress values and faster convergence rates, with the best hoop stress value reaching 175.7kPa, closely matching the theoretical minimum of ~175.6kPa from the symmetric minima graph (Figure 3). These solutions showed deviations of up to ±0.5m from the ideal symmetric foundation positions (each foundation 2.15m away from the tunnel centre), illustrated below in Figure 8.

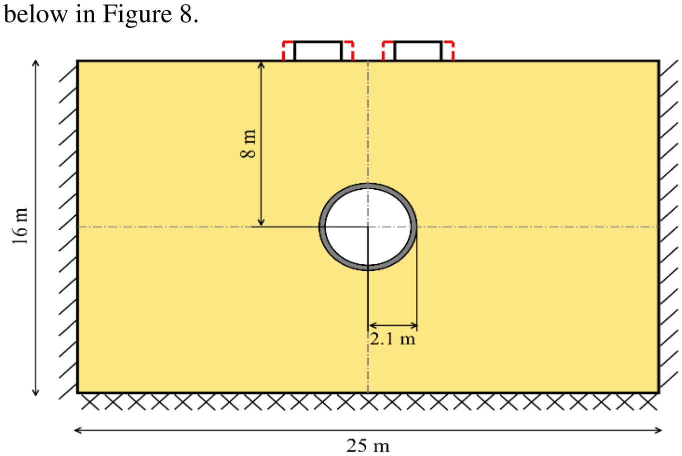

Figure 8: Illustration of the optimal foundation layout, with deviations of ±0.5m from the actual optimiser layouts shown in red.

In contrast, smaller populations exhibited less consistency in hoop stress reductions per generation, and therefore slower convergence rates and greater variability in final hoop stress values. This is attributed to the reduced amount of exploration and refinement that is possible in each generation and highlights the importance of sufficiently large population sizes to ensure robust optimisation outcomes.

Observed deviations arise again from the coarse tunnel lining mesh used to reduce simulation times, as seen in the fixed boundary experiment. Furthermore, the surrounding soil mesh is distorted by the moving foundation nodes, which further adds variability in hoop stresses within the tunnel lining, as also reflected in the symmetric minima graph.

4) Varying the retention-mutation ratio

To minimise fluctuations observed with smaller populations, a population size of 15 was selected to account for mutation randomness in each generation and better isolates the effects of varying retention rates.

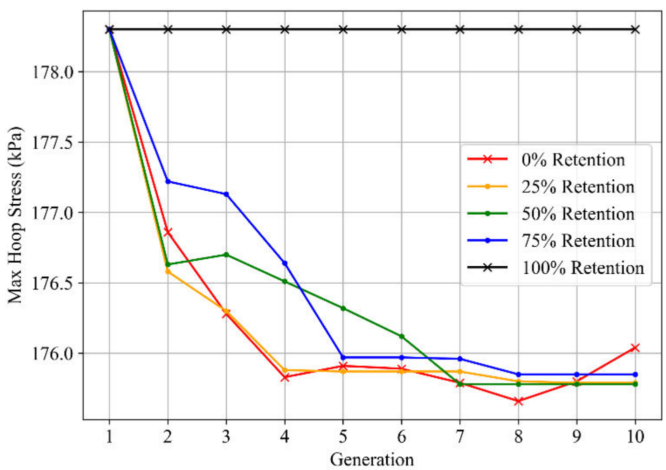

Figure 8: Graph showing the effects of varying the proportion of individuals that are retained in the next generation and individuals that are mutated in the next generation on maximum hoop stresses across the optimisation.

Figure 8 shows that retention ratios had a minor effect on final hoop stress values as well as convergence rates. The lowest maximum hoop stress achieved was also at 175.7kPa, with corresponding foundation coordinate layouts within ±0.5m of the symmetric optimum, as seen in the preceding tests.

No single retention ratio emerged as universally optimal, but retention rates lower than 50% generally improved convergence in the first few generations, by allowing a greater proportion of the generation to mutate and explore the solution space.

Higher retention rates limited mutation-based exploration in each generation, which reduced convergence rates. As expected, retaining all results would lead to no improvements at all, whilst no retention would lead to hoop stress values sometimes increasing in subsequent generations, highlighting the requirement of retention for the best foundation layouts from previous generations, to ensure continuous improvement occurs instead of regression.

5) Varying the mutation cutoff rate

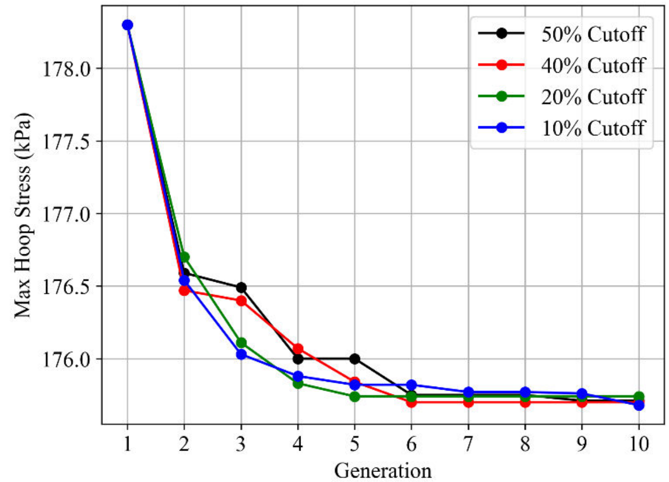

Figure 9: Graph showing the effects of varying the percentage of top performing foundation pairs carried over to generate the next generation, on maximum hoop stresses throughout the optimisation.

The mutation cutoff rate refers to the threshold percentage below which foundation pairs are selected for mutation and are gradually reduced from the default rate of 50% in this test. Figure 9 shows that final hoop stress values remain consistent across all mutation cutoff rates at 175.7kPa, with foundation layouts deviating up to ±0.5m, similar to the preceding tests. However, smaller mutation cutoff rates below 40% yield faster overall convergence rates, as the likelihood of selecting the best foundation pairs of the previous generation to mutate is reduced due to a larger proportion of the top foundation pairs being carried forward for mutation.

6) Reducing mutation range over time

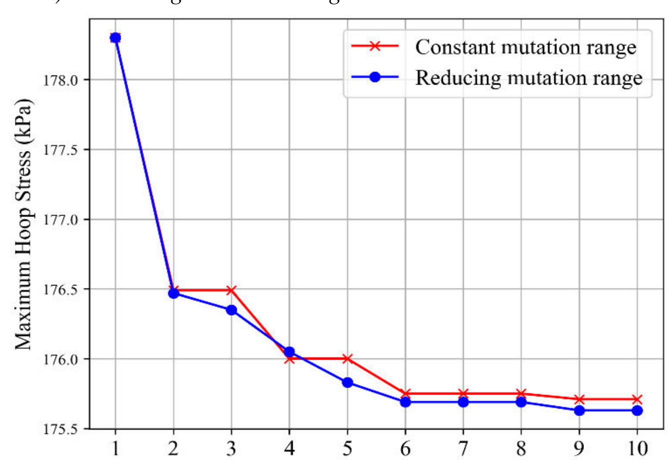

Figure 10: Graph comparing maximum hoop stresses between optimisation schemes employing a constant mutation range and a mutation range which halves after every three generations.

The mutation range was halved every 3 generations, with 3 generations being selected to ensure that solutions had gotten sufficiently close to the global minimum.

Figure 10 shows that the reducing mutation range achieved greater maximum hoop stress reductions, at 175.63kPa. Furthermore, the best foundation pair was within ±0.3m of the optimal symmetric layout, indicating an improvement in both accuracy and precision of hoop stress values despite persistent mesh-induced noise effects, as well as highlighting the optimiser's extended capabilities. However, smaller population sizes can negatively impact convergence rates, as the mutation range may decrease before sufficient progress towards the global minimum is made, leading to premature convergence on local minima or much slower convergence towards the global minimum. This therefore highlights the importance of the selection of an appropriate population size and mutation range reduction rate, to ensure accurate and efficient optimisation across a wider range of geometries. The optimal reduction rate will likely vary with problem geometry and warrants further testing, to ensure it aids the optimisation process instead of negatively affecting it. For more complex geometries, optimal foundation layouts are less likely to be reached in just a few generations, requiring more exploration of the solution space which a rapidly reducing mutation range would hinder, limiting the algorithm's ability to effectively navigate towards the global minimum.

B. Further discussion of single-tunnel geometry results

The optimiser achieved high accuracy, with final maximum hoop stress values within ±0.1kPa of the theoretical minimum for the single-tunnel problem, although real-world hoop stresses would vary due to the effects of the artificial confinement from the fixed boundaries of the soil block, which is expected to be a few kPa higher in reality due to the incomplete development of soil arching within the tunnel.

Foundation locations consistently converged towards the optimal layout with a foundation location precision within ±0.5m, limited by the coarse tunnel lining mesh, which, due to its thin geometry relative to the soil block, is highly sensitive to numerical effects, and may limit the optimiser's utility in detailed design with strict space constraints. However, the impact of positional inaccuracies – typically affecting hoop stresses by only a few 0.1kPa – is likely to be outweighed by real-world uncertainties such as soil anisotropy or construction imperfections, which could introduce greater variability in actual hoop stress values.

Systematic sampling was intentionally omitted in the tests to allow for greater improvements, in which foundation pairs were all initially far apart near the outer foundation location boundaries, to enable clearer observation of optimisation behaviour. In practice, systematic sampling would accelerate convergence by increasing the likelihood of initial foundation placements near the global minimum and reducing the minimum number of generations required to reach such minima.

In addition, the limited range of achievable improvements (~3.5kPa) in this problem reduced the sensitivity to changes in GA parameters. Even with suboptimal initial populations, the default mutation range of 1.05m would allow convergence to the theoretical optimum of 2.15m from the tunnel centre in 3 generations (starting near the outer bounds which were 4.25m from the tunnel centre), which rendered the effects of varying GA input parameters less pronounced.

C. Two-tunnel, three-foundation problem

The tunnels were of the same dimensions as those used in the single-tunnel problem and spaced 3m apart, centred within the soil block. Three foundations were used to prevent excessive spacing, with systematic sampling applied to ensure thorough exploration of the solution space in the first generation.

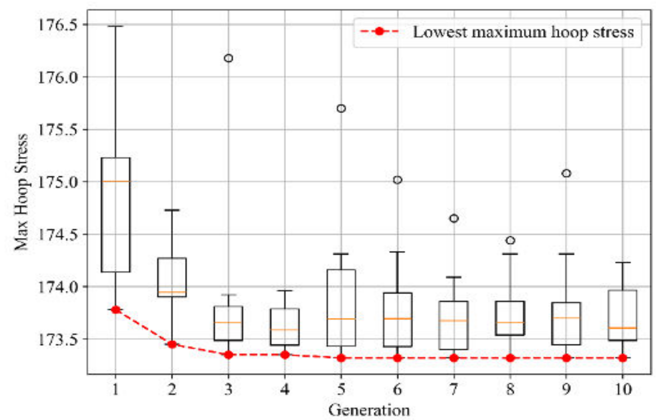

Figure 12: A box and line plot, showing the spread of maximum hoop stress values in each generation of the optimisation alongside the minimum hoop stress value achieved in that generation.

Figure 12 shows that the optimiser achieved reductions of about 0.5kPa to 173.23kPa, taking 5 generations to reach this value. The largest maximum hoop stress was found in the first generation, with a value just under 176.5kPa, suggesting that the range of possible improvement is approximately 3.2kPa, similar to that of the single-tunnel problem, and indicates that the use of systematic sampling reduces the amount of required improvement to reach optimal minimum values. Final foundation layouts did vary more than for the single-tunnel problem, yet still consistently resulted in similar hoop stresses, in the range of 173.2-173.3kPa, suggesting that the maximum hoop stress sensitivity towards foundation locations is similarly low to that of the single-tunnel problem. This also suggests that multiple configurations can yield similar performance.

This greater variability in optimal layouts arises from the increased complexity of arranging three foundations rather than two within the same location boundaries. As the tunnels are identical and symmetrically arranged, optimal foundation layouts are therefore able to be horizontally inverted. This may be advantageous in practical scenarios as it allows for greater flexibility in design.

Overall, the results reinforce that while the optimiser is successfully converging towards minimal optimisation values, it is still better suited for early-stage optimisation, due to uncertainties in foundation locations due to the mesh-induced noise impacting maximum hoop stress values.

IV. Conclusion

A. Significance of results

The results demonstrate that the optimiser tool can accurately minimise hoop stresses within tunnels through foundation layouts, achieving a location precision of ±0.3m and maximum hoop stress accuracy of ±0.1kPa with appropriate inputs and an appropriately selected adaptive mutation range for a single-tunnel problem. While coarse tunnel lining meshes introduce noise that prevents a finer precision, this trade-off enables faster optimisation times, with models for the single-tunnel problem taking approximately 30 seconds to setup, analyse, and calculate maximum hoop stresses from.

Although no tunnel data was available for validation of computed hoop stress values, corrections will likely be neccessary to account for fixed boundary effects in the FEA models, as the narrow soil width used limits the development of soil arching around the tunnel, reducing maximum hoop stresses obtained for the models. Given the narrow range of achievable hoop stresses in the tested geometries, real-world factors such as soil anisotropy and construction imperfections would likely have a greater influence on final hoop stress values.

As such, the main advantage of the optimiser lies in its automated FEA capabilities, avoiding the need for extensive manual setup, offering a simpler alternative to traditional geotechnical modelling workflows, only requiring navigation through the Jupyter Notebook script to use. However, its current precision in foundation locations may pose an issue if permissible foundation locations are very strict.

B. *Future Improvements and Suggestions*

To improve the optimiser's precision without significantly increasing simulation time, a more efficient approach to mesh refinement could be adopted. This involves manually refining the mesh within tunnel linings, where numerical sensitivity is highest. A coarser mesh will be manually specified around the tunnel using user-defined points, instead of allowing Gofer to automatically determine the mesh sizes around tunnels, as the surrounding soil is less numerically sensitive. Such a mesh configuration would significantly reduce the overall element count that would otherwise be used in the FEA for a small tunnel lining mesh, to reduce additional optimisation times whilst improving hoop stress accuracies and corresponding foundation location precision.

Another key improvement would involve quantifying and correcting the effects of the fixed soil boundaries. As it may be difficult to obtain hoop stress data on existing shallow tunnels in the current London Underground, similar shallow tunnels could be used for hoop stress value testing, which can then be modelled and tested on Gofer, with computed hoop stresses across the tunnel being compared with its real-world counterpart.

Further developments could also leverage more of Gofer's built-in geotechnical capabilities. Incorporating soil stratification, hydraulic gradients, or other modes of soil analysis would increase the tool's applicability to real-world conditions. The condition of the surrounding soil could also be explored through the incorporation of functions that minimise soil displacements for constructability and longevity of the building, alongside the main function of minimising maximum hoop stresses within the tunnels of the problem geometry.

Nonetheless, addressing mesh-induced noise remains the most critical priority. Without improved mesh precision, future modifications may still be undermined by the persistent mesh-induced noise that prevents the optimiser from converging towards the same foundation layout, even with carefully and appropriately selected input parameters.

ACKNOWLEDGEMENTS

I would like to thank Professor Augarde for his consistent guidance, perceptive feedback, and encouraging mentorship, which enabled me to navigate through the numerous challenges that arose in this project. Additionally, I would also like to express my deep gratitude to the helpful team at Oasys, for always patiently assisting with the extensive scope of queries I had whilst developing and testing the optimiser tool.

REFERENCES

[1] E. Ovando-Shelley. "CASE HISTORY APPROACH TO GEOTECHNICAL DESIGN ON HERITAGE STRUCTURES: MEXICO CITY´S METROPOLITAN CATHEDRAL (1989-2009)". [Online PDF]. Available: https://www.uvm.edu/~dwporter/ncptt/papers/ovandoShelley_Geotech_text%20only.pdf

[2] Bentley Systems, *Plaxis 2D.* (2024). Bentley Systems. [Online]. Available: https://www.bentley.com/software/plaxis-2d/

[3] K. Terzaghi, R. B. Peck, G. Mesri, *Soil Mechanics in Engineering Practice,* 3rd ed. John Wiley and Sons: Hoboken, New Jersey, USA, 1996.

[4] D. M. Potts, L. Zdravković, *Finite Element Analysis in geotechnical engineering – Application*, 1st ed. Thomas Telford Limited: London, UK, 2001.

[5] A. Lees, *Geotechnical Finite Element Analysis: A practical guide*, 1st ed. Thomas Telford Limited: London, UK, 2016.

[6] D. Boumezerane, "Recent Tendencies in the Use of Optimisation Techniques in Geotechnics: A Review", *geotechnics,* Vol. 2, pp. 114-132, 2022.

[7] B. Li, L. Zhang, J. Yuan, "A simple Monte Carlo simulation method for geotechnical reliability-based design", *Georisk*, Vol. 18(3), pp. 558-569, June 2023.

[8] D. R. Wijesinghe, A. Dyson, G. You, et al, "Simultaneous slope design optimisation and stability assessment using a genetic algorithm and fully automatic image-based analysis", *Numerical and analytical methods in Geomechanics.* Vol. 46(15), pp. 2868-2892, October 2022.

[9] M. S. Islam, M. Rokonuzzaman, "Optimized design of foundations: an application of genetic algorithms", *Australian Journal of Civil Engineering.* Vol. 16(1), pp. 46-52, March 2018.

[10] H. Rezk, A.G. Olabi, T. Wilberforce, et al. "Metaheuristic optimisation algorithms for real-world electrical and civil engineering application: A review," *Results in Engineering,* Vol. 23, Sep 2024.

[11] A. R. Kashani, A. H. Gandomi, "Facilitating Geotechnical Engineering Problems Using Metaheuristic Optimisation Algorithms," *Rocscience.* [Online]. Available: https://www.rocscience.com/learning/facilitating-geotechnical-engineering-problems-using-metaheuristic-optimization-algorithms

[12] H. Liu, L. Sun, D. Dias-da-Costa, "State-of-the-art review on the use of AI-enhanced computational mechanics in geotechnical engineering," *Artificial Intelligence Review*. Vol. 57, July 2024.

[13] A. Satyanaga, G. D. Aventian, T. Makenova, et al. "Building Information Modelling for Application in Geotechnical Engineering", *infrastructures*. Vol. 8, June 2023.

[14] H. Tian, Y. Wang, D. Zhang, "Real-time model updating and prediction of three-dimensional time-varying consolidation settlement using machine learning," *Journal of Rock Mechanics and Geotechnical Engineering*, vol. 16, Oct 2024.

[15] Oasys, *Gofer*. Oasys. [Online]. Available: https://oasysgofer.com/

[16] M. J. Tomlinson, "The general principles of foundation design", in *Foundations design and construction*, 7th ed. Prentice-Hall, London, UK: Prentice-Hall, 2001.

[17] Oasys, "Gofer API*"*. api.gofer.oasys-software.com. Accessed: Jan. 8, 2025. [Online]. Available: https://api.gofer.oasys-software.com/docs/#/

[18] A. F. Bower. "Vectors and Tensor Operations in Polar Coordinates." brown.edu. Accessed: Feb. 12, 2025. [Online.] Available: https://www.brown.edu/Departments/Engineering/Courses/En221/Notes/Polar_Coords/Polar_Coords.htm