

Comparison of tabular, support-vector, and neural network models for granular elastoplasticity

Biswajit Banerjee

Parresia Research Limited, Auckland, New Zealand

b.banerjee.nz@gmail.com

Monday 21st December, 2020

Abstract

Experiments and micromechanical simulations generate tabular data for material behavior. Typically, models are fit to these material data before engineering simulations can be performed. In this paper, we compare the response of elastic moduli models for a dry, poorly-graded, sand that use linear interpolation, support vector regression fits, and multilayer perceptron neural networks, respectively. The simulations are driven by deformation gradient and represent hydrostatic loading/unloading, uniaxial strain, and multiaxial strain. We observe that the directly interpolated elastic moduli produce the smallest stresses and the largest elastic strains. The support vector and neural network models produce stiffer responses and also take longer to compute.

1 Introduction

As more microscale simulation data become available, tabulated data and machine learning models fitted to these data will increasingly be used as inputs to elastoplastic simulation. In previous work, we have explored the use of support vector models (Banerjee, Fox, and Regueiro, 2020c) and neural network models (Banerjee, Fox, and Regueiro, 2020b) to generate functional representations of multi-variable tabular data for elastoplastic simulations. Various approaches to interpolating tabular data, e.g., linear interpolation, radial basis function interpolation, and kriging interpolation have been discussed in Banerjee, Fox, and Regueiro (2020a). An elastoplasticity algorithm that uses linear interpolation of tabular data has been described in detail in Banerjee, Fox, and Regueiro (2020d).

In this paper, we use the tabular elastoplasticity algorithm to compute the response of single material points to prescribed deformation gradient time histories. Hydrostatic, uniaxial strain, and multiaxial strain deformation histories are used to drive the simulations with the host code *Vaango* (Banerjee, 2014). Comparisons are made between the responses of tabular model with tabular elastic moduli, moduli computed from a support vector regression fit, and moduli from a neural network fit to the data.

This paper is organized as follows. Experimental data on a dry, poorly-graded, sand are presented

in Section 2 and the support vector and neural network fits to the data are also discussed briefly in that section. Single particle tests and comparisons are described in Section 3. Finally, some brief concluding remarks are given in Section 4.

2 Experimental data

The experimental data that have been used as a test-bed for the modeling process are for a dry, poorly-graded, concrete sand described by Fox et al. (2014) and tested at the University of Maryland.¹ Further details on the particular set used in this work can be found in Banerjee, Fox, and Regueiro (2020c). In this paper, compressive states are assigned positive values and tensile states are assigned negative values.

The hydrostatic loading-unloading data for that sand is presented as a plot of pressure (p) as a function of the total volumetric strain (ϵ_v) in Figure 1(a). The loading curve, shown in cyan, is used to fit a crush-curve model. The unloading curves are used to fit a bulk modulus model that depends on the plastic strain. Tangents to the unloading curves represent the bulk modulus and have been plotted in Figure 1(b). There is a strong dependence of the bulk modulus (K) on both the elastic (ϵ_v^e) and plastic (ϵ_v^p) volumetric strains.

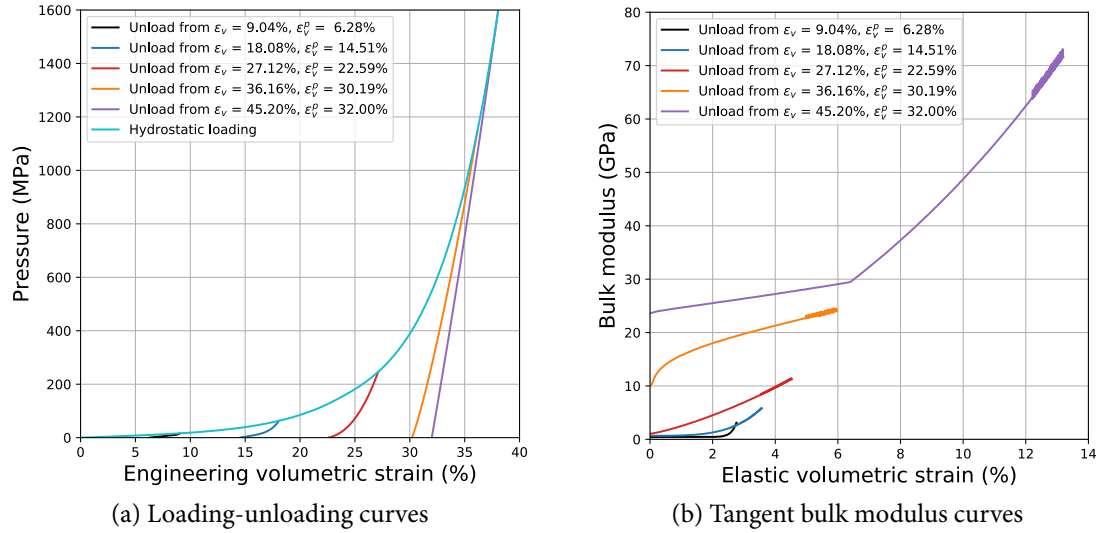


Figure 1 – Hydrostatic loading-unloading data and unloading bulk moduli for dry poorly-graded concrete sand.

The crush-curve extracted from the hydrostatic compression data is depicted in Figure 2(a). Since the term “crush-curve” is used more commonly to refer to the change in porosity as a function of pressure form, we shown this form of the curve in Figure 2(b). The porosity (ϕ) has been computed using $\phi = p_3 - \epsilon_v^p$ where ϵ_v^p is the volumetric plastic strain and $p_3 = 0.325$ is the volumetric plastic strain at which all pores have been crushed (Brannon et al., 2015).

¹Stephen Akers, 2018, Private communication, CCDC Army Research Laboratory, Aberdeen Proving Ground, MD, USA

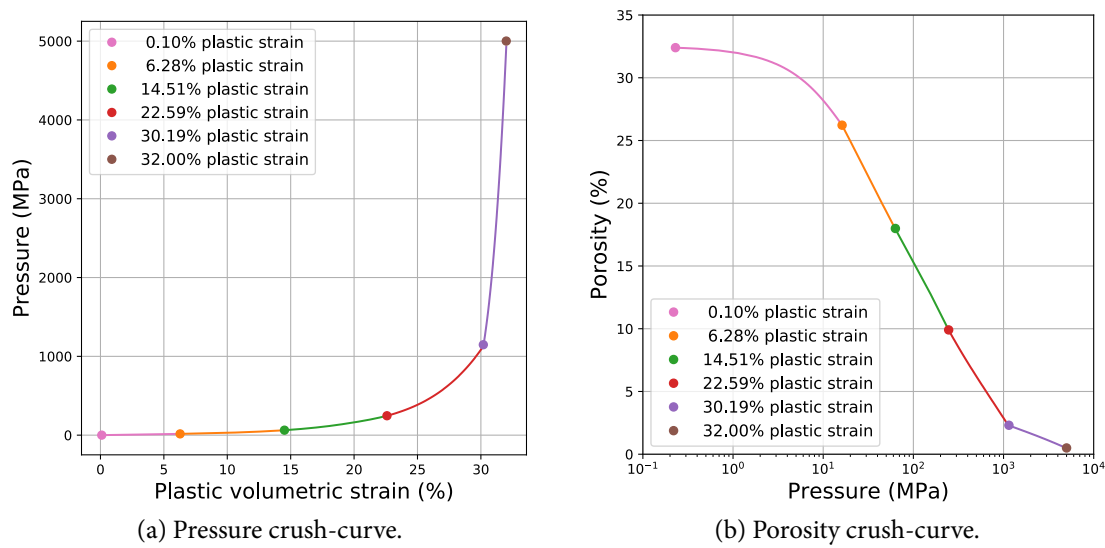


Figure 2 – Crush-curves for dry poorly-graded concrete sand.

The original yield function curve for the concrete sand intersected the deviatoric stress axis approximately 5 kPa (compression), indicating that a purely elastic response was not possible when the sand was loaded from zero strain. To ensure that simulations did not run into any difficulty, the function was shifted to the tensile regime by 5 kPa to provide a small amount of cohesion. A plot of the yield function in p - q space is shown in Figure 3. Here p is the mean stress, defined as $p = 1/3 \text{tr}(\sigma)$, q is the deviatoric yield stress given by $q = \sqrt{3J_2}$ where $J_2 = 1/2 \mathbf{s} : \mathbf{s}$, σ is the Cauchy stress, $\mathbf{s} = \sigma - p\mathbf{I}$ is the deviatoric part of the Cauchy stress, and \mathbf{I} is the second-order identity tensor. The crush curve suggests that nonlinear bulk moduli are not sufficient to explain the observed yield response and a compression cap is needed to model the material.




Figure 3 – Yield function for dry poorly-graded concrete sand.

In the simulations described in this paper, tabular models are used for the crush curve and yield function as described in (Banerjee, Fox, and Regueiro, 2020d) with linear interpolation of the data (Banerjee, Fox, and Regueiro, 2020a). Comparisons with support vector regression (Banerjee, Fox, and Regueiro, 2020c) and neural network models (Banerjee, Fox, and Regueiro, 2020b) are performed by varying the approach used to model the elastic bulk modulus function.

2.1 Tabular model

The bulk modulus vs. pressure data shown in Figure 4 are interpolated to determine the value of the tangent bulk and shear moduli at a given value of mean stress and volumetric plastic strain. Derivatives of the bulk modulus with respect to plastic strain are computed using a central difference scheme using two nearby values of plastic strain that differ by 2.0×10^{-6} . A constant Poisson's ratio of 0.189 is used to compute the shear modulus from the interpolated bulk modulus.




Figure 4 – Bulk modulus table expressed as a function of pressure.

The data are read as a JSON file (Crockford, 2006). The format of the file is described in Banerjee (2014).

2.2 Support vector regression model

The support vector regression (SVR) model used in these simulations is the pressure vs. total volumetric strain fits to the experimental data compute with a value of $C = 10$ and $\epsilon = 0.001$ as described in Banerjee, Fox, and Regueiro (2020c). Pressure curves predicted by the model are shown in Figure 5. Bulk modulus are extracted from the predicted pressures by analytical differentiation of the pressure SVR model.

The SVR model is also read into Vaango as a JSON file as described in Banerjee (2014).




Figure 5 – Support vector regression predictions of pressure as a function of the total volumetric strain.

2.3 Multiplayer perceptron neural network model

A dense, multilayer perceptron (MLP) model is used to model the elastic modulus data as described in Banerjee, Fox, and Regueiro (2020b). The three-layer model of the bulk modulus as a function of elastic volumetric strain is used in this paper. The model containing 64 neurons with ReLU activation in layer 1, 32 neurons with ReLU activation in layer 2, and 32 neurons with ReLU activation in layer 3. The model was trained for 800 epochs with batch sizes of 32. Predictions from that model are shown in Figure 6. Note that this model uses a linear extrapolation of the experimental pressure curves before computing the bulk modulus instead of assuming constant bulk moduli outside the range of the experimental data. The fitted model is saved as a HDF5 file (Koranne, 2011) by Tensorflow 2.0 (Abadi et al., 2016) and read into Vaango using the HDF5 library.

3 Single particle tests

The Vaango implementation of the tabular plasticity model in conjunction with the tabular elasticity model (Table), the support vector regression elasticity model (SVR), and the multilayer perceptron elasticity model (MLP) has been exercised on single particles driven by a deformation gradient curve. Three tests are discussed in this paper: hydrostatic strain compression with unloading, uniaxial strain compression with unloading, and multiaxial strain loading and unloading.

3.1 Hydrostatic strain loading and unloading

The deformation gradients used to drive the hydrostatic strain test problem are shown in Figure 7. The components F_{xx} , F_{yy} , and F_{zz} of the deformation gradient F are set to equation values. The shear deformation components, F_{xy} , F_{yx} etc., are set to zero. The material is compressed and then

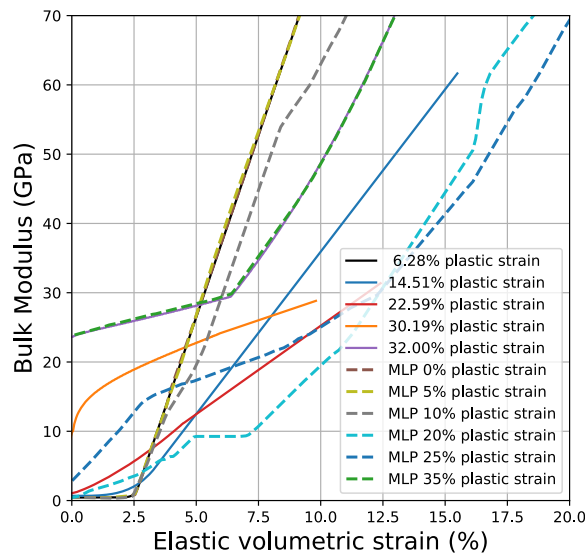


Figure 6 – Multilayer perceptron neural network predictions of bulk modulus as a function of the elastic volumetric strain.

unloaded in a sequence of steps that mimic the experiments described in Section 2.




Figure 7 – Input deformation gradient history for hydrostatic loading and unloading of a single particle.

The bulk modulus evolution for this deformation history is shown in Figure 8. The SVR model produces the smoothest bulk modulus evolution, while the values predicted by the other models tend to fluctuate considerably. The Table model predicts the smallest overall bulk moduli.

Components of the stress, σ , computed using the algorithm described in Banerjee, Fox, and Regueiro




Figure 8 – Evolution of bulk modulus for hydrostatic loading and unloading of a single particle.

(2020d) and the associated invariants, p and q , are shown in Figures 9(a) and (b), respectively. The stress components σ_{xx} , σ_{yy} and σ_{zz} are equal while the shear stress component, σ_{xy} etc., are zero as is the deviatoric stress invariant, q . The Table model produces the smallest peak normal stress while the MLP model produces the largest value. Stresses predicted by the SVR model are almost identical to those from the Table model until the final loading step. At that point the SVR and MLP models produce nearly identical results.

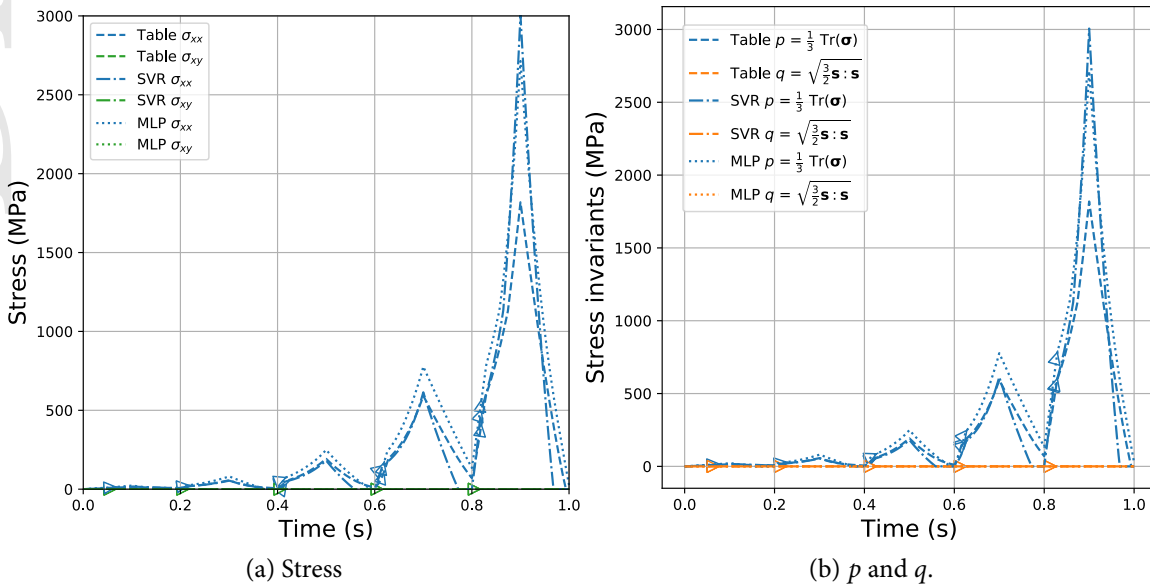


Figure 9 – Comparison of computed stress history for hydrostatic loading and unloading of a single particle.

Plots of the predicted elastic ($\boldsymbol{\varepsilon}^e$) and plastic strains ($\boldsymbol{\varepsilon}^p$) are shown in Figure 10. The MLP model produces the stiffest elastic response and, consequently, the smallest elastic strains, $\boldsymbol{\varepsilon}_{xx}^e = \boldsymbol{\varepsilon}_{yy}^e = \boldsymbol{\varepsilon}_{zz}^e$. The shear strain components are zero. The largest elastic strains are predicted by the SVR model while the Table model predicts strains that are close to the tabular model. On the other hand, because of the assumption $\boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}^e + \boldsymbol{\varepsilon}^p$, the plastic strains $\boldsymbol{\varepsilon}_{xx}^p = \boldsymbol{\varepsilon}_{yy}^p = \boldsymbol{\varepsilon}_{zz}^p$ predicted by the MLP model are the largest while the Table and SVR models produce nearly identical predictions. While the plastic strain components predicted by the Table model are monotonically increasing, both the SVR and MLP models show a decrease in plastic strain at the start of reloading after an unload step. This is because the deformation pushes the stress state to the tension regime during which the plastic strain tensor changes sign.




Figure 10 – Comparison of computed elastic and plastic strain history for hydrostatic loading and unloading of a single particle.

Plots of the normal components of stress as functions of the corresponding normal components of strain are presented in Figure 11. We observe that the algorithm produces the correct stress-strain symmetry because the plots of $\sigma_{xx} - \varepsilon_{xx}$ and $\sigma_{yy} - \varepsilon_{yy}$ are identical. The Table model produces results that are close to those predicted by the SVR model up to strains of around 13%. The MLP model predicts larger stresses after around 4% strain. However, beyond 13% strain, the SVR model produces a stiffer response than either the Table or the MLP model. Beyond 14% strain the SVR model produces the largest stress while the Table model predicts a value that is almost half that predicted by the SVR model.

Figure 12 shows plots of the pressure (mean stress) as functions of various components of the volumetric strain, ε_v . The experimental pressure-volumetric strain data are plotted as solid blue lines. The crush-curve extracted from experiments is shown as solid green line. All three models predict stresses that are lower than experimental values at peak compression. Table model predictions, shown as dashed lines, indicate a response that is softer than those from the MLP and SVR predictions. Plots of the pressure as a function of the volumetric plastic strain, ε_v^p , are shown in green. The models follow the crush-curve accurately. However, the peak plastic strains achieved by the Table model are smaller than those reached by the SVR and MLP models. Because of the rapid increase in stress at




Figure 11 – Comparison of computed stress-strain history for hydrostatic loading and unloading of a single particle.

these plastic strain values, small differences in plastic strain can lead to large differences in the stress state. The elastic response, the pressure as a function of the elastic strain (ϵ_v^e), is shown in orange in the figure. The initial response of the MLP model is stiffer than that of the Table and SVR models. However, as deformation progresses, the SVR model produces a stiffer response and catches up with the MLP model while the Table ode continues to predict a softer elastic response. These results indicate that the calibration of the elastic model and the crush curve from the experimental hydrostatic load-unload data requires consideration of deformation-induced anisotropy and elastic-plastic coupling for better matches to the experimental data are to be achieved.

The evolution of the yield surface for the three models is shown in Figure 13. Solid lines are used to represent the location of the cap at the values of ϵ_v at the peak of each loading step. The corresponding locations of the cap predicted by the Table, SVR, MLP models are shown as dashed, dash-dot, and dotted lines, respectively. The SVR model produces the best approximation to the experimental data. However, as observed in previous figures, the peak pressure achieved by the simulation is close to half that achieved in the experiments. The predicted values are relatively close to experimental values only for volumetric strains less than 27%. The Table and SVR models produce approximately the same response for strains less than 35%, while the MLP model generates higher values up to around 42% strain.

Figure 14(a) plots the evolution of the size of the timestep as the simulation proceeds. The timestep size decreases during the loading stage and increases during unloading because of the decrease in bulk modulus during unloading. However, since the bulk modulus continues to increase as compression increases with time, the timestep also decreases according to the computed p-wave speed. The MLP model initially requires the smallest timesteps but the size of the timestep required by the SVR model becomes the smallest after around 0.5s of simulation time. Memory usage evolution during the course of the simulation is shown in Figure 14(b). As the yield surface expands, the number of points needed to discretize the yield surface also increases and this is reflected in the memory

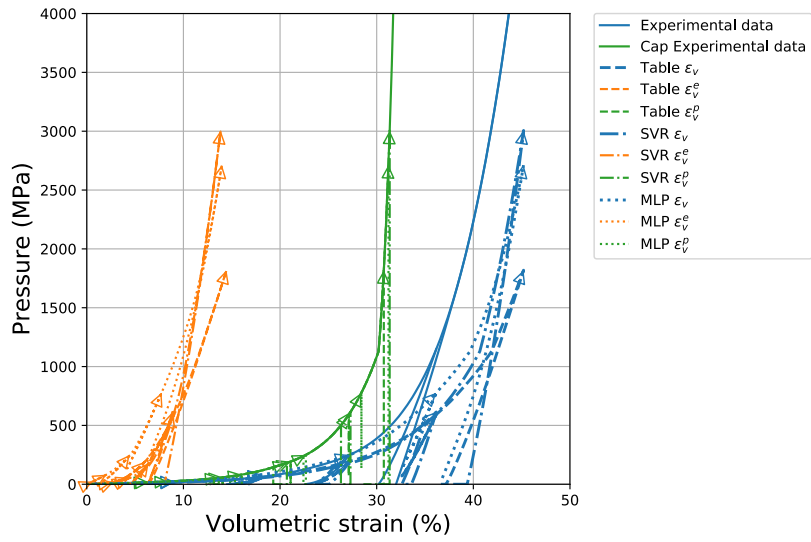


Figure 12 – Comparison of computed pressure-volumetric strain history for hydrostatic loading and unloading of a single particle.

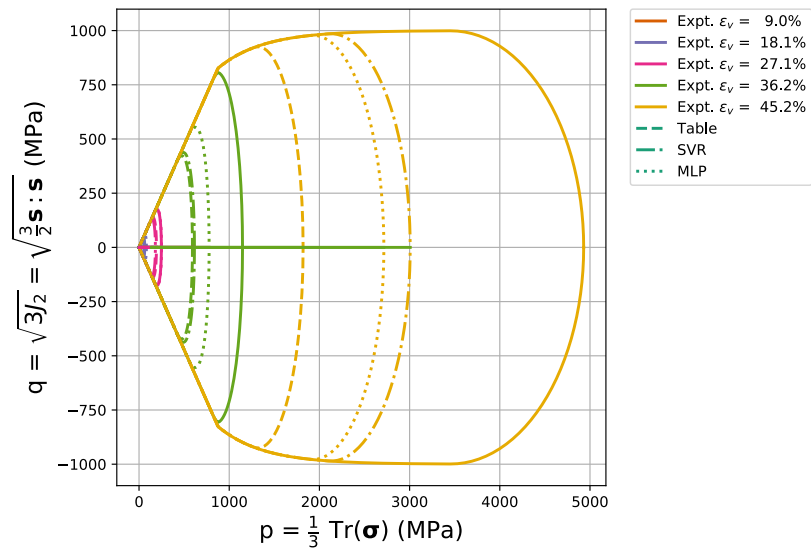


Figure 13 – Evolution of the yield surface for hydrostatic loading and unloading of a single particle.

usage. Memory leaks are avoided in Vaango by the use of smart pointers.

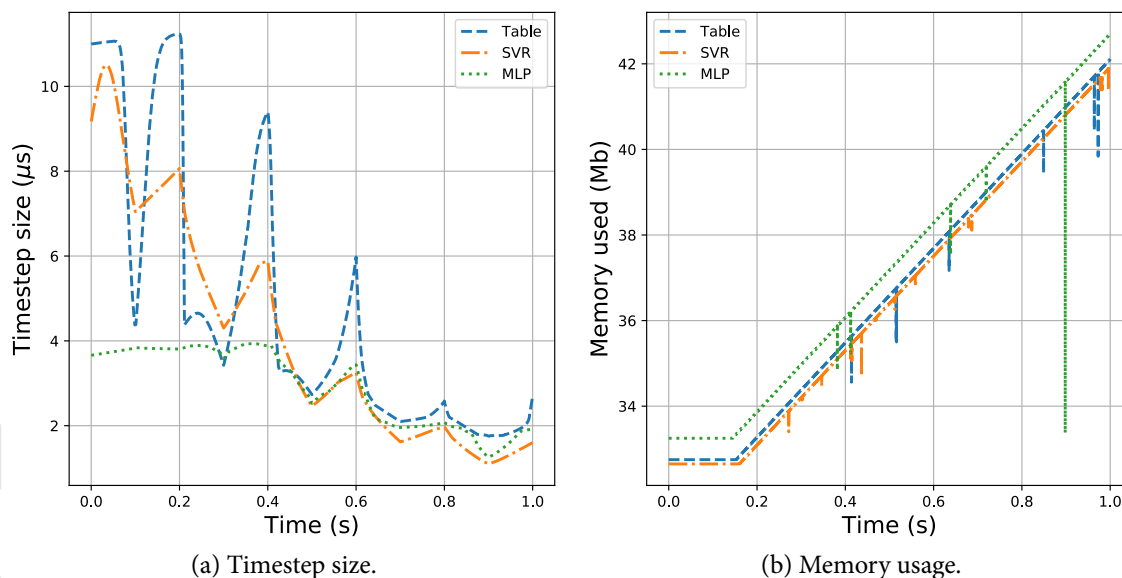


Figure 14 – Timestep size evolution and memory usage during hydrostatic loading and unloading of a single particle.

The time taken to compute each timestep is shown in Figure 15(a). All three models take approximately the same time to run during the course of the simulation, except for the SVR model which takes approximately twice as long per timestep towards the end of the simulation (1s). Figure 15(b) shows the time taken by each simulation. The Table model is the most efficient in terms of computation, followed by the MLP model. The SVR model is the least efficient in hydrostatic loading/unloading. The time taken is, of course, a function of the computed bulk modulus and the volumetric plastic strain which determine how efficiently closed points and normals to the yield surface can be evaluated.

3.2 Uniaxial strain loading and unloading

The hydrostatic test has been compared with experimental data. However, uniaxial strain experimental data are not available for the poorly graded dry concrete sand. Therefore, the results in this section are predictions from the tabular cap plasticity model that can potentially be compared to observations from split-Hopkinson pressure bar experiments. Figure 16 shows the deformation gradient that is used to drive the simulation. The deformation gradient component F_{xx} is decreased from 1.0 to 0.6 during the compression step and then increased during unloading until a small amount of tensile deformation is achieved. The components $F_{yy} = F_{zz}$ are kept constant at 1.0 while the shear components F_{xy} etc. are fixed at zero.

Bulk moduli for this deformation history are shown in Figure 17. As observed earlier, the Table model predicts the smallest values of bulk modulus. The SVR model produces the smoothest evolution of moduli while the MLP model exhibits the largest fluctuations.

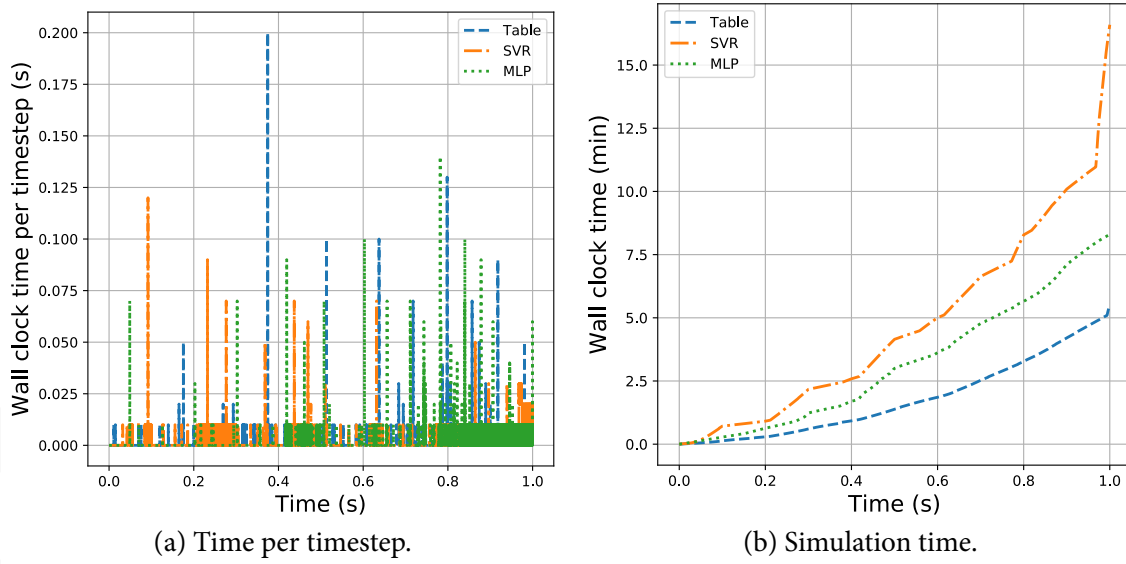


Figure 15 – Time per timestep and total simulation runtime for hydrostatic loading and unloading of a single particle.




Figure 16 – Input deformation gradient history for uniaxial strain loading and unloading of a single particle.

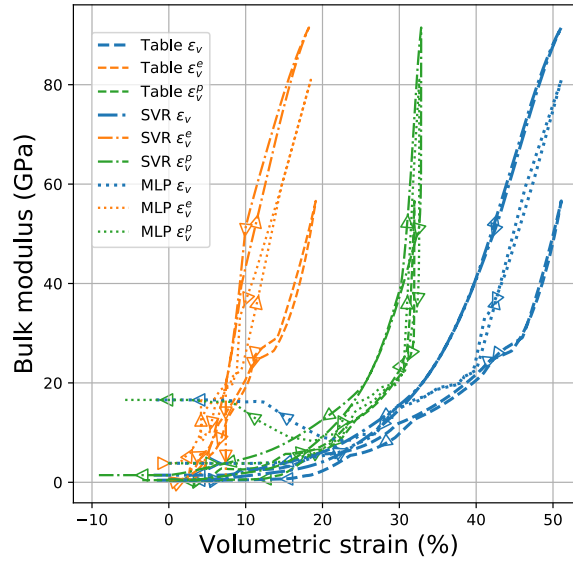


Figure 17 – Evolution of bulk modulus for uniaxial strain loading and unloading of a single particle.

The evolution of the predicted stress components σ_{xx} , σ_{yy} , and σ_{xy} during the course of the deformation are shown in Figure 18(a). We observe that the σ_{xx} values produced by the Table model are the smallest at the peak compression state. The SVR and MLP models produce approximately the same stresses at maximum compression. The unloading rate is slower for the Table model until around 1.25s, after which all three models unload at approximately the same rate. Transverse stress σ_{yy} and σ_{zz} are equal and typically lower than σ_{xx} during loading. However, the transverse stresses are higher than σ_{xx} during the unloading stage. Stress invariants p and q are plotted in Figure 18(b). The rate of increase of p during loading is larger than the rate of decrease during unloading. A similar trend is observed for the value of q (which has been given a sign based on the value of J_3 , the third invariant of the deviatoric stress). The sign of J_3 is negative during compressive loading and changes to positive during unloading.

Elastic and plastic strain components are plotted in Figure 19. The largest values of ϵ_{xx}^e are produced by the Table model, while the MLP model predicts the smallest values because of the large bulk modulus. During unloading, the elastic strain drops before setting to a nearly constant value. The corresponding values of ϵ_{yy}^e are shown with thinner lines and reach tensile states during unloading. The plastic strain component ϵ_{xx}^p is almost identical for the three models. However, dilation is observed in the ϵ_{yy}^p strain during the loading phase.

Figures 20(a) and (b) plot the stress components σ_{xx} and $\sigma_{yy} = \sigma_{zz}$ as functions of the axial strain ϵ_{xx} , respectively. As expected from the hydrostatic compression tests, the Table model produces the smallest peak compressive stresses, followed by the MLP model. The largest stresses are predicted by the SVR model.

The corresponding volumetric stress-strain plots are shown in Figure 21. Since the evolution of the cap is identical for the Table, SVR, and MLP models, the difference in predicted pressures is controlled by the stiffness of the volumetric elastic response (which also depends on the plastic strain).

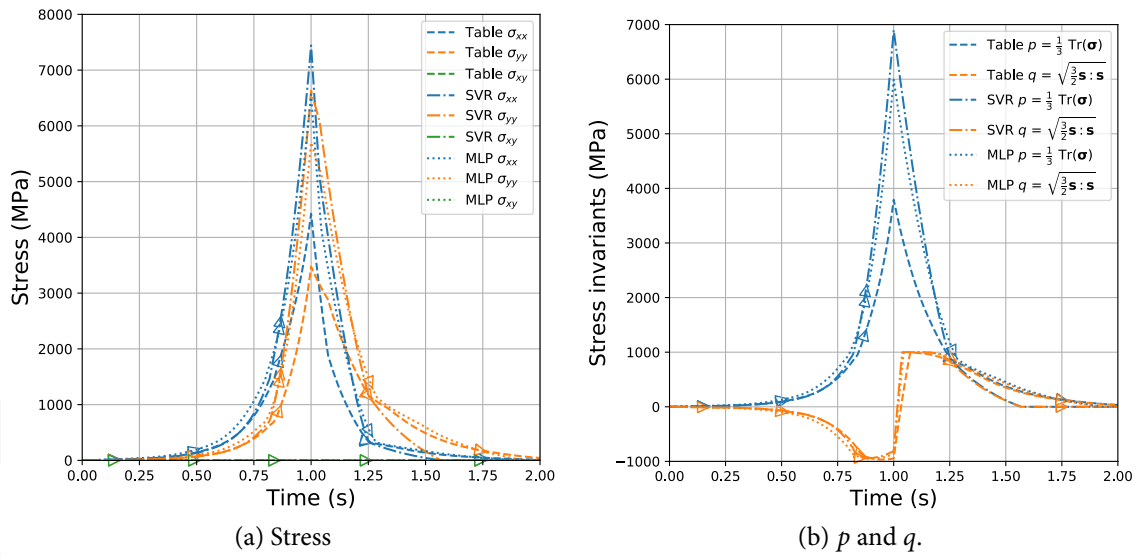


Figure 18 – Comparison of computed stress history for uniaxial strain loading and unloading of a single particle.

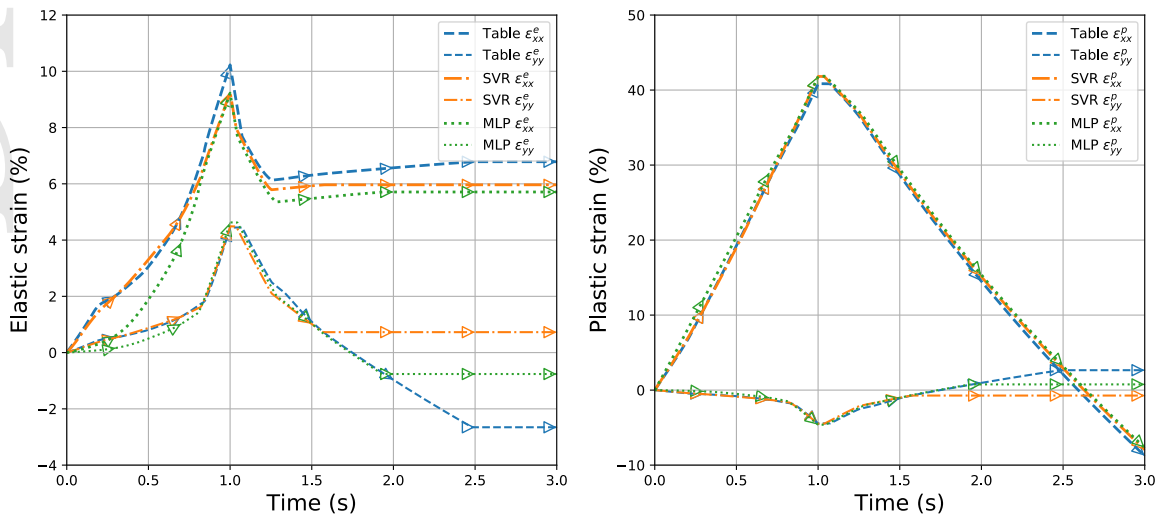


Figure 19 – Comparison of computed elastic and plastic strain history for uniaxial strain loading and unloading of a single particle.

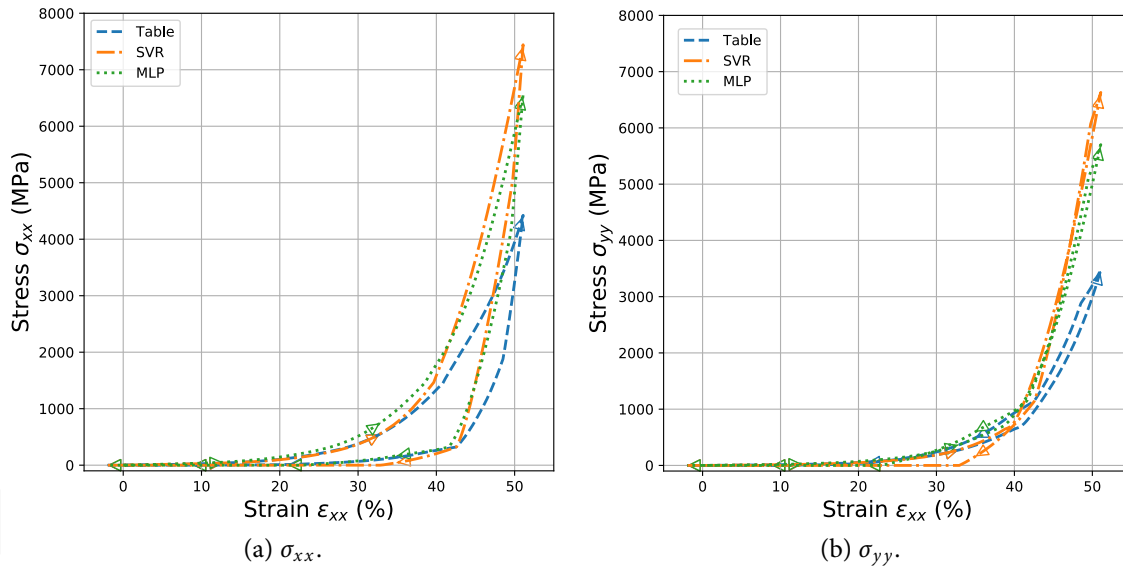


Figure 20 – Comparison of computed stress-strain history for uniaxial strain loading and unloading of a single particle.

The Table model predicts the least stiff response, followed by the MLP model. The SVR model predicts the stiffest response.

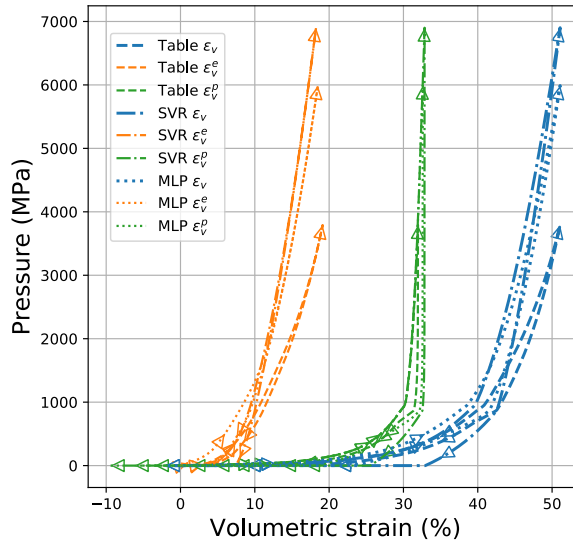


Figure 21 – Comparison of computed pressure-volumetric strain history for uniaxial strain loading and unloading of a single particle.

Plots of the yield surface and its evolution are shown in Figure 22. Though the yield surface cap location predicted by the three models is nearly identical at small compressive stresses, the locations differ by nearly 50% at peak compression. As observed earlier, the SVR model predicts the largest plastic volumetric strains and therefore the largest value of the compression cap location. The stress

state is also observed to push further into the cap region as the compressive stress increases, indicating that the stress state tends towards hydrostatic compression. Unloading is largely along the yield surface after an initial elastic unloading phase.




Figure 22 – Evolution of the yield surface for uniaxial strain loading and unloading of a single particle.

Figure 23(a) plots the evolution of the size of the timestep while memory usage during the course of the simulation is shown in Figure 23(b). The timestep decreases as the compression proceeds followed by an increase during the unloading phase. The tensile loading phase requires approximately the same time per timestep because the elastic moduli do not evolve and the yield surface remains fixed. The memory usage increases linearly after a constant phase, indicating a potential memory leak, probably caused by the XML or HDF5 libraries used in Vaango.

The time per timestep is shown in Figure 24(a) and Figure 24(b) shows the accumulated time taken to complete each simulation. As for the hydrostatic compression tests, the Table model takes the least time to compute, followed by the MLP model during most of the course of the simulation. However, toward the end of the simulation the time taken by the MLP model overtakes that taken by the SVR model.

3.3 Multiaxial strain loading and unloading

It is also informative to examine the response of the models to a bi-directional loading-unloading process as shown in the deformation gradient history in Figure 25. The shear components F_{xy} , etc. the deformation gradient are fixed at zero. The normal components F_{xx} , F_{yy} , and F_{zz} are identical compressive values until 1s after which F_{xx} continues in compression while $F_{yy} = F_{zz}$ are unloaded. At 2s, the F_{xx} component of the deformation gradient is unloaded until tensile strains are achieved, while the $F_{yy} = F_{zz}$ components are loaded back to compressive states until 3s. At that point the deformations are again reversed as shown in the figure.

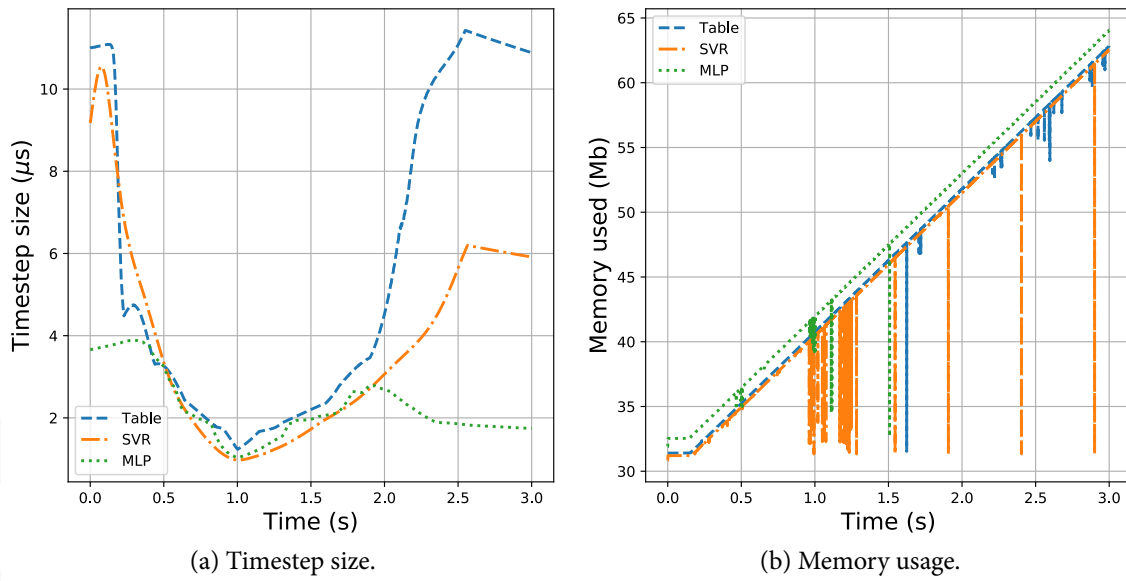


Figure 23 – Timestep size evolution and memory usage during uniaxial strain loading and unloading of a single particle.

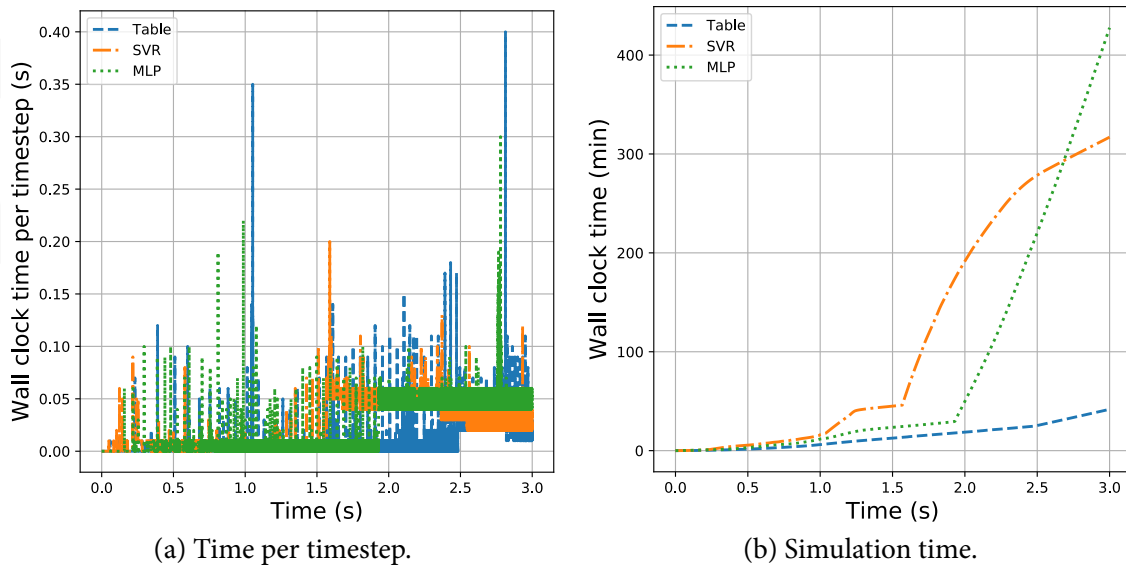


Figure 24 – Time per timestep and total simulation runtime for uniaxial strain loading and unloading of a single particle.




Figure 25 – Input deformation gradient history for multiaxial strain loading and unloading of a single particle.

Since the amount of compression is relatively small, the bulk modulus remains approximately constant for this deformation history as shown in Figure 26. A logarithmic scale has been used in the ordinate to highlight the fact that the MLP model produces the largest values of initial bulk modulus, almost an order of magnitude larger than those predicted by the Table and SVR models.

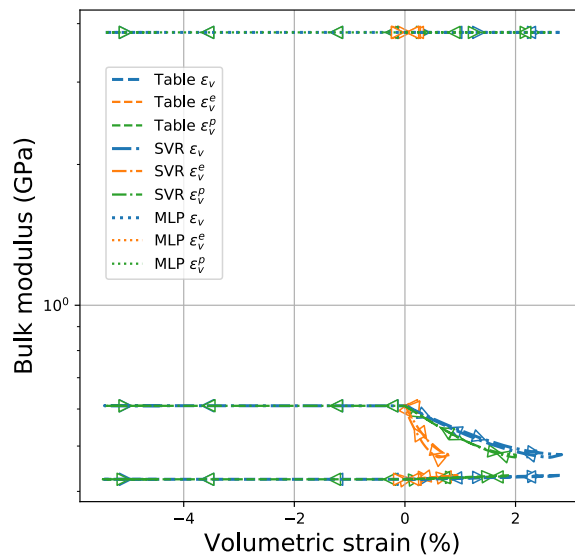


Figure 26 – Evolution of bulk modulus for multiaxial strain loading and unloading of a single particle.

The stress histories shown in Figure 27 reflect the fact that the bulk moduli are roughly constant for this simulation. The trends of the stresses are similar to those observed for the uniaxial strain

and hydrostatic loading simulations discussed earlier. We also note that the difference in predicted stresses between the three models is marginal for small deformations.

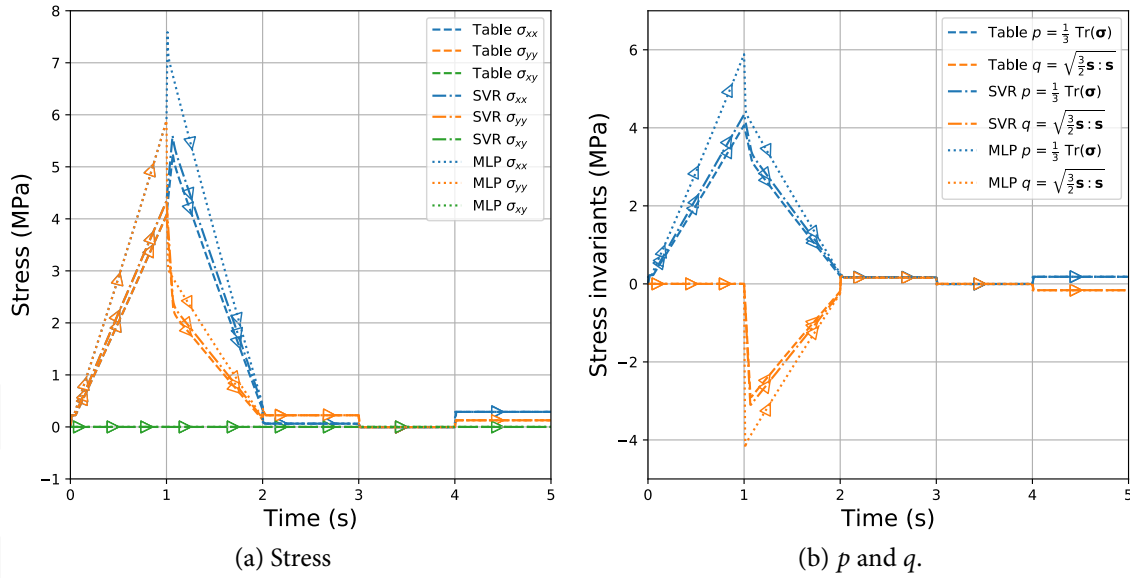


Figure 27 – Comparison of computed stress history for multiaxial strain loading and unloading of a single particle.

Elastic and plastic strain histories plotted in Figure 28 also indicate that the three models predict almost identical elastic and plastic strains. In particular, the plastic strain components attain both positive and negative values as the deformation continues.

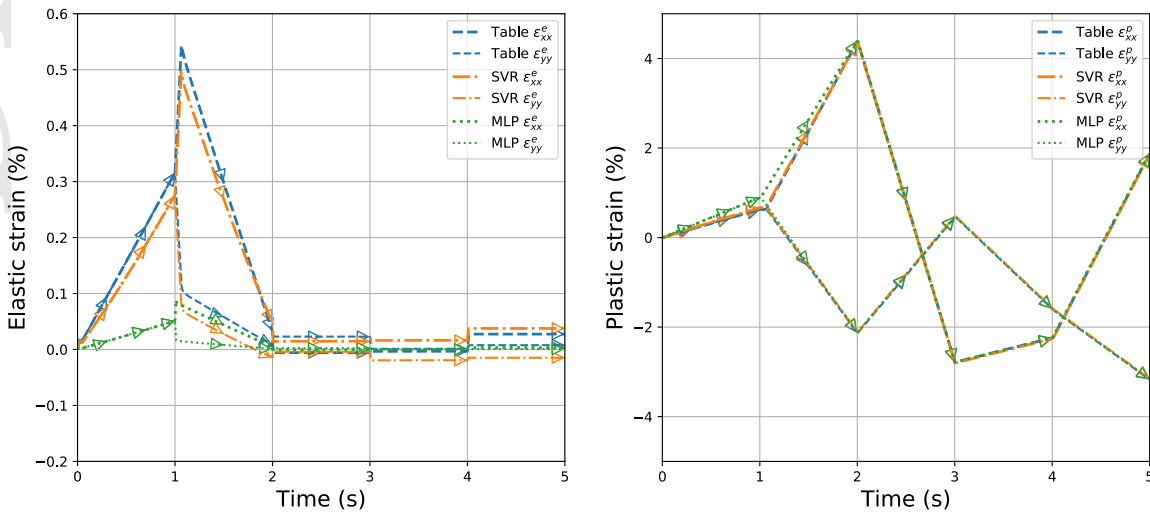


Figure 28 – Comparison of computed elastic and plastic strain history for multiaxial strain loading and unloading of a single particle.

If we consider the stress-strain response shown in Figure 29, we see that the MLP model produces

that largest stresses. The unloading paths are nearly identical for the three models and the cycling around tensile stress states produces reasonable results.

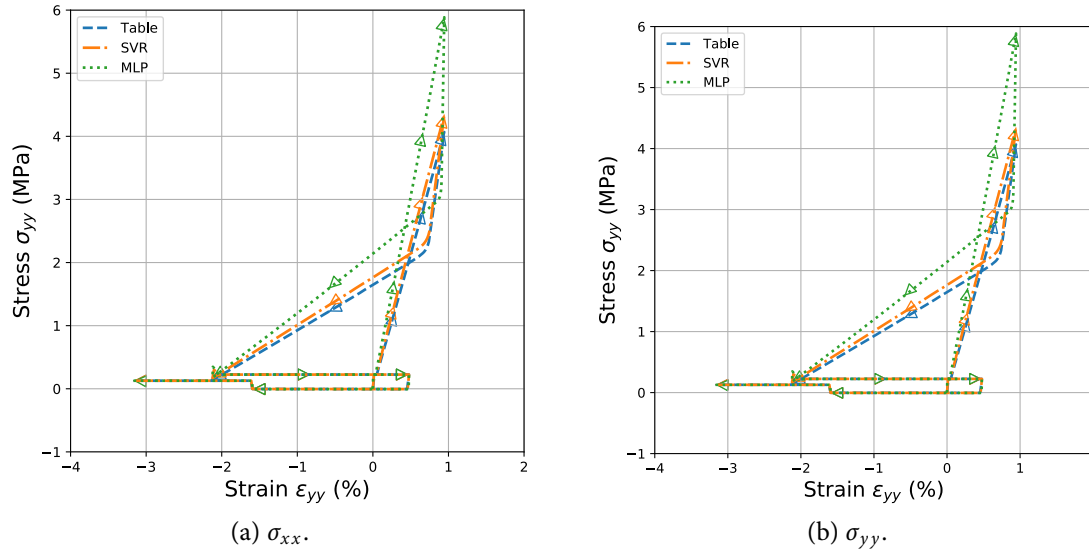


Figure 29 – Comparison of computed stress-strain history for multiaxial strain loading and unloading of a single particle.

The order-of-magnitude difference in the initial bulk modulus predicted by the MLP model is reflected in the volumetric stress-strain response plotted in Figure 30, especially when we closely examine the unloading curves for the three models. Though this difference is small compared to the magnitudes of the bulk moduli at large compressive strains, it is indicative of the need for better fits to the experimental data in the small-strain regime when that in the regime of interest.

Figure 31 shows the evolution of the yield surface and the stress paths taken by the three models. The MLP model produces the largest plastic strains and hence the largest elastic domain. The Table and SVR models produce approximately the same plastic volumetric strains. During unloading, the stress reaches a high curvature region of the yield surface where the iterative approach described in Banerjee, Fox, and Regueiro (2020d) converges slowly. The stress follows the yield surface until it reaches the tensile peak and cycles around that region during the last phase of deformation.

Figures 32(a) and (b) show the evolution timestep size and memory usage for the simulation, respectively. The timestep size remains nearly constant for the course of the simulation, with the Table model using the largest values because of the lower bulk modulus. The memory usage, once again, shows an increase with time even though the yield surface does not evolve much - indicating a potential memory leak.

The time per timestep and total simulation time are presented in Figures 24(a) and (b), respectively. All three models take approximately the same time per timestep. This is further demonstrated in the plots of the total simulation time which is almost identical for the Table and SVR models and slightly larger for the MLP model.




Figure 30 – Comparison of computed pressure-volumetric strain history for multiaxial strain loading and unloading of a single particle.




Figure 31 – Evolution of the yield surface for multiaxial strain loading and unloading of a single particle.

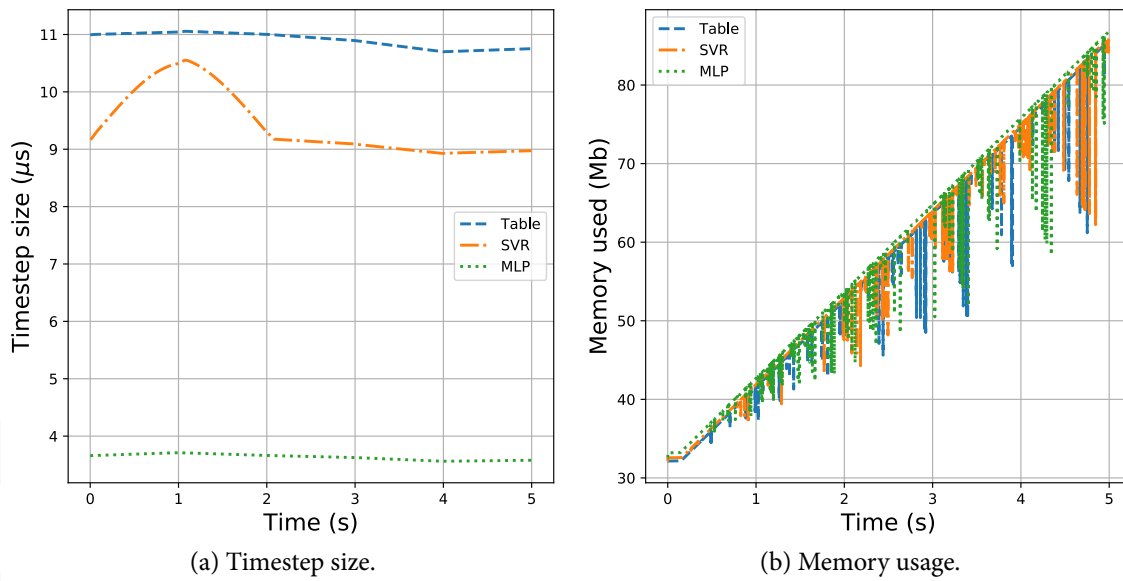


Figure 32 – Timestep size evolution and memory usage during multi-axial strain loading and unloading of a single particle.

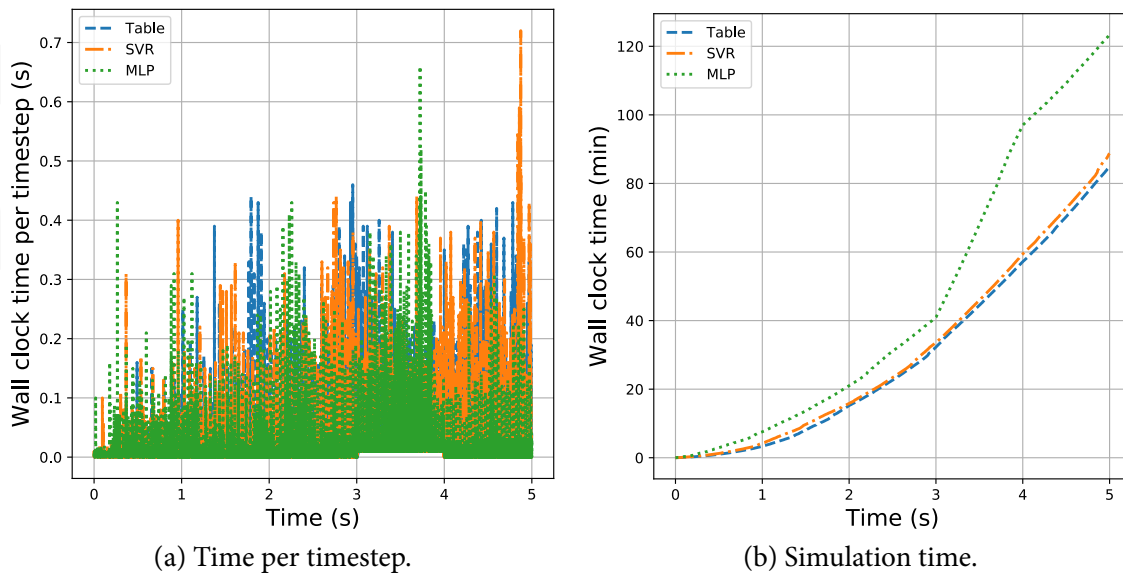


Figure 33 – Time per timestep and total simulation runtime for multi-axial strain loading and unloading of a single particle.

4 Concluding remarks

In this paper we have exercised the algorithm described in (Banerjee, Fox, and Regueiro, 2020d) on a dry, poorly-graded, concrete sand with three different models for the elastic moduli. The Table model used linear interpolation of bulk modulus vs. pressure data assuming that the modulus remained constant beyond the range of available data. The SVR model used analytical derivatives of a support vector regression fit to experimental pressure vs. volumetric strain data. The MLP model used a three layer neural network fit to bulk modulus vs. elastic volumetric strain.

Hydrostatic loading/unloading, uniaxial strain loading/unloading, and multiaxial strain loading/unloading simulations were used to compare the response of the three elastic modulus models. We observe that the SVR and MLP models predict larger bulk moduli and, therefore, a stiffer response of the material with larger overall peak compressive stresses. The softer response of the Table model leads to lower simulation times because the yield surface does not expand as much as for the other two models.

Acknowledgements

This research has been partially funded by the US Office of Naval Research PTE Federal award number N00014-17-1-2704.

References

- Abadi, Martín et al. (2016). “Tensorflow: A system for large-scale machine learning”. In: *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pp. 265–283 (cit. on p. 5).
- Banerjee, B. (2014). *Vaango User Manual: parallel simulation of high strain-rate solid mechanics with the Material Point Method*. <https://github.com/bbanerjee/ParSim/tree/master/Vaango> (cit. on pp. 1, 4).
- Banerjee, B., D. M. Fox, and R. A. Regueiro (2020a). *Interpolating tabular data for granular material models*. Tech. rep. PAR-10021867-092020-3. Parresia Research Limited (cit. on pp. 1, 4).
- (2020b). *Multilayer perceptron neural networks as multi-variable material models*. Tech. rep. PAR-10021867-092020-2. Parresia Research Limited (cit. on pp. 1, 4, 5).
- (2020c). *Support vector regression for fitting multi-variable material models*. Tech. rep. PAR-10021867-092020-1. Parresia Research Limited (cit. on pp. 1, 2, 4).
- (2020d). *Tabular models for high pressure and high strain-rate granular plasticity*. Tech. rep. PAR-10021867-092020-4. Parresia Research Limited (cit. on pp. 1, 4, 6, 20, 23).
- Brannon, R. M. et al. (2015). “KAYENTA: Theory and User’s Guide”. In: *Sandia National Laboratories report SAND2015-0803* (cit. on p. 2).
- Crockford, Douglas (2006). “JSON: The fat-free alternative to XML”. In: <http://www.json.org/xml.html> (cit. on p. 4).
- Fox, D. M. et al. (2014). “The effects of air filled voids and water content on the momentum transferred from a shallow buried explosive to a rigid target”. In: *International Journal of Impact Engineering* 69, pp. 182–193 (cit. on p. 2).

Koranne, Sandeep (2011). “Hierarchical data format 5: HDF5”. In: *Handbook of Open Source Tools*. Springer, pp. 191–200 (cit. on p. 5).

DRAFT