Granular Matter (2023) 25:2
https://doi.org/10.1007/s10035-022-01282-y

ORIGINAL PAPER

What is shape? Characterizing particle morphology with genetic algorithms and deep generative models

R. Buarque de Macedo[1] · S. Monfared[1] · K. Karapiperis[2] · J. E. Andrade[1]

Received: 8 December 2021 / Accepted: 5 September 2022

Abstract
Engineered granular materials have gained considerable interest in recent years. For this substance, the primary design variable is grain shape. Optimizing grain form to achieve a macroscopic property is difficult due to the infinite-dimensional function space particle shape inhabits. Nonetheless, by parameterizing morphology the dimension of the problem can be reduced. In this work, we study the effects of both intuitive and machine-picked shape descriptors on granular material properties. First, we investigate the effect of classical shape descriptors (roundness, convexity, and aspect ratio) on packing fraction ϕ and coordination number Z. We use a genetic algorithm to generate a uniform sampling of shapes across these three shape parameters. The shapes are then simulated in the level set discrete element method. We discover that both ϕ and Z decrease with decreasing convexity, and Z increases with decreasing aspect ratio across the large sampling of morphologies—including among highly non-convex grains not commonly found in nature. Further, we find that subtle changes in mesoscopic properties can be attributed to a continuum of geometric phenomena, including tessellation, hexagonal packing, nematic order and arching. Nonetheless, such descriptors alone can not entirely describe a shape. Thus, we find a set of 20 descriptors which uniquely define a morphology via deep generative models. We show how two of these machine-derived parameters affect ϕ and Z. This methodology can be leveraged for topology optimization of granular materials, with applications ranging from robotic grippers to materials with tunable mechanical properties.

Keywords Granular materials · Non-convex · Topology optimization · Deep generative models · Discrete element method · LS-DEM

1 Introduction

Topology optimization—the automated design of a material's structure to engineer a desired property—has demonstrated promise across multiple fields [49]. Due to exponential growth in computing power and algorithmic developments, computers have 'invented' functional structures, including: satellite brackets, cantilevers and lattice materials [7, 39, 49]. Nonetheless, topology optimization is inherently challenging due to the high dimension of the design space. As such, it is beneficial to parameterize the solutions with as few variables as possible, effectively reducing the dimension of the design space [28]. This process is prone to loss of critical information and it requires a deep understanding of the optimization problem. To overcome this, recent methods have leveraged deep learning to automatically discover a reduced number of parameters needed to uniquely define a class of solutions. For instance, Wang et al. utilized a deep learning architecture known as a variational autoencoder (VAE) to construct a finite dimensional vector space, or 'latent space' of microstructures [44]. Each vector in the latent space within certain bounds defined a unique and valid material microstructure. Further, the dimension of the latent space was significantly lower than that of the original binary image of the microstructure. Such a vector space is invaluable to topology optimization, as optimization algorithms can be easily run in the continuous, complete and relatively low dimensional latent space. In this

This article is part of a topical collection:Physics-informed artificial intelligence for granular matter.

✉ J. E. Andrade
jandrade@caltech.edu

[1] Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA

[2] Department of Mechanical and Process Engineering, ETH Zurich, Zurich, Switzerland

Published online: 04 November 2022

work, we apply recent developments in topology optimization to granular materials.

Granular materials—any collection of discrete-solid objects—are ubiquitous in both nature and industry [35]. Soils, ball bearings, and even asteroids can be classified as granular, and their behavior is approximated by rigid body dynamics [27, 30]. Thus, certain attributes persist across granular materials on multiple scales. One such phenomenon is jamming: when a collection of particles transition from a fluid-like flowing state to a solid-like locked state [4]. Recently, engineers and artists have attempted to design granular materials which take advantage of this unique jamming property for achieving desired functionality. Examples include self-supporting structures for housing, fabric which can transition from soft to stiff with pressurization and robotic grippers. [6, 26, 45]. A primary challenge in designing such substances is picking the shapes of the individual grains. Spherical particles tend to pack tightly, but do not have tensile strength when uncompressed. On the other-hand non-convex grains may entangle with each other and provide tensile strength, but pack loosely [10]. Consequently, finding the optimal morphology for a set of design specifications remains an open, yet crucial, problem.

Discovering the ideal grain structure for optimizing a macroscopic property such as tensile strength is a daunting task because of the infinite number of shapes one may consider. One method of exploring the high dimensional phase space of shapes is with genetic algorithms (GAs) [19]. Jaeger et al. [20, 32, 33] utilized GAs to evolve a particle morphology in order to maximize packing fraction in discrete element method (DEM) simulations [9]. Nonetheless, the technique was limited to clumps of spheres due to the limitations of traditional DEM. On the other hand, Makse et al developed a formula for estimating packing fraction and coordination number for arbitrary particle shapes by constructing the Voronoi volumes from sphere clusters [3]. Such an equation could be invaluable for use in topology optimization, as unlike simulations it can be quickly evaluated. Nonetheless, it assumes maximum packing density, and may not be accurate when meta-stable states such as arches occur in samples.

A way to reduce the computational requirement of automated grain design is by finding a function mapping from grain shape to material-scale behavior that can be quickly evaluated. One can then invert this 'forward' mapping from morphology-to-behavior to obtain an 'inverse' mapping from behavior-to-morphology. Creating this 'inverse' mapping is highly non-trivial and non-unique but of great scientific interest. In the last few years, such a mapping has been investigated for a multitude of shapes and behaviors using experiments and simulations, as analytical solutions for even simple shapes remain elusive. Commonly explored shapes in jammed particle simulations include spheres, cylinders, superballs, staples, ellipses and ellipsoids, sphere clusters, crosses, spherocylinders, tetrahedra, frustums, platonic solids, and realistically shaped grains—though for the latter there are few computational efforts due to technical limitations [1, 2, 10, 15, 17, 18, 21, 29, 31, 40, 47]. In all such cases, material properties continuously vary with particle shape.

The above studies suggests that a continuous function exists mapping particle shape to characteristic mesoscopic properties like average coordination number and packing fraction. However, each study only considers a small 'slice' of the entire function's domain. In this work, we demonstrate a methodology for constructing the general mapping from shape to material properties by parameterizing particle shape, and finding relationships between the chosen parameters and granular material properties. Tackling this herculean task is now possible due to recent technological advancements. These being: (1) the level-set discrete element method (LS-DEM), a DEM which can efficiently simulate particles of arbitrary shape [25], (2) granular cloning, which allows the generation of particle shapes with user-specified properties [11], and (3) neural-network based generative models, which can dramatically reduce the dimension of high-dimensional phase spaces by leveraging non-linear patterns in the data [12].

We first explore human chosen parameterizations of particle shape, and how these parameters affect packing fraction ϕ (ratio of solid to total area in a granular material) and coordination number Z (average number of grain-grain contacts per particle). In particular, we see how grain roundness, aspect ratio and convexity affect material properties. As these three parameters can not uniquely define a shape, we utilize a VAE to develop a 20-dimensional granular particle latent space. In such a space, each vector represents a unique particle shape. Finally, we demonstrate how such a latent space can be utilized for shape optimization, see Fig. 1. By attempting to define the shape of a particle with the minimum number of parameters, we ask: what is necessary to define a morphology? In other words, what is shape?

This paper is organized as follows. First, the methods of generating unique particles shapes is detailed in 2. Next, the simulation engine and methods are explained in 3. The results are then presented and analyzed in 4, followed by the conclusion and suggestions for future work in 5.

2 Particle generation methods

2.1 Generation with genetic algorithms

To start, we choose to continuously vary the convexity C, roundness R and aspect ratio A of the particles as it is well documented that such properties have a significant effect

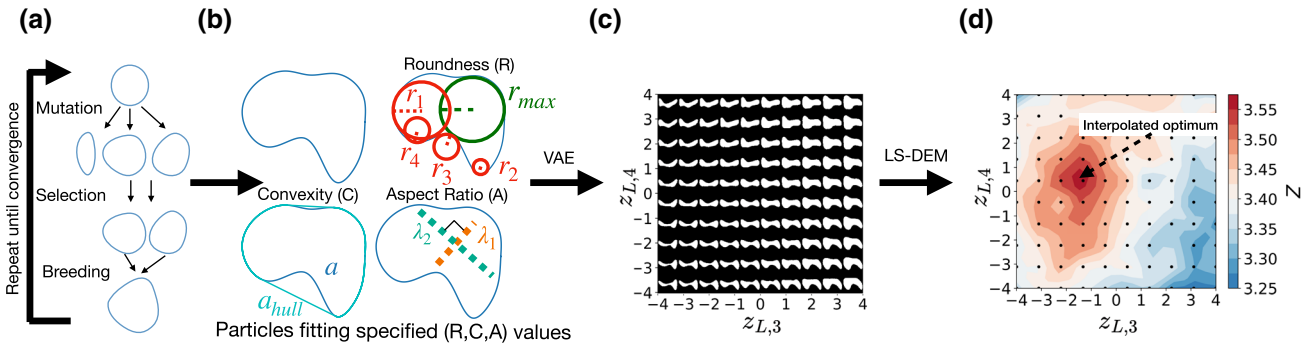

Fig. 1 Overview of topology optimization framework. **a** Obtain diverse collection of particle shapes from GAs which match specified (R, C, A) values. **b** Parameters for calculating (R, C, A). **c** Create continuous vector space of shapes with VAE. $z_{L,i}$ denotes the ith latent space vector, as described in Sect. 2.2. Sampling of shapes shown at evenly spaced vectors along a plane in latent space, with axes $z_{L,3}$ and $z_{L,4}$. The plane is defined by S (Eq. 6). **d** Run simulation for each shape at grid points (black dots) in latent space and obtain material property (Z). Use interpolation or optimization to find optimal particle shapes between grid points

on material behavior [21, 46]. Thus, each shape generated can be represented as a point in this three-dimensional parameter space, with coordinates given by (R, C, A). We investigate whether these dimensions are sufficient for predicting the properties at the mesoscale, or if more dimensions are required. These shape descriptors are defined for a given particle as follows:

$$R = \frac{\sum_{i=1}^{N} \frac{r_i}{N}}{r_{max}} \tag{1}$$

$$C = \frac{a_{hull}}{a} \tag{2}$$

$$A = \frac{\lambda_1}{\lambda_2} \tag{3}$$

where N is the number of corners on a particle, r_i is the radius of curvature r of the ith corner, r_{max} is the radius of the largest circle that can fit entirely within the particle, a and a_{hull} are, respectively, the area of the particle and the area of the convex hull of the shape and $\lambda_1 \leq \lambda_2$ are the principal component magnitudes of the points on the surface of the particle, see Fig. 1. Note all 3 parameters take values between 0 and 1, with 1 corresponding to the values for a disk.

In order to investigate how ϕ and Z change with these three dimensions, it is necessary to generate a wide range of particle shapes covering the (R, C, A) space. Generating new particle shapes autonomously, or 'granular cloning', is an active area of research [22, 42, 48]. Most methods require an external repository of particle shapes which are used as a blueprint for new shapes. However, we wish to generate particles from shape descriptors alone. The method proposed in [11] allows one to do this. Here, a GA morphs a particle's shape until the grain matches the shape descriptors to a certain tolerance. Clearly, multiple realizations of particle shape may correspond to the same (R, C, A) value. Hence we generate 10 particle realizations for each point on a grid in (R, C, A) space, evenly spaced by 0.1, with $R \in [0.2, 1.0]$, $C \in [0.7, 1.0]$ and $A \in [0.2, 1.0]$. We restrict our attention to $C \geq 0.7$ which still leads to pronounced nonconvexity (see Fig. 3) while ensuring that the genetic algorithm produces physically valid morphologies.

The engine used for the GE is python DEAP [14]. To summarize, the method deforms an ellipse to match a target set of morphological properties. In this case, the properties are a specified roundness, aspect ratio and convexity. The method begins by placing 8 equally spaced points along the perimeter of an ellipse, with the ellipse having the specified aspect ratio and area equal to 300. Then, a genetic algorithm perturbs the points until the morphological properties of the particle are within a tolerance of the specified properties. By using 8 points uniformly distributed across the surface to build the particle, we have chosen a 'scale' for the morphology of about particle diameter d/10. This is a common length scale for quantifying particle roundness [8].

The genetic algorithm consists of 3 steps: mutation, combination and selection. The algorithm begins with 50 equivalent ellipses, or 'individuals'. Random individuals are selected for mutation with probability 0.5. For individuals undergoing mutation, points on the particle boundary are selected with probability 0.5. In polar coordinates, the radial component r of the particle is moved in the radial direction by a random number sampled from the Gaussian $\Delta r \sim \mathcal{N}(0.0, 1.0)$ and change in polar angle $\Delta\theta \sim \mathcal{N}(0.0, 0.05)$.

Next, the 'cost' of each individual is calculated for selection. The cost of an individual is given by

$$\text{cost} = (C - C_{target})^2 + (R - R_{target})^2 + (A - A_{target})^2 + 100 \sum_{i=1}^{N} \delta\left(\frac{r_i}{r_{max}} < 0.08\right) + 100SI + 100BN \quad (4)$$

The subscript 'target' is the specified shape parameter value, and the squared difference is the squared error, a measure of the difference between the current individual and the target morphology. The fourth term in the cost function heavily penalizes the morphology for each corner with a normalized radius of curvature $\frac{r_i}{r_{max}}$ less than 0.08, thereby avoiding non-physical sharp edges. The *SI* term is equal to 1 if the particle self-intersects, heavily penalizing this non-physical behavior. The parameter is checked using the Bentley-Ottmann algorithm [5]. The *BN* term heavily penalizes 'bottlenecks' in the shape, i.e. when two points on opposite sides of the particles are squeezed close to one-another. It is equal to 1 when two opposite points are within a distance 7 of one-another. It is difficult to create consistent level sets from grains with bottlenecks. Minimizing this cost function produces a particle with the specified morphological parameters that is physical. In the selection step, pairs of individuals are randomly chosen, and the individual with a higher cost is eliminated. The remaining individuals are duplicated until the population size is back to 50.

Measuring particle roundness requires identifying corners on the particle at a relevant length scale. Given the points on the particle surface, the corners are identified as follows: First, a third-order spline is fit to the points. The spline is then smoothed, from which 500 ordered points on the smoothed particle surface are generated. The radius of curvature can be calculated at each of these points due to the smoothness of the spline. For each point, the radius of curvature of all points 20 ahead and behind are checked. If the radius of curvature of the current point is the minimum out of all the checked points, it is considered a corner. The radius of the maximum sized circle that can be fit in the particle is easily calculated from the level set: $r_{max} = |\min(\Phi)|$, see Fig. 2.

Finally, in the combination step pairs of individuals are randomly selected with probability 0.2. A new individual, or 'child' is created by randomly swapping boundary points on both grains.

This entire processes is repeated for 500 'generations', or until the minimum cost of an individual is below the tolerance $\epsilon = 0.0005$. After convergence, the individual with minimum cost is taken to be the solution. The location of the points are saved, and the level set for the particle is automatically generated. If generation 500 is reached and convergence has not occurred, the algorithm selects a different combination of target parameters. Note that all parameters can be modifieid to genereate different classes of shapes.

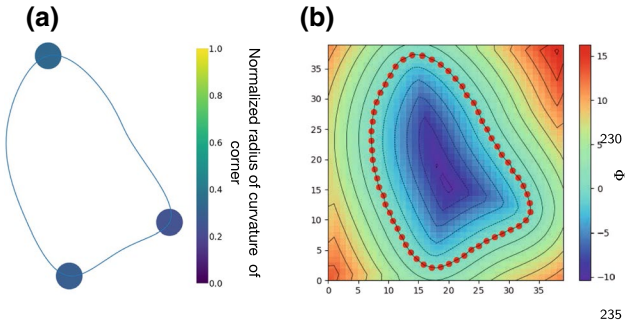

Fig. 2 **a** Corners detected by corner detection algorithm, with normalized radius of curvatures shown. **b** Points on the surface of a particle, with the grain's level set Φ given as a heatmap

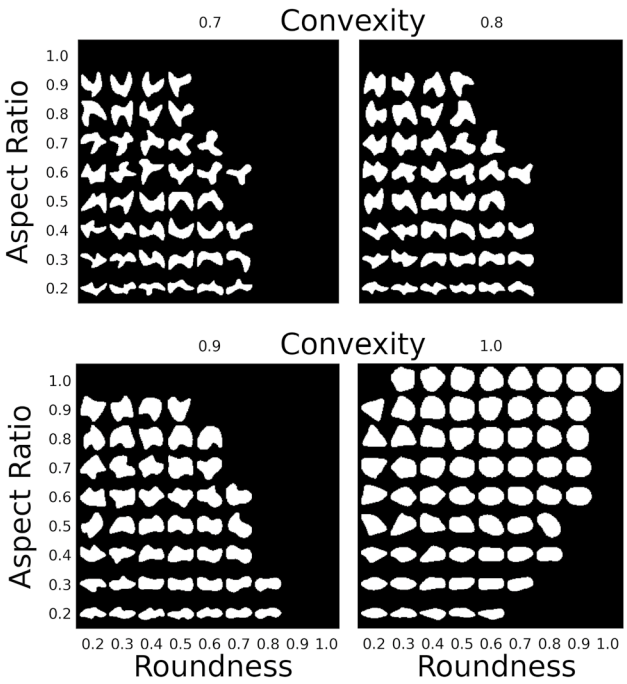

Fig. 3 Example of grains generated at grid points in *R*, *C*, *A* space. Empty boxes indicate the algorithm was unable to generate a particle with the prescribed (*R*, *C*, *A*) values. Each box represents a different convexity, with A as row and R as columns

For each realization of an (*R*, *C*, *A*) point, a simulation is performed to generate a packing and subsequently ϕ, Z are calculated. An example of a generated grain for each sampled point in (*R*, *C*, *A*) space is given in Fig. 3.

2.2 Generation with deep learning

While the classical, intuition based descriptors (*R*, *C*, *A*) can capture a significant amount of information about a particle's shape, they are not sufficient to fully quantify grain morphology. For instance, an infinite number of unique shapes could

have the same (R, C, A) coordinate. The limitations of these descriptors can be seen in the results for ϕ in Sect. 4 for high convexity and in Z for low convexity morphologies. Here, the dependence of ϕ on roundness and aspect ratio is complicated, implying there may be additional dimensions necessary to fully classify the grains. Consequently, we pose the question: what is the minimum number of dimensions necessary to sufficiently quantify 2D particle shape, in relation to its collective packing properties? Further, what are these dimensions?

Advances in machine learning offer tools for answering these questions. We utilize a VAE [36], which has recently been leveraged for generating unique and realistic granular particles [42]. The input to a VAE is an array of numbers, which in our case is the grey-values in a pixelated 64×64 2D image of a particle. The array is acted on by a series of matrix operations and non-linear functions, referred to as 'dense' layers. These operations include multiple convolutions, a commonly used operation in networks learning with image data. The transformed data is then fed into two arrays: μ and σ. Both of these vectors are the same length d, which is a parameter chosen by the user. This first half of the network is dubbed an 'encoder', as it compresses the data from the original image into d dimensions. Next, a 'sample' z is chosen by sampling from the distribution $z = \mu + \sigma \odot \epsilon$, where ϵ is a random standard normal variable sampled at evaluation time and $\odot$ is the element-wise dot product. The second half of the network is a series of convolutions that reconstructs the original image from z. This section of the network is called the 'decoder'. The loss function for the network is a measurement of the difference between the input image and the reconstructed image, and the similarity between the distribution of the training data in the latent space and a standard normal distribution. The parameters in the network are trained via back-propagation. Once trained, the network has learned how to compress the image into d dimensions such that the loss function is minimized.

The architecture of the encoder, which takes 64×64 black and white (single channel) images as input is given in Table 1, while the decoder architecture is given in Table 2 (Appendix). A general diagram of the architecture is displayed in Fig. 4.

The VAE parameters η and θ, representing the weights connecting neurons in the encoder and decoder network respectively, are tuned during training as to minimize the following loss function for each of the data points where x_i is the i^{th} training data point:

$$l_i = -\mathbb{E}_{\mathbf{z}\sim q_\theta(\mathbf{z}|x_i)}[\log(p_\eta(\mathbf{x}_i|\mathbf{z}))] + KL(q_\theta(\mathbf{z}|\mathbf{x}_i)||p(\mathbf{z})) \quad (5)$$

The first term is the 'reconstruction loss', and measures the difference between the input and the output image for data point x_i, while the second term (Kullback-Leibler divergence) encourages the distribution of data in the latent space $\mathbf{z}$ to be Gaussian. p and q represent probability distributions learned by the network. For details on all terms and derivation, see [12]. Code for variational autoencoder from [37].

We train a VAE on 10,000 unique images of particles generated by the genetic algorithm across a uniform sampling in (R, C, A) space. The grains are rotated and translated such that the principal axis of the grain is horizontal, and the grain centroid is in the middle of the image. The grain is converted into a black and white image, with the particle being white and the outside the particle black. The image is then subdivided into 62 pixels horizontally and vertically, with two empty pixels used to pad the sides of the image. These images are then used to train the network. We find that $d = 20$ is the lowest value of d necessary for successful reconstruction. Any lower, and the reconstruction loses too much structure when compared to the original image. Once trained, the encoder section of network can generate 2 unique d-dimensional vectors, μ and σ, for any given particle shape. Likewise, for any sample vector z the decoder can generate a unique particle shape. The d dimensional vector space of z is a latent space, where each vector defines a unique grain.

With the VAE, we achieve the goal of finding a complete set of dimensions for fully describing particle shape, such that the number of dimensions are minimized to avoid

Table 1 Architecture of encoder

Layer	Operation	Output dimension	Kernel size	Stride	padding	Activation function
1	2DConv	(64, 33, 33)	4	2	2	ReLU
2	2DConv	(128, 17, 17)	4	2	2	ReLU
3	2DConv	(256, 9, 9)	4	2	2	ReLU
4	2DConv	(512, 5, 5)	4	2	2	ReLU
6	2DConv	(1024, 3, 3)	4	2	2	ReLU
7	Avg. pool	(1024, 1)	N/A	N/A	N/A	None
8	Linear	(2048, 1)	N/A	N/A	N/A	None
9	Linear	(20, 1)	N/A	N/A	N/A	None

Input is a 64×64 black and white image (single channel) of the particle. Output is the μ vector and σ vector of latent space size d, 20. From here a sample z can be drawn

Table 2 Architecture of decoder

Layer	Operation	Output dimension	Kernel size	Stride	Activation function
1	linear	(1024,1)	N/A	N/A	N/A
2	2DConvTranspose	(512, 3, 3)	3	2	ReLU
3	2DConvTranspose	(256, 7, 7)	3	2	ReLU
4	2DConvTranspose	(128, 15, 15)	3	2	ReLU
5	2DConvTranspose	(64, 31, 31)	3	2	ReLU
5	2DConvTranspose	(64, 64)	4	2	Sigmoid

Input is sample z of size (d,1), drawn from normal distribution with mean μ and standard deviation σ, obtained from encoder. Output is a 64×64 reconstruction of the original image

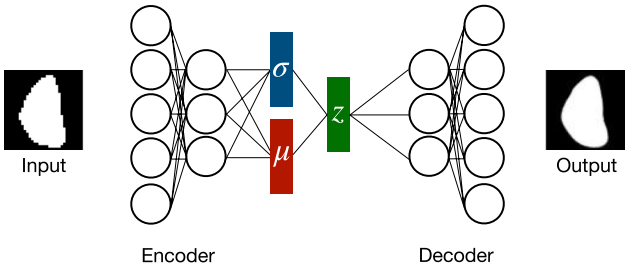

Fig. 4 Diagram of a VAE. On the left is the input, which is a pixelated image of a grain in this case. On the right is the reconstructed image that is output by the network. Note the network smooths the particle somewhat, representing a certain degree of lost information

redundancy. By discovering a function with this latent space as the domain, and mesoscale variables such as ϕ as the range, we would have a means of quickly predicting the mechanical properties of granular materials with arbitrary shaped grains.

We compute a z for each particle in the training data. The data is then scaled and rotated to reduce remaining correlations between the data points in the latent space using principal component analysis (PCA) [23]. The transformed z vectors are given by z_L, with ith components $z_{L,i}$. Each z_L lives in the d-dimensional latent space, with basis vectors $(b_1, ..., b_d)$. It is found that outside of the range $-4 \leq z_{L,i} \leq 4$, corresponding particles tend to be non-physical, i.e. containing disconnected regions. For computational tractability, only latent vectors in the somewhat arbitrary subspace S are simulated in this study, where:

$$S = \{(0, -4, z_{L,3}, z_{L,4}, 0, ..., 0) \in \mathbb{R}^{20}\} \tag{6}$$

However, the methods applied here are valid in any subspace of the latent space. S is chosen because of the high diversity of particle shapes defined by the space. A uniformly spaced 2D grid is defined in S with grid spacing 0.8 between bounds $-4 \leq z_{L,3}, z_{L,4} \leq 4$. A particle is generated from each z_L grid point for use in simulation. The particle shapes corresponding to each grid point are displayed in Fig. 1.

3 Simulation methodology

In order to measure ϕ, Z and any other mesoscopic quantity for a range of shapes, we leverage a specialized 2D DEM simulation. DEM was developed for disks that obey rigid body mechanics [9]. In DEM, disks are allowed to overlap by a very small amount. The contact force between two disks is a function of the overlap extent between two particles. Nonetheless, LS-DEM (level set DEM) allows DEM simulation of arbitrarily shaped grains, and thus is capable of simulating any physically possible shape created by the generative models.

LS-DEM [25] stores the location of a collection of points on the surface of the particle, in addition to the level set of each grain. The level set $\boldsymbol{\Phi}$ for a given particle surface is a function defined in all of space, such that its value at a given point is the signed distance to the surface of the grain. $\boldsymbol{\Phi}$ is positive outside the grain, and negative inside, pictured in Fig. 2. When the surface of one grain intersects another, the penetration extent is easily computed by looking up one particle's level-set value at the location of the other particle's surface points, and performing linear interpolation on the level-set value if necessary.

Each (R, C, A) simulation consists of 900 particles at a gas state, i.e. randomly distributed in an evenly spaced non-overlapping grid such that grains are not touching one-another with uniformly distributed random velocity and rotation. The horizontal and vertical grid spacing is 6 m, with 15 particles per row. The initial speed is a randomly chosen number between $-$ 20 and 20 ms^{-1} in a random direction, with 0 angular velocity. The initial conditions are chosen randomly in order to make sure the final sample is less dependent on initial conditions.

Each particle in the (R, C, A) simulation has the same area. Simulations for VAE generated grains contain 1,800 particles, as it is found for very non-convex particles generated by the VAE the representative volume element (RVE) requires more grains (see below). Particles generated from the VAE are scaled such that each pixel edge length in the 64×64 image is equal to 0.045 m. Particles are arranged

in a grid with the same spacing as in (R, C, A) simulations. Simulations parameters are given in Table 3.

The grains are left to fall under the influence of gravity until they reach a loosely packed jammed state. The simulation is run until the system is at rest. The packing fraction, in 2D, is calculated for the jammed state by first performing a Voronoi tessellation on the points on the particles' surface in order to calculate the area of the domain, then using the known area of the grains to calculate ϕ for a given RVE. Z is computed by measuring the average number of grains a given particle is contacting.

Calculation of material properties such as ϕ and Z requires defining a representative volume element (RVE), i.e. an area in the simulation domain over which these parameters can be found. The RVE size of a granular material requires calibration. If the RVE is too small, it is not representative of the mesoscopic properties. However, too large of an RVE will include the rigid wall boundaries, which will skew results. We take a circle of radius $D = 50$ m centered at the centroid of grain positions to be the RVE, from which Z and ϕ are calculated for each simulation. The value of $D = 50$ is determined by testing for convergence in mesoscale properties, see appendix Sect. 6.2.

Each simulation is repeated 5 times, with the same particles but different initial velocities, positions and orientations. The average ϕ, Z of these 5 simulations are recorded. Examples of simulations before and after pluviation are displayed in Fig. 5B.

4 Results and analysis

4.1 (R, C, A) simulations

ϕ and Z measurements are given in Fig. 5A. The data exhibits multiple patterns. The most obvious trend is that decreasing convexity decreases both packing fraction and coordination number. This is also evident from Fig. 5B and comparing Fig. 6A, D, G and J: as convexity decreases, the porosity of the material increases substantially. This is due to contact occurring at corners of highly non-convex grains leading to empty space within the grain's non-convex indent, or arches forming within the material around large voids (Fig. 6J) similar to cohesive granular materials [38].

For high convexity grains $C = 1$, ϕ gradually changes in the R, A plane, except for very low A, at which ϕ rapidly decreases. The high values of ϕ at $(R, C, A) = (1.0, 1.0, 1.0)$ and $(R, C, A) = (0.2, 1.0, 0.2)$ are due to the hexagonal-like (Fig. 6A) and tessellating (Fig. 6B) packings that occur with grain shapes near these locations, respectively.

Z is strongly affected by A for all $C >= 0.7$. Elongated grains exhibit local nematic order by tending to pack side-by-side (Fig. 6C, F, I and L). This packing is similar to what is observed with ellipses. So is the trend in Z: elongating the particles increases the average number of contacts. However, as convexity decreases, the elongated grains exhibit less orientational order and the trends in Z become more complex [18].

The patterns in ϕ in the $C = 0.9$, $C = 0.8$ and $C = 0.7$ plane are similar: maxima is achieved in the (R, A) plane at $(R, A) \approx (0.7, 0.4)$ (Fig. 6E, H) for $C = 0.8/0.9$, and $(R, A) \approx (0.6, 0.3)$ for $C = 0.7$ (Fig. 6K). Interlocking is present for both of these maxima, yet the prevalence decreases with decreasing convexity. Particles far from the maxima exhibit less interlocking and pack loosely, (Fig. 6D, G and J). As such, the sudden decrease in ϕ as R is decreased for low convexities is related to loss of interlocking.

In general, the (R, C, A) basis is effective at categorizing this diverse collection of shapes. The function mapping (R, C, A) to Z and ϕ appears smooth. Nonetheless, in the $C = 1$ plane the ϕ values vary more non-smoothly. Further, the space is full of 'holes' where a particle shape could not be found. These observations imply additional dimensions may be necessary to quantify shape in a continuous and

Table 3 Table of simulation parameters

Parameter	(R, C, A) sims	VAE sims
Box width lx (m)	108	64
Box height ly (m)	402	402
Particle area (m^2)	6	Variable
Density (kg/m^3)	2650	2650
Normal spring constant k_n (kg/s^2)	10^{11}	10^{11}
Tangential spring constant k_s (kg/s^2)	10^{11}	10^{11}
Coefficient of friction μ	0.5	0.5
Normal coefficient of restitution C^n_{res}	0.4	0.4
Shear coefficient of restitution C^s_{res}	0.5	0.5
Number of unique grain shapes per simulation	10	1
Total number of grains	900	1800

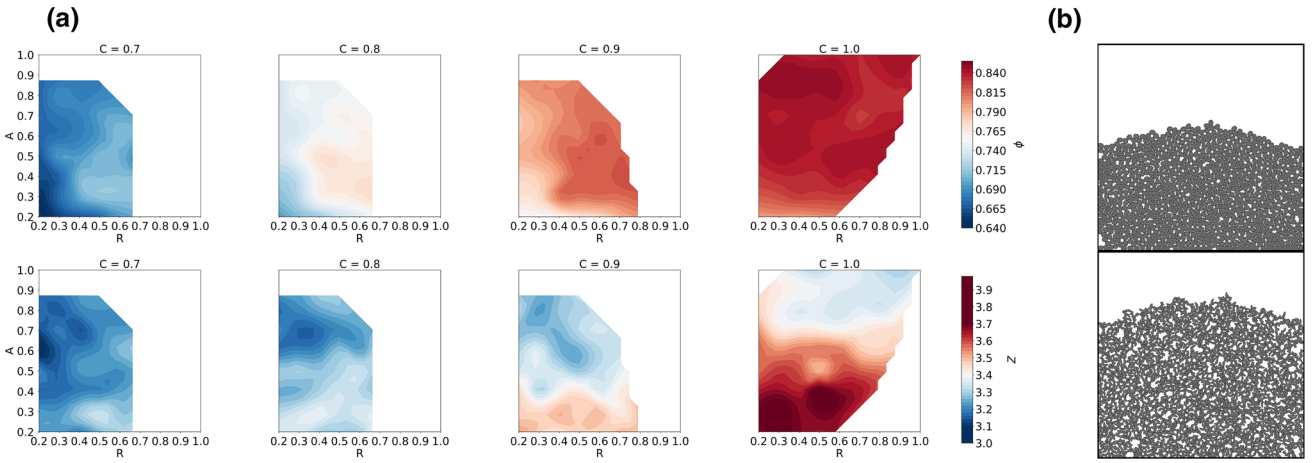

Fig. 5 **a** Average ϕ, Z in R, C, A space. Top row: ϕ values. Bottom row: Z values. Columns correspond to convexity, matching Fig. 3. White blocks are locations where there are too few points for interpolation. **b** final state of two R, C, A simulations at (0.7, 1.0, 0.7) (top) and (0.3, 0.7, 0.2) (bottom)

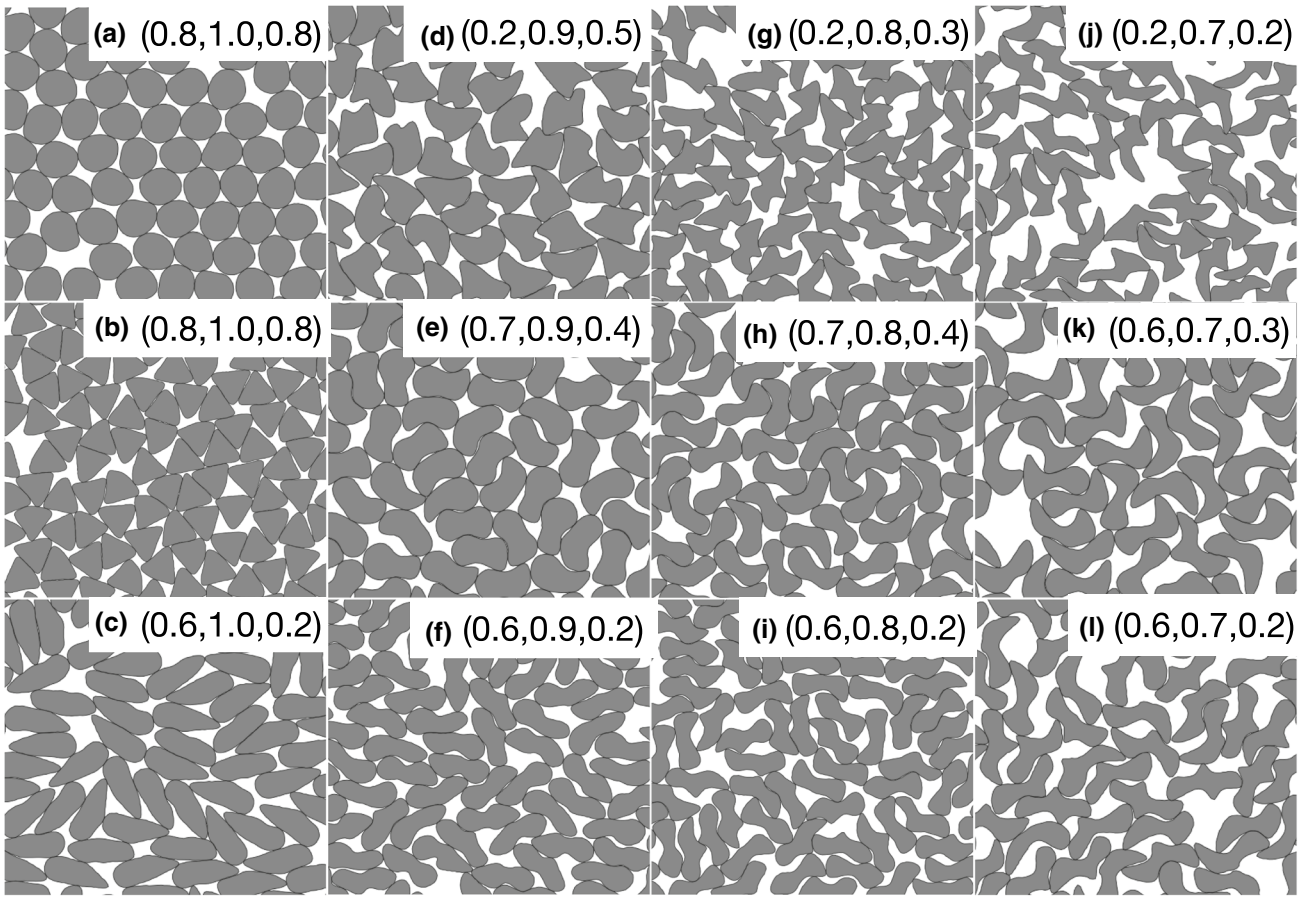

Fig. 6 Examples of pluviated packing configuration across (R, C, A) space. Corresponding (R, C, A) value for each packing given in top-right corner. Particles exhibit hexagonal packing (**a**), tessellation (**b**), nematic ordering (**c**, **f**), interlocking (**e**, **h**, **k**), arching (**j**) and propensity for corner contact (**d**, **g**, **j**)

complete sense. Exploring such additional coordinates is the subject of the next section.

4.2 VAE-generated particle simulations

For each particle on the uniformly spaced grid in S (see 2.) a pluviation simulation of 1800 identical copies of the particle is carried out, as detailed in 3. The results for ϕ and Z in the subspace S are given in Fig. 7. Unlike Fig. 5, the ϕ,Z function is defined for all of S, as a shape is guaranteed to exist for any vector—though the probability of a consistent shape decreases outside the $-4, 4$ bounds. In S, the particle with the highest packing fraction is predicted to be in the top right, or the (4, 4) location. This corresponds with a relatively convex shape, see inset. Meanwhile, the particle with the maximum coordination number is around the $(-1.5, 0.5)$, matching the elongated 'bar bell' shape in Fig. 1—a stretched, lower aspect ratio shape. Both optimums are consistent with observations in Sect. 4.1. As neither of these particles were in the training data for the VAE, the program has 'invented' an entirely new shape for optimizing a material attribute.

This research has focused on the 'forward' problem of granular design: mapping particle shape to *mesoscopic or* bulk properties. That being said, these tools can be applied towards the 'inverse' property-to-bulk problem. With an approximate forward mapping from latent space vector to bulk property in hand, one can search the latent space for areas of optimal bulk properties such as yield stress. Once an area in latent space has been identified, LS-DEM simulations can be run using particles sampled from this area to achieve a finer sample of bulk properties. This process can be repeated for other candidate areas of the latent space, thereby performing topology optimization (see Fig. 1). This is similar to the method in [16], which performs optimization within the lower dimensional latent space to reduce the complexity of the problem.

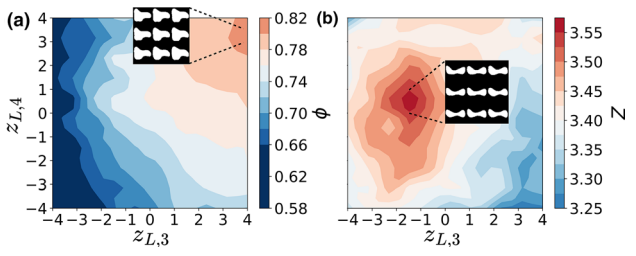

Fig. 7 Average ϕ (**a**) and Z (**b**) measured for uniform square grid in S space where z_L^i denotes the ith latent space vector (see Sect. 2.2) with interpolation between points. Particles close in latent space to maxima shown

5 Conclusions

In this work, we have investigated mappings from parameterized spaces of particle shape to mesoscale material quantities, Z and ϕ. We have shown how parameterizing morphology by roundness, sphericity and aspect ratio can successfully be used to predict material properties. In particular, we have found that ϕ and Z continuously vary throughout the shape space due to intuitive changes in packing geometry characterized by tessellation, local nematic ordering, hexagonal packing, arching and interlocking. We have also outlined a method for fully and automatically parameterizing particle shape via VAEs.

In future work, the techniques developed can be easily extended to 3 dimensions. Future work should also consider larger collections of particles, such to further minimize the effects of wall boundaries. Next, our results have only considered pluvation-prepared samples. However, mesoscale properties in granular materials are highly dependent on initial conditions [43]. Thus, it would be useful to understand how the ϕ and Z values corresponding to a particular shape change with preparation method. For instance, particles could be compressed or tapped to achieve higher packing fractions. Also, it would be valuable to see how results change with frictionless grains as qualitatively different packing configurations may be achieved. Further, the methodologies presented can be applied to any mesoscale or bulk material property—including tensile strength and critical state parameters. This mapping could also be used in coarse grained data-driven models, where model parameters are learned from lower-scale simulations [24]. Additionally, it could facilitate the 'engineering' of a confined fluid's thermodynamics [34]. Finally, our results suggest particle morphology can be fully characterized using a few number of variables. Exploring the meaning of the dimensions chosen by the generative model could shed light on the nature of shape.

This methodology contains certain limitations. The genetic algorithm will not converge for certain morphological parameters, leading to the 'gaps' in RCA space. Further, the genetic algorithm requires manual tuning, including mutation parameters and start conditions. The VAE overcomes many of these limitations as it can learn to generate particle shapes with minimal hand-tuning. Nonetheless, it is difficult to control the output of the autoencoder to only output particles of a certain class, such as with a constant convexity or roundness.

Previous research into granular materials was limited by the types of shapes one could simulate or obtain experimentally. Shape optimization of granular materials was limited to a small number of clustered of spheres, or relied upon theoretical results which may not always match DEM

simulations. With these results, we open research avenues for exploring high-dimensional spaces of arbitrary grain shapes is LS-DEM simulations.

6 Appendix

6.1 Frictionless ellipse validation

We perform simulations of frictionless ellipses with a range of aspect ratios to validate our findings against previous results, see Fig. 8. Results are generally very similar, with ϕ oscillating about 0.88 as the ellipse is elongated from a circle, followed by a decrease in ϕ that starts as A moves below 0.5. The insensitivity of the packing fraction to particle elongation in the aspect ratio range 0.5 to 1 in Fig. 8 differs from previous studies [13, 41] which observed a decrease in ϕ as an aspect ratio 1 is approached. However, these studies used bi-disperse particles to prevent crystallization, while in our simulations of identical, frictionless particles the tendency to crystallize keeps the packing fraction high. Indeed, we note that LS-DEM produces almost perfect hexagonal packing for frictionless circles (A = 1), with ϕ slightly below 0.9.

6.2 Determination of RVE size

To discover the minimum RVE size for a given complete simulation, circles of increasing diameter D centered at the average of all the particles' centroids are considered, where D_i corresponds to the ith circle of increasing diameter in steps of 1 m. Z_i is determined for the particles within each D_i diameter circle. A moving average of $\overline{Z_i}$ is computed by

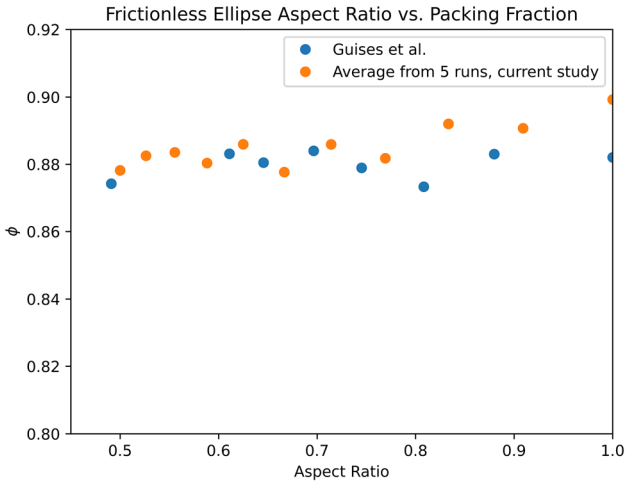

Fig. 8 Packing fraction of frictionless ellipses, comparing results from current study with Guises et al [18]

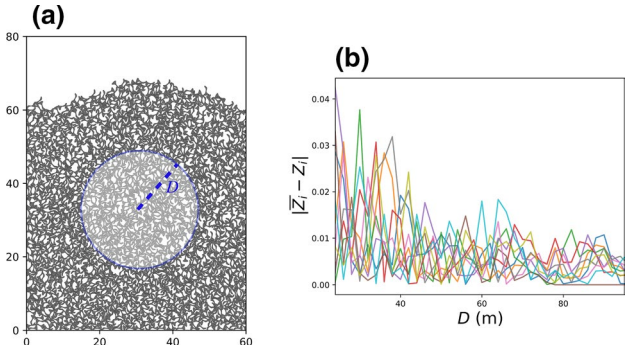

Fig. 9 **a** Finished simulation of 1800 identical particles with RVE of diameter $D = 50$ m bounds shown in blue. Axis ticks in m. **b** $|\overline{Z_i} - Z_i|$ as a function of D for 10 simulations, with values for each simulation given by a unique color, see Sect. 3. Convergence begins at about $D = 50$ m. Values after $D = 60$ m are unreliable as the RVE is greater than the box size

averaging Z_i with Z_{i-1} and Z_{i+1}. Then, $|\overline{Z_i} - Z|$ is calculated. This function quantifies the size of fluctuations in the value of Z as a function of D. As D is increased, this value is initially large due to the small RVE. However, the function oscillates about a smaller constant when D is sufficiently large enough to capture the general RVE behavior. As D is further increased, the RVE encompasses the boundary, and the value of the fluctuations suddenly change due to the boundaries' effect. The optimal value of D is when the domain is large enough to minimize the size of the fluctuations, but is minimally affected by the boundaries. This function of D becomes constant at about $D = 50$ m for the simulations, see Fig. 9. The convergence of $|\overline{Z_i} - Z|$ for 900 particles and 1800 for (R, C, A) and VAE simulations, respectively, in addition to the small variation in calculated ϕ and Z from the simulations indicates that the number of grains used in the simulations is satisfactory.

6.3 VAE parameters

Acknowledgements This work was supported by Army grant W911NF-19-1-0245.

Author Contributions *Conceptualization*: RBM, KK, SM, JEA. *Methodology*: RBM, KK, SM. *Writing—original draft*: RBM *Writing—review and editing*: SM, KK, JEA. *Formal Analysis/Investigation*: RBM, KK, SM. *Supervision/Funding Acquisition*: JEA

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Data availability Data are available upon reasonable request.

References

1. An, X., Li, C., Yang, R., Zou, R., Yu, A.: Experimental study of the packing of mono-sized spheres subjected to one-dimensional vibration. Powder Technol. **196**(1), 50–55 (2009)
2. Baker, J., Kudrolli, A.: Maximum and minimum stable random packings of platonic solids. Phys. Rev. E **82**(6), 061304 (2010)
3. Baule, A., Makse, H.A.: Fundamental challenges in packing problems: from spherical to non-spherical particles. Soft Matter **10**(25), 4423–4429 (2014)
4. Behringer, R.P., Chakraborty, B.: The physics of jamming for granular materials: a review. Rep. Prog. Phys. **82**(1), 012601 (2018)
5. Bentley, J.L., Ottmann, T.A.: Algorithms for reporting and counting geometric intersections. IEEE Trans. Comput. **28**(09), 643–647 (1979)
6. Brown, E., Rodenberg, N., Amend, J., Mozeika, A., Steltz, E., Zakin, M.R., Lipson, H., Jaeger, H.M.: Universal robotic gripper based on the jamming of granular material. Proc. Natl. Acad. Sci. **107**(44), 18809–18814 (2010)
7. Cheng, L., Bai, J., To, A.C.: Functionally graded lattice structure topology optimization for the design of additive manufactured components with stress constraints. Comput. Methods Appl. Mech. Eng. **344**, 334–359 (2019)
8. Cho, G.-C., Dodds, J., Santamarina, J.C.: Particle shape effects on packing density, stiffness, and strength: natural and crushed sands. J. Geotechn. Geoenviron. Eng. **132**(5), 591–602 (2006). https://doi.org/10.1061/(asce)1090-0241(2006)132:5(591)
9. Cundall, P.A., Strack, O.D.: A discrete numerical model for granular assemblies. Geotechnique **29**(1), 47–65 (1979)
10. de Graaf, J., van Roij, R., Dijkstra, M.: Dense regular packings of irregular nonconvex particles. Phys. Rev. Lett. **107**(15), 155501 (2011)
11. de Macedo, R.B., Marshall, J.P., Andrade, J.E.: Granular object morphological generation with genetic algorithms for discrete element simulations. Granular Matter **20**(4), 1–12 (2018)
12. Doersch, C.: Tutorial on variational autoencoders. arXiv preprint arXiv:1606.05908 (2016)
13. Donev, A., Connelly, R., Stillinger, F.H., Torquato, S.: Underconstrained jammed packings of nonspherical hard particles: ellipses and ellipsoids. Phys. Rev. E **75**(5), 051304 (2007)
14. Fortin, F.A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M., Gagné, C.: DEAP: evolutionary algorithms made easy. J. Mach. Learn. Res. **13**, 2171–2175 (2012)
15. Gan, J., Yu, A.: Dem simulation of the packing of cylindrical particles. Granular Matter **22**(1), 1–19 (2020)
16. Gladstone, R.J., Nabian, M.A., Keshavarzzadeh, V., Meidani, H.: Robust topology optimization using variational autoencoders. arXiv preprint arXiv:2107.10661 (2021)
17. Gravish, N., Franklin, S.V., Hu, D.L., Goldman, D.I.: Entangled granular media. Phys. Rev. Lett. **108**(20), 208001 (2012)
18. Guises, R., Xiang, J., Latham, J.-P., Munjiza, A.: Granular packing: numerical simulation and the characterisation of the effect of particle shape. Granular Matter **11**(5), 281–292 (2009)
19. Holland, J.H.: Genetic algorithms. Sci. Am. **267**(1), 66–73 (1992)
20. Jaeger, H.M., de Pablo, J.J.: Perspective: evolutionary design of granular media and block copolymer patterns. APL Mater. **4**(5), 053209 (2016)
21. Jerves, A.X., Kawamoto, R.Y., Andrade, J.E.: Effects of grain morphology on critical state: a computational analysis. Acta Geotech. **11**(3), 493–503 (2016)
22. Jerves, A.X., Kawamoto, R.Y., Andrade, J.E.: A geometry-based algorithm for cloning real grains. Granular Matter **19**(2), 30 (2017)
23. Jolliffe, I.T., Cadima, J.: Principal component analysis: a review and recent developments. Philosoph. Trans. R. Soc. A Math. Phys. Eng. Sci. **374**(2065), 20150202 (2016)
24. Karapiperis, K., Stainier, L., Ortiz, M., Andrade, J.: Data-driven multiscale modeling in mechanics. J. Mech. Phys. Solids **147**, 104239 (2021)
25. Kawamoto, R., Andò, E., Viggiani, G., Andrade, J.E.: Level set discrete element method for three-dimensional computations with triaxial case study. J. Mech. Phys. Solids **91**, 1–13 (2016)
26. Keller, S., Jaeger, H.M.: Aleatory architectures. Granular Matter **18**(2), 29 (2016)
27. Kollmer, J.E., Lindauer, S.M., Daniels, K.E.: Granular materials in space exploration. In: Earth and Space 2016. American Society of Civil Engineers (2016). https://doi.org/10.1061/9780784479971.021
28. Kumar, S., Kochmann, D.M.: What machine learning can do for computational solid mechanics. arXiv preprint arXiv:2109.08419 (2021)
29. Li, S., Lu, P., Jin, W., Meng, L.: Quasi-random packing of tetrahedra. Soft Matter **9**(39), 9298–9302 (2013)
30. Makse, H.A., Johnson, D.L., Schwartz, L.M.: Packing of compressible granular materials. Phys. Rev. Lett. **84**(18), 4160 (2000)
31. Meng, L., Lu, P., Li, S., Zhao, J., Li, T.: Shape and size effects on the packing density of binary spherocylinders. Powder Technol. **228**, 284–294 (2012)
32. Miskin, M.Z., Jaeger, H.M.: Adapting granular materials through artificial evolution. Nat. Mater. **12**(4), 326–331 (2013)
33. Miskin, M.Z., Jaeger, H.M.: Evolving design rules for the inverse granular packing problem. Soft Matter **10**(21), 3708–3715 (2014)
34. Monfared, S., Zhou, T., Andrade, J.E., Ioannidou, K., Radjaï, F., Ulm, F.-J., Pellenq, R.J.-M.: Effect of confinement on capillary phase transition in granular aggregates. Phys. Rev. Lett. **125**(25), 255501 (2020)
35. Oda, M., Iwashita, K.: Mechanics of Granular Materials. CRC Press, Florida (2020)
36. Pu, Y., Gan, Z., Henao, R., Yuan, X., Li, C., Stevens, A., Carin, L.: Variational autoencoder for deep learning of images, labels and captions. Adv. Neural. Inf. Process. Syst. **29**, 2352–2360 (2016)
37. Rath, S. R.: Generating fictional celebrity faces using convolutional variational autoencoder and pytorch, 2 2021. URL https://debuggercafe.com/generating-fictional-celebrity-faces-using-convolutional-variational-autoencoder-and-pytorch/
38. Rognon, P., Roux, J.-N., Wolf, D., Naaïm, M., Chevoir, F.: Rheophysics of cohesive granular materials. EPL (Europhysics Letters) **74**(4), 644 (2006)
39. Rosinha, I.P., Gernaey, K.V., Woodley, J.M., Krühne, U.: Topology optimization for biocatalytic microreactor configurations. Comput. Aided Chem. Eng. **37**, 1463–1468 (2015)
40. Salerno, K.M., Bolintineanu, D.S., Grest, G.S., Lechman, J.B., Plimpton, S.J., Srivastava, I., Silbert, L.E.: Effect of shape and friction on the packing and flow of granular materials. Phys. Rev. E **98**(5), 050901 (2018)
41. Schreck, C.F., Xu, N., O'Hern, C.S.: A comparison of jamming behavior in systems composed of dimer-and ellipse-shaped particles. Soft Matter **6**(13), 2960–2969 (2010)
42. Shi, J.-J., Zhang, W., Wang, W., Sun, Y.-H., Xu, C.-Y., Zhu, H.-H., Sun, Z.-X.: Randomly generating three-dimensional realistic schistous sand particles using deep learning: variational autoencoder implementation. Eng. Geol. **291**, 106235 (2021)
43. Staron, L., Hinch, E.: The spreading of a granular mass: role of grain properties and initial conditions. Granular Matter **9**(3), 205–217 (2007)
44. Wang, L., Chan, Y.-C., Ahmed, F., Liu, Z., Zhu, P., Chen, W.: Deep generative modeling for mechanistic-based learning and

design of metamaterial systems. Comput. Methods Appl. Mech. Eng. **372**, 113377 (2020)
45. Wang, Y., Li, L., Hofmann, D., Andrade, J.E., Daraio, C.: Structured fabrics with tunable mechanical properties. Nature **596**(7871), 238–243 (2021)
46. Yang, J., Luo, X.: Exploring the relationship between critical state and particle shape for granular materials. J. Mech. Phys. Solids **84**, 196–213 (2015)
47. Zhao, J., Li, S., Lu, P., Meng, L., Li, T., Zhu, H.: Shape influences on the packing density of frustums. Powder Technol. **214**(3), 500–505 (2011)
48. Zhou, B., Wang, J.: Generation of a realistic 3d sand assembly using x-ray micro-computed tomography and spherical harmonic-based principal component analysis. Int. J. Numer. Anal. Meth. Geomech. **41**(1), 93–109 (2017)
49. Zhu, J., Zhou, H., Wang, C., Zhou, L., Yuan, S., Zhang, W.: A review of topology optimization for additive manufacturing: status and challenges. Chin. J. Aeronaut. **34**(1), 91–110 (2020)