



HAL
open science

Efficient data structures for model-free data-driven computational mechanics

Robert Eggersmann, Laurent Stainier, Michael Ortiz, Stefanie Reese

► **To cite this version:**

Robert Eggersmann, Laurent Stainier, Michael Ortiz, Stefanie Reese. Efficient data structures for model-free data-driven computational mechanics. *Computer Methods in Applied Mechanics and Engineering*, 2021, 382, pp.113855. 10.1016/j.cma.2021.113855 . hal-03413194

HAL Id: hal-03413194

<https://hal.science/hal-03413194>

Submitted on 5 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Efficient Data Structures for Model-free Data-Driven Computational Mechanics

Robert Eggersmann^{a,*}, Laurent Stainier^b, Michael Ortiz^{c,d}, Stefanie Reese^a

^a*Institute of Applied Mechanics, RWTH Aachen University, Mies-van-der-Rohe-Str.1, D-52074 Aachen, Germany.*

^b*Institute of Civil and Mechanical Engineering, École Centrale de Nantes, 1 Rue de la Noë, F-44321 Nantes, France.*

^c*Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA.*

^d*Hausdorff Center for Mathematics, Rheinische Friedrich-Wilhelms-Universität Bonn, Endenicher Allee 62, D-53115 Bonn, Germany.*

Abstract

The data-driven computing paradigm initially introduced by Kirchdoerfer & Ortiz [1] enables finite element computations in solid mechanics to be performed directly from material data sets, without an explicit material model. From a computational effort point of view, the most challenging task is the projection of admissible states at material points onto their closest states in the material data set. In this study, we compare and develop several possible data structures for solving the nearest-neighbor problem. We show that approximate nearest-neighbor (ANN) algorithms can accelerate material data searches by several orders of magnitude relative to exact searching algorithms. The approximations are suggested by—and adapted to—the structure of the data-driven iterative solver and result in no significant loss of solution accuracy. We assess the performance of the ANN algorithm with respect to material data set size with the aid of a 3D elasticity test case. We show that computations on a single processor with up to one billion material data points are feasible within a few seconds execution time with a speed up of more than 10^6 with respect to exact k -d trees.

Keywords: data-driven computing, solid mechanics, nearest neighbor problem, approximate nearest-neighbor search, data structures, data science

*Corresponding author

Email addresses: robert.eggersmann@ifam.rwth-aachen.de (Robert Eggersmann), laurent.stainier@ex-nantes.fr (Laurent Stainier), ortiz@aero-caltech.edu (Michael Ortiz), reese@ifam.rwth-aachen.de (Stefanie Reese)

1. Introduction

The classical paradigm of computational mechanics is to use data from experimental tests to formulate material models by fitting to the data, then use the models in calculations. The process of modeling is often ill-posed and open-ended, results in loss of information relative to the original material data set and introduces epistemic uncertainty into the calculations. As material data becomes more plentiful, owing to advances in experimental science, multiscale modeling and other data sources, a new paradigm, model-free Data-Driven (DD) computational mechanics, suggests itself. The objective of the DD paradigm, first introduced in [1], is to make predictions regarding the behavior of physical systems directly from material data, without the intermediate step of material modeling.

DD entails a reformulation of the initial boundary value problems in which the solution is sought within a set of admissible states, or *constraint set*, and required to minimize distance to the material data set. The constraint set spans all states in phase space, e. g., stress and strain, that satisfy equilibrium and kinematic relations, as well as boundary conditions. DD has been extended to noisy data sets [2], dynamics [3], inelasticity [4], finite deformations [5–7], fracture [8], and second-order approximations [9]. Leygue et al. (2018) [10–13] proposed a closely-related inverse approach, termed Data-Driven Material Identification (DDMI), that takes strain fields from full-field optical measurements and corresponding loading boundary conditions to infer the corresponding stress field. DDMI can be used to generate very large material data sets directly from a small number of full-field microscopy measurements.

DD solvers, such as proposed in [1], iteratively minimize the distance between the material data set and the constraint set. From the standpoint of time efficiency, the central step in the solver is searching for the point in the material data set nearest to a given local state. Evidently, this search is an instance of the classical nearest-neighbor problem. The nearest-neighbor problem is also referred to as the post-office problem and can be solved by means of many algorithms. In general, the problem is defined by finding the closest point in a given set to a query point, where both the elements of the given set and the query point are points in the same space. A generalization of the problem is the k -nearest-neighbor problem, which is also closely related to the fixed-radius nearest-neighbor problem and the all nearest-neighbor problem.

The simplest algorithm for finding the exact solution to the nearest-neighbor problem is the linear search, where the query point is compared to every point in the data set. This naive approach is prohibitively expensive for very large data sets. Another possibility is the use of space partitioning algorithms where the branch and bound method [14] is applied. An early example is the k -d tree developed by Bentley & Friedman in 1975 [15, 16]. Here, k represents the dimension of the data set, and the search space is recursively partitioned along one dimension in each branching step. Another early algorithm is the k -means tree investigated by Fukunaga & Narendra in 1975 [17], where k stands for the

number of clusters built into each branch. Other algorithms include the Quad-tree [18], the R-tree [19], the metric-tree [20–22], the ball-tree [23], among many others, with numerous variants thereof. A deficiency of exact methods is that they cannot guarantee logarithmic search time, as shown in [20] or [24].

Conveniently, for many applications it is sufficient to find a good guess close to the true nearest-neighbor. These approximate nearest-neighbor (ANN) algorithms have recently attracted much attention since they are of practical importance for many applications in sales or social media. The main idea behind ANN algorithms is to exploit the trade-off between efficiency and accuracy. In an early contribution, Miclet & Dabouz [25] investigated a hierarchical search structure based on k -means with no backtrack. Arya *et al.* [26, 27] introduced the $(1 + \varepsilon)$ -nearest-neighbor criterion. A $(1 + \varepsilon)$ -nearest neighbor is a point, whose distance to the query point is less than $(1 + \varepsilon)$ times the distance to the true nearest neighbor. Brin [28] proposed a geometric near-neighbor access tree (GNAT) as an extension of k -means trees.

An approach based on randomized k -d trees was proposed in [29], where the search space is partitioned multiple times and no backtracking is performed. Muja and Lowe [30] further compared the randomized k -d trees with a hierarchical k -means tree using a bounded priority queue for backtracking. This idea of using priority queues for backtracking, also called the best-bin first (BBF) method, was initially proposed in [31] for k -d trees. The nodes to be checked during backtracking are stored in a sorted queue that checks the nodes with the smallest bin distances first. The length of the priority queue is then limited to be less than a prescribed number. Muja and Lowe also launched the FLANN-library [32], which is well-established today. Further investigations on the scalability of the algorithms were presented in [33].

More recent work has focused on proximity graph-based methods for high-dimensional problems. Those problems arise, e. g., in image recognition, computational linguistics, or product recommendation. Graph-based methods seem to be superior in those fields. Early graph-based searching methods were based on monotonic search networks introduced by Dearholt *et al.* [34] and on randomized neighborhood graphs by Arya and Mount in 1993 [26]. Hajebi *et al.* [35] do the nearest-neighbor search by means of a hill-climbing algorithm in a graph where every node is linked with its k nearest-neighbors. Dong *et al.* [36] proposed an algorithm to construct the k -NN graphs efficiently by an iterative procedure. Other efficient graph search algorithms are the navigable small world graph [37] and the hierarchical navigable small world (HNSW) graph [38]. In these methods, a small world graph is an approximation of a Delaunay graph. The HNSW graph has multiple layers. The bottom layer includes all points. With increasing layer number, the number of points decreases. These graph-search algorithms are compared with the diversified proximity graph, which is based on an existing k -NN graph, in [39]. In [40], a navigating spreading-out graph is proposed that enables billion-scale data set searches for online sales applications. The same authors introduce the satellite system graph (SSG) in [41]. This graph regulates its sparsity by a minimal angle of the vectors to its

neighbors. Recently, *Groh et al.* [42] proposed a graph-based GPU nearest-neighbor algorithm that betters the currently most efficient algorithms by more than a magnitude in search time.

In general, choosing the best search algorithm depends on many parameters. Data-related properties such as the dimension, number of examples, correlations, and density distributions need to be carefully considered. In addition, the objectives of the user must be taken into consideration. Relevant considerations are query time, accuracy, query workload, building time and memory usage.

Our present work focuses on the suitability and performance of different nearest-neighbor search algorithms in the context of Data-Driven computing with noise-free material data sets of up to a billion points. We investigate tree-based and graph-based methods. The aim of the work is to ascertain how the specific features of DD iterative solvers can be best exploited to accelerate searches. An examination of the iterative solvers suggests two lines inquiry:

1. We observe that it is sufficient to use rough guesses for the nearest-neighbor search in the initial DD iterations during which the accuracy of ANN algorithms can be set low. With decreasing distance to the solution, the accuracy of the searches needs to be steadily increased.
2. We additionally observe that the query points move relatively little when approaching convergence. In the context of graph search algorithms, this effect represents an increasing amount of knowledge that can be used to choose better starting points for navigating through the search graph.

We show that, by exploiting these features, approximate nearest-neighbor (ANN) algorithms can accelerate material data searches by several orders of magnitude relative to exact searching algorithms. We emphasize that the approximations are suggested by—and adapted to—the structure of the data-driven iterative solver and result in no significant loss of solution accuracy. We additionally assess the performance of ANN algorithms with increasing material data set size with the aid of a 3D elasticity test case. We show that computations on a single processor with up to one billion material data points are feasible within a few seconds execution time with a speed up of more than 10^6 with respect to exact k -d trees.

The paper is structured as follows. In Section 2, we recapitulate the DD paradigm and its iterative solution procedure. The efficiency of a 3D elastic solid example with varying data set sizes using a k -d tree is investigated in Section 3. In addition, we investigate the characteristics of the nearest-neighbor search that appear in the data-driven calculation. In Section 4, we study different ANN algorithms for accelerating computations and assess their performance on material data sets of up to one million points, with specific focus on the trade-off between accuracy and speed. A comparison of the different ANN algorithms with the most efficient parameters follows. These comparisons are based on data sets of up to 100 million points. Finally, we show that computations with billion-point data sets are possible within seconds. Section 5 closes with final conclusions and outlook.

2. The iterative solver of the data-driven problem

The data-driven formulation of the discretized initial boundary value problem of elasticity, as proposed in [1, 5], can be stated as follows. A system undergoes displacements $\mathbf{u} = \{\mathbf{u}_i\}_{i=1}^n$, with $\mathbf{u}_i \in \mathbb{R}^{n_i}$ being the displacement vector at node i with dimension n_i at all nodes $i = 1, \dots, n$, under the action of applied forces $\mathbf{f} = \{\mathbf{f}_i\}_{i=1}^n$. The vector $\mathbf{f}_i \in \mathbb{R}^{n_i}$ denotes the nodal force vector. The following minimization then defines the data-driven problem

$$\min_{\mathbf{y} \in C} \min_{\mathbf{z} \in D} d^2(\mathbf{y}, \mathbf{z}), \quad (1)$$

i. e., the objective is to find the state \mathbf{y} in the constraint set C , which is closest to the state \mathbf{z} in the data set D , which is again closest to the constraint set. Here, the constraint set C , as well as the data set D , are subsets of the global phase space Z . The constraint set contains all admissible states fulfilling equilibrium and kinematic relations. Experimental measurements or computations of microstructures sample the data set. The associated squared distance $d^2(\mathbf{y}, \mathbf{z})$ in phase space is defined by

$$d^2(\mathbf{z}, \mathbf{y}) = \sum_{e=1}^m \frac{1}{2} w^{(e)} d_e^2(\mathbf{z}^{(e)}, \mathbf{y}^{(e)}), \quad (2)$$

where w_e are volumes associated with the integration points $e = 1, \dots, m$ with $\mathbf{y}^{(e)} = (\boldsymbol{\varepsilon}^{(e)}, \boldsymbol{\sigma}^{(e)})$, $\mathbf{z}^{(e)} = (\boldsymbol{\varepsilon}'^{(e)}, \boldsymbol{\sigma}'^{(e)})$ being states in the local phase space $Z^{(e)} \in \mathbb{R}^{2M}$. The corresponding local distance is then defined by

$$d_e^2(\mathbf{y}^{(e)}, \mathbf{z}^{(e)}) = \mathbb{C}^{(e)}(\boldsymbol{\varepsilon}^{(e)} - \boldsymbol{\varepsilon}'^{(e)}) \cdot (\boldsymbol{\varepsilon}^{(e)} - \boldsymbol{\varepsilon}'^{(e)}) + \mathbb{C}^{(e)-1}(\boldsymbol{\sigma}^{(e)} - \boldsymbol{\sigma}'^{(e)}) \cdot (\boldsymbol{\sigma}^{(e)} - \boldsymbol{\sigma}'^{(e)}). \quad (3)$$

Here, the metric $\mathbb{C}^{(e)}$ is a symmetric positive definite matrix bringing the absolute values of stress and strain to an equal scale.

The solution scheme proposed in [1] then iteratively projects a state $\mathbf{z}_i \in D$ with $\mathbf{z}_i^{(e)} = (\boldsymbol{\varepsilon}_i^{*(e)}, \boldsymbol{\sigma}_i^{*(e)})$ to the closest point in the constraint set $\mathbf{y}^{i+1} \in C$, where the index i indicates the current iteration.

$$\mathbf{y}_{i+1} = P_C(\mathbf{z}_i). \quad (4)$$

The projection into the constraint set C is performed by solving the two equation systems

$$\underbrace{\mathbf{A}_{e=1}^m \{w^{(e)} \mathbf{B}^{(e)T} \mathbb{C}^{(e)} \mathbf{B}^{(e)}\}}_{\mathbf{K}} \mathbf{u} = \underbrace{\mathbf{A}_{e=1}^m \{w^{(e)} \mathbf{B}^{(e)T} \mathbb{C}^{(e)} \boldsymbol{\varepsilon}^{*(e)}\}}_{\mathbf{E}}, \quad (5a)$$

$$\underbrace{\mathbf{A}_{e=1}^m \{w^{(e)} \mathbf{B}^{(e)T} \mathbb{C}^{(e)} \mathbf{B}^{(e)}\}}_{\mathbf{K}} \boldsymbol{\eta} = \mathbf{f} - \underbrace{\mathbf{A}_{e=1}^m \{w^{(e)} \mathbf{B}^{(e)T} \boldsymbol{\sigma}^{*(e)}\}}_{\mathbf{S}}, \quad (5b)$$




Figure 1: Illustration of the data-driven iterative solver for a single truss problem with load F and cross section area A . Data set (black points) and constraint set (blue line) are given in the ε - σ -phase space. Left side: Projection P_C of state z_i in the constraint set. Right side: Projection P_D of state y_i in the data set D .

and computing the new values for stress and strain by

$$\boldsymbol{\varepsilon}^{(e)} = \mathbf{B}^{(e)} \mathbf{u}^{(e)}, \quad \forall e = 1, \dots, m, \quad (6a)$$

$$\boldsymbol{\sigma}^{(e)} = \boldsymbol{\sigma}^{*(e)} + \mathbb{C}^{(e)} \mathbf{B}^{(e)} \boldsymbol{\eta}^{(e)}, \quad \forall e = 1, \dots, m. \quad (6b)$$

where $\mathbf{B}^{(e)}$ denotes the classical strain-displacement operator.

Regarding the computational implementation it is reasonable to factorize the stiffness matrix \mathbf{K} directly after the first assembly. Then, for geometric linear and constant $\mathbb{C}^{(e)}$, only \mathbf{E} and \mathbf{S} have to be reassembled in every iteration.

The second part of a single iteration is to find the closest state in the data set to the previously calculated state in the constraint set. This second projection

$$\mathbf{z}_{i+1} = P_D(\mathbf{y}_{i+1}) \quad (7)$$

then specifies to minimize the local distances d_e for given states $\mathbf{y}_{i+1}^{(e)}$ at all integration points e . In other words, a nearest-neighbor problem has to be solved for each integration point, where we aim to find the state in the data set $\mathbf{z}_{i+1}^{(e)} \in D_e = \{\boldsymbol{\varepsilon}_j^{(e)}, \boldsymbol{\sigma}_j^{(e)}\}_{j=1}^N$ closest to $\mathbf{y}_{i+1}^{(e)}$ regarding the metric (3). The data set's size is specified by the cardinality N and dimension $\dim(Z^{(e)})$. The latter is the sum of the M stress and M strain components in Voigt notation.

One iteration of the DD solver can be finally expressed by

$$\mathbf{z}_{i+1} = P_D(P_C(\mathbf{z}_i)). \quad (8)$$

Starting the solver, the integration points are randomly associated with points from the data set. The final result of the problem is then chosen to be the state in the constraint set obtained when the global distance does not decrease any more. Equivalently, if the nearest-neighbor projection is exact, no change in the integration point’s association with the data points will finally be observed. Here, it should be remarked that the iterative solver with the minimum distance formulation does not find the best solution to the data-driven problem described in Eq. 1 in general. This is because the solver can stop at a local minimum if distances between points are too large. However, since this effect usually appears only quite close to the true solution, the obtained solution is usually a good approximation. It is possible to reduce this effect by using e. g., the maximum-entropy formulation [2] or a local second-order approximation described in [9] or a combination of both.

3. Example: 3D elastic solid with exact nearest-neighbor projections

In this work, all investigations are based on the problem of a three-dimensional elastic solid described in the following. Nevertheless, the phenomena, which will be discussed here, also occur with problems of lower-dimensional data sets like trusses or continua in plane strain or plane stress conditions. Furthermore, the following search algorithms can be applied on problems in dynamics [2] and inelasticity [4] in the same way.

The boundary value problem considered is a cube of side length 10 mm discretized by $20 \times 20 \times 20$ elements. The degrees of freedom at the bottom are fixed in all directions, whereas the cube’s upper surface is rotated by 2° , as depicted in Fig. 2. In total, 64000 integration points are evaluated in each iter-




Figure 2: Elastic solid cube with discretization and boundary conditions. Nodes at the lower blue surface are fixed. Nodes at the upper red surface are rotated around the centered z-axis and fixed in z-direction.

ation.

The reference material behavior is given by an isotropic but nonlinear material description of the form:

$$\boldsymbol{\sigma} = E(\boldsymbol{\varepsilon} + \alpha \boldsymbol{\varepsilon}^3 + 0.5(\text{tr}(\boldsymbol{\varepsilon}) + \alpha \text{tr}(\boldsymbol{\varepsilon})^3)\mathbf{I}), \quad (9)$$

where $E = 1000$ MPa and $\alpha = 500$ are material parameters.

Data sets $D^{(e)} = \{(\boldsymbol{\varepsilon}_i, \boldsymbol{\sigma}_i)\}_{i=1}^N$ of dimension twelve and N artificial measurements are randomly created on a range $[-0.025; 0.025]$ for all strain components in Voigt notation. The corresponding stresses are computed by the reference material model (Eq. 9).

Here, the constant matrix $\mathbb{C}^{(e)}$ is determined by using a principal component analysis applied to the data. Therefore, the first 6 principal component vectors are determined and written in a (12×6) -matrix \mathbf{A} . Then $\mathbb{C}^{(e)} = \text{sym}(\mathbf{A}_\varepsilon^{-1} \mathbf{A}_\sigma)$, where \mathbf{A}_ε is the upper and \mathbf{A}_σ the lower (6×6) -matrix of \mathbf{A} .

3.1. Performance studies using an exact 12-d tree.

Initially, we compare the DD-solver's computation times for increasingly fine data sets of 10^3 , 10^4 , 10^5 , and 10^6 points with 20 random samples each. The studies were performed using a AMD Ryzen™ 9 3900X 12-core processor. The c++ implementation makes use of the eigen3-library and the build-in Sparse-LU solver provided with this library. The nearest-neighbor problem is solved using a 12-d tree based on an own implementation according to [15, 16]. Times are measured for the assembly and LU decomposition of \mathbf{K} , the right-hand side assembly of \mathbf{E} and \mathbf{S} , solving the two equation systems (5a) and (5b), and finding the nearest-neighbors in the data sets. The results of this study are depicted in Fig. 3a. On the one hand, it can be observed that the total computational time increases heavily with the number of data points. This is majorly caused

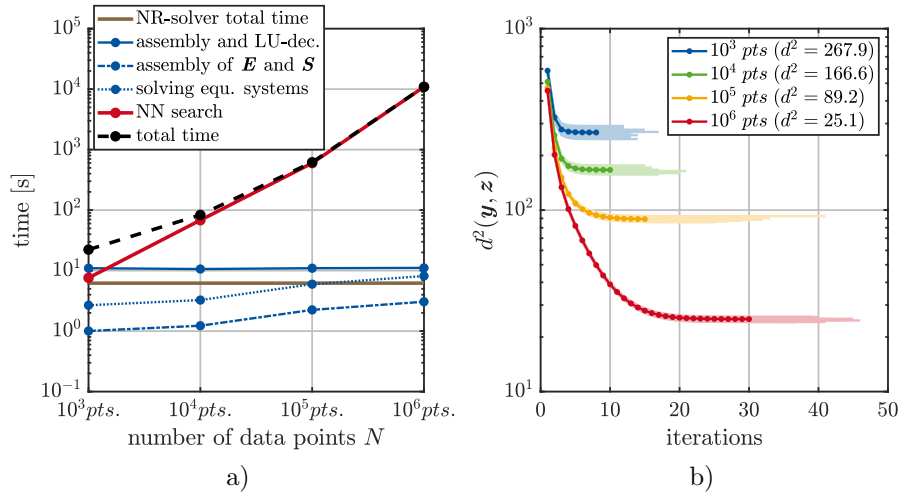


Figure 3: Elastic solid example: 12-d tree used as data structure for exact nearest-neighbor search. a) Comparison of computational times for different data set refinements. Times are averaged over the 20 samples. b) Global squared distance over solver iterations for different data set refinements. Averaged distances in bold, all results in light colors. Averaged final distances after 8, 10, 15 and 30 iterations stated in brackets.

by the fact that the time for the nearest-neighbor search increases drastically with the data set size. On the other hand, with an increasing fine data set, the remaining squared distance decreases, as shown in Fig. 3b. The variance in the results reduces as well, which can be observed from the curves the light colors representing all performed computations.

Additional effects can be observed considering the computational times for assembling the right-hand sides and solving the two equation systems. Here, the computational time for both parts slightly increases since the number of iterations needed to converge increases as well (see Fig. 3b). The time for the initial assembly and LU decomposition of the stiffness matrix remains constant, as expected. As an additional benchmark, we stated the time for computing the reference solution (6.2 s). Here, the Newton-Raphson solver uses the ConjugateGradient solver from the eigen3 library.

From this observation we can conclude that the DD-solver’s performance majorly depends on the time for solving the nearest-neighbor problem. One can further remark that it might be reasonable to use more advanced and more problem-specific linear solvers than those used in this study. Nevertheless, the study shows clearly that the crucial step remains the nearest-neighbor search. Thus, the main focus is set on the investigation of the latter.

The critical reader might wonder why the times for finding the nearest-neighbors increase in the order of $O(N)$ and not in the order $O(\log(N))$ as expected for k -d trees. The reason for this effect can be detected if one studies the single search times per iteration. These search times are depicted in Fig. 4a. For larger data sets, the main time effort is used in the first iterations. For example, more than half (51.6%) of the total search time is needed in the first three iterations of the 10^6 points computation. We can further study those observations regarding the number of distance comparisons that were made during the search. The comparisons over iterations are shown in Fig. 4b.

These results clearly correlate with the search times depicted in Fig. 4a. Another observation is that for the computations with 10^6 points, on average approx. 40% of all points have to be compared in the first iteration. For smaller data sets the percentage even increases. To further investigate this effect, we studied the correlation between the number of comparisons and the distances between the query point and its nearest-neighbor. In Fig. 5, the number of needed comparisons is plotted over the final query distance of a million points sample in iteration 1, 5, and 20 for all integration points. Here, the final query distances equal the final local distances $d_e^2(\mathbf{y}^{(e)}, \mathbf{z}^{(e)})$. Especially for iteration 5 and 20, a clear correlation between the distances and the number of comparisons can be observed. For some query points in iteration 1, almost every point has to be checked. A further observation is that many points move from the top right to the bottom left with increasing iterations. That means the query distances decrease over the iterations, and as a consequence, the number of comparisons decreases as well.

Then finally, we have to clarify the reason for this correlation. Therefore, we investigate the search in k -d trees in the following.




Figure 4: Elastic solid example: 12-d tree used as data structure for exact nearest-neighbor search. a) Averaged query times per iteration over iterations for different data set refinements. Averaged total query times stated in brackets. b) Distance comparisons per iteration over iterations for different data set refinements. Averaged results in bold, all results in light colors.




Figure 5: Elastic solid example: number of comparisons over final query distance for all integration points in iteration 1 (left), 5 (mid) and 20 (right). Results of one computation with 10^6 points.

3.2. Search in k -d trees

The construction of the k -d tree is based on the branch and bound method [14]. That means the data set is steadily divided into two branches according to the median coordinate of a certain branch dimension. This median coordinate defines the bound of the branch. Using this bound, additional checks in neighbored branches can be reduced. Fig. 6 shows a tree of a small data set with 18 points created according to the above procedure. The corresponding space partitioning is depicted as well.

The search in a k -d tree now works as follows. Starting from the root node on top of the tree, the branch with its corresponding node is recursively chosen where the query point is positioned in. The procedure ends when a chosen node has no children. Thereupon, the backtracking starts, which means that all non-chosen branches are checked according to the defined bound. As it is standard for k -d trees, non-chosen branches with according subbranches will be evaluated, if the following inequality is fulfilled:

$$\underbrace{|x_i - q_i|}_{d_b} > d_c, \quad (10)$$

where x_i and q_i are the coordinates in branch dimension i of the branch node and the query node, respectively. The absolute difference between both coordinates d_b is then the distance to the corresponding branch bound, and d_c represents the current best distance. In terms of DD distances (see Eq. 3) the coordinates above are computed in an intermediate mapping where, e. g., $\mathbf{q} = (\mathbb{C}^{(e)1/2}\boldsymbol{\varepsilon}_i \quad \mathbb{C}^{(e)-1/2}\boldsymbol{\sigma}_i)^T$.

For material data in the phase space a special situation occurs, which is not typical for nearest-neighbor problems. Since the material points typically lie in or close to a lower-dimensional manifold, two cases can be differentiated. In the first situation, the query point is relatively close to the points in the data set. The second situation deals with a query point which is far off the underlying manifold. An example of situation 1 and situation 2 is explained in Fig. 6a and Fig. 6b, respectively.

Situation 1 (close query point): First, we determine the sector where the query point is positioned in. In the current example, this means only four comparisons are needed, which is advantageous compared to a linear search with 18 comparisons. During backtracking, we find the situation that no additional branches have to be evaluated. As illustrated in Fig. 6a, figuratively speaking, this is because there is no intersection between a voxel boundary and the circle around the query point with the best distance radius. This means in the current example that the use of a 2-d tree would be advantageous compared to a linear search with 18 comparisons.

Situation 2 (far off query point): Again, we start following the same voxel’s query path as in the above situation. Here, the light purple root node is the current best point so far (see Fig. 6b). Now, there is the situation that numerous intersections between the dashed black line and the voxel boundaries exist. Thus, we have to investigate additional branches, and more comparisons have to be made. Finally, we compare 12 out of 18 nodes. All checked sectors are highlighted, and the corresponding children in the search tree are marked in red. For this higher ratio of comparisons, the 2-d tree might even be dominated by a linear search due to its additional overhead.

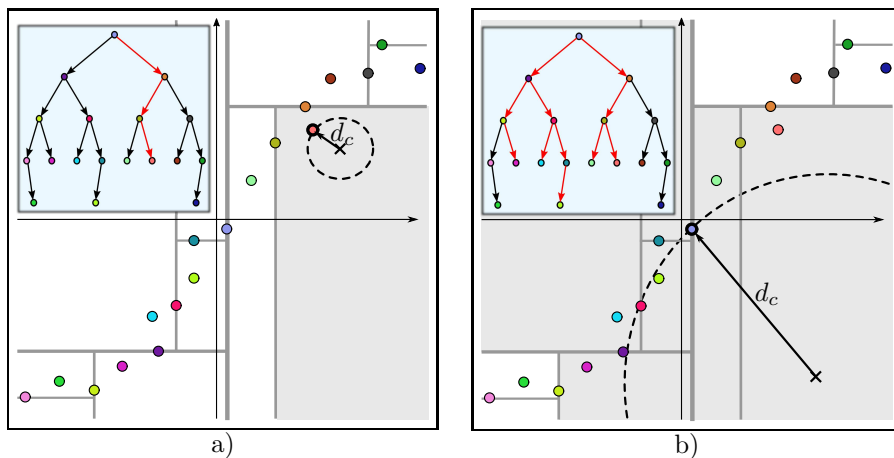


Figure 6: Space partitioning by a 2-d tree for a small data set with corresponding search tree. Sectors to be checked during the search for the closest point to the query point (black cross) are highlighted and corresponding search paths in the search trees are marked in red. Initial query distances d_c indicated by dashed lines. a) Situation where the distance between the query point and its closest point is relatively small. b) Situation where the distance between the query point and its nearest-neighbor is rather large.

This simple example gives an explanation for the correlation observed in Fig. 5. The effect intensifies for higher dimensional spaces because here, the distances between the points increase as well (see e. g. [20, 24]).

4. Nearest neighbor approximations

The previous chapter showed that solving the nearest-neighbor problem dominates the whole solution procedure for larger data sets. Therefore, we compared different approximate nearest-neighbor algorithms to speed up the search. In all presented approximation algorithms, the trade-off between computational efficiency and accuracy will be investigated.

As already mentioned in the introduction, the first central idea of this work is to control the accuracy so that it increases during the data-driven solver’s iterations. If we remember that the nearest-neighbors found in the first iterations

are only intermediate results, we can assume that it will be sufficient to demand only a low accuracy at the beginning. In the very first iterations, a very rough guess might even be good enough. That should lead to high computational time savings since those will be the most time-intensive ones, at least if we use a k -d tree.

A simple implementation could e. g. control an arbitrary search parameter f which influences the accuracy

$$f = \hat{f}(d(\mathbf{z}, \mathbf{y}), d_e(\mathbf{z}^{(e)}, \mathbf{y}^{(e)}), i), \quad (11)$$

where the parameter could be a function of the latest global or local distance and the iteration number i .

4.1. Approximate nearest-neighbors using k -d trees

All tree-based approximate nearest-neighbor algorithms have in common that they reduce additional branch evaluations during backtracking. For example, Fig. 7a shows a search tree for a specific query request. In this case, we perform a single run from top to bottom so that only the red paths are evaluated. To perform an exact search, where we can ensure to find the exact nearest-neighbor, the orange paths would have to be evaluated in addition. By neglecting the orange paths during backtracking, we accelerate the search because less distance comparisons have to be done. In contrast, the accuracy reduces since we might find a point which is just close, but not the closest. There are several ways to achieve this. One possibility is to limit the number of additional evaluations during backtracking. In [29] this procedure is investigated with the additional use of a priority queue on multiple randomized k -d trees. Several differently partitioned k -d trees are assessed, and those non-chosen branches closest to the query point are evaluated first. In [30], the authors compare this procedure to different methods and provide an implementation in the FLANN-library [32].

Another procedure, which we followed here, is proposed in [27]. The idea is to introduce a control parameter $f_d \in [0, 1]$, which is multiplied with the current best distance d_c so that inequality 10 is modified to

$$d_b > f_d \cdot d_c. \quad (12)$$

If we chose f_d to be one, the exact k -d tree will be recovered. Counter-wise, if $f_d = 0$, no additional branch will be evaluated. Thus, additional evaluations can be reduced because the number of bound intersections is decreased (see Fig. 7b).

Parameter studies were performed for the described implementation of the k -d tree with parameters $f_d = 0, 0.2, 0.4, 0.6, 1.0$. The same 20 samples of the 1 million points data sets, which we already used in the study of section 3.1, were investigated. These results are depicted in Fig. 8a and Fig. 8b showing the query times and distances over the iterations. For the sake of comparability, we




Figure 7: a) Illustration of reduced search paths in a tree structure. Red paths are evaluated, whereas the orange paths are neglected. b) Reduction of the number of checks by reducing the current best distance d_c by a factor f_d for backtracking.

consider the first 30 iterations only.

On the one hand, we can reduce the search times per iteration by almost 4 magnitudes. As mentioned before, for $f_d = 1.0$ the same results as for the exact k -d tree in Fig. 3 are retrieved. For the control parameter $f_d = 0.0$, we perceive query times of less than 0.04 seconds per iteration on average. Regarding the query times of the 20 samples, only very small fluctuations occur.

On the other hand, in Fig. 8b, we observe that the remaining distances after 30 iterations increase with decreasing values of f_d . But two exceptions from this can be seen. First, the remaining distances with $f_d = 0.0$ fall below those with $f_d = 0.2$. Second, the results with $f_d = 0.6$ show on average slightly lower distances than the computations with $f_d = 1.0$.

To explain these observations is difficult, but the authors have the following assumption. As mentioned above, the data-driven iterative solver might stop in local minima. The worse the nearest-neighbor algorithm’s accuracy, the higher the probability of running into a local minimum. However, a very “bad guess” for one integration point in an iteration can lead to a better solution if a larger step beyond a local minimum is done in the following iteration. This assumption would also explain the high variability observed for the computations with $f_d = 0.0$ because with lower f_d -values more local minima exist and larger local steps are more likely. The described effect seems to be intensified due to the k -d space partitioning, as we can conclude from the following section.

4.2. Approximate nearest-neighbors using k -means trees

As a second tree-based search algorithm, we investigate the k -means tree. This tree was proposed in the early work of Fukunga & Narendra [17]. It is based on the idea of recursively branching the data into k clusters until all clusters are smaller than a certain bucket size. This clustering can be performed

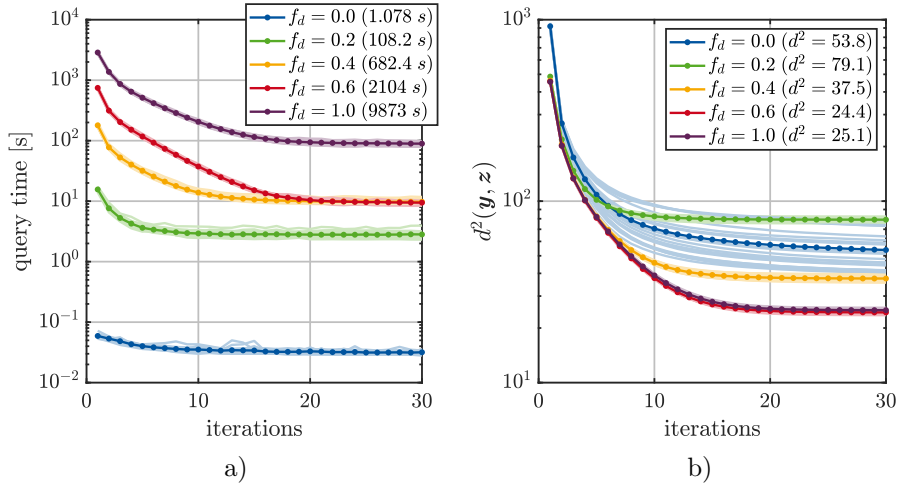


Figure 8: Elastic solid example: Approximate 12-d tree used for nearest-neighbor search on 20 one million points data sets. Results are studied for different parameters f_d . Averaged results are depicted in bold, all results in light colors. a) Query times per iteration over solver iterations. Total search times within the first 30 iterations stated in brackets. b) Global squared distance over solver iterations. Remaining squared distance after 30 iterations stated in brackets.

in principle by different clustering techniques. In [17], the k -means clustering algorithm was suggested, which we also used in our studies. Then, the node corresponding to a cluster bears the coordinates of the midpoint of all points in this cluster. The node is also equipped with the largest distance d_r from the cluster mean to all points included in the cluster.

The search in the k -means tree starts from the root node. Then, recursively all distances to its children are computed and sorted. The algorithm then chooses the child or cluster closest to the query point as the next node to evaluate until a leaf node is reached. During backtracking the non-chosen clusters are evaluated if

$$d(\mathbf{x}, \mathbf{q}) - f_d \cdot d_r < d_c, \quad (13)$$

where $d(\mathbf{x}, \mathbf{q})$ defines the distance between the non-chosen cluster midpoint \mathbf{x} and the query point \mathbf{q} . Here, the triangular inequality is exploited to reduce the number of necessary evaluations. Besides, the accuracy control parameter f_d is used in a similar manner as introduced before in section 4.1. However, here f_d reduces the radius d_r to ensure that the exact search is performed for $f_d = 1.0$ and no backtracking is done for $f_d = 0.0$.

According to the described k -means tree, a space partitioning is depicted in Fig. 9b with $k = 3$. First of all, the space partitioning is now similar to the data set's Voronoi tessellation (see Fig. 9a). This space partitioning appears to be

more natural to the problem compared to the before discussed 2-d partitioning. It is here sufficient to run a single search path to find the exact result regarding the depicted query situation. This is because there is no intersection between the circle around the query point with radius d_c and the circle around any non-evaluated cluster with radii d_r .

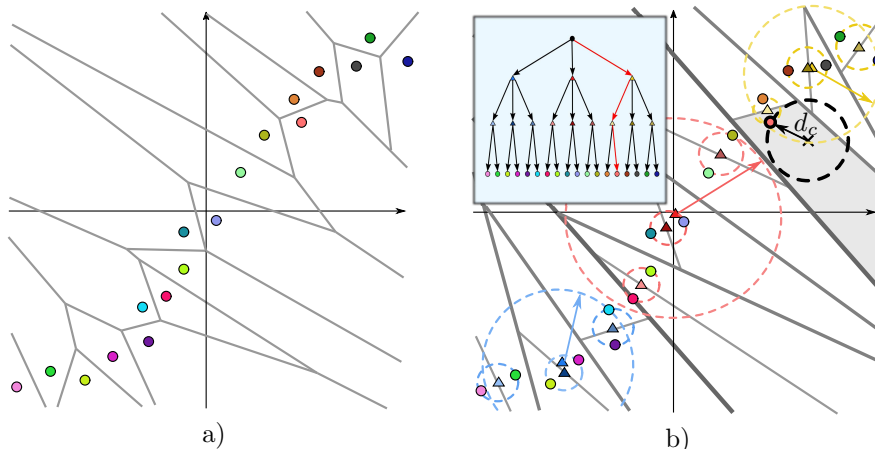


Figure 9: a.) Voronoi tessellation of a small data set. b) Space partitioning by a 3-means tree for the same data set with corresponding search tree. Sectors to be checked during the closest-point search to the query point (black cross) are highlighted and corresponding search paths in the search tree are marked in red. Final query distance indicated by black dashed line. Center points of clusters are indicated by triangles. Relevant cluster radii d_r for back checks are marked by colored dashed lines.

A parameter study according to the same procedure as in the section before was performed with values $f_d = 0.0, 0.2, 0.4, 0.6, 1.0$. Therefore, again the same 20 samples with one million points were used. Here, a 4-means tree was investigated with a bucket size of $k^2 = 16$. The choice of $k = 4$ was made here since the query times of studies with $k = 2$ showed similar results, but the needed memory is lower. For higher values of k , the query times significantly increased. The results of these studies are depicted in Fig. 10a and Fig. 10b.

Interestingly, the query times for the exact cases with $f_d = 1.0$ are about one magnitude smaller compared to those computed with the k -d tree. In contrast to this, the query times for $f_d = 0.0$ remain approximately at the same speed. However, for the computations with $f_d = 0.2$ and $f_d = 0.4$ the query times are relatively low.

Concentrating on the accuracy, already for $f_d = 0.4$ an excellent behavior is seen. Even the results with $f_d = 0.0$ are quite accurate and show only a small variation.

A further remark regarding the comparability to open-source algorithms should be added. We performed the same tests with the implementations of the FLANN-library and found that the results were comparable. However,

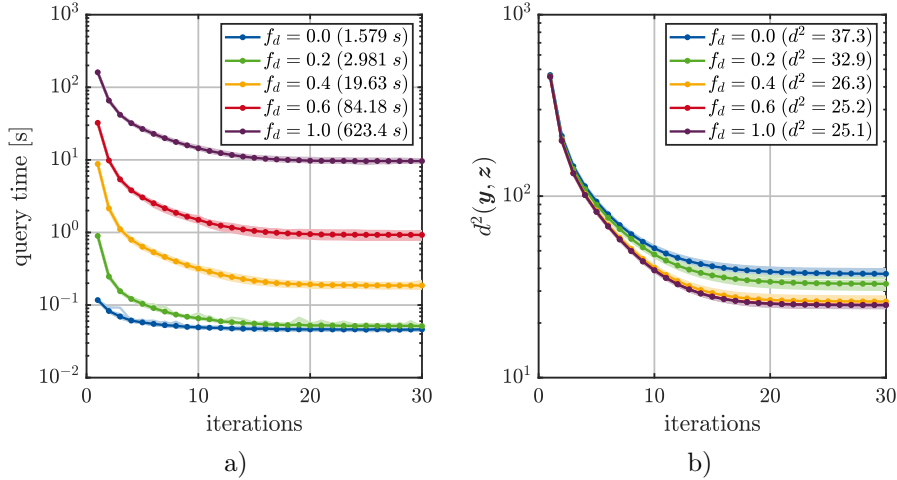


Figure 10: Elastic solid example: Approximate 4-means tree used for nearest-neighbor search on 20 one million points data sets. Results are studied for different parameters f_d . Averaged results are depicted in bold, all results in light colors. a) Query times per iteration over solver iterations. Total search times within the first 30 iterations stated in brackets. b) Global squared distance over solver iterations. Remaining squared distance after 30 iterations stated in brackets.

our implementations slightly dominated those results regarding query time and offered more insight. Besides, in the authors' opinion controlling accuracy is more intuitive in the way discussed above.

4.3. Approximate nearest-neighbors using k -nearest-neighbor graphs

Recently the category of graph searches has won much attention. For higher-dimensional data sets, this type of search algorithm seems to be superior. Therefore, the investigation of the latter is of the highest interest.

Here, we investigate the search in a k -NN graph initially introduced in [35]. In contrast to the tree search algorithms discussed before, graph search algorithms are not based on branch and bound algorithms. Instead, every point in the data set is connected with a couple of other points from the same set. Considering the k -NN graph, a network is built by connecting every node to its k nearest-neighbors. In this manner, a network is created as depicted in Fig. 11, where every point is connected with its $k = 3$ nearest-neighbors. The probably simplest strategy is then to choose a random starting point and use a greedy algorithm to find the nearest point. That means recursively, all distances between the neighbors of the current point and the query points are computed, and this point closest to the query point is selected as the next current. If the current point itself is the closest, the search will stop and the current point will be identified as the nearest-neighbor. For sure, this procedure will stop in a local minimum if the number of nearest-neighbors k in the network is set too




Figure 11: Graph structure of a small data set using a 3-nearest-neighbor graph (see enlargement). Query path from a starting point s_1 to the closest point p_1 of the query point q_1 is highlighted in red. Closest point p_1 is new starting point s_2 for query in second iteration.

small. This number k also has a large impact on the graph's properties like the sparsity, connectivity, or search time. Therefore, the accuracy can be controlled, e. g., by increasing or decreasing the number of nearest-neighbors or searching multiple times with different starting points.

This leads to the second central point regarding the combination of search algorithms and the data-driven solver. Since we observe convergence towards a solution with increasing iterations, the single query points $y^{(e)}$ will move less and less. This effect can be exemplarily seen in Fig. 1. Therefore, the knowledge about the nearest-neighbor of the previous iteration can be used. A simple approach is to select the previous solution as the starting point for the next iteration, as illustrated in Fig. 11. With ongoing iterations, the search path length should then decrease and thus also the query time.

The presented problem of a moving query point is also addressed, e. g., in [43–45]. Here, the problem of nearest-neighbor queries is applied to moving mobile devices in the 3D space. In [43], the authors prove that no additional search has to be performed if a query point moves less than a small distance δ . In our context, the relation can be estimated to

$$d(\mathbf{q}_i, \mathbf{q}_{i+1}) < \frac{d(\mathbf{q}_i, \mathbf{p}_{i,1}) - d(\mathbf{q}_i, \mathbf{p}_{i,2})}{2} = \delta, \quad (14)$$

where \mathbf{q}_i and \mathbf{q}_{i+1} are the query points in iteration i and $i + 1$, respectively. Further, $\mathbf{p}_{i,1}$ and $\mathbf{p}_{i,2}$ are the closest and second closest points, respectively, from the search in iteration i . Then, δ defines the maximum distance a query point

can move from one iteration to another so that no query has to be performed. We studied the described graph search algorithm’s behavior for its parameter k on the same 20 samples of the one million points data sets as used before. The neighborhood sizes were varied to $k = 10, 20, 50, 75, 100$. The graph structures were built using the approximate 4-means tree with $f_d = 0.6$, discussed in the previous section. The results of these studies are depicted in Fig. 12a and Fig. 12b.

Regarding the query times, a decrease of almost two amplitudes can be observed

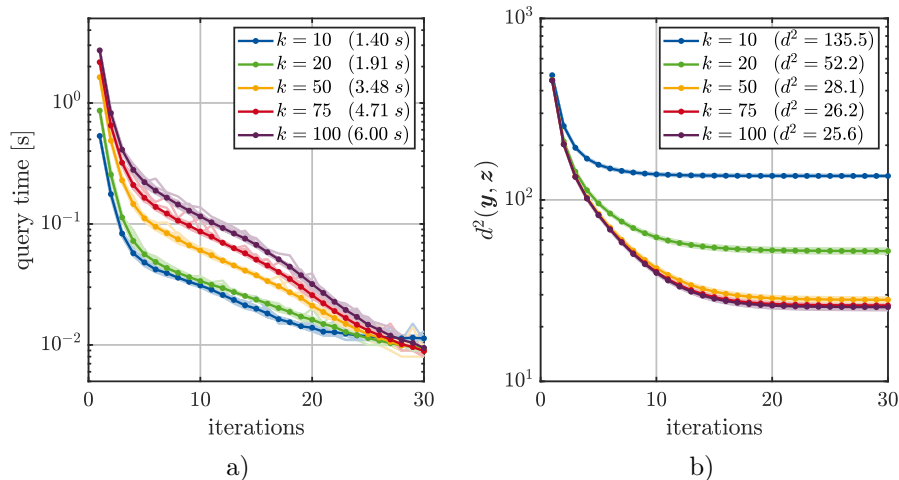


Figure 12: Elastic solid example: Approximate k -NN graph used for nearest-neighbor search on 20 one million points data sets. Results are studied for different parameters k . Averaged results are depicted in bold, all results in light colors. a) Query times per iteration over solver iterations. Total search times within the first 30 iterations stated in brackets. b) Global squared distance over solver iterations. Remaining squared distance after 30 iterations stated in brackets.

from iteration one to thirty for all parameter configurations. We further saw that the lower the neighborhood size k , the lower are also the needed search times. However, this goes along with a significant loss of accuracy, especially for $k = 10$ and $k = 20$. Increasing the neighborhood size from $k = 75$ to $k = 100$, almost no benefit regarding efficiency shows up.

The k -NN graph’s results discussed above are already quite promising, but still we see relatively long query times in the first iterations. Therefore, we performed an additional study where we limited the number of node changes to a bound f_s . Hereby, we aimed to reduce the number of comparisons in the first iterations, where the search paths are longest. The computations were performed with bounds $f_s = 10, 5, 3, 2, 1$.

In Fig. 13a, the query times for the first ten iterations are depicted. As intended, the query times in the first iteration can be reduced, whereas the query times in the following are slightly larger than before. In total, the query times can

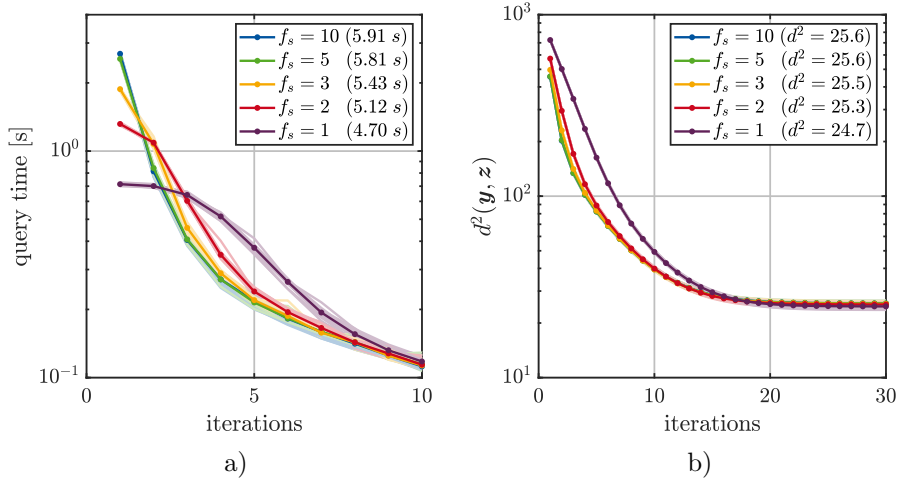


Figure 13: Elastic solid example: Approximate k -NN graph used for nearest-neighbor search on 20 one million points data sets. Results are studied for different parameters f_s . Averaged results are depicted in bold, all results in light colors. a) Query times per iteration over solver iterations. Total search times within the first 30 iterations stated in brackets. b) Global squared distance over solver iterations. Remaining squared distance after 30 iterations stated in brackets.

be reduced by more than 20% for $f_s = 1$ compared to $f_s = 10$. Interestingly, the remaining distances after 30 iterations for $f_s = 1$ slightly dominate the other computations (see. Fig. 13b). Further, the smaller f_s , the more gradual decreases the distance.

4.4. Comparisons of different algorithms on a 100 million points data set

The computations were repeated with the most promising parameter configuration on 100 million points data sets to compare the investigated search algorithms. Therefore, ten randomly sampled data sets were investigated. We considered the k -d tree with $f_d = 0$ and the k -means tree with $f_d = 0$ and $f_d = 0.3$. Additionally, we chose the k -NN graph with $k = 100$ and $f_s = 50$ as well as $k = 75$ and $f_s = 1$. The query times and squared distances over iterations are depicted in Fig. 14a and Fig. 14b. In Tab. 1 the total query times of the 30 iterations, as well as the remaining distances after 30 iterations, are summarized. In addition, the averaged times for building and the memory used for storing the data structure are remarked. Regarding the k -NN graph, the index was built using an approximate k -means tree with $f_d = 0.2$. Therefore, the graph structure is only an approximation of a k -NN graph as in the section above.

The results show that the query times for the k -d tree and the 4-means tree with $f_d = 0$ are the smallest with on average about 3 seconds. On the one hand, we observe that the k -d tree again shows rather bad results regarding the dis-




Figure 14: Elastic solid example: Comparison of different approximate nearest-neighbor algorithms used for nearest-neighbor search on 10 samples of 100 million points data sets. Averaged results are depicted in bold, all results in light colors. a) Query times per iteration over solver iterations. Total search times within the first 30 iterations stated in brackets. b) Global squared distance over solver iterations. Remaining squared distance after 30 iterations stated in brackets.

tances and accuracy. On the other hand, the accuracy of the 4-means tree with $f_d = 0$ is very promising. Regarding the final distances, only a small further improvement from 5.98 to 5.25 can be observed considering the results of the 4-means tree with $f_d = 0.3$. The smallest remaining distances are computed with the 4-means tree with $f_d = 0.3$, whereas the required query time is also the largest with 105.7 seconds on average. One further search structure, which seems to be quite efficient is the 75-NN graph with a maximum of two steps in one iteration. Compared to the 100-NN graph, the algorithm needs only half the time (13.2 seconds) and has mainly the same final squared distances. Fundamental differences can also be observed concentrating on the building time and needed memory capacity. On the one hand, the k -d tree uses only 111 seconds to build and 0.75 GB to store the index. On the other hand, the k -NN graph needed on average almost 17 hours for building and 37.3 GB for storing.

4.5. k -means tree on a 1 billion points data set

So far, we observed that the k -means tree shows excellent performance concerning query time and accuracy. Additionally, regarding building time and memory requirements, the tree also dominated the k -NN graph structure. Therefore, we performed a final study using 5 samples with one billion points. In contrast to the computations in Chapter 4.2, the control parameter f_d is not held constant but linearly increased within the first 20 iterations from zero to a final value of f_d^f . Computations were performed for $f_d^f = 0.0, 0.2, 0.4, 0.6$.

algorithm	parameters	total search time (s)	remaining sqrt. dist	building time (h)	index file size (GBytes)
<i>k</i> -d tree	$f_d = 0.0$	2.542	15.4	0.031	0.75
<i>k</i> -means tree	$k = 4$	3.245	5.98	0.888	1.5
<i>k</i> -means tree	$f_d = 0.0$ $k = 4$	105.7	5.25	0.888	1.5
<i>k</i> -NN graph	$f_d = 0.3$ $k = 100$	27.19	5.28	16.68	37.3
<i>k</i> -NN graph	$f_s = 50$ $k = 50$ $f_s = 2$	13.22	5.45	16.68	37.3

Table 1: Comparison of different indicators for selected search algorithms. Results are averaged over 10 samples.

The building time for these 4-means trees was on average 12.0 hours and took a memory of 15.5 GB to store a single index.

Fig. 15a and Fig. 15b show the query times and distances over iterations. In




Figure 15: Elastic solid example: Approximate 4-means tree used for nearest-neighbor search on 5 one billion points data sets. Results are studied for different parameters f_d^f . Averaged results are depicted in bold, all results in light colors. a) Query times per iteration over solver iterations. Total search times within the first 30 iterations stated in brackets. b) Global squared distance over solver iterations. Remaining squared distance after 30 iterations stated in brackets.

this study, the computations with $f_d^f = 0$ took on average 7.6 seconds. This is already a very impressive result considering that a single computation with the exact *k*-d tree would need approximately 300 days extrapolating from the

results obtained in chapter 3.1. That would mean a speedup of more than 10^6 . Nevertheless, a reduction of accuracy is observed compared to the results with $f_d^f = 0.4$ and $f_d^f = 0.6$. The latter show approximately the same remaining distances with 2.24 and 2.13. With a parameter of $f_d^f = 0.4$ we end up with approximately 33 seconds per computation.

5. Summary and concluding remarks

We have investigated three search structures for solving the nearest-neighbor problem in DD computing, namely, the k -d tree, the k -means tree, and the k -NN graph and shown how these search structures can be used for approximate nearest-neighbor searching. k -means search trees are found to be particularly well-suited for the task because the space partitioning is similar to a Voronoi tessellation. We find that, in the initial iterations of the DD solver, rough approximations of the nearest-neighbors are sufficient. As convergence is approached, the knowledge of the previous iteration can be used to accelerate the searches, e. g., using graph structures. We show that these strategies allow finite element simulations with data sets of up to billion points to be performed in the span of seconds.

Further improvements: With reference to graph search algorithms, we have observed that large benefits can be obtained by taking into account query results from previous iterations. The use of that information in the context of k -means trees naturally suggests itself. We have performed additional assessments of the satellite system graph [41] and the hierarchical navigable small world graph [38]. For the problems under consideration, we have not observed significant advantages, especially considering the additional effort required for building the indices. However, additional performance could be gained by means of parallel implementations or GPU computing.

Noisy data sets: We have limited our investigations to 3D elastic case problems with synthetic data sets free of scatter or noise. Such data sets might be generated, e. g., by simulations at a lower scale or by preprocessing or filtering measured data. However, experimental data sets are inevitably noisy and may contain outliers. In [2], a maximum-entropy solver was proposed that can deal effectively with noisy data and outliers. The resulting solvers are no longer minimum-distance solvers by aim to maximize a likelihood function. Nevertheless, the need to structure the data efficiently and perform fast searches persists. Given the maximum entropy characterization of the solutions, k -mean tree methods are likely to be best suited.

Connection to machine learning: The search data structures discussed in this paper may be regarded as a form of set-oriented unsupervised machine learning. Thus, in contrast to other forms of supervised machine learning, such as Artificial Neural Networks, here the outcome of the learning is not a functional relation, e. g., between stress and strain, but the data structures themselves. Such data structures reveal, or allow us to learn, how the data is organized. Unlike Artificial Neural Networks, the unsupervised learning afforded by the

search data structures is *lossless*, i. e., it entails no loss of information relative to the original data set. In [46], the connection between the k -NN problem and random decision forests [47] was already noted. The random decision forest, also known as random forest regression, is a machine learning technique based on multiple decision trees [48]. An application of such methods to DD computing immediately suggests itself. In a previous publication [9], the authors noted how other unsupervised machine learning techniques such as tensor voting [49] can be used to set up additional local data structures.

Data warehouse: Material data sets, such as generated by DDMI [10–12] can be exceedingly large and place onerous memory demands. This suggests centralization, or warehousing, of large material data sets in specialized servers for public access. In this data infrastructure, queries from a finite-element calculation performed on a local computer would be sent to the material data warehouse, which would perform near-neighbor searches and other operations on pre-structured data sets and send the results to the local user. Today, this procedure is already standard for many applications like web search or mobile navigation applications. By contrast, this form of data acquisition, management and exchange remains largely to be explored and developed in engineering applications.

Acknowledgments

MO gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft (DFG) through the Sonderforschungsbereich 1060 “The mathematics of emergent effects”. SR and RE gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft (DFG) through the project RE 1057/40-2 “Model order reduction in space and parameter dimension — towards damage-based modeling of polymorphic uncertainty in the context of robustness and reliability” within the priority program SPP 1886 “Polymorphic uncertainty modelling for the numerical design of structures”. Finally, all authors acknowledge the financial support of the DFG and French Agence Nationale de la Recherche (ANR) through the project “Direct Data-Driven Computational Mechanics for Anelastic Material Behaviours” (project numbers: ANR-19-CE46-0012-01, RE 1057/47-1) within the French-German Collaboration for Joint Projects in Natural, Life and Engineering (NLE) Sciences.

References

- [1] T. Kirchdoerfer, M. Ortiz, Data-driven computational mechanics, *Computer Methods in Applied Mechanics and Engineering* 304 (2016) 81–101.
- [2] T. Kirchdoerfer, M. Ortiz, Data driven computing with noisy material data sets, *Computer Methods in Applied Mechanics and Engineering* 326 (2017) 622–641.

- [3] T. Kirchdoerfer, M. Ortiz, Data-driven computing in dynamics, *International Journal for Numerical Methods in Engineering* 113 (11) (2018) 1697–1710.
- [4] R. Eggersmann, T. Kirchdoerfer, S. Reese, L. Stainier, M. Ortiz, Model-free data-driven inelasticity, *Computer Methods in Applied Mechanics and Engineering* 350 (2019) 81–99.
- [5] L. T. K. Nguyen, M.-A. Keip, A data-driven approach to nonlinear elasticity, *Computers & Structures* 194 (2018) 97–115.
- [6] A. Platzer, A. Leygue, L. Stainier, Assessment of data-driven computational mechanics in finite strain elasticity, in: *Constitutive Models for Rubber XI: Proceedings of the 11th European Conference on Constitutive Models for Rubber (ECCMR 2019)*, June 25-27, 2019, Nantes, France, CRC Press, 2019, p. 230.
- [7] S. Conti, S. Müller, M. Ortiz, Data-driven finite elasticity, *Archive for Rational Mechanics and Analysis* (2020) 1–33.
- [8] P. Carrara, L. De Lorenzis, L. Stainier, M. Ortiz, Data-driven fracture mechanics, *arXiv preprint arXiv:2006.03133* (2020).
- [9] R. Eggersmann, L. Stainier, M. Ortiz, S. Reese, Model-free data-driven computational mechanics enhanced by tensor voting, *arXiv preprint arXiv:2004.02503* (2020).
- [10] A. Leygue, M. Coret, J. Réthoré, L. Stainier, E. Verron, Data-based derivation of material response, *Computer Methods in Applied Mechanics and Engineering* 331 (2018) 184–196.
- [11] M. Dalémat, M. Coret, A. Leygue, E. Verron, Measuring stress field without constitutive equation, *Mechanics of Materials* (2019) 103087.
- [12] A. Leygue, R. Seghir, J. Réthoré, M. Coret, E. Verron, L. Stainier, Non-parametric material state field extraction from full field measurements, *Computational Mechanics* (2019) 1–9.
- [13] L. Stainier, A. Leygue, M. Ortiz, Model-free data-driven methods in mechanics: material data identification and solvers, *Computational Mechanics* (2019) 1–13.
- [14] A. H. Land, A. G. Doig, An automatic method for solving discrete programming problems, in: *50 Years of Integer Programming 1958-2008*, Springer, 2010, pp. 105–132.
- [15] J. H. Friedman, J. L. Bentley, R. A. Finkel, An algorithm for finding best matches in logarithmic expected time, *ACM Transactions on Mathematical Software (TOMS)* 3 (3) (1977) 209–226.

- [16] J. L. Bentley, Multidimensional binary search trees used for associative searching, *Communications of the ACM* 18 (9) (1975) 509–517.
- [17] K. Fukunaga, P. M. Narendra, A branch and bound algorithm for computing k-nearest neighbors, *IEEE transactions on computers* 100 (7) (1975) 750–753.
- [18] R. A. Finkel, J. L. Bentley, Quad trees a data structure for retrieval on composite keys, *Acta informatica* 4 (1) (1974) 1–9.
- [19] A. Guttman, R-trees: A dynamic index structure for spatial searching, in: *Proceedings of the 1984 ACM SIGMOD international conference on Management of data*, 1984, pp. 47–57.
- [20] J. K. Uhlmann, Satisfying general proximity/similarity queries with metric trees, *Information processing letters* 40 (4) (1991) 175–179.
- [21] P. N. Yianilos, Data structures and algorithms for nearest neighbor search in general metric spaces, in: *Soda*, Vol. 93, 1993, pp. 311–21.
- [22] P. Ciaccia, M. Patella, P. Zezula, M-tree: An efficient access method for similarity search in metric spaces, in: *Proceedings of the 23rd VLDB conference*, Athens, Greece, Citeseer, 1997, pp. 426–435.
- [23] S. M. Omohundro, *Five balltree construction algorithms*, International Computer Science Institute Berkeley, 1989.
- [24] R. F. Sproull, Refinements to nearest-neighbor searching ink-dimensional trees, *Algorithmica* 6 (1-6) (1991) 579–589.
- [25] L. Miclet, M. Dabouz, Approximative fast nearest-neighbour recognition, *Pattern Recognition Letters* 1 (5-6) (1983) 277–285.
- [26] S. Arya, D. M. Mount, Approximate nearest neighbor queries in fixed dimensions., in: *SODA*, Vol. 93, 1993, pp. 271–280.
- [27] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, A. Y. Wu, An optimal algorithm for approximate nearest neighbor searching fixed dimensions, *Journal of the ACM (JACM)* 45 (6) (1998) 891–923.
- [28] S. Brin, Near neighbor search in large metric spaces (1995).
- [29] C. Silpa-Anan, R. Hartley, Optimised kd-trees for fast image descriptor matching, in: *2008 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2008, pp. 1–8.
- [30] M. Muja, D. G. Lowe, Fast approximate nearest neighbors with automatic algorithm configuration., *VISAPP* (1) 2 (331-340) (2009) 2.

- [31] J. S. Beis, D. G. Lowe, Shape indexing using approximate nearest-neighbour search in high-dimensional spaces, in: Proceedings of IEEE computer society conference on computer vision and pattern recognition, IEEE, 1997, pp. 1000–1006.
- [32] M. Muja, D. Lowe, Flann-fast library for approximate nearest neighbors user manual, Computer Science Department, University of British Columbia, Vancouver, BC, Canada (2009).
- [33] M. Muja, D. G. Lowe, Scalable nearest neighbor algorithms for high dimensional data, *IEEE transactions on pattern analysis and machine intelligence* 36 (11) (2014) 2227–2240.
- [34] D. Dearholt, N. Gonzales, G. Kurup, Monotonic search networks for computer vision databases, in: Twenty-Second Asilomar Conference on Signals, Systems and Computers, Vol. 2, IEEE, 1988, pp. 548–553.
- [35] K. Hajebi, Y. Abbasi-Yadkori, H. Shahbazi, H. Zhang, Fast approximate nearest-neighbor search with k-nearest neighbor graph, in: Twenty-Second International Joint Conference on Artificial Intelligence, 2011.
- [36] W. Dong, C. Moses, K. Li, Efficient k-nearest neighbor graph construction for generic similarity measures, in: Proceedings of the 20th international conference on World wide web, 2011, pp. 577–586.
- [37] Y. Malkov, A. Ponomarenko, A. Logvinov, V. Krylov, Approximate nearest neighbor algorithm based on navigable small world graphs, *Information Systems* 45 (2014) 61–68.
- [38] Y. A. Malkov, D. A. Yashunin, Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs, *IEEE transactions on pattern analysis and machine intelligence* (2018).
- [39] W. Li, Y. Zhang, Y. Sun, W. Wang, M. Li, W. Zhang, X. Lin, Approximate nearest neighbor search on high dimensional data-experiments, analyses, and improvement, *IEEE Transactions on Knowledge and Data Engineering* (2019).
- [40] C. Fu, C. Xiang, C. Wang, D. Cai, Fast approximate nearest neighbor search with the navigating spreading-out graph, arXiv preprint arXiv:1707.00143 (2017).
- [41] C. Fu, C. Wang, D. Cai, Satellite system graph: Towards the efficiency up-boundary of graph-based approximate nearest neighbor search, arXiv preprint arXiv:1907.06146 (2019).
- [42] F. Groh, L. Ruppert, P. Wieschollek, H. Lensch, Ggnn: Graph-based gpu nearest neighbor search, arXiv preprint arXiv:1912.01059 (2019).

- [43] Z. Song, N. Roussopoulos, K-nearest neighbor search for moving query point, in: International Symposium on Spatial and Temporal Databases, Springer, 2001, pp. 79–96.
- [44] R. H. Güting, T. Behr, J. Xu, Efficient k-nearest neighbor search on moving object trajectories, *The VLDB Journal* 19 (5) (2010) 687–714.
- [45] Z. Yu, Y. Liu, X. Yu, K. Q. Pu, Scalable distributed processing of k nearest neighbor queries over moving objects, *IEEE Transactions on Knowledge and Data Engineering* 27 (5) (2014) 1383–1396.
- [46] Y. Lin, Y. Jeon, Random forests and adaptive nearest neighbors, *Journal of the American Statistical Association* 101 (474) (2006) 578–590.
- [47] T. K. Ho, Random decision forests, in: Proceedings of 3rd international conference on document analysis and recognition, Vol. 1, IEEE, 1995, pp. 278–282.
- [48] L. Breiman, J. Friedman, C. J. Stone, R. A. Olshen, Classification and regression trees, CRC press, 1984.
- [49] P. Mordohai, G. Medioni, Dimensionality estimation, manifold learning and function approximation using tensor voting., *Journal of Machine Learning Research* 11 (1) (2010).