



Original software publication

Particle-based and meshless methods with Aboria

Martin Robinson ^{a,*}, Maria Bruna ^b^a Department of Computer Science, University of Oxford, Wolfson Building, Parks Rd, Oxford OX1 3QD, United Kingdom^b Mathematical Institute, University of Oxford, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, United Kingdom

ARTICLE INFO

Article history:

Received 31 March 2017

Received in revised form 3 July 2017

Accepted 5 July 2017

Keywords:

Particle-based numerical methods

Meshless

C++

ABSTRACT

Aboria is a powerful and flexible C++ library for the implementation of particle-based numerical methods. The particles in such methods can represent actual particles (e.g. Molecular Dynamics) or abstract particles used to discretise a continuous function over a domain (e.g. Radial Basis Functions). Aboria provides a particle container, compatible with the Standard Template Library, spatial search data structures, and a Domain Specific Language to specify non-linear operators on the particle set. This paper gives an overview of Aboria's design, an example of use, and a performance benchmark.

© 2017 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Code metadata

Current code version	v0.4
Permanent link to code/repository used for this code version	http://github.com/ElsevierSoftwareX/SOFTX-D-17-00029
Legal Code Licence	BSD 3-Clause Licence
Code versioning system used	git
Software code languages, tools, and services used	C++
Compilation requirements, operating environments & dependencies	Tested on Ubuntu 14.04 LTS with the GCC compiler (version 5.4.1), and Clang compiler (version 3.8.0). Third-party library dependencies: Boost, Eigen (optional), VTK (optional)
If available Link to developer documentation/manual	https://martinrobinson.github.io/Aboria
Support email for questions	martin.robinson@cs.ox.ac.uk

1. Motivation and significance

Aboria is a C++ library that supports the implementation of particle-based numerical methods, which we define as having three key properties:

1. There is a set of N particles that have positions within a hypercube of d dimensions, with each dimension being either periodic or non-periodic.
2. The method can be described in terms of non-linear operators on the N particle positions and/or other variables associated to these particles.
3. These operators are defined solely by the particle positions and variables, and typically take the form of interactions

between pairs of closely spaced particles (i.e. neighbourhood interactions). There are no pre-defined connections between the particles.

This definition covers a wide variety of popular methods where particles are used to represent either physical particles or a discretisation of a continuous function. In the first category are methods such as Molecular Dynamics [1], Brownian and Langevin Dynamics [2]. The second category includes methods like Smoothed Particle Hydrodynamics (SPH) [3], Radial Basis Functions (RBF) for function interpolation and solution of Partial Differential Equations (PDEs) [4–6], and Gaussian Processes in Machine Learning [7].

To date, a large collection of software has been developed to implement these methods. Generally, each software package focuses on one or at most two methods. Molecular and Langevin Dynamics are well-served by packages such as GROMACS [8], LAMMPS [9], ESPResSo [10] or OpenMM [11]. SPHysics [12] is one of the best

* Corresponding author.

E-mail address: martin.robinson@cs.ox.ac.uk (M. Robinson).

known SPH solvers. There exists no large-scale package for RBF methods, but these can be implemented in [Matlab](#) [13], and are available as routines in packages such as [SciPy](#) [14].

The software listed above represents a considerable investment of time and money. The computational requirements of particle-based methods, such as the efficient calculation of interactions between particles, are challenging to implement in a way that scales well with the number of particles N , uniform and non-uniform particle distributions, different spatial dimensions and periodicity. Yet, to date there does not exist a general purpose library that can be used to implement these low-level routines, and so they are reimplemented again and again in each software package.

The situation with particle-based methods contrasts with mesh-based methods such as Finite Difference, Finite Volume or Finite Element Methods. These methods are supported by linear algebra libraries based on the [BLAS](#) [15] specifications, and today it would be very strange to implement a mesh-based method without using BLAS libraries. In general, particle-based methods cannot take advantage of linear algebra libraries, which are good for static nodes with connections that do not change during the simulation, as opposed to dynamic particles whose interaction lists are variable over the course of the simulation.

Aboria aims to replicate the success of linear algebra libraries by providing a general purpose library that can be used to support the implementation of particle-based numerical methods. The idea is to give the user complete control to define operators on the particle set, while implementing efficiently the difficult algorithmic aspects of particle-based methods, such as neighbourhood searches and fast summation algorithms. However, even at this level it is not a one-fits-all situation and Aboria is designed to allow users to choose specific algorithms that are best suited to the particular application. For example, calculating neighbourhood interactions for a uniform particle distribution is best done using a regular cell-list data structure, while for a highly non-uniform particle distribution a tree data structure like a k -d tree might be preferable [16]. For neighbourhood interactions that are zero beyond a certain radius, a radial search is the best algorithm to obtain interacting particle pairs, while for interactions that are amenable to analytical spherical expansions, the fast multipole method is an efficient fast summation algorithm [17].

1.1. Software capabilities

The [online documentation](#) provides a set of example programs to illustrate Aboria's capabilities. These include:

Molecular Dynamics Simulation of N Newtonian point particles interacting via a linear spring force.

Brownian Dynamics Simulation of N Brownian point particles moving through a set of fixed reflecting spheres that act as obstacles.

Discrete Element Model (DEM) Simulation of N granular particles with drag and gravity, interacting via a linear spring force with varying resting lengths.

Smoothed Particle Hydrodynamics (SPH) Simulation of a water column in hydrostatic equilibrium. SPH discretises the Navier–Stokes equations using radial interpolation kernels defined over a given particle set. See [18] for more details.

Radial Basis Function (RBF) Interpolation Computation of an RBF interpolant of a two-dimensional function using a multi-quadratic basis function. The linear algebra library [Eigen](#) [19] is used to solve the resulting set of linear equations.

Kansa Method for PDEs Application of RBFs (using Gaussian basis functions) to solve the Poisson equation in a square two dimensional domain with Dirichlet boundary conditions.

2. Software description

2.1. Software architecture

The high-level design of Aboria consists of three separate and complimentary abstraction levels (see [Fig. 1](#)). Aboria Level 1 contains basic data-structures that implement a particle set container, with associated spatial search capabilities. Level 2 contains efficient algorithms for particle-based methods, such as neighbour searches and fast summation methods. Level 3 implements a Domain Specific Language (DSL) for specifying nonlinear operators on the set of particles. While Levels 1 and 2 provide useful functionality for particle-based methods, the purpose of Level 3 is to tie together this functionality and to provide a easy-to-use interface that ensures that the capabilities of Levels 1 and 2 are used in the best possible way.

2.2. Software functionalities

2.2.1. Aboria Level 1

Aboria Level 1 implements a particle container class that holds the particle data ([Fig. 2](#)). This class is based on the Standard Template Library (STL) vector class (Level 0 in [Fig. 1](#)), which serves as the lowest-level data container. Three variables are stored by default: `position`, `id` and `alive`, representing respectively the particle's spatial position, unique identification number and a flag to indicate if the particle is marked for deletion. The user can specify the spatial dimension d ($d = 1$ in the example in [Fig. 2](#)), as well as any additional variables attached to each particle (an additional variable `velocity` is used in [Fig. 2](#)). The Level 1 particle set container will combine multiple Level 0 vectors to form a single data structure.

This particle set container generally follows the STL specification with its own iterators and traits (see [Fig. 2](#)). It supports operations to add particles (the STL `push_back` member function), remove particles (`erase`), and can return a single particle given an index i (`operator[]`). This index operation returns a lightweight type containing references to the corresponding index in the set of zipped Level 0 vectors. Individual variables can be obtained from this lightweight type via `get` functions provided by Aboria.

Level 1 also includes spatial search data structures, which can be used for fast neighbour searches throughout a hypercube domain with periodic or non-periodic boundaries. The particle set container interacts with the spatial search data structures to embed the particles within the domain, ensuring that the two data structures are kept synchronised while still allowing for updates to the particle positions.

The current version of the code implements only cell-list search data structures, which divide the domain into a square lattice of hypercube cells, each containing zero or more particles (see [Fig. 3](#)). Two cell-list implementations are provided, one which supports insertion of particles in parallel by reordering the particles in the set, and the other which only supports serial insertion. The latter is faster for serial use, and for when particles move rapidly relative to each other (e.g. Brownian dynamics). Both data structures support parallel queries. In future versions of Aboria a k -d tree data structure will also be added.

2.2.2. Aboria Level 2

Aboria Level 2 implements efficient search and fast summation algorithms using the particle set and spatial search data structures in Level 1. Currently Level 2 includes a box search algorithm around a given spatial position $\mathbf{r}_s$. This calculates which cell $\mathbf{r}_s$ is in and searches within that cell and its adjacent cells for potential neighbour particles within the given search box (see [Fig. 3](#)). This is useful for particle interactions that are zero beyond a certain radius,




Fig. 1. Overview of Aboria's design, showing the three different abstraction levels.




Fig. 2. Operations on an Aboria particle set with $d = 1$. The particle set is displayed as a grey box, with the default variables (position, id and alive) and a user-added variable velocity shown as horizontal vectors. The data associated to an individual particle correspond to a column in the set (shown as red dashed boxes). (Top left) A new particle p is added to the particle set `particles`. (Top right) The particle pointed to by the iterator i is erased from the particle set. (Bottom left) Indexing the particle set at index i returns a set of references to the values held at index i . (Bottom right) Individual references to variables can be obtained using `get` functions.




Fig. 3. Diagram showing the interaction of the two Level 1 data structures in two dimensions ($d = 2$). The particle set data structure (left) is kept synchronised with the spatial search data structure (right). The spatial search data structure is used to implement a box search algorithm in Level 2, which finds all the particles within a box region (red dashed line) around the query point (red point). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

such as the local force interactions of particles in the Molecular Dynamics method. Another usage is in the evaluation of compact kernels in the SPH method.

Currently we are implementing fast multipole methods for summation of global dense operators within Aboria, and this level will contain this and any other algorithms useful for particle-based methods.

2.2.3. Aboria Level 3

The highest abstraction level in Aboria implements a Domain Specific Language (DSL) for specifying non-linear operators on

the set of particles, using the [Boost.Proto](#) library. Users can use standard C++ operators (e.g. `*`, `+` or `/`) and a set of supplied functions (e.g. `sqrt`, `pow` or `norm`) to express the operator they wish to apply to the given particle set.

One key feature of the Aboria DSL is that the symbolic form of each operator is retained by Aboria, using C++ expression template techniques. Expression templates build up compile-time types that represent the symbolic form of expressions. For example, the expression `a*b + c` in C++, where `a`, `b` and `c` are of type `T`, could be represented using expression templates by the type `Add<Multiply<T,T>,T>`. In this case, the normal C++ operators do not calculate a specific value but instead produce a new type that encodes which operator was used.

Like most modern linear algebra libraries (e.g. [Eigen](#) [19] or [Blaze](#) [20]), Aboria uses expression templates to enable optimisations and error checking using the symbolic form of each operator. For example, Aboria can detect aliasing, which occurs when variables are being read and written to within the same operator, and can compensate by using a temporary buffer to store the operator results before updating the aliased variable. Aboria can detect when a Level 2 fast summation method is available, for example when the operator involves a summation over a local spatial range, or when the particle positions are altered, thus requiring an update of the spatial search data structure. As the symbolic form of the operator is known at compile-time, it can be efficiently inlined to all the low-level routines implemented in Level 1 and 2, ensuring zero overhead costs associated with this level.

Aboria's Level 3 DSL can be used to express vector operations on the variables within the particle set, or more complicated

non-linear operators that involve particle pair interactions, such as neighbourhood force interactions for the Molecular Dynamics method, or for the evaluation of RBFs.

3. Illustrative example

In this section we describe a complete step-by-step example of how the Aboria library can be used to implement a simple Molecular Dynamics simulation.

We consider N particles within a two-dimensional square domain with periodic boundary conditions, interacting via an exponential potential (with cutoff at r_{cut}). The force on particle i at position $\mathbf{r}_i$ due to particle j at position $\mathbf{r}_j$ is given by

$$\mathbf{f}_{ij} = \begin{cases} -c \exp(-\|\mathbf{dx}_{ij}\|) \frac{\mathbf{dx}_{ij}}{\|\mathbf{dx}_{ij}\|}, & \text{for } \|\mathbf{dx}_{ij}\| < r_{\text{cut}}, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where $\mathbf{dx}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ is the shortest vector between $\mathbf{r}_i$ and $\mathbf{r}_j$ and c is a constant. We use a semi-implicit Euler integrator with $\delta t = 1$ to evolve positions $\mathbf{r}_i$ and velocities $\mathbf{v}_i$ with accelerations $\mathbf{a}_i = \sum_j \mathbf{f}_{ij}$. This gives the following update equations for each timestep n

$$\mathbf{v}_i^{n+1} = \mathbf{v}_i^n + \sum_j \mathbf{f}_{ij}^n \quad (2)$$

$$\mathbf{r}_i^{n+1} = \mathbf{r}_i^n + \mathbf{v}_i^{n+1}. \quad (3)$$

This is implemented in Aboria as follows. First we define a new type `velocity` to refer to the two dimensional velocity variable $\mathbf{v}$ (using an Aboria `double2` type to store the 2D vector) as well as the particle set type, given by `container_type`, which contains the velocity variable and has a spatial dimension of $d = 2$ (specified by the second template argument). For convenience, we also define `position` as the `container_type::position` subclass (we will use this later on). Finally we create `particles`, an instance of `container_type`, containing N particles.

```
ABORIA_VARIABLE(velocity,double2,"velocity")
typedef Particles<std::tuple<velocity>,2> container_type;
typedef typename container_type::position position;

container_type particles(N);
```

Next we initialise the positions of the N particles to be uniformly distributed in the unit square (using the standard C++ random library), and initialise the velocities to zero.

```
std::uniform_real_distribution<double> uni(0,1);
std::default_random_engine g(seed);
for (int i = 0; i < N; ++i) {
    get<position>(particles)[i] = double2(uni(g),uni(g));
    get<velocity>(particles)[i] = double2(0,0);
}
```

We then initialise the Level 1 spatial search data structure, providing it with lower (0, 0) and upper (1, 1) bounds for the domain, and setting periodic boundary conditions. We set the width of the cells in the spatial data structure to r_{cut} .

```
particles.init_neighbour_search(double2(0,0),
                               double2(1,1),
                               r_cut,
                               bool2(true,true));
```

We now switch to using the Level 3 symbolic API to construct operators over the particle set we have created. We define two symbolic objects `p` and `v` representing the position and velocity variables. We also create two label objects `i` and `j` associated to the particle set `particles`. Note that we define two labels corresponding to indexes i and j symbols in Eq. (1), so that we can express the interaction force.

```
Symbol<position> p;
Symbol<velocity> v;
Label<0,container_type> i(particles);
Label<1,container_type> j(particles);
```

We also create `dx`, a symbolic object representing the shortest vector between particles i and j , and a symbolic accumulation operator `sum`, using the standard library `std::plus` structure to perform the accumulation.

```
auto dx = create_dx(i,j);
Accumulate<std::plus<double2> > sum;
```

Now we implement Eqs. (2) and (3). Combining a label with a symbol, for example `v[i]`, provides us with a concrete vector of variables, in this case all the velocity variables in the particle set. Similarly `p[i]` gives us all the position variables within the particle set.

Using expression templates, as described in Section 2.2.3, we can build an operator involving these objects, as well as the `dx` and `sum` objects described previously. The `sum` object acts as a summation operator over all the neighbours of particle i . It takes three arguments: (i) a label to accumulate over, (ii) a Boolean expression that is `true` for those particles that will be included in the summation, and (iii) an expression that returns the values to be added to the summation. Here we restrict the summation to all particles with $\|\mathbf{dx}_{ij}\| < r_{\text{cut}}$ and $\|\mathbf{dx}_{ij}\| > 0$, and sum the exponentially decaying force given in Eq. (1). We then accumulate this sum into `v[i]`, completing Eq. (2).

Listing 1: Velocity update

```
v[i] += c*sum(j, norm(dx)<r_cut && norm(dx)>0,
             -exp(-norm(dx))*dx/norm(dx));
```

Here, the `sum` operator detects that this is a summation incorporating only particles within a certain radius r_{cut} , and uses the Level 1 spatial search data structure and a Level 2 box search algorithm to speed up the summation. In the expression above, `norm` is a symbolic function provided by Aboria that returns the 2-norm of a vector.

Finally, we write Eq. (3) by incrementing the positions of the particles by the value of the velocity.

Listing 2: Position update

```
p[i] += v[i];
```

Inserting the individual timesteps written in Listings 1 and 2 in a time loop completes our implementation of the Molecular Dynamics example. The full source code is shown in Appendix A, and a visualisation movie of the positions and velocities of $N = 100$ particles over 1000 timesteps is shown in Appendix B.

To illustrate the efficiency of the Aboria implementation, we benchmark the most computationally demanding portion of the programme, the update operator for `v[i]` (Listing 1). A primary goal for Aboria was zero-cost abstraction: the data structures in Level 1 or the symbolic layer in Level 3 should not incur any additional overhead. take a Level 0 vector (i.e. a `std::vector`), and code by hand a given non-linear operator, this should run at the same speed as the Aboria equivalent.

We compare the Aboria operator given in Listing 1 with a hand-coded C++ version for a range of particle numbers N . Rather than re-implementing the neighbour search, we use the search provided in the [analysis tools](#) of the GROMACS library.¹ Other than the search facility, the comparison code uses standard C++

¹ Note that this search algorithm has not been optimised and is substantially different to the algorithm used by the GROMACS simulation engine.




Fig. 4. Comparison of the execution time T_e to calculate the velocity update (Eq. (2) and Listing 1) using Aboria versus a hand-coded C++/GROMACS implementation. The scaled inverse execution time T_e is shown versus the number of particles N (a higher value indicates better performance). (Left) Neighbourhood search cutoff is $r_{\text{cut}} = \sqrt{3/N}$ such that the number of neighbours is fixed. (Right) Cutoff is $r_{\text{cut}} = \sqrt{3/500}$ such that the number of neighbours increases with N .

(see Appendix B for the benchmark source code). The execution times T_e are computed as the average of $\lfloor 10^3/N \rfloor + 1$ evaluations of Listing 1, using a single Intel Core i5-6500T CPU at 2.50 GHz running Ubuntu 16.04 LTS and with 8 GB of RAM.

In Fig. 4 we plot N/T_e versus the number of particles N . In the left plot we set $r_{\text{cut}} = \sqrt{3/N}$ such that the number of neighbours is approximately constant in N (in two dimensions) and hence T_e is expected to scale linearly with N . In the right plot we choose a constant cutoff, $r_{\text{cut}} = \sqrt{3/500}$, so that the number of neighbours increases with N , and T_e is expected to grow by approximately N^2 . For very small $N < 10$ the C++/GROMACS version performs significantly better, as the GROMACS search facility turns off the neighbourhood search for small N and falls back to a more efficient brute-force N^2 method. This is not (yet) implemented in Aboria, which tries to do a spatial search, even for very small N . For $N > 10$, the two versions are comparable in performance and, for the left hand plot with $r_{\text{cut}} = \sqrt{3/N}$, exhibit the expected linear scaling (i.e. N/T_e constant). For this two dimensional benchmark, the Aboria version is noticeably faster than the C++/GROMACS version for $N > 10$, however, it should be noted that other benchmarks (described in the online documentation) between Aboria and the GROMACS search facility in three dimensions put the GROMACS version roughly equal to Aboria.

4. Impact

Aboria is designed to make the implementation of particle-based methods easy, particularly those involving particle pairwise interactions within local spatial neighbourhoods, in any number of dimensions. The design of Aboria was motivated by the lack of a general efficient and high quality library for particle-based methods – analogous to the linear algebra libraries available since the 1970s. With Aboria, new particle-based methods or adaptations to existing methods can be proposed and implemented quickly, without any sort of performance penalty. Thanks to its “zero-cost abstraction” design, Aboria can be used as a library within higher-level software packages, significantly reducing the implementation effort and encouraging the creation of new software that utilises particle-based methods.

Aboria has been used previously in a number of published studies. The core data structures and routines were used to implement a coupled SPH-DEM method to simulate bidisperse solid particles immersed in a fluid [21], as well as Smoluchowski Dynamics to simulate a set of particles interacting via chemical reactions, with coupling to a lattice-based Gillespie method [22]. More recently,

Aboria was used to simulate interacting elliptical particles in a molecular-scale liquid crystal model [23], diffusion through random porous media [24], and Brownian particles interacting via soft-sphere potentials [25].

Until now Aboria has been developed for the use within our research group, which has served as a proof of concept for a generic particle-based library. We are now expanding to reach other people interested in particle-based models, and other application areas. Our research focuses mainly on biological applications, in which the use of particle-based methods has become widespread in recent years (see e.g. [26]). Yet, we are also using Aboria in industrial applications projects involving heterogeneous materials where it is crucial to account for interactions at the microscale (e.g. membranes and batteries).

5. Conclusions

Particle-based numerical methods have already had a large impact on scientific progress, primarily through established methods such as Molecular Dynamics, and within the chemical physics and materials science communities. However, with recent developments in meshless methods, particularly the rise in popularity of Radial Basis Functions for interpolation and the solution of PDEs, or related developments in Gaussian Processes in Machine Learning, the landscape of different particle-based methods has widened significantly.

Aboria has been created to support the development of these types of numerical methods. We hope that this paper will encourage its wider use within the community, and that Aboria will have a positive impact by accelerating the research and development of particle-based methods, and their application to different fields in science and industry.

In the current version of Aboria, we have focussed on establishing the API and the three different abstraction levels. Future work will focus on improving performance on HPC platforms, and adding new spatial search data structures, starting with a k -d tree for non-uniform particle position distributions, and new fast summation algorithms, starting with the black-box fast multipole method [27].

Acknowledgements

This work was supported by EPSRC (grant EP/I017909/1), St John’s College Research Centre (grant 21138701) and the John Fell Fund (grant BLD10370). The authors wish to thank Professor Philip K. Maini for his guidance and support.

Appendix A. Code listing for Aboria example

```

void example(const size_t N,
             const unsigned seed,
             const bool write_out=false) {
    /*
     * set parameters
     */
    const int timesteps = 1e3;
    const double r_cut = std::sqrt(3.0/N);
    const double c = 1e-3;

    /*
     * Create a 2d particle container type with one
     * additional variable "velocity", represented
     * by a 2d double vector
     */
    ABORIA_VARIABLE(velocity,double2,"velocity")
    typedef Particles<std::tuple<velocity>,2> container_type;
    typedef typename container_type::position position;

    /*
     * create a particle set with size N
     */
    container_type particles(N);

    std::uniform_real_distribution<double> uni(0,1);
    std::default_random_engine gen(seed);
    for (int i = 0; i < N; ++i) {
        /*
         * set a random position, and initialise velocity
         */
        get<position>(particles)[i] = double2(uni(gen),uni(gen));
        get<velocity>(particles)[i] = double2(0,0);
    }

    /*
     * initiate neighbour search on a periodic 2d domain
     * of side length 1
     */
    particles.init_neighbour_search(double2(0,0),
                                   double2(1,1),
                                   r_cut,
                                   bool2(true,true));

    /*
     * create symbols and labels in order to use
     * the Level 3 API
     */
    Symbol<position> p;
    Symbol<velocity> v;
    Label<0,container_type> i(particles);
    Label<1,container_type> j(particles);

    /*
     * dx is a symbol representing the difference in
     * positions of particles i and j.
     */
    auto dx = create_dx(i,j);

    /*
     * sum is a symbolic function that sums
     * a sequence of 2d vectors
     */
    Accumulate<std::plus<double2> > sum;

    /*
     * perform timestepping
     */
    for (int io = 0; io < timesteps; ++io) {

```

```

/*
 * on every step write particle container to a vtk
 * unstructured grid file
 */
if (write_out) {
    vtkWriteGrid("aboria",io,particles.get_grid(true));
}

/*
 * leap frog integrator
 */
v[i] += c*sum(j, norm(dx)<r_cut && norm(dx)>0,
              -exp(-norm(dx))*dx/norm(dx)
            );
p[i] += v[i];
}
}

```

Appendix B. Supplementary data

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.softx.2017.07.002>.

References

- [1] Griebel M, Knapek S, Zumbusch G. Numerical simulation in molecular dynamics: numerics, algorithms, parallelization, applications, vol. 5. Springer Science & Business Media; 2007.
- [2] Lemons DS, Gythiel A. Paul Langevin's 1908 paper on the theory of Brownian motion [Sur la théorie du mouvement brownien, C. R. Acad. Sci. (Paris) 146, 530–533 (1908)]. Amer. J. Phys. 1997;65(11):1079–81.
- [3] Monaghan JJ. Smoothed particle hydrodynamics. Rep Progr Phys 2005;68(8):1703.
- [4] Hardy RL. Theory and applications of the multiquadric-biharmonic method 20 years of discovery 1968–1988. Comput. Math. Appl. 1990;19(8–9):163–208.
- [5] Buhmann MD. Radial basis functions: theory and implementations. Cambridge Monogr Appl Comput Math 2003;12:147–65.
- [6] Zhang X, Song KZ, Lu MW, Liu X. Meshless methods based on collocation with radial basis functions. Comput Mech 2000;26(4):333–43.
- [7] Rasmussen CE. Gaussian processes for machine learning. MIT Press; 2006.
- [8] Abraham MJ, Murtola T, Schulz R, Páll S, Smith JC, Hess B, et al. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. SoftwareX 2015;1:19–25.
- [9] Plimpton S. Fast parallel algorithms for short-range molecular dynamics. J. Comput. Phys. 1995;117(1):1–19.
- [10] Arnold A, Lenz O, Kesselheim S, Weeber R, Fahrenberger F, Roehm D, et al. Espresso 3.1: Molecular dynamics software for coarse-grained models. In: Meshfree methods for partial differential equations VI. Springer; 2013. p. 1–23.
- [11] Eastman P, Friedrichs MS, Chodera JD, Radmer RJ, Bruns CM, Ku JP, et al. Openmm 4: a reusable, extensible, hardware independent library for high performance molecular simulation. J. Chem. Theory Comput. 2012;9(1):461–9.
- [12] Gomez-Gesteira M, Rogers BD, Crespo AJ, Dalrymple RA, Narayanaswamy M, Dominguez JM. SPHysics—development of a free-surface fluid solver—Part 1: Theory and formulations. Comput Geosci 2012;48:289–99.
- [13] Fasshauer GE. Meshfree approximation methods with MATLAB, vol. 6. World Scientific; 2007.
- [14] Jones E, Oliphant T, Peterson P, et al. SciPy: Open source scientific tools for Python; 2001. <http://www.scipy.org/>. [Online; Accessed 20 June 2017].
- [15] BLAS Technical Forum. BLAS (Basic Linear Algebra Subprograms); 1979. <http://www.netlib.org/blas/>. [Online; Accessed 20 June 2017].
- [16] Bentley JL, Friedman JH. Data structures for range searching. ACM Comput. Surv. 1979;11(4):397–409.
- [17] Greengard L, Rokhlin V. A Fast algorithm for particle simulations. J. Comput. Phys. 1997;135(2):280–92.
- [18] Robinson M, Monaghan JJ. Direct numerical simulation of decaying two-dimensional turbulence in a no-slip square box using smoothed particle hydrodynamics. Internat. J. Numer. Methods Fluids 2012;70(1):37–55.
- [19] Jacob B, Guennebaud G. Eigen; 2008. <http://eigen.tuxfamily.org>. [Online; Accessed 20 June 2017].
- [20] Iglberger K, Hager G, Treibig J, Rude U. Expression templates revisited: a performance analysis of current methodologies. SIAM J. Sci. Comput. 2012;34(2):C42–69.
- [21] Robinson M, Ramaioli M, Luding S. Fluid-particle flow simulations using two-way-coupled mesoscale SPH-DEM and validation. Int. J. Multiph. Flow. 2014;59:121–34.
- [22] Robinson M, Flegg M, Erban R. Adaptive two-regime method: application to front propagation. J. Chem. Phys. 2014;140(12):124109.
- [23] Robinson M, Luo C, Farrell PE, Erban R, Majumdar A. From molecular to continuum modelling of bistable liquid crystal devices. Liq. Cryst. 2017;1–18.
- [24] Bruna M, Chapman SJ. Diffusion in spatially varying porous media. SIAM J. Appl. Math. 2015;75(4):1648–74.
- [25] Bruna M, Chapman SJ, Robinson M. Diffusion of particles with short-range interactions, SIAM J. Appl. Math. 2017 [submitted for publication] (arXiv:1703.09768).
- [26] Osborne JM, Fletcher AG, Pitt-Francis JM, Maini PK, Gavaghan DJ. Comparing individual-based approaches to modelling the self-organization of multicellular tissues. PLoS Comput. Biol. 2017;13(2):1–34.
- [27] Fong W, Darve E. The black-box fast multipole method. J. Comput. Phys. 2009;228(23):8712–25.