



Durham E-Theses

Colloidal Systems Confined to Curved Surfaces

LAW, JACK,OWEN

How to cite:

LAW, JACK,OWEN (2021) *Colloidal Systems Confined to Curved Surfaces*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/13864/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Colloidal Systems Confined to Curved Surfaces

Jack Owen Law



A thesis submitted for the degree of Doctor of Philosophy

Department of Physics

Durham University

UK

September 2020

Colloidal Systems Confined to Curved Surfaces

Jack Owen Law

Abstract

Surface curvature plays a vital role in many biological processes. Examples include the organising of proteins in cell membranes, the tiling of cells in epithelial layers and the growth of virus capsids. A major technological benefit of micro- and nanoscale curvature is that it can guide colloidal self-assembly, a property which is of major importance in fields such as drug-delivery, biosensor fabrication, and the development of meta-materials. Recent advances in fabrication techniques, such as 3D-printing, have made a rich library of geometries available to experimentalists and engineers seeking to realise these complex and fascinating systems.

Here, we use bespoke simulations and theoretical models to study the effects of surface curvature on two-dimensional systems of isotropically attractive colloids. We identify four important properties: the finite but boundary-free area of closed surfaces, the minimum perimeter of a patch with a given area, the difference between the Euclidean and geodesic separation of points on the surface, and the frustration of the hexagonal lattice in regions of non-zero Gaussian curvature. We also show that competition between these effects produces a range of novel behaviours.

Starting from the simplest example of a sphere, we show that surface curvature has a strong effect on the gas-liquid nucleation profile and the size of the critical nucleus, as well as destroying the equivalence of the canonical and grand canonical ensembles. Then, focussing on surfaces with non-uniform curvature, we use tori to demonstrate that the different thermodynamic phases are localised to specific regions of the surface, and the transitions between them involve the translation of the colloidal assemble. Finally, we investigate the cone, where there is no Gaussian curvature but the mean curvature varies. We find that surface curvature can stabilise chiral and achiral crystals, and the ground states depend on both the range of the potential and whether it acts through Euclidean space or along geodesics.

Contents

Acknowledgements	viii
Statement of Authorship	ix
Statement of Copyright	ix
Publications	x
1 Introduction	1
1.1 Colloids	2
1.1.1 Technological applications	3
1.1.2 Model particles	4
1.2 Colloids on spheres	6
1.2.1 Topology and defects on a sphere	7
1.2.2 Nucleation and growth	10
1.3 Non-spherical surfaces	12
1.3.1 Defining curvature	13
1.3.2 The generalised Thomson problem	14
1.3.3 Phase transitions	15
1.3.4 Complex interactions	16
1.4 Thesis objectives and outline	17
2 Methods	19
2.1 Introduction	19
2.2 Monte Carlo simulation	20
2.2.1 Statistical basis of Monte Carlo	20
2.2.2 Monte Carlo simulation	22

2.2.3	Rejection sampling	24
2.2.4	Metropolis Monte Carlo	25
2.2.5	Umbrella sampling	29
2.2.6	Monte Carlo in the grand canonical ensemble	31
2.2.7	Monte Carlo simulations on curved surfaces	32
2.3	Molecular dynamics	34
2.3.1	Molecular dynamics fundamentals	34
2.3.2	Molecular dynamics on curved surfaces	36
2.4	Basin hopping	37
2.4.1	The Basin hopping method	37
2.4.2	Basin hopping on curved surfaces	39
3	Nucleation on a sphere	
	The roles of curvature, confinement and ensemble	40
3.1	Introduction	40
3.2	Model and methods	42
3.2.1	Model system	43
3.2.2	Monte Carlo implementation	43
3.2.3	Multiple histogram reweighting	45
3.2.4	Measuring free energy barriers	46
3.3	Results in the grand canonical ensemble	47
3.3.1	Phase diagram	47
3.3.2	Nucleation	48
3.3.3	Critical nuclei	52
3.4	Results in the canonical ensemble	55
3.5	Conclusion	60
4	Localised states on surfaces of non-constant curvature	63
4.1	Introduction	63
4.2	Methods	65
4.2.1	Model System	65
4.2.2	Simulation techniques	67
4.3	Basic phenomenology	68

4.3.1	Localised states on a torus	69
4.3.2	Basin Hopping	70
4.3.3	Free energy and dynamics	71
4.4	Bond Order Histograms	72
4.5	Quantitative analysis	75
4.5.1	The L- to C+ transition	75
4.5.2	The C+ to X0 transition	76
4.6	Varying the number of particles and geometry	81
4.6.1	100 and 500 particles	81
4.6.2	A slender torus	85
4.6.3	Sinusoidal surface	85
4.7	Conclusion	86
5	Crystal structures on conical surfaces	88
5.1	Introduction	88
5.2	Geometry of a cone	92
5.3	Global optimisation on the cone	94
5.4	Minimum energy structures	96
5.5	Measuring chirality	101
5.6	Characterising crystals on cones	104
5.6.1	Hexagonal lattices on conical nets	105
5.6.2	Cone conformation criteria: Magic angles	108
5.6.3	Cone conformation criteria: Intermediate angles	108
5.7	Comparing through-space and geodesic potentials	110
5.8	Conclusions	115
6	Conclusions and outlook	116
6.1	Overview	116
6.2	Effects of curvature	116
6.2.1	Finite extent of the surface	116
6.2.2	Distance contraction	117
6.2.3	Line energy	118
6.2.4	Crystal frustration	119

6.2.5	Interplay between the effects	119
6.3	Future work	120
6.3.1	Potential applications	120
6.3.2	Exploring potential energy landscapes	120
6.3.3	Simulating flexible membranes	121
6.3.4	Non-spherical and polydisperse particles	122
	Bibliography	122

Acknowledgements

I would like to thank my supervisors Dr Halim Kusumaatmaja and Dr Mark Miller for their constant support, guidance and teaching over the last four years. Thanks also to Jacob Dean, Alex Wong and James Wolstencroft, who worked with me as Master's students. I am grateful to Prof. Colin Bain, Dr Mike Evans, Dr Johan Mattsson, Dr Stefan Paquay and Dr Patrick Warren for valuable discussions, and again to Stefan for helping me get to grips with LAMMPS.

I would also like to thank Prof. Vinothan Manoharan and his group (in particular Talha Rehman, Dr Nabila Tanjeem and Jessica Sun) at Harvard for hosting me for three months and continuing to support the cones project with fruitful discussions.

Special thanks go to everyone in the physics lunch club for making the last four years so enjoyable, and to SOFI cohort 2 for making the first six months endlessly entertaining. Finally, to my parents, who got me so far and then let me spend a year at home writing up my thesis and sheltering from the pandemic.

Statement of Authorship

I certify that all work described in this thesis is my own research unless otherwise acknowledged in the text or by references and has not been previously submitted for a degree in this or any other university.

Statement of Copyright

The copyright of this thesis rests with the author. No quotation from it should be published without the author's prior written consent and information derived from it should be acknowledged.

Publications

The following published works form the majority of Chapter 3 and chapter 4:

- J. O. Law, A. G. Wong, H. Kusumaatmaja, and M. A. Miller, “Nucleation on sphere: the roles of curvature, confinement and ensemble,” *Mol. Phys.*, vol 116, p 3008, 2018.
- J. O. Law, J. M. Dean, M. A. Miller, and H. Kusumaatmaja, “Phase transitions on non-uniform curved surfaces: Coupling between phase and location,” *Soft Matter*, vol 16, p 8069, 2020.

The author also made a significant contribution to:

- B. Guiselin, J. O. Law, B. Chakrabarti, and H. Kusumaatmaja, “Dynamic morphologies and stability of droplet interface bilayers,” *Phys. Rev. Lett.*, vol 230 p 238001, 2018.

Chapter 1

Introduction

Two-dimensional systems play a wide range of important roles in nature, as well as in engineering and industrial applications. There is now abundant evidence showing that the curvature of the surface on which such a system is confined can strongly influence the observed physical behaviours. As a simple example, old-style black and white footballs had to incorporate twelve pentagonal patches to cover the surface area and satisfy the topology of the ball. On a flat surface, the space can be tiled with only hexagons. The same topological constraint is observed in the self-assembly of virus capsids [1]. More complex and physically interesting examples come from the world of biology, including the growth of HIV-1 viruses, cellular endocytosis, and the organisation of lipids in bilayer membranes [1–3]. The importance of curvature has also received significant attention in materials science. For example, it has been reported that curvature affects not only the mechanical properties of graphene, but also its prized electronic properties [4–6].

In this thesis, I will focus on modelling colloidal systems embedded on curved surfaces. They are an excellent model system to isolate and quantify the effects of surface curvature. On one hand, to a good approximation, colloids interact classically with pairwise interactions, which make them amenable to computer simulations with several hundreds of particles [7, 8]. On the other hand, rapidly advancing laboratory techniques are allowing for the development of colloidal particles with complex shapes and interactions [7, 9–12] confined to ever more elaborately curved surfaces [13–16].

Indeed, numerous studies have been dedicated to understanding the behaviour of colloids on curved surfaces. However, they are primarily focused on ground states and simple geometries (e.g. spheres) [1, 17–20]. In contrast, here, I will provide novel insights into ther-

modynamics on curved surfaces, the interplay of different curvature effects on geometries with non-uniform curvature, the formation of chiral structures, and key differences between interaction potentials that act directly through Euclidean space and along the surface geodesics.

For the rest of the chapter, I will briefly introduce colloidal systems before discussing the state of the art of colloids on spherical and non-spherical curved surfaces. I will finish by presenting the main questions addressed in this thesis.

1.1 Colloids

Colloids are collections of micro- or nano-scale particles, usually suspended in a solvent. At this length scale, thermal fluctuations and Brownian motion play a large role but gravity can often be ignored, making colloids an example of soft matter. We are surrounded by colloidal systems every day. They are found in household items such as food and paint [21], and inside our bodies in the form of intracellular vesicles [2, 22], cells themselves [23, 24], and protein aggregates such as virus capsids [1, 19, 25, 26]. As our understanding grows, designer colloids are swiftly becoming a core component of many cutting edge technologies, including for drug and flavour delivery, photonic crystals, displays and biosensors [27–31].

Colloids also have two particularly advantageous properties that make them popular for scientific studies. The first is the size of the particles. Since the size of the colloidal particles is in the same range as the wavelength of visible light, they can be imaged by conventional microscopy in real time [7, 32, 33] and individually tracked by modern image processing tools [33, 34]. The second is that they can be manufactured with an immense range of properties, giving them the role of “designer atoms” [35–38]. They can be tuned to interact in a variety of ways, from simple hard-core repulsion to isotropic depletion or electrostatic forces [7, 11], to intricate anisotropic interactions [7, 36, 39]. Another way to manipulate the interactions of colloids is to shape the particles themselves. To provide some examples, colloidal particles with ellipsoidal, dumbbell, lock-and-key and cubic shapes have been produced [12, 40–42]. Thanks to recent advances, this can be achieved using a variety of methods including 3D printing [12, 16], seed polymerisation [29], and by buckling spherical particles [7].

1.1.1 Technological applications

The advantageous properties of colloids have been exploited for a number of important technological applications. In particular, one aspect that has proved transformative is the ability of colloidal particles to self-assemble. The standard approaches to making nanostructured materials, such as for nanoscale electronics, catalysis and soft tissue engineering, are with top-down methods like optical lithography, laser induced synthesis and evaporation [43–48]. Colloids offer the opportunity to switch to bottom-up manufacturing and design [29, 48–50]. Individual colloidal particles are treated as highly tuned building-blocks, which self-assemble into novel materials [28, 29, 48, 51]. This approach has the potential for easier fabrication and finer structural control [29]. Below, we describe two of the most significant colloidal technologies, photonic crystals and colloidosomes.

A leading goal in photonics is the development of crystals with a diamond-like structure and lattice spacings of the order of the wavelength of visible light. This is because such a system has a unique “photonic band gap” which could be used to make filters and resonant cavities of high quality [52–55]. This goal is difficult to achieve, as the diamond structure is open and mechanically flexible, rendering it unstable relative to denser crystal structures [28, 56]. Excitingly, recent progress has been made toward using colloids to create structures with diamond symmetries on the appropriate length scale. We describe two example approaches. In the first, the diamond structure is interlaced with a second crystal made of smaller colloidal particles [28]. This second lattice stabilises the growing crystal, and it is anticipated that by carefully removing the smaller colloids after assembly, optically active colloidal crystals could be produced [28]. The second method uses pre-assembled colloidal tetrahedra, interacting via DNA coatings to produce a range of exotic crystal structures including an open diamond lattice [57]. If the procedure could be replicated using colloids with appropriate material and optical properties, it has the potential to be adapted for a large number of photonic applications [57].

Another consequence of the similarity of length scales in colloidal systems and visible light is that some of these systems display structural colour. This occurs when the interaction of light with the structures in the colloidal systems favours constructive interference for a narrow window of wavelengths. Examples of colours produced by photonic crystal structures in nature include opals [58], the wings of butterflies, including the emerald-patched cattle-heart [59, 60] and the exoskeleton of certain insects [61]. Structural colours have also generated

interest for applications in colour printing and biometric recognition [29–31], among others.

Highly controlled drug delivery can be achieved by encapsulating the active molecules in a strong, permeable coating. This can be realised using colloidosomes: liquid droplets coated with colloidal particles [51,62]. These are generally manufactured by creating an oil-in-water or water-in-oil emulsion, with the droplets stabilised by colloidal particles adsorbed to the surface. Then, a wide variety of methods are used to stabilise the colloidal shell, so that the droplet template can be removed. These methods involve either binding the colloidal particles to their neighbours, e.g. through sintering or cross-linking [51,62–65], or trapping them in another phase such as a gel or a polymer [51,66,67]. These systems can be targeted to release a cargo at certain times or under certain mechanical conditions [51,68,69]. They can even cross the blood-brain barrier, and can be used for safe and effective delivery of molecules to the brain [27].

1.1.2 Model particles

In addition to their use in the design of novel materials, colloids have found use as model systems in numerous scientific areas, owing to their accessibility. For instance, colloidal systems are a rich ground to validate prevailing theories in statistical physics. This is possible because they can be made to closely replicate the simple interactions used in theoretical models, in particular hard-core repulsion. Some of the earliest computer simulations suggested that hard spheres can crystallise [8]. This is because in the crystal phase particles have more space to explore than in a liquid of the same density. Hence, the entropy of the crystal phase is higher than in the liquid phase [70]. This surprising result was confirmed for the first time using spherical colloids made of polymethylmethacrylate (PMMA) settling under gravity [71].

Another important statistical theory demonstrated using colloidal particles with hard-core repulsion is the two-step melting process for 2D crystals proposed by Kosterlitz, Thouless, Helperin, Nelson and Young (KTHNY theory) [72–75]. This theory predicts that a flat hexagonal crystal melts via a hexatic phase. The first phase transition in the melting process involves the appearance of dislocation defects, which disrupt the local positional order but not orientational order. Subsequently, the hexatic-liquid transition occurs through the unbinding of disclination pairs, which does destroy orientational order. This two-step melting has been observed in colloidal systems by analysing microscopy images of colloids trapped between glass plates at varying density [76–80].

A case in which the tunability of the interaction between colloidal particles is of great benefit is the study of glassy systems [35, 81]. A dense-packed assembly of hard-spheres can be used as a model for metallic glass [82]. Computer simulations have suggested a link between defect structures in densely-packed particles on negatively curved surfaces and icosahedral glasses in three-dimensional space [82, 83]. Colloidal experiments have shown that, counter-intuitively, adding a weak attraction between the particles can cause the glass to break down, while a stronger attraction causes a new glass to form [81]. Other experiments have shown that the dynamics of repulsive-sphere glasses are not completely frozen at long timescales. The timescale of their slow dynamics depends on the interactions between colloids. Glasses of colloids which interact electrostatically evolve on much shorter timescales than hard spheres [84, 85].

Colloidal particles can also be used as stand-ins for bioproteins when studying biological self-assembly. For example, the genetic material of a virus is protected by a hard capsid formed of a small number of protein monomers that self-assemble when the virus is created. The final step of its maturation is to exit via the cell membrane. The virus binds to the inside of the cell membrane, then becomes engulfed before budding off into the intracellular matrix [25, 26]. This passive process can be modelled using simple colloids in the place of the virus particle and a lipid membrane to stand in for the cell. These experiments have shown that the exit can be explained by the balance of the adhesion energy between the virus and the membrane and the elastic energy of the bending membrane [86].

In cells, membrane-binding particles are also used to trigger beneficial remodelling processes like tubulation and vesiculation [87–89]. Proteins in the membrane must act collectively to perform these tasks, and it has been suggested that this is aided by curvature mediated attraction [87, 90, 91]. This can be modelled using colloids in lipid vesicles [92]. When a particle is embedded in the membrane, it causes the membrane to distort [26]. As two particles close together cause less total distortion than two particles far apart, and distortions cost elastic energy, the particles feel an attraction towards one another [92, 93]. This mechanism could partially explain the aggregation of these membrane proteins [87, 92].

In materials science, colloids have also proved useful for generating new insights, by using colloidal particles as model atoms. Nucleation is one of the most important, and most disputed, phenomena in materials physics [94–96]. One of the main difficulties studying it experimentally at the molecular level is that the process is very fast, rare, and occupies very

short length scales [97]. Using colloidal particles allows nucleation to be imaged directly in real time, so the structure of the developing nucleus can be observed [97,98]. The most well used theory of nucleation is classical nucleation theory, which assumes that the structure of the nucleus is the same as the bulk phase, but this has been challenged [95,96,99]. Two-dimensional experiments with colloids have shown that the nucleus of the crystal phase has a liquid-like structure, which switches to a hexagonal crystal when the nucleus reaches a certain size, determined by the supersaturation [98]. As another example, nanoindentation is the process of pushing a sharp tip into the surface of a crystalline material and using the response to deduce the mechanical properties of the material [100,101]. Nanoindentation also nucleates defects in the crystal, and hence it can be used to measure material strength and failure processes [102,103]. To understand the defect patterns caused by indentation, colloidal analogues, such a crystal of silica spheres, have been used to monitor and trace the motion of the particles around the indenting tip [103,104].

1.2 Colloids on spheres

The most common surface geometry that has been employed to study the effects of curvature is the sphere. This field has a long history, with J. J. Thomson proposing the plum pudding model in 1904 to describe the distribution of electrons in an atom [105]. This amounts to finding the arrangement of point charges on the surface of a sphere with the lowest electrostatic energy. While this model was soon abandoned in that context, studying particles confined to spherical surfaces is highly useful for colloidal science. Examples in nature include the HIV-1 virus capsid, spherical bacterial surface layers, metazoan epithelia and the silicone skeletons of spherical radiolaria [1,17,106–111]. More broadly, the competition between local order for neighbouring particles and long-range constraints due to the curvature and geometry is also relevant for fullerene patterns of carbon clusters [112–115] and the surface ordering of electrons on liquid helium bubbles confined in Paul traps [116], among many others.

Experimentally, the spherical surface is also one of the easiest confining surfaces to realise. The simplest route is to employ Pickering emulsion [17,117]. Here, a water-in-oil emulsion droplet is stabilised by colloidal particles, such as polystyrene beads, adsorbed irreversibly over the entire surface [17,62]. Commonly, a short-range attractive interaction between the colloidal particles is then introduced using depletants [11,19,118,119], though other types

of interactions have also been used, such as screened Coulomb interactions [120–122], and capillary forces [123, 124].

This method, however, has a number of limitations. It is difficult to control the density of colloidal particles on the surface or eliminate the Coulomb and capillary interactions if they are unwanted [11, 125]. One way to solve this is by tethering colloidal spheres to a lipid-stabilised oil drop with DNA. This way, the colloids are confined by the surface and free to diffuse across it, but they are not physically in contact with the surface. The colloidal loading can be controlled by moderating the density of DNA tethers [11, 125].

1.2.1 Topology and defects on a sphere

In addition to its many applications in physics, chemistry, materials science and biophysics, the Thomson problem has also gained special attention in the mathematics and computer science communities because it appears in the 7th problem in Steven Smale’s list of eighteen unsolved problems for the 21st century [105, 126]. The Thomson problem is challenging to approach analytically, and to date, solutions have only been found for a handful of particle numbers, all below 12 [20]. This is in contrast to the ground state of particles on a plane, where the solution is simple: the particles arrange on a hexagonal lattice. This is, however, not possible on a sphere due to constraints imposed by the curvature and geometry.

According to Euler’s polyhedron formula, all convex polyhedra must have a characteristic χ of 2 [127, 128], where $\chi = V - E + F$, V is the number of vertices, E is the number of edges and F is the number of faces. A perfect hexagonal crystal cannot satisfy this criterion. Every hexagonal face has 6 vertices, each shared by 3 faces, so $V = 2F$. Each face has 6 edges, shared by 2, so $E = 3F$. In total $\chi = 0$, which does not match the Euler characteristic of a convex polyhedron. Thus, all convex polyhedra must contain non-hexagonal faces. The simplest solution is to include twelve pentagonal faces among the hexagons. These have 5 vertices and 5 edges so they contribute $12(5/3 - 5/2 + 1) = 2$ to Euler’s formula, satisfying Euler’s rule [127, 128].

If we consider each face as one cell of the Voronoi tessellation of a crystal on a sphere, we can see that the lattice must contain defects: particles with a number of nearest neighbours different from 6. Defects can be characterised by their topological charge. Any particle with n nearest neighbours has a charge of $6-n$. For example, a particle with 5 neighbours has a topological charge of $+1$, while a particle with 7 neighbours has a topological charge of -1 .

The general solution to Euler’s rule for a sphere is that the total topological charge on the surface must be $+12$ [17].

There is another important reason for the appearance of defects in ground state crystals embedded on curved surfaces. Unlike on a plane, parallel lines do not remain equidistant on curved manifolds. This introduces frustration in the crystalline lattice, which can be relaxed not just by isolated $+1$ or -1 defects (henceforth, termed disclinations), but also by the inclusion of lines of defects (henceforth, termed dislocations). Neutral dislocation lines with no excess $+1$ or -1 defects are called pleats, while charged dislocation lines with excess defects are called scars [17, 129]. These dislocation chains, also called grain boundaries, are found in planar crystals, but they always terminate at the edge of the sample, due to the prohibitive energy cost of terminating in the bulk. Uniquely, on a sphere the scars terminate freely in the bulk of the crystal [17].

As analytical solutions to the Thomson problem are elusive, much progress has been achieved using both computer simulations and experimental studies using colloids. For instance, Wales *et al.* [130] have tabulated putative global minima for up to $N = 972$ particles using the basin hopping algorithm, see for example Fig. 1.1A-B. Simulations using continuum elastic theory have been used to identify minima for up to $N \approx 2000$ particles, as depicted in Fig. 1.1C-D [131]. While isolated defects dominate the particle configurations for small N , defect lines and other more complex defect structures proliferate for large N . In all cases, scars are the preferred defect structure above $N = 400$, and pleats are not reported [130, 131]. To satisfy Euler’s rule, each scar has a net topological charge of $+1$ and there are twelve of them [132]. The scars are expected to follow geodesics and to arrange themselves on the vertices of an icosahedron [122, 131, 133, 134]. Alternatively, multiple scars can share a terminal $+1$ defect, resulting in “buttons” [131], as in Fig. 1.1D.

Similar observations are also found experimentally, as shown in Fig. 1.1. In panels E and F, only $+1$ dislocations are seen on a small droplet (around 60 particles), while for a sufficiently large number of particles (around 400 particles [132]), the dominant defect structure consists of 12 scars [17, 132, 132, 135] (see panels G and H), in agreement with the simulations described above.

More recently, interest in the Thomson problem has expanded to cases at finite temperatures, and it is clear that the defect dynamics are a rich ground for investigations. For example, it was found that thermal fluctuations allow for kinked scars and defect diffusion [135].




Figure 1.1: A-B: Putative ground states obtained from basin hopping for $N = 44$ (A) and $N = 812$ (B) particles showing $+1$ (red) and -1 (blue) disclinations arranged individually (A) and in scars (B). Adapted from [130] with permission. C-D: Putative minima for $N \approx 2000$ particles obtained from continuum elastic theory, showing $+1$ (red) and -1 (blue) disclinations. One minimum shows isolated scars (A) and the other buttons (B). Adapted from [131] with permission. E-F: Light microscopy images of a droplet coated with about 60 particles. F is overlaid with the Delaunay triangulation, showing a $+1$ disclination in red. G-H: Light microscopy images of a droplet coated with about 600 particles. H is overlaid with the Delaunay triangulation, showing $+1$ (red) and -1 (yellow) disclinations arranged in scars. E-F adapted from [17] with permission.

As motion perpendicular to the line of the scar only requires a local rearrangement of the lattice, this so-called “glide” occurs very readily. Motion along the axis of the scar, called “climb” has a much lower diffusion constant [135]. The particles themselves diffuse slowly far from the defects but are more mobile close to scars [122].

At a larger scale, beyond the individual defect dynamics, the freezing transition on the sphere has been systematically studied. On a plane, freezing proceeds by the rapid elimination of lattice defects over a small range of temperatures [72–75], leading to the emergence of long-ranged orientational order and quasi-long-ranged positional order. In contrast, on a sphere there are defects in the ground state, and orientational order is hard to define as lattice lines on the surface are curved. In addition, long-ranged order and quasi-long-ranged order are difficult to separate due to the limited extent of the surface. This freezing problem was tackled using a system of charged microspheres adsorbed onto an oil droplet by electrostatic image charge forces [122]. Guerra *et al.* found the reduction in the defect number on a sphere does occur but over a wider range of temperatures than on the plane. The surviving defects become concentrated at the vertices of an icosahedron [122, 134]. The emergence of long-ranged order can also be established by projecting particles onto the surface of this icosahedron and then unfolding the resulting net [122].

1.2.2 Nucleation and growth

According to classical nucleation theory, the free energy of a growing liquid nucleus in two dimensions is determined by the competition between two contributions: the decrease in energy associated with the growth of the stable phase, and the line energy due to the interface between the two phases. The nucleus will always minimise its perimeter in order to reduce the line energy. On a plane, this is achieved with a circular nucleus, where the line energy increases with the square-root of the area. On the sphere, the minimal surface is a spherical cap, with a different perimeter-area relationship, which has been shown to strongly influence the size of the critical nucleus [18].

If the growing phase is a crystal rather than a liquid, there is a third contribution to the free energy: the frustration of the lattice due to curvature. When the nucleus is small, the total curvature covered by the crystal is low enough that it can grow as a slightly strained spherical cap [1, 18, 136]. However, once the nucleus reaches a critical size, the stress dominates, and isotropic growth is no longer energetically favourable [1, 18, 136]. From here, the system




Figure 1.2: A: A sphere with 31 particles having undergone “hole implosion”. No particle can be added to the surface, although the ground state has not been reached. B: A sphere with $N = 410$ particles after “closure catastrophe”. There are large regions of disorder on the side of the sphere in view. A-B adapted from [142] with permission. C: A droplet partially coated in particles. Ribbon-like crystals can be seen. Adapted from [19] with permission.

can follow one of two pathways, depending on the nature of the interactions [1, 137, 138]. We will describe each in turn below.

If the attraction is relatively long ranged, the crystal begins to incorporate defects [139]. These can either be in the form of the isolated dislocations and scars, as discussed in Sec. 1.2.1 [140, 141], pleats [141], or pentagonal vacancies [139, 142]. The pleats which form on spherical caps always terminate at the boundary, and at low temperatures they are straight and oriented radially [141]. Pentagonal vacancies are holes in the crystal lined by five particles. They are unique to attractive systems [139, 142] and can have a topological charge of either $+1$, indicating that they disrupt the orientational order, or 0 , indicating that they only disrupt the positional order.

Then, as the crystal covers the majority of the surface, there are two kinetic traps which can prevent it from reaching the free energy minimum. If the surface is small, so that the global minimum would have fewer than 32 particles [142], the mechanism is “hole implosion” [142], see Fig. 1.2A. By prematurely closing the hole left by the growing crystal, the system eliminates all the remaining line energy at the cost of some strain. This prevents more particles from being added, trapping the system in a metastable state. On a larger surface, the Mullins-Sekerka instability causes the crystal to deviate from isotropic growth [143–145]. Instead, the cluster boundary fluctuates. When the sphere is almost covered these fluctuations collide, resulting in disordered patches on one hemisphere, see Fig. 1.2B. This is called a “closure catastrophe” [142].

If the attraction is short ranged, the crystal cannot incorporate defects as the energy cost is too high. Instead, the growth of the nucleus becomes anisotropic, splitting into a number of

thin ribbons radiating from the cap-shaped core [1, 19, 136]. This pathway represents a trade-off between the energy cost of forming large-perimeter but less-frustrated ribbons and a more compact but highly frustrated cap. An experimental example of this structure is depicted in Fig. 1.2C. Recent simulations suggest that branching and defects are not exclusive to each other. For certain interaction ranges and line tensions they can coexist [146].

1.3 Non-spherical surfaces

We have seen that confinement to a spherical surface strongly influences the behaviour of adsorbed colloids. An even richer set of phenomena can be observed on surfaces with varying curvature. An important motivation comes from biophysics, where cell membrane structures like tubular networks, Golgi apparatus and the endoplasmic reticulum display intricate and highly functional geometries [2, 147–149]. Proteins which sense curvature are partially responsible for vital processes such as endocytosis, vesiculation and tubulation, as well as stabilising the membrane surface [2, 87–89, 150, 150–153]. These mechanisms require complex chemistry, but simple colloidal particles can also be sensitive to curvature [15, 154–157]. There is in fact a growing interest in exploiting curved fluid interfaces to self-assemble colloidal particles [124, 157], with potential application in the fabrication of meta-materials with tunable mechanical, optical or electronic properties [157–160], medical physics [161], and biosensors [162, 163].

Fabricating surfaces with arbitrary curvature is much harder than making spheres, but recent experimental advances have enabled the creation of finely controlled geometries. One method is to employ non-spherical Pickering emulsion droplets, but such droplets often become unstable and reduce their surface energy by becoming spherical. As an important step forward, it has recently been demonstrated that toroidal droplets can be stabilised by a yield stress fluid [13, 164]. This is a material which acts as a solid below a certain stress threshold, and as a liquid above. If the yield stress of the medium is higher than the force directing the droplet to become more spherical, the droplet can be trapped in its non-spherical shape. Colloidal particles can then be adsorbed onto the interface between the yield stress fluid and the surrounding medium [13, 164].

Another method involves assembling a number of vertical micron-scale posts and pinning a fluid interface from them, to which the colloidal particles can adsorb [15, 154–156]. Here,

the fluid interface drapes from the top of the posts to the edge of the bath with its curvature governed by the Laplace equation. A further alternative that has been realised experimentally is to exploit capillary bridges forming between two opposing solid surfaces. By tuning the contact angle of the liquid on the solid, capillary bridges with positive or negative Gaussian curvature can be obtained [120,165].

Excitingly, a more general approach has also recently become available. Micron-scale 3D printing allows arbitrary solid curved surfaces to be created. In this approach, colloids can be adsorbed irreversibly onto these with depletion interactions [16,166]. At present, the downside of this approach is that the resulting solid surface can be quite rough, which in turn affects the diffusion of the adsorbed colloidal particles.

1.3.1 Defining curvature

Before we further discuss the behaviour of colloids on non-uniformly curved surfaces, it is important to establish our convention to describe the curvature of a surface. The curvature of a point on a surface is defined by its two principal radii of curvature R_1 and R_2 . These are the highest and lowest radii out of all the circles that can be placed in contact with that point such that they follow the curvature of the surface and lie in a plane containing the surface normal. To illustrate this, three examples are presented in Fig. 1.3. In panel (A), for a bowl-like geometry, the surface curves in the same direction in the two planes of principal curvature. Hence R_1 and R_2 have the same sign. In panel (C), for a saddle-like geometry, the surface curves in different directions, and hence R_1 and R_2 have opposite signs. In our convention, here R_1 will be the largest positive radius of curvature, and R_2 is the most negative radius of curvature. Finally, in panel (B), for a cylinder, one of the principal radii of curvature is infinite, i.e. the curvature is zero in that direction. These principal radii of curvature can be combined in various ways to yield important measures of curvature that are invariant under a change of coordinate system. First, the Gaussian curvature is given by $K = 1/(R_1 R_2)$. For instance, a sphere has positive Gaussian curvature, a saddle negative Gaussian curvature and a cylinder zero Gaussian curvature. Second, the mean curvature is given by $H = \frac{1}{2}(1/R_1 + 1/R_2)$. Hence, a cylinder has non-zero mean curvature despite having $K = 0$. Third, we define the deviatoric curvature $\Delta c = 1/R_1 - 1/R_2$, which is a measure of the anisotropy of the curvature.




Figure 1.3: A schematic of three curved surfaces. The planes of principal curvature are shown in grey. The black lines represent the intersections of the planes of principal curvature with the surface. At the point they cross, R_1 and R_2 are the radii of the circles in the planes of principal curvature which have the same curvature as the surface. Left: The principal curvatures have the same sign and magnitude, so the Gaussian and mean curvatures are positive, and the deviatoric curvature is zero. Middle: The principal curvature in one direction is flat, so the radius of curvature is infinite. Therefore, the Gaussian curvature is zero, and the mean curvature and the deviatoric curvature are equal and positive. Right: The principal curvatures have equal magnitude and opposite sign. The Gaussian curvature is negative, the mean curvature is zero and the deviatoric curvature is positive. Images courtesy of Mark Miller.

1.3.2 The generalised Thomson problem

One generalisation of the Thomson problem is to consider the ground states of particles on an arbitrary curved surface. The two most common geometries used to date are the torus and the capillary bridge. Depending on the mean curvature, the latter can belong to one of the three possible Delaunay surfaces, corresponding to catenoids (zero mean curvature), unduloids or nodoids (non-zero mean curvature). As in the spherical case, crystals on surfaces of non-constant curvature are frustrated, but the extent of this frustration can vary significantly. This can lead to defects favouring particular regions of the surface [107, 120, 167–171] and the stabilisation of defects not found on the sphere [120, 168, 169, 171].

While $+1$ disclinations and scars dominate the defect structures on spheres, on non-spherical surfaces, -1 disclinations and pleats also become widely observed. -1 disclinations are particles with seven neighbours, and they are typically found in regions of high negative Gaussian curvature [120, 168, 169]. They have been observed experimentally on unduloid surfaces [120], and on unduloids and tori in simulations [168, 171–173]. This is in contrast to $+1$ disclinations that favour regions of high positive Gaussian curvature [167, 168, 170, 171]. Pleats are chains of disclinations with an equal number of $+1$ and -1 defects, and so they

have no net topological charge. On capillary bridges, these are found when the Gaussian curvature is weak [120, 141, 169, 172], because they offer a much gentler relaxation of the frustration compared to isolated defects or scars [120, 169].

1.3.3 Phase transitions

When the surface has a non-uniform curvature, it may affect phase transitions in several different ways. Until now there are three main aspects which have received attention in the literature. First, we previously argued that, on a sphere, the nucleation and growth profile is determined by the perimeter-area relationship of a minimal patch. On a surface of non-constant curvature, this relationship is not only more complex, but also depends on where on the surface nucleation begins [18]. Generally, a small nucleus in a region of positive Gaussian curvature has a smaller perimeter than one in a region of negative Gaussian curvature, favouring nucleation in regions of positive curvature. The variation in curvature also leads the nucleus to grow inhomogeneously, as sections of the growing nucleus in positive Gaussian curvature regions outgrow those in negative Gaussian curvature regions. In fact, the growth of the nucleus can be arrested completely in negative Gaussian curvature regions due to the presence of an additional free energy barrier [18].

Second, when multiple phases exist, such as during phase separation, one of the phases may become localised to particular regions of the surface [3, 12, 174–177]. This has been studied most extensively in the context of lipid bilayer membranes [3, 174–176], where the liquid-ordered lipid phase with higher bending rigidity favours regions with low curvature, while the liquid-disordered phase with smaller bending rigidity occupies the high curvature regions. Furthermore, minimum energy configurations may contain multiple, separate domains on different parts of the surface, if no single region of favourable curvature can accommodate the entire patch area [177, 178]. This is in contrast to the plane, where the ground state is always a single domain.

Third, when colloidal crystals freeze from bulk liquid on a surface of varying curvature, defects must be nucleated, and eventually coarsen towards the ground state. The initial crystal nuclei appear on regions of low curvature, as this is where the frustration is lowest [179]. As the nucleus grows, defects are included. Due to the varying curvature, the distribution of defects is very inhomogeneous, with $+1$ defects concentrated in regions of positive Gaussian curvature, and -1 defects in regions of negative Gaussian curvature. Grain boundary scars

and defects occur both when separate nuclei collide and when different growing faces of a single nucleus meet each other, but are incommensurate due to the surface curvature [179]. The surface curvature also strongly affects the coarsening process. The number of isolated defects remains constant, while dislocations tend to diffuse towards regions of high curvature, where they collide and annihilate [167, 179].

1.3.4 Complex interactions

So far, we have discussed colloids which interact through pairwise, isotropic potentials. Colloids at oil-water interfaces also experience non-isotropic and single body forces, transmitted through the curvature of the surface itself. For instance, when a colloidal particle is placed on an oil-water interface, a contact line forms around it. The shape of this contact line is partially random, as it is pinned to rough patches on the surface of the colloid [93, 123, 180], and partially determined by the underlying curvature of the surface [15, 154–156]. The complex interactions between the pinned contact line and the underlying curvature lead to a one-body force on the particle which can act towards or away from the direction of increasing deviatoric curvature [15, 154–156], depending on the size of the particle.

When large numbers of such particles interact collectively, novel structures can be stabilised. As small particles interact through quadrupolar distortions of the surface, they tend to form square lattices when confined to a curved surface in large numbers [124]. What is notable here is that an uncommon crystal structure (in this case, a square lattice) can be formed without the use of difficult-to-manufacture particles with anisotropic interactions (patchy particles) or non-spherical geometry.

Another example is provided when colloids are attached to lipid membranes. The particle causes a pinch deformation of the membrane, which partially wraps around the colloid [181]. As with the oil-water interface, this deformation has an energy cost [181, 182]. In this case, the cost is minimised by the migration of the colloid to a region of negative Gaussian curvature [181–183]. Interestingly, when multiple colloids are present, not only can the colloids interact through membrane curvature to aggregate and self-assemble [182–185], but this colloidal self-assembly in turn also deforms the shape of the membrane itself [9, 183, 186].

1.4 Thesis objectives and outline

This thesis seeks to contribute answers to three open questions. *What are the effects of confinement to a curved surface on the free energy of a two-dimensional system of isotropically attractive particles? How do these effects interact with each other? What novel properties arise from these effects and their interactions?* Our approach to these questions is centred on the use of particle-based simulation methods. In chapter 2 we discuss Monte Carlo, molecular dynamics and basin hopping calculations and describe how they must be adapted to accommodate particles on curved surfaces. In the remaining chapters, we make progress on the three questions above, building on previous work and developing our understanding of curved surfaces in new directions.

In chapter 3, we describe the contributions of finite surface size and altered perimeter-area ratio to the nucleation profile of the gas-liquid transition on the surface of a sphere. We isolate the contribution of each effect by developing analytical models and comparing them with bespoke Monte Carlo simulations. One major consequence of finite surface area is that the canonical and grand canonical ensembles are not equivalent. Therefore, we treat each ensemble in turn and describe the differences between them.

In chapter 4, we turn our attention to unravelling the general effects of non-constant Gaussian curvature, and how these depend on temperature and the range of the inter-particle attraction. By focusing on a toroidal surface, we show that novel thermodynamic states are stabilised that adopt a specific shape and location on the surface, and we identify three free energy contributions that drive the transitions between these states. These are the perimeter-area relationship of a cluster, the geometrical frustration of the crystal phase and the effect of curvature on packing. To demonstrate that these effects are general, we also present phase diagrams for a second torus geometry and a sinusoidal surface.

In chapter 5, we investigate the ground states of attractive particles on cones, which have mean curvature but no Gaussian curvature. Using geometrical analysis and basin hopping simulations, we identify magic cone angles which host defect-free achiral crystals and show how chirality can arise and be quantified on cones with intermediate angles. In addition, we compare structures formed by particles which interact through Euclidean space and along surface geodesics and show that Euclidean interactions on curved surfaces frustrate the crystal lattice. Finally, in Chapter 6 we provide overarching conclusions and suggest a number of

possible directions for future work.

Chapter 2

Methods

2.1 Introduction

In this chapter, we present the theories and techniques used in this thesis. Sec. 2.2 introduces Monte Carlo (MC) simulations, which probe the equilibrium properties of many-body systems [187]. By taking advantage of the tools of statistical mechanics, Monte Carlo simulations allow us to take microscopic properties of particles in an ensemble (for example their interaction potentials) and return the values of macroscopic observables such as average potential energy, density, and crystalline order [188]. Monte Carlo simulations can further be used in conjunction with several advanced techniques to access the free energy profiles of these systems, which in turn can be used to map phase diagrams and understand phase transitions.

Monte Carlo simulations, however, provide no information about the dynamics of the system. If we want to set up a system in a given state and study how it evolves in time, another method is required. In Sec. 2.3 we discuss molecular dynamics (MD) simulations [189]. By numerically solving Newton's equation of motion for each particle in the system, MD yields trajectories that the system could take from a fixed starting point. It is worth noting that, as small changes in initial conditions may result in large deviations on the calculated trajectories, they serve best as examples of trajectories such systems can take, rather than robust predictions of the microscopic evolution of a given system.

As the temperature of the system is decreased, both Monte Carlo and MD simulations tend to perform poorly. The system can get trapped in metastable minima for long computational times, meaning properly equilibrated sampling cannot be achieved in a reasonable time. In Sec. 2.4 we introduce the basin hopping technique, which is useful for probing the zero-temperature states of a system [190]. In this limit, the free energy surface is reduced to a

potential energy surface, whose minima can be mapped to provide insight into the structures of the system at low temperatures.

For each of these three techniques, we introduce its theoretical basis, describe its implementation, discuss any variations on the technique that we will employ and describe how it must be adapted in the special case of simulating systems confined to curved surfaces.

2.2 Monte Carlo simulation

2.2.1 Statistical basis of Monte Carlo

When dealing with many-body systems, there can be a very large number of degrees of freedom (up to $DN - D$ in both position and momentum for N bodies in D dimensions). Although in colloidal systems it can be experimentally possible to measure every degree of freedom [7], simply looking at individual configurations (called microstates) rarely provides a full insight into the physics of a system. Instead, time-averages of macroscopic variables such as temperature, pressure, and crystallinity are both more useful and easier to measure.

As we wish to compare the results of our simulations to experimental systems, we need to be able to calculate the averages of their macroscopic observables for given external conditions. However, these model systems are typically defined only by their interactions and the set of possible microstates. Statistical mechanics provides the tool-set needed to bridge the gap between the microstates of a system and the averages of its macroscopic observables. The derivations presented in this section follow those of Frenkel and Smit [188].

The fundamental postulate of statistical mechanics is that, for a system with fixed energy E , *all microstates with that energy are equally likely to be observed*. A direct consequence of this is that for systems in thermal equilibrium with an external heat bath (i.e. in the canonical ensemble), the probability of finding the system in a state with energy E is given by the Boltzmann distribution [188]

$$P = \frac{\exp(-E/k_{\text{B}}T)}{\sum_j \exp(-E_j/k_{\text{B}}T)}.$$

where the sum is over all microstates. The denominator of this equation represents another important statistical entity, the partition function Z . The knowledge of this distribution allows for the calculation of the ensemble averages of collective variables (for example the

experimentally accessible quantities discussed above) by

$$\langle A \rangle = \frac{\sum_i A_i \exp(-E_i/k_B T)}{\sum_j \exp(-E_j/k_B T)},$$

where $\langle A \rangle$ is the ensemble average of collective variable A and A_i is the value of A for microstate i . In classical systems, microstates are not discrete but defined by continuous variables (usually the positions and velocities of all the particles). By changing the above sums to integrals, we arrive at

$$\langle A \rangle = \frac{\int_V \int_{-\infty}^{\infty} A(\mathbf{p}, \mathbf{r}) \exp[-E(\mathbf{p}, \mathbf{r})/k_B T] d\mathbf{p} d\mathbf{r}}{\int_V \int_{-\infty}^{\infty} \exp[-E(\mathbf{p}, \mathbf{r})/k_B T] d\mathbf{p} d\mathbf{r}}, \quad (2.1)$$

where \mathbf{p} and \mathbf{r} are the ND -dimensional position and momentum vectors of the whole system, and V is the volume of the system. For a conservative, many-body system, the energy of a particle can be split into kinetic and potential parts $E(\mathbf{p}, \mathbf{r}) = K(\mathbf{p}) + U(\mathbf{r})$. The kinetic contribution is given by the familiar form

$$K(\mathbf{p}) = \sum_{i=1}^N \frac{\mathbf{p}_i^2}{2m_i},$$

where \mathbf{p}_i is the momentum vector of particle i , whose mass is m_i . In the case that the observable of interest is only a function of particle position $A(\mathbf{r})$, it can be seen that the momentum part of both integrals in Eq. 2.1 can be separated out and cancelled, giving the velocity-independent form

$$\langle A \rangle = \frac{\int_V A(\mathbf{r}) \exp[-U(\mathbf{r})/k_B T] d\mathbf{r}}{\int_V \exp[-U(\mathbf{r})/k_B T] d\mathbf{r}}. \quad (2.2)$$

It is this equation which can be evaluated by Monte Carlo simulations. There is a final point to be considered. The above equation sums over all microstates, so the averages it measures are known as *ensemble averages*. However, in experiments, the system is usually prepared in a given state and observed over time, yielding *time averages*. It is assumed that for most systems, the time average will converge to the ensemble average if the measurement is taken over a sufficient period. It should be noted that for some systems such as glasses, a “sufficient time” can be impractically long. In addition, some systems, such as nearly harmonic solids or systems containing infinitely high energy barriers, are genuinely non-ergodic [188]. No such

systems shall be discussed further in this thesis.

Unfortunately, not all quantities of interest are simple averages over phase space. The most important exceptions are free energies. The Helmholtz free energy of a system is

$$F = -k_B T \ln Z(N, V, T) = -k_B T \ln \left(\frac{\int_V \int_{-\infty}^{\infty} \exp[-E(\mathbf{p}, \mathbf{r})/k_B T] d\mathbf{p} d\mathbf{r}}{\Lambda^{DN} N!} \right),$$

where Λ is the de Broglie wavelength. As this equation does not have the form of Eq. 2.2, it cannot be approximated by Monte Carlo simulations. In this thesis we are only interested in comparing how different thermodynamic states contribute to the total free energy of a system at a given state-point, so we only need to calculate free energies up to an additive constant. This will allow us to make the necessary transformation to render these free energy differences calculable.

Consider now that $A(\mathbf{r})$ could be a reaction coordinate for some phase transition (for example, a measure of crystallinity for a liquid-crystal transition). The free energy profile $F(A)$ measures the contribution of configurations with a given value of A to the total free energy

$$F(A) = -k_B T \ln \left(\frac{\int_V \int_{-\infty}^{\infty} \delta(A - A'(\mathbf{r})) \exp[-E(\mathbf{p}, \mathbf{r})/k_B T] d\mathbf{p} d\mathbf{r}}{\Lambda^{DN} N!} \right),$$

where δ is the Dirac delta function and A' is the measured value of A for a given configuration \mathbf{r} . By rearranging $F(A) - F$ (and separating the momentum parts of both integrals as before) one can obtain the equation

$$\exp[-(F(A) - F)/k_B T] = \frac{\int_V \delta(A - A'(\mathbf{r})) \exp[-U(\mathbf{r})/k_B T] d\mathbf{r}}{\int_V \exp[-U(\mathbf{r})/k_B T] d\mathbf{r}}, \quad (2.3)$$

the right hand side of which has the form of Eq. 2.2, and hence can be evaluated by simulation.

2.2.2 Monte Carlo simulation

Averages of the type in Eq. 2.2 must be calculated numerically for most systems. The simplest method to achieve this would be to apply some ND dimensional implementation of the trapezium rule to both integrals, but this is not computationally feasible for even small systems as the number of degrees of freedom is extremely large and the energy is very sensitive to the particle position, so the grid would have to be extremely small. Instead, we turn to Monte Carlo simulations, so named because of their prominent use of random numbers. It




Figure 2.1: An illustration of the basic Monte Carlo scheme. Left, an arbitrary shape embedded in a square box with sides of length L . Right: random points are generated in the square, allowing the area of the embedded shape to be estimated.

was first used by Enrico Fermi in the 1930s, as the basis for an analogue computer which predicted neutron diffusion [191], and was later rediscovered and popularised by Metropolis, Rosenbluth, Rosenbluth, Teller and Teller as one of the earliest computational methods [187].

The Monte Carlo scheme aims to avoid the problems described above by replacing the grid by randomly sampling points in phase space. The basic algorithm proceeds as follows. Consider the shape in Fig. 2.1, embedded in a box of area a_0 . To estimate its area, we could generate N_{tot} random points uniformly distributed in the box, and count how many fall inside the shape N_{in} . Our estimate for the area would be given by $a = a_0 N_{\text{in}} / N_{\text{tot}}$. We could try this method for solving Eq. 2.2 by generating random configurations of the system and averaging $A(\mathbf{r}^N) V^N \exp[-U(\mathbf{r}^N) / k_B T]$ for the numerator and $V^N \exp[-U(\mathbf{r}^N) / k_B T]$ for the denominator. However, this method has its own problems. For the vast majority of possible configurations, the Boltzmann factor is vanishingly small, so we would waste a lot of time calculating configurations which do not contribute much towards the average. A better method would spend more time in configurations with higher Boltzmann weights; this is called importance sampling.

The ideal solution would be to devise a scheme for generating random configurations from scratch sampled directly from the Boltzmann distribution. The solution to Eq. 2.2 would then be the simple average of $A(\mathbf{r}^N)$ over a large number of these configurations. Unfortunately, there is generally no efficient method available to generate such configurations. Nevertheless, it is instructive to introduce a method for sampling from an arbitrary distribution in low dimensions at this point, as it represents a logical step towards a useful Monte Carlo method and is also used in Sec. 2.2.7. This method is called rejection sampling.




Figure 2.2: A schematic of the rejection sampling algorithm. Points are uniformly generated in the dotted rectangle. The probability that a point generated at a given value of x being accepted (blue) is proportional to $g(x)$. Therefore, the set of x coordinates of the blue points correctly sample the distribution $g(x)$. Red points are rejected.

2.2.3 Rejection sampling

The aim of this technique is to generate a series of samples from an arbitrary, but bounded, distribution [192]. It is convenient to discuss rejection sampling for a one dimensional distribution $g(x)$, where x ranges from x_{\min} to x_{\max} and $g(x)$ ranges from 0 to g_{\max} . However, the technique generalises in principle to multiple dimensions. To obtain a sample:

1. Pick a value x_r at random from a uniform distribution in the range x_{\min} to x_{\max} .
2. Pick a value of g_r at random from a uniform distribution in the range 0 to g_{\max} .
3. If $g_r < g(x_r)$ then return x_r as the sample, otherwise go back to step 1.

The schematic shown in Fig. 2.2 illustrates how this works. The probability of finding the system with $x = x_s$ is proportional to the length of the line between the x axis and $g(x)$. It follows that a uniform, random scattering of points under the curve of $g(x)$ will have the correct distribution of x coordinates. This can be achieved by uniformly generating points in the rectangle defined by 0, g_{\max} , x_{\min} and x_{\max} and rejecting points that lie above the curve of $g(x)$.

We will use the rejection method to sample individual coordinates on some curved surfaces where the correct distribution cannot be generated analytically from a uniform deviate. However, this technique is a poor choice for sampling the Boltzmann distribution in a high-dimensional space for the same reason that the simple Monte Carlo method described above fails: most randomly generated configuration will have a very low Boltzmann weight. There-

fore, the rate of rejection will be very high and a large number of iterations will be required to attain a good average. On the other hand, because $g(x)$ and g_{\max} and g_r can be scaled by an arbitrary factor without changing the output of the algorithm, this method does avoid calculating the partition function by simply dividing the Boltzmann distribution through by Z .

2.2.4 Metropolis Monte Carlo

Let us assume that we have prepared the system in configuration o , with a high Boltzmann weight $\exp[-U(\mathbf{r})/k_B T]$. Configurations very similar to o will probably also have a high Boltzmann weight. In the name of importance sampling, it would be of clear advantage to use a version of the rejection sampling algorithm where, instead of generating the new configuration from scratch each time, we instead used a small perturbation to the previous configuration. This is the motivational basis for the Metropolis method [187].

The general form of such an algorithm would be

1. Pick a configuration t , close to configuration o , at random with probability $\alpha(o \rightarrow t)$.
2. Calculate the probability of accepting the move, $\text{acc}(o \rightarrow t)$.
3. Generate a number r at random from a uniform distribution in the range 0 to 1
4. If $r < \text{acc}(o \rightarrow t)$ then the next configuration (n) is t , otherwise it is o .

The next step is to derive a form for $\alpha(o \rightarrow t)$ and $\text{acc}(o \rightarrow t)$ such that the above algorithm samples the Boltzmann distribution [188]. Consider a very large number of Monte Carlo simulations M (much larger than the number of configurations of the system). If the Monte Carlo simulations have been working, the set of M configurations will approximate the Boltzmann distribution closely. Each Monte Carlo simulation then takes a step (i.e., performs a new iteration). After the step, the M new configurations should still approximate the Boltzmann distribution. That is to say, the average number of simulations that were in state o but have now left it should be the same as the number that were not in state o but now are. This “balance” condition is all that is necessary for a correct Metropolis algorithm. In practice, we use a stricter condition called detailed balance, which requires that the number of simulations that were in state o but are now in state n is equal to the number of simulations that were

in state n but are now in state o . This can be written as

$$P_o \pi(o \rightarrow n) = P_n \pi(n \rightarrow o), \quad (2.4)$$

where P_o is the Boltzmann factor for configuration o and $\pi(o \rightarrow n)$ is the probability of a move from state o to n taking place, called the transition matrix. The transition matrix can be split into two parts: the probability that the move from o to n is proposed $\alpha(o \rightarrow n)$ (step 1) and the probability that it is then accepted $\text{acc}(o \rightarrow n)$ (step 3 and 4): $\pi(o \rightarrow n) = \alpha(o \rightarrow n) \times \text{acc}(o \rightarrow n)$. Notice that this requires we resample the old configuration o in step 4, rather than trying again as we did in the rejection sampling algorithm. If we did repeat the trial moves until one was accepted, then $\pi(o \rightarrow n)$ would have to include terms that consider the chance that other trial moves were rejected before the $o \rightarrow n$ move was settled on. Substituting the above expression for $\pi(o \rightarrow n)$ into Eq. 2.4 gives,

$$\frac{\alpha(o \rightarrow n) \times \text{acc}(o \rightarrow n)}{\alpha(n \rightarrow o) \times \text{acc}(n \rightarrow o)} = \exp[-(U_n - U_o)/k_B T]. \quad (2.5)$$

In the original Metropolis scheme, the trial move matrix is chosen to be symmetrical, so $\alpha(o \rightarrow n) = \alpha(n \rightarrow o)$, this leaves the following condition for the acceptance probability

$$\frac{\text{acc}(o \rightarrow n)}{\text{acc}(n \rightarrow o)} = \exp[-(U_n - U_o)/k_B T].$$

There are many possible choices of $\text{acc}(o \rightarrow n)$ which satisfy this criterion. The most common choice is the Metropolis criterion given by

$$\text{acc}(o \rightarrow n) = \min\{1, \exp[-(U_n - U_o)/k_B T]\}.$$

All that remains is to define a set of symmetric trial moves. The simplest is to select a particle at random and shift it a short (and random) distance in a random direction. This will produce a chain of configurations that looks like a series of snapshots of Brownian motion. One of the advantages of Monte Carlo simulations, however, is that a more creative choice of moves can improve sampling performance. A common example is swapping the positions of two particles (for polydisperse systems). The additional moves used in this work are collective moves and replica exchange.

Collective moves

During normal Monte Carlo sampling, it is difficult for spatially separated aggregates of particles to come together, even if that would be thermodynamically favourable. This is because, for a cluster to move, each particle must individually move in the same direction, one particle at a time. It is therefore valuable to introduce an additional move type into the system that allows particles to move simultaneously.

An efficient technique for doing this was developed by Troisi *et al.* [193]. In this scheme, the system is partitioned into clusters stochastically, based on the interaction energy between pairs of particles. These clusters are then either rotated around their centre of mass or translated together in space. Note that for this method to work, the clusters must not be re-arranged internally during the move. Therefore, this method can only be applied to the small subset of surfaces with constant, isotropic, curvature. In practice, this means the plane and the sphere. The two stages of this process are explained in detail below.

Firstly, the system is partitioned into clusters. This procedure is carefully designed such that the average cluster size changes each time the clusters are assigned. This allows for collective motion over a range of length scales, which produces efficient sampling. The partitioning is performed as follows. Every particle is initially assumed to be its own cluster. Then, every pair of particles is considered in turn. The probability that the clusters containing particles i and j are directly joined into a single cluster is given by

$$P_{ij}(\text{link}) = 1 - \exp\left(\frac{\min(U_{ij}, 0)}{k_B\Theta}\right), \quad (2.6)$$

where V_{ij} is the energy of the pairwise interaction between the two particles and Θ is a pseudo-temperature which can be used to control the average cluster size. In the version of this method employed in this work, Θ is set to a random value between Θ_{\min} and Θ_{\max} each time the system is partitioned, so that the average cluster size also varies whilst maintaining detailed balance. In this way, the final set of clusters is built up by repeatedly merging clusters according to the above criteria.

The second part of this method is the cluster moves. These can be performed more often than the first part. Here, the cluster is either translated across the surface or rotated in the surface about its centre of mass. To ensure that the Boltzmann distribution is correctly




Figure 2.3: A schematic of the technique of Troisi *et al.* for a small system of attractive particles on a plane [193]. Left: The initial state. Region A is separated from the bulk by a line of defects, shown by the red line. The lattice vectors of region B are not aligned with the bulk and region C is not connected to the bulk. Middle: The system is partitioned into clusters stochastically, based on the probability described by Eq. 2.6. In this idealised case, regions A, B and C have each been grouped as single clusters. The bulk has been split into two clusters D and E. Right: A few collective moves have been performed on the system. Clusters A and C have been translated into more energetically favourable positions. Cluster B has been rotated to align with the bulk. Cluster D has not moved relative to cluster E, as any move would be towards a configuration with low Boltzmann weight, and was rejected. These moves are idealised and intended for illustration only. As moves are proposed randomly, it is quite possible that, for example, cluster C would move away from the bulk region.

sampled, the probability of accepting a cluster move is given by

$$P(\mathbf{r}'^N) = \min \left[1, \frac{p(c, n)}{p(c, o)} \exp \left(-\frac{U(\mathbf{r}'^N) - U(\mathbf{r}^N)}{k_B T} \right) \right].$$

where $p(c, n)$ is the probability of the current set of clusters c being formed in step 1 if the system was in its new state. $p(c, o)$ is the probability of the current set of clusters being formed if the system was in its old state. An illustration of this technique can be found in Fig. 2.3.

Replica exchange

The second technique for improving the efficiency of the simulations is replica exchange. At low temperature, systems of attractive particles have a large number of local free energy minima. This is especially the case when the interaction range is short [194]. Therefore, it is easy for the system to get stuck in a metastable state, which prevents it from efficiently exploring configuration space. To escape a local minimum, it is not sufficient to only consider small perturbations to the system (as in the traditional Monte Carlo moves), as then the free energy barrier must be climbed, which is very unlikely. It is, therefore, necessary to propose moves which completely rearrange the system. Unfortunately, a new configuration picked at

random will almost certainly have a very small Boltzmann weight, and therefore be rejected.

Replica exchange solves this problem by using configurations sampled from Monte Carlo simulations performed at a range of temperatures as proposed moves [195]. These collective moves greatly improve the sampling of configuration space and have the additional advantage of obtaining results for the simulations at other temperatures as well as the original. An example, demonstrating how this technique might help simulations escape from local minima is depicted in Fig. 2.4.

In practice, replica exchange is performed on an array of simulations running in parallel at a series of temperatures. A new trial move is added to the program: the exchange of the configuration between two simulations at neighbouring temperatures. This exchange is accepted with the probability

$$P(\text{acc}) = \min \left[1, \exp \left\{ \left(\frac{1}{k_B T_2} - \frac{1}{k_B T_1} \right) (U_1 - U_2) \right\} \right], \quad (2.7)$$

where T_1 is the temperature of the first simulation, T_2 is the temperature of the second simulation, U_1 is the total potential energy of the first simulation and U_2 is the total potential energy of the second simulation. Therefore, results are obtained simultaneously for all temperatures in the array, and these results are well sampled because of the configuration exchanges.

2.2.5 Umbrella sampling

The Monte Carlo technique can also be used to measure free energy profiles by approximating Eq. 2.3. The first practical issue is that this equation represents the average of a delta function, which is inconvenient. To bypass this, we can either select a discrete order parameter (in which case we replace the delta function with a Kronecker delta function) or measure a histogram H which approximates the right-hand side of Eq. 2.3, in effect replacing the delta function with a square function of width w and height $1/w$. w must be chosen so that it is wide enough that a sufficiently large number of configurations seen over the course of the simulation fall within each bin, but thin enough that the details of the distribution can be resolved.

The simplest method to measure $F(A) - F$ is to perform a Monte Carlo simulation, measure A at each step, accumulate the normalised histogram $H(A)$ and take $F(A) - F =$




Figure 2.4: A schematic of the replica exchange technique. Simulations are performed at three temperatures, $T_1 < T_2 < T_3$, and time increases from t_1 to t_3 , from left to right. Initially, the simulation at T_1 (light grey) is trapped in a local minimum. The simulations at T_1 and T_2 are then exchanged. The simulations are continued as normal. Now that the light grey configuration is at a higher temperature, moves which increase the potential energy are more likely to be accepted and it can escape the minimum. Later, the reverse exchange takes place and the light grey configuration will likely relax into another local or global minimum. As Equ 2.7 maintains detailed balance, the simulation at T_1 can be sampled at times t_1 , t_2 and t_3 to obtain correct estimates of the thermodynamic variables.

$-k_B T \ln(H)$. Here we run into another problem. If we are studying the free energy of a first order phase transition, we want H to be well sampled across the whole range of A , including at the top of any free energy barrier. However, due to importance sampling, the Monte Carlo algorithm will explicitly avoid sampling states in this region as they will likely have low Boltzmann weight or be rare. A technique called umbrella sampling allows us to overcome this limitation. The potential energy function is modified to bias the system towards configurations with certain values of A , and then the bias is factored out of the results [196, 197]. Here, the biasing potential is

$$U_{\text{bias}} = \frac{1}{2} k (A - A_{\text{target}})^2, \quad (2.8)$$

where k is a constant that determines the range of values of A likely to be observed during the simulation, A is the current value of the reaction coordinate, and A_{target} is the value of the order parameter we are biasing toward. To remove the effect of the biasing potential, the

histogram H is accumulated normally, and then scaled by

$$H(A) = H_0(A) \exp \left(\frac{U_{\text{bias}}(A)}{k_B T} \right), \quad (2.9)$$

where $H(A)$ is the true histogram and $H_0(A)$ is the measured histogram. As above, the free energy is then given by $F(A) - F = -k_B T \ln(H)$ where the total Helmholtz free energy F depends on the biasing potential [198]. We perform simulations over a range of values of A_{target} . Each resulting histogram will be well sampled over a small range of values of A , centred on its value of A_{target} . We ensure that adjacent values of A_{target} are close enough that their histograms overlap significantly. This way, a complete, well-sampled histogram over the whole range of A can be constructed. However, as the reference value for each of the free energy curves is different, the individual free energy curves for different values of A_{target} are offset from each other by an arbitrary amount. To match the curves a least-squares method is used to minimise the difference between the regions which overlap. The fitting algorithm is biased so that better-sampled points are weighted more heavily in the averaging.

The umbrella sampling technique can be made more efficient with the addition of replica exchange. Here, simulations are run across a range of values of A_{target} in parallel. These simulations exchange configurations with a probability

$$P(\text{acc}) = \min \left[1, \exp \left\{ -\frac{k}{k_B T} (A_1^{\text{target}} - A_2^{\text{target}})(A_1 - A_2) \right\} \right],$$

where A_1, A_2 are the current values of the order parameter in the first and second simulations and $A_1^{\text{target}}, A_2^{\text{target}}$ are the target values of the order parameter in the first and second simulations.

2.2.6 Monte Carlo in the grand canonical ensemble

In the grand canonical ensemble, the system can exchange particles as well as energy with the bath. In the thermodynamic limit (infinitely large systems), the canonical and grand canonical ensembles are equivalent. However, closed surfaces (such as a sphere or torus) are necessarily finite, and the canonical and grand canonical ensembles are not equivalent, so it is important to be able to simulate both ensembles.

Grand canonical simulations are achieved by adding a new type of trial move: the addition or removal of a particle from the system. The type of move is selected at random each cycle.

In the case of an addition move the new location is selected at random and accepted with probability

$$P(\mathbf{r}'^{N+1}) = \min \left[1, \frac{z(\mu)}{N+1} \exp \left(-\frac{U(\mathbf{r}'^{N+1}) - U(\mathbf{r}^N)}{k_B T} \right) \right]. \quad (2.10)$$

where N is the number of particles in the old state, μ is the chemical potential of the particle reservoir and $z(\mu)$ is some activity defined by

$$z(\mu) = \frac{A}{\Lambda^2} \exp \left(\frac{\mu}{k_B T} \right), \quad (2.11)$$

where A is the area of the system and Λ is the thermal de Broglie wavelength. In general, it is convenient to use z as a proxy for μ so as to absorb Λ into the control parameter.

In the case of a particle subtraction, a particle is selected at random to be removed. This move is accepted with probability

$$P(\mathbf{r}'^{N-1}) = \min \left[1, \frac{N}{z(\mu)} \exp \left(-\frac{U(\mathbf{r}'^{N-1}) - U(\mathbf{r}^N)}{k_B T} \right) \right]. \quad (2.12)$$

It is also possible to perform replica exchange on an array of grand canonical simulations at a range of chemical potentials, but constant temperature. The acceptance probability is then

$$P(\text{acc}) = \min \left[1, \frac{z_2}{z_1} \exp \left(-\frac{1}{k_B T} (N_1 - N_2) \right) \right],$$

where z_1 is the activity of the first simulation, z_2 is the activity of the second simulation, N_1 is the number of particles in the first simulation and N_2 is the number of particles in the second simulation [199]. It is also possible to perform replica exchanges two dimensions, for example making *diagonal* exchanges on a grid of simulations with varying temperatures and chemical potentials; a method known as hyperparallel tempering [199].

2.2.7 Monte Carlo simulations on curved surfaces

To perform Monte Carlo simulation on curved surfaces, we must be careful to sample the correct distribution. This is because both integrals in Eq. 2.2 and 2.3 are now surface integrals, not simple integrals over a rectangular volume in Cartesian coordinates.

In general, the location of a particle on a curved surface is represented by a pair of curvilinear coordinates u and v , related to the three-dimensional Cartesian coordinates by

$x = X(u, v)$, $y = Y(u, v)$ and $z = Z(u, v)$. In this case Eq. 2.2 becomes

$$\langle A \rangle = \frac{\int_{u_{\min}}^{u_{\max}} \int_{v_{\min}}^{v_{\max}} A(\mathbf{u}^N, \mathbf{v}^N) \exp[-U(\mathbf{u}^N, \mathbf{v}^N)/k_B T] G(\mathbf{u}^N, \mathbf{v}^N) d\mathbf{u}^N d\mathbf{v}^N}{\int_{u_{\min}}^{u_{\max}} \int_{v_{\min}}^{v_{\max}} \exp[-U(\mathbf{u}^N, \mathbf{v}^N)/k_B T] G(\mathbf{u}^N, \mathbf{v}^N) d\mathbf{u}^N d\mathbf{v}^N},$$

where $G(\mathbf{u}^N, \mathbf{v}^N)$ is the product of the square root of the determinant of the metric tensor $g(u, v)$ of the surface at the position of each particle [200]. Following through the derivation for the Metropolis criterion, we reach a new version of Eq. 2.5

$$\frac{\alpha(o \rightarrow n) \times \text{acc}(o \rightarrow n)}{\alpha(n \rightarrow o) \times \text{acc}(n \rightarrow o)} = \frac{G_n}{G_o} \exp[-(U_n - U_o)/k_B T]. \quad (2.13)$$

From this, we identify two ways to adapt the standard Metropolis Monte Carlo algorithm to curved surfaces.

The first method is to pick a set of curvilinear coordinates such that $g(u, v)$ is constant. In this case, we recover the standard Metropolis acceptance rule (Eq. 2.5). Then if we simply chose trial moves which are symmetrical in u and v , we can use the standard Metropolis algorithm. For simple surfaces, it is possible to find a set of curvilinear coordinates which obeys this criterion by following the method described by Brannon *et al.* [200]. For instance, for a sphere with azimuthal angle ϕ and polar angle θ , it can be shown that $u = \phi$, $v = \cos(\theta)$ has a constant metric, so this system of coordinates could be used with any symmetrical moves to simulate particles on a sphere.

It is not always possible to find a coordinate system with a constant metric. In this case, we select the most natural coordinate system and instead adapt the acceptance rule. If we again use only symmetrical trial moves, we obtain a new criterion on the acceptance probability

$$\frac{\text{acc}(o \rightarrow n)}{\text{acc}(n \rightarrow o)} = \frac{G_n}{G_o} \exp[-(U_n - U_o)/k_B T].$$

From this we can extract a new Metropolis criterion for curved surfaces:

$$P(\mathbf{r}'^N) = \min \left[1, \frac{G_n}{G_o} \exp \left(-\frac{U_n - U_o}{k_B T} \right) \right]. \quad (2.14)$$

Both of these methods require symmetrical trial moves. A general method for achieving this

is to select a particle at random and make trial moves uniformly in u and v

$$u' = u + d_u(2r - 1)$$

$$v' = v + d_v(2r - 1),$$

where u' and v' are the new coordinates, d_u and d_v are the maximum move lengths for each coordinate and r is a random number uniformly distributed between 0 and 1.

For grand canonical simulations, one must also be careful to maintain detailed balance in the particle insertion and removal moves. The particle removal move requires no adaptation, as the particles will already be correctly distributed on the surface thanks to the other trial moves, so proposing the removal of each particle with equal probability is acceptable. For particle additions, the trial insertion positions must be uniformly distributed across the surface. This can be done by sampling positions from the distribution $g(u, v)$ (rather than a uniform distribution in u and v). In this thesis, we do this with two-dimensional rejection sampling (see Sec. 2.2.3). It is acceptable to use the first version of step 3 of the rejection sampling algorithm here as the trial position of a particle insertion is completely uncorrelated with the current configuration of the system.

A bespoke Monte Carlo code, implementing all of the methods used in this section, was used throughout this thesis. The code is available from open data repositories [201, 202] associated with Refs. [203, 204].

2.3 Molecular dynamics

2.3.1 Molecular dynamics fundamentals

Molecular dynamics simulations calculate possible trajectories of a classical, many-body system through its phase space over a period of time [188]. This is achieved by attempting to solve Newton's laws for all the particles in the system. Early simulations focused on calculating the outcome of a series of discrete collisions of hard particles [205]. However, when long-ranged forces are involved, this is not possible. Thus, a new method was developed, that discretises time into a series of steps [189]. It is this time-driven method that we will use in this thesis.

Here, the relevant equation of motion is $m_i \ddot{\mathbf{r}}_i = \mathbf{F}_i$, where \mathbf{r}_i is the position vector of

particle i and \mathbf{F}_i is the total force on the particle i . In the simplest, microcanonical, case, the force is simply $\mathbf{F}_i = -\nabla_i U$ where U is the total potential energy felt by particle i due to all the other particles. There are several algorithms which can be used to integrate these equations. Each divides the simulation time into M time steps of length Δt and each represents a trade-off between short and long-time energy conservation and efficiency (i.e. the maximum length of Δt at which the algorithm retains accuracy; the biggest computational cost is calculating the forces, so using fewer, longer time-steps is always faster) [188]. The LAMMPS package, which we use in this work, implements the velocity Verlet integrator [206]. A particle i with momentum \mathbf{p}_i^n and position \mathbf{r}_i^n at time step n is propagated to the next time step in three stages

$$\begin{aligned}
(1) \quad \mathbf{p}_i^{n+\frac{1}{2}} &= \mathbf{p}_i^n + \frac{\Delta t}{2} \mathbf{F}_i^n \\
(2) \quad \mathbf{r}_i^{n+1} &= \mathbf{r}_i^n + \frac{\Delta t}{m_i} \mathbf{p}_i^{n+\frac{1}{2}} \\
(3) \quad \mathbf{p}_i^{n+1} &= \mathbf{p}_i^{n+\frac{1}{2}} + \frac{\Delta t}{2} \mathbf{F}_i^{n+1},
\end{aligned}$$

where \mathbf{F}_i^{n+1} is the force calculated using the new positions of all the particles. This algorithm is notable for its outstanding energy conservation [207, 208].

This scheme will accurately simulate a microcanonical (constant NVE) system. However, it is often more relevant to sample the canonical (NVT) ensemble. This can be done by including additional non-conservative forces. In LAMMPS, the Langevin thermostat is used to fix a constant temperature [209]. This works by adding a stochastic noise force and a friction force. The noise force represents random collisions with particles in the (implicit) solvent, while the friction force is a damping term that models the dissipation of energy due to friction between the modelled particles and the solvent. The total force is

$$\mathbf{F}_i = -\nabla_i U - \frac{\mathbf{p}_i}{\tau} + \sqrt{\frac{2k_B T}{\tau}} \mathbf{R}_i,$$

where τ has dimensions of time and sets the rate at which velocity autocorrelation decays. \mathbf{R} is a vector of random numbers sampled from the normal distribution with a mean of 0 and a variance of 1. This force explicitly depends on the momentum of the particle, so it cannot be integrated with the velocity Verlet algorithm shown above. Instead, an integrator derived

by Stoll *et al.* is used [209]

$$\mathbf{r}_i^{n+1} = \mathbf{r}_i^n + [\mathbf{r}_i^n - \mathbf{r}_i^{n-1}] e^{-\Delta t/\tau} - \frac{(\Delta t)^2}{m_i} e^{-\Delta t/2\tau} \left[\nabla_i U + \sqrt{\frac{m_i k_B T}{(\Delta t)^2}} \mathbf{R}_i \Theta \left(\frac{\Delta t}{\tau} - \alpha_i \right) \right],$$

where Θ is the Heaviside function and α_i is a random number uniformly distributed between 0 and 1. Unlike the Verlet formulation, this algorithm has a single step and does not deal with velocities, but instead uses the particle positions at the current and previous step to calculate the next step.

2.3.2 Molecular dynamics on curved surfaces

A scheme to adapt molecular dynamics and Langevin dynamics to work on curved surfaces has been developed by Paquay *et al.* [210]. This method both constrains the position vector of each particle to lie on the surface, and constrains the velocity vector to lie parallel to the surface. The algorithm is based on the RATTLE algorithm by Anderson, which was originally used to simulate fixed-length bonds [211]. A short description of the theoretical basis for this algorithm is given here, following that in reference [210]. For convenience, we will dispense with the thermostat and return to the microcanonical ensemble, although the friction and stochastic forces can easily be reintroduced.

The constraint that the particles must lie on the surface can be expressed by the function $g(\mathbf{r}_i)$, where $g(\mathbf{r}_i) = 0$ if particle i is on the surface. By construction, the surface normal is $\mathbf{n}_i = \nabla g(\mathbf{r}_i)$. Next, we show that this constraint also ensures that the particle velocities always lie parallel to the surface: The time derivative of $g(\mathbf{r}_i)$ is

$$\frac{dg(\mathbf{r}_i)}{dt} = \nabla g(\mathbf{r}_i) \cdot \dot{\mathbf{r}}_i + \frac{\partial g(\mathbf{r}_i)}{\partial t} = \dot{\mathbf{r}}_i \cdot \mathbf{n}_i,$$

where index i refers to the i th particle. If $g(\mathbf{r}_i) = 0$ is to hold at all times, its time derivative must also be zero. This means $(\dot{\mathbf{r}}_i \cdot \mathbf{n}_i) = 0$, which shows that the velocity of each particle will be perpendicular to the surface normal.

To handle the constraints, an additional energy term is added to the Hamiltonian of the system via the Lagrange multiplier λ

$$\mathcal{H}(\mathbf{p}^N, \mathbf{r}^N) = U(\mathbf{r}^N) + \sum_{i=1}^N \left(\frac{\mathbf{p}_i^2}{2m_i} + \lambda_i g(\mathbf{r}_i) \right).$$

According to the principles of Hamiltonian dynamics, the equations of motion for this system are

$$\begin{aligned}\dot{r}_{i\alpha} &= \frac{\partial \mathcal{H}}{\partial p_{i\alpha}} = \frac{p_{i\alpha}}{m_i} \\ \dot{p}_{i\alpha} &= \frac{\partial \mathcal{H}}{\partial r_{i\alpha}} = -\frac{\partial U(\mathbf{r}^N)}{\partial r_{i\alpha}} - \lambda_i n_{i\alpha},\end{aligned}$$

where α refers to one of the three Cartesian components of a given vector. This new system of equations is solved numerically by the RATTLE algorithm, which uses Newton iteration to determine the Lagrange multipliers at each time step, so that at every step all the constraints are satisfied and the new particle positions are consistent with the total force acting on them (including the force due to the Lagrange terms). As each constraint only depends on the position of one particle, this iteration scheme is relatively cheap to calculate, so the RATTLE-based method is efficient overall. It is also general, as it can easily be applied to any surface for which $g(\mathbf{r}_i)$ and the surface normal can be found.

2.4 Basin hopping

2.4.1 The Basin hopping method

The global potential energy minimum of a many-body system is the configuration with the lowest potential energy or, equivalently, the most stable configuration in the limit of zero temperature. Discovering global minima computationally presents several challenges. At low temperatures, the free energy barriers between states become high, so both Monte Carlo and molecular dynamics simulation can get stuck in local minima. As the potential energy landscape can be very complicated, with multiple funnels, these local minima can be very far from the true global minimum.

One method for circumventing this problem is basin hopping, first introduced by Li and Scheraga [190] and further developed by Doye and Wales [194, 212–214]. Basin hopping is a variation of the traditional Monte Carlo scheme that samples from a transformed potential energy surface

$$\tilde{U}(\mathbf{r}^N) = \min [U(\mathbf{r}^N)],$$

where $\min[f(x)]$ is a local downhill minimisation of f in x . This process associates each point on the potential energy surface with a local minimum. The set of points associated with a




Figure 2.5: A schematic of the transformation of the potential energy profile used in basin hopping simulation. The black line represents the original potential energy $U(x)$ and the red line the transformed potential $\tilde{U}(x)$. The potential shown has three basins of attraction.

given minimum are known as that minimum's basin of attraction. A schematic diagram of this process is shown in Fig. 2.5.

The simulation explores this surface by performing an energy minimisation after each proposed move, which incurs a computation cost. The advantage is that all moves within a basin of attraction are accepted, and barriers between adjacent minima are eliminated. This allows the simulation to explore configuration space more efficiently, so it is more likely to pass through the true global minimum, although this is never guaranteed. This scheme is less likely to get stuck in local minima away from the global minimum. However, this problem may still occur for potential energy surfaces with more than one deep funnel. This difficulty can be alleviated by using replica exchange, which is implemented exactly as described in Sec. 2.2.5.

As this method is only used to find low-energy configurations, and not sample the Boltzmann distribution, there is no need to maintain detailed balance (although a good choice of trial move can improve efficiency). In addition, the temperature that appears in the Metropolis-like criterion in basin hopping is not a physical quantity, but a simulation parameter that determines how likely a trial move into the basin of attraction of a higher energy minimum is accepted. In the work presented here, we use the GMIN program to perform all basin hopping simulations [215].

2.4.2 Basin hopping on curved surfaces

As mentioned above, trial moves for basin hopping do not need to maintain detailed balance. Therefore, no special accommodations are required to use basin hopping on curved surfaces. That said, it is often possible to find a more efficient trial move scheme than simply making uniform moves in curvilinear coordinates for one particle at a time. We construct a trial move by placing *all* particles at a random point on a small sphere centred on their original position and then projecting them back onto the closest points on the surface.

Chapter 3

Nucleation on a sphere

The roles of curvature, confinement and ensemble

By combining Monte Carlo simulations and analytical models, we demonstrate and explain how the gas-to-liquid phase transition of colloidal systems confined to a spherical surface depends on the curvature and size of the surface, and on the choice of thermodynamic ensemble. We find that the geometry of the surface affects the shape of the free energy profile and the size of the critical nucleus by altering the perimeter–area ratio of isotropic clusters. Confinement to a smaller spherical surface results in both a lower nucleation barrier and a smaller critical nucleus size. Furthermore, the liquid domain does not grow indefinitely on a sphere. Saturation of the liquid density in the grand canonical ensemble and the depletion of the gas phase in the canonical ensemble lead to a minimum in the free energy profile, with a sharp increase in free energy for additional growth beyond this minimum.

3.1 Introduction

In this chapter, we will investigate how two major effects of curvature—the perimeter–area relationship of a patch on a curved surface and the finite surface area of closed surfaces—influence systems of colloids confined to a spherical surface. These effects have a profound impact on a number of self-assembling systems in nature, such as the development of the HIV-1 virus capsid [1, 106–108] and the assembly of clathrin-coated pits on cell membranes [216, 217]. During spinodal decomposition, confinement to a curved surface can introduce new metastable states [39, 178] and even suppress typical spinodal decomposition in favour of Ostwald ripening [39]. This may influence the formation of lipid rafts in cell membranes [3, 39, 176], which are thought to be responsible for inter-cell interaction and may drive endocytosis [218–220]. In technological applications, spherical colloidal systems are now exploited for the development of novel smart materials, with applications including drug and

flavour encapsulation in colloidosomes [11, 17, 62], soft lithography [221, 222], defect functionalisation [223] and the creation of artificial cells [224, 225]. Recent research into these systems is further buoyed by the success of novel techniques for fine control over colloidal loading and interactions [11, 19, 62, 65], as well as the shape of the confining surface [120, 164, 169].

Here, we focus on the gas-liquid nucleation process of a system of attractive particles on a sphere. Such processes are interesting in their own right, but also allow us to isolate the two effects in question (the perimeter–area and finite-size effects) from a third effect, the frustration of the crystal phase, which has been discussed extensively in the literature, and summarised in Sec. 1.2. The first effect, discussed in Sec. 1.2, arises because the shape which minimises the perimeter of a nucleus depends strongly on the curvature of the underlying surface. This means that the scaling relationship between the minimum perimeter and area of a nucleus is set by the geometry of the surface. As this relationship is an important factor in determining the nucleation pathway, it is natural to expect that surface curvature will have a strong effect on nucleation phenomena [18]. On complex surfaces, this relationship can result in the stabilisation of metastable nuclei [18, 178], or even multiple nuclei [177]. This effect is also important even for a simple surface like a sphere. To illustrate this, consider a spherical cap: once it covers more than half of the sphere, increasing the area of the cap will *decrease* the length of the perimeter, in stark contrast to the case of a nucleus on a plane.

The second effect is that it is possible for a curved surface (such as the surface of a sphere) to have a finite area without being bounded by an edge. In contrast, a planar surface must either be notionally infinite or delimited by an edge that particles cannot cross. The edge-less yet finite space defined by some curved surfaces is a key feature [39, 226]. When dealing with phase transitions in computer simulations, the usual practice is to minimise finite-size effects by extrapolating to the thermodynamic limit [227, 228], but this is no longer appropriate in an intrinsically finite space. One consequence of having a finite system is that phase transitions which would normally be regarded as first-order are no longer sharp. Another is that finite size breaks the equivalence between the canonical and grand canonical ensembles, making it important to distinguish between open and closed experimental systems. In the grand canonical ensemble, finite size causes the saturation of the liquid density when the liquid cluster covers the whole surface, while in the canonical ensemble the gas phase eventually becomes depleted—an effect that has been exploited in the liquid–solid transition in three-dimensional systems to control polymorph crystallisation [229, 230].

We use Monte Carlo simulations to quantify these effects. In agreement with previous continuum models [18,136,177], we show the surface geometry primarily affects the nucleation barrier and critical nucleus. Our Monte Carlo simulation results further confirm Gómez’s prediction [18] for the dependency of the size of the critical nucleus on the curvature, except at low supersaturations, where the line tension becomes more sensitive to curvature. Unlike in the above continuum models we also see a minimum in the free energy profile which prevents the liquid domain from growing indefinitely.

The two effects of curvature are not always easy to disentangle, but using analytical models, we are able to separate their free energy contributions. By assuming the liquid phase to be a van der Waals-like fluid and the gas phase to be an ideal gas, we develop a model that captures the key features of the Monte Carlo simulation results in the grand canonical ensemble. We show that the perimeter effect is the primary driver of the changes to the shape of the nucleation barrier and the size of the critical nucleus, while the finite-size effects, ignored in the existing continuum models, cause the free-energy minimum found at larger nuclear sizes.

Finally, we shed some light onto the difference between open and closed colloidal systems by comparing simulations in the canonical and grand canonical ensembles. In particular, we observe that nucleation is absent for a Lennard-Jones system on a sphere in the canonical ensemble. However, we are able to characterise the free energy of a liquid cluster around its equilibrium size, and explain the shape of the curve by extending a free energy model for nucleation in finite systems developed by Rao and Berne [226] and by Reguera *et al.* [231].

This chapter is organised as follows. In Sec. 3.2 we describe the model system to be simulated and the methods employed to generate and analyse our results. In Sec. 3.3 we present and discuss our results for the grand canonical ensemble. This is followed by the results for the canonical ensemble in Sec. 3.4. Finally, in Sec. 3.5 we draw conclusions and consider directions for future work.

3.2 Model and methods

In order to understand the gas–liquid phase transition on the surface of a sphere we will study a model system of Lennard-Jones particles confined to a spherical surface. The following tools are required to analyse this model: (i) an efficient Monte Carlo algorithm for particles on

a sphere (see Sec. 2.2.7); (ii) the multiple histogram reweighting technique, which allows us to find the gas–liquid coexistence curve in the grand canonical ensemble, and (iii) a reaction coordinate, so that we can use the techniques described in Sec. 2.2 to measure the free energy profile of the nucleation process.

3.2.1 Model system

Our model system consists of soft spherical particles whose centres are confined to the surface of a sphere of radius R . The particles interact isotropically *via* the truncated, shifted and smoothed Lennard-Jones potential,

$$U(r) = \left[U_{\text{LJ}}(r) - U_{\text{LJ}}(r_{\text{cut}}) - (r - r_{\text{cut}}) \frac{dU_{\text{LJ}}}{dr} \Big|_{r_{\text{cut}}} \right] H(r_{\text{cut}} - r),$$

where r is the separation of the two particles, measured in three-dimensional space, rather than along the geodesic; r_{cut} is the cut-off distance, at which the potential is truncated; and $H(x)$ is the Heaviside function [1]. $U_{\text{LJ}}(r)$ is the Lennard-Jones potential, given by

$$U_{\text{LJ}}(r) = \varepsilon \left[\left(\frac{r_0}{r} \right)^{12} - 2 \left(\frac{r_0}{r} \right)^6 \right],$$

where ε is the depth of the potential well and r_0 is the separation at which the potential reaches its minimum value [232]. We choose to fix $r_{\text{cut}}/r_0 = 2.23$ throughout this work, in order to match the potential chosen by other authors [1, 140, 196, 226, 233]. We note that previous work on planar two-dimensional Lennard-Jones systems has shown that the phase diagram is significantly affected by the choice of potential truncation; the critical temperature of the truncated potential ($k_{\text{B}}T_{\text{c}}/\varepsilon = 0.459 \pm 0.001$) is significantly lower than for the full potential ($k_{\text{B}}T_{\text{c}}/\varepsilon = 0.515 \pm 0.002$) [234].

For the remainder of this chapter, all quantities are reported in reduced units where energy is measured in units of ε , distance in units of r_0 and temperature in units of ε/k_{B} , where k_{B} is Boltzmann’s constant.

3.2.2 Monte Carlo implementation

We use the Metropolis Monte Carlo algorithm to perform simulations in both the canonical (NVT) and grand canonical (μVT) ensembles as described in Sec. 2.2. There, we introduced the idea of using trial moves uniformly distributed in the polar angle ϕ and in $\cos(\theta)$, where




Figure 3.1: Schematic of trial Monte Carlo moves. The particle is initially at position A, on the spherical surface shown in grey. It is moved to a random point B on the surface of a small sphere of radius r_{\max} , shown here as a wire-frame. It is then projected radially back onto the confining surface. Therefore, its trial position for the Monte Carlo move is at position C.

θ is the azimuthal polar coordinate. Although this maintains detailed balance, it does not sample the surface very efficiently. Instead, we use the following scheme, illustrated in Fig. 3.1, which achieves a computational performance improvement of about a factor of two. Firstly, a randomly selected particle is moved to a random point on the surface of a small sphere of radius r_{\max} (see below), centred on its old position, where r_{\max} is the maximum move length. Then the particle is projected radially back onto the confining sphere. The value of r_{\max} is adjusted throughout the equilibration stage of the simulation until the probability of a move being accepted is approximately 0.25, which ensures that the sampling of phase space is efficient [235].

For the first step of this method we need to generate random points uniformly distributed on the surface of a unit sphere. We do this using the Marsaglia method [236]. We generate two values ξ_1 and ξ_2 uniformly on $(-1, 1)$ repeatedly, until $S = \xi_1^2 + \xi_2^2 < 1$ and then take $(2\xi_1\sqrt{1-S}, 2\xi_2\sqrt{1-S}, 1-2S)$.

To improve computational efficiency further, replica exchange is employed (see Sec. 2.2.4). During simulations aimed at locating the coexistence curve, exchange moves are made between temperatures. To improve the efficiency of the canonical simulations by reducing the number of moves required for a cluster of particles to explore the surface, collective moves are used (see Sec. 2.2.4).

3.2.3 Multiple histogram reweighting

Grand canonical Monte Carlo simulations can be used to trace the gas–liquid coexistence curve. As the particle number is not fixed, and the formation of an interface has an energy cost, structures with gas and liquid regions on the same sphere are very rare. Rather, the system switches rapidly between two (meta-)stable states—entirely covered by a liquid or entirely covered by a gas, with no interfaces in either case. Here, coexistence is defined as any environment in which both states have the same free energy, meaning the system spends equal time in both. As the two states both have a uniform density, and the surface area of the sphere does not change, we can simply use the number of particles on the whole surface to determine the density of each phase. To explain how this is done, we first recall that in the grand canonical ensemble, the system has a fixed temperature T , and activity z , which is defined as

$$z(\mu) = \frac{A_0}{\Lambda^2} \exp\left(\frac{\mu}{k_B T}\right), \quad (3.1)$$

where μ is the chemical potential of the particle reservoir, A_0 is the area of the confining sphere and Λ is the thermal de Broglie wavelength. For a given (z, T) pair, the probability density of observing the system containing N particles and having energy E is

$$h(N, E|z, T) = \frac{\Theta(N, E)}{\Xi(z, T)} z^N \exp(-E/k_B T), \quad (3.2)$$

where $\Theta(N, E)$ is the density of states and $\Xi(z, T)$ is the grand canonical partition function. We measure the density and energy histogram $h(N, E|z, T)$ for a number of temperature–activity pairs close to coexistence. These histograms can be collectively fitted to Eq. (3.2) to generate a self-consistent approximation of $\Theta(N, E)$. With this, Eq. (3.2) can be used to generate estimates of $h(N, E|z, T)$ for any (z, T) pair within the ranges measured. These are then integrated over energy to obtain $h(N|z, T)$. Close to coexistence this histogram will have two peaks, one for each of the gas and liquid phases. At coexistence these peaks will have equal weight. By searching for (z, T) pairs which yield such histograms, the coexistence curve can be traced [237–239]. Fig. 3.2 shows a histogram taken from a simulation, and a coexistence histogram generated by multiple histogram reweighting.

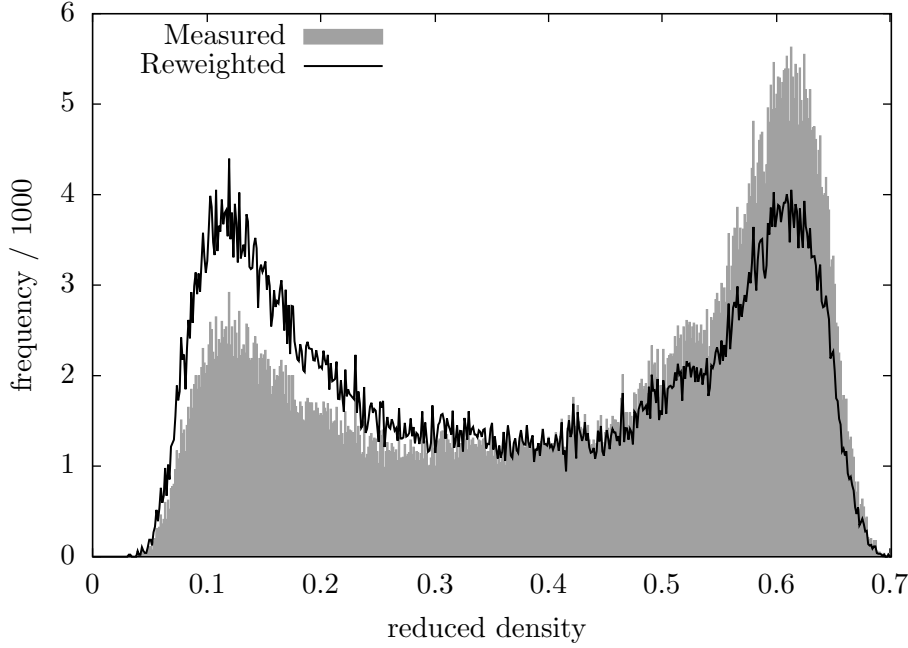


Figure 3.2: Histograms of the average packing fraction of particles on the sphere. The grey bars show one histogram measured by sampling the number of particles on the sphere over the course of a grand canonical simulation. The simulation was performed at $T = 0.51$ and $z = 27.54$. A number of such histograms, at a selection of different (T, z) combinations, are taken and used to estimate $\Theta(N, E)$. The black line is the outline of a histogram at coexistence, constructed using multiple histogram reweighting. Here the peaks of the histogram are at equal heights (to two significant figures). It corresponds to the histogram we would expect to measure at $T = 0.51$ and $z = 27.49$.

3.2.4 Measuring free energy barriers

In order to measure the free energy barrier to nucleation, a reaction coordinate must be chosen. Firstly, we will describe our choice of reaction coordinate for simulations in the context of the grand canonical ensemble, then discuss how it can be extended to the canonical ensemble. In the grand canonical ensemble, it can be assumed that the nucleation pathway will consist of the growth of liquid clusters against a background of gas, which maintains a constant density thanks to the particle bath. Therefore, the free energy change associated with the addition of a liquid cluster of size N is simply the free energy difference between a cluster of size N and area A and a gas occupying area A . Each cluster in the system can be considered independently. The reaction coordinate is defined for each cluster and is simply the number of particles in a cluster [196, 240, 241].

We use a definition based on pairwise separations to define a liquid cluster [240]: any pair of particles that can be linked by a chain of particles in which each link is shorter than a cut-off distance r_{cluster} are in the same cluster, while all other pairs are not. Here we use

$r_{\text{cluster}} = 1.45$, which was found to coincide with the first minimum in the radial distribution function¹. During the simulation, the cluster size distribution is measured every 2000 Monte Carlo steps. The free energy of a cluster containing N particles can be calculated using

$$F(N) = -k_{\text{B}}T \ln \frac{M_N}{\sum_{s=0}^Q N_{\text{Total}}(s)}, \quad (3.3)$$

where M_N is the total number of times a cluster of size N is observed, $N_{\text{Total}}(s)$ is the total number of particles present on the sphere at measurement s , and Q is the total number of measurements made during the simulation.

A standard Monte Carlo simulation does not sample M_N efficiently, especially close to the top of the barrier, so we use umbrella sampling, as described in Sec. 2.2.5. To improve efficiency further, we couple this with replica exchange (see Sec. 2.2.4), making exchange moves in the size of the target nucleus as described in Sec. 2.2.4. This approach works well for the grand canonical ensemble but must be modified for the canonical ensemble. These modifications are discussed in Sec. 3.4.

3.3 Results in the grand canonical ensemble

3.3.1 Phase diagram

Nucleation is likely to be seen in conditions close to coexistence, so knowledge of the phase diagram is necessary for measurements of the nucleation barrier to be made. We employed the multiple histogram reweighting technique to measure the gas–liquid coexistence curve of our model in the grand canonical ensemble. We chose to do this for a sphere of radius $R = 7$, which is small enough that the effects of curvature may be observed but large enough to hold sufficient particles for our measurements. We compare this phase diagram to that on a periodic square plane of the same area.

In Fig. 3.3B we present the coexistence curve for the plane and the sphere. It can be seen that they are very similar. There are two reasons for this. The first is that both the gas and liquid phases are amorphous so, unlike for elastic structures such as crystals, little structural frustration is introduced by the curvature [17, 19, 135–137, 141]. The second is that, for the overwhelming majority of the sampled states, the sphere is covered entirely in either liquid

¹This calculation was performed by Alex Wong.

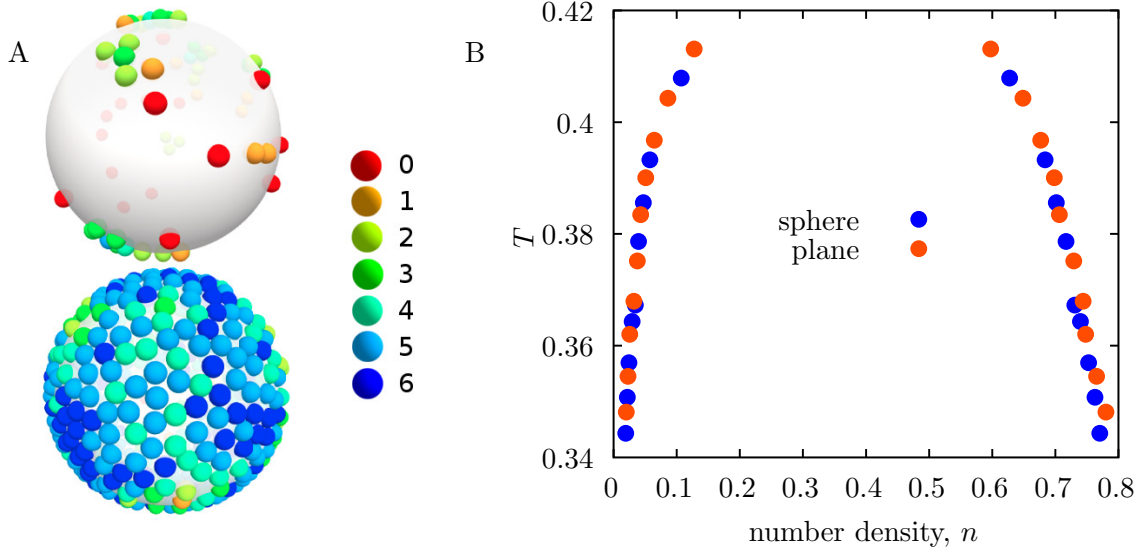


Figure 3.3: A: Snapshots of simulations on a sphere of radius $R = 7$, in the gas and liquid phases from grand canonical simulations. Particles are coloured by their coordination number as indicated in the legend. B: The coexistence curve of truncated, shifted and smoothed Lennard-Jones particles confined to the sphere (blue) and to a periodic square plane (orange) with equal area. It can be seen that the geometry of the surface does not strongly affect the gas–liquid coexistence curve in this system.

or gas, so there are no interfaces. Therefore, the overall geometry of the structure does not affect the phase behaviour (as it would if there were a phase boundary in the system, see Sec. 3.3.2). For radii other than 7 similar results were observed, except on spheres so small that they could only hold a few tens of particles. The two branches of the coexistence curve appear to be quite steep. This is at least partially due to the fact that they have been plotted over a wide temperature range, compressing the smooth broadening of the coexistence curve below the critical point into a narrow region. The shapes of the coexistence branches in the planar case (orange points in Fig. 3.3B) are comparable to those for a planar Lennard–Jones system presented in Fig. 4 of Ref. [234], despite a different method of calculation and a slight difference in potential truncation.

3.3.2 Nucleation

We measure the free energy barrier to the nucleation of the liquid phase on a sphere of radius $R = 7$ with umbrella sampling. Fig. 3.4 shows the resulting free energy as a function of the number of particles in the largest cluster, as defined in Sec. 3.2.4. The first feature of the free energy curve is a barrier, caused by the competition between the bulk free energy of the liquid phase and the line tension between the liquid and gas phases. The second feature is




Figure 3.4: An example nucleation curve taken from grand canonical simulations on a sphere of $R = 7$ at $z = 20.90$. Fits to the CNT model assuming planar geometry, the Gómez *et al.* model [18] and the full spherical model are also shown.

the minimum that is seen on the far side of the nucleation barrier. This is caused by the cluster completely covering the sphere. Adding further particles to the surface increases the density of the liquid phase above its equilibrium value, which has a high free energy cost.

In order to develop a model that explains these features, we start with the result from classical nucleation theory (CNT). In the grand canonical ensemble, the thermodynamic potential (free energy) is the grand potential $\Phi = F - \mu N$, where F is the Helmholtz free energy. The change in grand potential for the creation of a nucleus is

$$\Delta\Phi = \gamma P + A_L(P) \Delta\phi, \quad (3.4)$$

where γ is the line tension, $\Delta\phi = \phi_L - \phi_G$ is the bulk free energy of the nucleating phase relative to the parent phase per unit area, P is the perimeter of the nucleus and $A_L(P)$ is the area of the nucleus as a function of its perimeter. Fig. 3.4 includes an attempted fit of Eq. (3.4) to the simulated results, assuming the usual planar relation for a circular nucleus $A_L(P) = P^2/4\pi$. It can be seen that Eq. (3.4) not only fails to capture the increase in the free energy at large cluster sizes as expected, but it also cannot fit the shape of the initial barrier.

Recently, Gómez *et al.* developed a continuum Landau theory for phase nucleation on curved surfaces [18], which predicts that, in the absence of elasticity or finite thickness,

curvature affects the free energy barrier to nucleation by altering the relationship between the perimeter of the nucleus and its area. On a plane, assuming isotropic growth, $A(P) \propto P^2$. However, on surfaces with positive (negative) Gaussian curvature, the area grows faster (slower) as a function of P than on a plane. This will naturally change the shape of the free energy profile described in Eq. (3.4). On the sphere, the perimeter of a spherical cap of area A_L is given by

$$P(A_L) = \sqrt{4\pi \left(A_L - \frac{A_L^2}{A_0} \right)}, \quad (3.5)$$

where A_0 is the total surface area of the sphere. As shown in Fig. 3.4, this modification captures the shape of the barrier more satisfactorily.

In order to model the increase in free energy caused by the spherical confinement, a density-dependent free energy is needed. This model is built on the assumptions (i) that the free energy can be written as the sum of a bulk term and a term describing line tension; (ii) that the liquid phase can be modelled as a van der Waals-like fluid; (iii) that the gas phase can be modelled as an ideal gas; (iv) that the density of the gas phase is constant (i.e., the gas is in chemical equilibrium with the particle bath throughout the nucleation process); and (v) that the line tension between the gas and liquid phases is constant.

Under these assumptions, the free energy of a liquid cluster is

$$\Phi_L = -Nk_B T \left[1 + \ln \left(\frac{A_L - Nb}{N\Lambda^2} \right) \right] - \frac{aN^2}{A_L} - \mu_L N, \quad (3.6)$$

where N is the number of particles in the liquid cluster, A_L is the area of the liquid cluster, and a and b are the van der Waals parameters, representing the attraction between two particles, and the area occupied by a single gas particle respectively [242]. The first term represents the entropy due to the free space each particle has to explore within the droplet, where the effective area available to particles in the droplet is reduced from the total area of the nucleus A_L by the area of the particles themselves, Nb . The second is the van der Waals mean field approximation to the attraction energy due to the interaction between the particles and the third represents the coupling of the droplet to the particle bath. In Eq. (3.6), the chemical potential of the liquid droplet is denoted μ_L to allow for the fact that the droplet is not, in principle, in equilibrium with the particle bath. The corresponding activity [see Eq. (3.1)] is

$z_L = A_0 \Lambda^{-2} \exp(\mu_L/k_B T)$, so that Φ_L may also be written as

$$\Phi_L = -Nk_B T \left[1 + \ln \left(\frac{A_L - Nb}{NA_0} \right) \right] - \frac{aN^2}{A_L} - Nk_B T \ln(z_L).$$

It is the loss of area available to particles represented by the first logarithmic term in the expression for Φ_L that will ultimately lead to the sharp increase in free energy of large liquid clusters.

As the liquid phase grows, it replaces an equal area of ideal gas with free energy

$$\Phi_G = -pA_L = -n_G A_L k_B T = -\frac{z}{A_0} A_L k_B T,$$

where n_G is the number density of the gas phase and p is its pressure. As with CNT, the free energy cost of forming the interface between the gas and liquid phases is

$$\Phi_{\text{int}} = \gamma P(A_L).$$

The total free energy of formation of a liquid cluster containing N particles of size A_L is

$$\Delta\Phi = \Phi_L + \Phi_{\text{int}} - \Phi_G + d,$$

where d is a constant containing the translational entropy of the liquid cluster as a whole [243–245]. The full form of the equation is then

$$\begin{aligned} \Delta\Phi = & -Nk_B T \left[1 + \ln \left(\frac{A_L - Nb}{NA_0} \right) \right] - \frac{aN^2}{A_L} \\ & - Nk_B T \ln(z_L) + \gamma \sqrt{4\pi \left(A_L - \frac{A_L^2}{A_0} \right)} \\ & + zk_B T \frac{A_L}{A_0} + d. \end{aligned} \tag{3.7}$$

Eq. (3.7) has two mutually dependent variables: the cluster area A_L and the number of particles it contains N . In order to fit the function to a simulated free energy profile measured using Eq. (3.3), a function $A_L(N)$ is required. Rather than making a prediction for this function, we instead assume that for any N , the nucleus adopts the area that minimises its free energy, calculated by minimising Eq. (3.7) for $A_L(N)$. We find that the optimal A_L scales linearly with N , consistent with a constant liquid density n_L , until the liquid covers

Table 3.1: Fitting parameters for the grand canonical ensemble activity sweep shown in Fig. 3.5. The following fitting parameters were constrained to be identical across the whole set of simulations: $a = 1.26$, $b = 0.736$, $\gamma = 0.054$, $d = 1.78$.

$\ln(z)$	3.033	3.035	3.037	3.040	3.042
$\ln(z_L)$	3.603	3.605	3.607	3.610	3.611

the sphere. $A_L(N)$ then levels off at the total area A_0 of the sphere.

Having optimised A_L , the model has five fitting parameters: a , b , d , z_L and γ . Although the parameters a and b are related via the second virial coefficient B_2 by $a/k_B T = b - B_2$ in an exact van der Waals fluid, we find it necessary to treat a and b independently to capture deviation of the behaviour of the Lennard-Jones liquid from this ideal case.

Fig. 3.4 shows a fit of the model in Eq. (3.7) to the simulated nucleation curve at $z = 20.90$ and $T = 0.40$. It can be seen that this model successfully captures the full free energy profile. To demonstrate that we are not simply over-fitting the data, a series of simulations were performed with a range of applied activities, as shown in Fig. 3.5. We predict that the only fitting parameter that should be modified by changing the activity is z_L as this is related to the chemical potential of the out-of-equilibrium liquid droplet. Therefore, we can fit all the curves simultaneously, requiring that all fitting parameters, except z_L , are the same for each curve. The values of the fitting parameters can be found in Table 3.1. It can be seen that even with these restrictions, the model fits all the data well, suggesting that it does correctly describe the important features of gas-liquid nucleation on a sphere in the grand canonical ensemble. It is also worth noting that the fitted activity of the liquid domain z_L is close to that of the reservoir z . Furthermore, the liquid density n_L from the relation $A_L(N)$ is constant at 0.87 for all the curves in Fig. 3.5.

3.3.3 Critical nuclei

By comparing the geometric curvature of the perimeter of a spherical cap to that of a circle on a plane, Gómez *et al.* predicted that the radius of the critical nucleus depends on the radius of the spherical surface according to [18]

$$r_c(R) = R \arctan\left(\frac{r_c^0}{R}\right), \quad (3.8)$$

where $r_c(R)$ is the (geodesic) radius of the critical nucleus on a sphere of radius R , given that the critical nucleus of the equivalent system on the plane would be r_c^0 . This mapping between




Figure 3.5: The free energy barrier to the nucleation of the liquid phase for Lennard-Jones particles on the surface of a sphere of radius $R = 7$ in the grand canonical ensemble for an array of applied activities. The temperature for all simulations is $T = 0.40$. Solid lines correspond to simulation data and dashed lines to fits to the theoretical model described in Sec. 3.3.2.

the planar and spherical systems implicitly makes the assumption that the line tension and difference in free energy density between the phases are independent of the curvature of the surface.

To test the prediction in Eq. (3.8) over a range of conditions, we ran three sets of simulations. For each set, the temperature and chemical potential were kept constant while the radius of the confining sphere was varied. For reference, we quote the activity z_{ref} corresponding to the chosen chemical potential at radius $R = 7$, noting that Eq. (3.1) for the activity includes a factor of A_0 . The results for one such set of simulations are shown in Fig. 3.6.

As discussed in the previous subsection, for each free energy profile, we typically observe a constant liquid density n_L until the liquid covers the sphere. We use this liquid density to convert the measured number of particles in the critical nucleus to the radius of the critical nucleus, r_c , assuming that the critical nucleus is a spherical cap.

We estimate the uncertainty in r_c by repeating the simulation of a representative free energy curve five times, then calculating the uncertainty in the barrier height. By comparing this to the form of the nucleation curve at the top of the barrier, we estimate the uncertainty in the number of particles in a critical nucleus to be approximately ± 6 and from this calculate the corresponding uncertainty in the radius. Each set of simulations can then be fitted to Eq. (3.8) to obtain r_c^0 , the corresponding critical nucleus radius on the flat surface. We note




Figure 3.6: Free energy profile measured for grand canonical systems on spheres of different radii, but with the chemical potential fixed so that it corresponds to an activity of $z_{\text{ref}} = 20.90$ on a sphere of radius $R = 7$. The temperature for all simulations is $T = 0.40$.

Table 3.2: Line tensions extracted from the free energy profiles using Eq. (3.7) for the points plotted in Fig. 3.7. The top row corresponds to the red squares in Fig. 3.7. The second row corresponds to green circles and the third to blue diamonds.

(z_{ref}, T)	$R = 4$	$R = 5$	$R = 6$	$R = 7$	$R = 8$	$R = 9$	$R = 10$
(19.0, 0.391)	0.096	0.086	0.079	0.077	0.073	0.071	0.077
(20.6, 0.396)	0.088	0.079	0.075	0.077	0.076	0.075	0.072
(20.9, 0.400)	0.068	0.058	0.054	0.049	0.049	0.044	0.041

that, in several cases, the predicted value of r_c^0 would be impractically large to measure in a simulation.

For one activity–temperature pair (z_{ref}, T) , the simulation results can be fitted well to Eq. (3.8), and they fall close to a single line as shown in Fig. 3.7 (green circles). However, for the remaining two (z_{ref}, T) pairs, the fit breaks down for the points which correspond to two smallest radii: $R = 4$ and $R = 5$. To understand why this is the case, we fit Eq. (3.7) to the free energy profiles to extract the line tension in each case. For one of the sets (green circles), the line tension is almost independent of sphere radius, as shown in Table 3.2. In this set of data, the geometrical correction embodied in Eq. (3.8) is sufficient to relate the spherical and planar cases. In contrast, for the remaining sets (red squares and blue diamonds), which correspond to lower supersaturations, the fitted line tension decreases with increasing radius of curvature (Table 3.2), and this effect is particularly pronounced at low radii. This trend could reflect either increased sensitivity of the surface tension to curvature as the phase boundary is approached, or greater deviation from the idealised spherical cap shape that the




Figure 3.7: The size of the critical nucleus varies with the radius of the confining sphere. These data are plotted for three sets of grand canonical simulations at different temperatures and activities.

nucleus is assumed to adopt. Eq. (3.8) neglects any curvature dependence of the line tension or nucleus shape [18], and Fig. 3.7 demonstrates that such dependence can have a significant influence. These insights show the value of simulation data, as they are not bound by the assumptions of the Gómez theory.

3.4 Results in the canonical ensemble

In the canonical ensemble, the total number of particles, N_0 , is fixed in a given system. As the spherical surface we are considering is intrinsically finite in extent, the canonical and grand canonical ensembles are not equivalent. In particular, the density of the gas phase decreases when a liquid droplet forms in the canonical ensemble because the gas is not replenished by a particle bath as in the grand canonical ensemble.

Similar considerations apply to three-dimensional [246] and planar two-dimensional [247] systems with periodic boundary conditions in the canonical ensemble. In these cases, an evaporation–condensation transition is observed as a function of increasing density. First, a droplet condenses from the supersaturated vapour (Fig. 3.8B). At higher densities the condensed phase grows and a percolating slab geometry becomes preferable; this geometry reduces the surface area or perimeter of the condensed phase and is effectively stabilised by the periodic boundary conditions (Fig. 3.8C). Finally, for even denser cases, the vapour phase




Figure 3.8: The evaporation–condensation transition on a two-dimensional plane. A: A vapour. B: A condensed droplet. C: A percolating slab. D: A bubble.

is reduced to a bubble within a continuous condensed phase (Fig. 3.8D). In the spherically confined case, we observe states analogous to the initial gas state and to the droplet–gas equilibrium, as illustrated by the snapshots in Fig. 3.9A–B. The bubble state is continuously reached from the droplet by increasing the density, as the droplet simply grows to cover more than half of the sphere’s surface. The equivalent of the slab phase would be a condensed band, or spherical cap with a hole. This intermediate state does not arise on the sphere since, unlike the periodic planar case, the creation of the second interface always increases the total perimeter of the droplet and is therefore thermodynamically unfavourable.

The free energy methods described in Sec. 3.2.4 can only be applied to the canonical ensemble under certain conditions. In the canonical ensemble, as the nucleus grows, the gas phase is depleted. Therefore, if there is more than one nucleus in the system, the nuclei’s free energies are not independent and they should not be considered separately. This means that, in the canonical ensemble, the general definition of the free energy of a cluster in Sec. 3.2.4 can only be applied to situations in which it is very unlikely that more than one nucleus will be found in the system.

It is also difficult to use umbrella sampling correctly in the canonical ensemble. In order for the equilibrium state to contain a significant area of both phases, the system must be prepared as a deeply supersaturated gas, well into the spinodal region. One must be careful when applying umbrella sampling to such systems. Consider a system containing a small cluster, held within a narrow size range by umbrella sampling. The remaining particles are assumed to form a gas around the cluster. However, in the canonical ensemble, it is possible that this gas may still be supersaturated if the number of particles in the system is high enough. In this case the gas will collapse into additional liquid clusters. Therefore, the pathway mapped out by successive umbrella sampling runs will not represent the real transition pathway, which

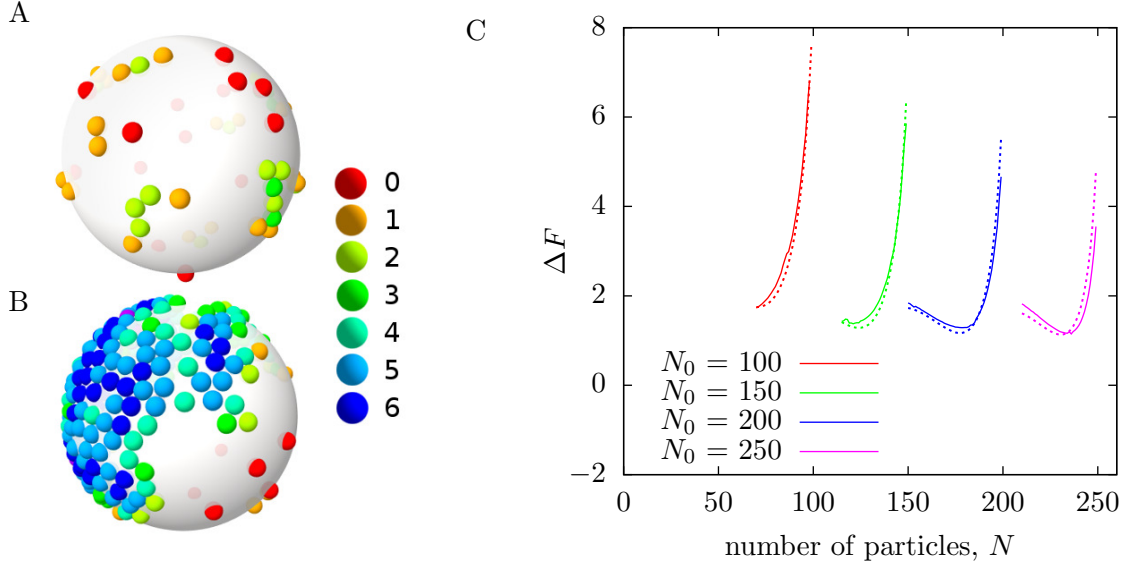


Figure 3.9: A-B: Snapshots of equilibrium configurations taken from simulations on a sphere of radius $R = 7$ at $T = 0.40$ in the canonical ensemble. At $N_0 = 50$ (A) the surface is covered by a homogeneous gas of isolated particles and small clusters. At $N_0 = 250$ (B) there is a single cap-shaped liquid drop in equilibrium with a gas. Particles are coloured according to their coordination number as indicated in the legend. C: The free energy variation of a liquid cluster around its equilibrium size on a sphere of radius $R = 7$ for four quantities of particles in the canonical ensemble. The temperature for all simulations is $T = 0.40$. Solid lines are taken from umbrella sampling simulations, while the dashed lines are fits to the modified Rao-Berne model. The values of the parameters of the model can be found in Table 3.3.

is simply spinodal decomposition. However, when the target nucleus is larger, so that the remaining gas is dilute enough that it does not spontaneously condense, umbrella sampling can be safely used. Therefore, one can safely map the free energy of liquid clusters around the equilibrium cluster size.

We have done this for a selection of values of N_0 . Fig. 3.9C shows the free energy of a system containing a single cluster embedded in the gas phase, around the equilibrium cluster size. In this case, the sharp increase in free energy with the size of the liquid cluster is not caused by high density across the surface. Instead, at large cluster sizes, the gas phase is depleted below its equilibrium density. As more gas particles are absorbed into the liquid cluster, the loss of entropy outweighs the energy gain of joining the liquid cluster. In contrast to the grand canonical ensemble, the position of the minimum is heavily dependent on the initial density of the parent gas phase in the canonical ensemble.

For sufficiently small N_0 (in the vicinity of 100 in Fig. 3.9C), the sphere is uniformly covered in a gas of isolated particles and small clusters, and the local minimum in the free energy corresponding to the droplet vanishes. Although it is hard to precisely locate the value

of N_0 at which the minimum disappears, this point, in principle, defines the density of the finite-size equivalent of the first-order phase transition. This threshold is an intrinsic property of a given sphere size and depends on its radius, but should approach the thermodynamic limit of the planar case as the radius of the sphere becomes very large.

In order to understand the free energy curves, we have developed a model for the free energy of a single liquid cluster surrounded by gas and confined to a sphere in the canonical ensemble. This model is based on one developed by Rao and Berne for finite, Euclidean, three-dimensional systems in the canonical ensemble [226]. A related approach was also adopted by Reguera *et al.* [231]. The model is based on the assumption that a single liquid cluster grows in the finite bath of the gas phase. Our version of the model accounts for the two-dimensional spherical geometry by introducing the line tension term described in Eq. (3.5). As in the original model by Rao and Berne [226], we choose to truncate the virial expansion at the second term so the vapour pressure p as a function of the density n_G of the gas in the space not occupied by the liquid cluster is

$$p(n_G) = n_G k_B T (1 + B_2 n_G), \quad (3.9)$$

where T is the temperature, B_2 is the second virial coefficient, which depends on the curvature of the surface and is evaluated by numerical integration, and n_G is

$$n_G = \frac{N_0 - N}{A_0 - N/n_{L\infty}}, \quad (3.10)$$

where N_0 is the total number of particles on the sphere of area A_0 , giving an overall density $n_0 = N_0/A_0$, $n_{L\infty}$ is the equilibrium liquid density, and N is the number of particles in the liquid phase. The chemical potential per molecule μ_G is

$$\mu_G(n_G) = \mu_G^0 + k_B T (\ln n_G + 2B_2 n_G), \quad (3.11)$$

where μ_G^0 is a reference chemical potential. The total Gibbs free energy G of a system containing a droplet surrounded by gas is

$$G = G_G + G_L + G_{\text{int}},$$

where

$$G_G = (N_0 - N)\mu_G, \quad (3.12)$$

$$G_L = N\mu_L, \quad (3.13)$$

and

$$G_{\text{int}} = \gamma P. \quad (3.14)$$

μ_L is the chemical potential per liquid particle, γ is the line tension and P is the perimeter of the cluster. μ_L can be calculated using the Gibbs-Duhem relation

$$\mu_L = \mu_L(n_{L\infty}) + \frac{1}{n_{L\infty}}(p - p_\infty), \quad (3.15)$$

where p_∞ is the corresponding reference pressure. Noting that at equilibrium

$$\mu_L(n_{L\infty}) = \mu_G(n_{G\infty}), \quad (3.16)$$

where $n_{G\infty}$ is the equilibrium gas density, $\mu_G(n_{G\infty})$ and hence $\mu_L(n_{L\infty})$ can be calculated from Eq. (3.11). Using Eq. 3.9, we can also write the equilibrium pressure as

$$p_\infty = n_{G\infty} k_B T (1 + B_2 n_{G\infty}). \quad (3.17)$$

The Helmholtz free energy F in two dimensions is

$$F = G - pA.$$

Starting from a gas of initial density $n_0 = N_0/A_0$, the free energy change associated with the formation of a liquid cluster containing N particles is

$$\Delta F = G - N_0\mu_G(n_0) - [p(n_G) - p(n_0)] A_0. \quad (3.18)$$

We now have all the components required to assemble the complete expression. Inserting Eqs. (3.12) to (3.14) into Eq. (3.18) and collecting terms gives

$$\begin{aligned} \Delta F = N_0 [\mu_G(n_G) - \mu_G(n_0)] + N [\mu_L - \mu_G(n_G)] - \\ [p(n_G) - p(n_0)] A_0 + \gamma P. \end{aligned}$$

Substitution of Eqs. (3.5), (3.9), (3.11), (3.15) and (3.17) for P , p , μ_G , μ_L and p_∞ then gives the full expression

$$\begin{aligned} \Delta F = & N_0 k_B T \left[\ln \left(\frac{n_G}{n_0} \right) + 2B_2 (n_G - n_0) \right] \\ & + N k_B T \left[\ln \left(\frac{n_{G\infty}}{n_G} \right) + 2B_2 (n_{G\infty} - n_G) + \frac{n_G - n_{G\infty} + B_2 (n_G^2 - n_{G\infty}^2)}{n_{L\infty}} \right] \\ & - A_0 k_B T [n_G + B_2 n_G^2 - n_0 - B_2 n_0^2] \\ & + \gamma \sqrt{4\pi \left(\frac{N}{n_{L\infty}} - \frac{N^2}{n_{L\infty} A_0} \right)} + d, \end{aligned} \quad (3.19)$$

The parameters in the model are $n_{L\infty}$, $n_{G\infty}$, γ , and d —a constant which contains the translational entropy of the cluster. All other symbols have their previous meanings. We refer to Eq. (3.19) as the modified Rao–Berne (MRB) model.

When applying the MRB model, we take the line tension γ obtained in the grand canonical simulations in Sec. 3.3 with the same sphere radius and temperature. The equilibrium liquid density, $n_{L\infty}$, is estimated as follows. A simulation is run without umbrella sampling. During the simulation, the total number of particles in a liquid-like environment (as defined in Sec. 3.2.4) is averaged. For each frame of the simulation, the surface of the sphere is divided into cells using a Voronoi decomposition of the particle positions [248]. The Voronoi cell of a particular particle is defined as all the points on the surface that are geodesically closer to that particle than any other. A histogram of the area of each Voronoi cell is constructed. The location of the peak of this histogram is an estimate for the inverse density of the liquid phase, as most particles are found in the bulk of the liquid cluster.

This leaves only the equilibrium reference gas density $n_{G\infty}$ and the offset d to be fitted. We constrain $n_{G\infty}$ to be the same for all runs, since it should be the same for a given temperature and sphere radius, and we allow d to be different for each curve. The results of these fits to the simulation data are shown as dashed lines in Fig. 3.9 and the corresponding values of the parameters are given in Table 3.3. It can be seen that the theory fits the simulation closely, for systems with a single large cluster.

3.5 Conclusion

We have shown that the thermodynamics of a two-dimensional soft system in disordered phases confined to the surface of a sphere is dominated by the effects of the surface geometry

Table 3.3: The calculated and fitted parameters for liquid cluster formation on a sphere of radius $R = 7$ at temperature $T = 0.40$ in the canonical ensemble. $n_{L\infty}$ and B_2 are calculated and have the values of $n_{L\infty} = 0.951$, $B_2 = -6.247$. γ is taken from grand canonical fits and has the value $\gamma = 0.054$. $n_{G\infty}$ and d are fitted. $n_{G\infty}$ is constrained to have the same value for all runs, and the best fit value is 0.120. d is fitted separately for each run.

N_0	d
100	-0.031
150	-11.40
200	-35.79
250	-75.23

on the cluster boundary and on the finite extent of the surface. Unlike the planar two-dimensional case, the finite size of these curved systems is a natural consequence of being embedded on a sphere, providing a physical realisation of an intrinsically finite system with neither hard boundaries nor artificial periodic boundary conditions.

In the grand canonical ensemble, the phase diagram of a system of Lennard-Jones particles on a sphere is indistinguishable from that of a periodic plane of the same area. This is because both phases are fluids and neither of the two (meta-)stable states of the system contains a phase boundary; the sphere is fully covered in either liquid or gas. However, the nucleation process is influenced by curvature. The free energy of the growing nucleus is initially modified by the geometry of the surface, and growth is later arrested when the surface is fully covered. A model, based on classical nucleation theory and modified to handle these additional considerations, is able to rationalise the simulation results. By investigating spheres of different radii over a range of conditions, we were also able to confirm the prediction for the curvature dependence of the critical nucleus size made by Gómez *et al.* [18] provided the sphere radius and supersaturation is not too low. At low supersaturation, we find that the line tension decreases monotonically with increasing radius of curvature, which conflicts with the assumption of constant line tension used in Ref. [18].

Due to the finite size of the system, ensemble equivalence does not hold for the system we are considering. Nucleation is not observed in the canonical ensemble of the Lennard-Jones particles. Instead, at sufficiently high densities a stable liquid cluster is established in equilibrium with a depleted gas phase, and the free energy for formation of the droplet at around its equilibrium size could be measured. In contrast to the analogous evaporation–condensation transition in a planar two-dimensional system [247], there is no intermediate slab-like phase because the creation of a second gas–liquid interface can never reduce the

total perimeter of a droplet on a sphere. A modified Rao–Berne model (with just two fitting parameters for a whole family of free energy profiles) was able to capture both the equilibrium liquid cluster size and the free energy of clusters close to the equilibrium size.

In this chapter we have focused on disordered phases confined to a surface of constant Gaussian curvature. In the next chapter we will introduce surfaces of non-constant Gaussian curvature, showing that this can have a dramatic effect on the phase behaviour of a colloidal system.

Chapter 4

Localised states on surfaces of non-constant curvature

We have seen that uniform surface curvature affects the structural, kinetic and thermodynamic properties of a two-dimensional system. Now we show that if the curvature is non-uniform, an even richer picture emerges. In this chapter, we demonstrate that surfaces of non-constant curvature can stabilise novel localised thermodynamic states in systems of attractive colloids. We argue that the existence and character of these states can be explained with reference to three contributions to the free energy of such systems. These are the perimeter–area ratio of a cluster, the frustration of the crystal lattice due to curvature and the change in second-nearest-neighbour separation due to an attractive potential which acts through space on particles on a curved surface. We also examine the transitions between these localised states, which are accompanied by a global, cooperative migration of matter across the surface. We present dynamic trajectories for these transitions and show that they are driven by competition between the three effects described above.

4.1 Introduction

In the previous chapter we discussed uniform curvature, but in nature two-dimensional systems confined to surfaces of non-uniform curvature also abound. For example, such regions are necessary in cellular membranes for key biological processes, including the sensing and trafficking properties of organelles such as the Golgi apparatus [183]. Non-uniform curvature is also found in the capsid of the torovirus, which infects agricultural animals [168], and in the cubic phases of lipids commonly used in the formulation and food industries [178]. Furthermore, engineered realisations of surfaces with desired curvatures are becoming increasingly achievable. Techniques include a rotating cuvette [13], lithography [14], suspending a liquid surface from a post [15], and 3D printing [16]. These surfaces could play host to a variety of two-dimensional systems, including membrane proteins [183], stress fibres [249], active and passive liquid crystals [9, 10], capsid protein shells [138, 168] and colloidal particles adsorbed onto a surface by depletion forces or tethered to one with DNA [11, 19]. Additionally, it

has been shown that the dimple patterns in buckled curved elastic bilayers form a crystal structure with similar properties to colloidal crystals on a curved surface [171].

The practical importance of these systems, as well as their rich and novel behaviour, mark them out for considerable scientific interest. As discussed in detail in Sec. 1.3 much of this behaviour emerges from the broken symmetry of the surface. For example, nucleating phases form preferentially in certain regions due to the underlying curvature, and the free energy profiles may include additional metastable minima [18,177]. Crystal defects feel a local potential that arises purely from the underlying curvature [167,171]. Topological defects also have preferential locations that depend on their type, as shown for nematic liquid crystals [13] and colloidal particles [120,169].

To date, however, there is still no complete picture of how non-uniform curvature affects the thermodynamics of two-dimensional systems. In this chapter, we identify three effects—relating to cluster perimeter, local strain and the energetics of packing—and their impacts on the phase behaviour of attractive colloids. The major consequences of these considerations are that the stable phases (gas, liquid, crystal) are located in different regions of the surface, and phase transitions must therefore involve global translation of the colloidal cluster. It is even possible for new states to emerge, producing states with the same phase of matter but different locations. It is remarkable that such a novel behaviours can be induced in particles with such simple interactions. While curvature sensing has previously been reported, in most cases it is due to the architecture or anisotropy of the particle or the molecule involved [2,250]. In contrast, here curvature sensing is due to a cooperative effect. Moreover, the effect is phase-dependent: it acts differently for liquid and crystal phases.

In this chapter we first present the particular methods that we use to study surfaces of non-uniform curvature. Then, taking as our primary example a torus occupied by 300 particles, we discuss the basic phenomenology of such systems. As discussed in Sec. 1.2.2, the behaviour of crystals of attractive particles on curved surfaces is known to depend on the softness of their interactions [19,138], so we vary the interaction range as well as the temperature. Using bespoke Monte Carlo and basin-hopping methods, we identify four localised states and explain each in terms of the three effects listed above. We present a phase diagram, showing at which temperatures and interaction-ranges each state is stable, and we measure the free energy barriers between them. Using molecular dynamics simulations, we plot dynamic trajectories which show the system switching between states *via* the translation of the colloidal assembly.

In Sec. 4.5, we quantify the three effects, showing that each contributes to the free energy, and that competition between them drives each of the transitions.

In Sec. 4.6, we show that the effects of non-constant curvature are universal by presenting simulation results for a variety of surfaces and particle densities. Firstly we investigate the example torus but with 100 and 500 particles, then a more slender torus and a sinusoidal surface with 300 particles. In each case we show that localised states occur, and though the details of the states vary from surface to surface, each can be explained by the same three effects listed above.

4.2 Methods

4.2.1 Model System

We model a system of attractive, spherical particles whose centres are confined to a curved surface. To achieve the adjustable range of interaction discussed in Sec. 4.1, we use a truncated, shifted and smoothed (tss) Morse potential, re-scaled to restore its well-depth ε from before the shift. This potential has an adjustable range parameter ρ . The pair potential energy is given by

$$\begin{aligned} U(r) &= -\varepsilon U_{\text{tss}}(r) / U_{\text{tss}}(r_0), \\ U_{\text{tss}}(r) &= \left[U_{\text{M}}(r) - U_{\text{M}}(r_{\text{cut}}) - (r - r_{\text{cut}}) \frac{dU_{\text{M}}}{dr} \Big|_{r_{\text{cut}}} \right] \Theta(r_{\text{cut}} - r), \\ U_{\text{M}}(r) &= \varepsilon e^{-\rho(r-r_0)} (e^{-\rho(r-r_0)} - 2), \end{aligned} \tag{4.1}$$

where r is the separation of the two particles, measured through three-dimensional space, rather than along the geodesic (see Fig. 4.1A); r_{cut} is the cut-off distance at which the potential is truncated; r_0 is the equilibrium separation of two particles; and $\Theta(x)$ is the Heaviside function [138]. For this work we have chosen to set $r_{\text{cut}}/r_0 = 2.23$, in line with comparable investigations [138, 140] and the potential used in chapter 3. A graph of this potential is shown in Fig. 4.1B.

As in chapter 3, all quantities are reported in reduced units. Except where notated otherwise, energy is measured in units of ε , distance in units of r_0 and temperature in units of ε/k_{B} , where k_{B} is Boltzmann's constant.

We have elected to use a toroidal surface as our primary example of a surface of varying

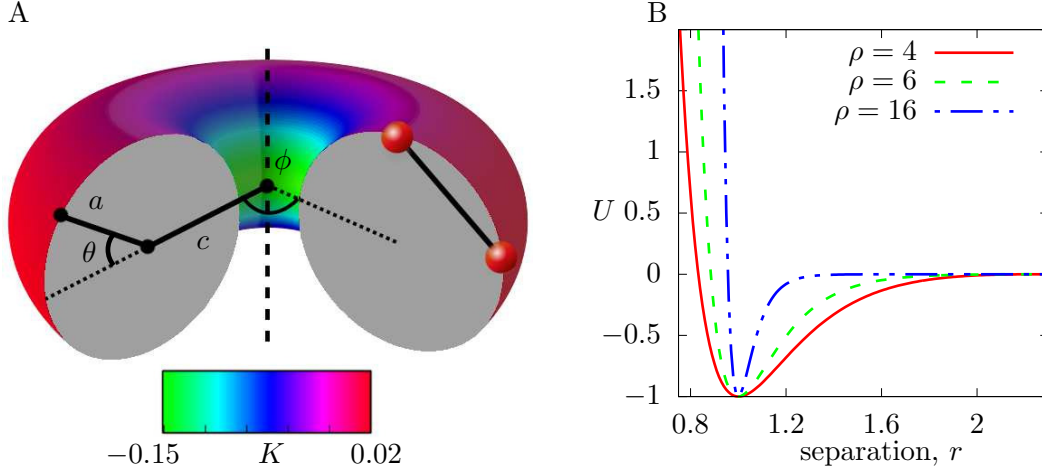


Figure 4.1: A: The torus surface, coloured by the local Gaussian curvature K in units of $1/r_0^2$. The two size parameters a and c are labelled on the cross section. In most of this work we use $a = 5$ and $c = 7$. The polar angles θ and ϕ are used to specify the location of particles on the surface. Also depicted are two particles and the line through which the Morse potential acts. B: The truncated, shifted, smoothed and scaled Morse potential for some representative values of ρ .

curvature for the following reasons: it is relatively simple, both mathematically and conceptually; it has regions of positive and negative Gaussian curvature; some progress is being made in reproducing it experimentally, for example with Pickering emulsions, as discussed in Sec. 1.3 [13, 251, 252]; and related surfaces can be found in nature (for example the torovirus capsid mentioned above [168]). Fig. 4.1A shows a schematic of the torus, defined by the equation

$$\left(c - \sqrt{x^2 + y^2}\right)^2 + z^2 - a^2 = 0,$$

where x, y, z are Cartesian coordinates, and a and c are geometrical parameters. We parameterise the surface with curvilinear coordinates θ and ϕ such that

$$\begin{aligned} x &= [c + a \cos(\theta)] \cos(\phi) \\ y &= [c + a \cos(\theta)] \sin(\phi) \\ z &= \sin(\theta). \end{aligned}$$

The metric g factor for this surface (required for the Monte Carlo simulations, see Sec. 2.2.7) is $g = c + a \cos(\theta)$.

In addition in Sec. 4.6.3, we present results for simulations performed on a sinusoidal

surface with periodic boundaries, defined by

$$z = h \sin\left(\frac{2\pi x}{L}\right) \sin\left(\frac{2\pi y}{L}\right),$$

where L is the length of the periodic box and h is the maximum height of the surface. In this case, we simply use x and y as our coordinate system. The metric g factor is

$$g = \sqrt{\left(\frac{2\pi h}{L}\right)^2 \left[1 - \cos\left(\frac{4\pi x}{L}\right) \cos\left(\frac{4\pi y}{L}\right)\right]} + 1.$$

4.2.2 Simulation techniques

In this chapter, we aim to identify unique thermodynamic states and investigate the transitions between them. To do this we make use of canonical Monte Carlo simulations, molecular dynamics and basin-hopping, each adapted to curved surfaces and described in detail in Sec. 2.2–2.4. Whilst the Monte Carlo and molecular dynamics simulations were performed using the potential above, the basin-hopping simulations used the untruncated Morse potential¹.

We have selected two reaction coordinates for our free energy calculations (see Sec. 2.2.1), so that we can track both the location and the phase of the system. As the states we wish to identify are associated with specific regions of the surface, which can be identified by their curvature, our first order parameter is the ratio of the density of particles in regions of positive Gaussian curvature to the overall density n_o/n . The second order parameter quantifies the structural order by counting the number of particles in a crystal-like environment, N_X . To determine whether a given particle is crystalline, we select all particles within $r = 1.45$ of it. These M neighbours are then gnomonically projected onto the plane tangent to the surface at the target particle [253]. In a gnomonic projection, a point on the surface is moved along the surface normal until it lands on the tangent plane. The magnitude of the bond-orientational order parameter is then calculated using

$$|\psi| = \left| \frac{1}{6} \sum_{j=1}^M e^{6i\theta_j} \right|,$$

where θ_j is the angle between the target particle, the j^{th} neighbour and some arbitrary

¹The basin-hopping calculations in this section were performed by Mark Miller

reference direction [73]. ψ does not depend on the choice of reference direction. Particles with $|\psi| > 0.6$ are considered crystalline. It is important to note that this definition of $|\psi|$ will not count particles on the edge of the cluster as crystalline. We construct the free energy surface by accumulating a two-dimensional histogram H in both order parameters and taking $F = -k_B T \ln(H)$ in this case, satisfactory convergence was achieved without umbrella sampling. We construct phase diagrams by measuring free energy surfaces at a range of ρ and T , and identifying the state associated with the lowest minima in each.

4.3 Basic phenomenology

Surface curvature can influence the free energy of a cluster on a surface in three ways:

1. As discussed in chapter 3, the length of the perimeter of a cluster of a given area depends on the underlying curvature, changing the line tension energy term [18]. In regions of positive Gaussian curvature, the perimeter is generally smaller than on a flat or negatively curved surface. We will refer to this as the “perimeter” effect.
2. A hexagonal crystal structure is stressed in regions of non-zero Gaussian curvature, penalising the free energy [19, 138]. We will call this the “strain” effect.
3. There is a small potential energy gain which affects both liquid and crystalline clusters on a curved surface by making next-nearest-neighbour interaction more favourable. For example, where the Gaussian curvature is large and negative the area within a given distance of a particle is larger than on a plane, allowing more particles to be found close to it. Furthermore, if, as in our model, the interactions act through real space rather than geodesically along the surface, any region with large principle curvature will bring its next-nearest-neighbours closer together. This effect is present even if the Gaussian curvature is zero (for example on a cylinder or a cone). We call this the “packing” effect.

In the remainder of this chapter, we will show that these effects come together to stabilise localised states in a number of systems confined to curved surfaces. We have developed a compact naming scheme for these states, summarised in Table. 4.1. We will show that the specific nature of these states is dependent on the details of the surface and the number of particles, but all can be explained with reference to the three effects listed above.

Phase	Position
G gas	+ positive curvature
L liquid	0 zero curvature (and vicinity)
C condensed	− negative curvature
X crystal	± both positive and negative curvature

Table 4.1: Each state apart from uniform gas is named with two symbols. The first represents the thermodynamic phase. The “condensed” phase has intermediate liquid and crystalline character as discussed in the text. The second symbol denotes the Gaussian curvature of the region where the state is found (see Fig. 4.1A for a map of the Gaussian curvature of the 5-7 torus). As the gas state typically covers the whole surface, it does not need a position specifier and is noted by the single symbol G.

4.3.1 Localised states on a torus

To explore the effects outlined above, we will first look in detail at a single example system. We will use a torus with $a = 5$ and $c = 7$, as shown in Fig. 4.1A. We refer to this as a 5-7 torus. There will be 300 particles absorbed onto the surface. We start with a brief survey of the phenomenology in this case study.

The three effects combine to stabilise the four states shown in Fig. 4.2A-D. These are G, L−, C+ and X0. At high temperatures the system exists in the gas(G) state. For long-range potentials, lowering the temperature leads to a gas–liquid transition. For a liquid, the packing effect described above dominates, so the cluster is always found in the centre of the torus, where the mean curvature is largest. This is the L− state. As the temperature is decreased further, the crystallinity of the cluster increases and the perimeter and strain effects become more important. Once a threshold of crystallinity is reached, the cluster adopts the C+ state, moving to the outside of the torus, where the Gaussian curvature is lower but where it can still adopt a compact shape to minimise the line tension energy. Close to the transition temperature, the cluster is still liquid-like, but as the temperature is decreased further, the cluster becomes continuously more crystalline. If the range of the potential is then decreased, another transition is reached, to the X0 state. Here, the strain effect dominates, since the elastic modulus depends sensitively on the curvature of the potential at its minimum. As a result, the crystal forms a ribbon on the top of the torus (where the Gaussian curvature, and therefore stress, is lowest). Heating this crystal causes it to sublime back to the gas(G) state. In Fig. 4.2E we present a phase diagram in temperature and interaction range, which shows the locations of the transitions between these states.




Figure 4.2: A-D: Snapshots of the four states of 300 particles on a 5-7 torus. A: G, at $\rho = 6$, $k_B T/\varepsilon = 0.72$; B: L-, at $\rho = 6$, $T = 0.42$; C: C+, at $\rho = 4$, $T = 0.32$; and D: X0, at $\rho = 18$, $T = 0.27$. Particles are coloured by the number of nearest neighbours. The states are named according to the convention in Table 4.1. E: Phase diagram for 300 Morse particles on a 5-7 torus. The saturation of the blue symbols represents the crystallinity, showing a steady increase as the C+ phase is cooled.

4.3.2 Basin Hopping

It can be seen in Fig. 4.2E that the phase boundary between the C+ and X0 states seems to be independent of temperature. This raises the question of how the system behaves as the temperature approaches zero. Lower temperatures are inaccessible to Monte Carlo simulations, due to the increased difficulty of escaping free energy minima. Instead, we turn to basin hopping parallel tempering (BHPT) to find global potential energy minima. We have performed BHPT runs for long- and short-ranged Morse potentials ($\rho = 4$ and $\rho = 18$, respectively). Snapshots of the global minima are found in Fig. 4.3 where it can be seen that the distinction between C+ and X0-like structures is maintained even in the ground state. For the long-ranged potential, the putative global minimum is found on the “side” of the torus, where it is compact. Such a configuration requires the inclusion of some defects, but this comes with a relatively low energy cost due to the softness of the potential. For the short-ranged potential, the energy cost for distortion of the hexagonal structure is high, so the crystal is found on the “top” or “bottom” of the torus, where the curvature is lowest, although this incurs a cost due to the longer perimeter. A quantitative breakdown of these contributions to the energy is given in Sec. 4.5.2.




Figure 4.3: Putative global potential energy minima for $\rho = 4$ (long-ranged Morse potential) and $\rho = 18$ for 300 particles on a 5-7 torus. Each row depicts the structure from two perspectives. Particles are coloured by coordination number. The potential energy V is given in units of ε and rounded to one decimal place.

4.3.3 Free energy and dynamics

Next, we investigate the character of the phase transition associated with the phase diagram in Fig. 4.2E. In Fig 4.4, we present the free energy surface at three different points on the phase diagram. Each chosen point lies close to a phase transition and shows two minima, indicating that two states coexist. Fig. 4.4A is taken close to the transition between the L- and C+ phases. The transition involves both a dramatic shift of particles towards the outside of the torus and a slight increase in crystallinity (although at this temperature, the C+ phase is still liquid-like, see Fig. 4.2E). Fig. 4.4B shows the vicinity of the transition from the C+ state to the X0 state. As the temperature here is cooler, the C+ phase is more crystalline. The figure shows that this transition is accompanied by a slight increase in crystallinity and a shift of particles away from the outside of the torus. Finally, Fig. 4.4C shows the coexistence between the G and X0 states, which only involves a change in crystallinity as both states are distributed fairly evenly between the inside and outside of the toroidal surface.

The presence of free energy minima separated by barriers implies that the system's dynamics should involve rapid motion within the basin of a minimum, with infrequent switches between minima. As the system switches quickly between localised states, we should observe




Figure 4.4: Free energy surfaces for 300 Morse particles on a 5-7 torus with minima highlighted by red dots. The insets show dynamic trajectories switching between the minima, with horizontal lines denoting the minima. Time is measured in reduced units $t^* = t\sqrt{(\varepsilon/m)}/r_e$, where t is the real time and m is the mass of a particle. A: At $\rho = 6.0$ and $T = 0.39$. The red dots indicate the L- (left) and C+ (right) states. B: At $\rho = 7.0$ and $T = 0.25$. The red dots indicate the X0 (left) and C+ (right) states. C: At $\rho = 20.0$ and $T = 0.30$. The red dots indicate the G (bottom) and X0 (top) states.

cooperative motion across the surface of the particles in the cluster. We used molecular dynamics simulations to observe transitions between the four coexisting states. For each free energy surface in Fig. 4.4, we performed a single simulation at fixed temperature and ρ and monitored the order parameter most relevant to the transition in question. The inset of each panel shows part of the resulting equilibrium time trace of the order parameter. In each case, the trace shows rapid switches of the system between two states, with comparatively long residence times within the states.

4.4 Bond Order Histograms

In this chapter we have used the local bond order parameter, defined in Sec. 4.2.2, to determine if a particle is in a crystalline environment. In this section, we discuss what the distribution of $|\psi|$ at a given state point can tell us about the microstructure of the system and how this changes around the state transitions.

Fig. 4.5A shows histograms of the bond order over a range of temperatures which span the G, L- and C+ states. It can be seen that each histogram has 7 peaks or steps. Each of these, from left to right, corresponds to particles with 0-6 neighbours, arranged in maximally crystalline order. Particles with the same number of nearest neighbours, but some deviation from hexagonal order, contribute to the left of that peak.

In the gas state (for example at $k_B T/\varepsilon = 0.56$), there is a sharp peak around $|\psi| = 0.33$,

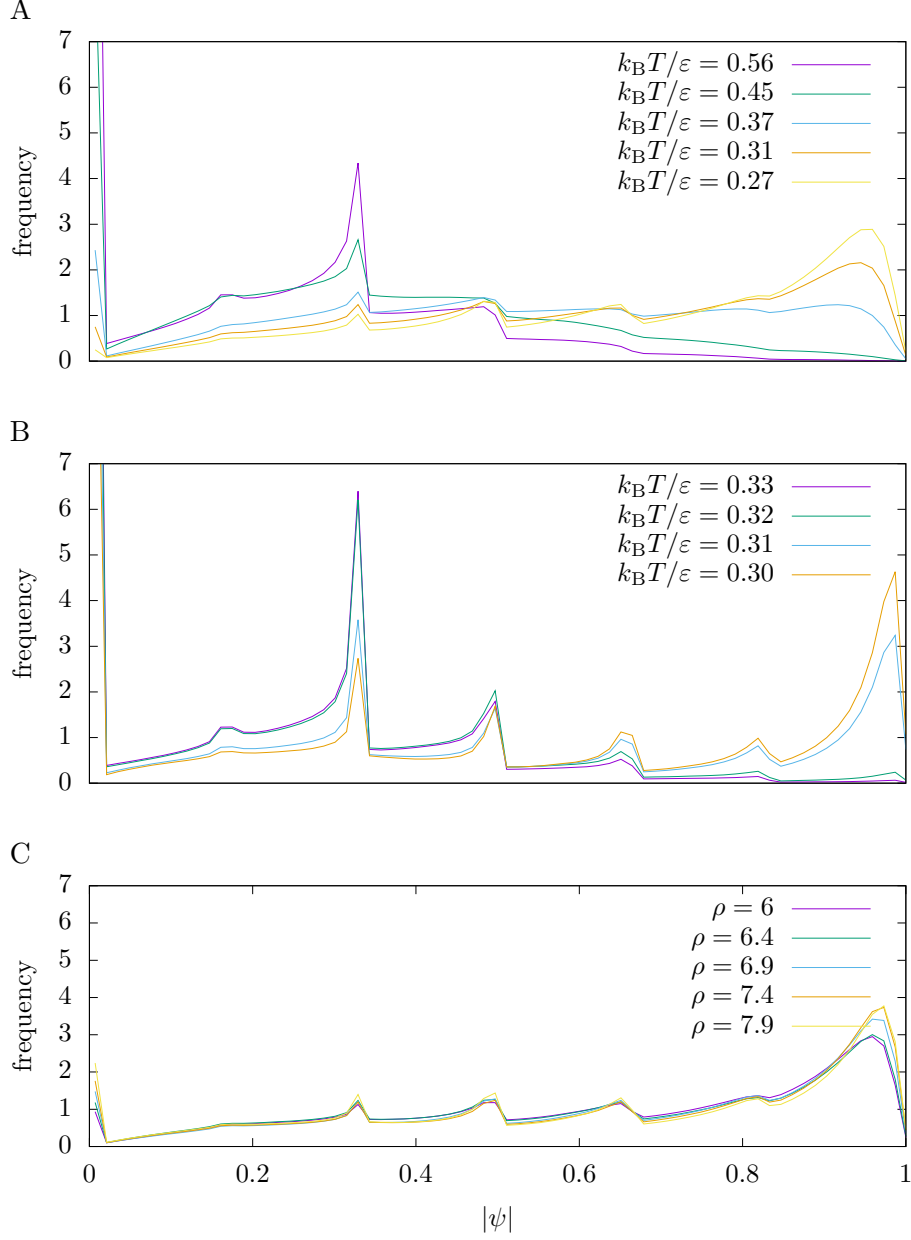


Figure 4.5: Normalised histograms of the per-particle bond order parameter defined in Sec. 4.2.2, taken from simulations of 300 particles on a 5-7 torus in the canonical ensemble. A: At $\rho = 5$ over a range of temperatures. The crystalline peak (just below $|\psi| = 1$), grows continuously and gradually. It is fairly broad. B: At $\rho = 16$ over a smaller range of temperatures. The crystalline peak emerges rapidly and is sharp. C: At $k_B T / \epsilon = 0.29$ over different potential ranges. The profiles are similar, although the peaks get slightly sharper.

which represents particles with two nearest neighbours in good order. This suggests that the second most common motif (other than isolated particles) is a group of three particles arranged in a triangle.

As the temperature is reduced, this peak decreases rapidly as the particles condense into the liquid phase (see histogram $k_B T/\varepsilon = 0.37$). The histogram becomes relatively flat, with less pronounced peaks as particles have a range of neighbours and little crystalline order. It is notable that this change happens smoothly, meaning the G to L- transition cannot be pinpointed from the bond order distribution.

As the temperature drops further, the peak at $|\psi| = 1$ becomes more dominant, indicating that the cluster has an increasingly dense, crystalline structure. In addition, all the peaks become sharper, suggesting the presence of particles with crystalline order but missing neighbours. These correspond to particles at the edges of the cluster. As the potential is soft, the peaks are quite broad, so the structure still contains some disorder. We would expect that at even lower temperatures, they would become sharper.

Fig. 4.5B shows histograms taken at $\rho = 16$ over a much smaller range of temperatures around the G to X0 transition point. It can be seen that this transition is much sharper. The profile of the gas phase is similar here to $\rho = 5$. The crystal phase features much sharper peaks, indicating strong crystalline order. Most particles have six neighbours, but there are also clear peaks for the boundary particles at three, four and five neighbours.

In addition, we see significant peaks at the zero and one neighbour positions, unlike for the crystal phase at $\rho = 5$. This suggests that the crystalline cluster is coexisting with a fair amount of gas. We would expect these peaks to vanish at lower temperatures.

Fig. 4.5C shows histograms taken at a variety of values of ρ around the C+ to X0 transition. Although the peaks become slightly sharper at higher ρ , the effect is small. This shows that the C+ to X0 transitions, which involves a major relocation and rearrangement of the structure, has little effect on the microstructure of the cluster.

We have seen that the bond-order distribution reflects the microstructure of the system, but, significantly, cannot be used in isolation to identify the state. This is why we need the second, positional, order parameter to understand these transitions (see Fig. 4.4).

4.5 Quantitative analysis

In this section, we will analyse and quantify how competition between the three curvature effects drives the localisation transitions outlined in Sec. 4.3.1.

4.5.1 The L− to C+ transition

In the liquid state, the two contributions to the free energy affected by curvature are the line energy (the perimeter effect) and the additional second-nearest neighbour energy due to mean curvature (the packing effect). In this section, we will examine how these two effects compete with each other and contribute to the L− to C+ transition on a 5-7 torus. We show that the number of particles on the surface strongly affects the interplay of these two contributions.

We acquired values for the perimeter of the C+ state, l_{C+} , and the L− state l_{L-} using the surface minimisation program Surface Evolver [254]. This allowed us to define a patch on the torus, constrained to maintain a fixed area A , and then optimise the shape to find the smallest possible perimeter. By carefully choosing the symmetry of the initial patch, we were able to obtain the perimeter of the liquid patch as a function of its area for the C+ and L− states. These are plotted in Fig. 4.6. It can be seen that the perimeter, and therefore line energy, of the C+ state is always lower than that of the L− state, which means that the perimeter effect works to destabilise the L− state. The difference $\Delta l = l_{C+} - l_{L-}$ is also plotted, showing that the perimeter-based penalty for the L− state increases roughly linearly with patch area.

We are able to quantify the potential energy contribution of the packing effect by using a grand canonical simulation [203]. Here, the number of particles is not fixed, but rather a constant chemical potential μ is applied. For convenience, we express this as an activity

$$z(\mu) = \frac{A_0}{\Lambda^2} \exp\left(\frac{\mu}{k_B T}\right), \quad (4.2)$$

where A_0 is the total area of the toroidal surface and λ is the thermal de Broglie wavelength, as in Sec. 3.2.3. This allows us to ensure that the surface is completely covered with a uniform, equilibrium liquid, so we can calculate bulk properties. At each measurement step, we bin the particles by their θ coordinate and calculate each particle's energy, collecting an average for each bin over the course of the simulation. The average energy of a particle as a




Figure 4.6: A: The minimal perimeter of the liquid–gas interface for the C+ and L– states as a function of the area of the liquid region on a 5-7 torus. B: The length of the C+ perimeter relative to that of L–, $\Delta l = l_{C+} - l_{L-}$. As this difference is always negative, the line energy of the C+ state is always lower than that of the L– state.

function of its θ coordinate $E(\theta)$ is shown in Fig. 4.7 for an example state point. It can be seen that particles at high θ (the region of high negative Gaussian curvature in the centre of the torus) have lower energy than those in the less curved regions. Hence, the packing effect favours the L– state.

The aforementioned arguments show that the perimeter and packing effects are in competition. However, to truly isolate the separate free energies associated with the perimeter and packing effects, we would need to compute the line tension for a truncated Morse potential in two-dimensions, the entropic contribution to the packing effect, and the emerging strain effect, which would be difficult to do accurately. Nonetheless, to illustrate the complex interplay between the two effects, we can perform overall free energy calculations (directly from the probability distributions of the order parameters) for a range of N from 100 to 300 for $\rho = 6$ and $T = 0.37$, where the L– and C+ states are in coexistence. Fig. 4.8 shows the free energy curves. Interestingly, these lines cross twice, emphasising the nonlinearity of the perimeter and packing effects with respect to N .

4.5.2 The C+ to X0 transition

We have shown that, at low temperature, as ρ increases, there is a transition between two crystal-like states, C+ and X0. We argue that this transition is driven by competition between line energy (the perimeter effect) and crystal frustration (the strain effect). As both states are crystalline and can exist at low temperatures and both effects are primarily energetic, it




Figure 4.7: The average potential energy of a particle in the liquid state as a function of its toroidal angle θ (see Fig. 4.1), taken from a grand canonical simulation with range parameter $\rho = 6$, reduced temperature $T = 0.48$ and activity $z = 57.544$ on a 5-7 torus.




Figure 4.8: The depth of the free energy minima of the L- and C+ states for $\rho = 6$ at reduced temperature $T = 0.37$ on a 5-7 torus. The origin of the vertical scale is arbitrary; only comparisons between the two lines matter.

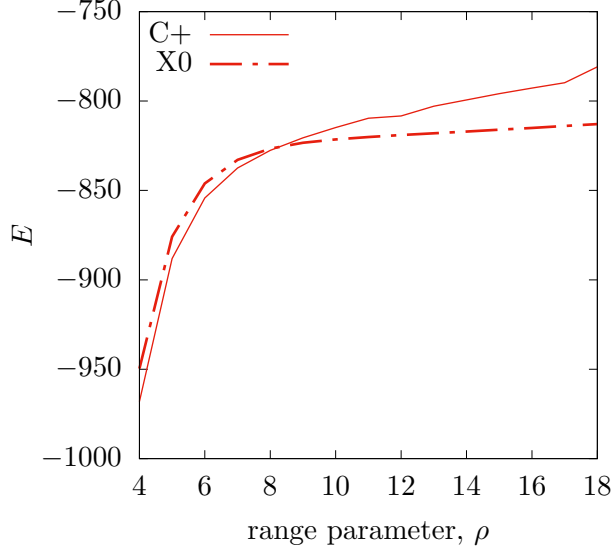


Figure 4.9: The energy of the C+ and X0 potential energy minima as a function of ρ for 300 particles on a 5-7 torus.

is convenient to study the zero-temperature case, as we can disregard entropic contributions. In this section, we quantify the energetic contributions to the ground-state structures on the 5-7 torus at $\rho = 4$ and $\rho = 18$ for the case of $N = 300$ particles, as presented in Fig. 4.3.

As seen in Fig. 4.3, the ground-state structure at $\rho = 4$ lies in the region of positive Gaussian curvature, and is characteristic of the C+ thermodynamic state, whereas at $\rho = 18$ it is in the flat region, like the X0 state. Starting from each of these two configurations, we have performed a simple local minimisation at values of ρ in the range $4 \leq \rho \leq 18$. This produces a local potential-energy minimum for both C+ and X0 structures at each value of ρ . These minimisations are performed carefully to ensure that we track the energy of the correct minimum without allowing it to collapse if it is no longer the global minimum. We observe that both minima continue to exist over the entire range of ρ studied.

The energies of these two sequences of structures are plotted as a function of ρ in Fig. 4.9. In line with our analysis, the energy of the C+ state is lower than the X0 state at lower ρ , but the lines cross and the X0 state becomes more stable at high ρ . Nevertheless, metastable minima of X0 and C+ structures persist in the two regimes, respectively.

The energy of each minimum can be broken down into the component contributions

$$E = -3N\varepsilon + E_{\text{line}} + E_{\text{defect}} + E_{\text{stress}} + E_{\text{non-nn}}, \quad (4.3)$$

which we now describe in turn.

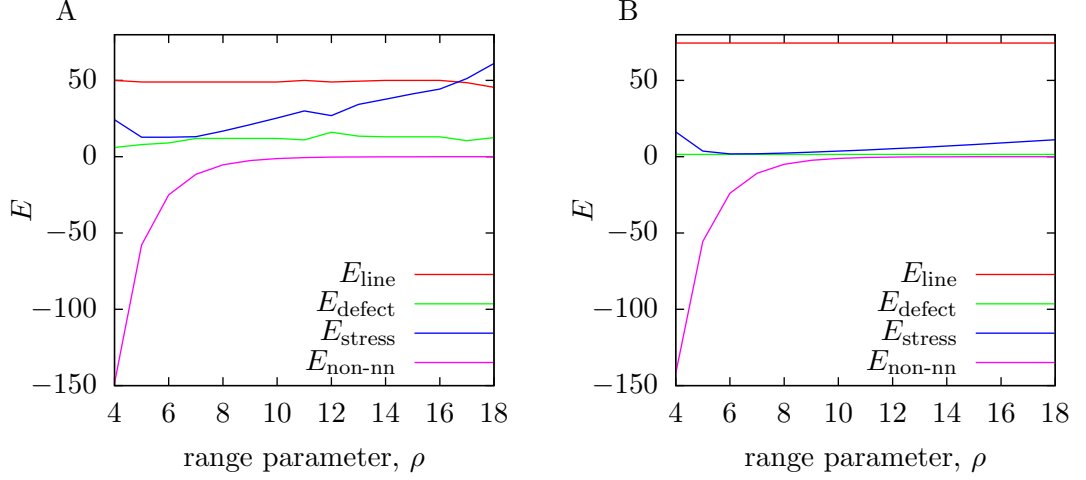


Figure 4.10: A: The components of the potential energy for the C+ state for $N = 300$ particles on the 5-7 torus, as a function of the potential range parameter ρ . B: The components of the potential energy for the X0 state.

In a perfect hexagonal crystal, the interaction energy of any given particle with its immediate neighbours would be -6ϵ . The first term on the right-hand side of Eq. 4.3 gives this ideal energy for N particles after correcting for double-counting of pairs.

The line tension energy term in Eq. 4.3 is given by

$$E_{\text{line}} = \frac{\epsilon}{2} \sum_{\text{edge}} (6 - N_{\text{nn}}), \quad (4.4)$$

where the sum is over particles on the edge of the cluster. N_{nn} is the number of nearest neighbours (nn) of each particle in the sum (*i.e.*, the number of particles lying closer than $r = 1.45$) and the factor of $1/2$ corrects for double-counting of bonds. We identify particles on the edge using a test particle method. For each pair of nearest neighbours, a test particle is inserted on the surface at the two locations where it is at a distance r_e from both real particles in the pair. If the test particle overlaps with another real particle in exactly one of these locations then both particles in the pair are designated as lying on the edge of the cluster.

The defect energy term in Eq. 4.3 corresponds to energy contributions from particles that have a number of nearest neighbours different from six,

$$E_{\text{defect}} = \frac{\epsilon}{2} \sum_{\text{not edge}} (6 - N_{\text{nn}}), \quad (4.5)$$

where the sum is over all particles not on the edge of the cluster.

The stress energy term in Eq. 4.3 captures the excess energy of nearest-neighbour pairs whose separation is different from the equilibrium distance,

$$E_{\text{stress}} = \sum_{\text{nn}} (U_{\text{M}} + \epsilon),$$

where the sum is over pairs of nearest neighbours.

Finally, the non-neighbours energy term in Eq. 4.3,

$$E_{\text{non-nn}} = \sum_{\text{non-nn}} U_{\text{M}},$$

captures the interactions between all pairs of particles that are not nearest neighbours.

This breakdown is presented for the C+ and X0 states as a function of ρ in Fig. 4.10. In Fig. 4.11 we also plot the energy difference between C+ and X0 for E_{line} (the perimeter effect), $E_{\text{defect}} + E_{\text{stress}}$ (the strain effect), and $E_{\text{non-nn}}$. In agreement with the discussion above, the X0 structures have no defect energy and are only weakly stressed, but at the expense of a more elongated shape with a higher line energy. In contrast, for the C+ state, the line energy is lower, but it has higher E_{defect} and E_{stress} . For both structures, the line energy in these zero-temperature structures does not vary much with ρ because the structures do not undergo significant rearrangement when relaxed at different ρ .

$E_{\text{non-nn}}$ varies with ρ , especially for low ρ where the range of interaction is larger than the typical equilibrium distance between any pair of neighbouring particles. However, $E_{\text{non-nn}}$ has a similar dependence for the X0 and C+ states, leading to a small $\Delta E_{\text{non-nn}}$. Thus, we conclude this term does not play a major role in the transition between the X0 and C+ states.

Overall, the increase in $E_{\text{defect}} + E_{\text{stress}}$ (the packing effect) against nearly static differences in E_{line} (the perimeter effect) and $E_{\text{non-nn}}$ drives the transition from C+ to X0 crystal-like structures with increasing ρ .




Figure 4.11: The change in potential energy, $\Delta E = E_{C+} - E_{X0}$, from the X0 to the C+ state for the perimeter effect (line energy), the strain effect (defects and stress), and the non-nearest-neighbour interactions ($\Delta E_{\text{non-nn}}$) for 500 particles on a 5-7 torus.

4.6 Varying the number of particles and geometry

In this section we demonstrate both the universality of the effects described above and the different ways they can manifest in different systems. We do this by exploring the phase diagrams of systems of 100 and 500 particles on the 5-7 torus, the phase diagram of 300 particles on a torus of the same area but different curvature, and the phase diagram of a sinusoidal surface. For each, we argue that the behaviour can be explained with reference to the same three effects, and explain the differences between each case.

4.6.1 100 and 500 particles

The phase diagram for a 5-7 torus with 100 Morse particles is shown in Fig. 4.12D, with snapshots of the three observed states shown in Fig. 4.12A-C. Although free energy calculations show that the L- state continues to exist as a metastable state, it has a higher free energy than the C+ state, and hence it does not appear in the phase diagram.

We find that, with decreasing temperature, the X0 state extends to lower values of ρ than in the case of 300 particles. This observation is consistent with the ground-state (*i.e.*, zero temperature limit) structures from BHPT shown in Fig. 4.13; at $\rho = 4$ the lowest-energy structure for $N = 100$ crystal is located in the vicinity of zero Gaussian curvature, whereas for $N = 300$ it has moved to the region of positive Gaussian curvature.

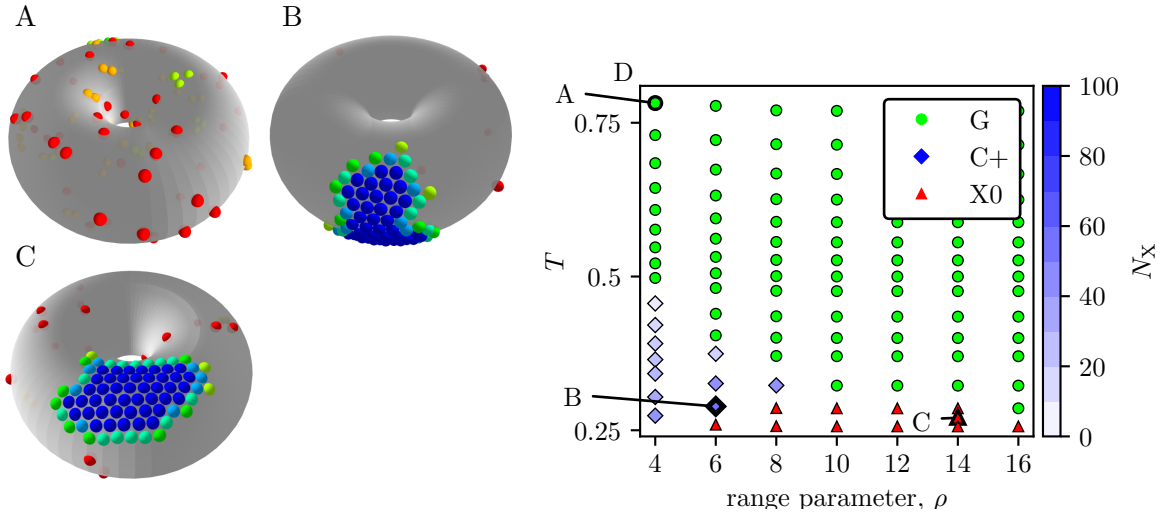


Figure 4.12: A-C: Snapshots of the states labelled in the phase diagram of Fig. 4.12D. A: G, at $\rho = 4$, $T = 0.73$; B: C+, at $\rho = 6$, $T = 0.28$; C: X0, at $\rho = 14$, $T = 0.26$. Particles are coloured by the number of nearest neighbours, using the key in Fig. 4.2. D: Phase diagram for 100 Morse particles on a 5-7 torus as a function of the potential range parameter ρ and the reduced temperature T .




Figure 4.13: Putative global potential energy minima for $\rho = 4$ (long-ranged Morse potential) and $\rho = 18$ for 100 particles on a 5-7 torus. Each row depicts the structure from two perspectives. Particles are coloured by coordination number. The potential energy V is given in units of ϵ and rounded to one decimal place.

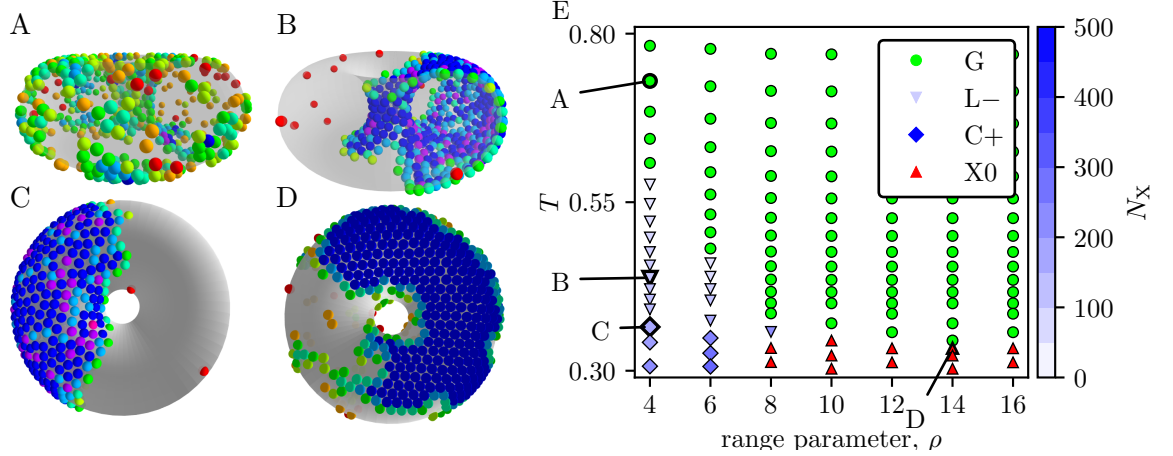


Figure 4.14: A-D: Snapshots of the states for $N = 500$ particles on a 5-7 torus. A: G, at $\rho = 4$, $T = 0.73$; B: L_{\pm} , at $\rho = 4$, $T = 0.44$; C: C+, at $\rho = 4$, $T = 0.36$; and D: X0, at $\rho = 14$, $T = 0.33$. Particles are coloured by the number of nearest neighbours, using the key in Fig. 4.2. E: Phase diagram for 500 Morse particles on a 5-7 torus as a function of the potential range parameter ρ and the reduced temperature T .

Moreover, unlike the case of 300 particles, the boundary between the C+ and X0 phases for $N = 100$ is not independent of temperature. To explain these observations, consider a C+ state on the side of the torus for a relatively low ρ . Decreasing the temperature drives the system towards increasing crystallinity. For a small number of particles, such as $N = 100$, this can be achieved by migrating to the “top” or “bottom” of the torus for two reasons: first, the crystal in the X0 state is less stressed due to Gaussian curvature. Secondly, the difference in perimeter, and hence line energy, between the C+ and X0 states of a small cluster is small. Thus, the strain effect can dominate over the perimeter effect. In contrast, for a larger number of particles (such as for $N = 300$), the line energy penalty becomes prohibitively expensive in the X0 state, since the shape of the cluster is highly anisotropic with a large aspect ratio. Therefore, for $N = 300$, the system instead creates defects to compensate for the increasing degree of crystallinity in the C+ state with decreasing temperature.

The transition from C+ to X0 can be induced by increasing ρ , since the elastic energy for stressing the crystal and/or creating defects becomes more expensive than the line energy penalty when the potential is short-ranged. The C+ to X0 phase boundary will tend to become more independent of temperature for larger crystals.

As an example of a larger cluster, we have also investigated the case of 500 Morse particles on the 5-7 torus. Once again, the key message is that we observe a localisation of the particles caused by the curvature of the surface, but the specific details are affected by the additional




Figure 4.15: Putative global potential energy minima for $\rho = 4$ (long-ranged Morse potential) and $\rho = 18$ for 500 particles on a 5-7 torus. Each row depicts the structure from two perspectives. Particles are coloured by coordination number. The potential energy V is given in units of ε and rounded to one decimal place.

particles. In Fig. 4.14A-D we show snapshots of the four observed states. Free energy calculations confirm that these phases are separated by barriers and are therefore distinct. At high temperature, we observe a gas (now rather dense) as usual. For the liquid phase, instead of just occupying the region of negative Gaussian curvature (L- phase), the additional particles lead to a second loop that wraps around the tube of the torus, forming what we term the $L\pm$ phase to signify the fact that it covers regions both of negative and of positive Gaussian curvature. The closure of this second loop reduces the total perimeter of the cluster in comparison to an L- phase that extends too far from the centre of the torus. The C+ and X0 phases appear similar to the 300 particle case. The transition between these two states is also independent of temperature for $N = 500$, for the reasons given above. The phase diagram for a 5-7 torus with 500 particles is shown in Fig. 4.14E.

Similar to the case of 300 particles, the ground state for the long-range potential is found on the “side” of the torus, and for the short-ranged potential it is found on the “top” or “bottom”. Snapshots can be found in Fig. 4.15.

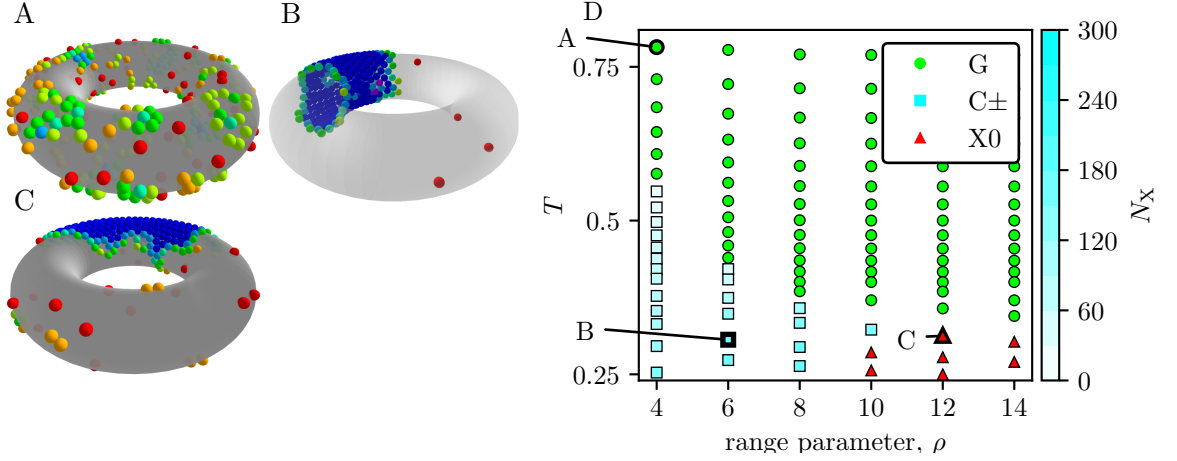


Figure 4.16: A-C: Snapshots of the states labelled in the phase diagram of Fig. 4.16. A: G, at $\rho = 4$, $T = 0.78$; B: $C\pm$, at $\rho = 6$, $T = 0.30$; and C: X0, at $\rho = 12$, $T = 0.30$. Particles are coloured by the number of nearest neighbours. D: Phase diagram for 300 Morse particles on a 3.5-10 torus as a function of the potential range parameter ρ and the reduced temperature T .

4.6.2 A slender torus

In this section, we investigate the 3.5-10 torus and show that its phase diagram has some different features from that of the 5-7 torus. The 3.5-10 torus has the same surface area as the 5-7 torus, but consists of a longer, thinner tube. We observe three states, shown as snapshots in Fig. 4.16A-C. These are a gas (G), a condensed state wrapped around tube of the torus ($C\pm$), and an X0 state. Fig. 4.16D shows the phase diagram. The loss of the L- and C+ state in favour of the wrapped state can easily be understood in terms of line energy. As the bore of the torus is so wide, an L- state would have a large perimeter; and as the tube is so thin, a C+ state would have a very large aspect ratio. However, the thinness of the tube stabilises the $C\pm$ state as its perimeter is relatively small (and also, notably, independent of N). The frustration of the crystal is also lower, as the Gaussian curvature of this torus is relatively low. Hence, the X0 state only becomes stable at higher ρ , compared to a 5-7 torus.

4.6.3 Sinusoidal surface

In order to demonstrate that the phenomena seen on the torus can be extended to other curved surfaces, we performed simulations on the sinusoidal surface described in Sec 4.2.1. We set $h = 7.5$ and $L = 30.75$ so that the area and maximum Gaussian curvature of this surface match the 5-7 torus. Snapshots of the states of this system are presented in Fig. 4.17A-C, and the phase diagram in Fig. 4.17D.

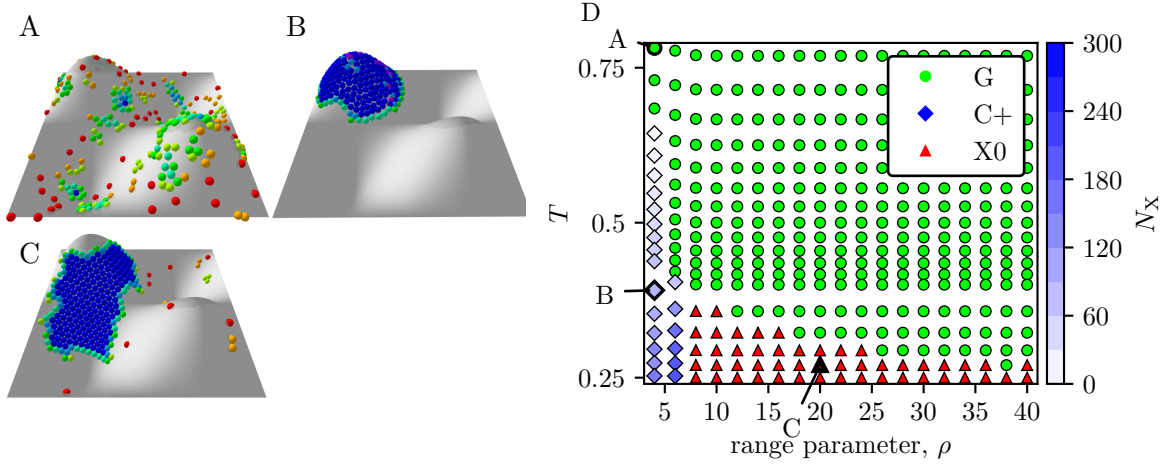


Figure 4.17: A-C: Snapshots of states on a periodic sinusoidal surface. A: G, at $\rho = 4$, $T = 0.78$; B: C+, at $\rho = 4$, $T = 0.39$; and C: X0, at $\rho = 20$, $T = 0.27$. Particles are coloured by the number of nearest neighbours (see Fig. 4.2 for scale). D: Phase diagram for 300 Morse particles on a periodic sinusoidal surface as a function of the potential range parameter ρ and the reduced temperature T .

This system has three states: G, C+ and X0. At high temperature, the system is found in the G state. For softer potentials, as the system is cooled the particles condense into a liquid phase which is always found on one of the peaks or troughs of the surface, the C+ state. This configuration is preferred for both its low perimeter length and high mean curvature (the perimeter and packing effects). As the system is cooled further, the condensed cluster becomes more crystalline. Although the contribution of frustration (in the crystal) to the free energy is increasing, at low ρ the cluster does not move again, unlike on the torus which has the L- to C+ transition. We suggest that this is because, unlike on the torus, the line tension is low enough that moving to a less curved region with a longer perimeter would not be favourable.

When cooled from the G state at higher values of ρ , the system crystallises around the flanks of the peaks, where the Gaussian curvature is lowest (the strain effect). As in the X0 phase on the torus, as well as in the branched phases of the work by Meng *et al.* [19], a longer perimeter is traded for less frustration in stiffer crystals.

4.7 Conclusion

In conclusion, we have shown that, in curved two dimensional systems, clusters of attractive crystals minimise their free energy by adopting specific shapes and translating to specific locations. These choices depend not only on the phase of the cluster, but also on the range

of the interaction potential and the underlying curvature of the surface.

We have identified three contributions to the free energy that drive this behaviour: the length of the perimeter, the packing frustration induced by the Gaussian curvature, and the length of the second-nearest-neighbour bond. These considerations explain the four states we observe on the 5-7 torus and the three states on the sinusoidal surface. Free energy calculations also showed the barriers between these states and molecular dynamics simulations confirmed the switching between them.

We hope our work will motivate experimental demonstrations of the cooperative curvature sensing effects shown in this chapter, both in biological and artificial systems. Unlike many of the curvature sensing effects found in nature, our particles are simple spherical objects with no internal structure. In nature, curvature sensing is used to sequentially recruit proteins to the cell membrane, helping to organise self assembly [2, 22]. Our work suggests that simple, easy to engineer colloids could be used to control self-assembly in artificial systems.

Due to the generality of the effects, the non-uniform curved surface need not be toroidal or sinusoidal. There are also a number of interesting avenues for future investigations. For instance, we can study the case where the non-uniform curved surface is coupled with an anisotropic interaction potential or polydisperse mixtures. Another interesting direction is to allow the confining surface to be flexible, where the curvature responds to the behaviour of the confined particles. This is relevant for a variety of clustering and aggregation phenomena on lipid membranes [183]. We expand on these suggestions in Sec. 6.3.

Chapter 5

Crystal structures on conical surfaces

We present putative global minima for model colloidal particles adsorbed onto conical surfaces. Unlike the surfaces examined in Chapters 3 and 4, cones have no Gaussian curvature and it is therefore possible for them to host defect-free crystals. Using geometrical and symmetry arguments, we show that perfect, achiral crystals arise for certain “magic” opening angles of the cone. Away from these special cases, defective structures emerge that depend on the range of the interaction potential. These imperfect structures may be achiral, and we introduce a new continuous measure of chirality in order to quantify this observation. The zero Gaussian curvature of cones also makes it computationally feasible to compare potentials that act directly through Euclidean space and along the surface geodesics. The geodesic case tends to produce less strained structures.

5.1 Introduction

In chapter 3, we discussed the ways in which confining a colloidal system to the surface of a sphere affects the nucleation and growth of the liquid phase, compared to colloids on an infinite flat plane. One of the important differences between the spherical and planar surfaces is that a sphere has finite area, but no boundaries. This means that the liquid cluster must stop growing when it has taken all the available space, and the growing edge has vanished. Essentially, the growing cluster collides with itself. Work by other authors has shown that this effect also has important implications for the development of solid phases on the sphere. For example, the ground states of spherical crystals can be difficult to achieve due to “hole implosion” and “closure catastrophe”, which are kinetic traps caused by the interaction of the growing crystal faces with each other when the surface is nearly covered. [142,255].

As discussed in Sec. 1.2.1, investigations of the Thomson problem have shown that the ground states of crystals on a sphere are also highly influenced by the topology and Gaussian curvature of the spherical surface [17,19]. However, the effects of self-collision can also be

observed on surfaces with no Gaussian curvature but with non-zero mean curvature, so long as they have no boundary in at least one dimension. Cylinders, cones and truncated cones all have these properties [256].

One of these effects is a set of selection criteria for the chirality of cylindrical crystals [256]. A chiral system is any configuration of objects with a nonsuperimposable mirror image. Chirality has long been of great interest to scientists [257, 258]. In nature, systems often have a high degree of symmetry [259]. Examples include crystals, virus capsids, hemoglobin and some polypeptides [108, 260, 261]. This suggests that high symmetry configurations tend to have lower free energy than low symmetry configurations. Using a variety of unconstrained, three-dimensional model systems, Wales showed that the minima of high symmetry configurations have a higher variance of energies than low symmetry configurations [262]. This means that there tend to be both high and low energy outlying minima with high symmetry. In contrast to this, tubular crystals, where the particles are constrained to a two-dimensional surface, are often are chiral (i.e. asymmetric). Examples include virus capsids, bacterial flagella and various types of microtubule [263]. We argue that this divergence between the symmetries of unconstrained three-dimensional structures and cylindrical systems is caused by the constraint on the crystal orientation imposed by the finite circumference of the cylinder, discussed below and illustrated in Fig. 5.1.

Formally, chirality is a binary property; any configuration with no superimposable mirror image is chiral, otherwise it is not. However, in many cases this stringency can obscure physically important near-symmetries. Consider, for example, an achiral configuration of a large number of particles, with a single particle added on one side of the mirror plane. Formally this configuration is now chiral, but intuitively we understand that it should be categorised as having near-symmetry. An example of near-symmetry would be an orange. The fruit has an underlying rotational symmetry due to its segmentation, but as each segment can contain a different number of seeds, this symmetry is disrupted for any given orange. To account for this it is helpful to formulate a continuous “degree of chirality” which measures how distant from symmetry a given configuration is [257]. In Sec. 5.5, we present such a measure tailored to quantifying chirality in colloidal systems on curved surfaces.

In this chapter, we discuss the generalised Thomson problem for conical surfaces. These are especially interesting as, unlike cylinders, they have a varying mean curvature, and there are a number of outstanding questions about the presence of chirality and the types of defects

observed. Throughout, we will directly compare the phenomenology of conical systems to previous results for cylindrical crystals [256]. In the lab, such systems have been realised as foams inside cylinders [264], or colloids packed inside tubes, as long as all the colloids touch the edge [265, 266]. Similar experiments can be performed with fullerenes packed into carbon nanotubes [267]. To make larger cylinders, for which tens of particles can fit around the circumference, colloids can be adsorbed onto cylindrical template by depletion interactions [256].

Crystals on cylinders can be divided into three cases, according to a set of “cylinder conformation rules”, which are visualised in Fig. 5.1 [256]. In the first case, the circumference of the cylinder is an integer multiple of the length of one of the lattice vectors $r = |m\mathbf{v}_1|$. This cylinder can host a perfect, achiral crystal. In the second case, the circumference is the length of an integer combination of both lattice vectors $r = |m\mathbf{v}_1 + n\mathbf{v}_2|$. Here, the ground state is a perfect but chiral crystal. Geometrically, the reason a perfect chiral crystal can be constructed on a cylinder is that cylinders can hosts helices with constant pitch that follow geodesics. As lattice vectors must follow geodesics, and rows of the crystal must be evenly spaced, this means that a single domain chiral crystal with no defects can exist. In contrast, on a cone there are no helices which meet these criteria. Therefore, all perfect crystals on cones must be achiral.

The third cylindrical case occurs when $r \neq |m\mathbf{v}_1 + n\mathbf{v}_2|$. The resulting crystal must contain defects. This is because, when the growing faces collide having encircled the cylinder, they are not commensurate. Recent simulation and theoretical work has shown that this frustration is often relaxed through the presence of line-slip defects [258, 268–270]. These take the form of a double-row of particles with only five neighbours. Away from the ground state, these line-slips develop kinks which have both thermodynamic and kinetic origin [256]. As the lattice vectors of the colliding crystal faces on a cylinder are parallel, line-slip defects can be infinitely long. This is not true on a cone. Recent unpublished experimental realisations show that a wedge-shaped defect occurs instead, in which the colliding faces drift apart from each other until a new row of particles can fit in the gap, resulting in a series of long, triangular gaps in the crystal [271]. These observations raise the questions: under what criteria can a cone host a perfect crystal, and what structures form when these criteria are not met?

Previously in this thesis, we have assumed that all potentials act through Euclidean space rather than geodesically along the surface. This is not necessarily the case in all real colloidal




Figure 5.1: Three schemes for constructing a cylindrical crystal by cutting a rectangular net from a sheet of planar, hexagonal lattice and rolling it into a cylinder. The black arrows are integer multiples of the lattice vectors $m\mathbf{v}_1$ and $n\mathbf{v}_2$. A: as $r = |m\mathbf{v}_1|$, an achiral crystal can be made. B: as $r = |m\mathbf{v}_1 + n\mathbf{v}_2|$, the resulting crystal is chiral. C: $r \neq |m\mathbf{v}_1 + n\mathbf{v}_2|$, meaning that the crystal is not commensurate across the join when the net is formed into a cylinder.

systems, e.g. those which interact through depletion [19, 256]. Unlike on other surfaces, for example the torus, calculating the geodesics of a cone is simple and computationally fast, as they are straight lines on the flat net of the surface. Therefore, it is feasible to compare these two types of potential in this case and observe how the choice of interaction pathway affects the resulting structures. This would also be possible on a cylinder, but the cone is especially suitable for this comparison because, for a pair of particles close to the peak, the difference between their Euclidean and geodesic separation is relatively high, which will emphasise the difference between the two potentials.

In this chapter, we will expand on the research into curved surfaces with zero Gaussian curvature by investigating the ground states of crystals confined to a cone. We will use basin hopping simulations to find ground states of Morse crystals on various shapes of cones and use these results to find a set of “cone conformation rules” which describes the conditions

under which conical crystals are defective or defect-free, chiral and achiral. We will show how the choice of potential affects the types of structure found, both in terms of the range of the potential and whether the potential acts geodesically or through Euclidean space.

5.2 Geometry of a cone

Away from the peak, cones have zero Gaussian curvature. As a result, it is possible to unwrap a conical surface into a planar net, by making a straight cut from the base to the peak. It is often more convenient to work in this plane view. In Fig. 5.2 we show how the cone is parameterised in each of these views. The parameters are: h , the height of the truncated cone; h_c , the height the untruncated cone; r_b , the radius of the base of the cone; r_t , the radius of the top face of the truncated cone; R_c , the distance from a point on the edge of the base of the cone to the peak of the untruncated cone; R_t , the distance from a point on the edge of the top face of the truncated cone to the peak of the untruncated cone; θ_p the opening angle of the cone at the peak; and θ_{\max} , the angle at the peak of the net of the cone. Throughout this work we use θ_{\max} , also called the cone angle, in preference to θ_p , as it is often more intuitive to work in the net view. Other works use θ_p , so we have noted both in Fig. 5.2.




Figure 5.2: A schematic of a conical surface and its net. The relationships between the geometric parameters are given in Equ. 5.1-4.

These constants are related by:

$$h_c = \frac{h}{1 - r_t/r_b}, \quad (5.1)$$

$$R_c = \sqrt{h_c^2 + r_b^2}, \quad (5.2)$$

$$R_t = \sqrt{(h_c - h)^2 + r_b^2}, \quad (5.3)$$

$$\theta_{\max} = \frac{r_b}{r_c}, \quad (5.4)$$

$$\theta_{\max} = 2\pi \sin\left(\frac{\theta_p}{2}\right). \quad (5.5)$$

To use the net view of the cone to analyse surface crystals, we must understand the relationships between rotations and reflections of a cone and its net. Fig. 5.3 summarises these transformations. At the top of Fig. 5.3A is a cone with a cut made from top to bottom. Below it is the same cone, unrolled and without rotating, as viewed from beneath. The image of the seagull shows how a figure on the cone relates to a figure on the net. In Fig. 5.3B the same cone is cut at a different point on its circumference. The same unwrapping procedure is followed and (again, viewed from below) this results in a different net for the same cone. In Fig. 5.3C, the same net as in B is achieved by rotating the cone before cutting along the same line as A, showing that cutting the cone along a different line is equivalent to rotating it. Finally Fig. 5.3D shows that cutting a segment from the right side of the net and appending it to the left side is equivalent to rotating the cone. This is significant as it allows us to generate all possible nets of a given cone from any one net, without having to recalculate the wrapped cone itself.

The reflection transformation is summarised in Fig. 5.3E, which shows that reflecting a cone in the plane which passes through the axis of the cone and the line of the cut is equivalent to reflecting the net along its midline. The operation of reflecting the cone in any other plane can be reproduced on its net by first performing the rotation operation in Fig. 5.3D to align the desired reflection line with the midline of the net, reflecting, and then performing the reverse rotation.




Figure 5.3: The transformations between a cone and its net. The top row shows side views of the same cone, and the bottom row shows the nets that can be obtained by unwrapping the cone in different ways. To form the net, a cut is made down the dotted line. The net is viewed from below, so the outside surface of the cone is visible. A: The cone is transformed into its net. B: Cutting the cone along a different line gives a different net. C: The same net as in B can be obtained by rotating the cone prior to cutting along the same line as in A. D: The net from B can also be achieved by taking the net from A, cutting a section from the left side and appending it to the right side. Importantly, transformations B, C and D are equivalent and will be used interchangeably throughout this chapter. E: A reflection of a cone in a vertical mirror plane is equivalent to a reflection of the net. Rotating the mirror plane is equivalent to performing transformation D on the net and reflecting along the centre line.

5.3 Global optimisation on the cone

To study the phenomenology of crystals confined to a cone, we perform basin hopping calculations on systems of 100 Morse particles confined to the surface of cones of various angles. For each simulation, we ensure that the height of the cone is large enough that the surface area greatly exceeds the cluster size, so that the crystal is never constrained by the bottom edge of the cone.

In this chapter, we use the untruncated Morse potential:

$$U_M(r) = \varepsilon e^{-\rho(r-r_0)} \left(e^{-\rho(r-r_0)} - 2 \right). \quad (5.6)$$

As in previous chapters, for most of the calculations we use an interaction which passes through three-dimensional space. However, later in this chapter, we will compare these results to those for a potential that acts geodesically along the surface of the cone. This is equivalent to performing simulations on the planar net of the cone, using two-dimensional




Figure 5.4: The geodesics connecting two particles on the surface of a cone, in the net view (top) and on the cone (bottom) [272]. A: $\theta_{\max} < \pi$. All pairs of particles are connected by two geodesics, one clockwise and one anticlockwise. B: $\theta_{\max} > \pi$. The yellow and green pair are connected by two geodesics, while the red and blue pair only have one.

Euclidean distances in the potential and cyclical boundary conditions to connect the slit edges. This makes the computation of the geodesic potential fast, compared to other surfaces. There is one important consideration: if $\theta_{\max} < \pi$, any pair of particles is connected by two geodesics, one clockwise and one anticlockwise. This is illustrated in Fig. 5.4A. In most cases, at least one of these distances is so large that its contribution to the potential is negligible. However, to maintain the continuity of the gradient in the potential when two particles are diametrically opposite each other, which is required for optimisation algorithms, we include the contribution to the potential of both geodesic distances between all pairs of particles. If $\theta_{\max} > \pi$ a pair of particles can share either one or two geodesics, depending on their location, see Fig. 5.4B. This introduces another discontinuity when a pair of particles switches from having one geodesic to two, so we restrict our simulations to $\theta_{\max} < \pi$ when using this potential.

We introduced the fundamentals of basin hopping calculations in Sec. 2.4.1. Here we will discuss how these have been adapted to the cone for the Euclidean and geodesic potentials. After each Monte Carlo move there is a downhill minimisation step, which moves each particle down the potential energy gradient until a local minimum is reached. To do this while keeping the particles confined to the surface of the cone, we parameterise the surface with a pair of coordinates θ and ψ and minimise the potential with respect to these coordinates using

$$\frac{dU_{\text{M}}(r)}{d\theta} = \frac{\partial U_{\text{M}}(r)}{\partial r} \left(\frac{\partial r}{\partial x} \frac{\partial x}{\partial \theta} + \frac{\partial r}{\partial y} \frac{\partial y}{\partial \theta} + \frac{\partial r}{\partial z} \frac{\partial z}{\partial \theta} \right), \quad (5.7)$$

$$\frac{dU_M(r)}{d\theta} = \frac{\partial U_M(r)}{\partial r} \left(\frac{\partial r}{\partial x} \frac{\partial x}{\partial \psi} + \frac{\partial r}{\partial y} \frac{\partial y}{\partial \psi} + \frac{\partial r}{\partial z} \frac{\partial z}{\partial \psi} \right). \quad (5.8)$$

where x , y and z are the three-dimensional Cartesian coordinates of the particles in the pair that determines the separation r . The local minimisation algorithm requires that these gradients be continuous. This means that the derivatives of θ and ψ coordinates must also be continuous. This presents a challenge as the cone has a boundary at the top and the bottom, defining hard walls that confine the particles. To solve this, we use a cyclic coordinate ψ to represent the height of the particle on the cone. As ψ increases the particle moves to the top of the cone, smoothly change direction and comes back down.

For the potential acting through Euclidean space, the particle positions used to calculate the derivatives in Equ. 5.7 and Equ. 5.8 are

$$\begin{aligned} x &= r_b \frac{h_c - (h/2) \cos(\psi) + h/2}{h_c} \cos(\theta) \\ y &= r_b \frac{h_c - (h/2) \cos(\psi) + h/2}{h_c} \sin(\theta) \\ z &= \frac{h}{2} \cos(\psi). \end{aligned}$$

For the potential acting along geodesics, we use the same cyclical coordinates θ and ψ , but instead convert to the positions on the net, which are given by

$$\begin{aligned} X &= R_c \frac{h_c - (h/2) \cos(\psi) + h/2}{h_c} \cos\left(\frac{r_b}{R_c} \theta\right), \\ Y &= R_c \frac{h_c - (h/2) \cos(\psi) + h/2}{h_c} \sin\left(\frac{r_b}{R_c} \theta\right). \end{aligned}$$

where X and Y are two-dimensional Cartesian coordinates on the plane of the net with the origin at the apex of the cone. For each particle, we calculate its interaction with each other particle, and the closest of the two periodic images of each other particle, located at $\theta' = \theta + \theta_{\max}$ and $\theta' = \theta - \theta_{\max}$. This way each particle interacts with each other particle twice, once along each of two geodesics which they share.

5.4 Minimum energy structures

We begin by looking at a fairly soft Morse potential, $\rho = 6$, interacting through Euclidean space. We have performed basin hopping simulations to find the ground states of systems of

100 particles confined to cones with various cone angles θ_{\max} . Each cone is truncated such that $r_t = 0.1r_0$. We used 70,000 basin hopping steps at a fictitious temperature of $1.0\varepsilon/k_B$. We checked for convergence by taking the energy of the lowest minimum for each cone every 20,000 basin hopping steps, and observing that it had stopped changing.

We find three classes of structure: a crystal at the peak of the cone with no defects, a crystal around the peak with defects, and a crystal with no defects on the flank of the cone. Examples of each can be found in Fig. 5.5. The energy and class of the lowest minima over a range of cone angles are shown in Fig. 5.6.

The first class, seen in Fig. 5.5C-D, has no defects. It is found at and near to $\theta_{\max} = n\pi/3$ with n an integer from 1 to 5. We refer to these as magic angles. At the magic angles, the resulting cone has a small amount of stress, mostly concentrated towards the peak. This is due to the frustration caused by the mismatch between the geodesic and Euclidean distances close to peak, which means that particles sitting on the geometrically expected lattice sites overlap in Euclidean space (See Fig. 5.5C). At angles close to, but not exactly at a magic angle, we see that the resulting crystal is stressed throughout (See Fig. 5.5D). As suggested in Sec. 5.1 all of these crystals have at least one line of near-mirror symmetry and therefore the underlying structure is achiral.

The second class, seen in Fig. 5.5A-B is split into two sub-categories: those like 5.5A, which have a straight line of symmetrical defects connecting the peak of the cone to the base and those like 5.5B with a ragged, asymmetrical defect line. The first type has a single plane of near-mirror symmetry (which passes through the axis of the cone and the defect line), while the second type is chiral. For many cone angles, both types of structures appear amongst the five lowest energy minima. The energy difference between them is very small, suggesting that there is no systematic preference for either type, and the eventual choice comes down to fine details of the structures. This is in contrast to the suggestion that higher symmetry systems tend to have lower energy ground states, although this is not especially surprising; the stronger prediction is that higher symmetry systems will have a higher variance of energies among their minima, which will only lead to the presence of lower energy states in certain cases [262].

At the largest angles, where the cone is close to a plane, the crystal moves away from the peak onto the side of the cone, as seen in Fig. 5.5E-F. As θ_{\max} increases, the line energy saving associated with occupying the peak rather than the side decreases, until it is less than

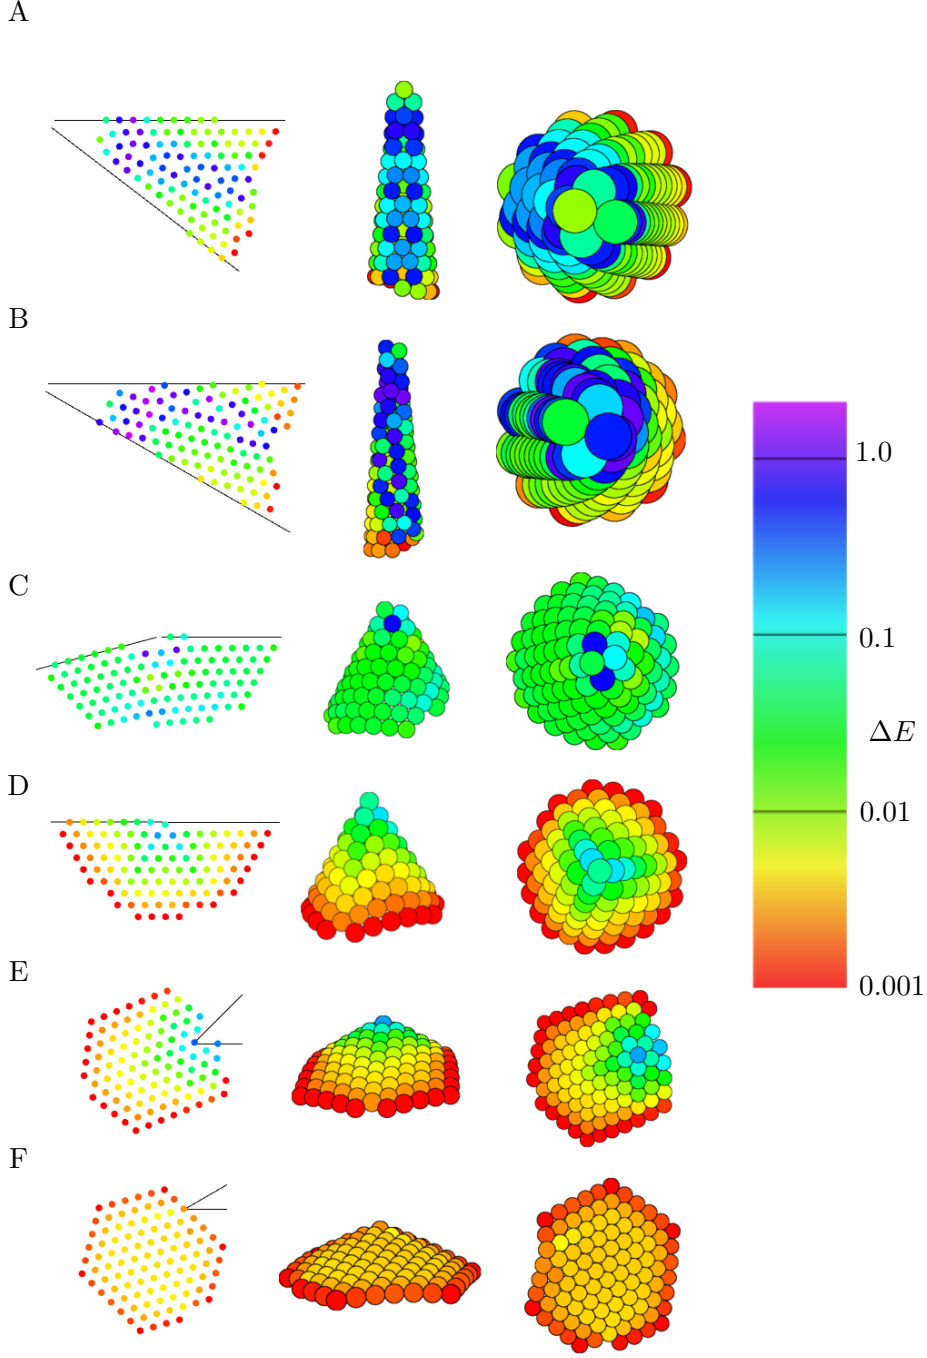


Figure 5.5: The putative global energy minimum for 100 Morse particles with $\rho = 6$ on cones with six cone angles, interacting through Euclidean space. We demonstrate the three main classes of structure. Left: The net of the crystal. Black lines show the edge of the net of the cone. Middle: Front view of the conical crystal. Right: Top view of the crystal. A-B: $\theta_{\max} = 37.5^\circ$ and $\theta_{\max} = 30.0^\circ$. These structures contain defects arranged in a line. The stress is concentrated around the defects and decays away from the defect line. C-D: $\theta_{\max} = 165^\circ$ and $\theta_{\max} = 180^\circ$. These structures have no defects. At non-magic angles (C), the frustration of the crystal lattice results in stress distributed uniformly throughout the crystal. At magic angles (D) there is no frustration and therefore only a little stress, which is concentrated close to the peak. E-F: $\theta_{\max} = 315^\circ$ and $\theta_{\max} = 330^\circ$. At higher cone angles, the structure is no longer centred on the peak of the cone. Either the cluster still covers the peak, but asymmetrically (E), or is completely detached from the peak (F).




Figure 5.6: The potential energy of the putative global energy minimum for 100 Morse particles with $\rho = 6$ and $\rho = 18$, interacting through Euclidean space on the surface of a cone as a function of the cone angle. Magic angles are indicated by vertical lines. The colour of the points refers to the class of structure. Black points denote structures containing defects (Fig. 5.5A-B), green points denote structures with no defects (Fig. 5.5C-E) and red points denote structures which do not occupy the peak of the cone (Fig. 5.5F). The labels A-F on the $\rho = 6$ line indicate the structures depicted in Fig. 5.5, and the labels A-D on the $\rho = 18$ line indicated the structures depicted in Fig. 5.7.

the stress cost of occupying the peak, so the crystal is driven away from the peak.

The lower (blue) line in Fig. 5.6 shows the energy of the lowest minimum over a range of cone angles from $\theta_{\max} = 15^\circ$ to $\theta_{\max} = 345^\circ$. Each minimum is also labelled by its class. We can see that the energy increases overall with θ_{\max} , but non-monotonically, with a periodic dip in the vicinity of the magic angles. The steady increase can be explained by the growing line energy as the cone angle increases, and the periodic variations by the inclusion of stress and defects away from the magic angles.

To understand how the range of the potential affects the global minima, we performed basin hopping simulations of 100 Morse particles with $\rho = 18$, which is a much shorter ranged interaction than at $\rho = 6$. Four of these structures are shown in Fig. 5.7, and their energies are shown by the upper (red) line in Fig. 5.6. We observe the same three classes of structure as for $\rho = 6$, but with some striking differences.

At the magic angles, we continue to see defect-free crystals, as in Fig. 5.7D. The stress energy here is higher than for $\rho = 6$ because for shorter-ranged potentials the same strain has a higher energy penalty. The four particles at the peak appear to be highly stressed because the pairs that lie opposite each other are close enough to be classed as neighbours, but have almost no attractive interaction due to the very short-ranged potential. This structure is also not

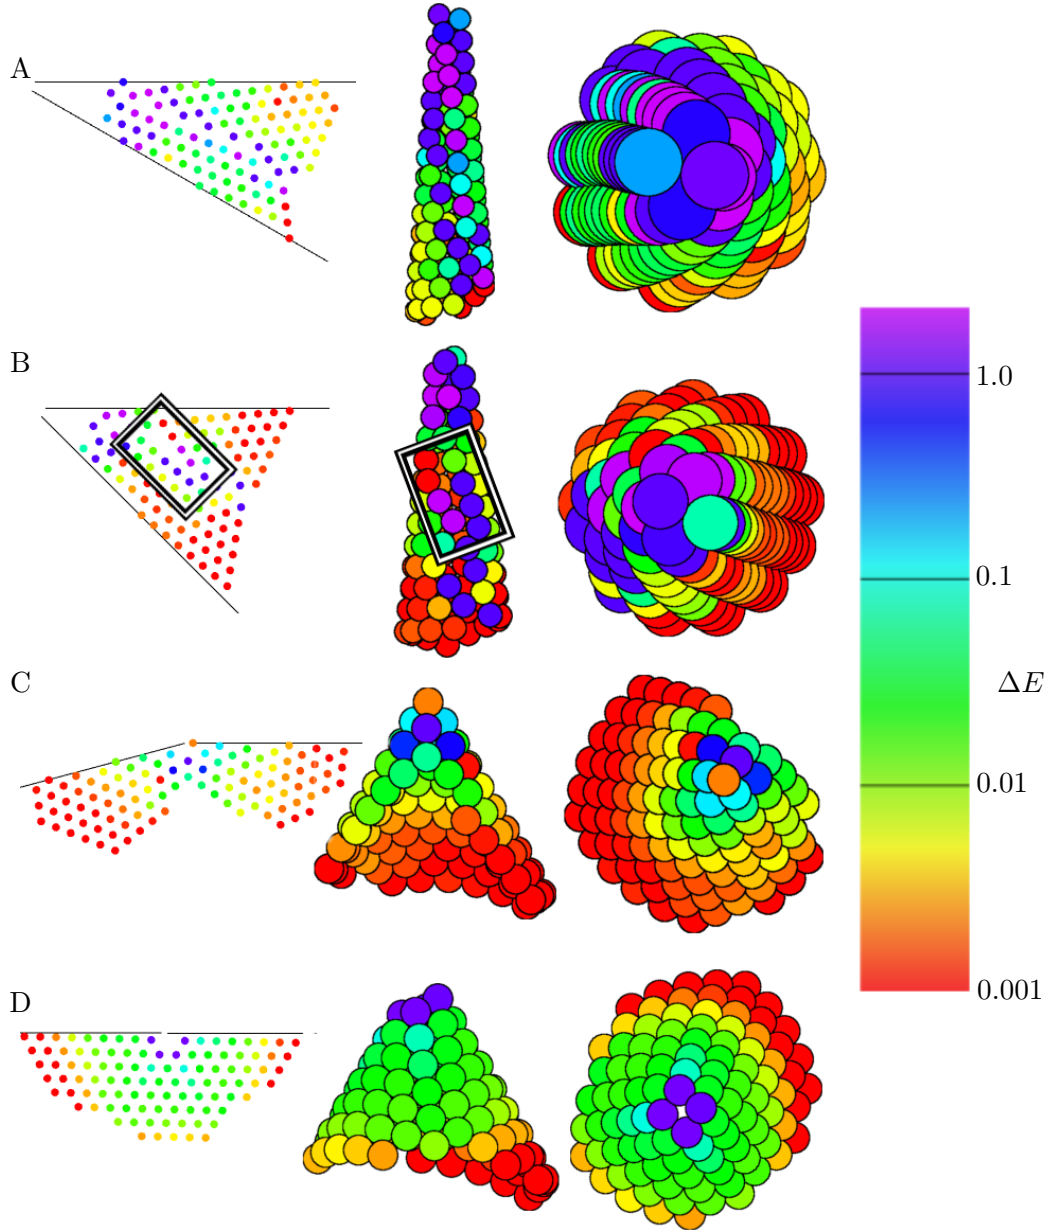


Figure 5.7: The putative global energy minimum for 100 Morse particles with $\rho = 18$ on cones with four cone angles, interacting through Euclidean space. Left: The net of the crystal. Black lines show the edge of the net of the cone. Middle: Front view of the conical crystal. Right: Top view of the crystal. The particles are coloured by the local stress. A: $\theta_{\max} = 30^\circ$ This structure contains defects arranged in a line. It is chiral. B: $\theta_{\max} = 45^\circ$ The underlying crystal structure is achiral, but there is a ragged defect line made of two short wedges. The longest is highlighted by a white frame. The stress is low compared to (A). C: $\theta_{\max} = 165^\circ$, the structure is not centred on the peak of the cone. D: $\theta_{\max} = 180^\circ$, this is a magic angle cone, and the structure is achiral and defect free.

centred on the peak of the cone. For short-ranged potentials, the energy barrier to rearranging particles on the edge of the cluster is very high, as bonds must be broken. However, the arrangement of particles itself has almost no effect on the total energy. Therefore, there are a series of similar energy near-ground states reflecting different arrangements of the outer particles. It is difficult to explore all these minima with basin hopping in a reasonable amount of simulation time, so the crystal in Fig. 5.7D may differ from the true ground state in the details of the perimeter structure, but likely has the same structure away from the boundary.

We also see a new type of defective crystal. Here, the underlying crystal structure appears to have plane of mirror symmetry, but the defect line is jagged. The defects consist of a pair of linked wedge-shaped vacancies, similar to those seen in recent experiments on colloids interacting through short-ranged attraction on cones [271]. They are oriented diagonally, rather than along a straight line connecting the peak of the cone to the base. This asymmetry contributes to an increase in the chirality. Such a structure is shown in Fig. 5.7B and should be compared to Fig. 5.5A. The structure in Fig. 5.7B is much less stressed, suggesting an energetic trade-off between the stress in the crystal and the structure of the defect line.

For the short-ranged potential we see the crystal moving away from the peak at much lower values of θ_{\max} compared to the long-ranged potential. We see that the crystal is always arranged to minimise the length of the defect line, as in Fig. 5.7C. The explanation for this is that the energetic cost of a defect line is much higher for a shorter ranged potential, so the competition between the line energy cost of moving off the peak and the gain associated with reducing the length of the defect line favours the latter more for shorter ranged potentials.

This also explains why the energy of the crystal does not increase as steeply in θ_{\max} . For the short-ranged potential, the increase in line energy at higher cone angles is more strongly compensated for by the reduction in the length of the defect line.

5.5 Measuring chirality

We have seen that conical crystals can adopt structures with both chiral and achiral characteristics, although some structures with underlying symmetry may be formally chiral because, for example, the relaxation of the defect line in Fig. 5.5A might destroy the exact symmetry, without fundamentally changing the structure of the rest of the crystal. Also, if the number of particles is not matched to the cone angle, then excess particles might arrange themselves

asymmetrically around the base of the crystal, which again would make the formal point group trivially C_1 even though the underlying structure should be understood as achiral. To tackle this, we will need to define a continuous degree of chirality whose value is zero for achiral structures and positive for chiral structures but can also take on intermediate values to reflect the extent of any deviations from perfect symmetry [257]. Note that it does not sort structures into left- and right-handed categories, and we see both in our calculations.

The most common chirality measure used in chemistry, the Hausdorff distance, measures the maximum distance between any particle and the closest particle in the mirror image. Therefore, it is not suitable for our purposes, as it is very sensitive to the existence of even one out-of-place particle [257]. In contrast, we need a measure which detects whether or not most particles have a mirror pair, but is robust, for example, to a few edge particles with no opposite partner.

Firstly, we take a net of the crystal and count all the particles without a mirror partner across the line bisecting the net. To do this we overlay the net with its own reflection and count the number of particles which are more than a distance l away from all the particles in the reflection. Here we use $l = 0.07r_0$, which we choose as the results are insensitive to small changes in this region.

This count gives a measure of the chirality for a single net of a given crystal, which is the same as testing one putative mirror plane of the conical structure. To test whether any mirror plane exists, we sweep the cutting line around the cone axis in increments of 0.036° and repeat the calculation at each step. As shown in Fig. 5.3E, rotating the cutting line is the equivalent of rotating the mirror plane about the same axis. This procedure allows us to detect any planes of mirror symmetry. The final measure of chirality is given by the lowest value of the count, as this represents the most symmetric net (i.e. the net where the highest number of particles have a mirror counterpart). If the crystal has no symmetries, the value of the chirality will be high, but if a mirror plane exists it will be low.

In Fig. 5.8, we plot the mean and the range of chiralities over the five lowest minima for 100 $\rho = 6$ Morse particles. At magic angles the structures are all nearly achiral, although no structure has perfect symmetry. Away from the magic angles, most cones have both high and low chirality structures among their lowest minima. To quantify this, we plot a histogram of the chiralities of the 200 lowest energy minima for $\theta_{\max} = 75^\circ$, in Fig. 5.9A. It can be seen that there are minima of all chiralities, but high chirality structures are favoured. The energy




Figure 5.8: The average chirality (purple line) and the range of chiralities (purple band) of the lowest 5 potential energy minima for 100 Morse particles with $\rho = 6$, interacting through Euclidean space on the surface of a cone as a function of the cone net angle. Magic angles are indicated by vertical lines. Close to magic angles the chirality is always low, whereas away from the magic angles, both high and low chirality structures are observed.

difference between the ground state, which has a chirality of 75%, and the first minimum with a chirality below 10% is only $E = 0.1\varepsilon$. This implies that reason that high chirality states are more common is simply that there are more possible chiral states than achiral ones, rather than because of a particular energetic advantage of chirality.

In contrast, at $\rho = 18$, below $\theta_{\max} = 150^\circ$, the range of chiralities in the lowest minima is small, see Fig. 5.10. A histogram of the chiralities of the 200 lowest energy minima for $\theta_{\max} = 75^\circ$ and $\rho = 18$ is shown in Fig. 5.9B. It can be seen that highly chiral structures dominate. What is more, the energy of the lowest minimum is $E = -267.5\varepsilon$, while the energy of the lowest minimum with a chirality less than 60% is $E = -262.0\varepsilon$. This suggests that highly chiral structures have a significant energy advantage over less chiral structures at this cone angle. However, the energy difference between chiral and achiral structures appears to depend on θ_{\max} , as the chirality has a very pronounced periodicity, with the least chiral crystals at $\theta_{\max} = 60^\circ$ and $\theta_{\max} = 120^\circ$ and the most chiral at $\theta_{\max} = 30^\circ$ and $\theta_{\max} = 90^\circ$.

This raises the question, why does the energy seem to depend strongly on chirality at $\rho = 18$, but less so at $\rho = 6$? We argue that, at $\rho = 18$, as the crystal cannot accommodate much strain and only near-direct contacts in the defect line contribute to the energy, the energy of the defect line is purely a function of the relative angles and positions of the colliding crystal faces, which in turn determines the chirality. In contrast, at $\rho = 6$, the crystal can relax the defect line without changing the chirality by including some strain elsewhere in the

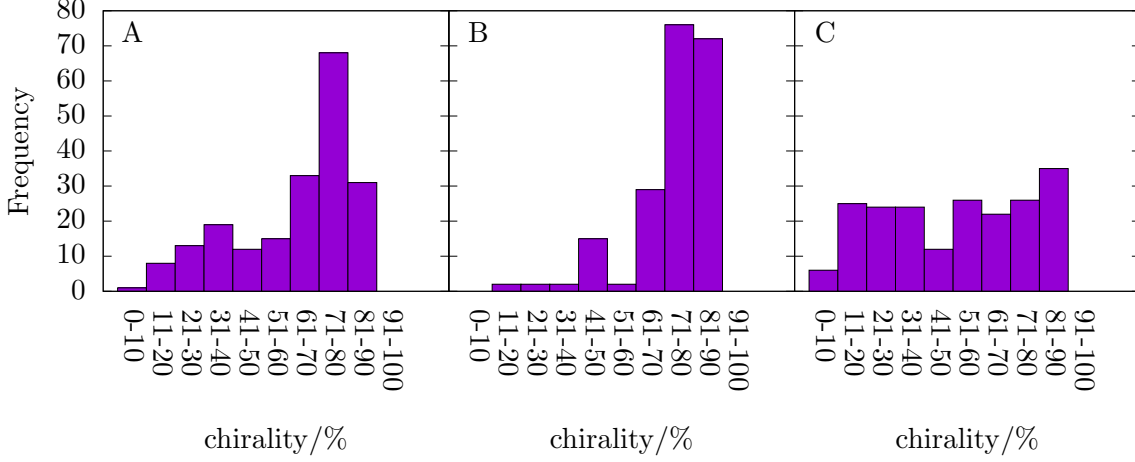


Figure 5.9: Histograms of the chirality of the 200 lowest energy minima for 100 particles interacting through Euclidean space. A: $\rho = 6$ and $\theta_{\max} = 75^\circ$. B: $\rho = 18$ and $\theta_{\max} = 75^\circ$. C: $\rho = 18$ and $\theta_{\max} = 195^\circ$.

crystal, or recovering energy from bonds stretched across the defect line. Both these effects serve to reduce the energy difference between faces colliding at different angles, and so reduce the energy dependence on chirality.

For $\rho = 18$, the picture changes completely above $\theta_{\max} = 150^\circ$. Away from the magic angles, the range of chiralities is large. We see an apparent continuum of structures ranging from those fully centred on the peak, to those like Fig. 5.7C, which are partially offset, to structures completely detached from the peak. In the latter two cases, as there is little or no collision between the growing faces, these structures can be chiral or achiral without a significant energy change. As a result, the histogram of chiralities for the lowest 200 minima at $\theta_{\max} = 195^\circ$ is nearly flat. Furthermore, the energies of these two types of structure are very similar. The ground state is offset from the peak, but the first minimum centred on the peak has an energy only $E = 0.5\epsilon$ higher. This suggests that the potential energy landscape is rough and glassy.

5.6 Characterising crystals on cones

In this section, we present a method for categorising each of the minima calculated from basin hopping with two vector parameters, by fitting them to geometrically-derived lattices. We show that this technique works well in most cases and use it to derive a set of “cone conformation rules”.




Figure 5.10: The average chirality (purple line) and the range of chiralities (purple band) of the lowest 5 potential energy minima for 100 Morse particles with $\rho = 18$, interacting through Euclidean space on the surface of a cone as a function of the cone net angle. Magic angles are indicated by vertical lines. Close to magic angles the chirality is always low. For $\theta_{\max} < 180^\circ$ the range is very small, but at higher angles both low and high chirality structures appear for each angle.

5.6.1 Hexagonal lattices on conical nets

As with the cylinder (Fig. 5.1), we can construct a geometrically-derived conical crystal by taking an infinite, planar hexagonal lattice, cutting out the net of the cone in question and wrapping it up. The character of the resulting structure will only depend on the position and orientation we choose for the net relative to the underlying lattice. Three examples are shown in Fig. 5.11. If we consider a fixed lattice, the placement of the net on it can be quantified using two vectors: \mathbf{p} , the vector connecting the point at the peak of the (untruncated) net to the closest particle in the lattice, and $\hat{\mathbf{a}}$ the unit vector which bisects the net. These are illustrated in 5.12B. When reporting these vectors we take the centre of the particle closest to the peak as the origin and pick the lattice vector of the crystal which makes the smallest angle with \mathbf{p} as the x-axis.

By fitting each basin hopping minimum to a crystal constructed geometrically, as described above, we can define \mathbf{p} and $\hat{\mathbf{a}}$ for that minimum. We use the following scheme. Firstly, we take the net of the basin hopping structure with the largest single crystal domain. This is usually best achieved by making the cut along the defect line. When there is no defect line (for example at a magic angle) we cut along a mirror plane. Next, we fit a perfect hexagonal lattice to the resulting planar structure, using a Monte Carlo algorithm.

The Monte Carlo algorithm is analogous to the approach for thermodynamic sampling




Figure 5.11: The construction of three conical crystals by cutting a net from a planar crystal. These structures are derived from pure geometry and do not make any allowance for interactions between the particles. A: A cone with $\theta_{\max} = 120^\circ$, which is a magic angle. In the left frame, the shaded area is the net of the cone. The dotted line bisects the net and follows a plane of symmetry in the underlying crystal. In the centre, the resulting cone is seen from the side, showing no defects. On the right, the cone is viewed from above, and the dotted lines indicate planes of symmetry, showing that the crystal is not chiral. B: A cone with $\theta_{\max} = 90^\circ$, cut from the same crystal as above. The dot-dash line which bisects the net does not follow a plane of symmetry. The resulting conical crystal has a line of defects at the join between the edges of the net. Viewed from the top, it can be seen that the cone has no planes of reflection and is therefore chiral. C: Another cone with $\theta_{\max} = 90^\circ$, but cut from a crystal rotated by 45° relative from those above. The line which bisects the cone now follows a symmetry plane so that, although the resulting conical crystal has defects, it also has a symmetry plane which slices through the defect line and so it is not chiral.

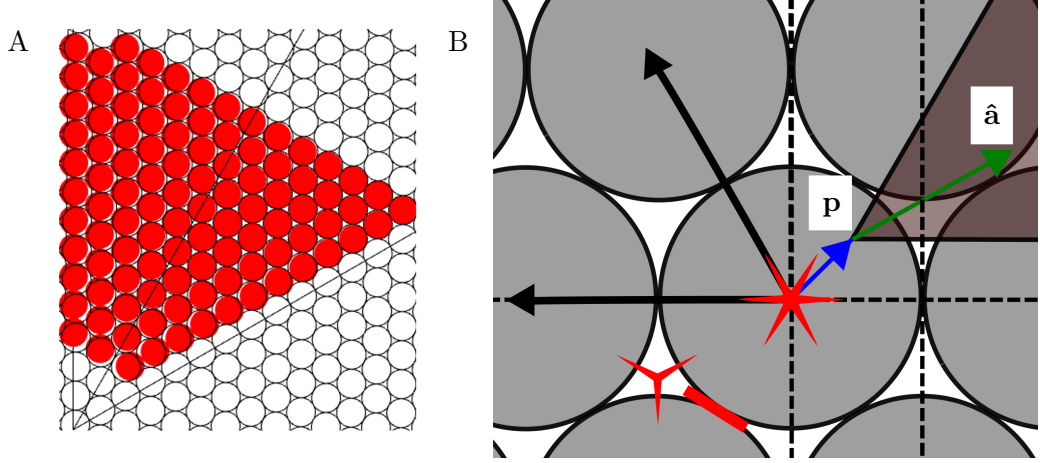


Figure 5.12: A) A perfect hexagonal lattice (black circles) fitted onto the net of a crystal obtained from basin hopping (red discs). B) A schematic outlining the characteristics of a conical crystal. The grey particles represent the perfect hexagonal lattice. The black arrows are the lattice vectors. The red shapes mark points of rotational symmetry on the lattice. There are three classes: a point of six-fold symmetry at the centre of each particle, a point of three-fold symmetry in the centre of each gap and a point of two-fold symmetry where the particles touch. The dashed lines represent some of the lines of symmetry in the lattice. The red figure is the peak of the net of the conical surface. The blue vector \mathbf{p} is the position vector of the peak of the lattice with respect to the closest particle. The unit vector $\hat{\mathbf{a}}$ bisects the net of the cone and represents its orientation with respect to the lattice.

that was described in Sec. 2.2.4. We place the perfect hexagonal lattice over the global minimum structure, and move it around by making uniform trial moves in \mathbf{p} and $\hat{\mathbf{a}}$. We then accept or reject each move with the probability

$$\text{acc}(\mathbf{o} \rightarrow \mathbf{n}) = \min \{1, \exp [-(D_{\mathbf{n}} - D_{\mathbf{o}})/2]\}.$$

where D is the sum of the squared distances from each particle in the calculated minimum to the closest particle on the perfect lattice. At each step we also measure the number of particles in the minimum within $l = 0.07r_0$ of a particle on the ideal lattice, M . After 10^6 Monte Carlo steps, we take the values of \mathbf{p} and $\hat{\mathbf{a}}$ that yielded the highest M (i.e the most overlaps between the global minimum and the perfect lattice) as the best fit. An example of such a fit is shown in Fig. 5.12A. A schematic of \mathbf{p} and $\hat{\mathbf{a}}$ is shown in Fig. 5.12B. It is important to note that, as it relies on fitting an ideal lattice, this procedure works poorly for any structure which contains a lot of strain, such as Fig. 5.5C, or has multiple domains.

5.6.2 Cone conformation criteria: Magic angles

Fig. 5.13 depicts the ground states taken from basin hopping simulations of 100 Morse particles with $\rho = 6$ interacting through Euclidean space for each magic angle. A schematic of \mathbf{p} and $\hat{\mathbf{a}}$, measured as described in Sec. 5.6.1, is also shown. $\hat{\mathbf{a}}$ points along a line of symmetry because that is where we have chosen to cut the cone when forming the net. Due to the underlying symmetry of the lattice, we could have chosen $\hat{\mathbf{a}}$ to point in any direction. In Fig. 5.13C-E, it can be seen that \mathbf{p} points to a centre of rotational symmetry.

This leads to the first case of the cone conformation rule: if \mathbf{p} points to a centre of m -fold rotational symmetry, and $\theta_{\max} = 2n\pi/m$ where n goes from 1 to $m - 1$, the resulting structure will be achiral, and defect-free. This rule is valid because the rotational symmetry of the lattice ensures that the lattice vectors on each side of the join will be commensurate when the net is rolled into a cone. The construction of such a crystal is illustrated in Fig. 5.11A.

The crystals in Fig. 5.13A-B contain a relatively high amount of strain. This is caused by the difference between the geodesic separation between particles, which determines their lattice sites, and the Euclidean separation, which is used to calculate the potential. As a result, these systems deviate from perfect symmetry and from the conformation rule described above. In Sec. 5.7, we show that when this frustration is removed, the calculated crystals fit the geometrical predictions more closely. The small amount of deviation seen in Fig. 5.13C-E could be explained by the truncation of the cone, as the particle which should be located exactly on the peak is slightly offset due to the truncation.

In order to minimise the energy by eliminating defects, we can predict that most real colloidal systems confined to magic angle cones will adopt this type of structure in the ground state. However, for a given colloidal system, the details of the potential will determine both which symmetry point is chosen and how far the structure will deviate due to strain. Cones with an angle close to the magic angles will also adopt achiral, defect-free structures by incorporating strain. In addition, recent experimental work has shown that kinetic traps can play an important roll in such colloidal systems, so the ground state may not be easily observed [256, 271].

5.6.3 Cone conformation criteria: Intermediate angles

Fig. 5.14 depicts the ground states and vectors \mathbf{p} and $\hat{\mathbf{a}}$ for three non-magic angle cones, one of which is achiral and two which are chiral. In this case, the choice of $\hat{\mathbf{a}}$ is not free but

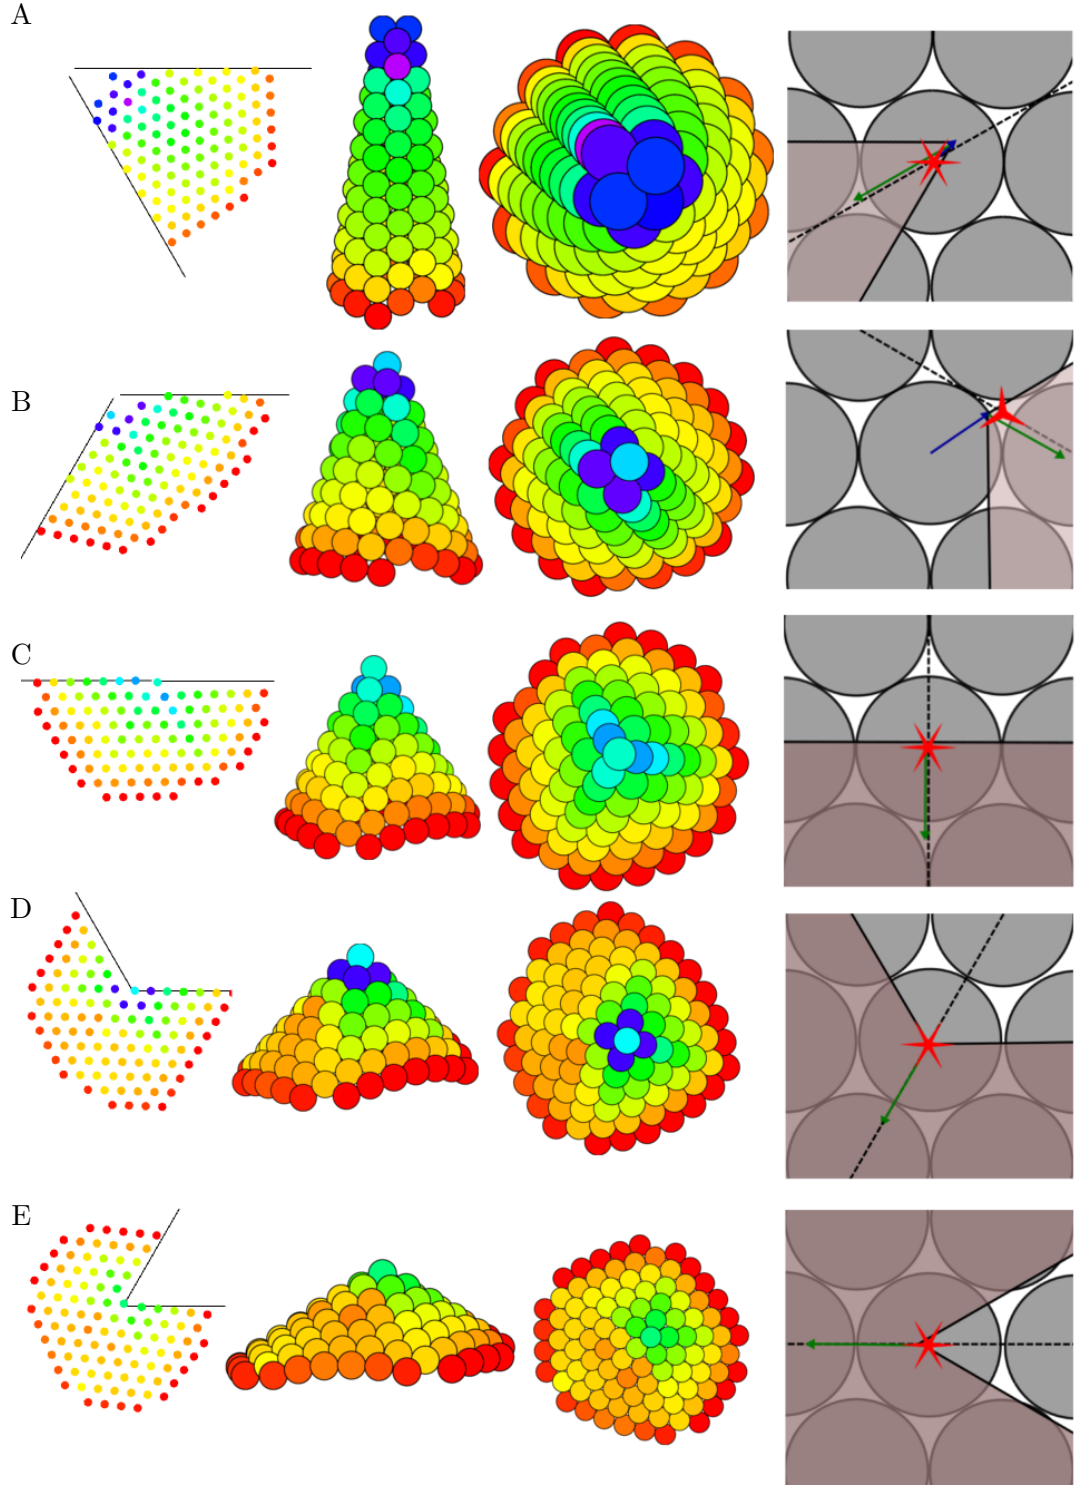


Figure 5.13: Putative global energy minima for 100 $\rho = 6$ Morse particles, interacting through Euclidean space on magic angle cones. For each structure, we show, from left to right: the net of the crystal, a front view of the crystal, a top view of the crystal, and a schematic of \mathbf{p} and $\hat{\mathbf{a}}$. Where applicable, particles are coloured by their local stress, see the colour bar in Fig. 5.5. The arrows, symbols and shaded area in the schematic have the same definition as in Fig. 5.12. A: $\theta_{\max} = 60^\circ$. B: $\theta_{\max} = 120^\circ$. C: $\theta_{\max} = 180^\circ$. D: $\theta_{\max} = 240^\circ$. E: $\theta_{\max} = 300^\circ$.

rather chosen by cutting the conical crystal along the defect line when making the net. In Fig. 5.14A, the achiral case, \mathbf{p} points to a location on a mirror plane of the lattice and $\hat{\mathbf{a}}$ is parallel to that symmetry line. For the chiral structures, Fig. 5.14B-C, this is not the case.

The presence of a defect line in a conical crystal means that any symmetry plane must pass through that line (and the axis of the cone), otherwise, one side of the “mirror” plane would have defects on it and the other would not, destroying any potential symmetry. As in Sec. 5.2, this is equivalent to arguing that the net of the cone must be symmetric around its centre line when the cut is made along the defect line. This gives us the second case of the cone conformation rule: if θ_{\max} is not at or close to a magic angle, then the conical crystal will contain defects. In that case, if \mathbf{p} points to a mirror line in the lattice, and $\hat{\mathbf{a}}$ is parallel to that line, then the conical crystal will be achiral, otherwise, it will be chiral. An example of the construction of an achiral crystal is shown in Fig. 5.11C, and a chiral crystal in Fig. 5.11B.

As seen in Fig. 5.8, at $\rho = 6$, low and high chirality structures for non-magic angles have a very similar energy, suggesting that the presence of chirality is likely to depend strongly on the detail of the potential and the surface. In contrast, for a short-ranged potential and low cone angle, all crystals at non-magic angles are significantly chiral.

5.7 Comparing through-space and geodesic potentials

The results presented so far have all used a potential that acts directly through three-dimensional space using Euclidean distances. For comparison, we have also performed basin hopping simulations of systems of 100 $\rho = 6$ Morse particles on a cone which interact along geodesics, for a range of cone angles. In Fig. 5.15 and Fig. 5.16 we present example structures, repeating the analyses which were developed in Sec. 5.5-5.6 of this chapter. In Fig. 5.17 we show how the energy of the ground state varies with θ_{\max} .

In addition to the defective and defect-free classes, we see a crystal with multiple domains at the lowest angle $\theta_{\max} = 15^\circ$. This is consistent with what is seen for colloidal systems interacting through depletion, although it should be noted that the potential we use is much longer ranged than the experimental interactions [271].

The remaining structures obey the cone conformation rules described in Sec. 5.6. In fact, at $\theta_{\max} = 60^\circ$ and $\theta_{\max} = 120^\circ$, the structures for the geodesic potential fit our predictions

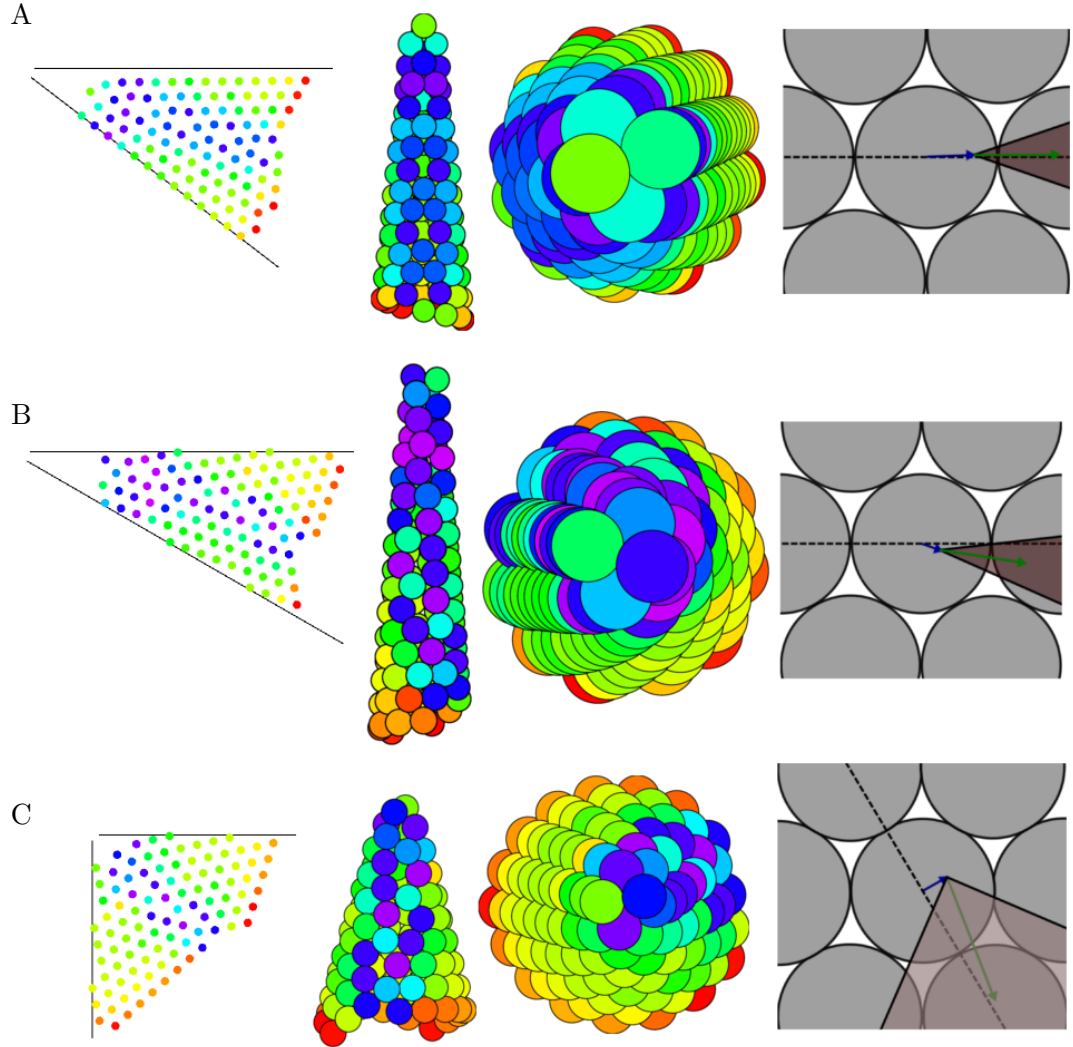


Figure 5.14: Putative global energy minima for 100 $\rho = 6$ Morse particles, interacting through Euclidean space on non-magic angle cones. For each structure, we show, from left to right: the net of the crystal, a front view of the crystal, a top view of the crystal, and a schematic of \mathbf{p} and $\hat{\mathbf{a}}$. Where applicable, particles are coloured by their local stress, see the colour bar in Fig. 5.5. The arrows, symbols and shaded area in the schematic have the same definition as in Fig. 5.12. A: $\theta_{\max} = 37.5^\circ$, the structure is highly achiral. B: $\theta_{\max} = 30^\circ$, the structure has a high degree of chirality. C: $\theta_{\max} = 90^\circ$, the structure has a high degree of chirality.




Figure 5.15: Putative global energy minima for 100 $\rho = 6$ Morse particles on magic angle cones. The potential acts along geodesics. For each structure, we show, from left to right: the net of the crystal, a front view of the crystal, a top view of the crystal, and a schematic of \mathbf{p} and $\hat{\mathbf{a}}$. Where applicable, particles are coloured by their local stress, see the colour bar in Fig. 5.5. A: $\theta_{\max} = 60^\circ$. \mathbf{p} points to a centre of sixfold rotational symmetry. B: $\theta_{\max} = 120^\circ$. \mathbf{p} points to a centre of threefold rotational symmetry. C: $\theta_{\max} = 180^\circ$. \mathbf{p} points to a centre of twofold rotational symmetry.

more closely that for the Euclidean potential, see Fig. 5.15. One of the most striking differences between the two potentials is the absence of stress around the peak of the cone at magic angles, compared to the Euclidean case. In the Euclidean interaction, this stress is caused by particles overlapping near the peak of the cone. For the geodesic potential, particles near the peak can overlap in 3D-space while still maintaining their geodesic separation.

Away from the magic angles, the behaviour of the two potentials is very similar. This can be seen in the chiralities, shown in Fig. 5.18 and the energies, shown in Fig. 5.17. In both the Euclidean and geodesic cases, the structures at magic angles show a low degree of chirality and low energy, while at intermediate angles both chiral and achiral structures are seen and the energy is higher.

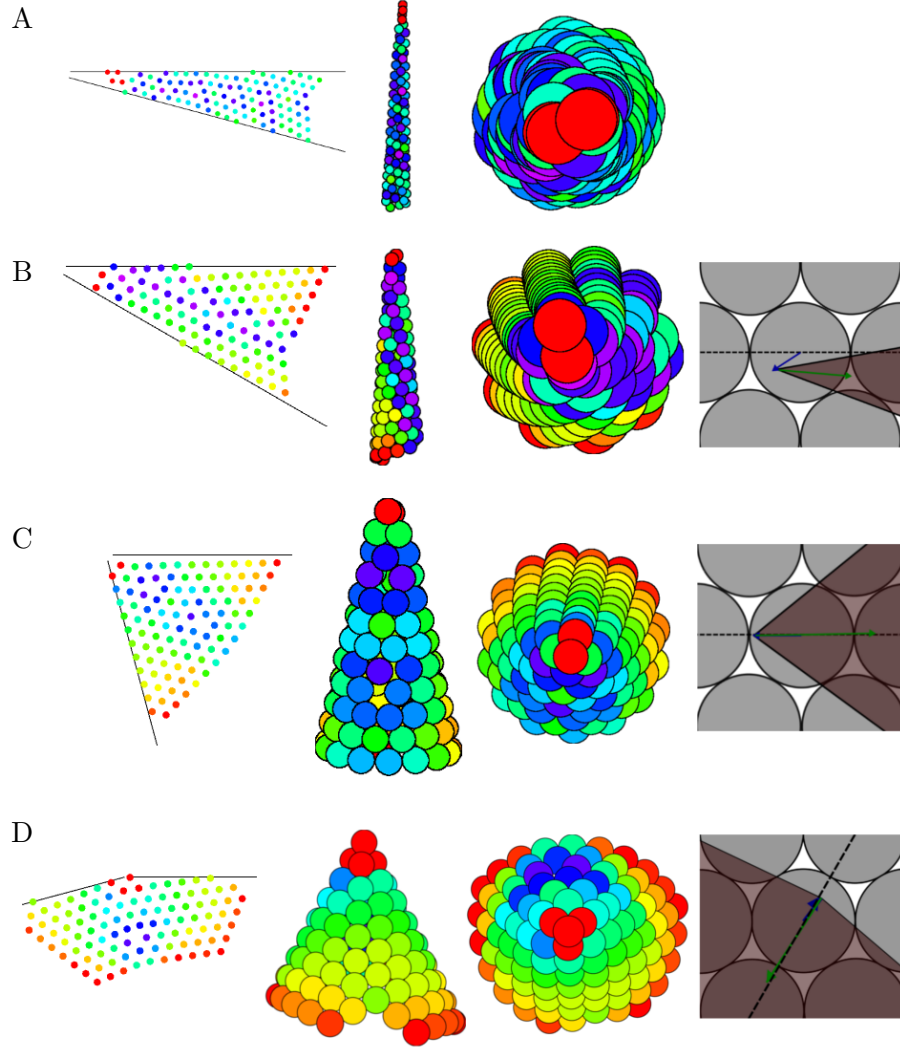


Figure 5.16: Putative global energy minima for 100 $\rho = 6$ Morse particles, interacting along geodesics on four cones. For each structure, we show, from left to right: the net of the crystal, a front view of the crystal, a top view of the crystal, and a schematic of \mathbf{p} and $\hat{\mathbf{a}}$. Where applicable, particles are coloured by their local stress, see the colour bar in Fig. 5.5. A: $\theta_{\max} = 15^\circ$. At this angle the crystal has three domains separated by a line slips. These are in addition to the defects usually seen at non-magic angles. \mathbf{p} and $\hat{\mathbf{a}}$ are not defined for multi-domain crystals, so that schematic is omitted. B: $\theta_{\max} = 30^\circ$. The crystal has defects and has a high chirality. C: $\theta_{\max} = 75^\circ$. The crystal has defects but is highly achiral. The “cone conformation criteria” is obeyed. D: $\theta_{\max} = 180^\circ$. The crystal is close to achirality.




Figure 5.17: The potential energy of the putative global energy minimum for 100 Morse particles with $\rho = 6$, interacting along geodesics on the surface of a cone as a function of the cone net angle. Magic angles are indicated by vertical lines. The colour of the points refers to the class of structure. Black points denote structures containing defects (Fig. 5.16B-D), green points denote structures with no defects (Fig. 5.15A-C) and yellow points denote structures with multiple domains (Fig. 5.16A).




Figure 5.18: The average chirality (purple line) and the range of chiralities (purple band) of the lowest 5 potential energy minima for 100 Morse particles with $\rho = 6$, interacting through along geodesics on the surface of a cone as a function of the cone net angle. Magic angles are indicated by vertical lines. Close to magic angles the chirality is always low, whereas away from the magic angles, both high and low chirality structures are observed.

5.8 Conclusions

In this chapter, we have investigated the ground states of crystals on conical surfaces for three potentials: a short- and long-ranged potential acting through Euclidean space and a long-ranged potential acting along geodesics. In each case, there are magic angles cones with $\theta_{\max} = n\pi/3$ that can host defect-free achiral crystals. The geodesic crystals are significantly less strained than the Euclidean crystals.

At intermediate angles, both the Euclidean and geodesic systems with long-ranged interactions show chiral and achiral structures among their lowest energy minima. In contrast, crystals formed of particles interacting through short-ranged potentials have a single class of structures which dominates the potential energy landscape for each cone angle below $\theta_{\max} = 150^\circ$. The chirality of these is periodic in θ_{\max} . Above $\theta_{\max} = 150^\circ$ at non-magic angles, the clusters are no longer always centred on the peak. This is because the energetic cost of defects and strain is higher for shorter ranged potentials, leading the crystal to trade an increased line tension for less frustration and fewer defects. While the ground state is always offset, a wide variety of structures are found at low energies, suggesting a glassy landscape.

We have shown that any single-domain hexagonal crystal on a cone can be categorised by two vectors \mathbf{p} and $\hat{\mathbf{a}}$, which relate the position of the peak of the cone to the closest particle in a perfect hexagonal lattice fitted to the real structure. A set of cone conformation rules have been developed which relate these vectors to the chirality of the structure. These state that for a cone with $\theta_{\max} = n\pi/3$, the peak of the cone will lie on a rotational symmetry point of the crystal. There will be no defects and the structure will be achiral. If θ_{\max} is near to $n\pi/3$, then the crystal may adopt a similar defect-free structure by incorporating some strain. If θ_{\max} is far from $n\pi/3$, the crystal can be chiral or achiral, but will contain a line of defects. If, when the crystal is cut along the defect line and unrolled into its net, the line which bisects the net of the conical surface lies directly on, and parallel to, a mirror line of the crystal lattice, then the crystal is achiral. Otherwise, it is chiral.

In this chapter we have investigated cones over the full range of θ_{\max} from $\theta_{\max} = 15^\circ$ and with a very short truncation. To bring these simulations closer to the experiments of Sun *et al.*, a continuation of this project would investigate cones with a narrower cone angle and larger truncation [271]. At very low cone angles, we would expect to see ring-like defect structures encircling the cone.

Chapter 6

Conclusions and outlook

6.1 Overview

In Sec. 1.4 we introduced three open questions: *What are the effects of confinement to a curved surface on the free energy of a two-dimensional system of isotropically attractive particles? How do these effects interact with each other? What novel properties arise from these effects and their interactions?* In subsequent chapters, we used a combination of computer simulations and analytical models to investigate a variety of surfaces with the aim of advancing our understanding of these questions. Our primary computational tools were Monte Carlo and molecular dynamics simulations, as well as energy minimisation with the basin hopping algorithm and the open source software Surface Evolver. Each method has been carefully adapted to simulate the curved surfaces of interest. We applied these to four surfaces: the sphere, torus, sinusoid and cone. In this chapter, we will first summarise the four general effects of surface curvature that we have referred to throughout this thesis and the main contributions of this thesis. Then, we will discuss how the work in this thesis could be refined and expanded upon in the future.

6.2 Effects of curvature

6.2.1 Finite extent of the surface

Many curved surfaces are intrinsically finite or periodic in extent. The finite area of a surface introduces so-called finite-size effects, which can influence structure, large-scale fluctuations, and the character of phase transitions. While the usual practice is to limit these effects by extrapolating to the thermodynamic limit, this is not appropriate for a finite curved surface. In fact, these finite-size effects are physically relevant. Furthermore, for a system confined

to a closed surface, the equivalence between the canonical and grand canonical ensembles breaks down. For instance, we quantified this effect for the gas-liquid transition on a sphere in chapter 3. In the grand-canonical ensemble, the nucleation of a new phase is arrested when the surface is saturated by that phase. This can be modelled mathematically by calculating the reduction in entropy associated with the increasing density of a liquid as particles are added once the surface is full. This entropy loss comes from a reduction of the available free space. In contrast, in the canonical ensemble, the nucleation process is instead arrested when the gas phase becomes depleted. This effect can be captured using a modified Rao-Berne equation, which explicitly models the free energy of the gas and liquid phases [226].

The periodicity of the surface gives rise to an interesting interplay between the periodic length of the surface (for example, the circumference of a cylinder) and the periodicity of any lattice formed by the colloidal particles confined to that surface [256]. In general, these two periodic length scales are not commensurate with each other. A clear example for this is provided by our work on the cone. As a crystal grows on the cone, it will eventually meet itself when it has encircled the surface. We showed in chapter 5 that, at certain magic cone angles, the crystal can exploit the rotational symmetry of the lattice to find a configuration in which the colliding crystal faces are commensurate. For other cone angles, however, this is impossible and the crystal is frustrated, causing a line of defects to form. The need to minimise the energy costs of the defects can lead to the formation of chiral ground state crystals.

6.2.2 Distance contraction

When two particles are confined to a curved surface, the Euclidean distance between them is reduced compared to the geodesic distance. Hence, if the particles interact through 3D-space, this effect will alter the potential energy of a curved system compared to its flat counterpart. In addition, for a particle with a given range of attraction, if the Gaussian curvature is more negative the area of the surface within the sphere of attraction of that particle is higher. In condensed phases, this means that each particle has more close neighbours and hence more interactions, lowering the potential energy.

On a torus, these effects contribute to clusters of colloids favouring regions with large negative Gaussian curvature, as we have seen in chapter 4 where the liquid phase is found in the hole at the centre of the torus, rather than on the outside of the torus. For instance,

on a 5-7 torus covered with $\rho = 6$ Morse particles, the energy difference per particle when located at the highest and lowest curved regions is about 5%. Interestingly, this effect is not simply proportional to either the mean or Gaussian curvature, but it is a non-local function of the shape of the surface.

While bringing distant particles closer together usually lowers the energy of attractive particles, the opposite is true if these particles come so close that they overlap. We investigated this effect in chapter 5, by comparing systems of particles that interact through Euclidean space and along geodesics on cones. The Euclidean interactions can cause frustration in the crystal when the crystal structure dictates that two particles overlap on the inside of the cone. This induces strain in the ground state crystal, especially at magic angles.

6.2.3 Line energy

The relationship between the perimeter and area of a patch on a curved surface depends on the underlying curvature and the shape of the patch. As the line energy of a patch of colloidal particles is proportional to its perimeter, the curvature of a surface can have a profound impact on the shape, location and free energy of a patch of colloids [18].

In classical nucleation theory (CNT), the nucleation and growth of a new phase depends on the relationship between the perimeter and area of the growing nucleus. On a plane, the minimal patch is disc-shaped, and the perimeter grows monotonically with the area. On a sphere, it is a spherical cap, and when the patch covers more than half of the sphere, the perimeter actually decreases with increasing area. In chapter 3 we simulated gas-liquid nucleation for Lennard-Jones particles on a sphere and showed that the free energy profile has a different form to that on the plane. We also found that this difference can be accounted for by inserting the correct perimeter-area relationship into CNT. In addition, we showed that the size of the critical nucleus is curvature dependent. However, the size of the critical nucleus is not well predicted by CNT at low supersaturations, as the line tension is not constant but depends on the radius of curvature.

On surfaces with varying curvature, the line energy of a patch affects where it will be found. A patch of a given area will – all other things being equal – adopt the shape and location which minimises its perimeter. For non-trivial curved surfaces, these cannot be calculated analytically, and we proceeded by employing Surface Evolver to minimise the line energy for a given patch area. In chapter 4 we found that, on a torus, the optimal position

can be on the “outside” for a sufficiently thick torus, or in a ring around the tube for a thin torus. On mountain-like surfaces, such as a sinusoid (chapter 4) or cone (see chapter 5), the optimal position is around the peak.

6.2.4 Crystal frustration

Crystals confined to topologically spherical surfaces must contain defects because they are required to carry a topological charge of $+12$. Furthermore, a crystal on any surface with non-zero Gaussian curvature is frustrated as parallel lines do not remain equidistant [17].

Previous works have shown that, for long-range potentials, this effect results in the presence of dislocations in the crystal, even in the ground state [17]. For shorter ranged potentials, where defects have a high energy cost, ribbon-like crystals form because this minimises the frustration of the structure, although at the cost of increased line energy [19].

On surfaces with varying curvature, this effect tends to drive crystalline clusters to regions of low Gaussian curvature. This effect becomes more prominent as the range of the potential decreases, which can drive localisation transitions. In chapter 4, we showed that crystals of Morse particles with $\rho < 8$ form patches with defects on the “outside” of a 5-7 torus, while at $\rho > 8$ the particles form ribbons around the flat top.

6.2.5 Interplay between the effects

The four aforementioned effects primarily contribute to the phenomenology of systems confined to curved surfaces by altering the free energy. In some cases, these effects act in isolation, but in many cases, it is the competition between two or more effects that drives the emergence of the novel phenomena observed. As a key example, we have shown in this thesis that the non-linear contributions of these effects can stabilise an array of new, localised states. Interestingly, the transitions between these states require the translation of the whole colloidal assemblies across the surface. This collective effect acts as a form of cooperative curvature sensing that emerges from isotropic particles, in contrast to other curvature sensing mechanisms that typically rely on anisotropic interactions [2, 124].

6.3 Future work

6.3.1 Potential applications

Alongside the recent advances towards manufacturing arbitrary curved surfaces at the colloidal scale [16, 166], we hope that our findings may provide insight to researchers working towards technological applications using colloids and curvature. Understanding the relationship between defect structures and geometry can aid in the creation of colloidal building-blocks with useful properties, such as tetravalent bonding [223, 273]. In addition, defect structures including pleats and scars have recently been identified in the arrangement of cells in spherical epithelial monolayers [110]. Many cellular structures are not spherical, but display complex curvature at the scale of the cell, for example at the junctions between small blood vessels [274]. Our work on complex curvature could help explain the role of defect structures on these surfaces.

Synthetic biology seeks to create minimal systems from artificial components which reproduce biological processes [275]. One key problem is the creation of systems that mimic cell division. One promising approach uses phase-separating lipids to drive the curvature required for a giant unilamellar vesicle to divide [276]. Our work on the thermodynamics of systems on curved surfaces could aid in the refinement of this process. Another approach to induce or sense membrane curvature is the use of specialised proteins [277], but our work suggests that simple colloidal membrane inclusions could have a similar effect. Synthetic biology also has potential health applications. Using carefully designed peptides, it is possible to produce artificial virus-like capsids with a specified structure. These can carry antimicrobial payloads with reduced risk of antibiotic resistance [278]. We hope that with a better understanding of the relationship between defect structures and geometry, a library of artificial capsids with a range of properties could become available.

6.3.2 Exploring potential energy landscapes

In chapters 4 and 5, we have studied configurations corresponding to low lying minima in the potential energy for colloids embedded on the surface of tori and cones using the Basin Hopping algorithm. We found cases where multiple states compete, such as crystals located on top or on the side of the tori in chapter 4. Furthermore, we observed that the roughness of the potential energy landscape may depend on two characteristics of the interaction potential,

namely the range of the potential, and whether it acts directly through three-dimensional space or along surface geodesics, as investigated in chapter 5.

A fruitful line of enquiry to characterise the potential energy landscape is to construct the so-called disconnectivity graph [279]. A disconnectivity graph maps the potential energy surface onto the set of local minima and the transition states (here, defined as first-order saddles on the potential energy surface) that connect them. Using this approach, we can investigate how easily different structures can interconvert, and whether distinct classes (or funnels) of minima exist. The structure of the disconnectivity graph also informs us whether the potential energy landscape is glassy (where numerous minima with similar energy exist and they are separated by large energy barriers) or it corresponds to a structure-seeker (in which there is a single steep funnel with low energy barriers) [280]. This approach has previously been employed for repulsive colloids on catenoids, unduloids and nodoids, showing that dislocation-dominated states and disclination-dominated states form two distinct funnels in the potential energy landscape, connected by a high energy transition state [169].

More specifically, building on the work in this thesis, potential energy landscape exploration can inform on key differences between magic and non-magic angle cones, including how the crystal structures on them depend on the interaction potentials used. It would also be interesting to find out whether chiral and achiral minima on the cones, as well as the different crystalline states on the tori, represent distinct minimum funnels and what the transition pathways between these states are.

6.3.3 Simulating flexible membranes

Throughout this thesis we have focused on particles confined to fixed-shape manifolds. In contrast, in many biological examples, the molecules that detect surface curvature are often also the ones that cause it. An important example is provided by the BAR and N-BAR domains [2, 281–283]. These molecules bind well to lipid membranes with a specific target curvature, but poorly otherwise. Interestingly, at low concentrations, these proteins simply detect curvature, while at higher concentrations, they can act to induce the membrane curvature [2]. BAR domain proteins play a key role in creating and stabilising the lipid bud during clathrin-mediated endocytosis [153, 284].

Thus, it would be valuable to extend the work in this thesis to study cases with deformable surfaces, in order to understand the “feedback-loop” relationship between the adsorbed col-

loids and the curved surfaces. This is, however, a challenging task computationally. One approach is to represent the surface using spherical harmonics [285], and to carry out Monte Carlo steps or minimisations on the weights of each spherical harmonic term. Another possible approach is to triangulate the surface [286,287] and then vary the positions of the vertices. In either case, the key challenges would be to combine the surface deformation and colloidal particle moves, and to manage the corresponding computational demand. I anticipate large computational cost would be incurred calculating the bending energy of the surface and the surface metric at each particle location, which would be required for many of the trial moves.

6.3.4 Non-spherical and polydisperse particles

It is noted in Sec. 1.3.4 that many curvature sensitive particles are anisotropic, or that they interact anisotropically with other particles. For example, there is growing interest in the adsorption of cylinders and other complex shape particles at fluid interfaces [182,288]. Furthermore, small colloidal spheres on oil-water interfaces can interact through quadrupolar capillary forces [124] and form square lattices.

It would be interesting to understand how the effects of curvature that we have discussed in this thesis apply to non-hexagonal crystal structures. What defect motifs are present when square crystals are embedded on a sphere? Since a square lattice has the same number of vertices (V) as faces (F), $V = F$ and the number of edges is twice the number of faces, $E = 2F$ for a square crystal, it has an Euler characteristic of $\chi = V - E + F = F - 2F + F = 0$. Thus, similar to the hexagonal crystal, defects must be present to satisfy the topological constraint of $\chi = 2$ for a sphere. However, it is an open problem whether this will be in the form of 12 particles with five neighbours and complex scars and pleats, or if other defect motifs will arise.

Additionally, recent experimental work has shown that scars form in the arrangement of cells on the surface of spherical metazoan epithelia [110]. Interestingly, they are seen at much lower particle number than predicted by the Thomson problem [130], and this has been attributed to the polydispersity in the cell size. Using the techniques established in this thesis, we could test this prediction and quantify the effect of polydispersity. This is also relevant for colloidal experiments, where polydispersity is often unavoidable and may affect the structures formed [16,19,271].

Bibliography

- [1] S. Paquay, *Physical principles in the self-assembly of colloids and virus capsids*, Ph.D. thesis, Technische Universiteit Eindhoven, 2016.
- [2] B. Antonny, “Mechanisms of membrane curvature sensing,” *Annu. Rev. Biochem.*, vol. 80, 101, 2011.
- [3] T. Baumgart, S. T. Hess, and W. W. Webb, “Imaging coexisting fluid domains in biomembrane models coupling curvature and line tension,” *Nature*, vol. 425, 821, 2001.
- [4] E.-A. Kim and A. H. C. Neto, “Graphene as an electronic membrane,” *Europhys. Lett.*, vol. 84, 57007, 2008.
- [5] F. Guinea, M. I. Katsnelson, and M. A. H. Vozmediano, “Midgap states and charge inhomogeneities in corrugated graphene,” *Phys. Rev. B*, vol. 77, 075422, 2008.
- [6] J. C. Meyer, A. K. Geim, M. I. Katsnelson, K. S. Novoselov, T. J. Booth, and S. Roth, “The structure of suspended graphene sheets,” *Nature*, vol. 446, 60, 2007.
- [7] V. N. Manoharan, “Colloidal matter: Packing, geometry, and entropy,” *Science*, vol. 349, 942, 2015.
- [8] B. Alder and T. Wainwright, “Phase transition for a hard sphere system,” *J. Chem. Phys.*, vol. 27, 1208, 1957.
- [9] F. C. Keber, E. Loiseau, T. Sanchez, S. J. DeCamp, L. Giomi, M. J. Bowick, M. C. Marchetti, Z. Dogic, and A. R. Bausch, “Topology and dynamics of active nematic vesicles,” *Science*, vol. 345, 1135, 2014.
- [10] T. Lopez-Leon, V. Koning, K. B. S. Devaiah, V. Vitelli, and A. Fernandez-Nieves, “Frustrated nematic order in spherical geometries,” *Nat. Phys.*, vol. 7, 390, 2011.
- [11] D. Joshi, D. Bargteil, A. Caciagli, J. Burelbach, Z. Xing, A. S. Nunes, D. E. P. Pinto, N. A. M. Araújo, J. Brujic, and E. Eiser, “Kinetic control of the coverage of oil droplets by DNA-functionalized colloids,” *Sci. Adv.*, vol. 2, e1600881, 2016.
- [12] M. Rinaldin, P. Fonda, L. Giomi, and D. J. Kraft, “Geometric pinning and antimixing in scaffolded lipid vesicles,” *Nat. Commun.*, vol. 11, 4314, 2020.
- [13] P. W. Ellis, D. J. G. Pearce, Y.-W. Chang, G. Goldsztein, L. Giomi, and A. Fernandez-Nieves, “Curvature-induced defect unbinding and dynamics in active nematic toroids,” *Nat. Phys.*, vol. 14, 85, 2017.
- [14] Q. Liu, B. Senyuk, M. Tasinkevych, and I. I. Smalyukh, “Nematic liquid crystal boojums with handles on colloidal handlebodies,” *Proc. Natl. Acad. Sci.*, vol. 110, 9231, 2013.

-
- [15] I. B. Liu, G. Bigazzi, N. Sharifi-Mood, L. Yao, and K. J. Stebe, “Curvature capillary repulsion,” *Phys. Rev. Fluids*, vol. 2, 100501, 2017.
- [16] T. Rehman and V. N. Manoharan, personal communication, 2018.
- [17] A. R. Bausch, M. J. Bowick, A. Cacciuto, A. D. Dinsmore, M. F. Hsu, D. R. Nelson, M. G. Nikolaides, A. Travesset, and D. A. Weitz, “Grain boundary scars and spherical crystallography,” *Science*, vol. 299, 1716, 2003.
- [18] L. R. Gómez, N. A. García, V. Vitelli, J. Lorenzana, and D. A. Vega, “Phase nucleation in curved space,” *Nat. Commun.*, vol. 6, 6856, 2015.
- [19] G. Meng, J. Paulose, D. R. Nelson, and V. N. Manoharan, “Elastic instability of a crystal growing on a curved surface,” *Science*, vol. 343, 634, 2014.
- [20] N. N. Andreev, “An extremal property of the icosahedron,” *East J. Approx.*, vol. 2, 459, 1996.
- [21] R. Mezzenga, P. Schurtenberger, A. Burbidge, and M. Michel, “Understanding foods as soft materials,” *Nat. Mater.*, vol. 4, 729, 2005.
- [22] H. T. McMahon and J. L. Gallop, “Membrane curvature and mechanisms of dynamic cell membrane remodelling,” *Nature*, vol. 438, 590, 2005.
- [23] T. Vissers, A. T. Brown, N. Koumakis, A. Dawson, M. Hermes, J. Schwarz-Linek, A. B. Schofield, J. M. French, V. Koutsos, J. Arlt, V. A. Martinez, and W. C. K. Poon, “Bacteria as living patchy colloids: Phenotypic heterogeneity in surface adhesion,” *Sci. Adv.*, vol. 4, eaao1170, 2018.
- [24] J. Arlt, V. A. Martinez, A. Dawson, T. Pilizota, and W. C. K. Poon, “Painting with light-powered bacteria,” *Nat. Commun.*, vol. 9, 768, 2018.
- [25] H. Garoff, R. Hewson, and D.-J. E. Opstelten, “Virus maturation by budding,” *Microbiol. Mol. Biol. Rev.*, vol. 62, 1171, 1998.
- [26] M. Deserno, “Elastic deformation of a fluid membrane upon colloid binding,” *Phys. Rev. E*, vol. 69, 031903, 2004.
- [27] M. Santander-Ortega, M. Plaza-Oliver, V. Rodríguez-Robledo, L. Castro-Vázquez, N. Villaseca-González, J. González-Fuentes, P. Marcos, M. Arroyo-Jiménez, and M. Lozano, “Colloids for drug delivery to the brain,” *J. Drug Deliv. Sci. Technol.*, vol. 42, 193, 2017.
- [28] Y. Wang, I. C. Jenkins, J. T. McGinley, T. Sinno, and J. C. Crocker, “Colloidal crystals with diamond symmetry at optical lengthscales,” *Nat. Commun.*, vol. 8, 14173, 2017.
- [29] F. Li, D. P. Josephson, and A. Stein, “Colloidal assembly: The road from particles to colloidal molecules and crystals,” *Angew. Chem.*, vol. 50, 360, 2011.
- [30] A. C. Arsenault, T. J. Clark, G. von Freymann, L. Cademartiri, R. Sapienza, J. Bertolotti, E. Vekris, S. Wong, V. Kitaev, I. Manners, R. Z. Wang, S. John, D. Wiersma, and G. A. Ozin, “From colour fingerprinting to the control of photoluminescence in elastic photonic crystals,” *Nat. Mater.*, vol. 5, 179, 2006.
- [31] H. Kim, J. Ge, J. Kim, S.-e. Choi, H. Lee, H. Lee, W. Park, Y. Yin, and S. Kwon, “Structural colour printing using a magnetically tunable and lithographically fixable photonic crystal,” *Nat. Photonics*, vol. 3, 534, 2009.

-
- [32] V. Prasad, D. Semwogerere, and E. R. Weeks, “Confocal microscopy of colloids,” *J. Phys. Condens. Matter*, vol. 19, 113102, 2007.
 - [33] A. van Blaaderen and P. Wiltzius, “Real-space structure of colloidal hard-sphere glasses,” *Science*, vol. 270, 1177, 1995.
 - [34] K. A. Rose, M. Molaei, M. J. Boyle, D. Lee, J. C. Crocker, and R. J. Composto, “Particle tracking of nanoparticles in soft matter,” *J. Appl. Phys.*, vol. 127, 191101, 2020.
 - [35] D. Frenkel, “Playing tricks with designer “atoms”,” *Science*, vol. 296, 65, 2002.
 - [36] Y. Wang, Y. Wang, D. R. Breed, V. N. Manoharan, L. Feng, A. D. Hollingsworth, M. Weck, and D. J. Pine, “Colloids with valence and specific directional bonding,” *Nature*, vol. 491, 51, 2012.
 - [37] W. B. Rogers and V. N. Manoharan, “Programming colloidal phase transitions with DNA strand displacement,” *Science*, vol. 347, 639, 2015.
 - [38] E. Zaccarelli, I. Saika-Voivod, S. Buldyrev, A. Moreno, P. Tartaglia, and F. Sciortino, “Gel to glass transition in simulation of a valence-limited colloidal system,” *J. Chem. Phys.*, vol. 124, 2006.
 - [39] M. C. Bott and J. M. Brader, “Phase separation on the sphere: Patchy particles and self-assembly,” *Phys. Rev. E*, vol. 94, 012603, 2016.
 - [40] C. C. Ho, A. Keller, J. A. Odell, and R. H. Ottewill, “Preparation of monodisperse ellipsoidal polystyrene particles,” *Colloid Polym. Sci.*, vol. 271, 469, 1993.
 - [41] S. Sacanna, W. T. M. Irvine, P. M. Chaikin, and D. J. Pine, “Lock and key colloids,” *Nature*, vol. 464, 575, 2010.
 - [42] F. Dekker, R. Tuinier, and A. P. Philipse, “Synthesis of hollow silica nanocubes with tuneable size and shape, suitable for light scattering studies,” *J. Colloid Interface Sci.*, vol. 2, 44, 2018.
 - [43] S. P. Beaumont, “III–V Nanoelectronics,” *Microelectron. Eng.*, vol. 32, 283, 1996.
 - [44] A. H. Chowdhury, N. Salam, R. Debnath, S. M. Islam, and T. Saha, “Chapter 8 - design and fabrication of porous nanostructures and their applications,” in “*Nanomaterials Synthesis*,” *Micro and Nano Technologies*, 265 – 294, Elsevier, 2019.
 - [45] A. Atiqah and M. Ansari, “Chapter 6 - nanostructure–polymer composites for soft-tissue engineering,” in “*Nanostructured Polymer Composites for Biomedical Applications*,” *Micro and Nano Technologies*, 105 – 115, Elsevier, 2019.
 - [46] V. Saikiran, M. H. Dar, R. Kuladeep, L. Jyothi, and D. N. Rao, “Chapter 6 - strategies in laser-induced synthesis of nanomaterials,” in “*Nanomaterials Synthesis*,” *Micro and Nano Technologies*, 149 – 199, Elsevier, 2019.
 - [47] Z. L. Wang, “Nanostructures of zinc oxide,” *Materials Today*, vol. 7, 26, 2004.
 - [48] A. Biswas, I. S. Bayer, A. S. Biris, T. Wang, E. Dervishi, and F. Faupel, “Advances in top–down and bottom–up surface nanofabrication: Techniques, applications & future prospects,” *Adv. Colloid Interface Sci.*, vol. 170, 2, 2012.
 - [49] S. Shaw, *Bottom-up approach to fabricate nanostructured thin films from colloidal nanocrystal precursors*, Ph.D. thesis, Iowa State University, 2016.

-
- [50] N. Mahynski, L. Rovigatti, C. Likos, and A. Panagiotopoulos, “Bottom-up colloidal crystal assembly with a twist,” *ACS Nano*, vol. 10, 5459, 2016.
 - [51] K. L. Thompson, M. Williams, and S. P. Armes, “Colloidosomes: Synthesis, properties and applications,” *J. Colloid Interface Sci.*, vol. 447, 217, 2015.
 - [52] E. Yablonovitch and T. J. Gmitter, “Photonic band structure: The face-centered-cubic case,” *Phys. Rev. Lett.*, vol. 63, 1950, 1989.
 - [53] E. Yablonovitch, “Inhibited spontaneous emission in solid-state physics and electronics,” *Phys. Rev. Lett.*, vol. 58, 2059, 1987.
 - [54] K. M. Ho, C. T. Chan, and C. M. Soukoulis, “Existence of a photonic gap in periodic dielectric structures,” *Phys. Rev. Lett.*, vol. 65, 3152, 1990.
 - [55] R. D. Meade, K. D. Brommer, A. M. Rappe, and J. D. Joannopoulos, “Photonic bound states in periodic dielectric materials,” *Phys. Rev. B*, vol. 44, 13772, 1991.
 - [56] A. B. Rao, J. Shaw, A. Neophytou, D. Morpew, F. Sciortino, R. L. Johnston, and D. Chakrabarti, “Leveraging hierarchical self-assembly pathways for realizing colloidal photonic crystals,” *ACS Nano*, vol. 14, 5348, 2020.
 - [57] É. Ducrot, M. He, G.-R. Yi, and D. J. Pine, “Colloidal alloys with preassembled clusters and spheres,” *Nat. Mater.*, vol. 16, 652, 2017.
 - [58] G. I. Waterhouse and M. R. Waterland, “Opal and inverse opal photonic crystals: Fabrication and characterization,” *Polyhedron*, vol. 26, 356, 2007.
 - [59] E. Yablonovitch, “Photonic crystals: Semiconductors of light,” *Scientific American*, vol. 285, 47, 2002.
 - [60] L. Biró, K. Kertész, Z. Vértessy, G. Márk, Z. Bálint, V. Lousse, and J.-P. Vigneron, “Living photonic crystals: Butterfly scales — nanostructure and optical properties,” *Mater. Sci. Eng. C*, vol. 27, 941, 2007.
 - [61] J. W. Galusha, L. R. Richey, J. S. Gardner, J. N. Cha, and M. H. Bartl, “Discovery of a diamond-based photonic crystal structure in beetle scales,” *Phys. Rev. E*, vol. 77, 050904, 2008.
 - [62] A. D. Dinsmore, M. F. Hsu, M. G. Nikolaidis, M. Marquez, A. R. Bausch, and D. A. Weitz, “Colloidosomes: Selectively permeable capsules composed of colloidal particles,” *Science*, vol. 298, 1006, 2002.
 - [63] S. Laib and A. F. Routh, “Fabrication of colloidosomes at low temperature for the encapsulation of thermally sensitive compounds,” *J. Colloid Interface Sci.*, vol. 317, 121, 2008.
 - [64] L. Croll and H. Stöver, “Formation of tectocapsules by assembly and cross-linking of poly(divinylbenzene-alt-maleic anhydride) spheres at the oil-water interface,” *Langmuir*, vol. 19, 5918, 2003.
 - [65] R. Fantoni, J. W. O. Salari, and B. Klumperman, “Structure of colloidosomes with tunable particle density: Simulation versus experiment,” *Phys. Rev. E*, vol. 85, 061404, 2012.
 - [66] O. J. Cayre, P. F. Noble, and V. N. Paunov, “Fabrication of novel colloidosome microcapsules with gelled aqueous cores,” *J. Mater. Chem.*, vol. 14, 3351, 2004.

-
- [67] Y. Chen, C. Wang, J. Chen, X. Liu, and Z. Tong, “Growth of lightly crosslinked phema brushes and capsule formation using pickering emulsion interface-initiated atp,” *J. Polym. Sci. A*, vol. 47, 1354, 2009.
- [68] Q. Sun, J.-F. Chen, and A. F. Routh, “Coated colloidosomes as novel drug delivery carriers,” *Expert Opinion on Drug Delivery*, vol. 16, 903, 2019.
- [69] R. Langer, “Drug delivery and targeting,” *Nature*, vol. 392, 5—10, 1998.
- [70] W. G. Hoover and F. H. Ree, “Melting transition and communal entropy for hard spheres,” *J. Chem. Phys.*, vol. 49, 3609, 1968.
- [71] P. N. Pusey and W. van Megen, “Phase behaviour of concentrated suspensions of nearly hard colloidal spheres,” *Nature*, vol. 320, 340, 1986.
- [72] J. M. Kosterlitz and D. J. Thouless, “Ordering, metastability and phase transitions in two-dimensional systems,” *J. Phys. A*, vol. 6, 1181, 1973.
- [73] B. I. Halperin and D. R. Nelson, “Theory of two-dimensional melting,” *Phys. Rev. Lett.*, vol. 41, 121, 1978.
- [74] D. R. Nelson and B. I. Halperin, “Dislocation-mediated melting in two dimensions,” *Phys. Rev. B*, vol. 19, 2457, 1979.
- [75] A. P. Young, “Melting and the vector Coulomb gas in two dimensions,” *Phys. Rev. B*, vol. 19, 1855, 1979.
- [76] C. A. Murray and D. H. Van Winkle, “Experimental observation of two-stage melting in a classical two-dimensional screened Coulomb system,” *Phys. Rev. Lett.*, vol. 58, 1200, 1987.
- [77] Y. Tang, A. J. Armstrong, R. C. Mockler, and W. J. O’Sullivan, “Free-expansion melting of a colloidal monolayer,” *Phys. Rev. Lett.*, vol. 62, 2401, 1989.
- [78] R. E. Kusner, J. A. Mann, J. Kerins, and A. J. Dahm, “Two-stage melting of a two-dimensional colloidal lattice with dipole interactions,” *Phys. Rev. Lett.*, vol. 73, 3113, 1994.
- [79] A. H. Marcus and S. A. Rice, “Phase transitions in a confined quasi-two-dimensional colloid suspension,” *Phys. Rev. E*, vol. 55, 637, 1997.
- [80] C. Eisenmann, U. Gasser, P. Keim, G. Maret, and H. H. von Grünberg, “Pair interaction of dislocations in two-dimensional crystals,” *Phys. Rev. Lett.*, vol. 95, 185502, 2005.
- [81] K. N. Pham, A. M. Puertas, J. Bergenholtz, S. U. Egelhaaf, A. Moussaid, P. N. Pusey, A. B. Schofield, M. E. Cates, M. Fuchs, and W. C. K. Poon, “Multiple glassy states in a simple model system,” *Science*, vol. 296, 104, 2002.
- [82] D. R. Nelson, “Liquids and glasses in spaces of incommensurate curvature,” *Phys. Rev. Lett.*, vol. 50, 982, 1983.
- [83] M. Rubinstein and D. R. Nelson, “Dense-packed arrays on surfaces of constant negative curvature,” *Phys. Rev. B*, vol. 28, 6377, 1983.
- [84] L. Cipelletti and L. Ramos, “Slow dynamics in glassy soft matter,” *J. Phys. Condens. Matter*, vol. 17, R253, 2005.

-
- [85] V. Viasnoff and F. m. c. Lequeux, “Rejuvenation and overaging in a colloidal glass under shear,” *Phys. Rev. Lett.*, vol. 89, 065701, 2002.
 - [86] Christian Dietrich, Miglena Angelova, and Bernard Pouligny, “Adhesion of latex spheres to giant phospholipid vesicles: Statics and dynamics,” *J. Phys. II France*, vol. 7, 1651, 1997.
 - [87] B. J. Reynwar, G. Illya, V. A. Harmandaris, M. M. Müller, K. Kremer, and M. Deserno, “Aggregation and vesiculation of membrane proteins by curvature-mediated interactions,” *Nature*, vol. 447, 461, 2007.
 - [88] K. Farsad, N. Ringstad, K. Takei, S. R. Floyd, K. Rose, and P. De Camilli, “Generation of high curvature membranes mediated by direct endophilin bilayer interactions ,” *J. Cell Biol.*, vol. 155, 193, 2001.
 - [89] T. Itoh, K. S. Erdmann, A. Roux, B. Habermann, H. Werner, and P. De Camilli, “Dynamamin and the actin cytoskeleton cooperatively regulate plasma membrane invagination by BAR and F-BAR proteins,” *Developmental Cell*, vol. 9, 791 , 2005.
 - [90] B. Antonny, “Membrane deformation by protein coats,” *Current Opinion in Cell Biology*, vol. 18, 386 , 2006.
 - [91] P. D. Blood and G. A. Voth, “Direct observation of Bin/amphiphysin/Rvs (BAR) domain-induced membrane curvature by means of molecular dynamics simulations,” *Proc. Natl. Acad. Sci.*, vol. 103, 15068, 2006.
 - [92] I. Koltover, J. O. Rädler, and C. R. Safinya, “Membrane mediated attraction and ordered aggregation of colloidal particles bound to giant phospholipid vesicles,” *Phys. Rev. Lett.*, vol. 82, 1991, 1999.
 - [93] D. Stamou, C. Duschl, and D. Johannsmann, “Long-range attraction between colloidal spheres at the air-water interface: The consequence of an irregular meniscus,” *Phys. Rev. E*, vol. 62, 5263, 2000.
 - [94] D. Kashchiev, “Preface,” in “Nucleation,” ix – xii, Butterworth-Heinemann, Oxford, 2000.
 - [95] A. Laaksonen, V. Talanquer, and D. W. Oxtoby, “Nucleation: Measurements, theory, and atmospheric applications,” *Annual Review of Physical Chemistry*, vol. 46, 489, 1995.
 - [96] E. Ruckenstein and Y. Djikaev, “Recent developments in the kinetic theory of nucleation,” *Advances in Colloid and Interface Science*, vol. 118, 51 , 2005.
 - [97] U. Gasser, E. R. Weeks, A. Schofield, P. N. Pusey, and D. A. Weitz, “Real-space imaging of nucleation and growth in colloidal crystallization,” *Science*, vol. 292, 258, 2001.
 - [98] T. Zhang and X. Liu, “Nucleation: What happens at the initial stage?” *Angewandte Chemie International Edition*, vol. 48, 1308, 2009.
 - [99] J. Zhu, M. Li, R. Rogers, W. Meyer, R. H. Ottewill, W. B. Russel, P. M. Chaikin, and S.-. S. S. Crew, “Crystallization of hard-sphere colloids in microgravity,” *Nature*, vol. 387, 883, 1997.
 - [100] C. Schuh and T. Nieh, “A nanoindentation study of serrated flow in bulk metallic glasses,” *Acta Materialia*, vol. 51, 87 , 2003.

-
- [101] A. Gouldstone, H. j. Koh, K. y. Zeng, A. E. Giannakopoulos, and S. Suresh, "Discrete and continuous deformation during nanoindentation of thin films," *Acta Mater.*, vol. 48, 2277, 2000.
 - [102] S. Suresh, T.-G. Nieh, and B. Choi, "Nano-indentation of copper thin films on silicon substrates," *Scripta Materialia*, vol. 41, 951, 1999.
 - [103] S. Suresh, "Crystal deformation: Colloid model for atoms," *Nat. Mater.*, vol. 5, 253, 2006.
 - [104] P. Schall, I. Cohen, D. A. Weitz, and F. Spaepen, "Visualizing dislocation nucleation by indenting colloidal crystals," *Nature*, vol. 440, 319, 2006.
 - [105] J. J. Thomson, "Xxiv. On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure," *Philos. Mag.*, vol. 7, 237, 1904.
 - [106] S. Li, C. P. Hill, W. I. Sundquist, and J. T. Finch, "Image reconstruction of helical assemblies of the HIV-1 CA protein," *Nature*, vol. 407, 409, 2000.
 - [107] L. E. Perotti, S. Dharmavaram, W. Klug, J. Marian, J. Rudnick, and R. F. Bruinsma, "Useful scars: Physics of the capsids of archaeal viruses," *Phys. Rev. E*, vol. 94, 012404, 2016.
 - [108] A. Zlotnick, "Theoretical aspects of virus capsid assembly," *J. Mol. Recognit.*, vol. 18, 479, 2005.
 - [109] C. J. Comerci, J. Herrmann, J. Yoon, F. Jabbarpour, X. Zhou, J. F. Nomellini, J. Smit, L. Shapiro, S. Wakatsuki, and W. E. Moerner, "Topologically-guided continuous protein crystallization controls bacterial surface layer self-assembly," *Nat. Commun.*, vol. 10, 2731, 2019.
 - [110] D. S. Roshal, K. Azzag, E. Le Goff, S. B. Rochal, and S. Baghdiguian, "Crystal-like order and defects in metazoan epithelia with spherical geometry," *Sci. Rep.*, vol. 10, 7652, 2020.
 - [111] T. Yoshino, A. Matsuoka, T. Kurihara, N. Ishida, N. Kishimoto, K. Kimoto, and S. Matsuura, "Application of Voronoi tessellation of spherical surface to geometrical models of skeleton forms of spherical radiolaria," *Forma*, vol. 27, 45, 2012.
 - [112] H. W. Kroto, J. R. Heath, S. C. O'Brien, R. F. Curl, and R. E. Smalley, "C₆₀: Buckminsterfullerene," *Nature*, vol. 318, 162, 1985.
 - [113] H. P. Schultz, "Topological organic chemistry. polyhedranes and prismanes," *J. Org. Chem.*, vol. 30, 1361, 1965.
 - [114] P. W. Fowler, S. Nikolić, R. De Los Reyes, and W. Myrvold, "Distributed curvature and stability of fullerenes," *Phys. Chem. Chem. Phys.*, vol. 17, 23257, 2015.
 - [115] P. Schwerdtfeger, L. N. Wirz, and J. Avery, "The topology of fullerenes," *Wiley Interdiscip. Rev. Comput. Mol. Sci.*, vol. 5, 96, 2015.
 - [116] V. Vadakkumbatt, E. Joseph, A. Pal, and A. Ghosh, "Studying electrons on curved surfaces by trapping and manipulating multielectron bubbles in liquid helium," *Nat. Commun.*, vol. 5, 4571, 2014.

-
- [117] R. Aveyard, B. P. Binks, and J. H. Clint, “Emulsions stabilised solely by colloidal particles,” *Adv. Colloid Interface Sci.*, vol. 100, 503, 2003.
 - [118] S. Asakura and F. Oosawa, “On interaction between two bodies immersed in a solution of macromolecules,” *J. Chem. Phys.*, vol. 22, 1255, 1954.
 - [119] R. Roth, B. Götzelmann, and S. Dietrich, “Depletion forces near curved surfaces,” *Phys. Rev. Lett.*, vol. 83, 448, 1999.
 - [120] W. T. M. Irvine, V. Vitelli, and P. M. Chaikin, “Pleats in crystals on curved surfaces,” *Nature*, vol. 468, 947, 2010.
 - [121] W. T. M. Irvine, M. J. Bowick, and P. M. Chaikin, “Fractionalization of interstitials in curved colloidal crystals,” *Nat. Mater.*, vol. 11, 948, 2012.
 - [122] R. E. Guerra, C. P. Kelleher, A. D. Hollingsworth, and P. M. Chaikin, “Freezing on a sphere,” *Nature*, vol. 554, 346, 2018.
 - [123] B. J. Park and E. M. Furst, “Attractive interactions between colloids at the oil–water interface,” *Soft Matter*, vol. 7, 7676, 2011.
 - [124] D. Ershov, J. Sprakel, J. Appel, M. A. Cohen Stuart, and J. van der Gucht, “Capillarity-induced ordering of spherical colloids on an interface with anisotropic curvature,” *Proc. Natl. Acad. Sci.*, vol. 110, 9220, 2013.
 - [125] A. Caciagli, R. Singh, D. Joshi, R. Adhikari, and E. Eiser, “Controlled optofluidic crystallization of colloids tethered at interfaces,” *Phys. Rev. Lett.*, vol. 125, 068001, 2020.
 - [126] S. Smale, “Mathematical problems for the next century,” *Math. Intell.*, vol. 20, 7, 1998.
 - [127] E. L., “*Elementa doctrinae solidorum*,” *Novi Commentarii academiae scientiarum Petropolitanae*, vol. 26, 109, 1750.
 - [128] P. Hilton and J. Pedersen, “The Euler characteristic and Pólya’s dream,” *The American Mathematical Monthly*, vol. 103, 121, 1996.
 - [129] D. R. Nelson, *Defects and Geometry in Condensed Matter Physics*, Cambridge University Press, Cambridge, 2002.
 - [130] D. J. Wales and S. Ulker, “Structure and dynamics of spherical crystals characterized for the Thomson problem,” *Phys. Rev. B*, vol. 74, 212101, 2006.
 - [131] M. J. Bowick, A. Cacciuto, D. R. Nelson, and A. Travesset, “Crystalline particle packings on a sphere with long-range power-law potentials,” *Phys. Rev. B*, vol. 73, 024115, 2006.
 - [132] T. Einert, P. Lipowsky, J. Schilling, M. J. Bowick, and A. R. Bausch, “Grain boundary scars on spherical crystals,” *Langmuir*, vol. 21, 12076, 2005.
 - [133] A. Travesset, “Ground state of a large number of particles on a frozen topography,” *Phys. Rev. E*, vol. 72, 036110, 2005.
 - [134] S. Giarritta, M. Ferrario, and P. Giaquinta, “Statistical geometry of hard particles on a sphere: analysis of defects at high density,” *Physica A*, vol. 201, 649, 1993.
 - [135] P. Lipowsky, M. J. Bowick, J. H. Meinke, D. R. Nelson, and A. R. Bausch, “Direct visualization of dislocation dynamics in grain-boundary scars,” *Nat. Mater.*, vol. 4, 407, 2005.

-
- [136] C. Köhler, R. Backofen, and A. Voigt, “Stress induced branching of growing crystals on curved surfaces,” *Phys. Rev. Lett.*, vol. 116, 135502, 2016.
 - [137] S. Schneider and G. Gompper, “Shapes of crystalline domains on spherical fluid vesicles,” *Europhys. Lett.*, vol. 70, 136, 2005.
 - [138] S. Paquay, G.-J. Both, and P. van der Schoot, “Impact of interaction range and curvature on crystal growth of particles confined to spherical surfaces,” *Phys. Rev. E*, vol. 96, 012611, 2017.
 - [139] Z. Yao, “Topological vacancies in spherical crystals,” *Soft Matter*, vol. 13, 5905, 2017.
 - [140] J.-P. Vest, G. Tarjus, and P. Viot, “Dynamics of a monodisperse Lennard–Jones system on a sphere,” *Mol. Phys.*, vol. 112, 1330, 2014.
 - [141] A. Azadi and G. M. Grason, “Emergent structure of multidislocation ground states in curved crystals,” *Phys. Rev. Lett.*, vol. 112, 225502, 2014.
 - [142] A. Luque, D. Reguera, A. Morozov, J. Rudnick, and R. Bruinsma, “Physics of shell assembly: Line tension, hole implosion, and closure catastrophe,” *J. Chem. Phys.*, vol. 136, 184507, 2012.
 - [143] M. Tribelsky and S. Anisimov, “Tuned Mullins-Sekerka instability: Exact results,” *Phys. Rev. E*, vol. 90, 042403, 2014.
 - [144] W. W. Mullins and R. F. Sekerka, “Morphological stability of a particle growing by diffusion or heat flow,” *J. Appl. Phys.*, vol. 34, 323, 1963.
 - [145] W. W. Mullins and R. F. Sekerka, “Stability of a planar interface during solidification of a dilute binary alloy,” *J. Appl. Phys.*, vol. 35, 444, 1964.
 - [146] L. Ma, X. Liu, A.-k. Soh, L. He, C. Wu, and Y. Ni, “Growth of curved crystals: competition between topological defect nucleation and boundary branching,” *Soft Matter*, vol. 15, 4391, 2019.
 - [147] H. T. McMahon and E. Boucrot, “Membrane curvature at a glance,” *J. Cell Sci.*, vol. 128, 1065, 2015.
 - [148] M. Schmick and P. Bastiaens, “The interdependence of membrane shape and cellular signal processing,” *Cell*, vol. 156, 2014.
 - [149] Y. Shibata, J. Hu, M. M. Kozlov, and T. A. Rapoport, “Mechanisms shaping the membranes of cellular organelles,” *Annu. Rev. Cell Dev. Biol.*, vol. 25, 329, 2009.
 - [150] J. E. Rothman and L. Orci, “Budding vesicles in living cells,” *Scientific American*, vol. 274, 70, 1996.
 - [151] J. Heuser, “Three-dimensional visualization of coated vesicle formation in fibroblasts,” *J. Cell Biol.*, vol. 84, 560, 1980.
 - [152] N. A. Bright, M. J. Gratian, and J. Luzio, “Endocytic delivery to lysosomes mediated by concurrent fusion and kissing events in living cells,” *Curr. Biol.*, vol. 15, 360, 2005.
 - [153] J. Liu, Y. Sun, D. G. Drubin, and G. F. Oster, “The mechanochemistry of endocytosis,” *PLOS Biol.*, vol. 7, 1, 2009.
 - [154] L. Yao, N. Sharifi-Mood, I. B. Liu, and K. J. Stebe, “Capillary migration of microdisks on curved interfaces,” *J. Colloid Interface Sci.*, vol. 449, 436, 2015.

-
- [155] N. Sharifi-Mood, I. B. Liu, and K. J. Stebe, “Curvature capillary migration of microspheres,” *Soft Matter*, vol. 11, 6768, 2015.
 - [156] I. B. Liu, N. Sharifi-Mood, and K. J. Stebe, “Capillary assembly of colloids: Interactions on planar and curved interfaces,” *Annu. Rev. Condens. Matter Phys.*, vol. 9, 283, 2018.
 - [157] M. Cavallaro, L. Botto, E. P. Lewandowski, M. Wang, and K. J. Stebe, “Curvature-driven capillary migration and assembly of rod-like particles,” *Proc. Natl. Acad. Sci.*, vol. 108, 20923, 2011.
 - [158] K. J. Stebe, E. Lewandowski, and M. Ghosh, “Oriented assembly of metamaterials,” *Science*, vol. 325, 159, 2009.
 - [159] N. Engheta, “Circuits with light at nanoscales: Optical nanocircuits inspired by metamaterials,” *Science*, vol. 317, 1698, 2007.
 - [160] Y. Liu and X. Zhang, “Metamaterials: a new frontier of science and technology,” *Chem. Soc. Rev.*, vol. 40, 2494, 2011.
 - [161] R. A. Petros and J. M. DeSimone, “Strategies in the design of nanoparticles for therapeutic applications,” *Nat. Rev. Drug Discov.*, vol. 9, 615, 2010.
 - [162] G. Emilsson, E. Röder, B. Malekian, K. Xiong, J. Manzi, F.-C. Tsai, N.-J. Cho, M. Bally, and A. Dahlin, “Nanoplasmonic sensor detects preferential binding of IRSp53 to negative membrane curvature,” *Frontiers in Chemistry*, vol. 7, 1, 2019.
 - [163] H. Lee and D. Kim, “Analysis of curvature effects on plasmon biosensing of molecular interactions,” in “*Plasmonics in Biology and Medicine XIV*,” vol. 10080, 84 – 88, International Society for Optics and Photonics, SPIE, 2017.
 - [164] E. Páram, H. Le, and A. Fernández-Nieves, “Stability of toroidal droplets inside yield stress materials,” *Phys. Rev. E*, vol. 90, 021002, 2014.
 - [165] R. Gillette and D. Dyson, “Stability of fluid interfaces of revolution between equal solid circular plates,” *Chem. Eng. J.*, vol. 2, 44, 1971.
 - [166] K. G. López, *Colloidal Crystals on Toroidal Curved Surfaces*, Master’s thesis, Leiden University, 2019.
 - [167] V. Vitelli, L. B. Lucks, and D. R. Nelson, “Crystallography on curved surfaces,” *Proc. Natl. Acad. Sci.*, vol. 103, 12312, 2006.
 - [168] L. Giomi and M. J. Bowick, “Defective ground states of toroidal crystals,” *Phys. Rev. E*, vol. 78, 010601, 2008.
 - [169] H. Kusumaatmaja and D. J. Wales, “Defect motifs for constant mean curvature surfaces,” *Phys. Rev. Lett.*, vol. 110, 165502, 2013.
 - [170] C. J. Burke, B. L. Mbanda, Z. Wei, P. T. Spicer, and T. J. Atherton, “The role of curvature anisotropy in the ordering of spheres on an ellipsoid,” *Soft Matter*, vol. 11, 5872, 2015.
 - [171] F. L. Jiménez, N. Stoop, R. Lagrange, J. Dunkel, and P. M. Reis, “Curvature-controlled defect localization in elastic surface crystals,” *Phys. Rev. Lett.*, vol. 116, 104301, 2016.
 - [172] V. Schmid and A. Voigt, “Crystalline order and topological charges on capillary bridges,” *Soft Matter*, vol. 10, 4694, 2014.

-
- [173] E. Bendito, M. J. Bowick, A. Medina, and Z. Yao, “Crystalline particle packings on constant mean curvature (Delaunay) surfaces,” *Phys. Rev. E*, vol. 88, 012405, 2013.
 - [174] T. Baumgart, S. Das, W. Webb, and J. Jenkins, “Membrane elasticity in giant vesicles with fluid phase coexistence,” *Biophys. J.*, vol. 89, 1067, 2005.
 - [175] S. T. Hess, M. V. Gudheti, M. Mlodzianoski, and T. Baumgart, *Shape Analysis of Giant Vesicles With Fluid Phase Coexistence by Laser Scanning Microscopy to Determine Curvature, Bending Elasticity, and Line Tension*, 367–387, Humana Press, Totowa, NJ, 2007.
 - [176] S. Semrau and T. Schmidt, “Membrane heterogeneity – from lipid domains to curvature effects,” *Soft Matter*, vol. 5, 3174, 2009.
 - [177] D. Marenduzzo and E. Orlandini, “Phase separation dynamics on curved surfaces,” *Soft Matter*, vol. 9, 1178, 2012.
 - [178] F. Paillusson, M. R. Pennington, and H. Kusumaatmaja, “Phase separation on bicontinuous cubic membranes: Symmetry breaking, reentrant, and domain faceting,” *Phys. Rev. Lett.*, vol. 117, 058101, 2016.
 - [179] N. A. Garcia, A. D. Pezzutti, R. A. Register, D. A. Vega, and L. R. Gomez, “Defect formation and coarsening in hexagonal 2D curved crystals,” *Soft Matter*, vol. 11, 898, 2015.
 - [180] D. M. Kaz, R. McGorty, M. Mani, M. P. Brenner, and V. N. Manoharan, “Physical ageing of the contact line on colloidal particles at liquid interfaces,” *Nat. Mater.*, vol. 11, 138, 2012.
 - [181] I. Koltover, J. O. Rädler, and C. R. Safinya, “Membrane mediated attraction and ordered aggregation of colloidal particles bound to giant phospholipid vesicles,” *Phys. Rev. Lett.*, vol. 82, 1991, 1999.
 - [182] J. Agudo-Canalejo and R. Lipowsky, “Uniform and Janus-like nanoparticles in contact with vesicles: energy landscapes and curvature-induced forces,” *Soft Matter*, vol. 13, 2155, 2017.
 - [183] A. Vahid, A. Saric, and T. Idema, “Curvature variation controls particle aggregation on fluid vesicles,” *Soft Matter*, vol. 13, 4924, 2017.
 - [184] R. Sarfati and E. R. Dufresne, “Long-range attraction of particles adhered to lipid vesicles,” *Phys. Rev. E*, vol. 94, 012604, 2016.
 - [185] A. Šarić and A. Cacciuto, “Fluid membranes can drive linear aggregation of adsorbed spherical nanoparticles,” *Phys. Rev. Lett.*, vol. 108, 118101, 2012.
 - [186] A. Bahrami and A. H. Bahrami, “Vesicle constriction by rings of janus nanoparticles and aggregates of curved proteins,” *Nanotechnology*, vol. 30, 345101, 2019.
 - [187] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of state calculations by fast computing machines,” *J. Chem. Phys.*, vol. 21, 1087, 1953.
 - [188] D. Frenkel and B. Smit, *Understanding Molecular Simulation*, Academic Press, San Diego, 2nd edition, 2002.
 - [189] A. Rahman, “Correlations in the motion of atoms in liquid argon,” *Phys. Rev.*, vol. 136, A405, 1964.

-
- [190] Z. Li and H. A. Scheraga, “Monte Carlo-minimization approach to the multiple-minima problem in protein folding,” *Proc. Natl. Acad. Sci.*, vol. 84, 6611, 1987.
- [191] N. Metropolis, “The beginning of the Monte Carlo method,” *Los Alamos Sci.*, vol. 15, 125, 1987.
- [192] L. Martino and J. Míguez, “Generalized rejection sampling schemes and applications in signal processing,” *Signal Process.*, vol. 90, 2981, 2010.
- [193] A. Troisi, V. Wong, and M. A. Ratner, “Self-assembly on multiple length scales: A Monte Carlo algorithm with data augmentation,” *J. Chem. Phys.*, vol. 122, 024102, 2005.
- [194] J. P. K. Doye, D. J. Wales, and R. S. Berry, “The effect of the range of the potential on the structures of clusters,” *J. Chem. Phys.*, vol. 103, 4234, 1995.
- [195] R. H. Swendsen and J.-S. Wang, “Replica Monte Carlo simulation of spin-glasses,” *Phys. Rev. Lett.*, vol. 57, 2607, 1986.
- [196] P. R. ten Wolde and D. Frenkel, “Computer simulation study of gas-liquid nucleation in a Lennard-Jones system,” *J. Chem. Phys.*, vol. 109, 9901, 1998.
- [197] L. Filion, M. Hermes, R. Ni, and M. Dijkstra, “Crystal nucleation of hard spheres using molecular dynamics, umbrella sampling, and forward flux sampling: A comparison of simulation techniques,” *J. Chem. Phys.*, vol. 133, 244115, 2010.
- [198] O.-Y. Zhong-can and W. Helfrich, “Instability and deformation of a spherical vesicle by pressure,” *Phys. Rev. Lett.*, vol. 59, 2486, 1987.
- [199] Q. Yan and J. J. de Pablo, “Hyper-parallel tempering Monte Carlo: Application to the Lennard-Jones fluid and the restricted primitive model,” *J. Chem. Phys.*, vol. 111, 9509, 1999.
- [200] R. M. Brannon, “A review of useful theorems involving proper orthogonal matrices referenced to three-dimensional physical space,” Technical report, the University of Utah, 2002.
- [201] J. O. Law, A. G. Wong, H. Kusumaatmaja, and M. A. Miller, open data, 2018, URL <https://doi.org/10.15128/r1jh343s29p>.
- [202] J. O. Law, J. M. Dean, M. A. Miller, and H. Kusumaatmaja, open data, 2020, URL <https://doi:10.15128/r1wh246s13w>.
- [203] J. O. Law, A. G. Wong, H. Kusumaatmaja, and M. A. Miller, “Nucleation on a sphere: the roles of curvature, confinement and ensemble,” *Mol. Phys.*, vol. 116, 3008, 2018.
- [204] J. O. Law, J. M. Dean, M. A. Miller, and H. Kusumaatmaja, “Phase transitions on non-uniformly curved surfaces: coupling between phase and location,” *Soft Matter*, vol. 16, 8069, 2020.
- [205] B. J. Alder and T. E. Wainwright, “Studies in molecular dynamics. I. general method,” *J. Chem. Phys.*, vol. 31, 459, 1959.
- [206] S. Plimpton, “Fast parallel algorithms for short-range molecular dynamics,” *J. Comput. Phys.*, vol. 117, 1, 1995.
- [207] R. D. Engle, R. D. Skeel, and M. Drees, “Monitoring energy drift with shadow hamiltonians,” *J. Comput. Phys.*, vol. 206, 432, 2005.

-
- [208] S. Toxvaerd, O. J. Heilmann, and J. C. Dyre, “Energy conservation in molecular dynamics simulations of classical systems,” *J. Comput. Phys.*, vol. 136, 224106, 2012.
 - [209] T. Schneider and E. Stoll, “Molecular-dynamics study of a three-dimensional one-component model for distortive phase transitions,” *Phys. Rev. B*, vol. 17, 1302, 1978.
 - [210] S. Paquay and R. Kusters, “A method for molecular dynamics on curved surfaces,” *Biophys. J.*, vol. 110, 1226, 2016.
 - [211] H. C. Andersen, “Rattle: A “velocity” version of the shake algorithm for molecular dynamics calculations,” *J. Comput. Phys.*, vol. 52, 24, 1983.
 - [212] D. J. Wales and J. P. K. Doye, “Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms,” *J. Phys. Chem. A*, vol. 101, 5111, 1997.
 - [213] J. P. K. Doye and D. J. Wales, “The effect of the range of the potential on the structure and stability of simple liquids: from clusters to bulk, from sodium to,” *J. Phys. B*, vol. 29, 4859, 1996.
 - [214] J. P. K. Doye and D. J. Wales, “The structure and stability of atomic liquids: From clusters to bulk,” *Science*, vol. 271, 484, 1996.
 - [215] D. J. Wales, “Gmin,” [Online; accessed 29-October-2019], URL <http://www-wales.ch.cam.ac.uk/GMIN/>.
 - [216] G. A. Vliegthart and G. Gompper, “Compression, crumpling and collapse of spherical shells and capsules,” *New J. Phys.*, vol. 13, 045020, 2011.
 - [217] K. Tamotsu, D. Kroll, and G. Gompper, “Budding of crystalline domains in fluid membranes,” *Phys. Rev. E*, vol. 68, 061905, 2004.
 - [218] P.-F. Lenne and A. Nicolas, “Physics puzzles on membrane domains posed by cell biology,” *Soft Matter*, vol. 5, 2841, 2009.
 - [219] D. Lingwood and K. Simons, “Lipid rafts as a membrane-organizing principle,” *Science*, vol. 327, 46, 2010.
 - [220] V. Y. Kisilev, D. Marenduzzo, and A. B. Goryachev, “Lateral dynamics of proteins with polybasic domain on anionic membranes: A dynamic Monte-Carlo study,” *Biophys. J.*, vol. 100, 1261, 2011.
 - [221] M. Park, C. Harrison, P. M. Chaikin, R. A. Register, and D. H. Adamson, “Block copolymer lithography: Periodic arrays of 1011 holes in 1 square centimeter,” *Science*, vol. 276, 1401, 1997.
 - [222] E. J. Kramer, “Condensed-matter physics: Melted by mistakes,” *Small*, vol. 437, 824, 2005.
 - [223] G. A. DeVries, M. Brunnbauer, Y. Hu, A. M. Jackson, B. Long, B. T. Neltner, O. Uzun, B. H. Wunsch, and F. Stellacci, “Divalent metal nanoparticles,” *Science*, vol. 315, 358, 2007.
 - [224] C. Mohrdieck, F. Dalmas, E. Arzt, R. Tharmann, M. Claessens, A. Bausch, A. Roth, E. Sackmann, C. H J Schmitz, J. Curtis, W. Roos, S. Schulz, K. Uhrig, and J. P Spatz, “Biomimetic models of the actin cytoskeleton,” *Small*, vol. 3, 1015, 2007.

-
- [225] L. Limozin, A. Roth, and E. Sackmann, "Microviscoelastic moduli of biomimetic cell envelopes," *Phys. Rev. Lett.*, vol. 95, 178101, 2005.
 - [226] M. Rao and B. J. Berne, "Nucleation in finite systems: Theory and computer simulation," *Astrophys. Space Sci.*, vol. 65, 39, 1979.
 - [227] J. M. Caillol, "Critical-point of the Lennard-Jones fluid: A finite-size scaling study," *J. Chem. Phys.*, vol. 109, 4885, 1998.
 - [228] N. B. Wilding, "Critical-point and coexistence-curve properties of the Lennard-Jones fluid: A finite-size scaling study," *Phys. Rev. E*, vol. 52, 602, 1995.
 - [229] C. Chen, O. Cook, C. E. Nicholson, and S. J. Cooper, "Leapfrogging Ostwald's rule of stages: Crystallization of stable γ -glycine directly from microemulsions," *Cryst. Growth Des.*, vol. 11, 2228, 2011.
 - [230] C. E. Nicholson, C. Chen, B. Mendis, and S. J. Cooper, "Sable polymorphs crystallized directly under thermodynamic control in three-dimensional nanoconfinement: A generic methodology," *Cryst. Growth Des.*, vol. 11, 363, 2011.
 - [231] D. Reguera, R. K. Bowles, Y. Djikaev, and H. Reiss, "Phase transitions in systems small enough to be clusters," *J. Chem. Phys.*, vol. 118, 340, 2003.
 - [232] J. E. Jones, "On the determination of molecular fields. ii. from the equation of state of a gas," *Proc. Royal Soc. A*, vol. 106, 463, 1924.
 - [233] N. Gribova, A. Arnold, T. Schilling, and C. Holm, "How close to two dimensions does a Lennard-Jones system need to be to produce a hexatic phase?" *J. Chem. Phys.*, vol. 135, 054514, 2011.
 - [234] B. Smit and D. Frenkel, "Vapor-liquid equilibria of the two-dimensional Lennard-Jones fluid(s)," *J. Chem. Phys.*, vol. 94, 5663, 1991.
 - [235] R. D. Mountain and D. Thirumalai, "Quantative measure of efficiency of Monte Carlo simulations," *Physica A*, vol. 210, 453, 1994.
 - [236] G. Marsaglia, "Choosing a point from the surface of a sphere," *Ann. Math. Stat.*, vol. 43, 645, 1972.
 - [237] N. B. Wilding, "Computer simulation of fluid phase transitions," *Am. J. Phys.*, vol. 69, 1147, 2001.
 - [238] S. Weerasinghe and F. G. Amar, "Absolute classical densities of states for very anharmonic systems and applications to the evaporation of rare gas clusters," *J. Chem. Phys.*, vol. 98, 4967, 1993.
 - [239] M. A. Miller and D. Frenkel, "Phase diagram of the adhesive hard sphere fluid," *J. Chem. Phys.*, vol. 121, 535, 2004.
 - [240] F. H. Stillinger Jr., "Rigorous basis of the frenkel-band theory of association equilibrium," *J. Chem. Phys.*, vol. 38, 1486, 1963.
 - [241] M. Santra, S. Chakrabarty, and B. Bagchi, "Gas-liquid nucleation in a two dimensional system," *J. Chem. Phys.*, vol. 129, 234704, 2008.
 - [242] D. C. Johnston, *Advances in Thermodynamics of the van der Waals Fluid*, 2053-2571, Morgan & Claypool Publishers, 2014.

-
- [243] S. Ryu and W. Cai, “Validity of classical nucleation theory for ising models,” *Phys. Rev. E*, vol. 81, 030601, 2010.
- [244] I. J. Ford, “Nucleation theorems, the statistical mechanics of molecular clusters, and a revision of classical nucleation theory,” *Phys. Rev. E*, vol. 56, 5615, 1997.
- [245] S. L. Girshick and C.-P. Chiu, “Kinetic nucleation theory: A new expression for the rate of homogeneous nucleation from an ideal supersaturated vapor,” *J. Chem. Phys.*, vol. 93, 1273, 1990.
- [246] M. Schrader, P. Virnau, and K. Binder, “Simulation of vapor–liquid coexistence in finite volumes: A method to compute the surface free energy of droplets,” *Phys. Rev. E*, vol. 79, 061104, 2009.
- [247] M. Rovere, D. W. Heermann, and K. Binder, “The gas-liquid transition of the two-dimensional Lennard-Jones fluid,” *J. Phys. Condens. Matter*, vol. 2, 7009, 1990.
- [248] F. Aurenhammer, R. Klein, and D. T. Lee, *Voronoi Diagrams and Delauney Triangulations*, World Scientific publishing co. Pte. Ltd, London, 2013.
- [249] N. D. Bade, T. Xu, R. D. Kamien, R. K. Assoian, and K. J. Stebe, “Gaussian curvature directs stress fiber orientation and cell migration,” *Biophys. J.*, vol. 114, 1467, 2018.
- [250] H. T. McMahon and J. L. Gallop, “Membrane curvature and mechanisms of dynamic cell membrane remodelling,” *Nature*, vol. 438, 590, 2005.
- [251] J.-K. Kim, E. Lee, Z. Huang, and M. Lee, “Nanorings from the self-assembly of amphiphilic molecular dumbbells,” *Am. Chem. Soc.*, vol. 128, 14022, 2006.
- [252] V. Haridas, A. R. Sapala, and J. P. Jasinski, “Self-assembling triazolophanes: from croissants through donuts to spherical vesicles,” *Chem. Commun.*, vol. 51, 6905, 2015.
- [253] J. P. Snyder, “Map projections: A working manual,” Technical report, USGS Publications Warehouse, 1987.
- [254] K. A. Brakke, “The surface evolver,” *Exp. Math.*, vol. 1, 141, 1992.
- [255] T. C. T. Michaels, M. M. J. Bellaiche, M. F. Hagan, and T. P. J. Knowles, “Kinetic constraints on self-assembly into closed supramolecular structures,” *Sci. Rep.*, vol. 7, 12295, 2017.
- [256] N. Tanjeem, W. H. Wilkin, D. A. Beller, C. H. Rycroft, and V. N. Manoharan, “Crystallization on a cylinder,” preprint, 2020.
- [257] P. Fowler, “Quantification of chirality: Attempting the impossible,” *Symmetry Cult. Sci.*, vol. 16, 321, 2005.
- [258] G. T. Pickett, M. Gross, and H. Okuyama, “Spontaneous chirality in simple systems,” *Phys. Rev. Lett.*, vol. 85, 3652, 2000.
- [259] P. Curie, *Oeuvres de Pierre Curie*, Gauthiers-Villars, Paris, 1908.
- [260] P. Wolynes, “Symmetry and the energy landscapes of biomolecules,” *Proc. Natl. Acad. Sci.*, vol. 93, 14249–14255, 1996.
- [261] J. Lee, S. I. Blaber, V. K. Dubey, and M. Blaber, “A polypeptide “building block” for the β -trefoil fold identified by “top-down symmetric deconstruction”,” *J. Mol. Biol.*, vol. 407, 744, 2011.

-
- [262] D. J. Wales, "Symmetry, near-symmetry and energetics," *Chem. Phys. Lett.*, vol. 285, 330, 1998.
- [263] R. O. Erickson, "Tubular packing of spheres in biological fine structure," *Science*, vol. 181, 705, 1973.
- [264] J. Winkelmann, B. Haffner, D. Weaire, A. Mughal, and S. Hutzler, "Simulation and observation of line-slip structures in columnar structures of soft spheres," *Phys. Rev. E*, vol. 97, 012610, 2017.
- [265] G. Wu, H. Cho, D. A. Wood, A. D. Dinsmore, and S. Yang, "Confined assemblies of colloidal particles with soft repulsive interactions," *J. Am. Chem. Soc.*, vol. 139, 5095, 2017.
- [266] M. Lohr, A. Alsayed, B. Chen, Z. Zhang, R. Kamien, and A. Yodh, "Helical packings and phase transformations of soft spheres in cylinders," *Phys. Rev. E*, vol. 81, 040401, 2010.
- [267] A. N. Khlobystov, D. A. Britz, A. Ardavan, and G. A. D. Briggs, "Observation of ordered phases of fullerenes in carbon nanotubes," *Phys. Rev. Lett.*, vol. 92, 245507, 2004.
- [268] A. Mughal, H. K. Chan, and D. Weaire, "Phyllotactic description of hard sphere packing in cylindrical channels," *Phys. Rev. Lett.*, vol. 106, 115704, 2011.
- [269] A. Mughal, H. K. Chan, D. Weaire, and S. Hutzler, "Dense packings of spheres in cylinders: Simulations," *Phys. Rev. E*, vol. 85, 051305, 2012.
- [270] D. A. Wood, C. D. Santangelo, and A. D. Dinsmore, "Self-assembly on a cylinder: a model system for understanding the constraint of commensurability," *Soft Matter*, vol. 9, 10016, 2013.
- [271] J. Sun and V. N. Manoharan, 2020.
- [272] R. Kountanya, "Boundary value problems for cone geodesics," [Online; accessed 20-August-2020] at demonstrations.wolfram.com/BoundaryValueProblemsForConeGeodesics, licensed under CC BY-NC-SA: creativecommons.org/licenses/by-nc-sa/3.0/.
- [273] D. R. Nelson, "Toward a tetravalent chemistry of colloids," *Nano Lett.*, vol. 2, 1125, 2002.
- [274] R. K. Assoian, N. D. Bade, C. V. Cameron, and K. J. Stebe, "Cellular sensing of micron-scale curvature: a frontier in understanding the microenvironment," *Open Biol.*, vol. 9, 190155, 2019.
- [275] P. Schuille and S. Diez, "Synthetic biology of minimal systems," *Crit. Rev. Biochem. Mol. Biol.*, vol. 44, 223, 2009.
- [276] Y. Dreher, K. Jahnke, E. Bobkova, J. P. Spatz, and K. Göpfrich, "Division and regrowth of phase-separated giant unilamellar vesicles," *Angew. Chem. Int. Ed.*, vol. Accepted Author Manuscript., 2020.
- [277] T. Jones, A. Liu, and B. Cui, "Light-inducible generation of membrane curvature in live cells with engineered bar domain proteins," *ACS Synth. Biol.*, vol. 9, 893, 2020.

-
- [278] E. De Santis, H. Alkassem, B. Lamarre, N. Faruqui, A. Bella, J. E. Noble, N. Micale, S. Ray, J. R. Burns, A. R. Yon, B. W. Hoogenboom, and M. G. Ryadnov, “Antimicrobial peptide capsids of de novo design,” *Nat. Commun.*, vol. 8, 2263, 2017.
- [279] O. M. Becker and M. Karplus, “The topology of multidimensional potential energy surfaces: Theory and application to peptide structure and kinetics,” *J. Chem. Phys.*, vol. 106, 1495, 1997.
- [280] D. J. Wales, M. A. Miller, and T. R. Walsh, “Archetypal energy landscapes,” *Nature*, vol. 394, 578, 1998.
- [281] B. J. Peter, H. M. Kent, I. G. Mills, Y. Vallis, P. J. G. Butler, P. R. Evans, and H. T. McMahon, “BAR domains as sensors of membrane curvature: The amphiphysin BAR structure,” *Science*, vol. 303, 495, 2004.
- [282] A. Frost, V. M. Unger, and P. De Camilli, “The BAR domain superfamily: Membrane-molding macromolecules,” *Cell*, vol. 137, 191, 2009.
- [283] J. L. Gallop, C. C. Jao, H. M. Kent, P. J. G. Butler, P. R. Evans, R. Langen, and H. T. McMahon, “Mechanism of endophilin N-BAR domain-mediated membrane curvature,” *EMBO J.*, vol. 25, 2898, 2006.
- [284] B. J. Galletta, O. L. Mooren, and J. A. Cooper, “Actin dynamics and endocytosis in yeast and mammals,” *Curr. Opin. Biotechnol.*, vol. 21, 604, 2010.
- [285] K. Khairy, E. Reynaud, and E. Stelzer, “Detection of deformable objects in 3D images using Markov-chain Monte Carlo and spherical harmonics,” in “Medical Image Computing and Computer-Assisted Intervention – MICCAI 2008,” 1075–1082, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [286] N. Ramakrishnan, P. B. Sunil Kumar, and J. H. Ipsen, “Monte Carlo simulations of fluid vesicles with in-plane orientational ordering,” *Phys. Rev. E*, vol. 81, 041922, 2010.
- [287] J.-S. Ho and A. Baumgärtner, “Simulations of flexible manifolds,” *Mol. Sim.*, vol. 6, 163, 1991.
- [288] E. P. Lewandowski, J. A. Bernate, P. C. Searson, and K. J. Stebe, “Rotation and alignment of anisotropic particles on nonplanar interfaces,” *Langmuir*, vol. 24, 9302, 2008.