1 Introduction

1.1 Discontinuous Fiber Composites

A composite is any material with two or more constituents that retain their identities in the final material. Everyday examples range from plywood, which is layers of wood veneer glued together, to concrete, a mixture of gravel and sand bonded by Portland cement.

Many composites contain fibers, often glass or carbon, embedded in a polymer matrix. When the material must be as strong, stiff and light as possible, and when high manufacturing cost can be tolerated, the fibers are continuous and are either aligned in thin layers much like plywood, or woven into a fabric. Continuous fiber composites are widely used in aircraft, wind-turbine blades, boat hulls, and in sporting goods such as tennis rackets, bicycles and golf clubs.

When manufacturing costs must be low or when the final component has a complex shape, the fibers are cut or chopped into shorter lengths to create a *discontinuous fiber composite* (DFC). The fibers and matrix polymer are combined when the matrix is in a liquid state, then formed into final shape by processes like injection or compression molding. Discontinuous fiber composites are used in business machines, power tools and many other products. In automobiles they are used for body panels, structural components and under-the-hood components. Since the 1970s the automotive industry has been the largest user of fiber reinforced composites [GGOS20], and almost all automotive composites have discontinuous fibers. Discontinuous fiber composites are the focus of this book.

Many discontinuous fiber composites have thermoplastic matrices and are processed by injection molding. These are often subdivided into short fiber/thermoplastics (SFTs) and long fiber/thermoplastics (LFTs). Pellets of the short fiber materials are prepared by mixing the fibers and matrix in a compounding extruder, and average fiber lengths are less than 1 mm. Long fiber materials are prepared by impregnating a continuous fiber tow with the matrix resin, usually by an extrusion process, then cutting the coated fibers to a length of approximately 12 mm to make pellets. Pellets of either type are processed in a conventional injection molding machine, where they are re-melted in the plasticizing screw and injected into the mold. For LFT materials this substantially reduces the fiber length, but final lengths greater than 1 mm are possible with careful processing. We examine both SFTs and LFTs in detail in this book.

Another important type of discontinuous fiber composite is sheet molding compound (SMC), which has a thermosetting polymer matrix with chopped fibers. The molding compound is initially produced by combining chopped fiber bundles, 25 to 50 mm long, with a liquid resin, particulate filler and other chemicals at room temperature. The resin is allowed to thicken until the sheet is can be handled but is still soft. Layers of this material are cut and stacked in a specified shape, called the charge, then placed into a hot compression mold. Closing the mold squeezes the charge, causing it to flow and fill the mold cavity, while heat from the mold initiates a curing reaction that hardens the polymer. After one to two minutes, the mold is opened and the solid part is removed.

There are many other types of discontinuous fiber composites. In direct in-line processing of long fiber/thermoplastic composites (D-LFTs) the polymer is melted in a compounding extruder, and continuous glass fibers are added at a port partway down the extruder barrel. The hot fiber/resin mixture is extruded from a die, cut into a strand, and placed immediately into a mold where it is compression molded to final shape. This process can preserve fiber lengths as long as 60 mm in the finished part.

With glass mat thermoplastic (GMT) composites, a mat of reinforcing fibers is impregnated with a polymer resin and solidified into an initial sheet form. The fibers may be continuous, woven or chopped. In a subsequent step, blanks cut from the sheets are re-heated, placed into a compression mold and formed into final shape.

Discontinuous fiber composites are also used with additive manufacturing or 3-D printing. This is most commonly done via fused-filament fabrication (FFF) or fused deposition modeling (FDM), in which a bead of molten polymer is extruded from a heated die and laid down on the bed of the 3-D printer, where it solidifies. By moving the die and/or the printer bed, beads of polymer are laid down in rows in the desired pattern. Subsequent layers are deposited on top to create the final shape. By adding discontinuous fibers to the feed material, a fiber reinforced material can be printed [RHJS18, MGB+18, OBW+20].

Some composites are reinforced with disk-shaped particles such as glass flakes or ceramic platelets. These materials are often processed by injection or compression molding. Because the particles can orient under flow, these materials have the same connection between processing, microstructure and properties as discontinuous fiber composites. The orientation and mechanical property models in this book can be applied to these composites as well.

1.2 Processing, Microstructure and Properties

While discontinuous fiber composites offer an attractive combination of mechanical properties, light weight and low cost, designing with them presents a challenge. This is because the properties of a DFC component depend on how the final part is made.

Figure 1.1 shows an example. The two boxes in this figure were injection molded from the same glass-fiber/polyamide SFT material using the same mold. However, one box is straight and the other is badly warped. The difference is the location of the gate where the molten fiber/polymer mixture first enters the mold cavity. The box on the left uses a film gate across the near end, while the warped box on the right uses a small tab gate partway down the left side. The two gate locations produce different mold filling patterns, which create different patterns of fiber orientation in the two moldings.

For both moldings, the material is above room temperature when it solidifies, then cools to room temperature after it is removed from the mold. The fibers reduce the thermal contraction of the polymer matrix along their length, but have only a small effect in other directions. If the fibers have a preferential orientation, then as the material cools it will contract less in the fiber direction but more in the perpendicular direction. In the box on the right in Fig. 1.1, the direction of fiber orientation is very different from one location to another, and this causes

Figure 1.1 Two injection molded boxes. Photo by the author.

the box to warp as it cools. In the box on the left, the fibers are oriented in nearly the same direction at all locations, so this box does not warp. Processing has affected the properties of the material in these boxes.

A second example appears in Fig. 1.2, which shows stress-strain curves for samples of sheet molding compound. All of the curves are for the same batch of material and were molded to the same final thickness in the same rectangular plaque mold. However, some samples were made by cutting the charge to nearly fit the mold, while other samples used charges with smaller area but more layers. The charges were all designed to flow in only one direction as they were compressed, as in Fig. 4.28(a). Also, the tensile specimens were cut from the plaques so that some were loaded parallel to the flow direction (0°), while some were loaded perpendicular to the flow direction (90°).

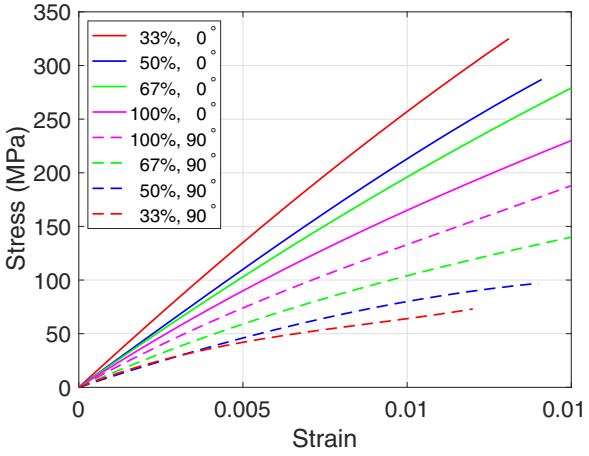

Figure 1.2 Tensile stress-strain curves for samples of sheet molding compound with 65% by weight of glass fibers. The legend gives the percentage of the mold area initially covered by the charge and the direction of the tensile test relative to the flow. Data from [CT84].

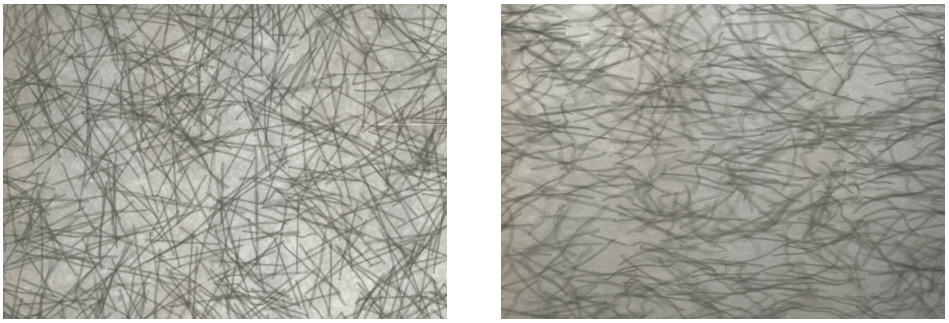

(a) 100% charge **(b)** 33% charge

Figure 1.3 X-ray images of two of the samples of sheet molding compound from Fig. 1.2. The flow direction is horizontal and the fibers are 25 mm long [Che84].

These variations in processing and testing produce the wide range of properties seen in Fig. 1.2. The strongest specimen has an ultimate tensile strength of 325 MPa, five times greater than the strength of the weakest specimen, 73 MPa. The elastic moduli show similar large variations.

These property variations follow a pattern. For each initial charge shape, the material is stronger and stiffer parallel to the flow, weaker and less stiff perpendicular to the flow. The smaller the charge, the greater the difference between the two directions.

These differences in mechanical properties arise because of differences in fiber orientation. Figure 1.3 shows X-ray images from two of the plaques. In these samples, some of the fibers were made from a glass with a high lead content and are opaque to X-rays, so they appear as dark fibers in the images. The fibers in the 100% charge sample are almost randomly oriented, while fibers in the 33% charge sample have a strong horizontal (flow-direction) orientation. Fibers reinforce the polymer matrix primarily along their axes. This makes the 33% charge sample much stronger and stiffer than the 100% charge sample in the flow direction, but weaker and less stiff perpendicular to flow. The other two charges have intermediate degrees of orientation and intermediate properties. Again, processing has a major effect on properties.

Discontinuous fiber composites present a striking contrast to materials like steel and aluminum, whose elastic moduli can be looked up in any materials handbook and are the same in any direction. The properties of a discontinuous fiber composite depend on the amount of flow, on the loading direction, and may vary from one location to another within a part.

How can one design with these materials? If we use only the minimum properties that can occur in a molded part, the design will be overly conservative and the part will be heavy and expensive. To use DFCs effectively, a designer must have good information about the actual properties at each location in the component. This can depend on processing conditions and part geometry, as well as on the intrinsic properties of the fiber and the matrix.

In these examples the connection between properties and processing is fiber orientation. If we can predict how the fibers are oriented in a molded part and relate that orientation to the mechanical properties at each location, we can determine how processing will affect the properties. We can then design components that use discontinuous fiber composites effectively. The required calculations are complex, but computer-aided engineering software is available to do most of the work.

1.3 Computer-Aided Engineering Workflow

The most widely used computer-aided engineering tools for DFCs are for injection molding. Many of these software packages can also simulate compression molding and can treat both thermoplastic and thermosetting matrices.

With these software tools, each analysis starts from the final part geometry, represented by a CAD solid model. This determines the shape of the mold cavity. An initial choice for the material is specified, and the analyst adds some manufacturing-related information, such as runners and gates for an injection mold, or the charge shape and location for compression molding. The analysis then proceeds in three major steps:

1. A **mold filling simulation** calculates the velocity, pressure and temperature of the polymer/fiber mixture as a function of time during filling, at all locations in the mold. Fiber orientation predictions are integrated into this calculation, and also proceed step by step as the mold fills. The fiber length distribution may be calculated, and the simulation may extend into the packing and cooling phases of the process. Output from this step includes the fiber orientation (and possibly fiber length distribution) at each location in the final part.
2. **Mechanical property predictions** use the fiber and matrix properties and the microstructural data from the first step to find direction-dependent properties (stiffness, thermal expansion, etc.) at each location in the final part.
3. A **structural analysis** is performed on the solid part using the finite element method. This calculation incorporates the mechanical properties from the previous step on an element-by-element basis. Multiple load cases may be analyzed to determine dimensional stability of the part (warping and shrinkage), deflections under specified loads, local stresses and loads at failure.

Figures 1.4 and 1.5 illustrate these steps for a water pump housing. The initial CAD model is shown in Fig. 1.4(a). In this figure, gates have been added, and are indicated by the yellow cones. A finite element mesh is then generated, dividing this geometry into small, three-dimensional elements. A portion of this mesh is shown in Fig. 1.4(b). The mold filling analysis can then proceed. Figure 1.4(c) shows a portion of the filling pattern from this analysis. At this instant in time the colored portion of the mold has been filled, and the colors indicate where the flow front was at previous times.

A portion of the predicted fiber orientation pattern is shown in Fig. 1.5(a). The colored lines show the main directions of orientation, with the red and yellow lines indicating highly aligned fibers, and blue or green lines indicating more random orientations. A cutting plane has been used to provide a view across the part thickness, and orientation does vary across the thickness.

This orientation information is converted into an elastic modulus value for each element in the mesh, as shown in Fig. 1.5(c). Each element has direction-dependent properties, and the figure shows the tensile modulus in the main fiber direction from Fig. 1.5(a). This modulus information is used, together with an orientation-dependent thermal expansion for each element, to calculate the size and shape of the part after it has cooled to room temperature. Deflections from this calculation are shown in Fig. 1.5(c), with the deflections amplified 10× to make them more visible. The base of the part is slightly curved, and the large hole in the center is no longer perfectly round.

All of these analyses may be done in a single integrated software package, as in the example above, or separate packages may be used with information flowing between them.

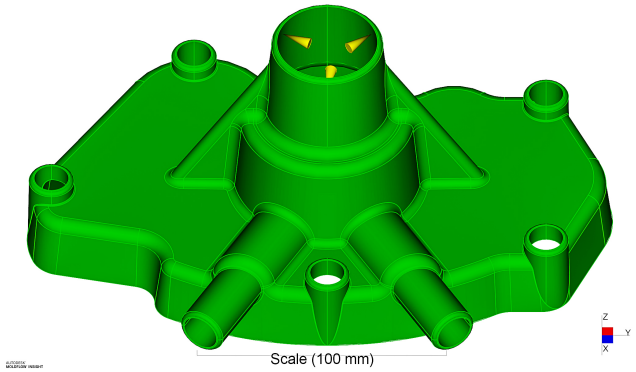

(a) CAD model

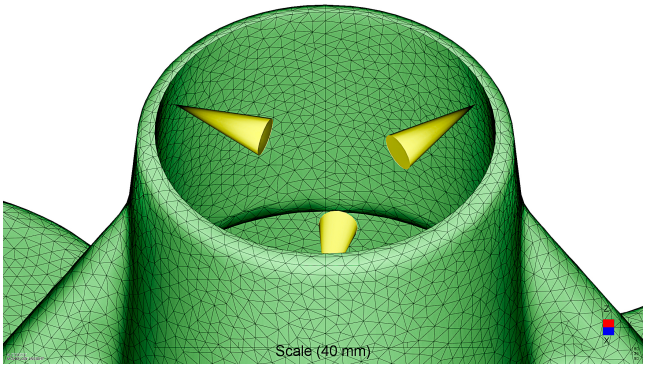

(b) Finite element mesh near the gates

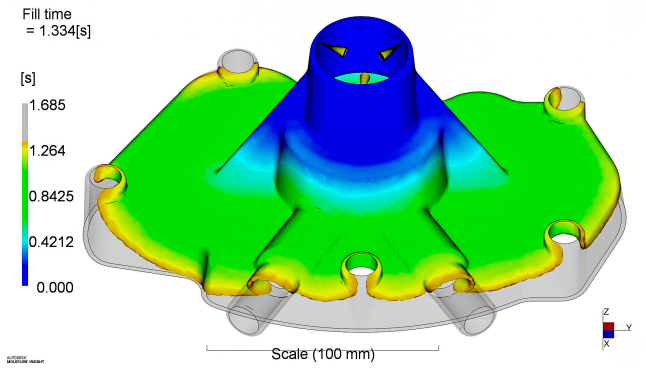

(c) Mold filling pattern

Figure 1.4 Initial steps in the CAE workflow for an injection molded, fiber reinforced part. Images provided by Autodesk, Inc.

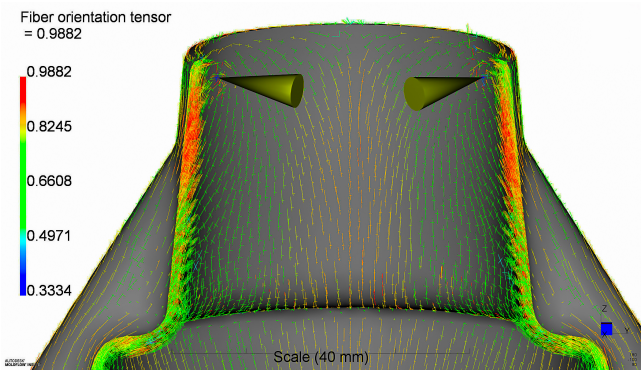

(a) Fiber orientation near the gates (cutaway view)

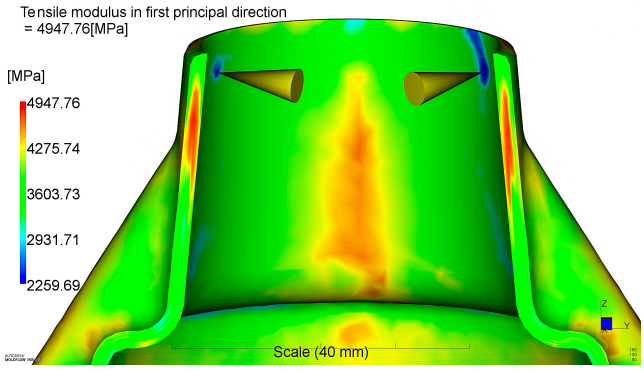

(b) Tensile modulus in the main orientation direction

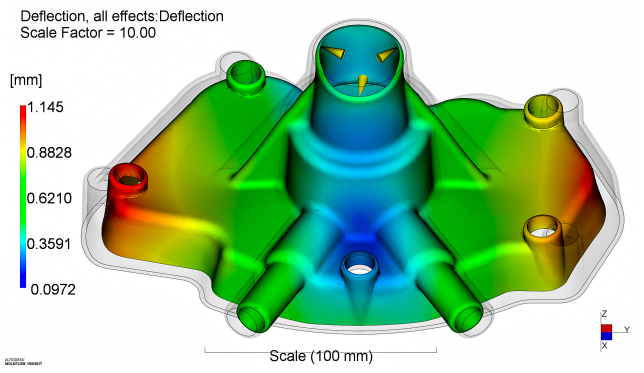

(c) Deflection after cooling

Figure 1.5 Additional steps in the CAE workflow for an injection molded, fiber reinforced part. Images provided by Autodesk, Inc.

At any step, detailed results are viewed and interpreted. If the results do not meet design requirements, one can revise the processing conditions, material choice or part design, then start the analysis again. This cycle of analysis, interpretation and revision continues until a suitable combination of part design, material and processing conditions has been found.

1.4 Scope and Organization of the Book

This book is written for users of CAE software who want to know how fiber orientation calculations work and how to use their software effectively. The primary focus is on calculation of fiber orientation and fiber length within a mold filling simulation. We also discuss the simplest mechanical property calculations—linear elastic stiffness and thermal expansion—to show how fiber orientation is connected to properties. The goal is to cover all the models that are used in commercial software, explain their underlying concepts, and show when to use each option and how to set the parameters for each model.

Another key audience is graduate students who are beginning research in this field. For these readers the book provides an organized presentation of the state of the art as of this writing, with selected references to recent research.

Flow and heat transfer models for mold filling are not discussed, as that information is available elsewhere [Ken95, ZTF11]. Similarly, we do not discuss finite element structural analysis.

While it is easy to see that fiber orientation affects properties, predicting those properties quantitatively takes some effort. An essential step is to choose the right variables to describe the microstructure. The best microstructural variables meet the following requirements:

- They have a precise mathematical definition.
- They can be measured experimentally.
- The ways they change during processing can be modeled using physical principles.
- Their effect on mechanical properties can be quantified, also using physical principles.

All of these requirements are met by the fiber orientation and length variables used in CAE software, and this book is designed around them:

- Chapter 2 defines the variables used to describe fiber orientation and fiber length. Chief among these is the second-order orientation tensor. This is the principal tool used for fiber orientation, and without it practical CAE tools for discontinuous fiber composites would not exist. Fiber length distributions and fiber length averages are also presented.
- Chapter 3 discusses how fiber orientation and fiber length are measured. For orientation we cover both traditional measurements on planar sections and modern micro-CT measurements. Methods for fiber length include burning off the matrix and micro-CT.
- Chapter 4 discusses the orientation of single fibers and Jeffery's equation. It also reviews basic types of deformation, and shows how two simple rules can qualitatively explain the orientation patterns in almost all situations.

- Chapter 5 explains how the orientation statistics for groups of fibers are predicted, and examines the models that are used in CAE software to make quantitative orientation predictions.
- Fibers can also alter the flow, and ways to model this are discussed in Chapt. 6. This requires modeling the rheology of fiber suspensions.
- Fiber length modeling is presented in Chapt. 7. These models are especially useful for LFT materials, to help preserve as much fiber length as possible in the final part.
- The connection between fiber orientation and mechanical properties is discussed in Chapt. 8. The focus here is on linear elastic behavior and thermal expansion. Property models for nonlinear stress-strain behavior, damage and failure also exist, but those topics are outside the scope of the book.
- Chapter 9 provides an overview of two advanced models from recent research. One, shear-induced fiber migration, causes the fiber volume fraction to vary from point to point. Models of this type are now appearing in commercial software. The other, direct simulation of multiple, interacting fibers, has the potential to affect how fiber orientation is modeled in the future.

Readers are assumed to be familiar with fluid mechanics and material behavior at the level of a bachelor's degree in engineering, as well as calculus of multiple variables, differential equations and matrix algebra.

Tensors are the natural mathematical language for fiber orientation modeling and are used extensively here. Readers who are less familiar with tensors will find the discussion of orientation tensors and the examples in Section 2.3 helpful. Appendix A summarizes the types of tensor notation used in this book and shows how to perform mathematical operations involving tensors. In practice, all of these calculations are done numerically by a computer, and Appendix A shows how vector and tensor operations are carried out in MATLAB.

The cited references are either foundational papers or representative of recent research. No attempt has been made to comprehensively review the literature, and there are many excellent papers on fiber orientation and related topics that are not cited here.

Key ideas and equations appear in shaded boxes, and each chapter ends with a summary.

1.5 Accompanying Software

For readers who want to explore fiber orientation and mechanical property models in more detail, a set of MATLAB functions is available at *https://github.com/charlestucker3/Fiber-Orientation-Tools*. Most of the calculations in this book were performed using this toolkit, and its capabilities are mentioned throughout the text. Documentation at that website gives more information. Fiber orientation models in Python are provided by Meyer [Mey21].

2 Describing Fiber Orientation and Fiber Length

The first step in dealing with fiber orientation is to describe it mathematically. We begin with single fibers, then discuss statistical distributions of groups of fibers. For this, there are two main options: probability density functions (also called fiber orientation distribution functions) and orientation tensors. The tensor description is a key tool both for predicting orientation and for presenting experimental data, and it plays a central role in the chapters that follow.

Fiber length plays a role in properties as well. Here we discuss fiber length distributions and different average values of fiber length.

2.1 Single Fiber Orientation

We primarily consider fibers that are straight and rigid (i.e., they do not bend by any noticeable amount). Both glass and carbon fibers have circular cross-sections, so each fiber is a circular cylinder. The orientation of any fiber is then described by the orientation of its symmetry axis. The angles θ and ϕ, shown in Fig. 2.1, provide one way to do this.

An alternate description, which we will use extensively, is a unit vector $\mathbf{p}$ that points along the fiber axis. This vector is also shown in Fig. 2.1. The components of this vector are p_1, p_2 and p_3, and they are related to the angles θ and ϕ according to

$$\begin{Bmatrix} p_1 \\ p_2 \\ p_3 \end{Bmatrix} = \begin{Bmatrix} \sin\theta\cos\phi \\ \sin\theta\sin\phi \\ \cos\theta \end{Bmatrix} \tag{2.1}$$

How can $\mathbf{p}$ have three components when only two angles are needed to describe the orientation of a fiber? $\mathbf{p}$ also has unit length, which requires that

$$\mathbf{p}\cdot\mathbf{p} = 1 \qquad \text{or} \qquad p_1^2 + p_2^2 + p_3^2 = 1 \tag{2.2}$$

Thus, only two components of $\mathbf{p}$ are independent.

While the vector $\mathbf{p}$ has a head and a tail, a fiber does not. As Fig. 2.2 shows, for any fiber there are two equally valid ways to choose $\mathbf{p}$. If the unit vector $\mathbf{p}_1$ describes a fiber, then so does the vector $\mathbf{p}_2 = -\mathbf{p}_1$. Similarly, if a fiber has angles θ_1 and ϕ_1, we could also describe it as having angles (θ_2, ϕ_2) where $\theta_2 = \pi - \theta_1$ and $\phi_2 = \phi_1 + \pi$. All equations that deal with fiber orientation must respect these relationships.

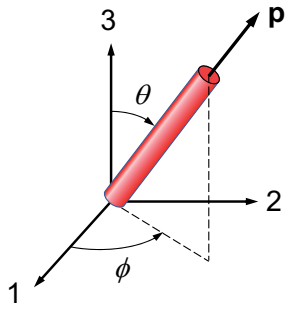

Figure 2.1 The angles θ and ϕ and the unit vector $\mathbf{p}$ that define the orientation of a single, rigid fiber with respect to Cartesian coordinate axes.

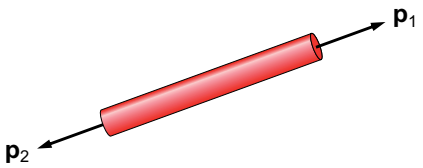

Figure 2.2 The orientation of a single fiber is described by either $\mathbf{p}_1$ or $\mathbf{p}_2$, where $\mathbf{p}_2 = -\mathbf{p}_1$.

2.2 Distributions of Fiber Orientation

An injection molded short fiber composite can easily have 10,000 fibers per mm^3, so for realistic composites we must consider many fibers. Describing the orientation of each individual fiber would be a huge task, and is usually unnecessary. Instead, one describes the statistics of the fiber orientation. This is most easily understood when all fibers lie in a plane, so we will treat that case first, then generalize to three-dimensional fiber orientation.

2.2.1 Planar Orientation

Consider a composite where all fibers lie in the 1–2 plane. This is a good approximation for compression molded sheet molding compound (SMC), where the fibers are long compared to the final part thickness. For planar orientation all fibers have $\theta = \pi/2$, and each fiber's orientation is described by a single angle ϕ. Similarly, the orientation vector $\mathbf{p}$ always has $p_3 = 0$, so p_1 and p_2 describe an individual fiber and $p_1^2 + p_2^2 = 1$.

Figure 1.3 showed two X-ray images of SMC samples. The orientations appear nearly random in the sample on the left, and have a strong horizontal alignment in the sample on the right. (The fibers are also not straight, an issue we will ignore for now and return to in Chapt. 9.)

One can measure the angle ϕ for many fibers and collect the data in a histogram, like those in Fig. 2.3. Note that the data has been restricted to $-90° < \phi \leq 90°$. This is a typical way to treat experimental data for planar orientation.

In Fig. 2.3 the histogram on the left indicates a mostly random orientation, while the histogram on the right shows a strong tendency for fibers to have ϕ close to zero. This corresponds to our

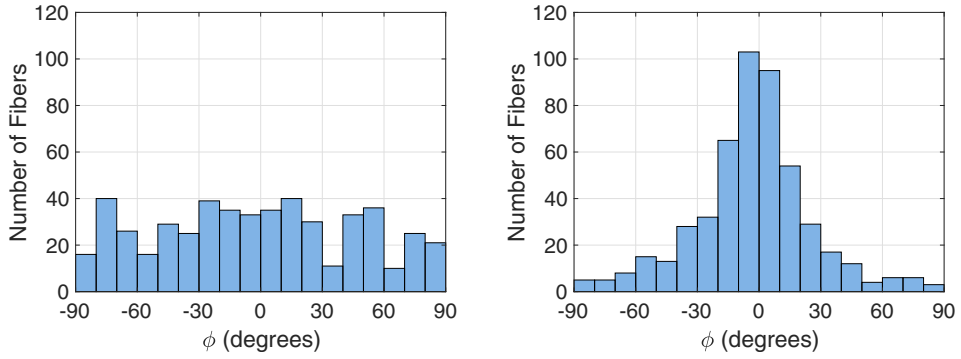

(a) 100% charge **(b)** 33% charge

Figure 2.3 Histograms of orientation angle ϕ for 500 fibers measured from each of the samples in Fig. 1.3 [CT84].

visual impressions of the X-ray images. Looking closely, the histogram on the left does show a slight preference for the 0° direction compared to the 90° direction, a feature that is not easily detected by looking at the X-ray image.

A histogram is a discrete approximation to the underlying probability density function for orientation. For planar orientation states we will use $\psi_\phi(\phi)$ to denote this function, which is also called the *orientation distribution function.* This function is defined such that the probability of any fiber having a ϕ value between some angle ϕ^* and $\phi^* + \mathrm{d}\phi$ is

$$P(\phi^* \le \phi < \phi^* + \mathrm{d}\phi) = \psi_\phi(\phi^*)\,\mathrm{d}\phi \tag{2.3}$$

The function ψ_ϕ must satisfy some requirements for any orientation state. First, every fiber must have some orientation, so the integral of ψ_ϕ over all values of ϕ must equal unity:

$$\int_0^{2\pi} \psi_\phi(\phi)\,\mathrm{d}\phi = 1 \tag{2.4}$$

We will refer to this as the *normalization* requirement. Also, because ϕ is equivalent to $\phi + \pi$, the orientation distribution function must be *periodic* with period π:

$$\psi_\phi(\phi + \pi) = \psi_\phi(\phi) \tag{2.5}$$

If the fibers are randomly oriented in the plane, then all values of ϕ are equally likely and $\psi_\phi(\phi)$ is constant. The normalization requirement (2.4) then tells us that

$$\psi_\phi(\phi) = \frac{1}{2\pi} \qquad \text{random-in-plane orientation} \tag{2.6}$$

2.2.1.1 Matching ψ_ϕ to Data

In later sections we will write equations for $\psi_\phi(\phi)$ and develop models to predict it, so we will want to be able to compare those predictions to experimental data like that in Fig. 2.3. This is done by scaling the histogram data on the vertical axis. Let ϕ_i, $i = 1, \ldots, n$, be the angles at the centers of the histogram bins, where n is the number of bins. Also assume that the bins span a

total angle of π. Let $\Delta\phi = \phi_{i+1} - \phi_i$ be the bin width, which is the same for all bins, and let N_i be the number of fibers assigned to the bin with ϕ_i. We know that the distribution function at the center of bin i should be proportional to N_i,

$$\psi_\phi(\phi_i) = kN_i \tag{2.7}$$

where k is some constant of proportionality to be determined. We want the distribution function to satisfy the normalization condition, Eqn. (2.4), such that

$$\int_0^{2\pi} \psi_\phi(\phi)\ \mathrm{d}\phi \approx 2\sum_{i=1}^{n} \psi_\phi(\phi_i)\,\Delta\phi = 1 \tag{2.8}$$

The factor of two appears in the second expression because the ϕ_i values span an angular range of π, but the normalization requirement is written over the full range of 2π. Combining Eqns. (2.7) and (2.8) gives the value of k and allows us to write the correspondence between histogram data and the orientation distribution function:

$$\psi_\phi(\phi_i) = \frac{N_i}{2\Delta\phi \sum_{j=1}^{n} N_j} \tag{2.9}$$

While we often plot data with ϕ in degrees, the value of $\Delta\phi$ in this equation must be in radians.

2.2.2 Three-Dimensional Orientation

For three-dimensional orientation the set of all possible orientations (θ,ϕ) or $\mathbf{p}$ is the surface of a unit sphere, as shown in Fig. 2.4. We will call this surface the *orientation space*.

The orientation distribution function for 3-D orientation, written as $\psi(\theta,\phi)$ or $\psi(\mathbf{p})$, is defined such that the probability of a fiber lying between θ^* and $\theta^* + \mathrm{d}\theta$ and between ϕ^* and $\phi^* + \mathrm{d}\phi$ is

$$P(\theta^* \le \theta < \theta^* + \mathrm{d}\theta, \phi^* \le \phi < \phi^* + \mathrm{d}\phi) = \psi(\theta^*,\phi^*)\sin\theta^*\,\mathrm{d}\theta\,\mathrm{d}\phi \tag{2.10}$$

The factor of $\sin\theta^*$ appears, as in any problem using spherical coordinates, because an increment $\mathrm{d}\theta\,\mathrm{d}\phi$ covers different amounts of area at different values of θ^*.

The integral of any function $f(\theta,\phi)$, or equivalently $f(\mathbf{p})$, over the orientation space can be written in either of two equivalent ways:

$$\int_{\theta=0}^{\pi}\int_{\phi=0}^{2\pi} f(\theta,\phi)\ \sin\theta\ \mathrm{d}\theta\ \mathrm{d}\phi = \oint f(\mathbf{p})\,\mathrm{d}\mathbf{p} \tag{2.11}$$

In the latter expression, $\mathrm{d}\mathbf{p}$ is a small solid angle corresponding to the small area on the surface of the sphere $\sin(\theta)\,\mathrm{d}\theta\,\mathrm{d}\phi$, as shown in Fig. 2.4. The $\mathbf{p}$ notation is compact and will be useful in describing concepts, while the (θ,ϕ) notation is useful for deriving results algebraically or for numerical calculations.

The normalization requirement for 3-D orientation is

$$\oint \psi(\mathbf{p})\,\mathrm{d}\mathbf{p} = 1 \qquad \text{or} \qquad \int_{\theta=0}^{\pi}\int_{\phi=0}^{2\pi} \psi(\theta,\phi)\ \sin\theta\ \mathrm{d}\theta\ \mathrm{d}\phi = 1 \tag{2.12}$$

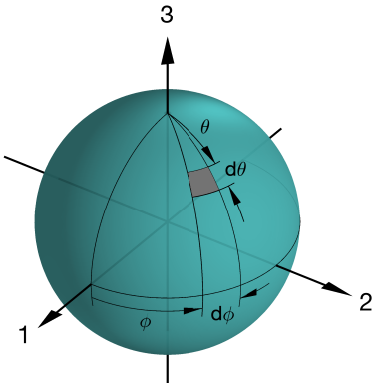

Figure 2.4 The unit-sphere orientation space for three-dimensional orientation. The gray shaded region has area $\sin(\theta)\,d\theta\,d\phi = d\mathbf{p}$.

If the fibers are oriented randomly in space then all directions $\mathbf{p}$ are equally likely and $\psi(\mathbf{p})$ is constant. The surface area of a sphere with unit radius is 4π, so the normalization requirement (2.12) tells us that

$$\psi(\mathbf{p}) = \frac{1}{4\pi} \qquad \text{3-D random orientation} \tag{2.13}$$

We saw in Section 2.1 that any fiber can be represented by two different $\mathbf{p}$ vectors or by two different combinations of (θ,ϕ). This produces the *periodicity* requirement for 3-D orientation, which is

$$\psi(\mathbf{p}) = \psi(-\mathbf{p}) \qquad \text{or} \qquad \psi(\theta,\phi) = \psi(\pi-\theta,\phi+\pi) \tag{2.14}$$

Any point on the unit sphere must have the same value of ψ as the polar opposite point. If we know $\psi(\mathbf{p})$ on a hemisphere of the orientation space then we have complete information about ψ.

In principle, $\psi(\mathbf{p})$ can take on any form that satisfies the normalization and periodicity conditions. In practice, the distribution functions for discontinuous fiber composites vary smoothly with $\mathbf{p}$ and are reasonably well approximated by certain mathematical forms[1]. Figure 2.5 shows the orientation distribution functions for three examples:

(a) The fibers are strongly oriented along the 1 axis, and the orientation is symmetric about that axis. This is the type of alignment that occurs in a converging nozzle, where the flow is in the 1 direction.

(b) Many fibers lie close to the 1–2 plane, but the orientation is not perfectly planar. The orientation is symmetric about the 3 axis, with no preferred direction in the 1–2 plane. A biaxial squeezing flow, as seen in some compression moldings, would create this type of alignment by squeezing along the 3 axis.

(c) There is strong alignment in the flow (1) direction, with some spread in the crossflow (2) direction but less spread in the thickness (3) direction. The direction of greatest orienta-

[1] These forms are discussed in Section 2.4.

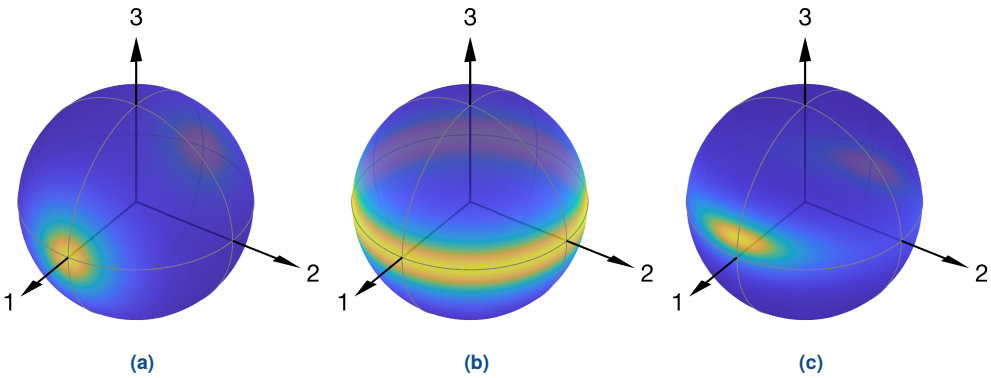

Figure 2.5 Examples of 3-D orientation distribution functions. Lowest and highest values of ψ are indicated by dark blue and yellow, respectively. (a) Axisymmetric alignment in the 1 direction. (b) Nearly planar alignment, uniformly distributed in the 1–2 plane. (c) Orientation similar to the shell layer in an injection molded composite. The corresponding orientation tensors appear in Eqns. (2.40–2.42).

tion is also tilted out of the 1–2 plane. This orientation state is similar to the shell layer in an injection molded composite. It is created by simple shear flow in the 1 direction with velocity changing in the 3 direction.

In all three examples in Fig. 2.5 the pattern on the front of the sphere is mirrored on the back of the sphere, as required by the periodicity condition.

2.3 Orientation Tensors

2.3.1 Introduction

The distribution function $\psi(\mathbf{p})$ contains a great deal of information, and even if we can calculate or measure the full function, interpreting it is not a simple task. It would be very helpful to have some kind of average, perhaps analogous to the mean or the standard deviation for scalar data, to help describe and interpret orientation distributions.

The average value of $\mathbf{p}$ is not helpful in this regard. The sign of $\mathbf{p}$ is arbitrary and, because of periodicity, each $\mathbf{p}$ should be cancelled by a corresponding $-\mathbf{p}$, so the average value of $\mathbf{p}$ is always equal to zero.

We could deal with the sign of $\mathbf{p}$ by averaging $\mathbf{p}$ squared, but there are several ways to multiply a vector by itself. The dot product is not helpful because $\mathbf{p}\cdot\mathbf{p} = 1$ for every $\mathbf{p}$, so an average of dot products will aways equal unity. The cross product is also not helpful, since $\mathbf{p}\times\mathbf{p} = 0$ for any $\mathbf{p}$, so an average of the cross products will aways equal zero.

A way to multiply $\mathbf{p}$ with itself that *is* useful here is the *dyadic product*. We will write this product as $\mathbf{pp}$, without an intervening symbol. This way of multiplying two vectors produces a tensor whose ij component is the product $p_i p_j$, for all combinations of i and j. In matrix