Universidad de los Andes | School of Engineering and Applied Sciences

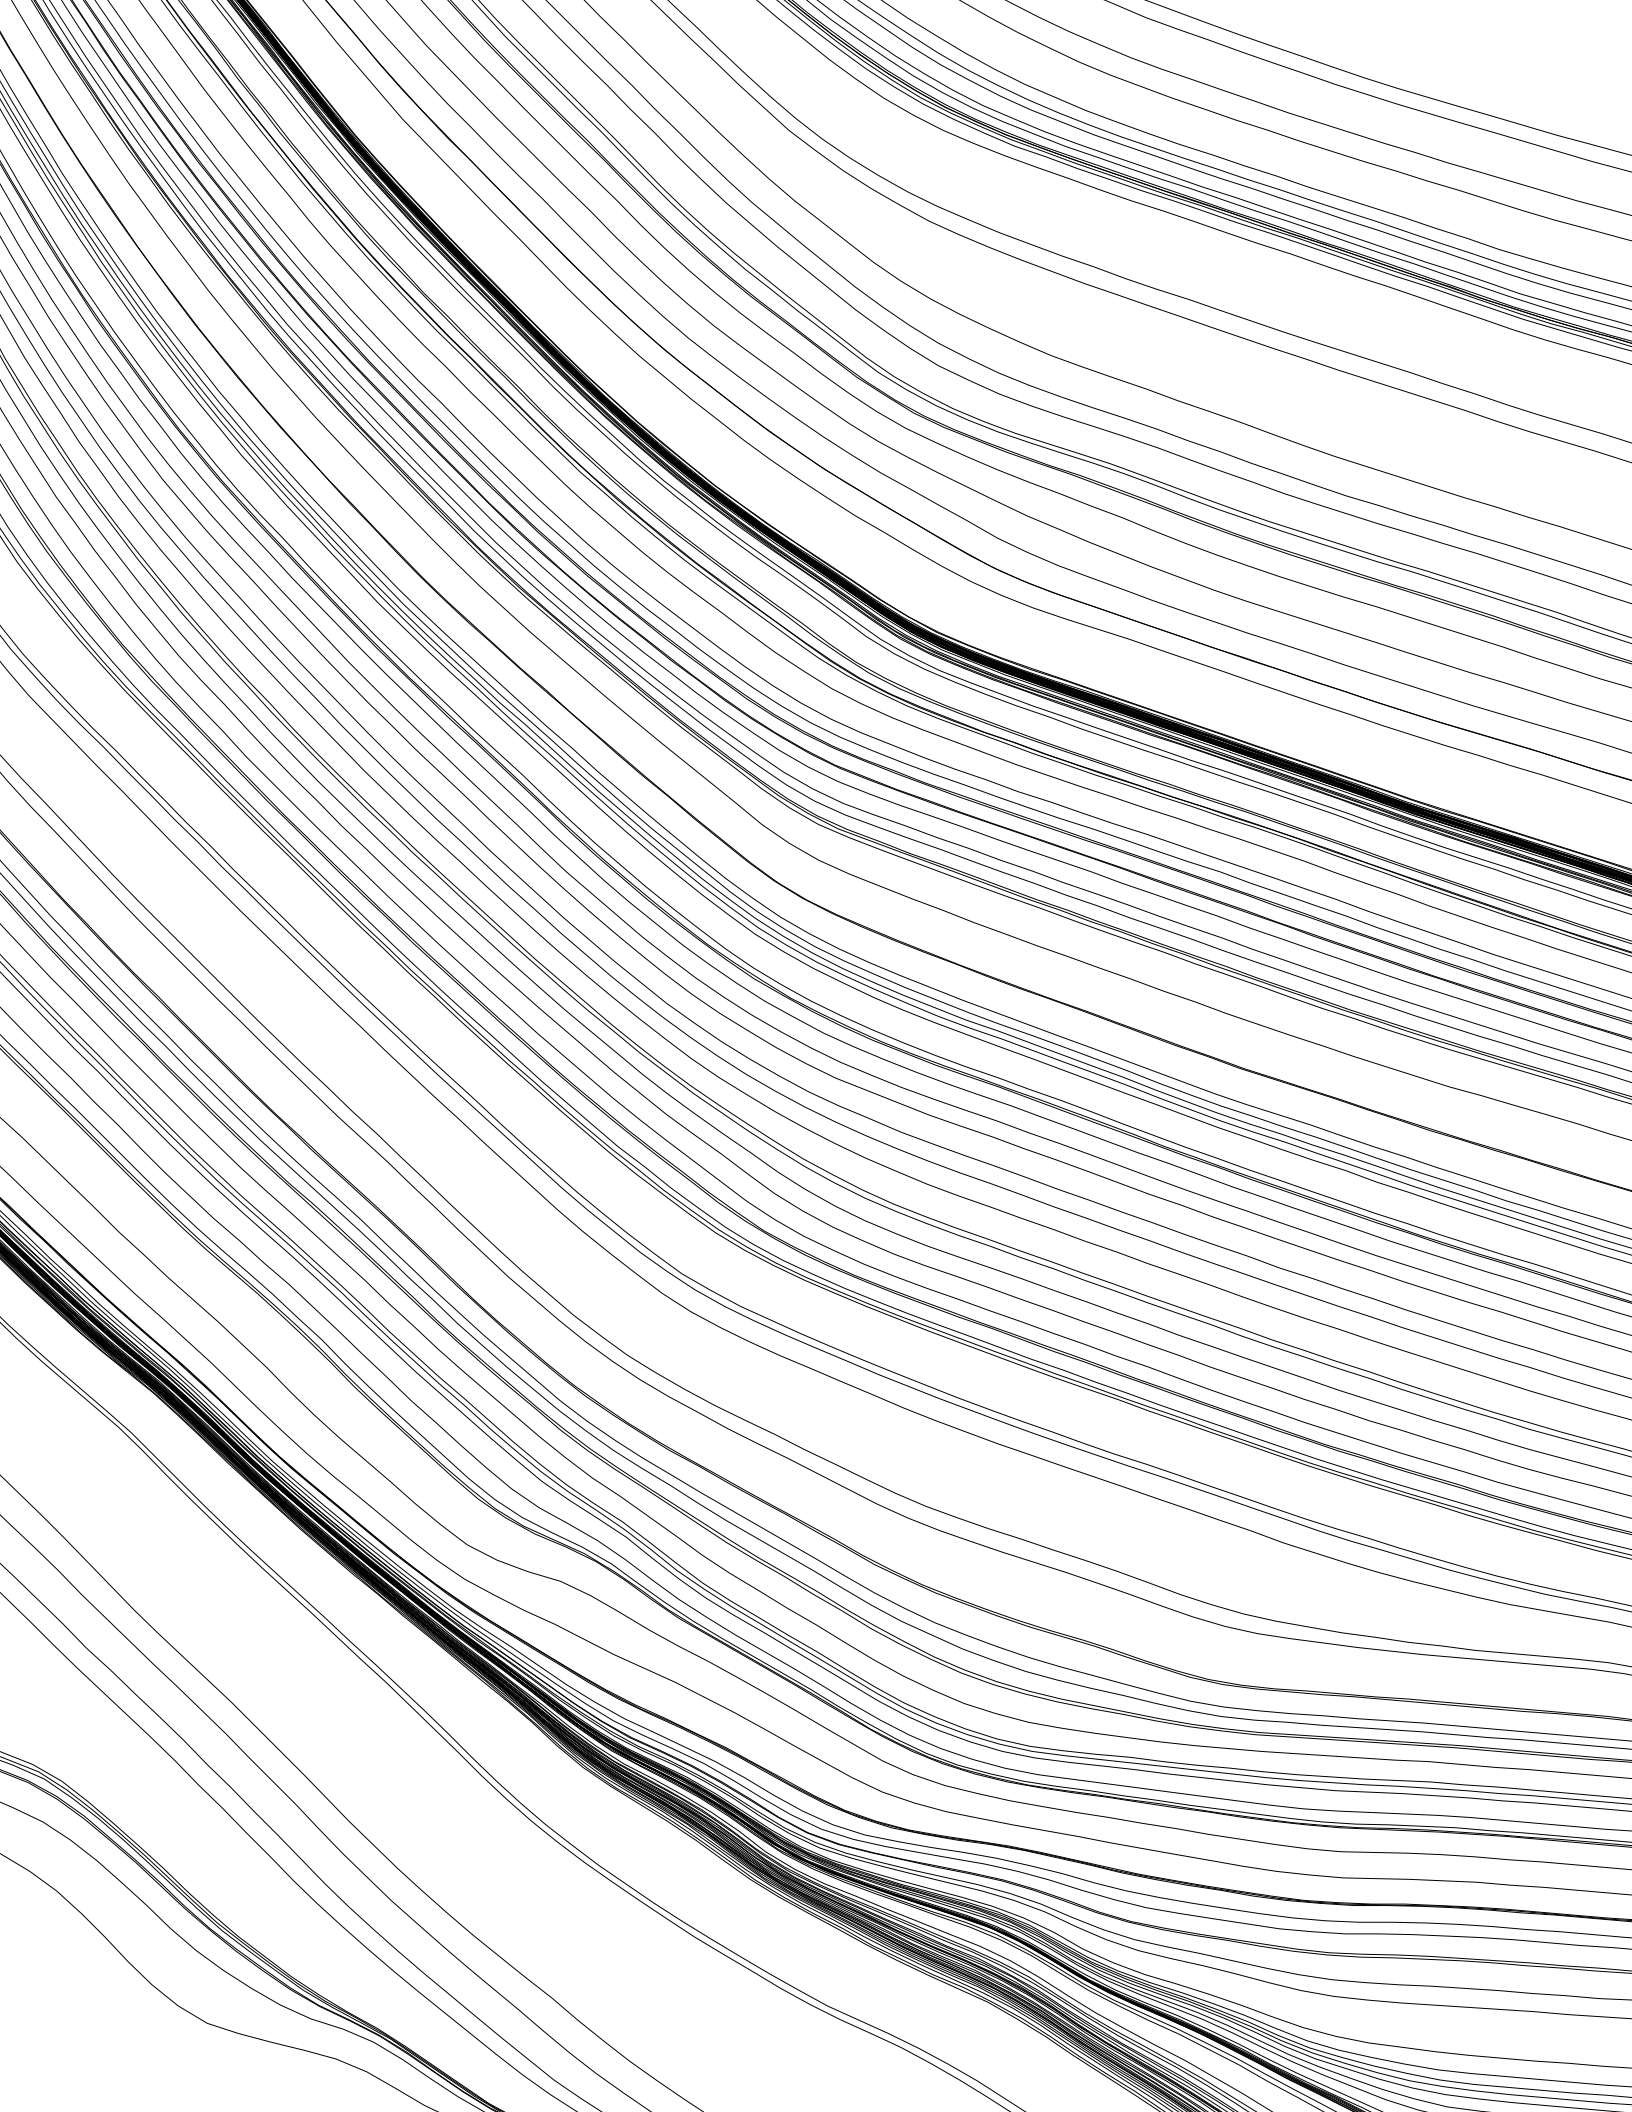
A Predictive Model for Fiber Orientation Dynamics in Viscoplastic Flows from Experimental PIV and PTV on Carbopol as UHPC Analogue

Lukas Wolff Casanova

A Thesis for the degree of
Master of Science in Engineering

Advisor: Patricio Moreno Casas
Co-Advisor: Álvaro Paul Quiroz
Co-Advisor: Jorge Gómez Mir

Santiago, Chile, August 2026



Design is not just what it looks like and feels like.
Design is how it works.

Steve Jobs

The grand aim of all science is to cover the greatest number of empirical facts by logical deduction from the smallest number of hypotheses or axioms.

Albert Einstein

All models are wrong, but some are useful.

George Box

I certify that I have read this thesis and that in my opinion its scope and quality are completely to be considered a thesis leading to the Master of Science in Engineering Degree.

Patricio Moreno, Ph.D.

Álvaro Paul, Ph.D.

Jorge Gómez, Ph.D.

José Abell, Ph.D.

Fabián Bombardelli, Ph.D.

Jaime Cisternas, Ph.D.

Agradecimientos

Desocupado investigador:

Tú y yo sabemos que investigar es sumergirse en un mundo plagado de gigantes. Habrá quien nos llame locos, sea por las ideas dementes o por la cantidad de libros leídos. Pero es ahí, en la manía de develar lo que ignoramos, donde está nuestra cordura: con una explicación suficientemente simple, muchos gigantes terminan siendo molinos de viento.

Citamos a los grandes, a quienes caminaron estas tierras mucho antes que nosotros, con la esperanza secreta de escribir alguna vez una frase digna de ser citada. Pero mis referentes no están en los libros: estuvieron siempre a mi lado. A Patricio Moreno, Álvaro Paul y Jorge Gómez les debo la entrada a esta aventura, la confianza en mis habilidades y en mi locura, y el consejo dado siempre con paciencia.

No hay aventura sin un lugar al que volver. La mayoría de las veces volvemos derrotados y con las manos vacías, preguntándonos qué podríamos haber hecho distinto para descubrir la verdadera identidad de algún gigante. A mi familia y a los distintos hogares que me recibieron en este camino, en especial a mi mamá, mis hermanos Mathias y Pablo Wolff, Norman Zambrano, Verónica Mericq, Betina Wolff y Walter Riegel, les agradezco haberme devuelto al camino cada vez con un nuevo consejo.

Y finalmente, todos tenemos a ese amigo que cabalga a nuestro lado, que nunca duda de nuestra locura. Vicente Cueva, Bernardo Caprile, Catalina Vicencio, Felipe Vicencio, Pedro Valenzuela, Sebastián Sepúlveda, Isidora Linacre, Cristóbal Maggi y Mauricio Correa: ustedes no solo me acompañaron en la aventura, también son los mejores recuerdos que me llevo.

Y a ti, desocupado investigador, te invito a disfrutar esta aventura, en la que logré develar la identidad de un pequeño gran gigante...

Acknowledgements

Idle researcher:

You and I both know that to do research is to plunge into a world teeming with giants. There will be those who call us mad, whether for our deranged ideas or for the sheer number of books we have read. But it is precisely there, in that mania for uncovering what we do not know, that our sanity lies: given a simple enough explanation, many giants turn out to be nothing but windmills.

We cite the greats, those who walked these lands long before us, harbouring the secret hope of one day writing a line worth citing. But my true references were never in the books: they were always at my side. To Patricio Moreno, Álvaro Paul and Jorge Gómez I owe my entry into this adventure, their faith in my abilities and in my madness, and advice always given with patience.

There is no adventure without a place to return to. More often than not we come back beaten and empty-handed, wondering what we might have done differently to uncover the true identity of some giant. To my family and to the many homes that took me in along the way, especially my mother, my brothers Mathias and Pablo Wolff, Norman Zambrano, Verónica Mericq, Betina Wolff and Walter Riegel, thank you for setting me back on the road each time with new counsel.

And finally, we all have that friend who rides alongside us and never doubts our madness. Vicente Cueva, Bernardo Caprile, Catalina Vicencio, Felipe Vicencio, Pedro Valenzuela, Sebastián Sepúlveda, Isidora Linacre, Cristóbal Maggi and Mauricio Correa: you did not merely keep me company on the adventure, you are the best memories I take from it.

And to you, idle researcher, I invite you to enjoy this adventure, in which I managed to uncover the identity of one small, great giant...

Resumen

Las estructuras de hormigón armado modernas están sometidas a exigencias de mayor vida de servicio, desempeño y rendimiento. Adicionalmente, la alta y creciente demanda de hormigón ha impulsado a la industria de la construcción a buscar materiales cementicios avanzados. El Hormigón de Ultra Alto Desempeño (UHPC por sus siglas en inglés) es un material que, mediante la optimización del empaquetamiento de su matriz, el uso de aditivos hiperplastificantes y la adición de fibras metálicas, alcanza no solo un gran rendimiento mecánico, sobre 150 MPa en compresión y del orden de 15 MPa en tracción, sino también una vida útil que supera los 200 años. La literatura señala que las fibras cumplen un rol fundamental, donde la alineación y dispersión de las fibras respecto de la dirección de carga son factores fundamentales que gobiernan la respuesta mecánica. Ambos parámetros derivan del proceso de colocación del hormigón, y dependen directamente de la reología de la mezcla y del contenido de fibras. Velocimetría de imágenes de partículas (PIV) fue implementada sobre Carbopol, un fluido transparente no newtoniano viscoplástico empleado para replicar las propiedades reológicas del UHPC, con el fin de estudiar el comportamiento euleriano del flujo. Adicionalmente, la metodología, velocimetría de seguimiento de partículas (PTV) fue implementada para estudiar el comportamiento lagrangiano de las fibras. La campaña experimental empleo Carbopol en concentraciones de 0,2% y 0,5%, y cuatro contenidos de fibra: un control sin fibras y dosificaciones equivalentes a 0,72, 1,48 y 2,92 veces el límite diluido. Las mediciones de PIV y PTV determinaron que los factores dominantes son la geometría de la sección y la reología de la mezcla. Sin embargo, se observó que a mayor concentración de fibras, éstas experimentan una mayor velocidad angular mientras que la vorticidad del flujo permanece prácticamente inalterada. Además, el alineamiento final de las fibras disminuye con la concentración. Todos estos efectos son predichos por el modelo de Advani-Tucker para flujos newtonianos. De esta manera, esta tesis propone una reducción del modelo de Advani-Tucker a una ley exponencial de aproximación al equilibrio con dos parámetros, el alineamiento alcanzable $\eta_{1,\max}$ y la deformación característica Γ^*, aplicable a flujos no newtonianos viscoplásticos como método de predicción de la orientación de las fibras a partir de la deformación acumulada inducida por el flujo. El modelo se calibró con los datos provenientes de las mediciones de PTV en Carbopol 0,2%, y luego se puso a prueba alimentándolo con la deformación acumulada recuperada por PIV en el flujo sin fibras. La predicción resultó correcta dentro de una tolerancia de $\pm 5°$, de modo que se determina que el modelo de Advani-Tucker, o una forma reducida de este, permite predecir la alineación final de grupos de fibras según la deformación acumulada que se espera que estas experimenten.

Abstract

Modern reinforced concrete structures are subject to increasing demands in terms of service life, performance and efficiency. In addition, the high and growing demand for concrete has driven the construction industry to seek advanced cementitious materials. Ultra High Performance Concrete (UHPC) is a material that, through the optimization of the packing of its matrix, the use of superplasticizing admixtures and the addition of steel fibers, achieves not only a high mechanical performance, above 150 MPa in compression and on the order of 15 MPa in tension, but also a service life exceeding 200 years. The literature indicates that fibers play a fundamental role, where the alignment and dispersion of the fibers with respect to the loading direction are the governing factors of the mechanical response. Both parameters arise from the casting process and depend directly on the rheology of the mixture and on the fiber content. Particle Image Velocimetry (PIV) was implemented on Carbopol, a transparent non-Newtonian viscoplastic fluid used to replicate the rheological properties of UHPC, in order to study the Eulerian behavior of the flow. Additionally, the Particle Tracking Velocimetry (PTV) methodology was implemented to study the Lagrangian behavior of the fibers. The experimental campaign employed Carbopol at concentrations of 0.2% and 0.5%, and four fiber contents: a control without fibers and dosages equivalent to 0.72, 1.48 and 2.92 times the dilute limit. The PIV and PTV measurements determined that the dominant factors are the cross section geometry and the rheology of the mixture. However, it was observed that at higher fiber concentrations the fibers experience a higher angular velocity, while the vorticity of the flow remains practically unaltered. Furthermore, the final alignment of the fibers decreases with concentration. All of these effects are predicted by the Advani-Tucker model for Newtonian flows. On this basis, this thesis proposes a reduction of the Advani-Tucker model to an exponential approach-to-equilibrium law with two parameters, the attainable alignment $\eta_{1,\max}$ and the characteristic strain Γ^*, applicable to non-Newtonian viscoplastic flows as a method for predicting fiber orientation from the accumulated strain induced by the flow. The model was calibrated with the data obtained from the PTV measurements in 0.2% Carbopol, and was then tested by feeding it with the accumulated strain recovered by PIV in the flow without fibers. The prediction was correct within a tolerance of $\pm 5°$, so it is concluded that the Advani-Tucker model, or a reduced form of it, allows the final alignment of groups of fibers to be predicted from the accumulated strain they are expected to experience.

Contents

Agradecimientos vii
Acknowledgements ix
Resumen xi
Abstract xiii

1 Introduction **1**
1.1 Specific Objectives 8
1.2 Thesis Structure 9

2 Theoretical Background **11**
2.1 Ultra High Performance Concrete 12
2.1.1 Fiber Influence on Mechanical Performance 12
2.2 Rheology of UHPC and Carbopol Analog 14
2.2.1 Viscoplastic Constitutive Models 14
2.2.2 Rheological Properties of Carbopol 16
2.2.3 Concentration Regimes and the Semi-Dilute Limit 17
2.3 Non-Intrusive Experimental Techniques 19
2.3.1 Particle Image Velocimetry 20
2.3.2 Particle Tracking Velocimetry 24
2.3.3 Order and Dispersion Metrics 26
2.4 L-Shape Model 30
2.4.1 Optical Configuration of the Mold 30

3 Methodology **33**
3.1 Experimental Setup 34
3.2 Steel Fibers 36
3.3 Carbopol Mixing Protocol 37
3.3.1 Rheological Characterization 39
3.4 PIV Implementation 40
3.5 PTV Implementation 43
3.5.1 Fiber Detection and Segmentation 43
3.5.2 High Concentration Detection 46
3.5.3 Fiber Characterization and Correlation 47
3.5.4 PTV Algorithm Validation Tests 52

3.6 Experimental Campaign . . . 53
3.7 Analysis Metrics and Significance Criteria . . . 56
3.7.1 Measured Quantities . . . 56
3.7.2 Effect Size . . . 57
3.7.3 The Sampling Control . . . 59

4 Experimental Characterization of Flow and Fiber Kinematics 61
4.1 PTV Algorithm Validation . . . 62
4.1.1 Segmentation Performance . . . 62
4.1.2 Tracking Performance . . . 63
4.2 Domain Zones . . . 66
4.2.1 Flow Stages by Zone . . . 67
4.3 PIV and PTV Global Analysis . . . 69
4.3.1 PIV Measurement Error . . . 70
4.3.2 PIV Results . . . 70
4.3.3 PTV Tracking Efficiency at High Concentration . . . 73
4.3.4 PTV Results . . . 75
4.3.5 Discussion . . . 77
4.4 Fiber's Final State Analysis . . . 79
4.4.1 Hierarchy of Factors . . . 79
4.4.2 Isolated Effect of Fiber Content . . . 81
4.5 Lagrangian Reconstruction of the Fiber Paths . . . 83

5 Fiber Dynamics 87
5.1 Jeffery's Orbits . . . 88
5.2 Folgar-Tucker Diffusion Model . . . 90
5.3 Advani-Tucker Model . . . 91
5.4 The Scalar Order Parameter . . . 94
5.4.1 Reduction of the Orientation Tensor . . . 94
5.4.2 Equivalence with the Measured Order Parameter . . . 96
5.5 Equilibrium State . . . 98
5.6 Exponential Function of Approach to Equilibrium . . . 100

6 Prediction of the Final Orientation 103
6.1 Correspondence Between the Models and the Experiment . . . 104
6.2 Confrontation with the Measured Orientation . . . 105
6.2.1 The Angle Predictor . . . 108

6.3 Test Against the Fiber-Free Flow . . . 109

7 Conclusions **113**
7.1 The Model and What It Predicts . . . 114
7.2 The Effect of the Fiber Content . . . 115
7.3 Limits of the Measurement . . . 116
7.4 Future Work . . . 116

Nomenclature **119**

Appendix **135**

A Derivation of the Orientation Saturation Law **135**
A.1 Tensor Objects and Operations . . . 136
A.1.1 Identity With the Measured Order Parameter . . . 140
A.2 From the Single Fiber to the Population . . . 142
A.3 Projection Onto the Scalar Order Parameter . . . 145
A.4 The Accumulated Shear Strain as the Independent Variable . . . 149
A.5 The Relaxation Equation and its Explicit Solution . . . 150
A.5.1 The Predicted Magnitude of the Characteristic Strain . . . 153
A.6 Interpretation and Identifiability of the Parameters . . . 154
A.7 Planar Reduction and What it Costs . . . 157
A.8 Assumptions and Domain of Validity . . . 158

B PIV and PTV Block-Structure Sensitivity Analysis **161**
B.1 PIV . . . 162
B.2 PTV . . . 163
B.3 Computational Cost . . . 164

C Carbopol Mixtures **168**

D Set-Up Information **170**
D.1 Domain Dimensions . . . 170
D.2 Camera Positions . . . 170
D.3 Renders of the Experimental Setup . . . 171

E Zone-Resolved Statistical Analysis **174**
E.1 Protocol . . . 174
E.2 Eulerian Variables . . . 175

E.3 Lagrangian Variables 175
E.4 Reading the Distribution 177

F Final State of the Fiber Population **180**
F.1 Descriptors of the Deposited Population 180
F.1.1 The Orientation 181
F.1.2 The Spatial Arrangement 182
F.2 The Effect of the Rheology 183
F.3 The Effect of the Fiber Content on the Spatial Arrangement 184
F.4 Photographs of the Arrested Elements 185

G Design **188**

List of Figures

1.1 Comparison between UHPC and conventional concrete in terms of compressive strength, tensile strength, durability, CO_2 emissions, and material requirements. While CO_2 emissions increase per unit, as the required material is reduced and the service life is extended, the overall environmental impact of UHPC is lower than that of conventional concrete. . . 3

1.2 Expected fiber distribution and alignment versus anticipated failure mode of the element. 4

2.1 (a) Tensile stress-strain curves of UHPC specimens with different fiber orientations. (b) Effect of fiber dosage on the rheological properties of UHPC. 13

2.2 Relation between the fiber volume fraction ϕ and the number density n. (a) Test fiber of length L and diameter d_f centered at c, the sphere of diameter L swept by its rotation (shaded), and the centers of the neighboring fibers (dots) within a volume L^3. (b) Number of neighbors within L^3, $nL^3 = (4/\pi)\,\phi\, r_f^2$, against the fiber volume fraction for $r_f = 65$, with the dilute limit $\phi_{sd} = r_f^{-2}$ and the concentrated limit $\phi_c = r_f^{-1}$. N_1, N_2 and N_3 are the fiber contents of the experimental campaign, described in Section 3.6 and Table 3.6, and the points labeled UHPC correspond to practical mixes with 1% and 2% of fibers by volume. . . 18

2.3 Experimental setups: (left) PIV configuration with planar laser illumination and tracer particles seeding the Carbopol solution for velocity field measurement; (right) PTV configuration with volumetric LED illumination for tracking individual fiber motion in three dimensions. 20

2.4 PIV workflow: (i) image pre-processing (background subtraction and contrast enhancement), (ii) dynamic masking (Carbopol boundary detection), (iii) multi-pass cross-correlation for velocity-vector calculation, and (iv) post-processing (outlier removal and interpolation). 21

2.5 PTV workflow, from image preprocessing (background subtraction and contrast enhancement) to fiber detection, feature extraction, and trajectory reconstruction. 25

2.6 L-shape benchmark geometry showing the vertical inlet channel, ninety-degree bend, and horizontal outlet channel. The container is tilted thirty degrees from vertical to induce flow via gravity, and the rear face is finished in matte black to provide a dark-field background for the optical measurements. 31

3.1 Multi-camera setup for simultaneous PIV and PTV measurements. Cameras 1–3 provide wide-field coverage of the L-shape container (Cam 1) and beam container from orthogonal views (Cams 2–3). Camera 4 captures high-speed close-up footage of the L-shape outlet region to resolve initial flow development. 35

3.2 Render of the experimental setup, showing the multi-camera system, laser sheet configuration, and the acrylic containers for the L-shape and beam tests. The optical table allows precise positioning and alignment of all components. 37

3.3 Fiber size relative to a Chilean 10 peso coin (diameter of 21 mm). . . . 38

3.4 Temporal block structure for PIV analysis. The image sequence is divided into blocks with progressively increasing correlation time intervals ($\Delta t_1 < \Delta t_2 < \Delta t_3$) to adapt to decreasing flow velocities. Intermediate frames (transparent) are skipped to reduce computational cost while preserving temporal resolution during critical flow phases. 42

3.5 Representative examples of non-uniform illumination along individual steel fibers. The reflected intensity varies as a function of fiber orientation relative to the LED sources and local surface characteristics, producing three distinct cases: (I) a poorly visible fiber with low overall reflectance, (II) a well-illuminated fiber with consistent brightness, and (III) a fiber with irregular brightness distribution along its length due to surface unevenness. This variability is further compounded at high fiber concentrations by inter-particle occlusion, posing significant challenges for conventional segmentation techniques. 44

3.6 Representative examples from the training dataset: (a) image with a sparse distribution of well-separated fibers, and (b) image with a denser distribution where overlapping fibers and bubbles are more prominent. 46

3.7 Workflow of the high-concentration fiber detection algorithm. The original image is upscaled by a factor of two to enhance feature resolution, then processed through overlapping 640×640 pixel patches for YOLO-based segmentation. Detected fiber masks are reconstructed into the original image space to yield improved segmentation performance in dense fiber clusters. 47

3.8 (a) Workflow of the α-β-γ filter for fiber state prediction and cosine similarity-based matching for fiber association across frames. (b) Geometric features extracted from detected fibers: orientation θ, length ℓ, and centroid coordinates (c_x, c_y), which are used to construct the feature vectors for similarity evaluation. 48

3.9 Temporal block structure for PTV analysis. The image sequence is divided into blocks with progressively increasing frame-skipping intervals to adapt to decreasing fiber velocities and maintain tracking accuracy while reducing computational cost during less dynamic flow phases. . . 51

3.10 PTV test images with varying fiber loadings: 25, 100, 200, 400 and 800 fibers respectively. 52

3.11 The L-shaped device prior to an experimental run at each fiber dosage, 750, 1500, and 3000 fibers, shown for both Carbopol formulations. The images illustrate the increase in fiber density and spatial overlap with the loading. 56

4.1 PTV tracking performance across fiber loadings. (a) Cumulative number of distinct track identities over the 600-frame sequence, normalized by the known fiber count. Every fiber crosses the field of view within this window, so an ideal tracker would approach unity; values above it indicate trajectory fragmentation rather than spurious detection. (b) Number of fibers detected per frame, showing a detection rate that is stable in time at every loading. 65

4.2 Spatial zones used throughout the analysis. The L-shape is partitioned into the vertical arm (Z1), the elbow (Z2), and the outlet (Z3); the beam is partitioned into a two-by-three grid of cells (r1c1–r2c3), with the lower row adjacent to the L-shape outlet. The figure is drawn to the proportions of the mold and is scaled uniformly, so relative dimensions are preserved. 67

4.3 Flow stages by zone for the fiber-free control of each Carbopol matrix (Carbopol 0.2%, left; Carbopol 0.5%, right). Each horizontal bar is a zone, ordered from top to bottom along the downstream filling sequence, from the L-shape zones Z1–Z3 to the beam cells r$ic j$, and is segmented into the onset, transition and quasi-steady stages, shaded from light to dark within each matrix. Zones marked with an asterisk inherit their boundaries from their geometric region, too few control-case measurements being available to resolve them individually. 69

4.4 Spatial distribution of $|\mathbf{u}|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom). 72

4.5 Spatial distribution of $|\zeta|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom). 72

4.6 Spatial distribution of the fiber speed $|\mathbf{v}|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom). 76

4.7 Spatial distribution of the fiber angular velocity $|\omega|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom). 76

4.8 Spatial distribution of the orientation order η_1 across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom). 77

4.9 Final orientation order parameter η_1 in each beam cell, grouped by fiber content N and rheological matrix (blue tones: Carbopol 0.2%; orange tones: Carbopol 0.5%). Error bars span one standard deviation across the fibers of each group; a colour missing from a cell indicates that no fibers of that matrix were tracked there. 80

4.10 Final centroid dispersion index Λ in each beam cell, grouped by fiber content N and rheological matrix, with the same colour convention as Figure 4.9. The dotted line marks $\Lambda = 1$, the value of a spatially uniform arrangement; lower values indicate clustering. Error bars span one standard deviation across the fibers of each group. 80

4.11 Routes followed by the fibers reaching each cell of the beam, reconstructed by stitching the PTV trajectory fragments. Each row is one destination cell, and the marked zones are those its fibers traversed on the way. Left: Carbopol 0.2%. Right: Carbopol 0.5%. The three dosages are pooled. 84

4.12 Orientation order η_1 along the mean path, by zone, against the shear strain Γ accumulated from the entry to the L-shape, for the three fiber dosages. Zones Z1 to Z3 lie in the L-shape and the cells r2c1 onward in the beam, separated by the dashed line. The order rises to a peak at Z3 and drops at the first beam cell, which sits at larger accumulated strain, the drop marking the disordering impact where the material falls from the outlet into the beam. The value shown for each zone is that of the fibers arrested or observed within it, not of the landing that precedes the first beam cell, which the planar measurement does not resolve. The strain is referred here to the L-shape entry, whereas the saturation law of Chapter 6 is fitted against strain referred to the beam entry, downstream of this reset. Left: Carbopol 0.2%. Right: Carbopol 0.5%. 85

6.1 Orientation order η_1 against the shear strain Γ_{beam} accumulated between the entry of the material into the beam and its arrest, for the three fiber dosages in Carbopol 0.2%. Large markers are the arrest points of each beam cell, curves are the fits of Equation (5.27) to those points, and shaded bands are the confidence intervals of the fits. Small markers and arrows show the reconstructed Lagrangian passage of the fiber population through the upstream cells of each destination. 106

6.2 Effective misalignment δ_{eff} of the arrested fiber population against the accumulated shear strain Γ_{beam}, for the three fiber dosages in Carbopol 0.2%. The curves are the angle predictor of Equation (6.3), obtained by reparametrizing the fits of Figure 6.1: the asymptotic angle δ_∞ widens with the fiber content while the strain scale Γ^* is carried over unchanged. Markers and conventions as in Figure 6.1. 109

6.3 Final orientation angle of the fibers in each beam cell and at each fiber dosage. Coloured boxes give the interquartile range of the measured per-fiber angle; the white marker is the population angle predicted from the accumulated strain of the fiber-free flow through Equation (6.3). The prediction is a property of the population, not of the individual fiber, whose angle spans the full range within every cell. Carbopol 0.2%. . . . 110

B.1 PIV inter-pair skip sensitivity for Carbopol 0.2%, camera 1. Left: average velocity under four inter-pair skips, showing convergence as the flow stabilizes. Right: percentage of spurious vectors identified by the UOD test, on a logarithmic axis. Each skip minimizes the spurious rate over its own interval of the filling, which fixes the transition to the next. . . 162

B.2 PIV inter-pair skip sensitivity for Carbopol 0.5%, camera 1, over the longer window the stiffer matrix requires. Axes and conventions as in Figure B.1. 163

B.3 PTV skip sensitivity for Carbopol 0.2%. Pair displacement Δx, the distance a fiber travels between analyzed frames, under four skips, against the floor of 0.25 mm (dashed). Each skip is drawn bold over the interval in which it is adopted; below the floor the displacement is indistinguishable from the localization noise. 164

B.4 PTV skip sensitivity for Carbopol 0.5%Pair displacement Δx, the distance a fiber travels between analyzed frames, under four skips, against the floor of 0.25 mm (dashed). Each skip is drawn bold over the interval in which it is adopted; below the floor the displacement is indistinguishable from the localization noise. 164

D.1 Dimensions of the L-shaped device and of the beam. All measurements are in millimetres. 171

D.2 Positions and viewing directions of the four cameras relative to the mold. Distances are in millimetres. 172

D.3 Render of the assembled setup, showing the mold, the cameras, and the illumination mounted on the optical table. 172

D.4 Render of the assembled setup from a second viewpoint. 173

F.1 Arrested state of the beam in Carbopol 0.2%. (a) Run M74, 750 fibers. (b) Run M73, 1500 fibers. 185

F.2 (a) Run M70, 3000 fibers in Carbopol 0.2%, which completes the series of the softer matrix. (b) Run M86, 750 fibers in Carbopol 0.5%, which begins the series of the stiffer one. The two are shown together so that the contrast in filling extent is immediate: the softer matrix reaches the far end of the beam at its highest dosage, while the stiffer one does not reach it at its lowest. 186

F.3 Arrested state of the beam in Carbopol 0.5%. (a) Run M94, 1500 fibers. (b) Run M83, 3000 fibers. The material arrests before filling the third column of the beam at every dosage of this matrix. 186

G.1 Colour palette used throughout this work. Carbopol 0.2% is drawn in blues and Carbopol 0.5% in oranges; within each matrix, the three fiber dosages are drawn in a light, medium and dark tone, the darkest denoting the highest fiber content. 189

List of Tables

2.1 Concentration regimes for a suspension of rigid rods of aspect ratio r_f. 19

3.1 Camera specifications and operating parameters. 36

3.2 Physical and mechanical properties of the Dramix® OL 13/.20 steel fibers used in the experiments. The numerical density is the manufacturer's value and is the conversion factor by which fibers are counted by weighing. 37

3.3 Preparation parameters for Carbopol Ultrez 10 solutions at 0.2% and 0.5% concentrations. NaOH additions are sequential: the primary dose initiates gelation and the secondary dose achieves the target pH range. 38

3.4 Herschel-Bulkley parameters for the two Carbopol Ultrez 10 formulations. The yield stress τ_0 was determined from the arrested free-surface deposit to avoid wall-slip underestimation; the consistency index K and flow index n were refitted from flow-curve data with τ_0 held fixed. The critical shear rate $\dot{\gamma}_c = (\tau_0/K)^{1/n}$ separates the yield-dominated from the viscous-dominated response. 40

3.5 PIV interrogation parameters by camera and Carbopol concentration. Window sizes follow a three-pass multi-pass scheme; the final window size determines spatial resolution. 41

3.6 Fiber dosages of the experimental campaign, set against the boundaries of the concentration regimes and against the dosages of a real UHPC mix. The number in the L-shape is the subset deposited within the device at a fill height of 22.4 mm, about 73% of each batch, and is informative only; the volume fractions and their ratios are computed over the full working suspension of 1.802 L, from the nominal fiber volume of 0.408 mm^3. Two regime boundaries follow from the aspect ratio $r_f = 65$: the dilute limit $\phi_{sd} = r_f^{-2}$, below which the fibers do not interact, and the concentrated limit $\phi_c = r_f^{-1}$, above which they are in permanent contact; the semi-dilute regime lies between them. The campaign occupies the lower edge of that regime, while a real UHPC mix sits at or above its upper edge. Reference rows are not tested conditions. 54

3.7 Quantities used throughout the analysis. 57

3.8 Interpretation bands for the partial eta squared, after Cohen, and the thresholds adopted in this work. The relevance threshold sits just above the small-effect boundary; the grading thresholds span the range observed in this dataset and do not mark new statistical categories. . . . 58

4.1 Segmentation performance for rod-like particle detection. Metrics for YOLOv11 and SAM2 are computed at the instance mask level; the Hough transform IoU is computed at the pixel level and is not directly comparable. 63

4.2 PTV tracking performance across fiber loadings, evaluated over the 600-frame window within which every fiber crosses the field of view. The detection ratio is the instantaneous fraction of the fiber population segmented per frame; trajectory lengths are given in frames; "complete" denotes the fraction of tracks lasting ≥ 500 frames, that is, spanning at least five sixths of the window. 64

4.3 PIV total relative error by camera and Carbopol concentration. Percentage errors are relative to the mean flow speed $\bar{u}$. 70

4.4 Effect of fiber content on the two Eulerian variables, by flow stage and matrix. Cells report the peak η^2 across all spatial zones, with the zone of maximum effect in parentheses. The bottom rows report η^2 computed over the pooled domain. * $\eta^2 \geq 0.02$; ** $\eta^2 \geq 0.10$; *** $\eta^2 \geq 0.30$. . . . 71

4.5 PTV tracking efficiency by nominal fiber concentration and Carbopol matrix, computed as the ratio of detected to expected fibers within the field of view during the quasi-steady window. 74

4.6 Effect of fiber content on the three Lagrangian variables, by flow stage and matrix. Cells report the peak η^2 across all spatial zones, with the zone of maximum effect in parentheses. The bottom rows report η^2 computed over the pooled domain. * $\eta^2 \geq 0.02$; ** $\eta^2 \geq 0.10$; *** $\eta^2 \geq 0.30$. 75

4.7 Effect size η^2 of the principal determinants of the fiber final state, from the global Analysis of Covariance over the pooled domain. The relevance threshold of Section 3.7 is $\eta^2 = 0.02$. 79

4.8 Net effect of fiber content on the final state, isolated from the geometric effect within each rheological matrix. The relevance threshold of Section 3.7 is $\eta^2 = 0.02$, which neither orientation entry reaches. 82

6.1 Agreement between the predicted and the measured final angle, by fiber dosage, over all beam cells. The prediction uses the accumulated strain of the fiber-free flow. The bias is the mean signed departure, the RMSE its root-mean-square, and p tests the bias against zero; the practical tolerance is $\pm 5°$. 110

A.1 Assumptions underlying the orientation saturation law. 159

B.1 PIV block structure by camera and matrix. Each block is defined by its inter-pair skip s_i (frames between the two frames of a correlated pair) and its final skip s_f (frames between one pair and the next); the block length $s_\mathrm{i} + s_\mathrm{f} + 2$ sets the time step $\Delta t = \mathrm{block/fps}$. Cameras 1–3 record at 220 fps, camera 4 at 660 fps. 166

B.2 PTV block structure by camera and matrix, for the three cameras used in the fiber tracking. Each block is defined by its skip (frames between analyzed frames) and the maximum linking distance Δx_max. All three cameras record at 220 fps. 167

B.3 Reduction in computational load relative to processing every frame: for PIV, the reduction in the number of correlated pairs; for PTV, the reduction in the number of frames analyzed. PTV uses cameras 1–3 only. 167

C.1 Average pH measured during the three stages of the Carbopol preparation, for every batch of the campaign. pH_1 is measured after dispersion of the polymer and before neutralization; pH_2 after the first dose of 1M NaOH; pH_3 after the second dose, and is the value against which the target range of 6.8–7.2 is checked. Each entry is the average of five points sampled across the mixture. 169

E.1 Effect of fiber content N on the Eulerian variables, Carbopol 0.2%, by zone and flow stage. Each variable reports the partial eta squared and the covariate-adjusted p-value of the fiber-content term. Bold marks a relevant and significant effect ($\eta^2 \geq 0.02$ and $p < 0.05$); asterisks mark the effect-size bands alone. $^{*}\,\eta^2 \geq 0.02$; $^{**}\,\eta^2 \geq 0.10$; $^{***}\,\eta^2 \geq 0.30$. . . . 176

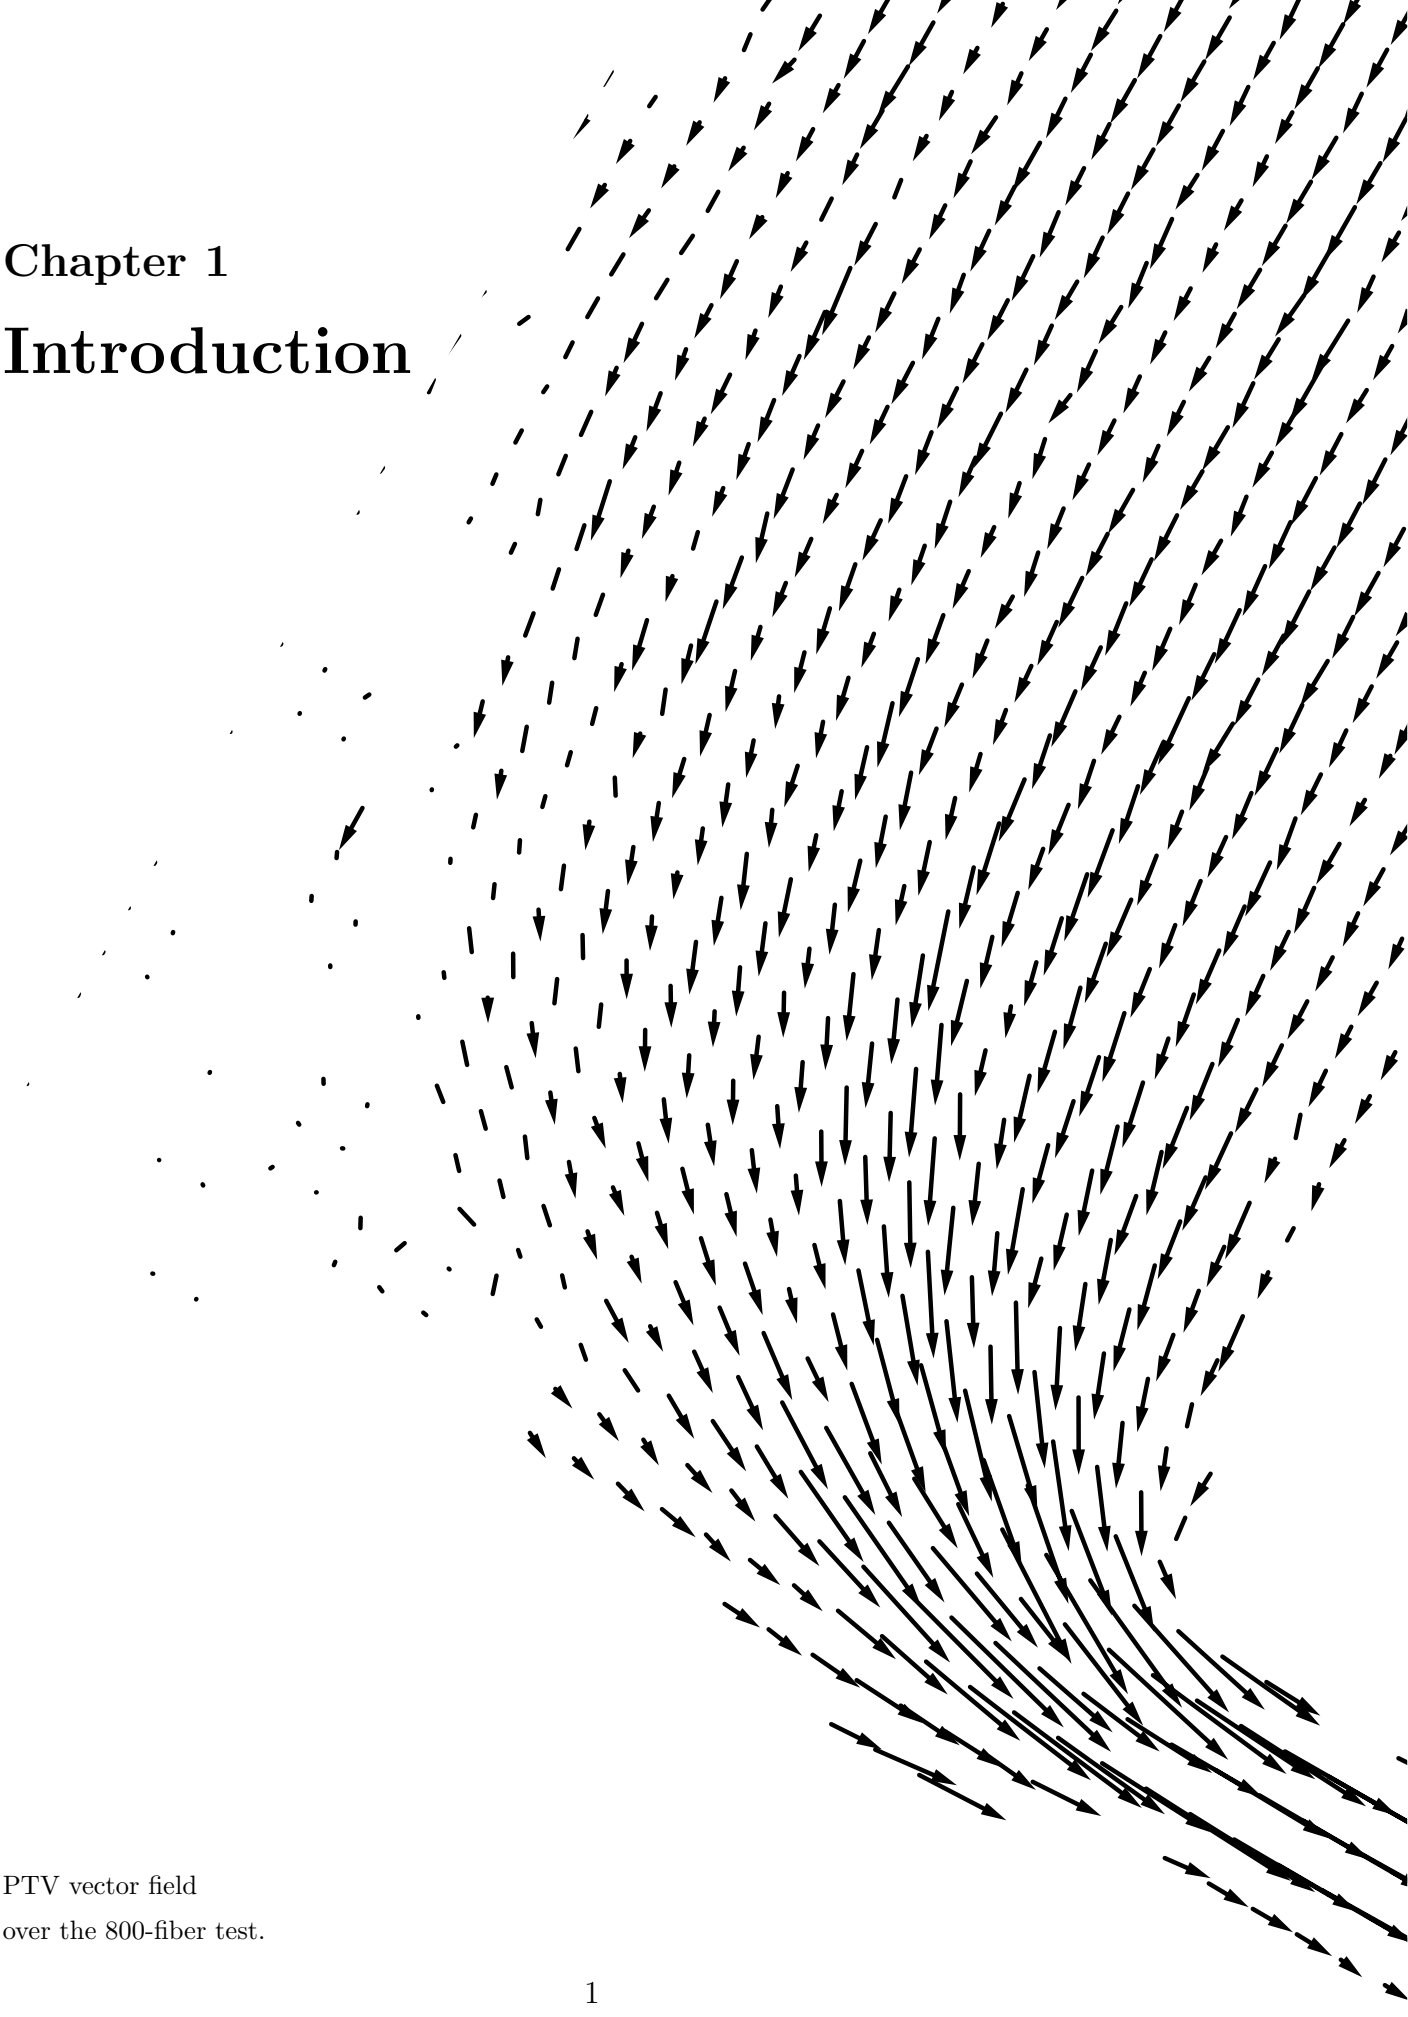
E.2 Effect of fiber content N on the Eulerian variables, Carbopol 0.5%, by zone and flow stage. The stiffer matrix arrests before filling the third column of the beam, so its distal cells receive no material and are absent from the table. Conventions as in Table E.1. 177

E.3 Effect of fiber content N on the Lagrangian variables, Carbopol 0.2%, by zone and flow stage. The table covers the same zones as the Eulerian analysis of Table E.1; the distal cell the flow reaches only in sparse transient passages, unable to support the test, is omitted. Conventions as in Table E.1. 178
E.4 Effect of fiber content N on the Lagrangian variables, Carbopol 0.5%, by zone and flow stage. As in the Eulerian analysis, the distal columns the arresting flow never fills are absent from the table; the six zones listed are those in which the carrier establishes a resolved field and a countable fiber population is tracked. Conventions as in Table E.1. 179

F.1 Final-state descriptors of the deposited fiber population in each beam cell, by matrix and nominal fiber content. n: tracked fibers in the cell; $\hat{\eta}_1$: debiased orientation order; $\bar{\theta}$: mean director; R_g: radius of gyration of the centroids, against a uniform reference of $R_g^{\mathrm{unif}} = 30.8$ mm; $\Lambda = R_g/R_g^{\mathrm{unif}}$: spatial dispersion index, with $\Lambda < 1$ denoting clustering. . . . 181
F.2 Effect of the rheological matrix on the final state, by zone and nominal fiber content. $\Delta\hat{\eta}_1 = \hat{\eta}_1^{0.2} - \hat{\eta}_1^{0.5}$, with its permutation p-value; r_Λ is the rank-biserial effect on the dispersion, with its Mann-Whitney p-value; p_{WW} is the Watson-Williams p-value for the difference in director. Bold marks $p < 0.05$. 183
F.3 Tests on the final centroid arrangement, by zone. The effect of the fiber content is tested within each matrix and reported as η^2; the effect of the rheology is tested at each nominal dosage and reported as a rank-biserial correlation r. Bold marks $p < 0.05$. Cells absent from a matrix are omitted. 187

Chapter 1

Introduction



PTV vector field
over the 800-fiber test.

Concrete is ubiquitous: it can be found in every house, town, city, and country, regardless of the level of development, wealth, or political foundation of the region. It is such a versatile material, initially fluid, but hardening over time into a rock like solid, that it can be used in countless applications, from the foundation of a small house to the construction of bridges and skyscrapers. Concrete is produced from a mixture of components: cement or cementitious materials, water, chemical admixtures, and aggregates (such as sand, gravel, or stone). Together, these elements constitute the most widely used construction material worldwide, with global demand projected to increase by approximately 12%-23% by 2050 (International Energy Agency, 2018).

Although concrete exhibits high compressive strength, its tensile strength is relatively weak. Consequently, steel reinforcement is added to improve the tensile capacity of concrete structures (Guo, 2014). However, durability issues arise due to the corrosion of steel reinforcement, which can lead to structural failure and reduced service life. Chlorides can penetrate the porous structure and microcracks of the concrete, leading to steel corrosion, reduction in structural integrity, and increased maintenance costs (Liu et al., 2020; Angst et al., 2009). In fact, developed countries allocate approximately 40-50% of their construction budgets solely to the maintenance and repair of existing structures (Chang et al., 2021). Thus, the construction industry not only seeks to develop more environmentally sustainable materials and methodologies, but also aims to improve the durability and mechanical properties of structures and materials to reduce long-term maintenance and repair costs (Habert et al., 2020).

Ultra-High-Performance Concrete (UHPC) has been developed to address these challenges, with the aim of improving durability and extending the lifespan of concrete structures. UHPC is characterized by very high compressive strength (around 150 MPa) (Bahmani et al., 2024), extremely low permeability, extended service life exceeding 200 years with minimum maintenance (Li et al., 2020), and excellent resistance to environmental factors such as freeze-thaw cycles and chemical attacks (Kang et al., 2025). These properties are achieved by incorporating high-quality materials such as silica fume, superplasticizers, and steel fibers, which greatly enhance the mechanical and durability performance of the composite. Figure 1.1 shows a comparison between UHPC and conventional concrete.

Fibers are considered the most essential component of UHPC (Shi et al., 2024). They

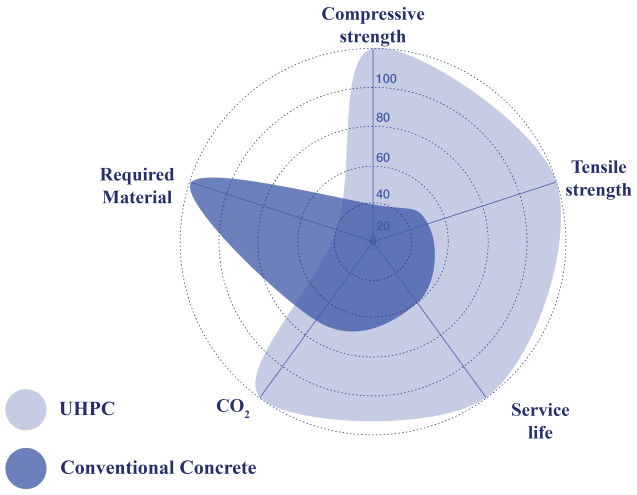


Figure 1.1: Comparison between UHPC and conventional concrete in terms of compressive strength, tensile strength, durability, CO_2 emissions, and material requirements. While CO_2 emissions increase per unit, as the required material is reduced and the service life is extended, the overall environmental impact of UHPC is lower than that of conventional concrete.

do not merely replace steel reinforcement but also provide distributed reinforcement throughout the matrix, improving the tensile behavior of UHPC (typically around 15 MPa) and enhancing the post-cracking performance of the material (Wille et al., 2011). This generates strain-hardening behavior, allowing UHPC to sustain higher loads and larger deformations without brittle failure. Moreover, fibers prevent explosive failure and enhance durability by bridging microcracks and reducing crack widths, which limits the ingress of harmful substances such as chlorides and moisture, thereby protecting embedded steel reinforcement from corrosion (Wang et al., 2015). Nevertheless, the effectiveness of steel fibers strongly depends on their orientation and dispersion within the cementitious matrix. Undesirable alignment or clustering can significantly limit their contribution, reducing the tensile strength and overall efficiency of UHPC (Kang et al., 2010). Figure 1.2 illustrates the expected fiber distribution and alignment versus the anticipated failure mode of the element.

In fact, several studies have shown that fiber concentration and alignment can significantly influence the mechanical properties of UHPC, with higher fiber content and better alignment leading to improved tensile strength and durability (Abrishambaf et al., 2017). Furthermore, fiber concentration and geometry significantly influence

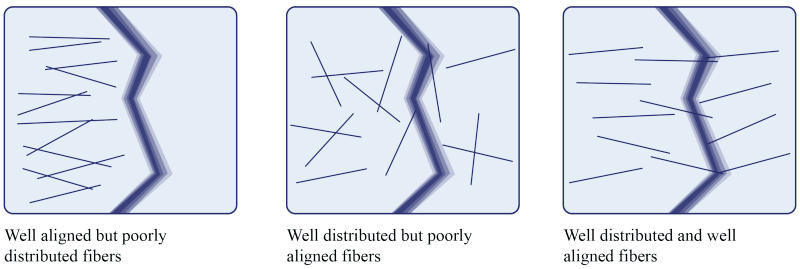


Figure 1.2: Expected fiber distribution and alignment versus anticipated failure mode of the element.

the rheological properties of the fresh mixture, which in turn affects the flow behavior and final fiber orientation (Ravichandran et al., 2022). Therefore, a four-way coupling system emerges between rheology, flow, fiber orientation, and mechanical performance: the rheological properties of the mixture influence fiber behavior, fibers modify the rheological properties of the mixture, and fibers interact with one another (Ferrara et al., 2008). Thus, it is essential to understand not only how individual fibers interact with the surrounding fluid but also how their collective behavior alters the global flow dynamics and, consequently, the final orientation of the fibers.

Fiber behavior in UHPC is studied today almost entirely in the hardened state, by sectioning specimens for mechanical testing or by scanning them with X-ray computed tomography (Zhang et al., 2021; Konrád et al., 2025). Both record the arrangement the fibers came to occupy, and neither records the process by which they came to occupy it: the orientation is set during the seconds in which the material is placed, and the flow responsible for it has long stopped by the time an element can be cut or scanned. Coupled computational fluid dynamics-discrete element method (CFD-DEM) simulation offers the route the hardened state cannot, since it resolves the flow of the matrix together with the motion of each individual fiber and therefore delivers the whole history rather than its endpoint, and it does so at a cost that full-scale casting cannot approach, opening the range of mixes, placement techniques, and mold geometries that can be explored. A CFD-DEM model of UHPC is under development in parallel by the research group, and the simulation itself lies outside the present scope. What lies within it is what the simulation requires before it can be trusted. A model that reproduces a history cannot be validated against the arrangement the fibers finally reached, since

many histories terminate in the same arrangement, and it must therefore be confronted with the flow and the fibers as they move. The paired fluid-phase and fiber-phase data produced in this thesis are the measurements against which that model will be validated.

The continuum route already has a foothold in concrete research: the Folgar-Tucker and Advani-Tucker orientation tensors have been coupled to full flow simulations of fiber-laden self-compacting concrete and confronted, qualitatively, with the final state of an L-box casting of fibers suspended in a transparent gel (Svec et al., 2012; Gudžulić et al., 2019). What these simulations compare, however, is a simulated end state against a photographed end state; neither the strain delivered to the material nor the orientation of individual fibers is measured as the flow develops, so the confrontation is with an outcome, not with the history that produced it. The same gap is named, rather than closed, by the research group that has come closest to it: having established an orientation-tensor model numerically in a concrete rheometer, Gerland et al. (Gerland et al., 2024) identify the simultaneous optical measurement of flow and time-resolved fiber orientation in a transparent surrogate as the step still missing, and Vaupel et al. (Vaupel et al., 2025) supply exactly that measurement one year later, but in a Newtonian silicone oil, where no yield stress divides the domain into a plug and a sheared layer. What remains open, and what the present work supplies, is that same simultaneous measurement carried into a matrix that does yield.

Observing the fibers as they move is what no opaque material permits, since fibers travelling inside fresh UHPC cannot be seen at all. Carbopol, a transparent Herschel-Bulkley fluid combining a yield stress with a shear-thinning post-yield response, therefore replaces the mixture as the matrix of the experiment while allowing light to pass through it (Auernhammer et al., 2020). The property the two matrices are required to share is the yield stress, since it is the yield stress alone that divides a casting into a moving, sheared region and a rigid, unyielded one, and it is that division, not the exact curvature of the flow curve beyond it, that fixes how unevenly the strain is delivered to the fibers. The curvature itself is not a settled property of fresh UHPC: it has been reported as shear-thinning at the shear rates encountered during placement (de Larrard et al., 1998), and as shear-thickening in mixes with a high content of fine powder (Feys et al., 2009), so no single post-yield exponent can be claimed on the matrix's behalf.

Carbopol's own shear-thinning response is therefore adopted here as a property of the working fluid, not asserted as a property the UHPC casting is thereby known to share. Particle Image Velocimetry (PIV) recovers the velocity field of the fluid, and from it the rate of strain at every point and instant, and therefore the strain a material element accumulates along its path. A deep learning-enhanced Particle Tracking Velocimetry (PTV) framework, reported in a companion study (Wolff et al., 2026), recovers the position, the orientation, and the rotation of each individual fiber. Thus, the deformation delivered to a fiber and the alignment it acquired become known together.

Connecting the deformation to the alignment it produces requires a description of how flow orients the fibers, and the one adopted here rests on a competition. The flow orders: its deformation tilts each fiber toward the direction along which the material is being stretched (Jeffery, 1922). The fibers disorder: once packed closely enough to disturb one another, every encounter deflects both in a direction the flow does not determine, eroding the alignment the flow has imposed (Folgar et al., 1984). Neither prevails, because the ordering is strongest when the population is disordered and the randomizing is strongest when it is aligned. Both therefore balance at a partially aligned state, which the population approaches exponentially. Reducing the tensorial orientation model of Advani and Tucker (Advani et al., 1987) to a single scalar, as developed in Chapter 5 and derived in Appendix A, gives the law that organizes the whole of the present analysis,

$$\eta_1(\Gamma) = \eta_{1,\max}\left(1 - e^{-\Gamma/\Gamma^*}\right), \tag{1.1}$$

where η_1 is the alignment of a group of fibers, running from zero at complete disorder to one at perfect alignment, and Γ is the shear strain accumulated along the path the material has travelled. The law compresses the orientation problem into two numbers: how far the fibers can align, $\eta_{1,\max}$, and how much deformation they need in order to get there, Γ^*.

Neither had ever been measured in a cementitious material. The law was built for polymer melts and calibrated in rheometers, on flows that are steady, homogeneous, and Newtonian, and where the deformation reaches every fiber evenly (Eberle et al., 2010; Cieslinski et al., 2016; Egelmeers et al., 2025). A casting is none of these. It is transient and heterogeneous, and its matrix is viscoplastic: it carries a yield stress

that must be exceeded before any part of it flows at all, so the deformation reaches a plug-bound core and a sheared layer in radically unequal amounts. Whether a law calibrated under steady, homogeneous, yield-free conditions can survive their absence is the question this thesis examines. To that end it recovers, on the same flow and at the same instant, the strain delivered to a fiber and the alignment it acquires, across a range of fiber concentrations spanning the upper dilute and the lower semi-dilute regimes, and confronts the resulting orientation with the law. The paired velocity and fiber data, and the two constants recovered from them, are the contribution of this work.

Hypothesis

In an unsteady viscoplastic fluid, two groups of fibers may accumulate entirely different strains even as they travel the same path. The yield stress alone is responsible for this: it splits the material into a sheared region and an unyielded plug that gathers almost no deformation, so the strain delivered to a fiber is set by the part of the flow it crosses and not by the distance it covers, for reasons developed in Chapters 2 and 5. The final orientation state of a group of fibers therefore depends not on the accumulated strain alone, but on the deformation history that produced it and on the fiber content of the suspension, whose interactions the Folgar-Tucker model accounts for.
On this premise rests the hypothesis of this thesis:

The law of exponential approach to equilibrium, obtained by reducing the Advani-Tucker orientation model to a single scalar, predicts the final angular state of a group of fibers in an unsteady viscoplastic flow from the accumulated strain Γ *alone, through two parameters: the attainable alignment* $\eta_{1,\max}$ *and the characteristic strain* Γ^**.*

Both parameters are expected to shift with the fiber content, since the interaction coefficient of the Folgar-Tucker model that sets them grows with the concentration: $\eta_{1,\max}$ should fall as the population is packed more densely. The test follows from a single asymmetry: the deformation a fiber accumulates can be read from the flow before any fibers are added, whereas its alignment can only be read from the fibers themselves, so

PIV recovers the strain of the fiber-free flow and the angle the model predicts from it is confronted with the angle PTV measures, at concentrations spanning the upper dilute and lower semi-dilute regimes. The hypothesis is rejected if the fitted $\eta_{1,\max}$ shows no monotonic decrease with fiber content, or if the predicted angle departs systematically, beyond measurement error, from the angle PTV records at any tested concentration.

General Objective

To establish whether the reduced two-parameter model, derived from the Advani-Tucker orientation model, predicts the final orientation of fibers in a Carbopol casting used as a UHPC analog, from experimental measurements.

1.1 Specific Objectives

Answering this required first standardizing a low-cost, synchronized PIV-PTV acquisition protocol and optimizing the PIV methodology for an unsteady free-surface flow; the deep learning-enhanced PTV framework itself is reported separately (Wolff et al., 2026). With that instrumentation in place, the specific objectives of this thesis are to:

a) To test the applicability of the Folgar-Tucker and Advani-Tucker models to an unsteady viscoplastic shear-thinning flow across a range of fiber concentrations.

b) To determine the coefficients $\eta_{1,\max}$ and Γ^* of the reduced Advani-Tucker model in an unsteady viscoplastic shear-thinning flow, at different fiber concentrations.

c) To determine the coefficient δ_∞ arising from the angular decomposition of the reduced model, and through it interpret the final angle of the fibers relative to a global axis.

d) To validate the angle prediction of the reduced model against the orientations measured by PTV for the several fiber groups and concentrations, using the deformations recovered by PIV in the fiber-free flow, over the concentration range reaching up to the lower limit of the semi-dilute regime.

1.2 Thesis Structure

This thesis is organized into seven chapters. The present chapter has set out the problem, the hypothesis, and the objectives that frame the work. Chapter 2 assembles the theoretical background, from the role of fiber orientation in the mechanical response of UHPC to the constitutive models, the use of Carbopol as a transparent analog, and the optical techniques employed. Chapter 3 describes the methodology, covering the experimental setup, the Carbopol matrices, and the implementation of the Particle Image Velocimetry and Particle Tracking Velocimetry frameworks. Chapter 4 reports the measurements that precede the test of the law, bounding what the tracking can recover, fixing the partition of the domain and the statistical protocol, and establishing how the fiber content acts on the flow and on the fibers it carries. Chapter 5 develops the dynamics of the fibers and reduces the orientation tensor to the two-parameter law of exponential approach to equilibrium that organizes the analysis. Chapter 6 tests that law, confronting the orientation it predicts from the fiber-free flow with the orientation measured across the range of fiber concentrations the experiment reaches. Chapter 7 draws the conclusions and sets out the limits of the measurement together with what the CFD-DEM simulation these data are gathered to validate is positioned to lift. The derivation of the saturation law and the supporting analyses are collected in the appendices.

Chapter 2

Theoretical Background



Fiber tracking example.

This chapter assembles the theoretical elements required to understand the design of the experiment and the results that follow from it. The material is presented in the order in which it is needed. Section 2.1 establishes why the orientation of the fibers governs the mechanical response of UHPC and why that orientation is decided during casting rather than during mix design, which is what makes the flow the object of study. Section 2.2 introduces the constitutive models that describe the fresh state of the material, justifies the use of Carbopol as a transparent analog for it, and delimits the concentration regime in which the fibers of this study operate. Section 2.3 presents the two optical techniques employed, PIV for the kinematics of the flow and PTV for the state of the fiber population, together with the error sources that bound each measurement and the metrics through which the raw output of both is reduced to interpretable quantities. Section 2.4 describes the geometry in which the flow is produced and the reasons it was adopted. The dynamics of the fibers themselves, and the models that predict their reorientation under a given flow, are developed separately in Chapter 5.

2.1 Ultra High Performance Concrete

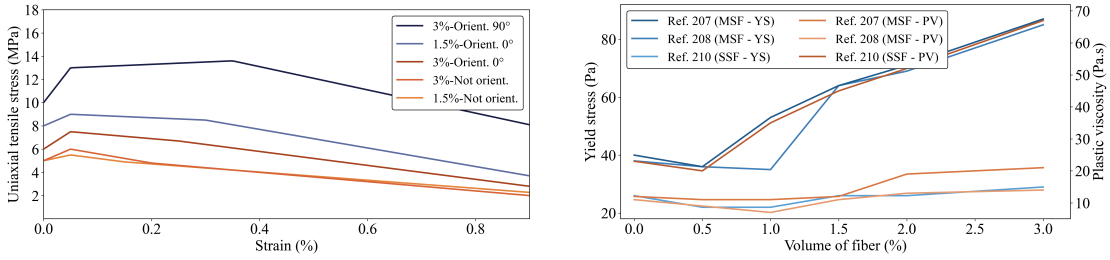
UHPC is an advanced cementitious composite that attains compressive strengths above 150 MPa and, through the incorporation of short steel fibers, a tensile capacity that persists after cracking (Richard et al., 1995; Wille et al., 2011; Graybeal, 2014). It is that post-cracking tensile response, and not the compressive strength, that the fibers govern, and this section establishes how far it is decided by their orientation and spatial arrangement, and how far these in turn are decided during casting rather than during mix design.

2.1.1 Fiber Influence on Mechanical Performance

The tensile behavior of UHPC is controlled by the fibers that bridge the crack, and the contribution of a single fiber to that bridging depends on the angle it forms with the crack plane. Fibers lying close to the direction of the applied tension transfer stress across the crack efficiently, whereas fibers inclined at large angles carry a smaller share of the load and are more prone to matrix spalling at the exit point, so a population that is well aligned with the loading direction develops a higher tensile strength and a

more pronounced strain-hardening response than one that is not (Abrishambaf et al., 2017; Zhou et al., 2017; Kang et al., 2011). Figure 2.1(a) illustrates this dependence for specimens of identical composition that differ only in the orientation of their fibers. The magnitude of the effect is substantial: variations in casting procedure alone, with no change in mix design or dosage, have been reported to alter the post-cracking tensile strength by a factor of two or more (Ferrara et al., 2011; Wille et al., 2014). Orientation is therefore not a secondary characteristic of the cast element but a property of the same order of importance as the dosage itself.

The dosage, in turn, does not act on the composite in isolation. Increasing the fiber content raises the yield stress and the plastic viscosity of the fresh mix, since the fibers obstruct the flow of the paste and, above a certain content, begin to interfere with one another and to form clusters that the flow can no longer disperse (Ravichandran et al., 2022; Martinie et al., 2010; Roussel, 2006b). Figure 2.1(b) shows this effect for a UHPC mix over a range of fiber dosages. The consequence is that the dosage reaches the mold through the flow it has itself modified, so that a higher fiber content is placed under a fresh state that is less able to orient it. The gain in tensile capacity that a higher dosage promises is thus partially offset by the loss of alignment that the accompanying change in rheology produces, and the two effects act against one another rather than in series (Ferrara et al., 2011; Alberti et al., 2017).



(a) (b)

Figure 2.1: (a) Tensile stress-strain curves of UHPC specimens with different fiber orientations. (b) Effect of fiber dosage on the rheological properties of UHPC.
Source: (Abrishambaf et al., 2017) and (Ravichandran et al., 2022).

Two consequences follow, and together they frame the experimental problem addressed in this work. The first is that the orientation of the fibers is a design variable and not an incidental outcome, since the same mix, the same dosage, and the same mold yield different tensile capacities depending on how the material was placed (Ferrara

et al., 2011; Wille et al., 2014; Alberti et al., 2017). The second is that the dosage cannot be raised indefinitely in the expectation of a proportional gain, because the fiber content modifies the rheology of the fresh mix and, through it, the orientation the fibers acquire while the mold is being filled (Martinie et al., 2010; Alberti et al., 2017). Both consequences point to the same object, namely the flow that carries the fibers into the mold. It is that flow, and the deformation it imposes on the fibers as they travel through it, that the experiments of this work set out to measure.

2.2 Rheology of UHPC and Carbopol Analog

The flow behavior of fresh UHPC and its interaction with embedded fibers is governed by the rheological properties of the cementitious matrix. To investigate fiber dynamics in a controlled experimental setting, Carbopol, a transparent shear-thinning viscoplastic fluid, is employed as an analog for UHPC (Auernhammer et al., 2020). This section reviews the constitutive models used to describe viscoplastic fluids, examines the rheological characteristics of both materials, and establishes the concentration regime in which the fibers of this study operate.

2.2.1 Viscoplastic Constitutive Models

Viscoplastic materials exhibit a yield stress τ_0 below which the material behaves as a solid and above which it flows as a viscous fluid (Bingham, 1922; Bonn et al., 2017). The Bingham model represents the simplest viscoplastic behavior, assuming a linear relationship between shear stress τ and shear rate $\dot{\gamma}$ after yielding,

$$\tau = \tau_0 + \mu_p \dot{\gamma} \quad \text{for } \tau > \tau_0, \tag{2.1}$$

where μ_p is the plastic viscosity. This model is widely used to characterize fresh cementitious materials, since it captures the essential yielding behavior with two parameters and both of them can be extracted from a standard rheometer test (de Larrard et al., 1998; Roussel, 2006a). It does not, however, reproduce the curvature of the flow curve that fresh pastes exhibit at the shear rates encountered during placement, and fitting a straight line to a curved response can return a yield stress that depends on the range

of shear rates over which the fit is performed (de Larrard et al., 1998).

The Herschel-Bulkley model extends this formulation by incorporating a power-law dependence of the stress on the shear rate through the exponent n, capable of describing a shear-thinning or a shear-thickening post-yield response depending on its value,

$$\tau = \tau_0 + K\dot{\gamma}^n \quad \text{for } \tau > \tau_0, \tag{2.2}$$

where K is the consistency index and n is the flow behavior index (Herschel et al., 1926). This is the form that has been shown to describe fresh cementitious pastes over the range of shear rates encountered during placement (de Larrard et al., 1998). The two ingredients of the model have distinct physical consequences for the flow. The yield stress τ_0 divides the material into two coexisting states: wherever the local stress exceeds τ_0 the material yields and flows with a finite shear rate, whereas wherever the stress remains below τ_0 the material does not deform and moves as a rigid unyielded block, commonly referred to as a plug (Balmforth et al., 2014). The exponent n governs the post-yield response: when $n < 1$ the effective viscosity decreases as the shear rate increases, a shear-thinning behavior characteristic of Carbopol solutions, while $n = 1$ recovers the linear Bingham response. When $n > 1$ the model instead describes a shear-thickening response, in which the effective viscosity rises with the shear rate; this has been reported in UHPC mixes with a high content of fine powder, where the densely packed particles form transient hydrodynamic clusters that stiffen the suspension as the shear rate grows (Feys et al., 2009). The value of n is therefore not a fixed property of fresh cementitious matrices but depends on the mix design, and the post-yield curvature of a given UHPC cannot be assumed without characterization. Equating the two terms of Equation (2.2) defines a shear rate $\dot{\gamma}_c = (\tau_0/K)^{1/n}$ at which the viscous contribution matches the yield stress, below which the local stress stays close to τ_0 and the material behaves as a near-plug, and above which the power-law term dominates.

The coexistence of yielded and unyielded regions is the feature of this constitutive law with the most direct bearing on the present work. In a Newtonian fluid the deformation is distributed throughout the domain, and every fluid element accumulates strain as it flows. In a yield-stress fluid it is not. The deformation concentrates in the sheared regions near the walls and around the obstacles, while the core of the material translates

as a solid body and accumulates almost no strain (Coussot, 2014; Balmforth et al., 2014). Two fibers that travel the same distance through the same mold may therefore experience histories of deformation that differ by an order of magnitude, depending on whether they traveled through the shear layers or within the plug. Since it is the deformation accumulated along the path, and not the distance traveled, that governs the orientation the fibers acquire, the strain history becomes a quantity that must be measured rather than inferred from the trajectory.

2.2.2 Rheological Properties of Carbopol

Carbopol is a cross-linked polyacrylic acid polymer that, once dispersed in water and neutralized with a base such as NaOH, forms a transparent gel whose consistency can be tuned through the polymer concentration (Piau, 2007). At rest the neutralized microgel behaves as a soft jammed solid, and it begins to flow only once the applied stress exceeds a finite yield stress, above which it exhibits shear-thinning behavior consistent with the Herschel-Bulkley description introduced in the previous section (Piau, 2007; Dinkgreve et al., 2016). This combination of a well-defined yield stress and a shear-thinning response is what makes Carbopol a suitable analog for the fresh state of cementitious materials such as UHPC.

Beyond its rheological similarity, Carbopol offers two experimental advantages. The first is optical transparency: unlike opaque cementitious pastes, a neutralized Carbopol gel allows light to pass through it, enabling the direct visualization of embedded fibers and seeded tracer particles and therefore the application of optical techniques such as PIV and PTV (Auernhammer et al., 2020). The second is temporal stability: because the gel does not undergo hydration or setting, its rheological properties remain essentially constant over the duration of an experiment, so the time-dependent evolution that complicates the testing of fresh concrete is absent. These two properties allow the flow of a yield-stress fluid to be characterized under controlled and repeatable conditions, which is not feasible with the opaque and continuously evolving cementitious system it represents.

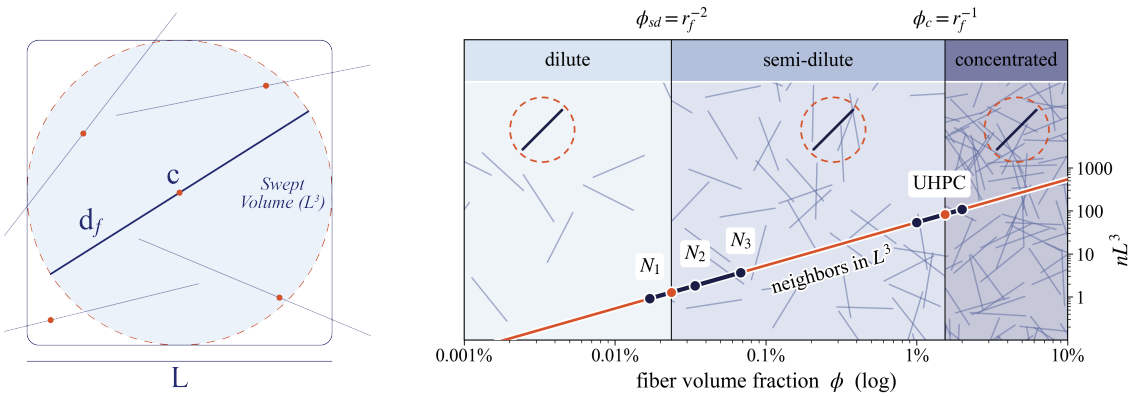
The substitution is not free of cost, and its limits should be stated alongside its advantages. Carbopol is a homogeneous microgel, whereas fresh UHPC is a suspension of aggregate particles in a cement paste, so the analog reproduces the continuum rheology

of the mix but not its granular structure. It carries no aggregate, so the interference between the fibers and the coarse particles of the matrix is absent, and it does not set, so the thixotropic build-up that stiffens a cementitious paste at rest has no counterpart (Roussel, 2006b). A further limitation concerns the boundary condition: Carbopol is prone to slip at smooth walls, where a thin solvent layer forms and the gel translates rather than shears, so the near-wall velocity gradient measured in a smooth mold may be smaller than the one the constitutive law would predict (Meeker et al., 2004; Dinkgreve et al., 2016). What the analog does preserve is the feature that governs the kinematics of the flow: the existence of a yield stress that separates the plug from the sheared regions. The shear-thinning response Carbopol exhibits above that threshold is a property of the working fluid, adopted here for its own sake and not asserted, per the qualification above, as a curvature the UHPC it represents is thereby known to share. Since it is the kinematics of the flow, and not the chemistry of the matrix, that reorients the fibers, the substitution is legitimate for the purpose at hand, and its validity is bounded by the fidelity with which the rheological parameters of the gel are matched to those of the mix it represents.

2.2.3 Concentration Regimes and the Semi-Dilute Limit

The behavior of a suspension of rods is not determined by the volume fraction alone, since a slender rod occupies a volume far smaller than the region it sweeps when it rotates. Figure 2.2(a) illustrates this geometry for a test fiber of length L and diameter d_f, drawn with a thicker line, whose center is c. The fiber occupies a volume $v_f = \pi d_f^2 L/4$, whereas a rotation about c sweeps a sphere of diameter L, shaded in the figure, whose volume $\pi L^3/6$ is of order L^3. The ratio between the two volumes, $(2/3)\, r_f^2$, grows as the square of the aspect ratio $r_f = L/d_f$ and reaches about 2800 for the fibers of this work. Whether a fiber can turn without meeting a neighbor is therefore decided by the overlap of these swept volumes and not by the overlap of the fibers themselves (Doi et al., 1988), and the swept volumes begin to overlap at volume fractions at which the suspension would still be considered extremely dilute by any measure based on volume alone (Martinie et al., 2010).

This observation is made quantitative by counting the neighboring fibers whose centers, marked as dots in Figure 2.2(a), fall within a volume L^3 around c. Let n be the number



(a) (b)

Figure 2.2: Relation between the fiber volume fraction ϕ and the number density n. (a) Test fiber of length L and diameter d_f centered at c, the sphere of diameter L swept by its rotation (shaded), and the centers of the neighboring fibers (dots) within a volume L^3. (b) Number of neighbors within L^3, $nL^3 = (4/\pi)\,\phi\, r_f^2$, against the fiber volume fraction for $r_f = 65$, with the dilute limit $\phi_{sd} = r_f^{-2}$ and the concentrated limit $\phi_c = r_f^{-1}$. N_1, N_2 and N_3 are the fiber contents of the experimental campaign, described in Section 3.6 and Table 3.6, and the points labeled UHPC correspond to practical mixes with 1% and 2% of fibers by volume.

density, the number of fiber centers per unit volume. Since each fiber occupies a volume v_f, the fiber volume fraction is $\phi = n\, v_f$, and the number of neighbors within a volume L^3 around c is

$$n\, L^3 = \phi\, \frac{L^3}{v_f} = \frac{4}{\pi}\, \phi\, r_f^2, \tag{2.3}$$

which shows that the relevant grouping is the product $\phi\, r_f^2$ and not ϕ by itself. The onset of the semi-dilute regime is conventionally placed where this count becomes of order unity, that is, where a fiber first acquires a neighbor within the volume it needs in order to rotate (Doi et al., 1988; Petrie, 1999),

$$\phi_{sd} = \frac{1}{r_f^2}, \tag{2.4}$$

and the upper bound of the regime is placed where the mean separation between fibers falls to the order of the fiber diameter, at $\phi_c \approx 1/r_f$ (Doi et al., 1988; Petrie, 1999). The three resulting regimes are summarized in Table 2.1. The configuration of Figure 2.2(a), with several neighbor centers inside the swept volume, belongs to the semi-dilute regime. Figure 2.2(b) places these regimes on the line of Equation (2.3) for the fibers of this

Table 2.1: Concentration regimes for a suspension of rigid rods of aspect ratio r_f.

Regime	**Condition**	**Physical picture**
Dilute	$\phi < r_f^{-2}$	Fewer than one neighbor within the rotation volume. Each fiber turns as though alone in the fluid, and its motion is fixed by the local velocity gradient.
Semi-dilute	$r_f^{-2} < \phi < r_f^{-1}$	The rotation volumes overlap. Fibers disturb one another through the fluid they displace, but the mean separation stays large compared with the fiber diameter and direct contact is not the governing mechanism.
Concentrated	$\phi > r_f^{-1}$	The mean separation approaches the fiber diameter. Direct contact, excluded volume, and the formation of a contact network govern the response.

Source: (Doi et al., 1988) and (Petrie, 1999).

work, $r_f = 65$, for which $\phi_{sd} = 0.024\%$ and $\phi_c = 1.54\%$. The three fiber contents of the experimental campaign, N_1 to N_3 (Section 3.6), bracket the dilute limit, with between about one and four neighbors within L^3, whereas a practical UHPC mix, with 1 to 2% of fibers by volume, holds some 50 to 110 neighbors and lies at the boundary of the concentrated regime.

2.3 Non-Intrusive Experimental Techniques

The experiment must deliver two quantities of different nature. The first is Eulerian and belongs to the fluid: the velocity field, from which the local rate of strain, the shear rate, and the strain accumulated along a material trajectory are obtained. The second is Lagrangian and belongs to the fibers: the position and the orientation of each rod, from which the orientation state of the population is constructed. The first characterizes what the flow does to the fibers; the second characterizes the state the fibers reach as a result.

Two complementary techniques are employed. PIV resolves the velocity field of the fluid by correlating the collective displacement of tracer particles seeded within the gel, supplying the kinematics of the flow. PTV resolves the trajectory and the orientation of each fiber individually, supplying the state of the fiber population. Figure 2.3 illustrates

both configurations as applied to the L-shape geometry discussed in Section 2.4.

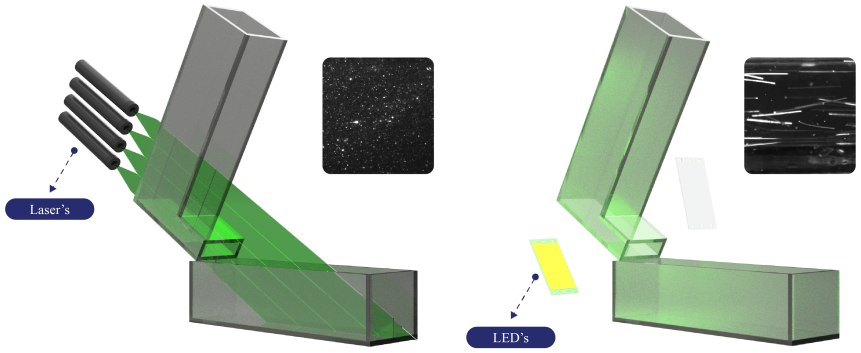


Figure 2.3: Experimental setups: (left) PIV configuration with planar laser illumination and tracer particles seeding the Carbopol solution for velocity field measurement; (right) PTV configuration with volumetric LED illumination for tracking individual fiber motion in three dimensions.

2.3.1 Particle Image Velocimetry

PIV measures instantaneous velocity fields by tracking the displacement of neutrally-buoyant tracer particles between consecutive image pairs (Adrian, 1991; Raffel et al., 2018). A laser sheet illuminates a planar cross-section of the flow, and a high-speed camera captures the scattered light from the suspended particles at controlled time intervals Δt. Cross-correlation algorithms applied to small sub-regions of the image, called interrogation windows, determine the most probable particle displacement vector within each region, yielding a two-dimensional velocity field $\mathbf{u}(x, y, t)$ whose spatial resolution is set by the window size and overlap, and whose temporal resolution is determined by the camera frame rate (Raffel et al., 2018). Figure 2.4 shows the workflow.

Multi-Pass Cross-Correlation

The core of any PIV algorithm is the normalized cross-correlation function between two interrogation windows I_1 and I_2 extracted from consecutive frames,

$$C(\mathbf{s}) = \frac{\iint I_1(\mathbf{X})\, I_2(\mathbf{X}-\mathbf{s})\, \mathrm{d}\mathbf{X}}{\sqrt{\iint I_1^2(\mathbf{X})\, \mathrm{d}\mathbf{X} \iint I_2^2(\mathbf{X})\, \mathrm{d}\mathbf{X}}}, \tag{2.5}$$

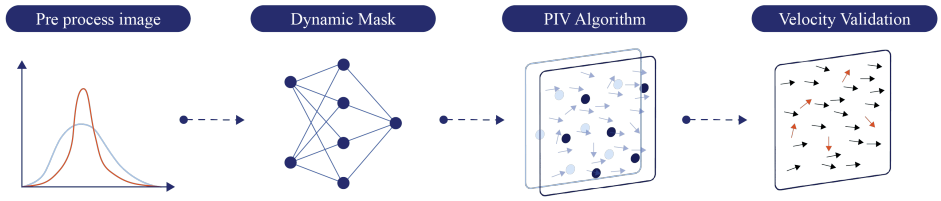


Figure 2.4: PIV workflow: (i) image pre-processing (background subtraction and contrast enhancement), (ii) dynamic masking (Carbopol boundary detection), (iii) multi-pass cross-correlation for velocity-vector calculation, and (iv) post-processing (outlier removal and interpolation).

where $\mathbf{s}$ is a trial displacement and $\mathbf{X}$ denotes the spatial coordinates within the window. The peak of $C(\mathbf{s})$ gives the integer-pixel displacement, and sub-pixel accuracy is recovered by fitting a two-dimensional Gaussian to the neighborhood of the peak (Westerweel, 1997; Raffel et al., 2018).

A multi-pass strategy with successively decreasing window sizes is adopted to improve spatial resolution and reduce measurement bias (Scarano, 2002; Raffel et al., 2018). In the first pass, a coarse window (for example, 128×128 px) yields a reliable first estimate of the displacement field, which is then used as a predictor offset for the next, finer pass. This hierarchical approach allows the algorithm to handle a broader dynamic range of velocities while converging toward the finest spatial scales accessible with the seeding density available.

Error Analysis

Three independent error sources are quantified for each PIV measurement. The sub-pixel displacement error arises from the finite precision of the Gaussian peak fit to the correlation maximum. Its magnitude depends on the diameter of the particle image, and it is bounded by $\varepsilon_{\text{px}} \approx 0.1$ px when that diameter spans two pixels or more; below that value the fit collapses toward integer displacements and the error becomes a systematic bias rather than a random one (Westerweel, 1997; Raffel et al., 2018). It propagates into velocity space as

$$\delta u_{\text{subpx}} = \frac{\varepsilon_{\text{px}}\, S_c}{\Delta t}, \tag{2.6}$$

where $S_c\,[\text{mm px}^{-1}]$ is the spatial scale factor and $\Delta t\,[\text{s}]$ is the inter-frame time. The calibration error stems from the uncertainty δr_{px} in the pixel span r_{px} of a reference

target of known physical length ℓ_{ref}, giving $\delta S_c = (\ell_{\mathrm{ref}}/r_{\mathrm{px}}^2)\,\delta r_{\mathrm{px}}$, which propagates as

$$\delta u_{\mathrm{cal}} = \frac{|\overline{\Delta x}_{\mathrm{px}}|\,\delta S_c}{\Delta t}, \tag{2.7}$$

where $\overline{\Delta x}_{\mathrm{px}} = \bar{u}\,\Delta t/S_c$ is the mean particle displacement in pixels. The velocity-gradient bias error arises when non-uniform flow within an interrogation window of physical size d_I broadens and shifts the correlation peak; following (Westerweel, 2008), its leading-order term is

$$\varepsilon_{\mathrm{grad}} = \frac{1}{24}\,d_I^2\left|\frac{\partial^2 u}{\partial x^2}\right|, \tag{2.8}$$

where the second derivative is estimated by central differences from the measured velocity field. The three contributions are of different nature and are combined accordingly: the two random terms are added in quadrature, since they are statistically independent, and the systematic bias is added linearly, since it does not average out over repeated measurements (Coleman et al., 2009),

$$\delta u_{\mathrm{total}} = \sqrt{\delta u_{\mathrm{subpx}}^2 + \delta u_{\mathrm{cal}}^2} \;+\; |\varepsilon_{\mathrm{grad}}|. \tag{2.9}$$

As a complementary diagnostic, the criterion of Keane and Adrian (Keane et al., 1992) is applied at each vector location, requiring that the displacement variation across the window satisfy $M\,|\Delta u|\,\Delta t/d_I < 0.03$, where M is the optical magnification of the imaging system; vectors exceeding this threshold are flagged as potentially invalid.

The gradient bias of Equation (2.8) requires particular attention in a yield-stress fluid, for a reason that is specific to the material rather than to the technique. The deformation in such a fluid is not distributed but concentrated in the sheared layers that separate the plug from the walls, so the velocity gradient is small over most of the domain and large within a narrow band. The second derivative in Equation (2.8) is largest within that band, which is where the interrogation window is least able to resolve the field it is averaging. The bias is therefore not uniform across the image. It concentrates in the regions where the strain is being accumulated, which are the regions the measurement most needs to resolve, and for this reason its magnitude is estimated locally rather than reported as a single figure for the field as a whole.

From the Velocity Field to the Strain History

Differentiating $\mathbf{u}(x, y, t)$ with respect to the spatial coordinates gives the velocity gradient $\nabla\mathbf{u}$, whose symmetric part is the rate-of-strain tensor $\mathbf{D}$ and whose antisymmetric part is the rotation tensor $\mathbf{W}$,

$$\mathbf{D} = \tfrac{1}{2}\left(\nabla\mathbf{u} + \nabla\mathbf{u}^{\mathsf{T}}\right), \qquad \mathbf{W} = \tfrac{1}{2}\left(\nabla\mathbf{u} - \nabla\mathbf{u}^{\mathsf{T}}\right). \tag{2.10}$$

The rotation tensor should be distinguished from the vorticity, which is the vector $\boldsymbol{\zeta} = \nabla \times \mathbf{u}$. The two carry the same information, because an antisymmetric second-order tensor has only three independent components, and with the definitions of Equation (2.10) they are related by

$$W_{ij} = -\tfrac{1}{2}\,\epsilon_{ijk}\,\zeta_k, \qquad \mathbf{W}\cdot\mathbf{a} = \tfrac{1}{2}\,\boldsymbol{\zeta}\times\mathbf{a}, \tag{2.11}$$

for any vector $\mathbf{a}$, where ϵ_{ijk} is the Levi-Civita permutation symbol. The components of $\mathbf{W}$ are therefore minus one half of those of the vorticity, and a fluid element rotates as a rigid body at half the local vorticity. In the plane of the PIV measurement only the out-of-plane component $\zeta = \partial v/\partial x - \partial u/\partial y$ is resolved, and $W_{12} = -W_{21} = -\zeta/2$. From $\mathbf{D}$ follows the scalar shear rate $\dot\gamma = \sqrt{2\,\mathbf{D}:\mathbf{D}}$ at every point and instant of the flow (Macosko, 1994). Integrating $\dot\gamma$ along the pathlines reconstructed from the same field yields the strain accumulated by a material element as it travels through the geometry,

$$\Gamma = \int_0^t \dot\gamma\left(\mathbf{x}(t'),\, t'\right)\,\mathrm{d}t', \tag{2.12}$$

which is a Lagrangian quantity, attached to the material and not to the point through which it happens to pass. The accumulated strain is the variable against which the orientation of the fibers is examined in Chapter 5, and Equation (2.12) is how the experiment supplies it.

Three limitations bound this reconstruction. The first is that PIV in the planar configuration resolves two components of the velocity in a single plane, so the out-of-plane contributions to $\mathbf{D}$ are not measured and $\dot\gamma$ is evaluated from the in-plane components alone. Since the omitted terms enter the double contraction $\mathbf{D}:\mathbf{D}$ as squares, their exclusion lowers the value of $\dot\gamma$, so the planar estimate is a lower bound on the true

shear rate rather than an unbiased estimate of it (Raffel et al., 2018). The measurement plane is therefore chosen to contain the principal direction of the flow, where the out-of-plane contribution is smallest, and the resulting bias is assessed rather than assumed to be negligible. The second is that Equation (2.12) requires the trajectory and not merely the instantaneous field, so its accuracy depends on the temporal resolution with which the field is sampled, and a pathline integrated from a sparsely sampled sequence accumulates the integration error along with the strain. The third is that Γ is a scalar, so it records the magnitude of the deformation but not the direction in which it was applied. Two material elements that accumulate the same Γ under velocity gradients oriented differently are not equivalent, and this restricts the interpretation of Γ as a control variable to flows in which the direction of the deformation does not change substantially along the path.

2.3.2 Particle Tracking Velocimetry

While PIV provides Eulerian velocity fields averaged over interrogation windows, PTV tracks individual particles in a Lagrangian frame, resolving the motion of each particle with high temporal resolution (Maas et al., 1993; Ouellette et al., 2006). This distinction is essential when the particle behavior deviates from the bulk fluid motion owing to anisotropic geometry, inertia, or preferential alignment, as is the case for the steel fibers studied in this work. A fiber is not a tracer. It has a length and an orientation, and the orientation is the quantity of interest, so a technique that returns a velocity field averaged over a window cannot supply it. The PTV workflow consists of three core stages: particle detection and segmentation, feature extraction of the centroid, orientation, and length of each fiber, and particle linking across frames to reconstruct trajectories. Unlike the neutrally-buoyant tracer particles used in PIV, PTV tracks the steel fibers themselves through their light-reflective properties. Figure 2.5 illustrates the workflow, from image preprocessing through fiber detection, feature extraction, and trajectory reconstruction.

Fiber Detection and Segmentation

The detection stage is the most critical determinant of PTV performance, since a fiber that is not segmented in a given frame cannot be tracked. Conventional image-

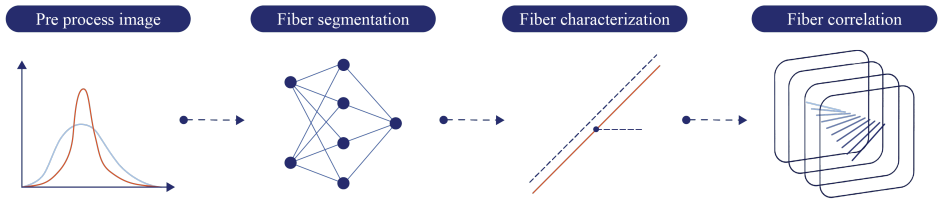


Figure 2.5: PTV workflow, from image preprocessing (background subtraction and contrast enhancement) to fiber detection, feature extraction, and trajectory reconstruction.

processing techniques, such as the Hough transform, intensity thresholding, and Canny edge detection, rely on manually tuned parameters and deterministic geometric assumptions, which make them sensitive to noise, particle overlap, and illumination variability (Plaksyvyi et al., 2023). These limitations are particularly severe for steel fibers in a translucent gel: their reflected intensity varies along the fiber length depending on orientation and surface properties, their projected width approaches the resolution limit of the imaging system, on the order of two pixels, and inter-particle occlusions become increasingly frequent as fiber concentration rises. Under these conditions, classical methods either miss a large fraction of the fibers or generate numerous false detections.

The bias introduced by a missed detection is not random, and this is what makes the detection stage a measurement problem rather than a preprocessing step. A fiber lying at an unfavorable angle to the illumination reflects less light toward the camera than one lying favorably, so the fibers a detector fails to find are not a random sample of the population but are drawn preferentially from certain orientations. A detection rate below unity therefore distorts the orientation distribution itself, and it distorts it toward the orientations that are easiest to see. Any metric computed from the detected population inherits that bias, including the order parameter of Section 2.3.3, which is why the segmentation stage is treated here on the same footing as the quantities derived from it.

To overcome these limitations, deep learning-based instance segmentation has emerged as a robust alternative, capable of detecting elongated, low-contrast features under noisy and non-uniformly illuminated conditions (Qamar et al., 2024). Instance-level architectures assign a distinct mask to each detected object, which is a requirement for separating individual fibers within dense clusters. Two families of models are relevant in this context: single-stage detectors such as YOLOv11 (Khanam et al., 2024), which

integrate boundary delineation and classification within a unified network optimized for real-time inference, and prompt-driven foundation models such as the Segment Anything Model (Ravi et al., 2025), pretrained on large-scale generic datasets to generalize across imaging domains. The relative performance of these approaches for rod-like particle detection, together with their adaptation to the specific optical conditions of this study, is developed in the methodology (Section 3.5).

Particle Correlation

Once fibers are detected and characterized, they must be linked across consecutive frames to reconstruct their trajectories. This correspondence problem becomes increasingly ambiguous as fiber concentration rises, since multiple fibers within a small spatial neighborhood may present similar geometric features, and temporary occlusions interrupt otherwise continuous trajectories (Jaqaman et al., 2008). A fixed-radius nearest-neighbor search is inadequate under these conditions, because fibers accelerate, rotate, and follow curved paths, violating the assumption of uniform inter-frame motion (Crocker et al., 1996).

To address this, predictive tracking schemes estimate the expected state of each fiber in the subsequent frame from its kinematic history, constraining the search region and reducing correspondence ambiguity. Filters of the Kalman family, and their simplified α-β-γ variant, are suited to this task, as they simultaneously estimate position, velocity, and acceleration while suppressing high-frequency segmentation noise (Tenne et al., 2000; Welch et al., 1995). The combination of deep learning-based instance segmentation with such a predictive filter constitutes the core of the PTV framework employed in this work. Its complete formulation, parameter selection, and validation are presented in the methodology and in the companion study where the framework was originally developed (Wolff et al., 2026).

2.3.3 Order and Dispersion Metrics

The output of the PTV stage is a list of fibers, each carrying a centroid $\mathbf{x}_i$ and an in-plane angle θ_i, within a given region of interest and at a given instant. This list is not itself a result. It must be reduced to a small number of quantities that summarize how the fibers are oriented and how they are arranged in space, and that can be

compared across regions, across dosages, and against the prediction of a model. Two complementary families of metrics serve this purpose, one describing the orientation of the population and one describing its spatial arrangement.

Orientational Order

The orientation of a fiber is defined modulo 180°, since a rod has no distinguishable head or tail and the angles θ and $\theta + 180°$ describe the same physical state. Any summary of the orientation of a population must respect this symmetry, and an ordinary arithmetic mean of the angles does not. A set of fibers all lying horizontally, half of them recorded as pointing left and half as pointing right, would average to a direction that no fiber occupies and would be reported as disordered when it is perfectly aligned. The symmetry is enforced by doubling the angle before averaging, which maps θ and $\theta + 180°$ onto the same point of the unit circle (Mardia et al., 2000), and the resulting quantity is the nematic order parameter,

$$\eta_1 = \left|\left\langle e^{i2\theta} \right\rangle\right| = \sqrt{\langle \cos 2\theta \rangle^2 + \langle \sin 2\theta \rangle^2}, \tag{2.13}$$

where the angular brackets denote the average over the fibers detected within the region (de Gennes et al., 1993). The parameter η_1 is bounded between zero, corresponding to an isotropic distribution of orientations with no preferred direction, and one, corresponding to perfect alignment of all fibers along a common axis. It is the planar reduction of the second-order orientation tensor of Advani and Tucker (Advani et al., 1987), and its identification with the scalar order parameter of that framework is established in Chapter 5. Here it is introduced as a measured quantity, computed directly from the angles returned by the tracking stage.

The same complex average supplies the direction about which the population is ordered, the local director, obtained from the argument rather than the modulus and halved to undo the doubling,

$$\theta_{\mathrm{dir}} = \tfrac{1}{2} \arg \left\langle e^{i2\theta} \right\rangle = \tfrac{1}{2} \operatorname{atan2}(\langle \sin 2\theta \rangle \,,\ \langle \cos 2\theta \rangle) \,. \tag{2.14}$$

While η_1 summarizes the strength of the orientational order and θ_{dir} its direction, neither describes how the individual fibers are distributed about that direction. For this purpose

an angular deviation δ is defined for each fiber with respect to the local director,

$$\delta = \left|((\theta - \theta_{\text{dir}} + 90^\circ) \bmod 180^\circ) - 90^\circ\right|, \tag{2.15}$$

where the modular wrapping accounts for the same 180° periodicity used in the definition of η_1. This formulation bounds δ between 0°, indicating perfect alignment with the local director, and 90°, indicating an orientation perpendicular to it, and provides a per-fiber measure of orientational coherence that complements the population-level summary given by η_1. The distribution of δ across the fibers of a region distinguishes two situations that a single value of η_1 cannot: a population in which every fiber deviates moderately from the director, and one in which most fibers lie close to it while a minority lies far from it.

Spatial Dispersion

Beyond orientation, the spatial arrangement of fibers within a given region provides additional information on flow-induced clustering or segregation. A common approach in point pattern analysis is to compare an observed spatial distribution against the null hypothesis of complete spatial randomness, under which points are assumed to be uniformly distributed across the domain of interest (Baddeley et al., 2016). Departures from this null hypothesis indicate either clustering, when observed points are more concentrated than expected under uniformity, or regularity, when points are more evenly spread than expected (Diggle, 2013). The statistic adopted here is the radius of gyration of the fiber centroids within a region,

$$R_g = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \|\mathbf{x}_i - \bar{\mathbf{x}}\|^2}, \tag{2.16}$$

where $\mathbf{x}_i$ is the centroid of fiber i, $\bar{\mathbf{x}}$ is the mean centroid position over the N fibers in the region, and $\|\cdot\|$ denotes the Euclidean norm. Taken alone, R_g is not interpretable, since it grows with the size of the region regardless of how the fibers are arranged within it, and it must therefore be referred to the value that a uniform filling of the same region would produce. For a rectangular region of width w and height h this reference value

follows from the variance of a uniform distribution along each side,

$$R_g^{\text{unif}} = \sqrt{\frac{w^2 + h^2}{12}}, \tag{2.17}$$

and is computed independently for each geometry considered, so that the resulting index is comparable across regions of different size and shape without being biased by the geometry of the region itself.

The comparison requires one further correction, of the same nature as the one applied to η_1. Equation (2.16) measures the dispersion about the sample mean $\bar{\mathbf{x}}$ rather than about the true center of the region, and the sample mean is itself displaced toward the points it summarizes, so a finite uniform sample yields a radius of gyration systematically smaller than Equation (2.17)

$$\Lambda = \frac{R_g}{R_g^{\text{unif}}} \sqrt{\frac{N}{N-1}}, \tag{2.18}$$

which is unity in expectation for a uniform filling of the region, falls below unity when the fibers occupy a smaller effective area than expected and are therefore clustered, and rises above unity when they are displaced toward the periphery of the region.

Two limitations of Equation (2.18) should be recorded, since together they bound the conclusions that may be drawn from it. The first is that the radius of gyration is a single-scale summary of the point pattern, so it does not distinguish a single dense cluster from several smaller ones dispersed across the region, and both may return the same value of Λ. The index is therefore used here to detect the presence of clustering and to compare its severity across regions and dosages, and not to characterize its structure, for which a scale-resolved statistic such as the Ripley K function would be required (Baddeley et al., 2016; Diggle, 2013). The second is that Λ is computed from the fibers the detector returns and not from the fibers present, so the orientation-dependent detection bias described in Section 2.3.2 propagates into the spatial statistic as well, and a region whose fibers are predominantly oriented unfavorably with respect to the illumination will yield a point pattern thinned in a manner that is not spatially uniform.

2.4 L-Shape Model

The flow must be produced in a geometry that satisfies three requirements simultaneously. It must impose on the fibers a kinematic history rich enough to be informative, which is to say a sequence of flow conditions that varies along the path rather than a uniform shear that every fiber experiences identically. It must be optically accessible along the whole of that path, so that the velocity field and the fiber orientations can be measured wherever the material goes. It must also be reproducible, so that the effect of varying the fiber dosage can be separated from the effect of varying the flow.
The L-shape geometry proposed by (Huang et al., 2018) satisfies these three requirements. It induces preferential fiber alignment at the outlet through the change in flow direction and the velocity gradients that accompany it, and it provides a scenario in which the effect of the rheological properties and of the fiber concentration on the final orientation distribution can be examined under conditions representative of casting. Figure 2.6 illustrates the model, which consists of a vertical inlet channel transitioning into a horizontal outlet channel through a ninety-degree bend, with the entire container tilted thirty degrees from vertical so that the flow is driven by gravity alone and no pump or piston introduces a disturbance of its own.

2.4.1 Optical Configuration of the Mold

The mold is built from transparent acrylic, and a matte black PLA sheet is inserted within the container, behind the volume through which the flow is observed, so that it forms the rear boundary of the fluid. The reason is optical rather than structural. Both techniques of Section 2.3 operate in dark-field, the tracer particles of the PIV and the steel fibers of the PTV being visible because they return the incident light to the camera while everything else in the scene is background, and a transparent mold placed in front of a bright or structured background transmits that background into the image, raising its mean gray level and reducing the contrast on which both measurements rest. The black sheet removes it, since light that is not returned by a particle or a fiber is absorbed rather than transmitted, and the matte finish rather than a polished one suppresses the specular reflections that the walls and the surrounding room would otherwise return as stationary bright artifacts. What this preserves is not merely the appearance of the

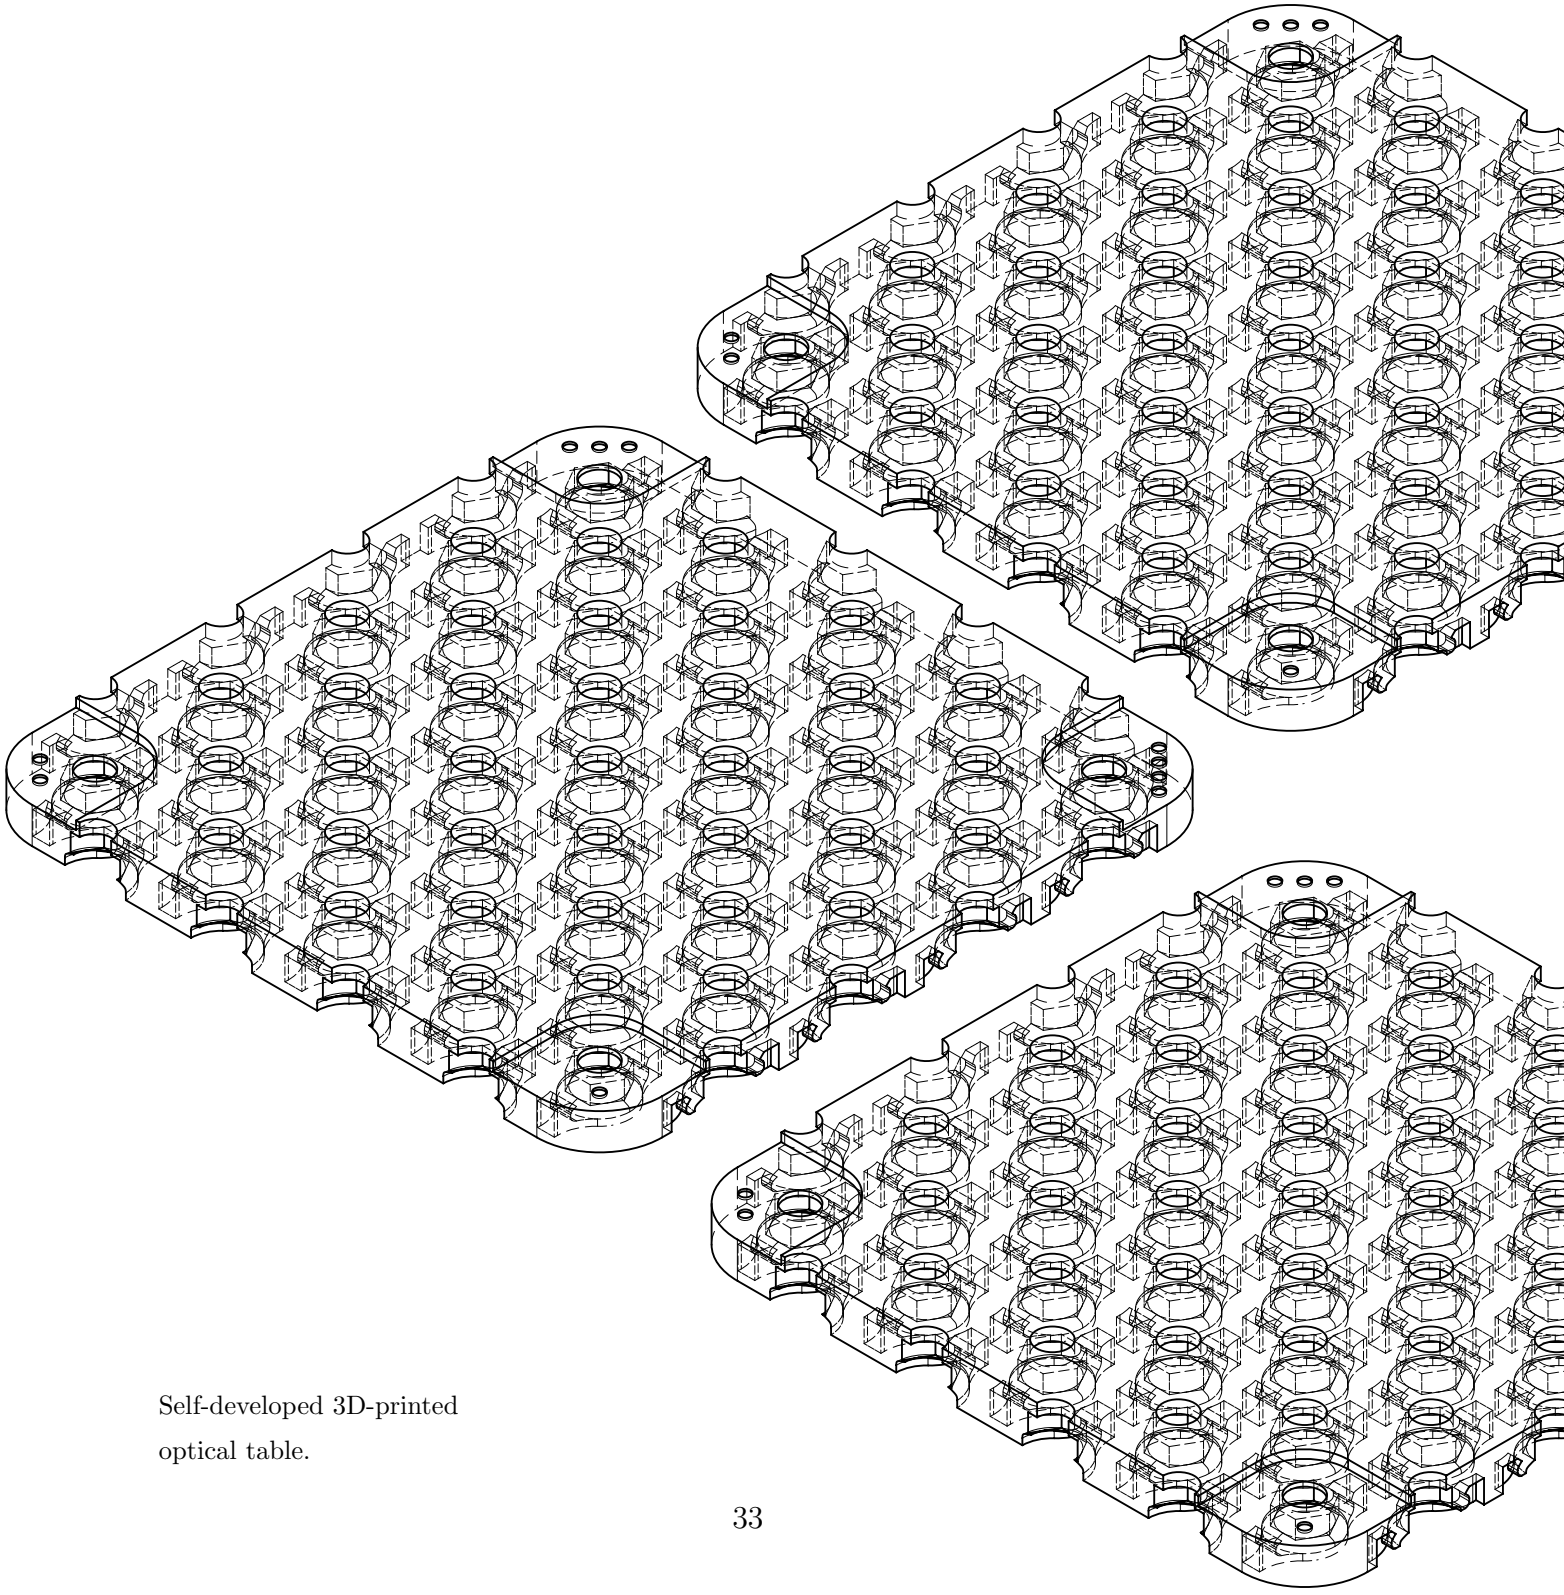


Figure 2.6: L-shape benchmark geometry showing the vertical inlet channel, ninety-degree bend, and horizontal outlet channel. The container is tilted thirty degrees from vertical to induce flow via gravity, and the rear face is finished in matte black to provide a dark-field background for the optical measurements.

images but the quantities extracted from them: the correlation of Equation (2.5) is computed from the intensity fluctuation about the local mean, so a bright background enters its denominator without contributing to its peak, while the detection of a fiber depends on the contrast it holds against its surroundings, and steel fibers are specular reflectors whose apparent brightness falls with their inclination to the light source.

Chapter 3

Methodology



Self-developed 3D-printed optical table.

This chapter describes the experiment through which the deformation delivered to a fiber and the alignment it acquires are measured on the same flow and at the same instant. Section 3.1 presents the multi-camera arrangement and the optical table on which it is mounted, and records the resolution at which a fiber is seen, from which the uncertainty of every quantity reported in this work descends. Section 3.2 specifies the steel fibers and establishes the concentration bounds their aspect ratio fixes. Section 3.3 gives the mixing protocol of the transparent analog and its rheological characterization, and locates the flow with respect to the yield stress, which is what confines the deformation to a small fraction of the domain. Sections 3.4 and 3.5 describe the implementation of the two optical techniques, the first supplying the accumulated strain against which the saturation law is written and the second the orientation the law predicts, together with the measurement problems each poses in a transient, viscoplastic flow carrying specularly reflecting rods. Section 3.6 sets out the campaign, which varies the fiber content at fixed geometry and fixed rheology, and states the dosages at which the hypothesis is put to the test and the bounds those dosages impose on what it can establish.

3.1 Experimental Setup

A new experimental setup was designed from scratch based on the configuration proposed by (Maggi, 2023), with modifications aimed at improving measurement repeatability, image quality, and spatial resolution. A multi-camera system was implemented to simultaneously capture both global flow patterns and localized high-speed phenomena. Figure 3.1 illustrates the camera distribution around the L-shape and beam containers. Camera 1 captures the region of interest (ROI) of the L-shape container to track overall flow progression and fiber distribution, while Cameras 2 and 3 record the beam container from orthogonal angles. Camera 4 was positioned to capture a close-up view of the L-shape outlet at higher frame rate, resolving the rapid fiber reorientation occurring during the transition from vertical to horizontal flow. Table 3.1 summarizes the specifications of each camera.

All cameras were equipped with 24 mm $f/1.8$ lenses. The wide aperture ($f/1.8$) enables short exposure times, reducing motion blur, at the cost of reduced depth of field. The

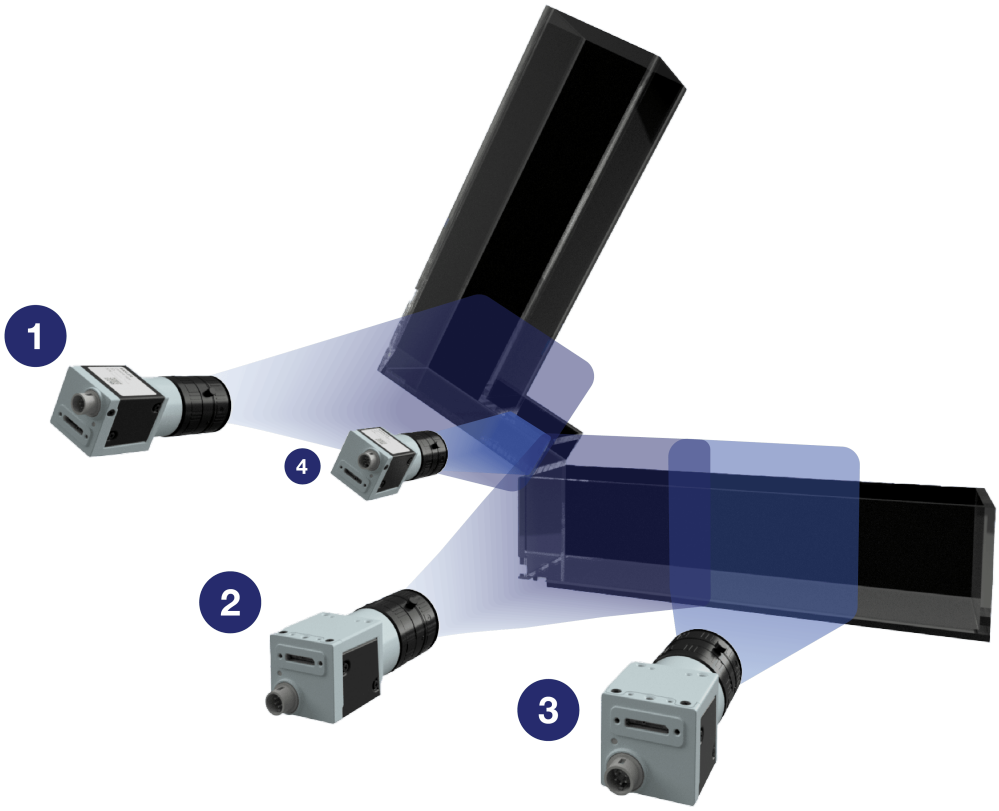


Figure 3.1: Multi-camera setup for simultaneous PIV and PTV measurements. Cameras 1–3 provide wide-field coverage of the L-shape container (Cam 1) and beam container from orthogonal views (Cams 2–3). Camera 4 captures high-speed close-up footage of the L-shape outlet region to resolve initial flow development.

shallow depth of field is acceptable because the planar laser sheet for PIV and the narrow measurement volume for PTV ensure that only features within the focal plane are illuminated and contribute to the measurements. Because the analysis pairs the fluid and fiber fields at a common instant, the cameras were synchronized in hardware rather than by software timestamp: an external trigger signal drove the exposure of every camera, and the signals were aligned on an oscilloscope so that all cameras were triggered from a common time base. This ensures that a given frame index corresponds to the same physical instant across all views, which is the condition under which the velocity field measured by one camera and the fiber state measured by another can be compared at the same point in time.

The scale factors of Table 3.1 fix the resolution at which a fiber is seen, and it is worth

Table 3.1: Camera specifications and operating parameters.

Camera	**Frame Rate (fps)**	**Resolution (pixels)**	**Scale Factor (mm/px)**
1	220	1024 × 1024	0.1250
2	220	1024 × 1024	0.1282
3	220	1024 × 1024	0.1282
4	660	384 × 384	0.0935

recording what they imply, since the uncertainty of every quantity reported in this work descends from them. At 0.125 $\mathrm{mm\,px^{-1}}$, a fiber of 13 mm length and 0.2 mm diameter projects onto approximately 104 px along its axis and 1.6 px across it. The fiber is therefore well resolved in length and barely resolved in width, which is the condition that makes its orientation recoverable to a fraction of a degree while making its detection sensitive to illumination, and it is the reason the detection stage is treated in Section 3.5 as a measurement problem rather than as a preprocessing step. The distances from which these factors follow, together with the dimensions of the two containers, are given in Appendix D.

Figure 3.2 shows a render of the complete experimental setup, and further views are given in Appendix D.3. A custom-designed optical table was fabricated to mount the cameras, laser, LED panels, and acrylic containers, providing precise positioning control through a modular grid of threaded mounting points. This rigid platform ensures consistent component alignment between experiments, and it does more than improve repeatability.

3.2 Steel Fibers

The steel fibers used in all experiments were Dramix® OL 13/.20 (Bekaert), a commercially available straight steel fiber designed for fiber-reinforced concrete applications. Their physical and mechanical properties are summarized in Table 3.2, and a size reference relative to a Chilean 10 peso coin is shown in Figure 3.3.

The aspect ratio $r_f = 65$ is the parameter that fixes the concentration regime, and it does so through the bounds established in Section 2.2.3. The dilute limit falls at $\phi_{sd} = r_f^{-2} = 0.024\%$ by volume and the concentrated limit at $\phi = r_f^{-1} = 1.54\%$.

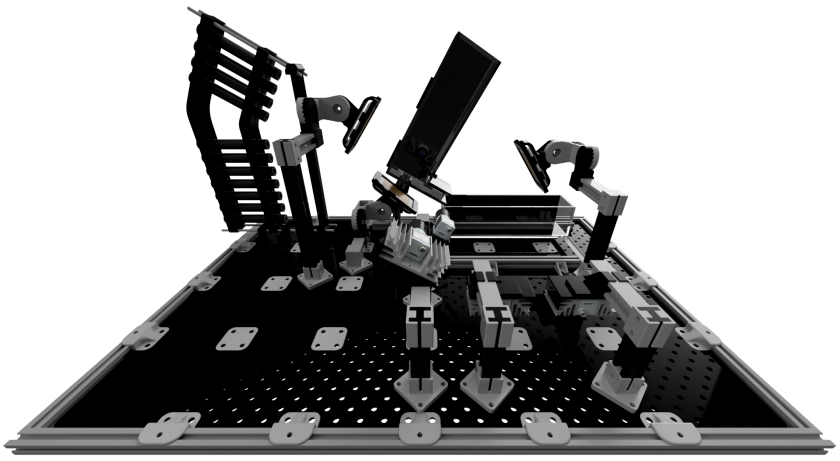


Figure 3.2: Render of the experimental setup, showing the multi-camera system, laser sheet configuration, and the acrylic containers for the L-shape and beam tests. The optical table allows precise positioning and alignment of all components.

Table 3.2: Physical and mechanical properties of the Dramix® OL 13/.20 steel fibers used in the experiments. The numerical density is the manufacturer's value and is the conversion factor by which fibers are counted by weighing.

Property	**Value**	**Unit**
Total length (L)	13	mm
Diameter (d_f)	0.2	mm
Aspect ratio ($r_f = L/d_f$)	65	–
Elastic modulus	200	GPa
Tensile strength	2750	MPa
Numerical density	282 556	fibers/kg

3.3 Carbopol Mixing Protocol

The preparation of the Carbopol Ultrez 10 polymer (Lubrizol) solutions followed the procedure described by (Tapia, 2026). Two concentrations were prepared, 0.2% and 0.5% (carbopol/water), differing in yield stress by approximately a factor of two. The pair was chosen for the model the experiment is designed to validate: a CFD-DEM simulation is not verified by reproducing a single flow, since a model can be tuned to any one of them, but by reproducing the change a flow undergoes when a known property of the matrix is altered, and the yield stress is one of the properties that governs the kinematics of a casting. The two formulations therefore supply a contrast



Figure 3.3: Fiber size relative to a Chilean 10 peso coin (diameter of 21 mm).

in the one parameter against which the constitutive law of the simulated matrix must be checked. For each batch, deionized water was first measured and its initial pH verified to be approximately 7.0. The corresponding amount of Carbopol powder was then gradually added over one minute while mixing with a mechanical stirrer at 150 rpm. Stirring was maintained for one hour to ensure complete and homogeneous dispersion of the polymer.

Following dispersion, the pH was measured at five points throughout the mixture. Subsequently, the first dose of 1M NaOH solution was added and the mixture was stirred for 30 minutes. After this period, the pH was measured again and the second dose of 1M NaOH was added, followed by an additional 30 minutes of stirring. Finally, the pH was measured at five new points to confirm that the average value fell within the accepted range. The specific quantities of water, Carbopol, and NaOH for each concentration, along with the accepted final pH range, are summarized in Table 3.3. The resulting gel was stored in a sealed container at 20 °C for 24 hours to ensure complete hydration and structural stabilization.

Table 3.3: Preparation parameters for Carbopol Ultrez 10 solutions at 0.2% and 0.5% concentrations. NaOH additions are sequential: the primary dose initiates gelation and the secondary dose achieves the target pH range.

Carbopol	**Water (mL)**	**Carbopol (g)**	**$NaOH_1$ (mL)**	**$NaOH_2$ (mL)**	**Target pH**
0.2%	1996	4.0	26	3	6.8 – 7.2
0.5%	1990	10.0	70	20	6.8 – 7.2

Once stabilized, the gel was centrifuged in a Hettich Universal 320R centrifuge at 2000 rpm and 20 °C for 10 minutes to remove entrapped air bubbles, yielding a clear

and optically homogeneous gel suitable for flow visualization experiments. The removal of bubbles is not cosmetic: an elongated bubble in a translucent gel presents to the detector of Section 3.5 an object of the same aspect ratio and comparable contrast as a fiber, and is a source of false detections that no amount of training can wholly eliminate. Each prepared batch was used for a maximum of two experimental runs, and the same protocol was applied to both PIV and PTV measurements. The pH recorded at each of the three stages of the protocol, for every batch of the campaign, is given in Appendix C.

3.3.1 Rheological Characterization

The rheological behavior of both Carbopol formulations is described through the Herschel-Bulkley model introduced in Section 2.2.1, which requires three parameters: the yield stress τ_0, the consistency index K, and the flow index n. Their determination is complicated by wall slip, a phenomenon in which a thin, polymer-depleted layer forms against smooth surfaces and allows the gel to slide before it yields in bulk. On the smooth walls of a conventional rheometer this effect systematically underestimates the yield stress of Carbopol suspensions (Younes et al., 2020), so a measurement strategy that avoids it is required for the present work.

To obtain a slip-free estimate of the yield stress, τ_0 was determined from the geometry of an arrested free-surface flow rather than from rotational rheometry (Coussot et al., 1996). When a yield-stress fluid is left to spread under gravity, it comes to rest once the gravitational driving stress can no longer overcome the yield resistance, so that the shape of the arrested deposit encodes τ_0 directly and the measurement is made on the bulk material rather than at a wall. Applying the thin-film lubrication balance to the arrested layer (Balmforth et al., 2007) gives

$$\tau_0 = \tfrac{1}{2}\,\rho\, g\, h_{\max}\, \tan\theta_\infty, \tag{3.1}$$

where ρ is the fluid density, g the gravitational acceleration, $h_{\max}$ the maximum height of the arrested fluid, and θ_∞ the inclination of the free surface at rest, the prefactor following from the balance between the gravitational driving stress and the yield resistance integrated over the depth of the layer. This procedure yields $\tau_0 = 40.6$ Pa for the

0.2% formulation and $\tau_0 = 84.0$ Pa for the 0.5% formulation.

With the yield stress fixed by the deposit method, the consistency index K and the flow index n were obtained by refitting steady-state flow-curve data to the Herschel-Bulkley model while holding τ_0 at its slip-free value, so that the shear-thinning branch of the response is captured without reintroducing the bias associated with wall slip. The complete set of parameters for both formulations is summarized in Table 3.4.

Table 3.4: Herschel-Bulkley parameters for the two Carbopol Ultrez 10 formulations. The yield stress τ_0 was determined from the arrested free-surface deposit to avoid wall-slip underestimation; the consistency index K and flow index n were refitted from flow-curve data with τ_0 held fixed. The critical shear rate $\dot{\gamma}_c = (\tau_0/K)^{1/n}$ separates the yield-dominated from the viscous-dominated response.

Carbopol	τ_0 **(Pa)**	K **(Pa·s^n)**	n **(–)**	$\dot{\gamma}_c$ **(s^{-1})**
0.2%	40.6	7.70	0.532	22.8
0.5%	84.0	23.96	0.442	17.1

Source: Research group.

The flow index n below unity in both cases confirms the shear-thinning character of the two gels, while the higher yield stress and consistency of the 0.5% formulation reflect its denser polymer network. The critical shear rate deserves particular attention, because it locates the flow with respect to the mechanism the hypothesis rests on. Below $\dot{\gamma}_c$, which is of the order of 20 s^{-1} in both formulations, the local stress remains close to τ_0 and the material behaves as a near-plug that translates without deforming; above it, the power-law term dominates and the material shears.

3.4 PIV Implementation

The role of PIV in this work is to supply the independent variable of the saturation law. Differentiated in space, the velocity field gives the rate of strain and the shear rate, and integrated along the trajectories it defines, it gives the strain a material element has accumulated by the time it reaches a given station of the mold. It is measured on every run, and not only on the fiber-free ones, because the strain against which the orientation is read must be the strain that same suspension delivered, and a mix carrying more fibers advances more slowly than one carrying fewer.

Two features of the present flow place demands on the technique that a steady, Newtonian flow would not. The first is that the flow is transient: the material accelerates as it enters, reaches its maximum velocity in the bend, and decelerates as it spreads along the outlet until it arrests, so that a correlation interval appropriate to the first phase produces displacements too small to correlate in the last. The second is that the flow is viscoplastic: the deformation is not distributed through the domain but concentrated in the sheared layers that separate the plug from the walls, so the velocity gradient is small over most of the field and large within a narrow band. Both features are addressed below, the first through a block-analysis strategy and the second through a local, rather than global, treatment of the gradient bias.

As described in Section 2.3.1, the PIV algorithm employs cross-correlation to estimate velocity fields by tracking tracer particle displacements between consecutive image pairs using interrogation windows. Table 3.5 summarizes the interrogation window parameters used for each camera and Carbopol concentration. Cameras 1 and 2 required a larger final window of 32 px for the Carbopol 0.2% case due to the higher flow velocities recorded in those views, which would otherwise cause particles to exit smaller windows between frames. For all remaining configurations a final window of 16 px was used, providing finer spatial resolution consistent with the lower velocities and the more viscous character of Carbopol 0.5%.

Table 3.5: PIV interrogation parameters by camera and Carbopol concentration. Window sizes follow a three-pass multi-pass scheme; the final window size determines spatial resolution.

Camera	**Carbopol [%]**	**Window sizes [px]**
1	0.2	$128 \rightarrow 64 \rightarrow 32$
1	0.5	$128 \rightarrow 64 \rightarrow 32$
2	0.2	$128 \rightarrow 64 \rightarrow 32$
2	0.5	$64 \rightarrow 32 \rightarrow 16$
3	0.2	$64 \rightarrow 32 \rightarrow 16$
3	0.5	$64 \rightarrow 32 \rightarrow 16$
4	0.2	$64 \rightarrow 32 \rightarrow 16$
4	0.5	$64 \rightarrow 32 \rightarrow 16$

To address the transient nature of the gravity-driven Carbopol flow, a block-analysis approach was implemented. The recorded image sequence is divided into temporal

blocks corresponding to distinct flow regimes, as illustrated in Figure 3.4. During the initial rapid flow phase, characterized by high velocities and large displacements, image pairs are analyzed at short time intervals to prevent particles from exiting the interrogation windows. As the flow transitions to a slower, quasi-steady regime, the correlation time interval is progressively increased to maintain adequate particle displacement and preserve the signal-to-noise ratio of the correlation despite the reduced velocities.

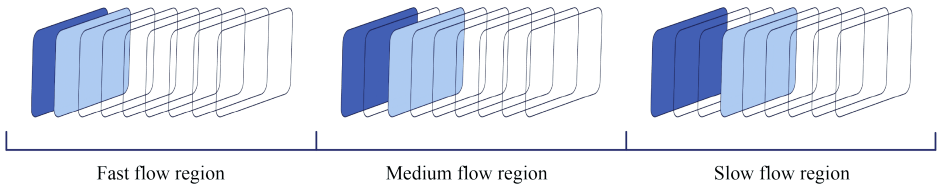


Figure 3.4: Temporal block structure for PIV analysis. The image sequence is divided into blocks with progressively increasing correlation time intervals ($\Delta t_1 < \Delta t_2 < \Delta t_3$) to adapt to decreasing flow velocities. Intermediate frames (transparent) are skipped to reduce computational cost while preserving temporal resolution during critical flow phases.

Since the Carbopol flow remains laminar throughout the experiment, with a Reynolds number below unity when evaluated with the apparent viscosity at the characteristic shear rate of the bend, intermediate frames between analyzed pairs are skipped without loss of relevant flow information. This frame-skipping strategy reduces computational cost while maintaining sufficient temporal resolution to capture the evolution of the velocity field during the critical flow transitions. The block boundaries were not prescribed a priori, but were determined based on the diagnostic described below.

To ensure the reliability of the velocity fields under these varying conditions, a validation step based on the Universal Outlier Detection (UOD) method, also known as the Normalized Median Test, was implemented (Westerweel et al., 2005). This metric is particularly effective for laminar flows with strong velocity gradients, such as non-Newtonian Carbopol flows, as it evaluates vector validity from local physical coherence rather than from global statistics. The algorithm computes a normalized residual $R_{i,j}$ by comparing each vector $\vec{V}_{i,j}$ with the median of its 3×3 neighborhood,

$$R_{i,j} = \frac{|\vec{V}_{i,j} - \vec{V}_m|}{r_m + \epsilon}, \tag{3.2}$$

where $\vec{V}_m$ is the median vector of the neighbors, r_m is the median of the absolute

residuals within the neighborhood, and ϵ is a small normalization factor accounting for measurement noise. A vector is classified as spurious if $R_{i,j}$ exceeds a predefined threshold. In this study, the evolution of the percentage of spurious vectors was used as the primary criterion for locating the transitions between temporal blocks: an abrupt increase in the UOD failure rate indicates that the particle displacement has fallen below the level at which the correlation can be resolved, and triggers the transition to a larger Δt.

The block structure adopted for every camera and both matrices, together with the spurious-vector diagnostic that fixes each inter-pair skip and the resulting reduction in the number of correlated pairs, is given in Appendix B.

3.5 PTV Implementation

Where PIV supplies the independent variable of the saturation law, PTV supplies the dependent one. The orientations of the fibers detected within a region reduce to the order parameter η_1, the quantity the law predicts, while the same rod tracked across successive frames yields in addition its angular velocity. Both quantities are recovered from the same detections, so the detection stage is treated here as a measurement problem rather than as a preprocessing step. The reason is that the fibers a detector fails to find are not a random sample of the population: a rod inclined unfavorably with respect to the illumination returns less light to the camera than one inclined favorably, so a detection rate below unity biases the measured distribution toward the orientations that are easiest to see. This section describes the implementation of the framework, covering the detection and segmentation pipeline, the training dataset, and the characterization and tracking algorithm. The framework was originally developed and validated in a companion study (Wolff et al., 2026).

3.5.1 Fiber Detection and Segmentation

As Figure 3.5 illustrates, the intensity reflected from each fiber depends on its orientation relative to the LED sources and the camera, as well as on the local surface properties of the fiber itself, resulting in spatially non-uniform brightness along the fiber length. Furthermore, owing to the small physical diameter of the steel fibers relative to

the imaging resolution, the projected width of each fiber spans only approximately two pixels in the image plane, approaching the resolution limit of the imaging system. At high fiber concentrations, these effects are compounded by frequent inter-particle occlusions, where partially overlapping fibers produce ambiguous intensity patterns that further compromise detection reliability.

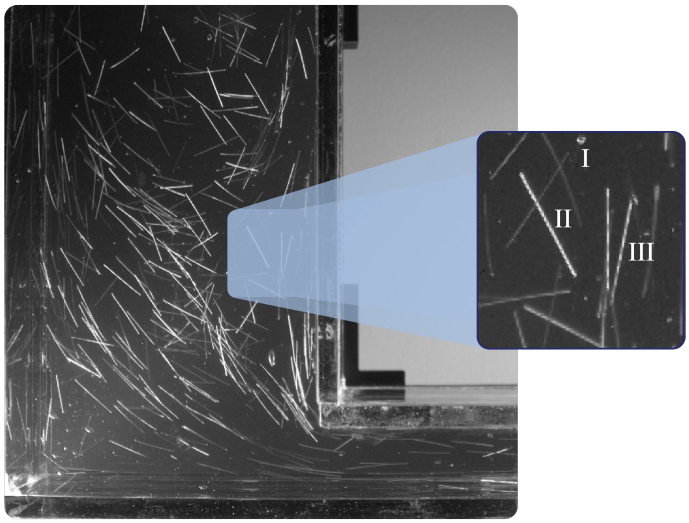


Figure 3.5: Representative examples of non-uniform illumination along individual steel fibers. The reflected intensity varies as a function of fiber orientation relative to the LED sources and local surface characteristics, producing three distinct cases: (I) a poorly visible fiber with low overall reflectance, (II) a well-illuminated fiber with consistent brightness, and (III) a fiber with irregular brightness distribution along its length due to surface unevenness. This variability is further compounded at high fiber concentrations by inter-particle occlusion, posing significant challenges for conventional segmentation techniques.

Case (I) of Figure 3.5 is the one that carries the measurement consequence. A fiber of low reflectance is not merely harder to detect; it is harder to detect *because* of the orientation it happens to hold with respect to the illumination, and the same orientation is held by every other fiber in the same condition. Missing it therefore removes a specific direction from the measured distribution rather than a random fiber from the count. Quantifying the fraction of the population that is lost in this way, and demonstrating that the loss does not vary systematically with the fiber dosage, is a requirement of the validation reported in Section 3.5.4, since a detection rate that fell with concentration would produce an apparent change in the orientation statistics with no change in the fibers themselves.

Two deep learning models capable of instance-level segmentation were evaluated and compared against the Hough transform baseline: YOLOv11 (Khanam et al., 2024) and the Segment Anything Model (SAM2) (Ravi et al., 2025). Both models were fine-tuned on the same domain-specific dataset using the Roboflow platform.

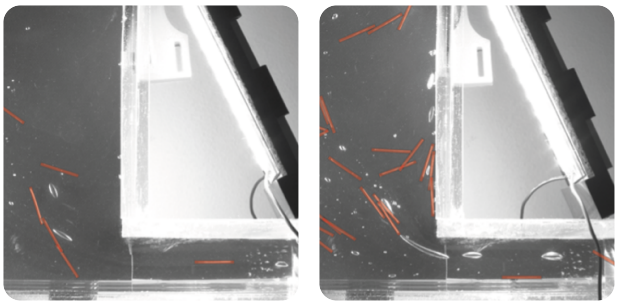
Training Dataset

The training dataset was constructed to reflect the two primary imaging scenarios encountered in the experiments, shown in Figure 3.6: sparse fiber distributions, where individual fibers are well-separated and clearly visible, and dense fiber distributions, where overlapping fibers, partial occlusions, and background artifacts such as elongated bubbles are prevalent. Including both scenarios was essential to ensure that the segmentation model learns to distinguish individual fibers from background noise and to handle the ambiguity introduced by fiber clustering and occlusion.

A total of 267 images were manually annotated at the instance mask level using the Roboflow platform. To increase dataset diversity and improve model generalization, each image was augmented using horizontal and vertical flips, 90° and 270° rotations, random rotations within ±15°, and horizontal and vertical shearing of ±15%, yielding a final augmented dataset of 673 images. The dataset was partitioned into training (90%), validation (8%), and testing (2%) subsets. Both YOLOv11 and SAM2 were fine-tuned on identical training and validation splits to ensure a fair performance comparison.

Following fine-tuning, both models were evaluated on the held-out validation set using standard instance-level segmentation metrics, namely precision, recall, F1 score, and intersection over union. YOLOv11 substantially outperformed both the Hough transform baseline and SAM2, the latter failing to detect any fiber despite identical fine-tuning conditions; the detailed comparison is reported in Section 4.1.1. On this basis YOLOv11 was adopted as the detector for the companion study, in which the PTV framework was developed and validated.

The experimental campaign of this thesis reaches fiber loadings higher than any represented in the original dataset and PTV validation tests, where overlap and occlusion grow severe and detection becomes correspondingly harder. To cover this regime the training set was extended with additional frames drawn from the high-concentration runs of Section 3.5.4, so that the dense conditions the campaign explores enter the



(a) (b)

Figure 3.6: Representative examples from the training dataset: (a) image with a sparse distribution of well-separated fibers, and (b) image with a denser distribution where overlapping fibers and bubbles are more prominent.

training distribution rather than being left to extrapolation. These frames were annotated in a model-assisted loop, in which YOLOv11 produced an initial segmentation of each frame that was then verified and corrected by hand, so that the annotations rest on human judgment and not on the predictions of the earlier detector. The extended set was augmented with the same parameters as before, and the dataset used to train YOLO26 comprised 679 images partitioned into training (90%), validation (7%), and testing (3%) subsets.

On this dataset the detector was updated to YOLO26, released after the companion study and, by then, the state of the art in instance segmentation. For the densest frames a dedicated detection methodology was built around it, in which the field is upscaled and tiled before segmentation so that fibers packed into a cluster can still be resolved individually, detailed in Section 3.5.2. The high-concentration detection of the present work therefore rests on YOLO26 and the extended dataset assembled for it, while the validation of the framework reported in the companion study rests on YOLOv11 and the original dataset.

3.5.2 High Concentration Detection

Although the instance segmentation model performs well at low and moderate fiber concentrations, its performance degrades at higher concentrations due to the increased frequency of occlusions and the resulting ambiguity in fiber boundaries, as Figure 3.11 shows. To address this, an algorithm was developed to improve the detection of individual fibers within dense clusters.

Figure 3.7 illustrates the steps of the algorithm, which proceeds as follows. First, the preprocessed image is upscaled by a factor of two in resolution, thereby increasing the number of pixels spanned by each fiber and, in particular, raising the projected fiber width from approximately two pixels to approximately four, which is the dimension on which the segmentation of adjacent, partially overlapping rods depends. Subsequently, overlapping 640×640 patches are generated with a 10% overlap and fed into the YOLO model for fiber segmentation, so that the network operates at a scale for which it was trained rather than on a downsampled view of the whole field. Finally, the detections are reconstructed onto the original image using a non-maximum suppression strategy to resolve the duplicate predictions arising from the patch overlap. A dedicated model was trained on a dataset processed through the same upscaling and tiling logic, so that the training and inference distributions coincide.

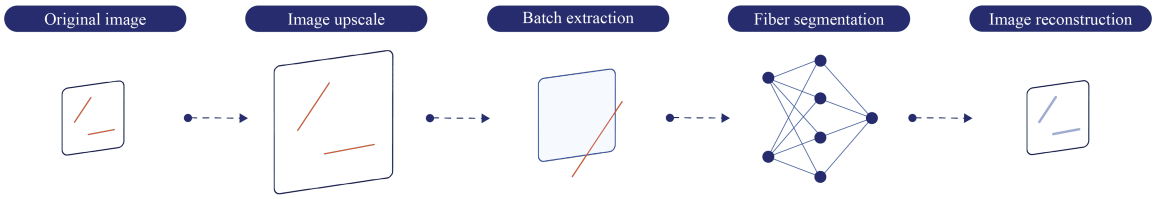


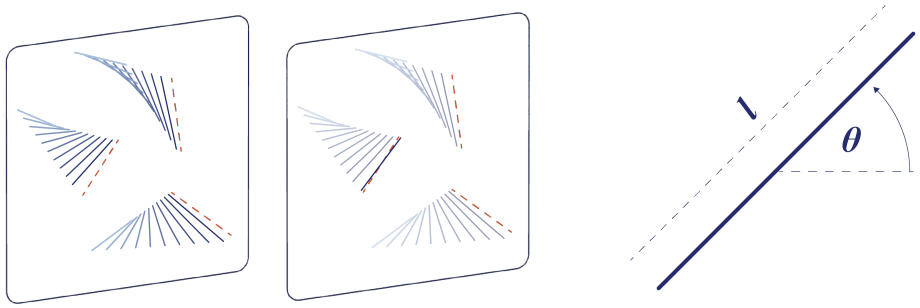
Figure 3.7: Workflow of the high-concentration fiber detection algorithm. The original image is upscaled by a factor of two to enhance feature resolution, then processed through overlapping 640×640 pixel patches for YOLO-based segmentation. Detected fiber masks are reconstructed into the original image space to yield improved segmentation performance in dense fiber clusters.

3.5.3 Fiber Characterization and Correlation

Once fibers are detected and segmented, their geometric properties, namely the orientation θ, the length ℓ, and the centroid (c_x, c_y), are extracted to enable tracking across frames. A two-stage prediction-correction approach is employed for robust fiber association: first, an α-β-γ filter (Tenne et al., 2000; Welch et al., 1995) predicts the position and orientation of each fiber in the next frame from its kinematic history, providing a constrained search region for potential matches and reducing ambiguity in dense suspensions; second, a cosine similarity metric evaluates the candidate pairings from their geometric feature correspondence to select the most likely matches. This combined approach maintains tracking accuracy under partial occlusion and fiber overlap, which

are common in a semi-dilute suspension. Figure 3.8(b) illustrates the fiber features used for correlation.

Figure 3.8(a) shows how the two stages operate on a single frame transition. From the established tracks, each carrying the kinematic history of a fiber, the α-β-γ filter projects the expected state of every fiber into the current frame, drawn as the predicted fibers (dashed orange). The fibers actually detected in the current frame (blue) are then matched against these predictions: each detection is compared with the predicted fibers that fall within its search region, and the cosine similarity of their geometric features selects the correspondence. A detection that matches a prediction continues the existing track, while the predicted state supplies the position and orientation of a fiber that the segmentation misses in that frame, so that a temporary occlusion interrupts neither the trajectory nor the count.



(a) (b)

Figure 3.8: (a) Workflow of the α-β-γ filter for fiber state prediction and cosine similarity-based matching for fiber association across frames. (b) Geometric features extracted from detected fibers: orientation θ, length ℓ, and centroid coordinates (c_x, c_y), which are used to construct the feature vectors for similarity evaluation.

α-β-γ Filter for Kinematic Prediction

The α-β-γ filter is a simplified Kalman filter that maintains a kinematic state estimate for each tracked fiber, consisting of the position X, the velocity $\dot{X}$, and the acceleration $\ddot{X}$, with analogous quantities for the angular coordinate. The filter alternates two steps: a correction of the current prediction using the new measurement, followed by a propagation of the corrected state to the next time step. The notation $\hat{X}_{m,n}$ denotes the estimate of the state at frame m given the measurements available up to frame n.

When a measurement Z_n, the detected fiber centroid or orientation, becomes available at frame n, the residual between it and the prediction made at the previous step is

$$r_n = Z_n - \hat{X}_{n,n-1}, \tag{3.3}$$

and the state is corrected in proportion to it,

$$\hat{X}_{n,n} = \hat{X}_{n,n-1} + \alpha\, r_n, \tag{3.4}$$

$$\dot{\hat{X}}_{n,n} = \dot{\hat{X}}_{n,n-1} + \frac{\beta}{\Delta t}\, r_n, \tag{3.5}$$

$$\ddot{\hat{X}}_{n,n} = \ddot{\hat{X}}_{n,n-1} + \frac{\gamma}{2\,\Delta t^2}\, r_n \tag{3.6}$$

where α, β, and γ are gains controlling the responsiveness of the position, velocity, and acceleration estimates to the new measurement. The corrected state is then propagated forward under constant-acceleration kinematics to produce the prediction for the next frame,

$$\hat{X}_{n+1,n} = \hat{X}_{n,n} + \dot{\hat{X}}_{n,n}\,\Delta t + \tfrac{1}{2}\,\ddot{\hat{X}}_{n,n}\,\Delta t^2, \tag{3.7}$$

$$\dot{\hat{X}}_{n+1,n} = \dot{\hat{X}}_{n,n} + \ddot{\hat{X}}_{n,n}\,\Delta t, \tag{3.8}$$

$$\ddot{\hat{X}}_{n+1,n} = \ddot{\hat{X}}_{n,n}. \tag{3.9}$$

The prediction $\hat{X}_{n+1,n}$ defines the expected state of the fiber in frame $n+1$ and hence the search region for the matching step. The same equations are applied independently to the translational coordinates (c_x, c_y) and to the rotational coordinate θ.

Cosine Similarity Metric for Fiber Matching

To associate the predicted fibers with the fibers detected in the new frame, each fiber is represented by a feature vector encoding its geometric properties,

$$\mathbf{v} = \begin{pmatrix} w_1 \cos 2\theta \\ w_2 \sin 2\theta \\ w_3\, \ell \\ w_4\, c_x \\ w_5\, c_y \end{pmatrix}, \tag{3.10}$$

in which the double-angle representation ensures continuity for the periodic orientation coordinate, since a fiber at 0° and one at 180° describe the same physical state, and the coefficients w_k normalize each feature to a comparable scale. The similarity between a predicted fiber $\mathbf{q}_i$ and a detected fiber $\mathbf{d}_j$ is quantified by the cosine of the angle ψ_{ij} between their feature vectors,

$$s_{ij} = \frac{\mathbf{q}_i \cdot \mathbf{d}_j}{\|\mathbf{q}_i\|\, \|\mathbf{d}_j\|} = \cos \psi_{ij} \in [-1, 1], \tag{3.11}$$

where $s_{ij} = 1$ indicates identical feature vectors and $s_{ij} = 0$ indicates orthogonality. Associating the M predicted fibers with the N detected fibers is then a linear assignment problem. With a binary variable m_{ij} equal to one when $\mathbf{q}_i$ is matched to $\mathbf{d}_j$ and zero otherwise, the association maximizes the total similarity,

$$\max_{m_{ij}} \sum_{i=1}^{M} \sum_{j=1}^{N} s_{ij}\, m_{ij} \quad \text{subject to} \quad \sum_{j=1}^{N} m_{ij} \leq 1, \qquad \sum_{i=1}^{M} m_{ij} \leq 1, \tag{3.12}$$

so that each fiber is matched to at most one counterpart. The pairs are selected jointly. A greedy rule that fixes the most similar pair first can leave the remaining fibers with poor matches, a situation that becomes frequent in dense suspensions, where neighboring fibers have similar features. Problem (3.12) is solved with the Hungarian algorithm (Kuhn, 1955; Munkres, 1957), which finds the global optimum without evaluating every possible assignment, in a time that grows with the cube of the number of fibers. The algorithm operates on the equivalent cost matrix with entries $1 - s_{ij}$ and relies on a simple property: subtracting a constant from a whole row or column changes the cost

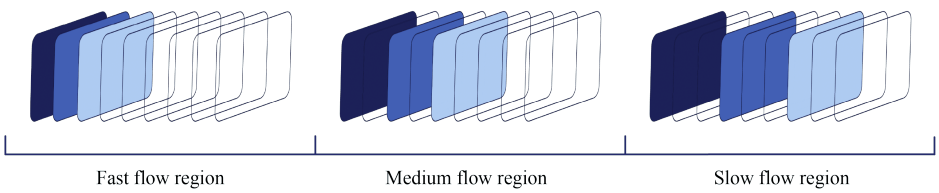


Figure 3.9: Temporal block structure for PTV analysis. The image sequence is divided into blocks with progressively increasing frame-skipping intervals to adapt to decreasing fiber velocities and maintain tracking accuracy while reducing computational cost during less dynamic flow phases.

of every complete assignment by the same amount, so the optimal assignment is unchanged. Rows and columns are reduced in this way until an assignment can be formed from zero entries only, one per row and one per column, which is then optimal because no entry of the reduced matrix is negative. When $M \neq N$, the matrix is completed with dummy rows or columns. A prediction left without a detection corresponds to a fiber missed by the segmentation in that frame, whose state is then supplied by the filter prediction, and a detection left without a prediction starts a new track. Pairs farther apart than the maximum linking distance $\Delta x_{\max}$ of Appendix B receive a prohibitive cost and are never matched.

PTV Block Structure

A block-analysis approach analogous to that used for PIV is applied to PTV, but it means something different. PIV correlates image pairs, so its block structure sets the time interval between the two frames of a pair. PTV analyzes each frame independently and then links the detections, so its block structure sets the temporal sampling rate, that is, how many frames are skipped between the frames that are analyzed.

During the initial rapid flow phase, when fibers experience high velocities and large displacements, frames are analyzed at short intervals so that the kinematic prediction of the α-β-γ filter remains accurate and no fiber travels beyond the search region between consecutive analyzed frames. As the flow slows, the analysis interval is progressively increased, reducing the number of segmentation and matching operations while preserving sufficient temporal resolution to follow the evolution of the fiber orientation. Figure 3.9 illustrates this adaptive structure.

The choice of the sampling interval is constrained from both sides, and the constraint is

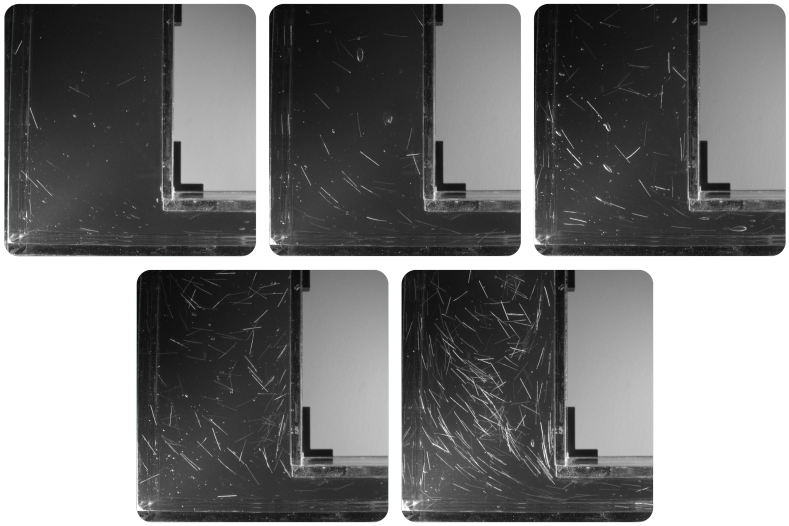


Figure 3.10: PTV test images with varying fiber loadings: 25, 100, 200, 400 and 800 fibers respectively.

tighter here than in PIV. An interval too long breaks the trajectories, since a fiber that leaves the search region between analyzed frames is lost to the tracker. An interval too short reduces the fiber displacement between analyzed frames toward the localization noise, so that the change in position carries less signal than the error in measuring it, and the angular velocity computed from it degrades. The block structure adopted for every camera and both matrices, together with the displacement diagnostic that fixes each interval and the resulting reduction in computational cost, is given in Appendix B.

3.5.4 PTV Algorithm Validation Tests

To assess the performance of the framework across the range of conditions encountered during the experiments, six fiber loadings were tested in the fixed-volume L-shaped device: 25, 50, 100, 200, 400, and 800 fibers, as shown in Figure 3.10. This range brackets the transition from sparse, well-separated fibers, for which individual detection and tracking are expected to be reliable, to dense suspensions dominated by occlusion and overlap, for which trajectory continuity is progressively challenged. Characterizing the framework across this span establishes the loadings within which Lagrangian trajectory reconstruction remains viable and identifies those at which the analysis must be restricted to frame-by-frame detection.

The 50-fiber test exhibited anomalously low performance metrics, departing from the

monotonic degradation with density observed across the remaining loadings. This was attributed to a suboptimal fiber alignment relative to the camera viewing angle during that particular run, which left a significant fraction of the fibers poorly illuminated and of low contrast. The 50-fiber test was accordingly excluded from the quantitative performance analysis, and the validation proceeded with the remaining five loadings, namely 25, 100, 200, 400, and 800 fibers.

3.6 Experimental Campaign

The campaign measures the motion of the carrier and the motion of the fibers on the same flow, and it is designed to separate the effect of the fiber content from that of the rheology of the carrier. It therefore varies one factor at a time. The geometry and the filling procedure are held fixed, and within each series the rheology is fixed as well, so that the fiber content is the only quantity that changes between runs and any difference in the orientation it produces is attributable to it. Two series are conducted, one in each Carbopol formulation, which provide a contrast in yield stress. Fiber-free runs are included in both and serve as the carrier reference against which the perturbation the fibers exert on the flow is measured. The paired fluid-phase and fiber-phase data obtained in this way constitute experimental evidence on their own of how a viscoplastic casting flow orients steel fibers, and they also provide a reference dataset for the validation of the CFD-DEM model under development in the research group.

The concentration regime of Section 2.2.3 sets the dosages. With $r_f = 65$ the dilute limit falls at $\phi_{sd} = r_f^{-2} = 0.024\%$ by volume, and the design places one dosage below it and two above, bracketing the onset of the interaction mechanism rather than probing a single point of it. Each Carbopol batch is prepared to a nominal 2 L and loses approximately 198 mL during dispersion, neutralization, centrifugation, and transfer, leaving a working suspension of 1.802 L, at which volume the dilute limit corresponds to some 1023 fibers. The device is filled to a height of 22.4 mm and receives a fixed fraction, about 73%, of each batch, the remainder staying in the reservoir, so the number deposited within the measurement domain follows from the nominal count by that same fraction and is reported alongside it in Table 3.6. Fibers were counted by weighing, using the numerical density of Table 3.2 as the conversion factor and a balance

whose resolution is at least an order of magnitude finer than the mass of a single fiber, approximately 3.5 mg.

Table 3.6: Fiber dosages of the experimental campaign, set against the boundaries of the concentration regimes and against the dosages of a real UHPC mix. The number in the L-shape is the subset deposited within the device at a fill height of 22.4 mm, about 73% of each batch, and is informative only; the volume fractions and their ratios are computed over the full working suspension of 1.802 L, from the nominal fiber volume of 0.408 mm^3. Two regime boundaries follow from the aspect ratio $r_f = 65$: the dilute limit $\phi_{sd} = r_f^{-2}$, below which the fibers do not interact, and the concentrated limit $\phi_c = r_f^{-1}$, above which they are in permanent contact; the semi-dilute regime lies between them. The campaign occupies the lower edge of that regime, while a real UHPC mix sits at or above its upper edge. Reference rows are not tested conditions.

	N° of Fibers	**Weight Fract. (%)**	**Volume Fract. ϕ (%)**	**N° in L-shape**	ϕ/ϕ_{sd} (–)	ϕ/ϕ_c (–)
Experimental campaign						
N_0	0	0.000	0.000	0	0.000	0.000
N_1	750	0.140	0.017	547	0.720	0.011
N_2	1500	0.270	0.035	1094	1.480	0.023
N_3	3000	0.500	0.069	2188	2.920	0.045
Regime boundaries ($r_f = 65$)						
Dilute limit $\phi_{sd} = r_f^{-2}$	1023	0.185	0.024	746	1.000	0.015
Concentrated limit $\phi_c = r_f^{-1}$	66 525	12.000	1.538	48 509	65.000	1.000
Real UHPC mix						
UHPC (1% vol.)	43 241	7.800	1.000	31 531	42.300	0.650
UHPC (2% vol.)	86 482	15.600	2.000	63 062	84.500	1.300

The reference rows place the campaign on the map of concentration regimes, and the distance they reveal is the context for a result that recurs throughout this work. A practical UHPC mix carries 1 to 2% of fibers by volume, 0.65 to 1.30 times the concentrated limit $\phi_c = r_f^{-1}$ (Luo et al., 2023): the fibers of a real element are in near-permanent contact, at the boundary of the concentrated regime or already within it. The dosages of this campaign reach at most 0.045 times that limit, one to two orders of magnitude lower, at the lower edge of the semi-dilute regime and barely across the dilute limit below which the fibers would not interact at all. The interaction mechanism the experiment sets out to observe is therefore present but faint: the fibers are just close enough to disturb one another, and no closer.

This distance is not a shortcoming of the design but a constraint it cannot escape,

and it has a definite consequence for the results. The interactions are weakest at these dosages, so the effect of the fiber content on every quantity measured here is the smallest it can be while remaining nonzero. This is why the concentration effects reported in Chapter 4 are modest in size even where they are systematic in direction, and why any effect recovered at all is a lower bound on what a real mix would produce. The constraint is optical. A suspension dense enough to reproduce a real element is one in which the fibers occlude one another and can no longer be resolved individually, so raising the dosage to close the gap would defeat the measurement itself. The gap is the price of seeing the fibers at all. Closing it is left to the CFD-DEM simulation these data are gathered to validate, which resolves each fiber without needing to see it and can be run at the concentrations the experiment cannot image.

Each formulation-dosage combination was replicated with two independently prepared batches, so that the batch-to-batch variability of the gel is separated from the effect of the fibers, and every batch supplies two acquisitions, the limit imposed by the mixing protocol of Section 3.3. In the runs carrying fibers the two acquisitions are a PTV followed by a PIV: the fiber dynamics are recorded first, tracer particles are then seeded into the same gel, and the pour is repeated, so that the velocity field and the fiber orientations are obtained on the same batch and under identical conditions. In the fiber-free runs, the replicates of N_0, both acquisitions are PIV. Across the four dosages each series therefore comprises eight batches and sixteen acquisitions, yielding ten PIV and six PTV datasets: four PIV from the two N_0 batches, and one PIV and one PTV from each of the six batches carrying fibers.

The runs are referred to throughout this work by their nominal fiber count, 0, 750, 1500, and 3000 fibers, although the number effectively present within the measurement domain is lower, since the gel is not poured in its entirety into the device. The values of Table 3.6 are the ones used in every calculation.

With the experimental campaign fixed, one task remains before the results of Chapter 4 can be reported: defining the quantities measured, the estimators applied to them, and the thresholds that separate a meaningful effect from noise, so that the effect sizes compared across zones, stages, matrices, and variables throughout that chapter rest on a single, common protocol.



Figure 3.11: The L-shaped device prior to an experimental run at each fiber dosage, 750, 1500, and 3000 fibers, shown for both Carbopol formulations. The images illustrate the increase in fiber density and spatial overlap with the loading.

3.7 Analysis Metrics and Significance Criteria

This section fixes the quantities that describe the flow and the fibers, the estimators used to test them, and the thresholds that separate a meaningful effect from noise. Two questions recur in every section that follows: whether the rheology of the carrier affects a given quantity, and whether the fiber content does. Both are answered with a common protocol, so that the effect sizes reported across zones, stages, matrices, and variables are directly comparable.

3.7.1 Measured Quantities

The fluid is measured with PIV and described in an Eulerian frame, through the speed $|\mathbf{u}| = \sqrt{u^2 + v^2}$ and the vorticity magnitude $|\zeta|$, with $\zeta = \partial v/\partial x - \partial u/\partial y$ the out-of-plane component of the vorticity vector $\boldsymbol{\zeta} = \nabla \times \mathbf{u}$ defined in Section 2.3.1. The symbol ζ is reserved for the fluid so as not to collide with the fiber angular velocity ω. The two are not equivalent rates: by Equation (2.11) a fluid element rotates at $\zeta/2$, so a fiber turned only by the rotation of the fluid would have $\omega = \zeta/2$.

The fibers are measured with PTV and described in a Lagrangian frame, through the fiber speed $|\mathbf{v}|$ and the angular velocity ω of the fiber axis. Three further descriptors summarize the state of the population within a zone, all defined in Section 2.3.3: the order parameter $\eta_1 = |\langle e^{i2\theta} \rangle|$, which measures how aligned the fibers are; the angular

deviation δ, which measures the misalignment of each fiber from the director of its zone, so that a rise in δ signals a loss of local order; and the dispersion index Λ, which compares the radius of rotation of the fiber centroids with that of a uniform arrangement over the same cell. Table 3.7 collects them.

The order parameter and the angular deviation describe the same state from opposite ends, and the distinction matters for reading what follows. The order parameter is a property of the population, computed as a single number from all the fibers of a cell, so it is the natural response when the question concerns the degree of alignment a cell has reached. The angular deviation is a property of each fiber, so it supplies one observation per fiber rather than one per cell, and it is the natural response when the question requires a test with enough degrees of freedom to be resolved within a single cell. A rise in δ and a fall in η_1 describe the same physical event, namely a loss of order.

Table 3.7: Quantities used throughout the analysis.

Symbol	**Quantity**	**Phase**	**Range and meaning**
$\lvert\mathbf{u}\rvert$	Fluid speed	PIV	≥ 0
$\lvert\zeta\rvert$	Fluid vorticity	PIV	≥ 0; twice the local rotation rate of the fluid
$\lvert\mathbf{v}\rvert$	Fiber speed	PTV	≥ 0
ω	Fiber angular velocity	PTV	Rotation rate of the fiber axis
$\hat{\eta}_1$	Orientation order	PTV	0 isotropic $\rightarrow$ 1 fully aligned
δ	Angular deviation	PTV	$[0°, 90°]$; misalignment from the director
Λ	Spatial dispersion	PTV	<1 clustered; 1 uniform; >1 over-dispersed

3.7.2 Effect Size

The filling is transient, so the effect of a factor cannot be read from a simple mean: the same zone passes through very different flow conditions as the mold fills. Every test is therefore an Analysis of Covariance with time as a continuous covariate, so that the shared temporal trend is removed before the factor of interest, the fiber content or the matrix, is evaluated. The strength of the factor is reported as the partial eta squared,

$$\eta_p^2 = \frac{SS_{\text{factor}}}{SS_{\text{factor}} + SS_{\text{error}}}, \tag{3.13}$$

the share of variance the factor explains once the variance attributable to the other terms has been removed. Unlike the ordinary eta squared, whose denominator is the total variance, the partial form does not shrink as further factors are added, which is what makes the effect sizes of different models comparable. It is written simply η^2 below.

The effect size answers how large an effect is, not whether it is distinguishable from zero, and the two questions are settled by separate numbers: the p value of the ANCOVA term at $\alpha = 0.05$ for the second, and the bands of Table 3.8 for the first (Cohen, 2013). A relevance threshold of $\eta^2 \geq 0.02$ is adopted throughout, placed just above Cohen's small-effect boundary so that it acts as a noise floor: values below it are reported but not interpreted, however small their p value. The threshold guards against a specific failure. Most tests below are run on the populated zones, where the observations number in the hundreds or thousands, so the standard error is minute and an effect explaining less than one percent of the variance is still flagged as significant; the threshold is what keeps such an effect from being read as a result.

The opposite failure occurs where the data are sparse, and is guarded against separately. The distal cells of the beam, and every cell of the stiffer matrix beyond the second column, hold too few fibers to support a test, and a fluctuation computed over a few dozen observations can return a large effect size on a p value that establishes nothing. Such cells are not reported: the tables of Appendix E are restricted, in each matrix, to the zones the flow actually reaches, a zone being retained only where the carrier establishes a resolved velocity field and a countable population of fibers is tracked.

Table 3.8: Interpretation bands for the partial eta squared, after Cohen, and the thresholds adopted in this work. The relevance threshold sits just above the small-effect boundary; the grading thresholds span the range observed in this dataset and do not mark new statistical categories.

η^2	**Cohen**	**Used in this work**
< 0.01	negligible	below the noise floor
0.01–0.06	small	≥ 0.02: relevance threshold
0.06–0.14	medium	≥ 0.10: moderate
≥ 0.14	large	≥ 0.30: strong

3.7.3 The Sampling Control

One control is applied on top of the effect size wherever the fiber content is the factor, and it is not optional, because the measurement and the hypothesis push the same quantity in the same direction.

The detection efficiency falls with the loading (Section 4.3.3), so a cell of a richer mix is described by a smaller sample of fibers than the same cell of a leaner one. The order parameter carries a finite-sample bias that falls as the sample grows, $\mathbb{E}[\eta_1^2] \simeq 1/n$ for a population of n tracked fibers drawn from an isotropic distribution, so a cell holding fewer fibers returns a higher spurious order. The bias therefore acts against the hypothesis at high dosage and would, if uncorrected, mask the effect rather than manufacture it; but the bias also varies from cell to cell with the count, and an uncorrected comparison across dosages would confound the physical trend with the sampling. Two measures address this. The order parameter is reported throughout in its debiased form, obtained by subtracting the leading term of the finite-sample bias,

$$\hat{\eta}_1^2 = \frac{n\,\eta_1^2 - 1}{n - 1}, \tag{3.14}$$

which is unbiased to leading order for an isotropic population and reduces to η_1 as the count n grows. And where a trend with the dosage is claimed, the tracked fibers are bootstrap-resampled to an equal count across the three levels and the trend is retained only if it survives in the majority of resamples; a trend that does not survive is an artifact of unequal sampling and not a physical effect.

Chapter 4

Experimental Characterization of Flow and Fiber Kinematics

Velocity profiles from CFD solution,
each correspond to 1 second.

This chapter provides what the central result of the thesis requires before it can be established. That result, the confrontation of the saturation law with the measured orientation, is reserved for Chapter 6. What precedes it is organized around three tasks. The first is to establish what the measurement can and cannot recover, since the orientation of a fiber that the detector fails to find is not missing at random, and the analysis must be bounded by what the instrument delivers. The second is to fix the frame within which every subsequent comparison is defined, namely the partition of the mold into zones and of the filling into stages, together with the statistical protocol that makes an effect measured in one zone comparable with one measured in another. The third is to deliver the three facts on which the dynamics of the fibers depends: that the fiber content perturbs the flow, and does so regionally rather than uniformly; that the fibers rotate faster as their number grows while the fluid carrying them does not, which locates the additional rotation outside the mean flow; and that the material falls into the beam rather than flowing into it, which fixes the origin from which the accumulated strain is counted and the rheology on which the law can be traced.

4.1 PTV Algorithm Validation

The PTV framework was characterized along two axes, using the validation tests of Figure 3.10: the accuracy of the detection stage, which determines how reliably a fiber present in a frame is segmented within it, and the robustness of the tracking stage with fiber loading, which determines whether the data support an analysis at the level of the individual fiber or only at the level of the population. The complete assessment is reported in the companion study (Wolff et al., 2026).

4.1.1 Segmentation Performance

Three methods were compared on the held-out validation set: the Hough transform as a classical baseline, and YOLOv11 and SAM2 fine-tuned on identical data (Table 4.1). The Hough transform recovers roughly half of the fibers but returns a large number of false detections for every true one, at a precision of 0.041. SAM2 recovered nothing under identical fine-tuning conditions. The cause is not the fine-tuning, nor an inability to segment structures this thin: given the fiber location as a prompt, SAM2 reaches a

Table 4.1: Segmentation performance for rod-like particle detection. Metrics for YOLOv11 and SAM2 are computed at the instance mask level; the Hough transform IoU is computed at the pixel level and is not directly comparable.

Method	Model type	Precision	Recall	F1 Score	IoU
Hough Transform	Traditional	0.041	0.557	0.076	0.739*
SAM2	Deep Learning	0.000	0.000	0.000	0.000
YOLOv11	Deep Learning	0.947	0.982	0.963	0.848

*Pixel-level IoU; not directly comparable to instance-level values.

mean IoU of 0.71 . SAM2 is a promptable model with no object-proposal stage, so in automatic mode the prompts come from a fixed lattice laid over the image rather than from the image content itself. Fibers approximately two pixels wide occupy too little of the frame to fall reliably on that lattice, and most are never queried. The failure lies in finding the fibers rather than in segmenting them. YOLOv11, already adopted in Chapter 3 on the strength of these numbers, reaches a recall of 0.982 and a precision of 0.947; the recall is measured against the manually annotated ground truth, so it reports the fraction of *visible* fibers the detector finds, not the fraction of *present* fibers the imaging system reveals.

4.1.2 Tracking Performance

Tracking was evaluated over a window of 600 frames, which is the interval within which every fiber of the batch crosses the field of view of the camera. The window fixes the reference against which the tracking is judged: a tracker that recovered each fiber once, as a single unbroken trajectory, would issue exactly as many identities as there are fibers, and the cumulative count of identities normalized by the known fiber population would therefore approach unity at the end of the sequence. Any excess over unity is a fiber that was lost and reacquired under a new identity, and is a direct measure of the fragmentation of the trajectories. Two quantities are reported within this window: the instantaneous detection ratio, the fraction of the known fiber population segmented in a given frame, and the trajectory length, the number of consecutive frames over which a fiber identity is maintained (Table 4.2).

The two quantities degrade in different ways, and the distinction between them governs how the data may be used. The detection ratio falls sharply between the sparsest

Table 4.2: PTV tracking performance across fiber loadings, evaluated over the 600-frame window within which every fiber crosses the field of view. The detection ratio is the instantaneous fraction of the fiber population segmented per frame; trajectory lengths are given in frames; "complete" denotes the fraction of tracks lasting ≥ 500 frames, that is, spanning at least five sixths of the window.

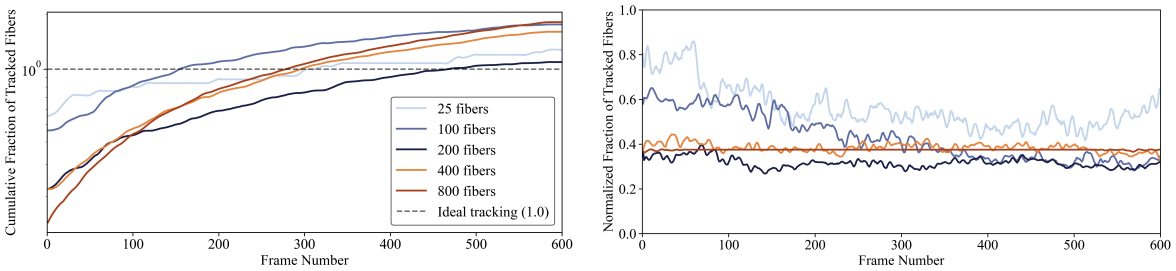
Fibers	**Detection ratio**	**Mean track length**	**Median**	**Complete** [%]
25	0.64	197.6	96	18.8
100	0.35	116.0	58	5.7
200	0.32	140.9	72	5.9
400	0.33	96.0	50	1.6
800	0.38	67.8	41	0.6

loading and the next, from 0.64 to 0.35, and then holds roughly constant, between 0.32 and 0.38, over the remaining fourfold increase in the fiber count. Its level is set by three causes that operate at any loading: not every fiber lies within the field of view at a given moment, fibers inclined toward the optical axis present a reduced projection and are intermittently lost, and mutual occlusion removes a further share. Its flatness beyond the second loading indicates that these causes saturate, so the framework recovers about a third of the population at any instant almost irrespective of how many fibers that population contains.

The trajectory statistics do not saturate. The mean track length falls from 197.6 frames at 25 fibers to 67.8 at 800, and the fraction of complete tracks collapses from 18.8% to 0.6%, so that at the highest loading essentially no fiber is followed from its entry to its arrest. This is the central limitation of the measurement, and it is the reason a fiber cannot simply be tracked from the mold inlet to the beam: what the framework delivers is a sequence of short fragments, and the provenance of a fiber must be reconstructed by stitching them, as Section 4.5 does.

Figure 4.1(a) reports the cumulative count of distinct identities against the unit reference established above, and it exceeds unity at every loading, which is the signature of fragmentation rather than of spurious detection: a tracker inventing objects that are not there would climb without bound, whereas one that merely breaks trajectories issues a bounded excess of identities over fibers. That the excess remains modest is therefore the evidence that the detections are real and the failures are of continuity.

Figure 4.1(b) confirms that the detection rate is stable throughout the sequence rather than degrading as the flow develops, so the fragmentation is not the accumulation of a drift.



(a) (b)

Figure 4.1: PTV tracking performance across fiber loadings. (a) Cumulative number of distinct track identities over the 600-frame sequence, normalized by the known fiber count. Every fiber crosses the field of view within this window, so an ideal tracker would approach unity; values above it indicate trajectory fragmentation rather than spurious detection. (b) Number of fibers detected per frame, showing a detection rate that is stable in time at every loading.

Fragmentation is one of two limitations of the measurement, and the second is inherent to its geometry rather than to the loading. The tracking recovers each fiber from a single camera, so it locates the fiber within the plane of the image but not its depth along the optical axis, and it measures the projection onto that plane of an orientation that is in general three-dimensional. A fiber inclined out of the plane appears shorter than it is and at a projected angle that is not the angle it holds in space, so the population reported in any cell is a projection of a three-dimensional arrangement onto two dimensions. This is not removed by any amount of training or by widening the tracking tolerances, since the information is absent from a single view and would require a second camera or a volumetric acquisition to recover. Its consequence for the order parameter is examined where the planar reduction of the orientation model is set out, in Appendix A.7: the projection biases the measured alignment, and the bias has a definite sign, but it acts in the same direction at every dosage and therefore displaces the curves of the saturation law together rather than manufacturing a separation between them. Because the flow conditions a fiber experiences change markedly along its path through the mold, a fiber descending the vertical arm is subject to a different kinematic history than one settling at the beam outlet. A single domain-wide average would conflate re-

gions of distinct behavior. The following section partitions the geometry into a fixed set of spatial zones that follow the sequence of flow conditions along the casting direction, so that the effect of fiber concentration and carrier rheology can be resolved locally rather than averaged across the entire domain.

4.2 Domain Zones

Every quantity reported hereafter is local. A fiber accumulates strain unevenly along its path, so the order parameter and the accumulated strain both vary across the mold, and an average over the whole domain would conflate regions in which the material has been strongly deformed with regions in which it has barely deformed at all. The partition into zones is what allows a single casting to supply a range of accumulated strains rather than a single value, and it is therefore not a convenience of presentation but the instrument through which the saturation law of Equation (5.27) is traced.

Figure 4.2 shows the partition, defined once and applied identically to every Carbopol formulation and fiber dosage. The L-shape is divided into three zones following the sequence of flow conditions along the casting path: the vertical arm Z1, where the flow is predominantly translational and the material advances largely as a plug; the elbow Z2, where the flow turns through ninety degrees and both the deformation and the vorticity reach their maxima; and the outlet Z3, where the flow passes from the L-shape into the beam. The beam is partitioned into a two-by-three grid of cells, r1c1 to r2c3, with the lower row adjacent to the outlet. The figure preserves the proportions of the mold, so the relative sizes of the zones are those they occupy in the experiment. Z1, Z3, and the beam cells follow from the physical boundaries of the mold, but the elbow does not: the rotation of the flow decays gradually into the adjoining straight sections rather than stopping at a wall, so Z2 has no natural geometric definition and must be given one. It was fixed by a sweep over the height and width of its bounding box, computing for each candidate the fraction of the total rotational content of the region it captures, using the fiber-free control of each matrix so that the zone is defined by the flow alone and not by the fibers whose behavior within it is to be measured. The adopted Z2 is the most compact box that captures the target fraction in both formulations simultaneously, so that a single definition serves the stiff and the soft

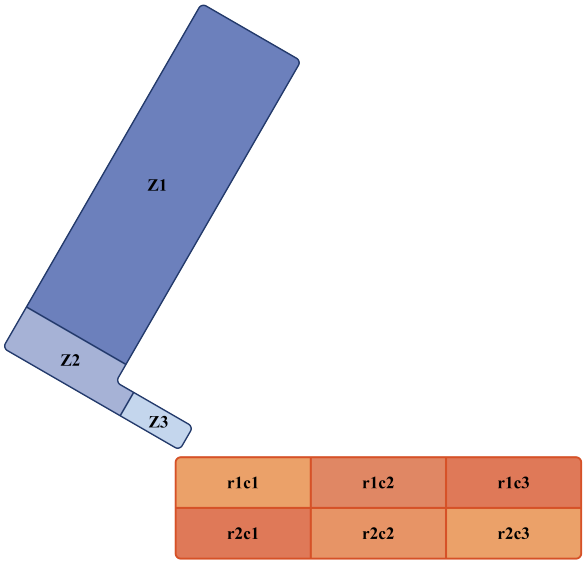


Figure 4.2: Spatial zones used throughout the analysis. The L-shape is partitioned into the vertical arm (Z1), the elbow (Z2), and the outlet (Z3); the beam is partitioned into a two-by-three grid of cells (r1c1–r2c3), with the lower row adjacent to the L-shape outlet. The figure is drawn to the proportions of the mold and is scaled uniformly, so relative dimensions are preserved.

matrix without favoring either.

The filling is transient, so each zone passes through a sequence of regimes as the mold fills, and treating the whole filling time as one observation would average together periods of very different fiber dynamics. Three stages are distinguished within each zone: the onset, as the front arrives and the velocity rises from rest to its peak; the transition, as the flow decelerates sharply after that peak; and the quasi-steady stage, during which the flow settles for the remainder of the filling. The boundaries are located independently for each zone and each matrix, as described in the following section.

4.2.1 Flow Stages by Zone

The stage boundaries are read from the local velocity field rather than imposed globally, because the zones do not fill in unison: the vertical arm reaches its peak velocity while the beam is still empty, and the far end of the beam is arresting while the elbow is still yielding. A stage defined by the clock would therefore group together material in entirely different kinematic conditions, which is precisely what the partition is meant to avoid.

The criterion is applied to the fiber-free control of each matrix, so that the stages are fixed by the carrier alone and not by the fibers whose behavior within them is to be measured. For each zone the median velocity magnitude across the measurement points is computed at every timestamp and smoothed with a centred moving average of 0.25 s, and the peak is located as the global maximum of the smoothed series after the first 0.5 s, which is excluded because it is dominated by the transient of the gate release rather than by the flow the analysis concerns. The three stages are then delimited as fractions of that local peak. The onset runs from the arrival of the front to the first instant at which the velocity falls below 60% of the peak; the transition runs from there to the instant at which it falls below a second, lower fraction; and the quasi-steady stage occupies the remainder of the filling. Defining the boundaries as fractions of the local peak rather than as absolute velocities or fixed times is what allows a single criterion to serve a zone in which the material descends at some $40\,\mathrm{mm\,s^{-1}}$ and a cell in which it creeps forward at less than 2.

One parameter of that criterion is set separately for the two matrices, and the reason is physical rather than procedural. The lower threshold, which marks the entry into the quasi-steady stage, is placed at 15% of the peak for Carbopol 0.2% and at 30% for Carbopol 0.5%. The stiffer gel carries a higher yield stress and arrests earlier, so its flow comes to rest at a residual velocity that remains a larger fraction of the peak it once reached, and a 15% threshold would locate the onset of the quasi-steady stage beyond the point at which the material has already stopped moving. The 30% threshold captures the same physical transition in that matrix that 15% captures in the softer one, and the asymmetry of the two numbers is what keeps the stages comparable rather than what breaks their comparability.

Two conventions complete the procedure. A minimum onset duration of 0.5 s is enforced, so that a zone whose velocity decays very rapidly is not assigned a degenerate first stage of a few frames. And a zone holding fewer than fifty valid control-case timestamps cannot resolve its own boundaries reliably, so it inherits them from the larger geometric region to which it belongs; such zones are marked with an asterisk in Figure 4.3. Zones whose velocity never falls below the lower threshold within the recorded window are assigned a quasi-steady stage extending to the end of that window.

Figure 4.3 shows the resulting segmentation and highlights a key feature relevant to the

subsequent analysis: the two matrices evolve on different timescales. Carbopol 0.5% takes roughly twice as long to reach the same state as Carbopol 0.2%, which is why the observation window of the former is 40 s against the 20 s of the latter. This difference also motivates defining the stages as fractions of the time to the local peak rather than at fixed times.

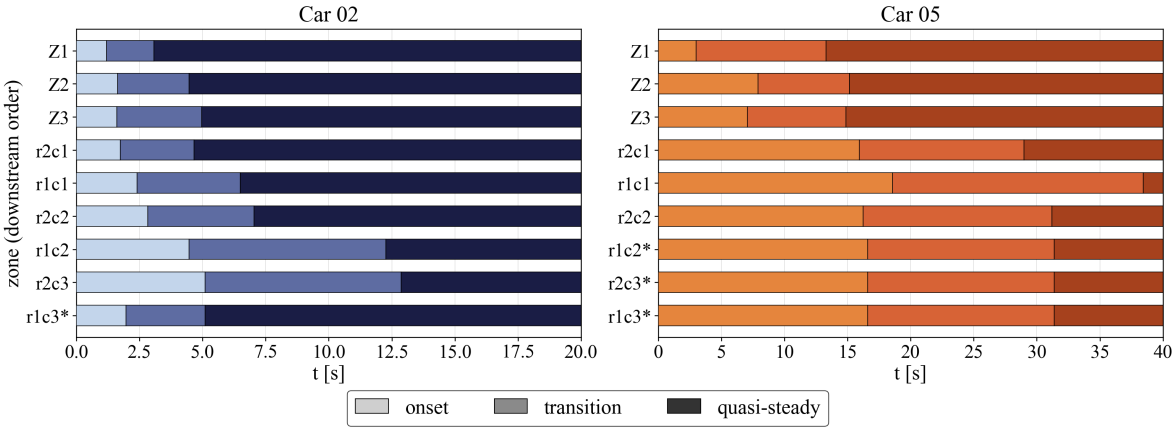


Figure 4.3: Flow stages by zone for the fiber-free control of each Carbopol matrix (Carbopol 0.2%, left; Carbopol 0.5%, right). Each horizontal bar is a zone, ordered from top to bottom along the downstream filling sequence, from the L-shape zones Z1–Z3 to the beam cells r$ic j$, and is segmented into the onset, transition and quasi-steady stages, shaded from light to dark within each matrix. Zones marked with an asterisk inherit their boundaries from their geometric region, too few control-case measurements being available to resolve them individually.

With the spatial zones and their flow stages established, the analysis can examine how the fiber content and the rheology of the carrier act within each zone-stage combination, using the effect-size protocol and significance criteria set out in Section 3.7.

4.3 PIV and PTV Global Analysis

This section characterizes the first-order effects of carrier rheology, mold geometry, and fiber concentration on the multiphase flow, before the influence of each variable is isolated within specific zones and flow stages. The Eulerian PIV fields describe how the fluid velocity and vorticity respond to rheology and fiber loading across the full domain, while the Lagrangian PTV data describe how the fibers themselves move and orient. Taken together, these two perspectives address the central question of the chapter:

how fiber concentration affects the suspension flow and through which mechanisms these effects arise

4.3.1 PIV Measurement Error

Table 4.3 summarises the total relative error for each camera and Carbopol concentration. The sub-pixel displacement error dominates in all cases, so the relative uncertainty is highest where mean velocities are lowest: Camera 3 reaches 36.2% for Carbopol 0.2%, attributable to light-sheet attenuation by the suspended fibers across the full mold width, while Camera 4 achieves the lowest uncertainty (10.0%) owing to the highest mean speed in the dataset. These figures bound the confidence of the velocity comparisons that follow and are consistent with the laminar, low-speed character of the flow.

Table 4.3: PIV total relative error by camera and Carbopol concentration. Percentage errors are relative to the mean flow speed $\bar{u}$.

Camera	**Carbopol [%]**	$\bar{u}$ [mm s^{-1}]	δu_{total} [%]
1	0.2	27.7	10.6
1	0.5	6.1	14.8
2	0.2	12.6	21.8
2	0.5	2.4	13.3
3	0.2	9.4	36.2
3	0.5	1.7	16.1
4	0.2	39.5	10.0
4	0.5	15.0	11.3

4.3.2 PIV Results

The Eulerian field is examined for the effect of the fiber content on two quantities: the fluid speed and the vorticity magnitude. An Analysis of Covariance with time as a covariate was run for each zone and flow stage, reporting the effect of the fiber content as the partial eta squared against the relevance threshold of Section 3.7. Table 4.4 summarizes the outcome by reporting, for each variable and stage, the largest effect found anywhere in the domain together with the zone in which it is found; the effect in every zone is given in Appendix E.2.

Table 4.4: Effect of fiber content on the two Eulerian variables, by flow stage and matrix. Cells report the peak η^2 across all spatial zones, with the zone of maximum effect in parentheses. The bottom rows report η^2 computed over the pooled domain. * $\eta^2 \geq 0.02$; ** $\eta^2 \geq 0.10$; *** $\eta^2 \geq 0.30$.

		Peak η^2 (zone)		
Carbopol	**Variable**	**Onset**	**Transition**	**Quasi-steady**
0.2 %	$\lvert\mathbf{u}\rvert$	0.165^{**} (r1c1)	0.102^{**} (Z2)	0.049^{*} (Z2)
	$\lvert\zeta\rvert$	0.283^{**} (Z1)	0.228^{**} (Z1)	0.061^{*} (r2c3)
0.5 %	$\lvert\mathbf{u}\rvert$	0.292^{**} (Z2)	0.434^{***} (Z2)	0.525^{***} (r1c1)
	$\lvert\zeta\rvert$	0.094^{*} (Z3)	0.095^{*} (Z3)	0.160^{**} (r2c2)
0.2 % (pooled)	$\lvert\mathbf{u}\rvert$	0.002	0.004	0.005
	$\lvert\zeta\rvert$	0.000	0.000	0.001
0.5 % (pooled)	$\lvert\mathbf{u}\rvert$	0.007	0.004	0.008
	$\lvert\zeta\rvert$	0.001	0.001	0.003

The fluid slows as the fiber content rises, in both matrices and at every stage, which is what the added hydrodynamic resistance of the fibers would produce. The effect size of Equation (3.13) gives only the magnitude of this response, not its sign, so the direction is read from the velocity fields of Figure 4.4. Where the effect concentrates it is strong, reaching $\eta^2 = 0.525$ in the cell r1c1 of the stiffer matrix. The pooled rows of Table 4.4 look much smaller, but they do not contradict this. Pooling averages the whole domain at once, and the domain contains regions as different as a plug descending the arm and a nearly arrested cell of the beam; that spread is far larger than the effect of the fibers, so it dominates the pooled figure and drives it toward zero. The pooled value therefore measures how small the fiber effect is next to the spatial variation of the flow, which is a property of the geometry and not of the fibers. The zonal values are the ones that answer the question, and they show the fibers slowing the carrier wherever it moves.

The vorticity behaves differently. Where the speed responds most strongly, in the stiffer matrix, the vorticity responds least: the effect on $|\mathbf{u}|$ rises to 0.525 while the effect on $|\zeta|$ never exceeds 0.160 anywhere in that matrix (Figure 4.5). The rotational field of the carrier is therefore largely indifferent to how many fibers it contains, even where its translational field is not.

The two matrices channel the effect differently, and the difference follows from their resistance to yielding. Carbopol 0.2% yields readily across most of the domain, so its

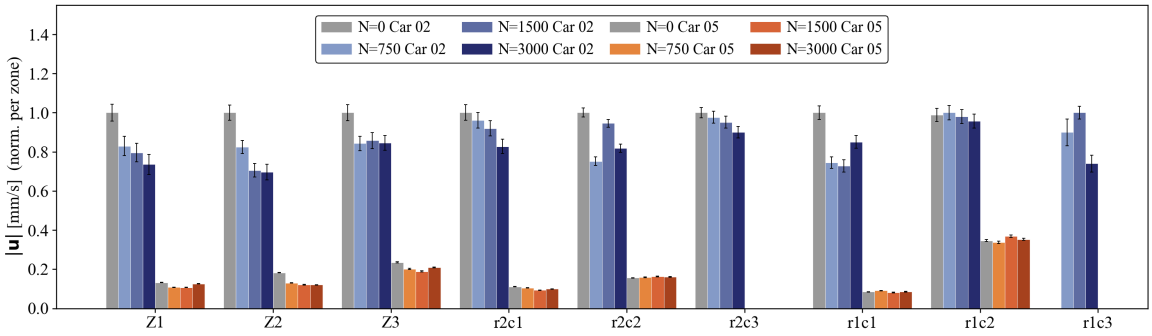


Figure 4.4: Spatial distribution of $|\mathbf{u}|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom).

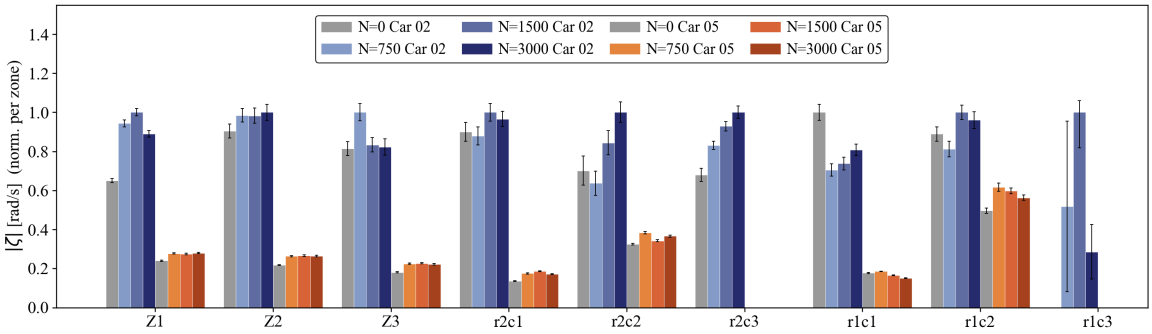


Figure 4.5: Spatial distribution of $|\zeta|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom).

translational field is set by gravity and is comparatively insensitive to the loading, and the fibers perturb instead the local rotation, which is why its largest effects appear in the vorticity of the upper arm where the descending flow first develops. Carbopol 0.5% operates closer to its yield threshold over wider regions and sustains larger plugs, so that adding fibers raises the effective consistency enough to shift whole regions between the yielded and the arrested state, and thereby to alter the local advance velocity itself. In the compliant matrix the fibers modulate the rotation; in the stiff one they act on the transport.

This reduction in velocity has important implications for the interpretation of the Lagrangian results presented in the following sections. Because strain accumulates more slowly in a slower flow, a mixture with a higher fiber concentration reaches a given position in the mold having undergone less deformation than a mixture with a lower fiber concentration. Consequently, comparing two mixtures at the same spatial location does

not necessarily correspond to comparing them at the same accumulated strain. Analyzing fiber orientation as a function of position would therefore confound two distinct effects: the direct effect of fiber concentration on fiber alignment and the indirect effect associated with differences in the amount of deformation accumulated by the flow. For this reason, Chapter 6 reports the orientation against the accumulated shear strain of Equation (2.12), computed for each run from its own velocity field, rather than against position in the mold.

PIV Reproducibility

The fiber-free mixtures are the only cases for which two PIV acquisitions were performed and therefore provide the only basis for assessing measurement repeatability. This assessment is important for establishing the level of variability against which the concentration effects reported above should be interpreted. An ANCOVA of the form $var \sim t + C(\text{take})$ was fitted to each cell to determine whether two acquisitions of the same batch produced distinguishable fields after accounting for their common temporal trend.

The two matrices show markedly different levels of repeatability. Carbopol 0.2% is highly reproducible, with significant differences in only about 15% of the cells. Carbopol 0.5% shows substantially greater variability, with vorticity differing between acquisitions in 44% of the cells and velocity magnitude in 28%. This lower repeatability is consistent with the greater sensitivity of the stiffer gel to its shear history and to the manner in which the plug is released. These results therefore suggest greater caution when interpreting the concentration effects observed for the stiffer matrix.

4.3.3 PTV Tracking Efficiency at High Concentration

The tracking efficiency of the experimental campaign is reported separately from that obtained in the validation tests of Section 4.1.2, as the two metrics represent different quantities. The validation tests measured the instantaneous detection ratio in a device containing a known and relatively small number of fibers, ranging from 25 to 800, and established the detection performance of the tracking framework under controlled conditions. In contrast, the efficiency reported here is defined as the fraction of fibers expected within the field of view that are successfully tracked during the quasi-steady

window of an experimental run at the fiber dosages considered in the campaign. It therefore indicates the effective sample size available for the analysis of each run. The resulting tracking efficiencies are reported in Table 4.5.

Table 4.5: PTV tracking efficiency by nominal fiber concentration and Carbopol matrix, computed as the ratio of detected to expected fibers within the field of view during the quasi-steady window.

Carbopol	**Nominal** N	**Detected fibers**	**Efficiency** [%]
0.2 %	750	162.5	29.7
0.2 %	1500	191.7	17.5
0.2 %	3000	168.7	7.7
0.5 %	750	227.6	41.6
0.5 %	1500	264.1	24.1
0.5 %	3000	244.7	11.2

The tracking efficiency decreases with fiber concentration in both matrices, from approximately 30–42% at the lowest dosage to below 12% at the highest. Occlusion, already identified in Section 4.1.2 as a cause of reduced detection efficiency, is compounded here by two additional effects associated with the denser experimental conditions. First, the larger number of candidate detections increases the ambiguity of the linking step. Second, increasing fiber concentration attenuates the incident light, reducing the illumination reaching individual fibers.

The Carbopol 0.5% matrix is tracked more efficiently than the Carbopol 0.2% matrix at every dosage, by approximately ten percentage points. Its lower flow velocity reduces the distance traveled by the fibers between analyzed frames, thereby decreasing linking ambiguity. In addition, the Carbopol 0.5% mixture spreads less, allowing a larger fraction of the fiber population to remain within the field of view. This difference reflects the measurement conditions rather than an intrinsic difference in tracking behavior between the suspensions and is one reason why the two matrices are not compared based on the number of recovered fibers.

The decrease in tracking efficiency must be accounted for in the analysis, because the fibers recovered at high concentrations constitute a smaller and potentially biased sample of the total population. This limitation is addressed through the resampling procedure described in Section 3.7.3, which is applied whenever trends with fiber dosage

are evaluated.

4.3.4 PTV Results

The same protocol from Section 4.3.2 was applied to the three Lagrangian variables, the fiber speed $|\mathbf{v}|$, the angular velocity ω, and the orientation order η_1. Table 4.6 summarizes the outcome, and the effect in every zone is given in Appendix E.3.

Table 4.6: Effect of fiber content on the three Lagrangian variables, by flow stage and matrix. Cells report the peak η^2 across all spatial zones, with the zone of maximum effect in parentheses. The bottom rows report η^2 computed over the pooled domain. * $\eta^2 \geq 0.02$; ** $\eta^2 \geq 0.10$; *** $\eta^2 \geq 0.30$.

		Peak η^2 (zone)		
Carbopol	**Variable**	**Onset**	**Transition**	**Quasi-steady**
0.2 %	$\lvert\mathbf{v}\rvert$	0.165^{**} (r2c2)	0.279^{**} (Z1)	0.396^{***} (r2c3)
	$\lvert\omega\rvert$	0.426^{***} (Z1)	0.502^{***} (Z1)	0.358^{***} (r2c2)
	η_1	0.284^{**} (Z2)	0.161^{**} (r2c1)	0.713^{***} (r2c1)
0.5 %	$\lvert\mathbf{v}\rvert$	0.161^{**} (Z1)	0.091^{*} (Z1)	0.368^{***} (r2c1)
	$\lvert\omega\rvert$	0.333^{***} (Z1)	0.327^{***} (Z1)	0.430^{***} (r2c1)
	η_1	0.407^{***} (Z1)	0.523^{***} (r2c1)	0.384^{***} (Z2)
0.2 % (pooled)	$\lvert\mathbf{v}\rvert$	0.001	0.032^{*}	0.022^{*}
	$\lvert\omega\rvert$	0.051^{*}	0.050^{*}	0.062^{*}
	η_1	0.017	0.000	0.045^{*}
0.5 % (pooled)	$\lvert\mathbf{v}\rvert$	0.006	0.005	0.007
	$\lvert\omega\rvert$	0.039^{*}	0.050^{*}	0.110^{**}
	η_1	0.035^{*}	0.028^{*}	0.009

The angular velocity is the variable most consistently modulated by the fiber content. It carries a relevant effect in nearly every zone and stage, reaching $\eta^2 = 0.502$ in the vertical arm of the softer matrix during the transition and $\eta^2 = 0.430$ in the cell r2c1 of the stiffer matrix once the flow has settled. It is also the only variable whose pooled effect survives the relevance threshold at every stage and in both matrices, which is to say that the fiber content raises the rotation of the fibers not merely in particular regions but across the domain as a whole. The orientation order climbs higher still where it climbs at all, to $\eta^2 = 0.713$ in the cell r2c1 of the softer matrix in the quasi-steady stage, which is the single largest response recorded anywhere in the dataset, but

it does not survive the pooling, and neither does the fiber speed: the speed vanishes when pooled in the stiffer matrix and the order parameter does so in the quasi-steady stage of the same matrix (Figures 4.6 to 4.8).

This rotation does not track the rotation of the carrier. The PIV of Section 4.3.2 showed the vorticity to be largely indifferent to the fiber content, with a pooled η^2 never exceeding 0.003, whereas the pooled effect on the fiber angular velocity reaches 0.110, a difference of more than an order of magnitude. The two quantities therefore respond to the fiber content on entirely different scales, and a fiber rotation supplied by the turning of the fluid, which proceeds at half the local vorticity, would require the fluid to turn faster as fibers are added, which it does not. The additional rotation is not inherited from the carrier, and its magnitude grows with the fiber content.

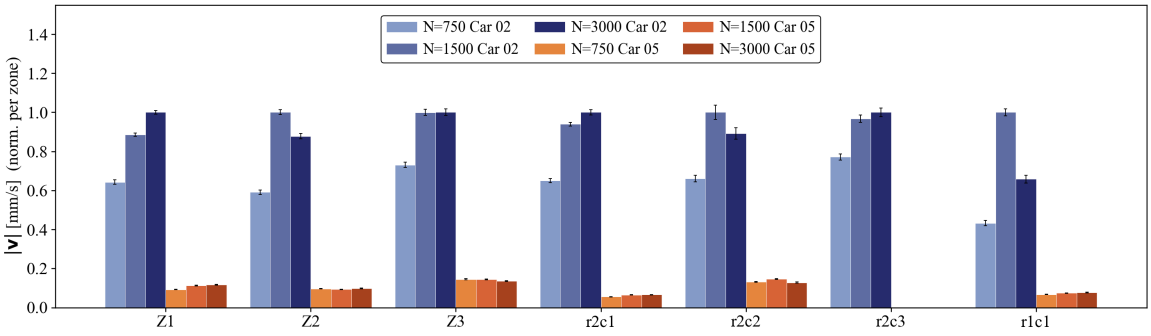


Figure 4.6: Spatial distribution of the fiber speed $|\mathbf{v}|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom).

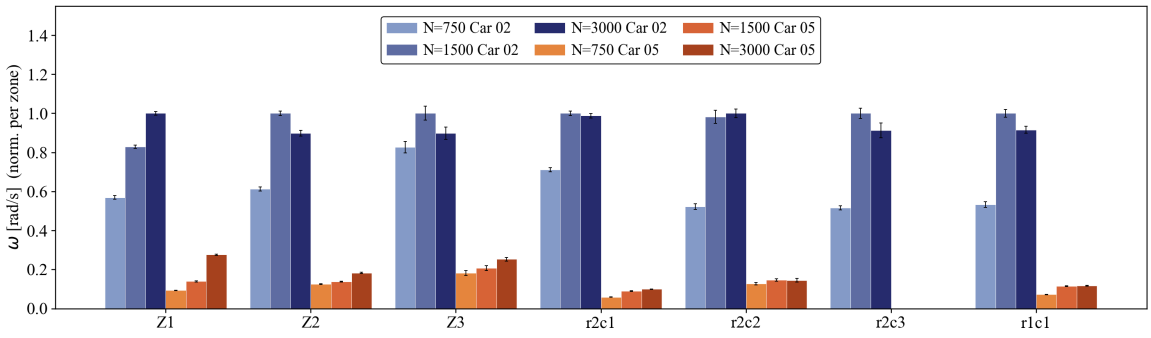


Figure 4.7: Spatial distribution of the fiber angular velocity $|\omega|$ across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom).

One feature of Figure 4.6 requires care. The fiber speed rises with the fiber content, which is the opposite of the fall the PIV records in the fluid speed of the same runs.

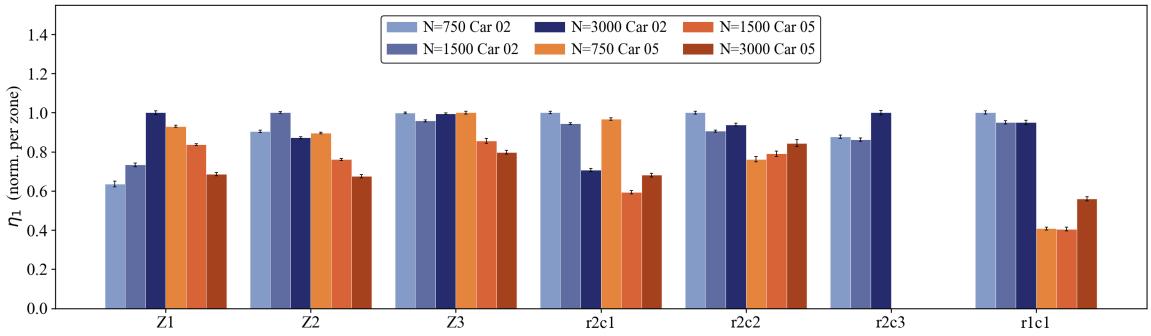


Figure 4.8: Spatial distribution of the orientation order η_1 across zones and fiber contents during the transition stage, for Carbopol 0.2% (top) and 0.5% (bottom).

Taken at face value this would have the fibers moving faster than the fluid that carries them, which cannot hold for particles that follow the flow passively, so the discrepancy is instrumental and not physical. The two techniques do not sample the same material. The PIV resolves the single plane lit by the laser sheet, whereas the PTV recovers fibers through the whole depth of the mold without resolving where in that depth each one lies, so the fibers are drawn in part from regions the velocity field never measures, including the faster core. The two speeds therefore cannot be compared directly, and the fiber speed is reported here only as a relative quantity across dosages within the PTV data itself. Whether the fibers lead or lag the fluid, and by how much, is a question the present measurement cannot settle, since it would require the velocity of the fluid and the velocity of the fiber at the same point and the same depth, which only a three-dimensional measurement or the CFD-DEM simulation the experiment is built to validate can supply. Establishing the true velocity coupling between the phases is accordingly left to that simulation.

4.3.5 Discussion

Three patterns emerge from the Eulerian and Lagrangian analyses together.

The first concerns the scale at which the fiber content acts. Within a single zone its effect is large, reaching $\eta^2 = 0.525$ on the fluid speed and 0.713 on the orientation order of the fibers. Averaged over the whole domain it nearly vanishes. The two figures do not conflict, because the domain spans conditions as different as a plug descending the arm and a nearly arrested cell of the beam, and that variation is far larger than the

effect of the fibers, so it dominates the average and drives it toward zero. The effect of the fiber content is thus regional, concentrated where the geometry forces a change in the flow, and a measurement averaged over the domain would miss it.

The second is a decoupling between the rotation of the fluid and the rotation of the fibers. The fiber content leaves the vorticity of the carrier almost untouched, with a pooled η^2 below 0.003 in both matrices, while raising the angular velocity of the fibers by a pooled η^2 of up to 0.110, a difference of more than an order of magnitude between two quantities that both measure rotation on the same flow. The fibers turn faster as their number grows, and the fluid that carries them does not. This is the central observation of the global analysis, and it locates the additional rotation outside the mean flow.

The third concerns the two matrices, which do not respond alike. In Carbopol 0.2% the fiber content acts chiefly on the rotational field, its largest Eulerian effect falling on the vorticity of the vertical arm; in Carbopol 0.5% it acts chiefly on the translational field, its largest effect falling on the speed at the elbow and in the first cells of the beam. The stiffer matrix operates closer to its yield threshold over wider regions, so an increment in effective consistency shifts whole regions between the yielded and the arrested state and alters the local advance velocity, whereas the softer matrix, which yields readily, has a translational field set by gravity and comparatively insensitive to the loading. The rheology therefore does not determine whether the fibers perturb the flow, but which component of it they perturb.

Taken together, these results establish that the fiber content is not a negligible parameter of the filling: it slows the carrier, it raises the rotation of the fibers, and it does both regionally rather than uniformly. What they leave open is what survives. The quantities examined so far are instantaneous, and an elevated rotation during filling does not by itself imply an altered orientation once the material has come to rest. The following section examines the final state of the fiber population, and relates it to the deformation the flow accumulated before the material arrested.

4.4 Fiber's Final State Analysis

This section measures the state the fibers settle into once motion has ceased. The orientation and the spatial arrangement of the population are recorded after the material has arrested, and the analysis is descriptive: it establishes which factors govern the final state, while the deformation that produces it is taken up in Chapter 6. The descriptors of every populated cell, the tests reported in summary below, and photographs of the arrested elements for every dosage and both matrices are given in Appendix F.

4.4.1 Hierarchy of Factors

An Analysis of Covariance was performed with the final zone, the rheological matrix, and the fiber content as predictors of the angular deviation δ and the dispersion index Λ. The angular deviation replaces the order parameter here for the reason given in Section 3.7.1: δ supplies one observation per fiber and so admits a test within a single cell, whereas η_1 supplies one per cell. The two are inverse descriptions of the same state, so a rise in the effect on δ is a loss of order; the figures report η_1 because it is the quantity the saturation law predicts. Table 4.7 reports the result.

Table 4.7: Effect size η^2 of the principal determinants of the fiber final state, from the global Analysis of Covariance over the pooled domain. The relevance threshold of Section 3.7 is $\eta^2 = 0.02$.

Factor	η^2 **orientation** (δ)	η^2 **dispersion** (Λ)
Final zone (geometry)	0.0310	0.1880
Rheological matrix	0.0014	0.0290
Fiber content	0.0007	0.0001

The hierarchy is unambiguous and is the same for both responses. The zone in which a fiber comes to rest dominates, exceeding every other factor by at least a factor of six and, for the orientation, by more than twenty. The matrix is second, and is the only other factor to cross the relevance threshold, which it does for the dispersion alone. The fiber content, pooled over the dataset, is negligible for both. The pooled figure understates it, however, for the reason given in Section 4.3.5: an influence confined to a few zones is averaged away when the domain is pooled, and Section 4.4.2 therefore

repeats the test within each matrix.

Figures 4.9 and 4.10 show the hierarchy cell by cell. The variation is between cells and between matrices more than between dosages: the bars of the two matrices sit at clearly different levels within most cells, while the three bars of a single matrix, which differ only in fiber content, stay closer together.

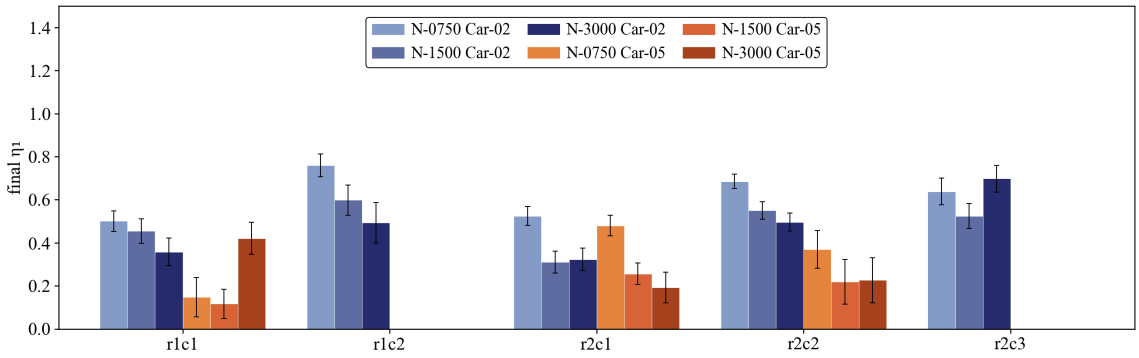


Figure 4.9: Final orientation order parameter η_1 in each beam cell, grouped by fiber content N and rheological matrix (blue tones: Carbopol 0.2%; orange tones: Carbopol 0.5%). Error bars span one standard deviation across the fibers of each group; a colour missing from a cell indicates that no fibers of that matrix were tracked there.

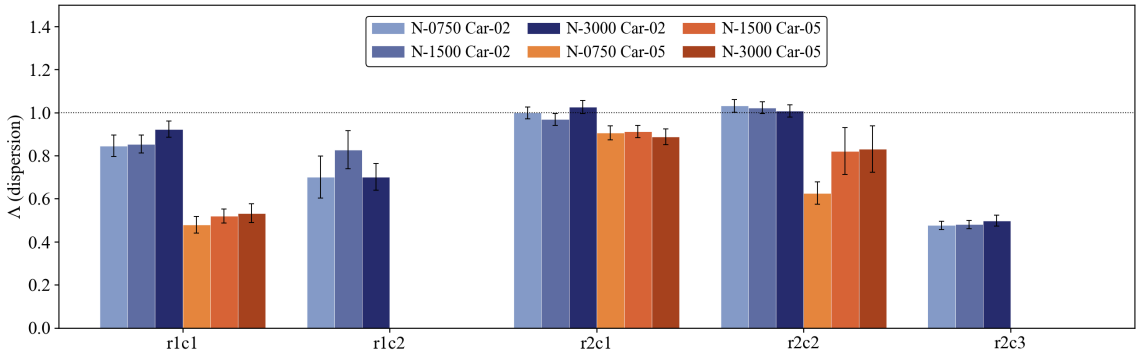


Figure 4.10: Final centroid dispersion index Λ in each beam cell, grouped by fiber content N and rheological matrix, with the same colour convention as Figure 4.9. The dotted line marks $\Lambda = 1$, the value of a spatially uniform arrangement; lower values indicate clustering. Error bars span one standard deviation across the fibers of each group.

Within the beam, the spatial dispersion is governed by the rheology. Carbopol 0.2% reaches Λ at or slightly above the uniform value in the cells of the lower row, r2c1 and r2c2, whereas Carbopol 0.5% remains below unity in every populated cell, and within

the beam the fiber content moves neither. The contrast follows from how each matrix arrests. The stiffer gel stops flowing while still advancing and locks the fibers wherever they happen to be, preserving whatever clustering they carried out of the L-shape. The softer gel keeps flowing for longer, so the fibers are carried and rearranged more completely before the material sets, and the widening flow spreads them slightly past a uniform arrangement, pushing Λ above one. That value is not an artifact: it indicates fibers sitting closer to the cell boundaries than a uniform set would.

The vertical arm is the exception, and it is stated rather than folded into the general conclusion. In Z1, and in the softer matrix alone, the fiber content carries an effect on the dispersion of $\eta^2 = 0.302$ with $p = 0.003$, the largest effect the fiber content exerts on any quantity anywhere in this dataset, and it is not compatible with the statement that the arrangement of the centroids is a matter of rheology alone (Appendix F.3). Z1 is also the only zone in which the rheological effect on the dispersion reverses sign with the dosage. It is the zone in which the material descends before it has begun to spread, in which the fibers are still arranged much as they were introduced, so it is the region in which the initial distribution is least modified by the flow. Whether the effect measured there describes the deposition or the introduction is not settled by these data, and the claim that the reinforcement is distributed by the rheology and not by its own quantity is therefore made of the beam, where the element is cast, and not of the arm through which it passes on the way.

4.4.2 Isolated Effect of Fiber Content

The negligible pooled effect of Table 4.7 does not establish that the fiber content is irrelevant, for the reason already given in Section 4.3.5: geometry is an order of magnitude stronger, and an influence confined to a few zones is averaged away once the dataset is pooled. The Analysis of Covariance was therefore repeated separately within each matrix, with the final zone retained as a covariate, so that the effect of the fiber content is measured net of the dominant geometric one. Table 4.8 reports the outcome. The two responses separate. The dispersion does not respond to the fiber content within either matrix: no significant effect, and an effect size indistinguishable from zero. The orientation does respond. The effect grows by an order of magnitude over its pooled value and is significant in both matrices, and the angular deviation rises with the fiber

Table 4.8: Net effect of fiber content on the final state, isolated from the geometric effect within each rheological matrix. The relevance threshold of Section 3.7 is $\eta^2 = 0.02$, which neither orientation entry reaches.

	Orientation (δ)		**Dispersion (Λ)**	
Carbopol	η^2	p	η^2	p
0.2 % (soft)	0.0083	4×10^{-5}	0.0003	n.s.
0.5 % (stiff)	0.0041	0.017	0.0000	n.s.

content at fixed rheology, which is to say that the fibers align less well as their number grows. The trend survives the resampling control of Section 3.7.3, so it is not an artifact of the unequal detection statistics.

This effect must be read in two respects, one of size and one of significance, and the two do not point the same way. Its magnitude falls below the relevance threshold of $\eta^2 = 0.02$ that Section 3.7 sets, so by the criterion adopted throughout this work the pooled effect is not interpreted on the strength of its size alone. Its significance, its sign, and its consistency are another matter. The direction is the same in both matrices, the effect is significant in both, and it is the direction the interaction mechanism predicts, so what the aggregate establishes is not the size of the effect but the reality and the sign of its trend. The size is recovered, and the threshold crossed, only once the final orientation is read against the deformation that produced it rather than against the cell in which it was measured, which is the confrontation of Chapter 6; the present section establishes the direction of the effect, and the next establishes its magnitude in the variable that resolves it.

The magnitude must further be read against the dosages that produced it. The fiber contents examined here span from 0.7 to 2.9 times the dilute limit of Equation (2.4), which places them barely across the threshold at which fiber-fiber interactions begin to operate, and one to two orders of magnitude below the dosages of a real UHPC mix. That the effect is small is therefore what this regime should produce, and a small but consistent and significant effect measured at the lower edge of the semi-dilute range is a lower bound on the effect to be expected at practical fiber contents, where the interactions are correspondingly stronger.

Having established what the final state consists of, that geometry and rheology govern it, and that the fiber content leaves a small trace on the orientation and none on the

arrangement, one question remains open, and the rest of the analysis depends on it. Every result reported so far is attached to a place: the alignment measured in a beam cell, the rotation measured in the elbow, the effect of the dosage within a zone. But a fiber does not acquire its orientation in the cell where it comes to rest; it acquires it along the path that carried it there, and the deformation that set it was accumulated over that whole route and not at its end. Relating the final orientation to the deformation that produced it therefore requires the path of each arrested fiber, from where it entered to where it stopped.

4.5 Lagrangian Reconstruction of the Fiber Paths

The PTV does not supply these paths directly. Since the fraction of fibers followed from entry to arrest falls below one percent at the loadings of the campaign (Section 4.1.2), each fiber is recovered as a series of short trajectory fragments, and its path is reconstructed by stitching those fragments forward in time, joining a fragment that ends to one that begins shortly afterwards whenever the two are consistent in the interval between them, in the position reached by advecting the fiber forward from its measured velocity, in the orientation predicted from its measured angular velocity, and in the detected length that fixes the fiber size. The procedure is repeated until a chain spans the path from the vertical arm, through the elbow and the outlet, to the beam.

The tolerances are set conservatively everywhere except across the transition from the L-shape into the beam, where detection is interrupted and the gap between fragments is wider, so that a doubtful join is refused rather than forced: a missing link ends a path early, whereas a wrong one would fabricate a trajectory no fiber followed. The reconstructed paths are consequently shorter than the true ones and biased toward the fibers easiest to follow, but the provenance they establish is reliable, agreeing on the same population with two independent attributions of each fiber's origin, the modal zone in which its fragments were observed and the zone reached by integrating its final position backward along the measured PIV field.

Figure 4.11 shows the routes recovered in this way, one row per destination cell. The L-shape is a common trunk, since by the geometry of the mold no fiber reaches the beam without traversing it in the order Z1, Z2, Z3, and the routes separate only afterwards, by

destination. The contrast between the matrices governs what follows: Carbopol 0.2% populates five of the six beam cells and reaches the far end, whereas Carbopol 0.5% populates only three, clustered near the entry, and never reaches the third column, because the stiffer gel arrests while the beam is still half empty.

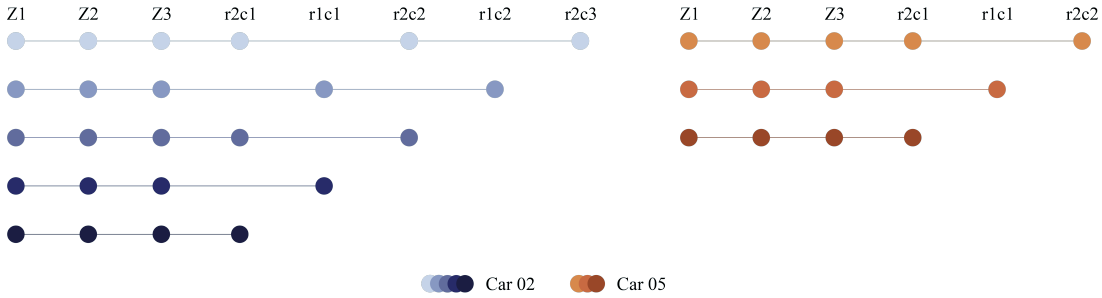


Figure 4.11: Routes followed by the fibers reaching each cell of the beam, reconstructed by stitching the PTV trajectory fragments. Each row is one destination cell, and the marked zones are those its fibers traversed on the way. Left: Carbopol 0.2%. Right: Carbopol 0.5%. The three dosages are pooled.

Two consequences follow, and together they fix the domain within which the saturation law can be traced. The first concerns the order in which the fibers enter the beam. The material does not flow smoothly from the outlet but falls and lands, and the impact disorders the population, breaking whatever alignment the fibers carried out of the L-shape as they strike the accumulating deposit. The upstream deformation therefore does not carry forward, and the clock of the saturation law starts at the beam entry rather than at the mold inlet, which licenses the two-parameter form in which the population is taken to enter disordered, $\eta_{1,0} \approx 0$. That value is an assumption of simplicity and identifiability rather than a measurement: the tracking is planar while the angular state at the landing is three-dimensional, so the entry order cannot be observed, and with so few arrest points per dosage it could not be identified alongside the ceiling and the strain scale in any case.

The evidence for the reset is in Figure 4.12, which supports it without measuring the entry order itself. It follows the orientation order along the mean path against the strain accumulated from the L-shape entry: the order climbs to a peak at Z3, the last zone before the outlet, and drops at the first beam cell even though that cell sits at larger accumulated strain. A population that carried its Z3 alignment across the transition would begin high and relax slowly, whereas one that begins lower, and at the higher

dosages recovers through the beam from that lower start, is a population that lands disordered and realigns within the beam. The order shown in each beam cell is that of the fibers arrested there and not of the landing that precedes it, since between landing and arrest the fibers accumulate strain and regain some alignment, so the true entry state lies below the first cell and is not itself resolved by the figure.

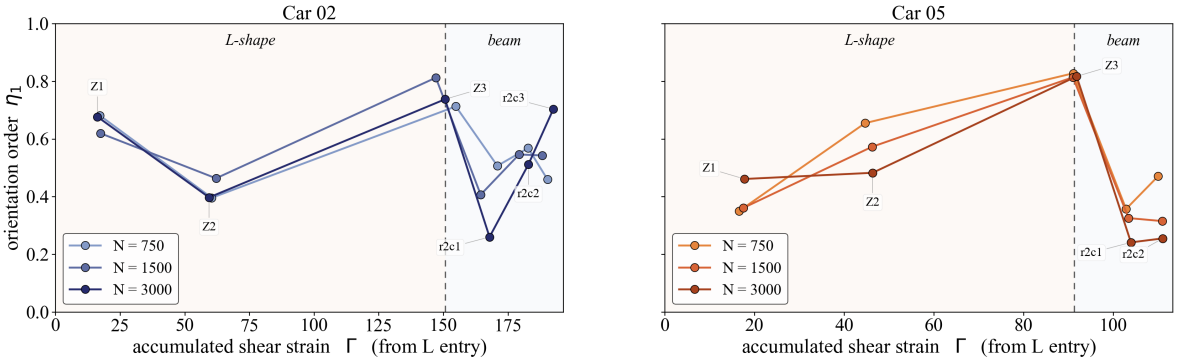
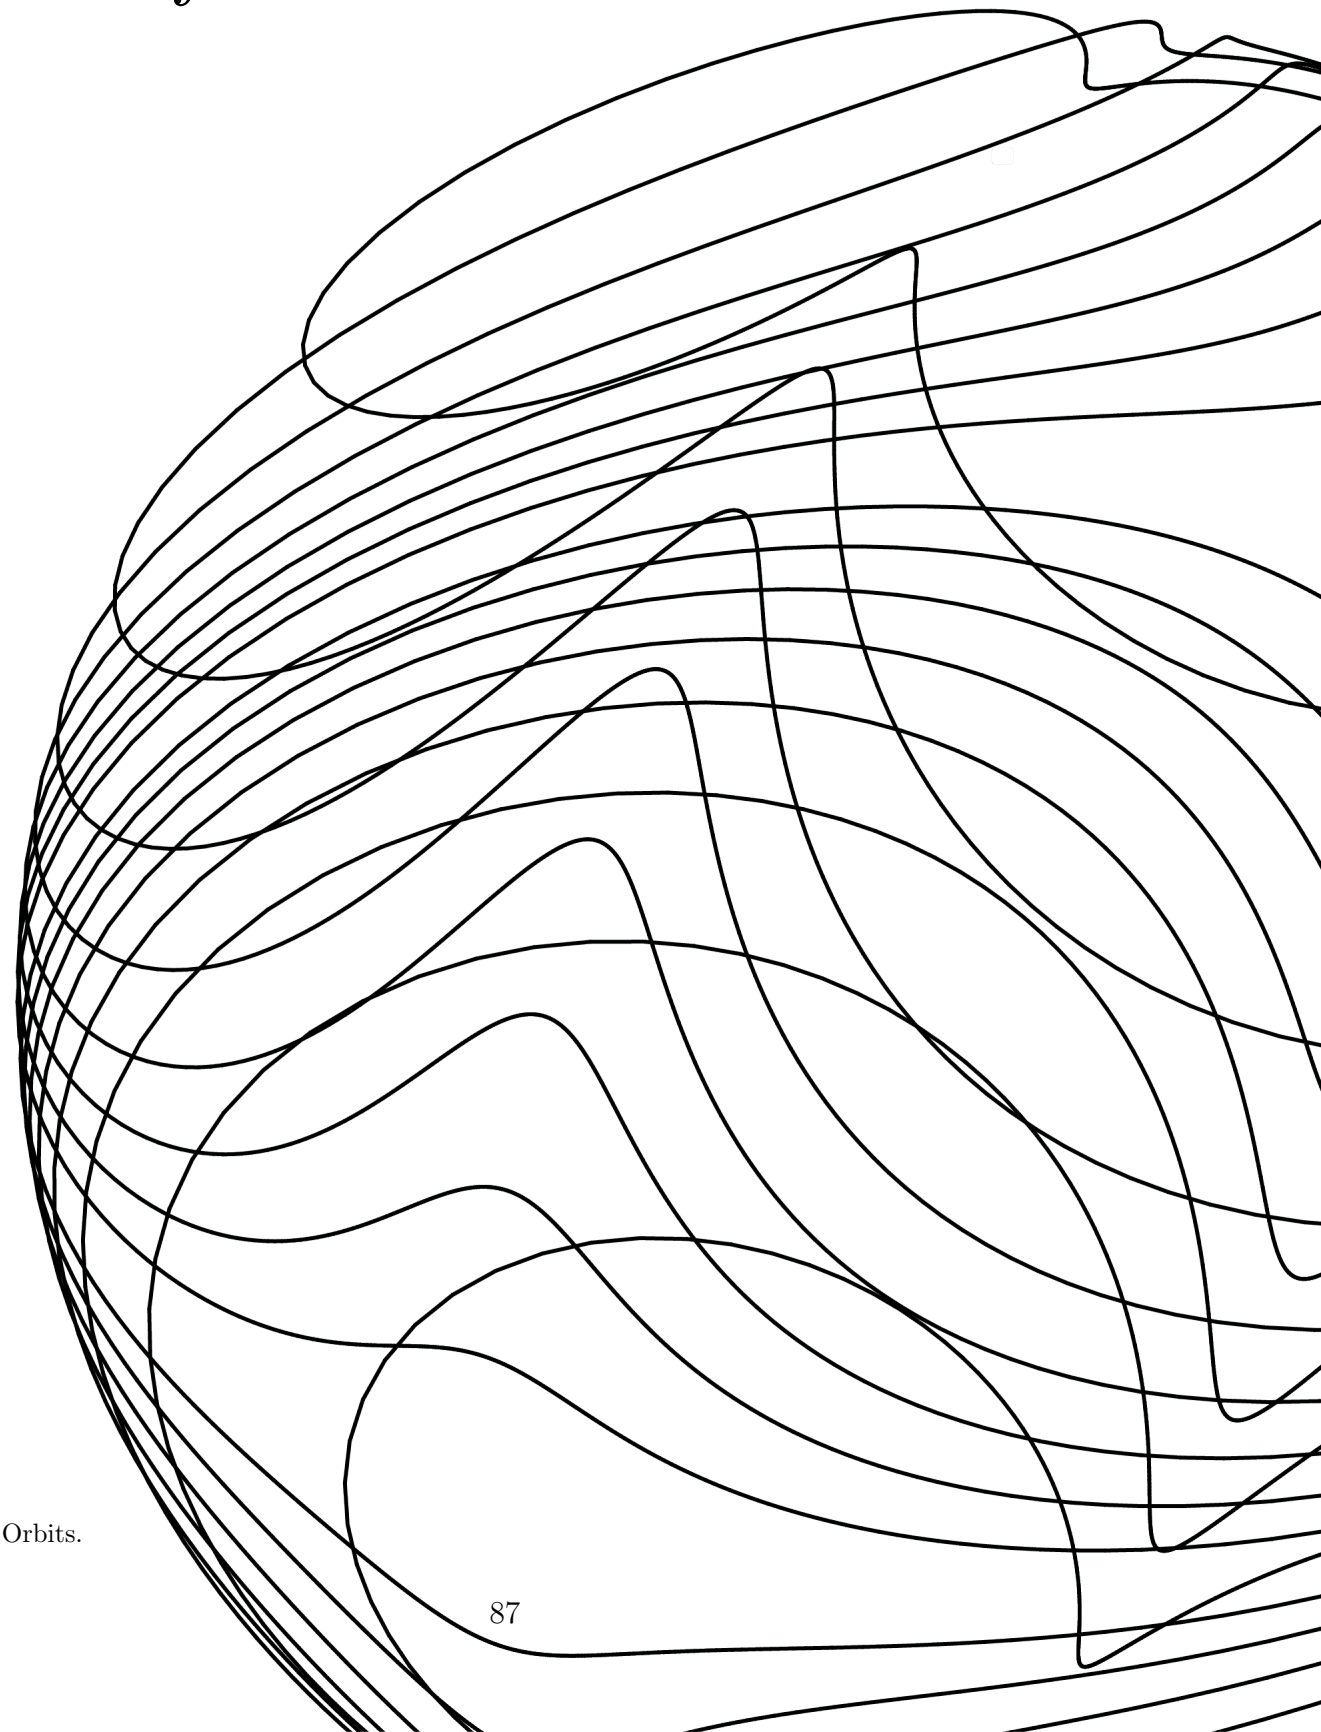


Figure 4.12: Orientation order η_1 along the mean path, by zone, against the shear strain Γ accumulated from the entry to the L-shape, for the three fiber dosages. Zones Z1 to Z3 lie in the L-shape and the cells r2c1 onward in the beam, separated by the dashed line. The order rises to a peak at Z3 and drops at the first beam cell, which sits at larger accumulated strain, the drop marking the disordering impact where the material falls from the outlet into the beam. The value shown for each zone is that of the fibers arrested or observed within it, not of the landing that precedes the first beam cell, which the planar measurement does not resolve. The strain is referred here to the L-shape entry, whereas the saturation law of Chapter 6 is fitted against strain referred to the beam entry, downstream of this reset. Left: Carbopol 0.2%. Right: Carbopol 0.5%.

The second consequence concerns the range of deformation attained by each matrix. Only the Carbopol 0.2% matrix spans a sufficiently wide range of accumulated strain to characterize the saturation curve. Its fibers reach five of the six beam cells and extend over the full length of the beam, allowing the arrest points to sample the curve from its initial rise toward saturation. In contrast, the Carbopol 0.5% matrix arrests before the beam is fully filled, so its arrest points are confined to the low-strain region of the curve and provide insufficient information to characterize its overall shape. Together with the limited repeatability of the Carbopol 0.5% baseline (Section 4.3.2), this provides two independent reasons for restricting the accumulated-deformation analysis to Carbopol 0.2%, as done in Chapter 6. The analysis therefore compares fiber dosages within a single matrix, under fixed rheological and geometric conditions.

Chapter 5

Fiber Dynamics



Jeffery Orbits.

Chapter 4 showed that the fibers rotate faster as their number grows while the vorticity of the carrier remains nearly unchanged, and that their final alignment decreases with the fiber content. This chapter develops the theoretical framework that explains both observations. Starting from the Jeffery equation for a single fiber, it adds the Folgar-Tucker diffusion and the Advani-Tucker orientation tensor, reduces that tensor to the measured order parameter η_1, and derives the exponential law of approach to equilibrium whose two parameters, $\eta_{1,\max}$ and Γ^*, are tested in Chapter 6. The complete derivation is given in Appendix A.

5.1 Jeffery's Orbits

The description of fiber orientation in a flowing suspension begins with the motion of a single fiber immersed in a fluid. Jeffery (Jeffery, 1922) solved this problem for a rigid ellipsoidal particle suspended in a Newtonian fluid at vanishing Reynolds number, obtaining the rotation of the particle axis as a function of the local velocity gradient alone. Representing the fiber by a unit vector $\boldsymbol{p}$ directed along its axis, and splitting the velocity gradient into its symmetric part, the rate-of-strain tensor $\boldsymbol{D}$, and its antisymmetric part, the rotation tensor $\boldsymbol{W}$, whose components are minus one half of those of the vorticity vector, as given by Equation (2.11), the Jeffery equation reads:

$$\frac{\mathrm{d}\boldsymbol{p}}{\mathrm{d}t} = \boldsymbol{W}\cdot\boldsymbol{p} + \xi\left(\boldsymbol{D}\cdot\boldsymbol{p} - (\boldsymbol{p}\cdot\boldsymbol{D}\cdot\boldsymbol{p})\,\boldsymbol{p}\right), \tag{5.1}$$

where $\xi = (r_f^2 - 1)/(r_f^2 + 1)$ is a shape factor built from the fiber aspect ratio r_f, which approaches unity for slender fibers such as those employed in UHPC.

The two terms on the right-hand side act in different ways, and the distinction between them governs everything that follows. The deformation term tilts the fiber toward the direction along which the fluid is being stretched, so it is this term, and not the rotation of the fluid, that aligns the fiber with the flow. The rotation term spins the fiber at the local rotation rate of the fluid, half the vorticity, as if it were embedded in a rigidly rotating fluid, and by itself selects no preferred direction: it moves the fiber without ordering it.

The competition between the two produces the behavior that gives the model its name. The deformation term drives the fiber toward the direction of principal stretching, but

it does so at a rate that vanishes as the fiber approaches it, while the rotation term continues to turn the fiber at an undiminished rate however it is oriented. Alignment is therefore never quite reached: the fiber slows as it swings toward the flow direction, lingers there while the deformation term has almost nothing left to do, and is then carried through the transverse orientation by the rotation of the fluid, at which point the cycle begins again. The trajectory is closed and periodic, and it is known as a Jeffery orbit. The period,

$$T = \frac{2\pi}{\dot{\gamma}} \left(r_f + \frac{1}{r_f} \right), \tag{5.2}$$

depends only on the shear rate and on the aspect ratio, and the slenderness of the fiber enters it directly: a long rod spends most of its period nearly parallel to the flow and flips through the transverse orientation in a small fraction of it. The particular orbit a fiber follows is fixed by the orientation it held when the flow started, and no subsequent deformation can move it to a different one. A dilute Newtonian suspension therefore never reaches a steady orientation state: it tumbles without end, and it never forgets how it was set in motion.

Two of the hypotheses on which this result rests are not met by the system studied in this work. The first concerns the suspending fluid, since the Jeffery solution assumes a Newtonian medium of constant viscosity, whereas the fluid studied here is viscoplastic and does not flow at all below a yield stress, a property fresh UHPC shares with the Carbopol that stands in for it; the Carbopol tested is, in addition, shear-thinning, its apparent viscosity falling as the local shear rate rises. The second concerns the concentration, since the solution describes an isolated fiber in an unbounded fluid, an idealization valid only in the dilute regime, whereas the fibers of a real casting lie close enough for their disturbances of the surrounding flow to overlap and for their motions to cease being independent. The dosages of a practical UHPC mix sit near the boundary of the concentrated regime, at $\phi \gtrsim r_f^{-1}$, where the fibers are in near-permanent contact; the experiments of this work operate far below that, at the lower edge of the semi-dilute regime, barely across the dilute limit $\phi_{sd} = r_f^{-2}$, where the interactions are present but weak. Both lie above the dilute idealization, and both therefore depart from Jeffery in the same way, though to very different degrees. Both departures act on the same feature of the solution. The closed orbit does not survive a real casting, and neither does the memory of the initial orientation it carries: the encounters between fibers perturb

the trajectory at every turn, and no periodic motion is recovered. What does survive, and what remains the physical basis of the model developed herein, is the mechanism carried by the deformation term, namely that the fiber is driven toward the direction of principal stretching at a rate set by the rate of strain of the surrounding fluid.

A usable model must therefore retain the alignment mechanism while accounting for the loss of the deterministic orbit. Folgar and Tucker (Folgar et al., 1984) proposed exactly this, extending Jeffery's description to concentrated suspensions by superimposing a stochastic contribution on the deterministic rotation. Rather than resolving individual fiber-fiber interactions, their model represents their cumulative effect through a single interaction parameter, which introduces an orientation-space diffusion that counteracts the ordering imposed by the flow and progressively erases the memory of a deterministic Jeffery orbit.

5.2 Folgar-Tucker Diffusion Model

The idea behind the Folgar-Tucker extension is best understood as a competition between flow-induced ordering and interaction-induced randomization (Folgar et al., 1984). In a fiber suspension, order corresponds to fibers preferentially aligning along particular directions, whereas disorder corresponds to their orientations being more broadly distributed in orientation space. In the dilute limit, the motion of an isolated fiber is governed by the local flow and follows a deterministic Jeffery orbit. As the fiber concentration increases, however, neighboring fibers increasingly disturb this motion. Each encounter produces a small perturbation whose direction cannot be predicted from the local velocity gradient alone. Accumulated over the many interactions experienced by a fiber as the suspension flows, these perturbations act as a randomizing mechanism on the population, continually competing with the alignment imposed by the flow.

Folgar and Tucker represented this agency as a rotary diffusion, in direct analogy with the Brownian diffusion of molecules, but driven by hydrodynamic encounters between fibers rather than by thermal agitation. Its strength is set by a single dimensionless interaction coefficient C_I, which measures how strongly the fibers interfere with one another. Since the frequency of the encounters grows as the fibers are packed more closely, C_I increases with the fiber concentration, and vanishes in the dilute limit,

where the fibers no longer meet and the deterministic Jeffery description is recovered. The magnitude of the resulting diffusion is written as $C_I\dot{\gamma}$, with the scalar shear rate

$$\dot{\gamma} = \sqrt{2\,\boldsymbol{D}:\boldsymbol{D}}, \tag{5.3}$$

so that the randomizing influence, like the ordering one, is proportional to the intensity with which the fluid is being deformed. A suspension at rest neither aligns its fibers nor randomizes them: both mechanisms are switched on by the flow and switched off with it.

The consequence for a cast UHPC element is direct. As the fiber dosage increases, fiber-fiber interactions become more frequent and the effective interaction coefficient C_I is correspondingly expected to increase, strengthening the rotary diffusion of the orientation distribution. A larger fraction of the alignment generated by the flow is therefore counteracted by interaction-induced randomization before a strongly oriented state can develop. For a given deformation history, a suspension with a higher fiber content consequently tends to attain a lower degree of alignment. This provides a physical explanation for an effect that is well documented experimentally: increasing the fiber content does not necessarily enhance alignment with the flow, but can instead reduce it (Ferrara et al., 2008; Ravichandran et al., 2022).

The rotary diffusion describes the fate of a population rather than the trajectory of any individual fiber, which makes the single-fiber equation a poor vehicle for it. Advani and Tucker (Advani et al., 1987) recast the Jeffery rotation together with the Folgar-Tucker diffusion as an evolution equation for a statistical descriptor of the whole population, and it is in this form that the model is used in the present work.

5.3 Advani-Tucker Model

Rather than tracking every fiber, the Advani-Tucker formulation (Advani et al., 1987) summarizes the orientation state of the population at a point through the second-order orientation tensor, defined as the average of the outer product of the fiber direction with itself,

$$\boldsymbol{A} = \langle \boldsymbol{p} \otimes \boldsymbol{p} \rangle \,, \qquad A_{ij} = \langle p_i \, p_j \rangle \,, \tag{5.4}$$

where $\langle \cdot \rangle$ denotes the average taken over all the fibers present at that point. The outer product, and not $\boldsymbol{p}$ itself, is the object averaged, because a rod has no distinguished end: $\boldsymbol{p}$ and $-\boldsymbol{p}$ describe the same physical fiber, and any descriptor built from them must be blind to the sign, which $\boldsymbol{p} \otimes \boldsymbol{p}$ is and $\boldsymbol{p}$ is not. Averaging the vector itself would return zero for a population of fibers lying along a common axis but pointing in opposite senses, which is a perfectly aligned population.

Written out in three dimensions,

$$\boldsymbol{A} = \begin{pmatrix} \langle p_1^2 \rangle & \langle p_1 p_2 \rangle & \langle p_1 p_3 \rangle \\ \langle p_1 p_2 \rangle & \langle p_2^2 \rangle & \langle p_2 p_3 \rangle \\ \langle p_1 p_3 \rangle & \langle p_2 p_3 \rangle & \langle p_3^2 \rangle \end{pmatrix} , \tag{5.5}$$

each entry admits a direct reading. The diagonal terms are the fractions of orientation pointing along each coordinate direction, so that A_{11} measures how much of the population lies along the first axis, A_{22} along the second, and A_{33} along the third. Because every fiber is a unit vector, these fractions always add up to one, $\operatorname{tr} \boldsymbol{A} = 1$, and the diagonal behaves as a fixed budget of orientation distributed among the available axes. The off-diagonal terms measure the correlation, or the tilt, between pairs of directions: a nonzero A_{12} indicates that the fibers are inclined within the 1-2 plane instead of lying along either of its axes. The tensor is symmetric by construction, since $\langle p_i p_j \rangle = \langle p_j p_i \rangle$, so it carries six distinct components in three dimensions and three in the plane – five and two independent, respectively, once the unit-trace constraint is imposed.

Two limiting states fix the meaning of the descriptor, and they are written here for a space of d dimensions, with $d = 3$ in the bulk and $d = 2$ in the plane to which the measurements of this work are confined. If every fiber is perfectly aligned with the first coordinate direction, the whole orientation budget is assigned to that axis and the correlations vanish, so that the leading diagonal entry is unity and every other entry is zero. If instead the fibers point in random directions, no axis is preferred, the budget

is shared equally, and the correlations vanish again by symmetry,

$$\boldsymbol{A}_{\text{iso}} = \frac{1}{d}\,\boldsymbol{I}. \tag{5.6}$$

Any real suspension lies between these two extremes, and the displacement of the leading diagonal entry from $1/d$ toward unity is the quantitative measure of how far the flow has succeeded in aligning the fibers.

The evolution of this descriptor under flow is governed by the Advani-Tucker equation, which combines the deterministic Jeffery rotation with the Folgar-Tucker diffusion,

$$\frac{\mathrm{D}\boldsymbol{A}}{\mathrm{D}t} = \underbrace{(\boldsymbol{W}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\boldsymbol{W})}_{\text{rotation}} + \underbrace{\xi\,(\boldsymbol{D}\cdot\boldsymbol{A} + \boldsymbol{A}\cdot\boldsymbol{D} - 2\,\mathbb{A}\!:\!\boldsymbol{D})}_{\text{strain rate}} + \underbrace{2\,C_I\,\dot{\gamma}\,(\boldsymbol{I} - d\,\boldsymbol{A})}_{\text{interaction}}, \tag{5.7}$$

where $\mathrm{D}/\mathrm{D}t = \partial/\partial t + \boldsymbol{u}\cdot\boldsymbol{\nabla}$ is the material derivative, the rate of change following a fluid element, and $\mathbb{A} = \langle \boldsymbol{p}\otimes\boldsymbol{p}\otimes\boldsymbol{p}\otimes\boldsymbol{p}\rangle$ is the fourth-order orientation tensor, whose presence requires a closure approximation expressing it as a function of $\boldsymbol{A}$. The tensor algebra of each term, together with the closure, is developed in Appendix A. In words, Equation (5.7) states that the orientation of the population changes for three reasons: the fluid rotates it, the fluid stretches it, and the fibers disturb one another.

The structure of the interaction term follows from the reading of $\boldsymbol{A}$ given above. A diffusion must drive the system toward its most disordered configuration, which is the isotropic state $\boldsymbol{A}_{\text{iso}}$ of Equation (5.6), and it must therefore be proportional to the distance separating the current state from it, $\left(\frac{1}{d}\boldsymbol{I} - \boldsymbol{A}\right)$, which is $\frac{1}{d}\,(\boldsymbol{I} - d\,\boldsymbol{A})$ once the constant is absorbed into $2C_I\dot{\gamma}$. This is the origin of the identity tensor and of the factor d appearing in Equation (5.7). Substituting the isotropic state $\boldsymbol{A} = \frac{1}{d}\boldsymbol{I}$ gives $\boldsymbol{I} - d\cdot\frac{1}{d}\boldsymbol{I} = \boldsymbol{0}$, so the term switches itself off once the population is fully disordered, as any diffusion must. Conversely, the more ordered the suspension, the further $\boldsymbol{A}$ lies from $\frac{1}{d}\boldsymbol{I}$ and the harder the interactions push back. The strain-rate term acts in the opposite sense, driving the population toward alignment with the flow, while the rotation term only turns it as a whole, so that Equation (5.7) is a statement of competition: the flow orders, the interactions disorder, and both act at once on the same tensor.

The fiber content enters this competition through C_I alone, and its effect can be read directly from the equation. As the dosage increases, the interaction term gains weight relative to the two flow terms, the disordering contribution claims a larger share of

the balance, and the same casting flow yields a lower degree of alignment. Wherever the population carries some order, the interaction term is nonzero and grows with C_I, so the magnitude of $\mathrm{D}\boldsymbol{A}/\mathrm{D}t$ increases with the dosage: the orientation state evolves faster, which is to say that the fibers reorient faster, although the additional rotation is randomizing rather than ordering. Concentration does not slow the fibers down, it redirects them away from the flow direction. The rotation term, by contrast, is left untouched, since $\boldsymbol{W}$ is a property of the velocity field and the fiber content enters the model only through C_I; the interactions modify the balance between order and disorder without altering the rigid rotation that the flow imposes on the population.

The two influences pull in opposite directions and act simultaneously, one supplied by the flow and the other by the fibers themselves. Neither can prevail indefinitely, and an equilibrium orientation is to be expected, at which the ordering and the randomizing exactly cancel and the average orientation of the population ceases to change however long the suspension keeps flowing.

5.4 The Scalar Order Parameter

To connect this tensor description with the experiment, the orientation state must now be reduced to a scalar quantity that can be obtained directly from the measured fiber orientations. This section establishes that connection by showing how the order parameter η_1, introduced in Section 2.3.3 and recovered from the PTV measurements, can be expressed directly in terms of the orientation tensor $\boldsymbol{A}$. The experimental and tensor descriptions therefore characterize the same underlying orientation state at different levels: one is constructed from the measured orientations of individual fibers, whereas the other is their continuum representation. This correspondence provides the link needed to compare the law derived below directly with the experimental data of Chapter 4..

5.4.1 Reduction of the Orientation Tensor

Because $\boldsymbol{A}$ is symmetric, it is diagonalizable: its eigenvectors define the principal directions of orientation, while its eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$ quantify the degree of orientation associated with each of those directions. By the unit-trace constraint,

$$\sum_{i=1}^{d} \lambda_i = 1. \tag{5.8}$$

The leading eigenvalue λ_1 is therefore a natural scalar measure of the strength of the preferential orientation. It reaches unity when the fibers are perfectly aligned along a single direction and takes the value $1/d$ for an isotropic population. Hence,

$$\frac{1}{d} \leq \lambda_1 \leq 1. \tag{5.9}$$

Although λ_1 already measures the degree of alignment, it is less convenient as a normalized descriptor because it does not span the interval $[0, 1]$ and its isotropic reference value depends on the dimensionality of the orientation space.

This motivates the use of the nematic tensor, commonly used in the theory of liquid crystals (de Gennes et al., 1993; Doi et al., 1988), which removes the isotropic part of the orientation tensor and normalizes the remaining anisotropy,

$$\boldsymbol{Q} = \frac{d\,\boldsymbol{A} - \boldsymbol{I}}{d-1}, \tag{5.10}$$

this is an affine rescaling of $\boldsymbol{A}$ whose trace vanishes identically, $\operatorname{tr}\boldsymbol{Q} = (d{\cdot}1{-}d)/(d{-}1) = 0$, so that $\boldsymbol{Q}$ measures departure from isotropy and nothing else. Because $\boldsymbol{Q}$ is an affine function of $\boldsymbol{A}$, the two tensors share the same eigenvectors, while their eigenvalues are related by the same rescaling. The scalar order parameter is defined as the leading eigenvalue of $\boldsymbol{Q}$,

$$\eta_1 = \frac{d\,\lambda_1 - 1}{d-1}, \tag{5.11}$$

which vanishes at isotropy, where $\lambda_1 = 1/d$, and equals unity at perfect alignment, where $\lambda_1 = 1$, in any dimension. In three dimensions this gives $\eta_1 = (3\lambda_1 - 1)/2$ and in the plane $\eta_1 = 2\lambda_1 - 1$, the difference in the coefficients being a consequence of the isotropic state carrying a leading eigenvalue of $1/3$ in the bulk and $1/2$ in the plane, and not a difference in the quantity being measured.

5.4.2 Equivalence with the Measured Order Parameter

The PTV analysis used in this work recovers the fiber orientations projected onto the imaging plane. The orientation tensor constructed from these measured directions is therefore two-dimensional, so that $d = 2$. Experimentally, the order parameter is not obtained from the eigenvalues of this tensor, but from the modulus of the circular average of the measured fiber angles,

$$\eta_1 = \left|\left\langle e^{\,i2\theta}\right\rangle\right|, \tag{5.12}$$

where the factor of two accounts for the axial nature of fiber orientation: the angles θ and $\theta + \pi$ describe the same physical orientation. The two definitions are shown below to be equivalent.

Let θ denote the angle between a fiber and the first coordinate axis, so that

$$\boldsymbol{p} = \begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix}. \tag{5.13}$$

The corresponding planar orientation tensor is

$$\boldsymbol{A} = \begin{pmatrix} \langle\cos^2\theta\rangle & \langle\cos\theta\sin\theta\rangle \\ \langle\cos\theta\sin\theta\rangle & \langle\sin^2\theta\rangle \end{pmatrix}, \tag{5.14}$$

whose trace is unity, as required. Using the double-angle identities $\cos^2\theta = \frac{1}{2}(1+\cos 2\theta)$, $\sin^2\theta = \frac{1}{2}(1-\cos 2\theta)$, and $\cos\theta\sin\theta = \frac{1}{2}\sin 2\theta$, the tensor can be separated into its isotropic part and a traceless anisotropic contribution:

$$\boldsymbol{A} = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \frac{1}{2}\begin{pmatrix} \langle\cos 2\theta\rangle & \langle\sin 2\theta\rangle \\ \langle\sin 2\theta\rangle & -\langle\cos 2\theta\rangle \end{pmatrix}. \tag{5.15}$$

The doubled angle is not introduced here merely as a mathematical convenience, but arises naturally from the tensor algebra, and it is the same doubling that appears in Equation (5.12). Its origin is the axial symmetry described below Equation (5.4): a fiber oriented at θ and one oriented at $\theta + 180°$ represent the same physical orientation. Under the mapping $\theta \mapsto 2\theta$, these two orientations are mapped onto the same

point of the unit circle, since they differ by 360° after doubling. Any descriptor of the orientation population must therefore be invariant under this reversal. The tensor representation achieves this automatically because it is quadratic in $\boldsymbol{p}$, whereas the circular representation enforces the same invariance explicitly through the doubled angle.

Substituting Equation (5.15) into the definition of the nematic tensor, Equation (5.10) with $d = 2$, the isotropic contribution cancels identically and only the anisotropic part remains:

$$\boldsymbol{Q} = 2\boldsymbol{A} - \boldsymbol{I} = \begin{pmatrix} \langle \cos 2\theta \rangle & \langle \sin 2\theta \rangle \\ \langle \sin 2\theta \rangle & -\langle \cos 2\theta \rangle \end{pmatrix}, \tag{5.16}$$

which is a symmetric, traceless 2×2 tensor. Its two eigenvalues are equal in magnitude and opposite in sign. From the characteristic equation,

$$\lambda^2 - \langle \cos 2\theta \rangle^2 - \langle \sin 2\theta \rangle^2 = 0, \tag{5.17}$$

they are

$$\lambda_{1,2}^{(\boldsymbol{Q})} = \pm\sqrt{\langle \cos 2\theta \rangle^2 + \langle \sin 2\theta \rangle^2}. \tag{5.18}$$

The leading eigenvalue of $\boldsymbol{Q}$, which is the scalar order parameter of Equation (5.11), is therefore the modulus of the complex number $\langle \cos 2\theta \rangle + i \langle \sin 2\theta \rangle = \langle e^{i2\theta} \rangle$,

$$\eta_1 = 2\lambda_1 - 1 = \sqrt{\langle \cos 2\theta \rangle^2 + \langle \sin 2\theta \rangle^2} = \left| \langle e^{i2\theta} \rangle \right|, \tag{5.19}$$

which is precisely the definition given in Equation (5.12). The equivalence is therefore exact. The circular average of the doubled fiber angles, computed directly from the PTV measurements, yields the same scalar order parameter as that obtained from the planar orientation tensor $\boldsymbol{A}$ evolved by the Advani-Tucker formulation.

Two consequences follow, and they are what make the remainder of this chapter possible. The first is that η_1 inherits the properties of a tensor invariant: it does not depend on the coordinate frame in which the angles happen to be measured, so the alignment reported in a cell of the mold is a property of the fiber population and not of the axes chosen to describe it. The second is that the argument of the complex average carries the direction of that alignment, since $\frac{1}{2} \arg \langle e^{i2\theta} \rangle$ is the eigenvector of $\boldsymbol{Q}$ and therefore the principal direction of orientation, the director against which the angular deviation

δ of Section 2.3.3 is defined. The modulus measures how ordered the population is and the argument measures the direction of its order, and both are recovered from the same average.

The second is that the argument of the same complex average contains the direction of that alignment. Specifically,

$$\theta_{\mathrm{dir}} = \frac{1}{2} \arg \left\langle e^{i2\theta} \right\rangle , \tag{5.20}$$

gives the angle of the principal director, defined modulo π, and therefore the direction of the eigenvector associated with the largest eigenvalue of $\boldsymbol{Q}$.

The law derived below is accordingly written in η_1 throughout, without further distinction between the measured quantity and the theoretical one.

5.5 Equilibrium State

Projected onto the scalar order parameter, the competition contained in Equation (5.7) reduces to two contributions of opposite sign. The flow supplies order at a rate proportional to the disorder that remains, since a population already aligned offers nothing left to align, while a population in complete disorder offers the largest room for improvement,

$$\left(\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} \right)_{\mathrm{order}} = a\,(1 - \eta_1), \tag{5.21}$$

and the interactions supply disorder at a rate proportional to the order that is present, since a population in complete disorder offers nothing left to randomize, while an aligned population is the one most exposed to being disturbed,

$$\left(\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} \right)_{\mathrm{disorder}} = -\,b\,\eta_1, \tag{5.22}$$

where a is the effective ordering rate carried by the deformation term and b the effective randomizing rate carried by the interaction term, so that b grows with C_I and hence with the fiber concentration. Both rates are imprinted on the fibers by the deformation that the flow delivers to them, and both are consequently measured against the same clock. The rotation term does not appear in the balance: rotating the orientation pattern as a whole turns the eigenvectors of $\boldsymbol{Q}$ without altering its eigenvalues, so it changes the direction of the alignment and leaves its degree, which is η_1, untouched. The projection

of each of the three terms, together with the closure on which the deformation term depends, is carried out in Appendix A.

That clock is not time but the accumulated shear strain,

$$\Gamma = \int \dot{\gamma}\, \mathrm{d}t, \tag{5.23}$$

which is the natural independent variable of the problem. The reason lies in the structure of Equation (5.7): the shear rate multiplies the whole of its right-hand side, the deterministic terms and the diffusive one alike, so it sets the pace at which the orientation evolves but not the state toward which it evolves. Pouring the same mix slowly or quickly through the same geometry does not change the orientation it acquires; only the amount of deformation accumulated along the way does. Two flows that accumulate the same Γ reach the same orientation, whatever the speed at which the strain was applied, and Γ is the variable against which the alignment of the fibers is properly measured.

Adding the two contributions gives the scalar balance of the orientation,

$$\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} = a\,(1 - \eta_1) - b\,\eta_1, \tag{5.24}$$

in which the first term pushes η_1 upward and the second pulls it downward, each weighted by how much order is already present.

The equilibrium state is the orientation at which the two exactly cancel. Setting $\mathrm{d}\eta_1/\mathrm{d}\Gamma = 0$ in Equation (5.24) gives $a\,(1 - \eta_{1,\mathrm{eq}}) = b\,\eta_{1,\mathrm{eq}}$, and solving,

$$\eta_{1,\mathrm{eq}} = \frac{a}{a + b}, \tag{5.25}$$

the equilibrium alignment is the share of the ordering rate within the total of the two rates. It is neither a state of complete alignment nor one of complete disorder, but a partially aligned state, reached not because either mechanism has stopped acting but because the two have come into balance. When the flow dominates, at low fiber content, $a \gg b$ and $\eta_{1,\mathrm{eq}}$ approaches unity; when the interactions dominate, at high fiber content, b grows and $\eta_{1,\mathrm{eq}}$ falls.

Grouping the terms in η_1 in Equation (5.24) and factoring the sum of the two rates,

$$\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} = (a+b)\left[\frac{a}{a+b} - \eta_1\right], \tag{5.26}$$

exposes the structure of the balance: the orientation changes at a rate proportional to the distance that still separates it from its equilibrium value, and it changes more slowly the closer it comes to it. This is a first-order relaxation, and its solution is the exponential approach to equilibrium developed in the following section.

5.6 Exponential Function of Approach to Equilibrium

Integrating Equation (5.26) for a population that enters the mold in a disordered state, $\eta_1(\Gamma = 0) = 0$, yields the closed form

$$\eta_1(\Gamma) = \eta_{1,\max}\left(1 - e^{-\Gamma/\Gamma^*}\right), \tag{5.27}$$

which is the orientation saturation law proposed in this work. The integration, the general solution for a nonzero entry order, and the conditions under which it collapses to Equation (5.27) are given in Appendix A.

The two parameters of the law are read from the solution rather than postulated. The first,

$$\eta_{1,\max} = \frac{a}{a+b}, \tag{5.28}$$

is the equilibrium alignment obtained in the previous section, and it is the ceiling of the curve: the highest degree of alignment the population can reach, no matter how long the mix continues to flow. The second,

$$\Gamma^* = \frac{1}{a+b}, \tag{5.29}$$

is the characteristic strain, the scale against which the approach to that ceiling is measured. At $\Gamma = \Gamma^*$ the population has covered a fraction $1 - e^{-1} \approx 0.63$ of the distance to $\eta_{1,\max}$, and by $\Gamma \approx 3\,\Gamma^*$ it is saturated for practical purposes. The law

compresses the orientation problem into two numbers: how far the fibers can align, and how much deformation they need in order to get there.

The form of Equation (5.27) is not new in physics. A quantity relaxing toward an equilibrium fixed by two competing rates, at a pace proportional to the distance that remains, is the structure that governs the charging of a capacitor through a resistor, the cooling of a body toward the temperature of its surroundings, and the approach of a reversible chemical reaction to its equilibrium concentration. What is new is the object to which it is applied. The reduction of the Advani-Tucker equation to a two-parameter saturation law has not been used to describe the orientation of steel fibers in a non-Newtonian viscoplastic flow, where the Jeffery hypotheses fail, the orbit is lost, and the orientation of the population cannot be predicted from the single-fiber solution.

The fiber concentration does not appear explicitly in Equation (5.27). It enters indirectly, through the interaction coefficient C_I, which sets the randomizing rate b and is therefore contained within both $\eta_{1,\max}$ and Γ^*. Two predictions follow. The first concerns the ceiling: as the dosage increases, b grows, the share $a/(a+b)$ falls, and $\eta_{1,\max}$ decreases, so that the entropy of the system carries a larger weight and the maximum alignment attainable by any casting flow is lowered. The second concerns the strain required to reach it: as the dosage increases, Γ^* grows, and the population must accumulate more deformation before it saturates, because the flow has to work for longer against a stronger randomizing agency before the balance is struck.

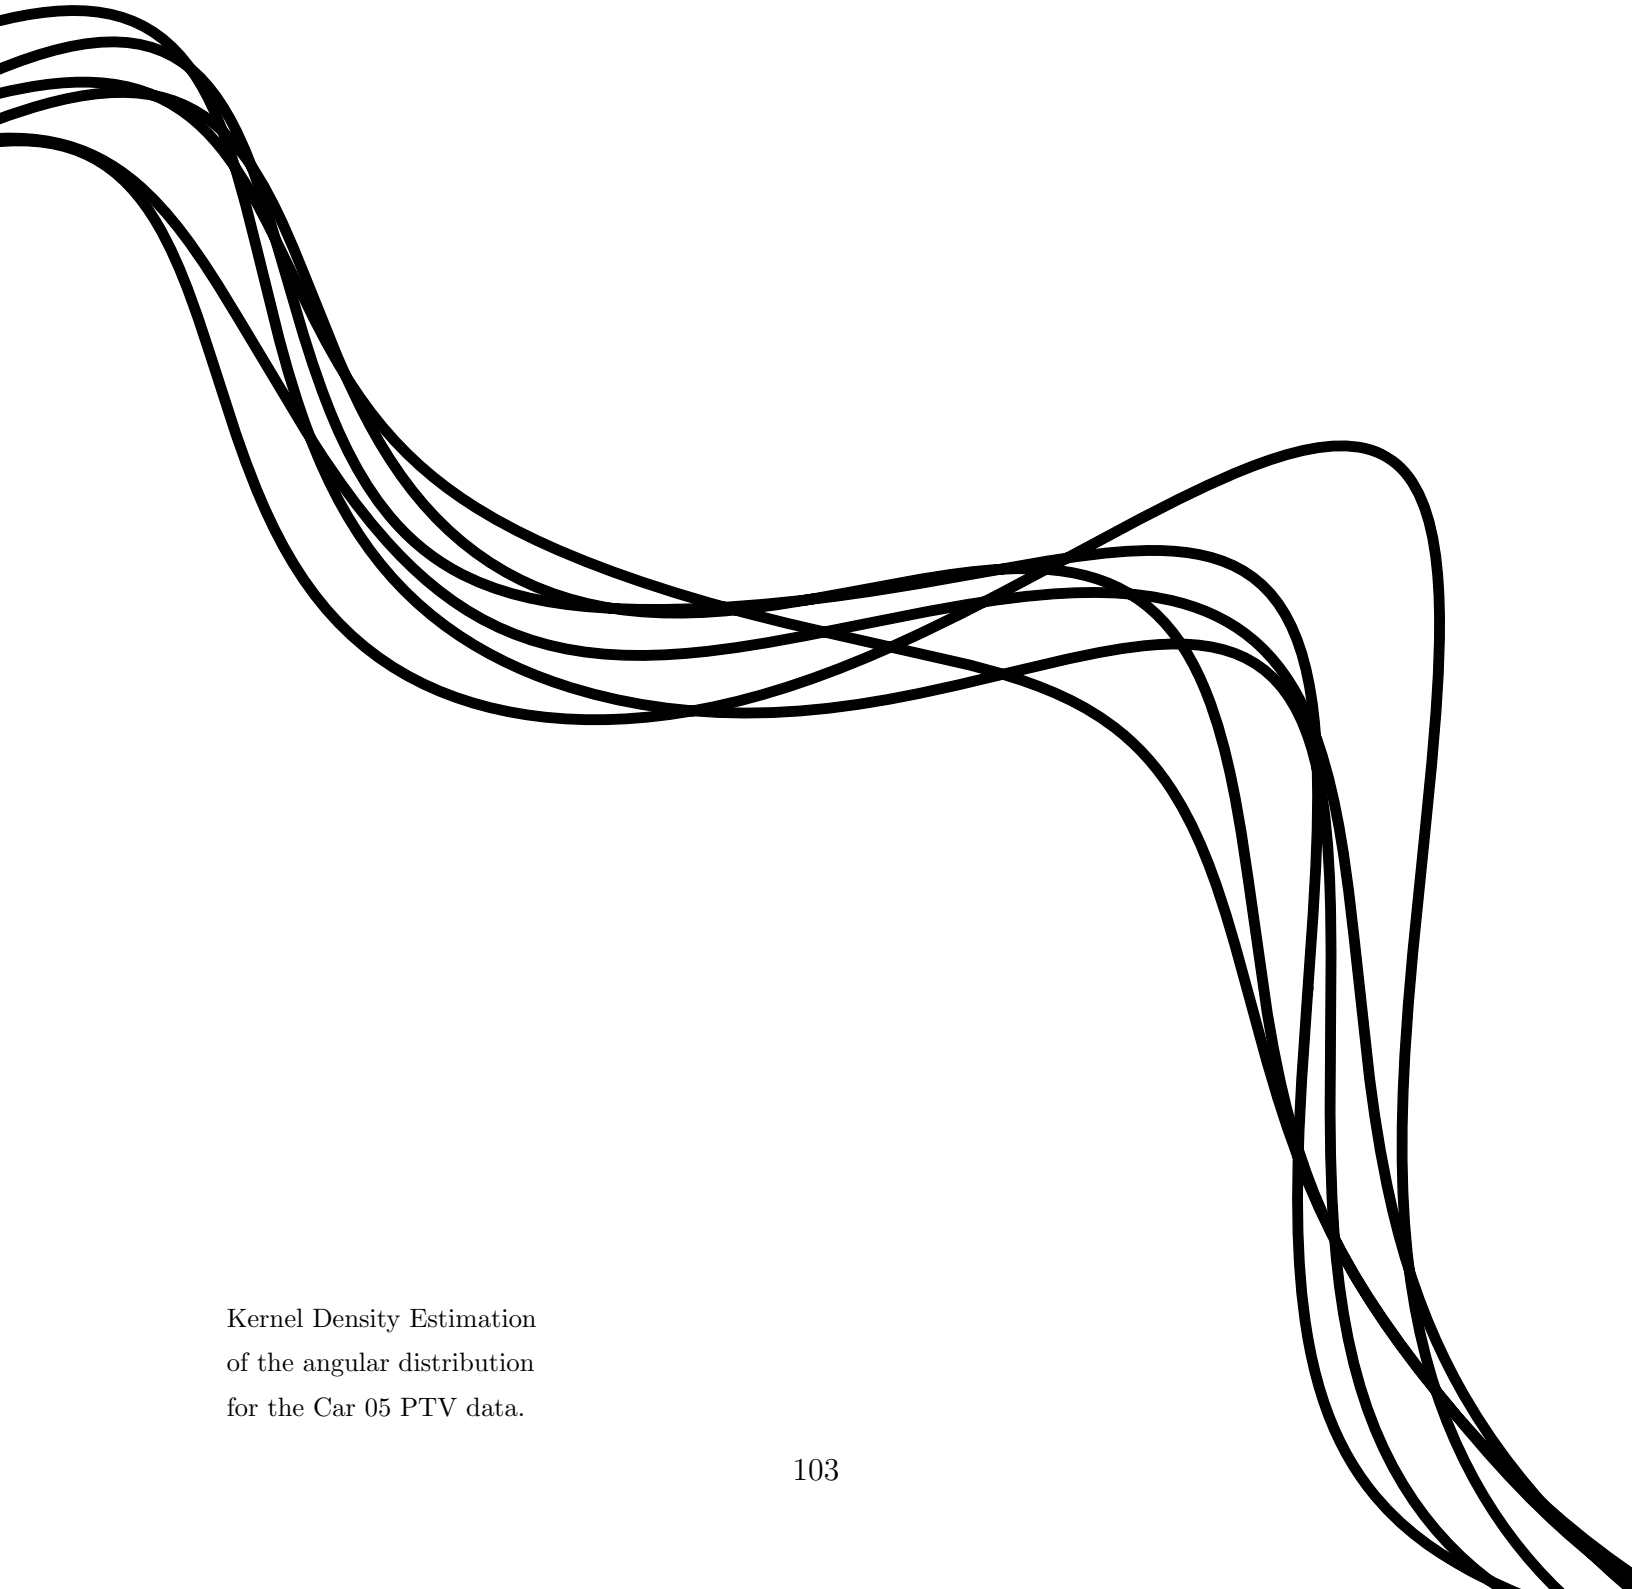
These two predictions are not independent within the two-constant form, and the tension between them is instructive. Since $\eta_{1,\max}$ and Γ^* share the denominator $(a+b)$, interactions that acted only by raising b would lower $\eta_{1,\max}$ and lower Γ^* together, which is to say that a richer mix would align less but saturate sooner. An increase of Γ^* with the fiber content therefore carries a physical meaning of its own: it indicates that the interactions do not merely randomize the population but also degrade the effective ordering rate a supplied by the flow, retarding the alignment in addition to opposing it. This slow-kinetics behavior is documented for concentrated fiber suspensions, which align more slowly with respect to strain than the single-coefficient model predicts (Wang et al., 2008). For this reason $\eta_{1,\max}$ and Γ^* are treated in what follows as phenomenological parameters, fitted directly to the measured orientation, rather than derived from a single value of C_I; their identifiability and their relation to a and b are

examined in Appendix A.

Reduced to this form, the orientation problem acquires a compact structure. The entire history of the casting flow enters the law through a single variable, the accumulated shear strain Γ, and the entire effect of the fiber dosage enters through two constants, $\eta_{1,\max}$ and Γ^*. Chapter 6 confronts Equation (5.27) with the orientation measured in the beam, and determines how its two parameters respond to the fiber content.

Chapter 6

Prediction of the Final Orientation



Kernel Density Estimation
of the angular distribution
for the Car 05 PTV data.

Chapter 4 showed that the fiber content acts on the flow and on the fibers it carries, slowing the carrier flow where the geometry forces a change in its course and introducing an additional reorientation that cannot be attributed to the flow alone. Chapter 5 then reduced the Advani-Tucker description of this reorientation to a single scalar quantity and integrated its evolution over the accumulated strain, yielding a law of exponential approach to equilibrium governed by two parameters. What remains is to confront that law with the orientations measured in the beam and then use it predictively: to determine whether the deformation history of the flow alone, together with the interaction parameters of the suspension, can reproduce the degree of alignment at which the fibers ultimately come to rest. This prediction is the central test of the present chapter.

6.1 Correspondence Between the Models and the Experiment

The Folgar-Tucker and Advani-Tucker models describe the reorientation of the fibers as a competition between two actions, both made quantitative in Chapter 5. The flow orders the population, tilting each fiber toward the direction along which the material is stretched. The fibers, once close enough to disturb one another, randomize that alignment through a rotary diffusion whose strength is set by an interaction coefficient, and the models tie that coefficient to the fiber content: the denser the population, the more frequent the encounters between fibers, and the stronger the randomizing term relative to the ordering one.

Two predictions of this mechanism are confirmed by the measurements of Chapter 4, one during filling and one at rest. A higher fiber content raises the interaction coefficient and with it the rate of change of the orientation tensor $\mathrm{D}\boldsymbol{A}/\mathrm{D}t$, so the fibers reorient faster while the flow that carries them does not, which is the separation the global analysis found between the angular velocity of the fibers and the vorticity of the carrier. The same stronger randomizing term balances the flow at a less aligned equilibrium, so the final state grows more disordered with the dosage, which is the rise in angular deviation the arrested fibers showed at fixed rheology. The models therefore apply to the unsteady, viscoplastic, shear-thinning flow studied here, outside the Newtonian

dilute setting for which they were derived, and reducing them to a compact law for this flow is warranted rather than assumed.

The consequence the experiment can test most directly, and the one on which the hypothesis of this work rests, is that the alignment the population reaches at equilibrium falls as the fiber content rises, since a stronger randomizing term balances the ordering of the flow at a less aligned state. The strain required to reach that equilibrium is governed by the second constant of the law, whose response to the fiber content is less simple to anticipate and is left to the measurements to settle. The reduction of Chapter 5 collects both into a two-parameter law, so that the trend of either constant with the fiber content is a direct reading of the models against the experiment.

In an unsteady viscoplastic shear-thinning flow, then, the saturation law is the form in which the models meet the measurements: it predicts the final angular state of a group of fibers from the strain that group has accumulated, through two constants that the fiber content is expected to displace. The remainder of this chapter fits those constants to the measured orientations and tests the prediction they encode.

6.2 Confrontation with the Measured Orientation

Figure 6.1 relates the orientation order measured in each beam cell to the shear strain accumulated by the fibers in that cell at the time of arrest, for the three fiber dosages in Carbopol 0.2%. Carbopol 0.5% arrests over too narrow a range of accumulated strain to resolve the saturation curve and is therefore not fitted here. Each large marker represents the arrested state of one beam cell, positioned by the final orientation order attained by the fiber population and by the accumulated deformation at arrest. The curves are the corresponding fits of the saturation law derived in Chapter 5, Equation (5.27), with one fit for each fiber dosage. The comparison must be interpreted subject to two limitations that constrain what can be inferred from only five cells.

The first limitation is the number of points. Each curve is fitted to only five arrested cells, which is insufficient to estimate its two constants with high precision; the fits are therefore shown together with their confidence bands. The trends are nevertheless consistent across the measured range, although the exact parameter values remain uncertain. A validated CFD-DEM simulation could extend this analysis over deformation

histories far beyond those accessible in a physical casting experiment, providing a substantially wider range over which the proposed saturation law can be tested. The values reported here should therefore be regarded as initial estimates of the two constants for a material of this kind rather than as definitive material parameters.

The second limitation concerns the cell farthest downstream, r2c3. This is the cell reached last by the advancing flow and the one that fills least completely. At arrest, the local flow is still developing against the far wall of the beam, and camera 3 does not capture the cell in full. The orientation recovered there is therefore inferred from an incomplete view of a flow that has itself not fully developed. Its arrested state is consequently the least reliable of the five and is treated with caution throughout: it is shown in every figure but is never used on its own to support a conclusion. Its behavior is examined further among the final-state measurements in Appendix F.

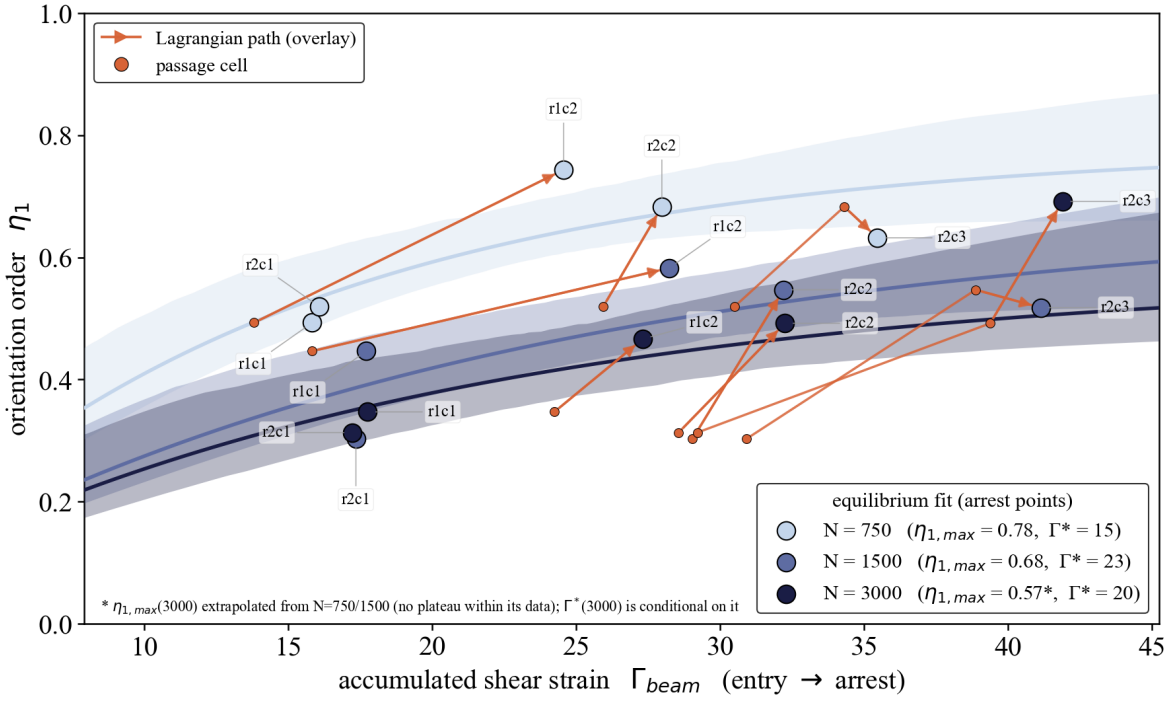


Figure 6.1: Orientation order η_1 against the shear strain Γ_{beam} accumulated between the entry of the material into the beam and its arrest, for the three fiber dosages in Carbopol 0.2%. Large markers are the arrest points of each beam cell, curves are the fits of Equation (5.27) to those points, and shaded bands are the confidence intervals of the fits. Small markers and arrows show the reconstructed Lagrangian passage of the fiber population through the upstream cells of each destination.

Read across the three dosages, the fits displace the equilibrium alignment in the sense

the models anticipate. The ceiling $\eta_{1,\max}$ falls as the fiber content rises, from 0.78 at the lowest dosage to 0.68 and then to 0.57, which is the randomizing action of the fibers made visible: a denser population balances the ordering of the flow at a less aligned equilibrium. This is the central prediction of the hypothesis, and the measurements bear it out.

The strain scale Γ^* is less straightforward to interpret. Between the lowest and intermediate fiber concentrations, for which both fitting parameters are constrained by the data, Γ^* increases from 15 to 23. The higher fiber concentration therefore not only leads to a lower asymptotic degree of alignment, but also makes the evolution toward that state more gradual, requiring a larger accumulated deformation. At the highest fiber concentration, however, Γ^* cannot be determined independently. The cells with sufficiently reliable orientation measurements do not extend far enough into the high-strain regime for the orientation state to approach its asymptote. Consequently, $\eta_{1,\max}$ cannot be independently constrained by the data and is prescribed based on the trend established by the two lower fiber concentrations. The strain scale Γ^* is then fitted conditionally, with $\eta_{1,\max}$ held fixed. The only cell extending substantially into the high-strain range is r2c3, but its orientation measurements are the least reliable of the five cells and therefore cannot, on their own, constrain the asymptotic behavior. The reported value of Γ^* at the highest fiber concentration should therefore be interpreted as conditional on the assumed value of $\eta_{1,\max}$.

So far, the model describes the degree of alignment attained by the fiber population through the scalar order parameter η_1. For the structural motivation of this work, however, an angular description is also useful. Fiber orientation is quantified here by the angle measured relative to the horizontal direction, providing a consistent geometric reference for the analysis. Its structural significance, however, lies in the orientation of the fibers relative to the crack plane, or equivalently to its normal direction, since this determines how effectively the fibers intersect and bridge a developing crack. The next section therefore recasts the order parameter in terms of a characteristic angular deviation, providing an angular interpretation of the same orientation state.

6.2.1 The Angle Predictor

The model rests on two facts. First, the orientation order η_1 of a beam cell approaches its equilibrium exponentially in the accumulated strain Γ; since the impact at the beam entry resets the fibers to a near-isotropic state ($\eta_{1,0} \approx 0$),

$$\eta_1(\Gamma) \;=\; \eta_{1,\max}\left(1 - e^{-\Gamma/\Gamma^*}\right). \tag{6.1}$$

Second, the order parameter maps directly to an angle through the identity $\eta_1 = \cos(2\delta_{\text{eff}})$. Taking the longitudinal beam axis as the angular reference, θ_{eq} denotes the offset of the population director from this axis. In the tensile region under bending, this axis approximately coincides with the principal tensile stress direction and is normal to the expected crack plane.

$$\cos\!\left(2\,\delta_{\text{eff}}^{\text{beam}}\right) \;=\; \eta_1\,\cos 2\theta_{\text{eq}}. \tag{6.2}$$

Inserting Equation (6.1) into Equation (6.2), the two constants $\eta_{1,\max}$ and $\cos 2\theta_{\text{eq}}$ merge into a single one, the asymptotic angle $\delta_\infty \equiv \frac{1}{2}\arccos\!\left(\eta_{1,\max}\cos 2\theta_{\text{eq}}\right)$, and the predictor reduces to two parameters per dosage:

$$\delta_{\text{eff}}^{\text{beam}}(\Gamma) \;=\; \tfrac{1}{2}\arccos\!\Big[\cos(2\delta_\infty)\left(1 - e^{-\Gamma/\Gamma^*}\right)\Big] \tag{6.3}$$

Figure 6.2 shows the same predictor expressed in angular form and compared with the measured angular deviations, again with one curve for each fiber concentration The strain scale Γ^* is the same parameter obtained from the order-parameter fit, so the change of variable introduces no additional fitting freedom. Instead, it reveals the asymptotic angular deviation δ_∞, which provides an angular measure of the residual dispersion retained by the fiber population as the accumulated strain grows without bound. This angle increases with fiber content, from 19.3° to 24.0° and then to 27.7°, with the last value conditional in the same sense as the asymptotic order parameter from which it is derived. Because the principal direction of deformation is nearly horizontal in this region, a larger δ_∞ corresponds to a broader angular distribution about the preferred flow direction. The decrease in equilibrium alignment is therefore recovered, in equivalent form, as an increase in the asymptotic angular deviation.

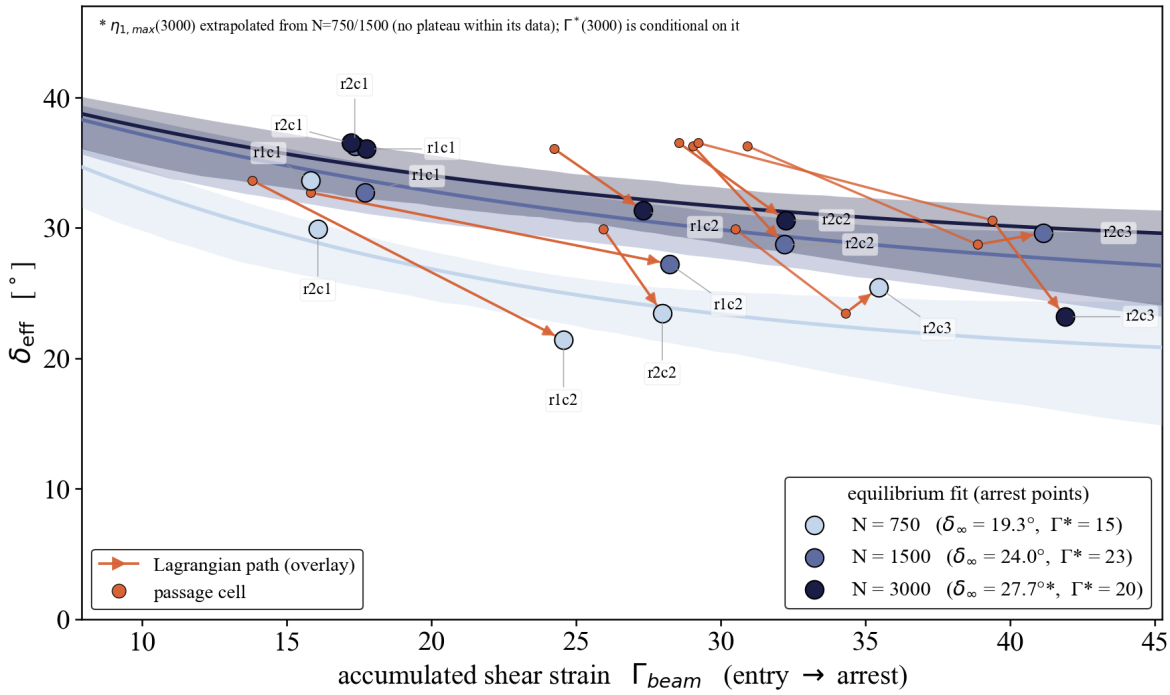


Figure 6.2: Effective misalignment δ_{eff} of the arrested fiber population against the accumulated shear strain Γ_{beam}, for the three fiber dosages in Carbopol 0.2%. The curves are the angle predictor of Equation (6.3), obtained by reparametrizing the fits of Figure 6.1: the asymptotic angle δ_∞ widens with the fiber content while the strain scale Γ^* is carried over unchanged. Markers and conventions as in Figure 6.1.

6.3 Test Against the Fiber-Free Flow

This section reverses the calibration procedure in order to test the predictive capability of the model independently. For each fiber concentration, the two constants of the reduced model, $\eta_{1,\max}$ and Γ^*, were calibrated using the strain accumulated by the fibers themselves. The present test evaluates whether the calibrated model correctly predicts the measured orientation angle when the accumulated strain is instead obtained from PIV measurements of the fiber-free flow.

For each spatial cell of the beam grid, the strain accumulated by the fiber-free flow over the full duration of the casting is computed from PIV, and the corresponding predicted angle is obtained from Equation (6.3). This predicted angle is compared with the angle measured for the fiber population in the same cell once the flow has arrested. Figure 6.3 presents the comparison on a cell-by-cell basis, and Table 6.1 summarizes the agreement for each dosage relative to a tolerance of ±5°.

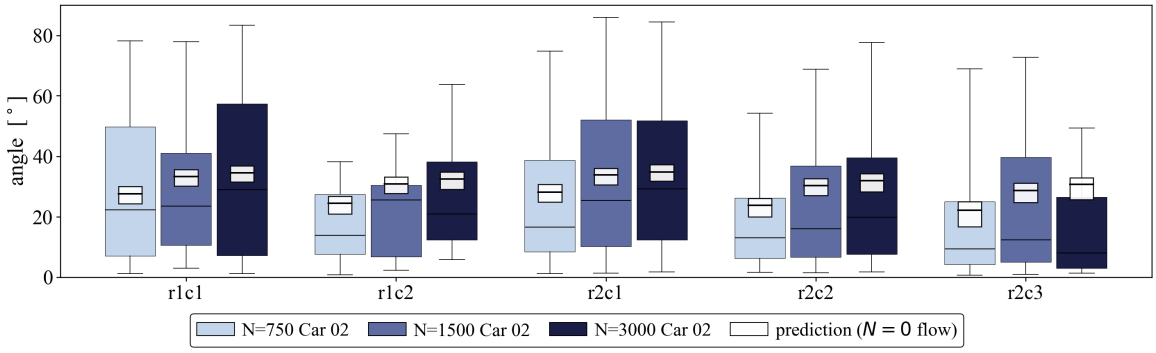


Figure 6.3: Final orientation angle of the fibers in each beam cell and at each fiber dosage. Coloured boxes give the interquartile range of the measured per-fiber angle; the white marker is the population angle predicted from the accumulated strain of the fiber-free flow through Equation (6.3). The prediction is a property of the population, not of the individual fiber, whose angle spans the full range within every cell. Carbopol 0.2%.

Table 6.1: Agreement between the predicted and the measured final angle, by fiber dosage, over all beam cells. The prediction uses the accumulated strain of the fiber-free flow. The bias is the mean signed departure, the RMSE its root-mean-square, and p tests the bias against zero; the practical tolerance is $\pm 5°$.

N [fibers]	regime	bias [°]	RMSE [°]	p
750	dilute	$+1.3$	3.5	0.46
1500	semi-dilute	-0.7	2.3	0.56
3000	semi-dilute	-1.7	3.7	0.37

Across the three fiber concentrations, the predictions remain within the prescribed tolerance. The mean difference between the predicted and measured angles remains below 2° for all fiber concentrations, while the root-mean-square difference does not exceed 3.7°. In addition, none of the three concentrations exhibits a systematic bias distinguishable from zero. The agreement spans the full range examined, from the dilute regime at the lowest fiber concentration to the lower limit of the semi-dilute regime at the highest concentration.

The accumulated strain of the fiber-free flow closely matches that experienced by the fibers in most cells, with the largest discrepancies occurring in some downstream cells at the higher fiber concentrations. Despite these local differences, the angular prediction remains within the prescribed 5° tolerance across all three concentrations. The RMSE ranges from 2.3° to 3.7°, with no systematic deterioration as the fiber concentration increases. The mean bias also remains small, varying from $+1.3°$ to $-1.7°$, and is not

statistically distinguishable from zero at any concentration ($p = 0.46$, 0.56, and 0.37). Taken together, these results indicate that the use of the fiber-free deformation field does not introduce a systematic error in the predicted orientation over the concentration range examined, and that it provides a sufficiently accurate approximation of the deformation experienced by the fiber-laden suspension.

Cell r2c3 is retained in the analysis and therefore contributes to the aggregate statistics reported in Table 6.1. Figure 6.3 shows its prediction alongside those of the other cells and also makes clear that r2c3 is among the cells for which the agreement is least favorable, particularly at the highest fiber concentration. This behavior is consistent with r2c3 being the least reliably observed of the five cells. Importantly, its inclusion does not alter the overall result: across the three fiber concentrations, the bias remains small and statistically indistinguishable from zero, while the RMSE does not exceed 3.7°, remaining within the prescribed ±5° tolerance. Excluding r2c3 improves the aggregate agreement but does not change this conclusion. The cell is therefore retained in the reported analysis rather than excluded on the basis of its poorer individual agreement.

Across the three fiber concentrations tested, the predicted angle agrees with the measured angle within the adopted tolerance. The mean difference between the predicted and measured angles remains below 2° for every fiber concentration, the root-mean-square difference does not exceed 3.7°, and none of the three concentrations exhibits a bias statistically distinguishable from zero. This agreement holds across the full range examined, from the dilute regime at the lowest fiber concentration to the lower part of the semi-dilute regime at the highest concentration. It therefore extends beyond the degree of alignment alone to the angular state predicted for the fiber population.

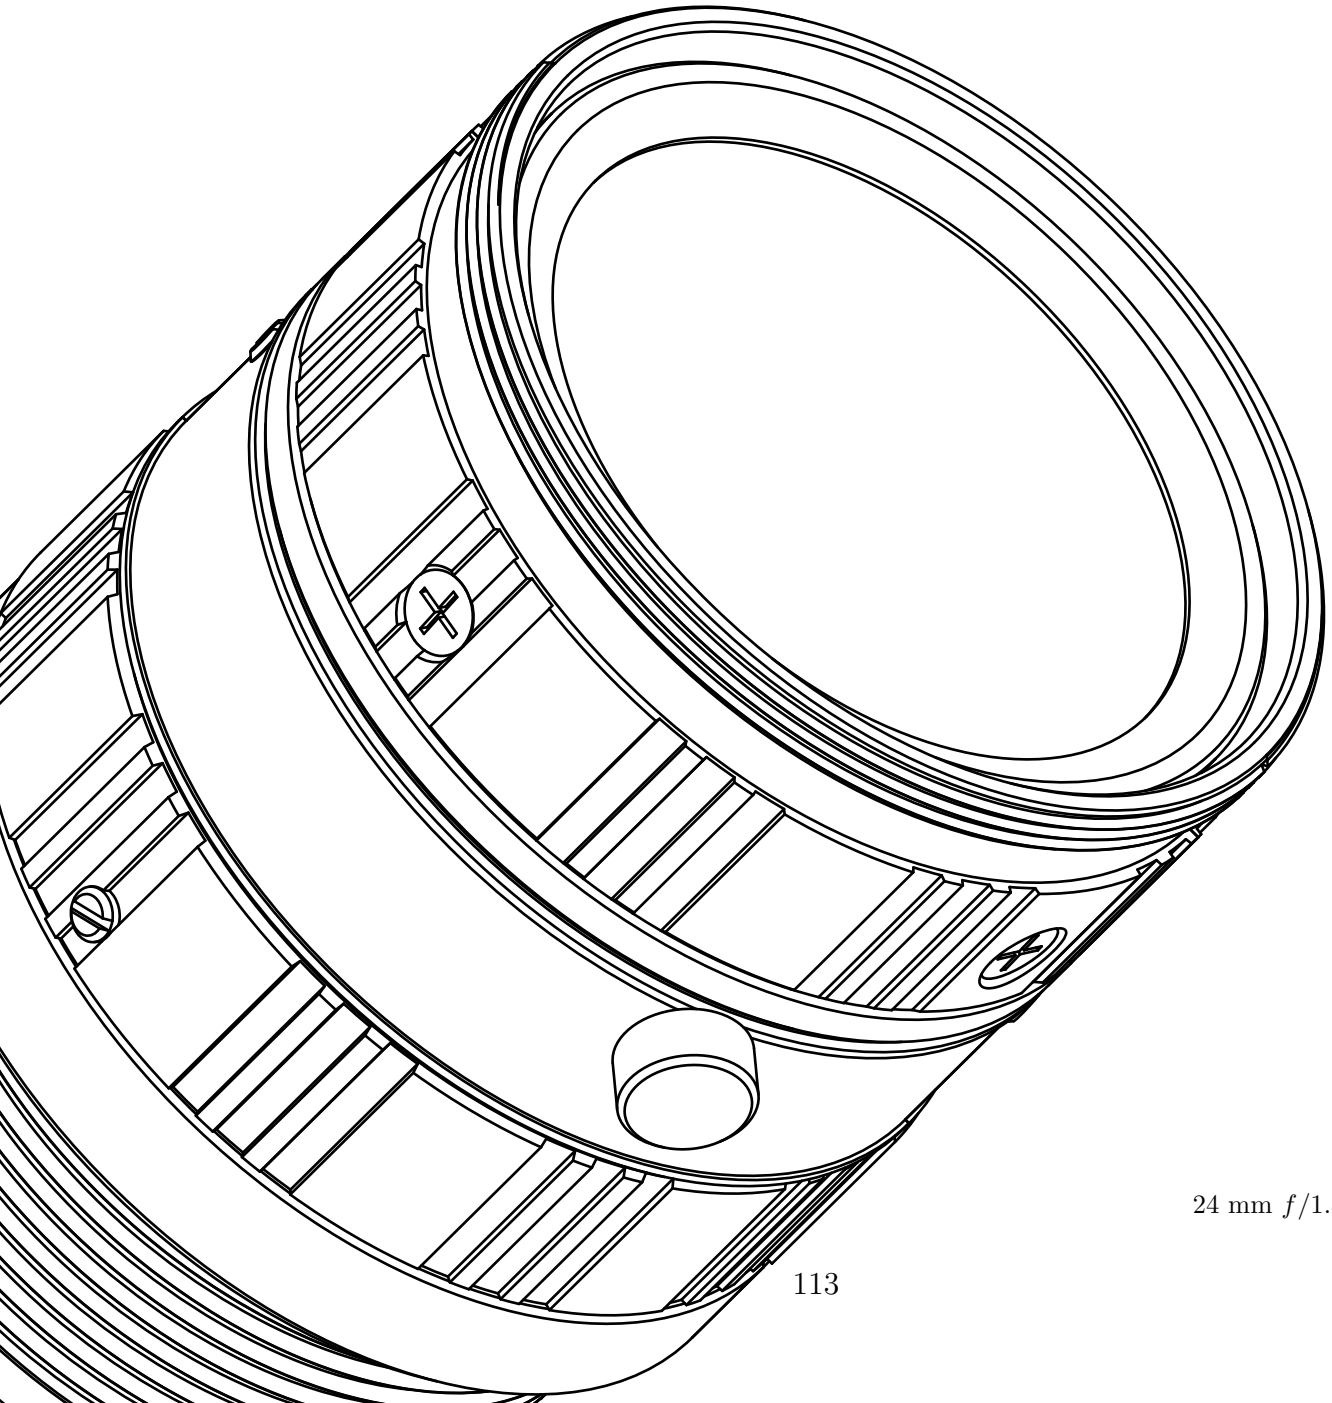
The hypothesis of this work is therefore supported over the range covered by the experiments. In an unsteady Herschel–Bulkley viscoplastic flow, the reduced two-parameter model predicts the final angular state of a fiber population from the accumulated strain of the carrier flow, from the dilute regime into the lower part of the semi-dilute regime. Its two parameters also vary systematically with fiber concentration in the direction anticipated by the hypothesis: as the fiber population becomes more concentrated, the attainable degree of alignment decreases and the residual angular dispersion increases.

Two limitations of the present experimental validation indicate how this analysis should

next be extended. First, the tracking method is planar, so the angle it recovers is a projection of a reorientation that is inherently three-dimensional; the complete trajectory of a fiber cannot be reconstructed from a single camera view. Second, the fiber concentrations accessible to the optical measurements are limited by the need to resolve and track individual fibers. They remain one to two orders of magnitude below those encountered in practical UHPC mixtures, so the dependence of the model parameters on fiber concentration is established only over the dilute and lower semi-dilute regimes rather than across the full concentration range of engineering interest. The CFD–DEM simulations that these experiments were designed to validate provide a route beyond both limitations: they track every fiber in three dimensions and can access fiber concentrations for which individual optical tracking is no longer feasible. The two parameters measured experimentally here therefore provide quantitative benchmarks against which the simulations can now be evaluated.

Chapter 7

Conclusions



24 mm $f/1.8$ lens.

7.1 The Model and What It Predicts

When this work began, no compact reduced-order model had been established for predicting the orientation attained by a population of fibers under the unsteady Herschel-Bulkley viscoplastic flow conditions considered here. The challenge arises from the coupled nature of the problem. Yield-stress rheology produces strongly heterogeneous deformation, with highly sheared regions coexisting with regions that accumulate little deformation; the fibers can in turn modify the carrier flow that transports them; and, as the fiber concentration increases, fiber-fiber interactions progressively hinder alignment and increase the dispersion of the orientation distribution. The resulting orientation state therefore cannot be inferred from the dynamics of an isolated fiber alone, but emerges from the accumulated deformation experienced by the population together with the interactions among its members. This coupling has made the orientation of fibers in fresh cementitious suspensions considerably more difficult to predict than to characterize experimentally after the flow has occurred.

Building on the Folgar-Tucker and Advani-Tucker descriptions of fiber reorientation, and expressing the orientation state through a single scalar order parameter, this work proposed a reduced two-parameter model in which the fiber population approaches an asymptotic orientation state as accumulated deformation increases. The two parameters have direct physical interpretations: one defines the maximum degree of alignment attainable by the population, while the other sets the characteristic deformation scale over which that state is approached. The model was characterized from the dilute regime into the lower part of the semi-dilute regime, the concentration range over which the optical techniques employed here, Particle Image Velocimetry for the carrier flow and Particle Tracking Velocimetry for the fibers, can still provide sufficiently resolved measurements of the flow and individual fiber orientations.

The central conclusion is that, within this range, the reduced model predicts the final orientation state of the fiber population using as input only the accumulated deformation of the fiber-free flow. When evaluated with the strain field measured in the absence of fibers, the model reproduces both the characteristic angular deviation and the degree of alignment attained by the fiber population, with angular errors remaining within the adopted 5° tolerance across all three fiber concentrations. Its predictive scope is delib-

erately limited: it describes the final orientation state of the population, rather than the individual fiber trajectories, translational velocities, or spatial distribution, which lie outside the two-parameter formulation. The fact that the fiber-free deformation field provides an adequate input is itself an important result. It indicates that, over the concentration range examined, fiber-induced modifications of the deformation field remain sufficiently limited for the fiber-free flow to serve as an effective approximation of the deformation experienced by the fiber-laden suspension.

7.2 The Effect of the Fiber Content

In the Folgar-Tucker description, the effect of increasing fiber concentration is represented through the interaction coefficient, which accounts for the additional rotational motion generated by fiber-fiber interactions. As the fiber concentration increases, this interaction mechanism becomes progressively more important relative to the ordering action of the carrier-flow deformation. The fibers therefore exhibit greater rotational activity, but this does not imply a faster approach toward alignment. On the contrary, the additional rotation acts to broaden the orientation distribution, reducing the asymptotic degree of alignment and increasing the accumulated deformation required for the population to approach that state.

The measurements provide direct evidence of this distinction between the rotation imposed by the carrier flow and that associated with the presence of other fibers. While the fluid vorticity changes only weakly with fiber concentration, the angular velocity of the fibers increases appreciably. The additional rotational activity therefore cannot be explained by a corresponding increase in the local rotation of the carrier flow alone and is consistent with an increasing contribution from fiber-fiber interactions. In this sense, the experiments provide direct evidence, under the Herschel-Bulkley viscoplastic flow conditions examined here, of the interaction-induced rotary motion represented phenomenologically by the Folgar-Tucker model, observed from the motion of the fibers themselves rather than inferred solely from the orientation distribution they leave behind.

7.3 Limits of the Measurement

Three limitations bound these conclusions, and all three follow from the optical access rather than from the law. The velocity field recovered by PIV is that of the central plane of the domain, whereas the fibers tracked by PTV move in three dimensions and are recovered as their projection onto that plane, so the deformation delivered to a fiber is taken to lie in the plane in which it is measured, and the depth of each fiber, together with any motion across the plane, is not resolved. The number of fibers that can be followed is likewise bounded by what the illumination and the concentration allow to be detected. Against this background a decoupling is observed between the linear velocities of the fluid and of the fibers, which may be an artifact of comparing measurements taken at different depths, or may reflect a genuine lag of the fibers behind the carrier; the distinction cannot be settled optically, and it is separate in kind from the rotational decoupling reported above, which is a physical finding and not a limitation.

7.4 Future Work

These are the limitations that the coupled CFD-DEM simulations, for whose validation the present measurements were obtained, are positioned to address. Once validated against these data, the simulations can resolve the motion and reorientation of each fiber in three dimensions, extending the projected orientations measured experimentally to their full three-dimensional counterparts. They can also resolve the velocities of the carrier and fiber phases throughout the depth that is collapsed in the planar optical measurements, allowing the apparent kinematic decoupling between the two phases to be examined directly. Finally, the simulations can access fiber concentrations beyond those for which individual-fiber optical tracking remains feasible, allowing the evolution of the two model parameters to be followed from the dilute and lower semi-dilute regimes examined here toward concentrations representative of practical UHPC mixtures.

Whether the orientation behavior identified experimentally persists as those concentrations are approached, and whether the resulting orientation state favors the fiber arrangement required for effective crack bridging and structural performance, remain open questions. The validated CFD-DEM framework may provide a means of address-

ing them, extending the chain established here from the flow that places the material, through the deformation that governs fiber orientation, to the mechanical role of the fibers in the hardened element.

Nomenclature

Symbol	Description	Units
Roman symbols		
a	Effective ordering rate supplied by the deformation term	—
$\mathbf{A}$	Second-order fiber orientation tensor, $\langle \mathbf{p} \otimes \mathbf{p} \rangle$	—
$\mathbb{A}$	Fourth-order fiber orientation tensor, $\langle \mathbf{p} \otimes \mathbf{p} \otimes \mathbf{p} \otimes \mathbf{p} \rangle$	—
b	Effective randomizing rate supplied by the interaction term	—
C_I	Fiber–fiber interaction coefficient (Folgar–Tucker)	—
(c_x, c_y)	Centroid coordinates of a detected fiber	mm
d	Spatial dimension ($d = 2$ in the plane, $d = 3$ in the bulk)	—
d_f	Fiber diameter	mm
$\mathbf{D}$	Rate-of-strain tensor, symmetric part of $\nabla \mathbf{u}$	s^{-1}
$\mathbf{e}_1$	Eigenvector of $\mathbf{D}$ giving the principal stretching direction	—
h	Height of a rectangular region of interest	mm
$\mathbf{I}$	Identity tensor	—
K	Consistency index (Herschel–Bulkley)	$\mathrm{Pa\,s}^n$
L	Fiber length	mm
$\mathbf{M}$	Traceless remainder of the planar orientation tensor	—
n	Flow behavior index (Herschel–Bulkley)	—
n	Fiber number density, $4\phi/(\pi d_f^2 L)$	mm^{-3}
N	Nominal fiber count of a dosage; number of fibers in a region	—
$\mathbf{p}$	Unit vector directed along the fiber axis	—

Symbol	Description	Units
$\mathbf{Q}$	Nematic tensor, $(d\mathbf{A} - \mathbf{I})/(d-1)$	—
r_f	Fiber aspect ratio, L/d_f	—
r_Λ	Rank-biserial effect size on the dispersion index	—
R_g	Radius of gyration of the fiber centroids in a region	mm
R_g^{unif}	Radius of gyration of a uniform filling of the same region	mm
s_i	Inter-pair skip of a PIV block (frames within a correlated pair)	frames
s_f	Final skip of a PIV block (frames between one pair and the next)	frames
SS	Sum of squares (factor and error terms of the ANCOVA)	—
t	Time	s
T	Period of a Jeffery orbit	s
$\mathbf{u}$	Fluid velocity field, of components u and v	$\mathrm{mm\,s^{-1}}$
$\lvert u\rvert$	Fluid speed, $\sqrt{u^2+v^2}$	$\mathrm{mm\,s^{-1}}$
$\bar{u}$	Mean flow speed (reference for the PIV relative error)	$\mathrm{mm\,s^{-1}}$
$\lvert v\rvert$	Fiber speed	$\mathrm{mm\,s^{-1}}$
w	Width of a rectangular region of interest	mm
$\mathbf{W}$	Rotation tensor, antisymmetric part of $\nabla\mathbf{u}$ (Equation (2.11))	$\mathrm{s^{-1}}$
$\mathbf{x}_i$	Centroid of fiber i; $\bar{\mathbf{x}}$ the mean centroid of the region	mm
Greek symbols		
$\dot{\gamma}$	Scalar shear rate, $\sqrt{2\,\mathbf{D}:\mathbf{D}}$	$\mathrm{s^{-1}}$
$\dot{\gamma}_c$	Critical shear rate, $(\tau_0/K)^{1/n}$	$\mathrm{s^{-1}}$
Γ	Accumulated shear strain, $\int \dot{\gamma}\,dt$	—

Symbol	Description	Units
Γ^*	Characteristic strain of the saturation law, $1/(a+b)$	—
Γ_{beam}	Strain accumulated between entry into the beam and arrest	—
δ	Angular deviation of a fiber from the local director	deg
δ_{eff}	Effective misalignment of the population from the beam axis	deg
δ_∞	Asymptotic misalignment, $\frac{1}{2}\arccos(\eta_{1,\max}\cos 2\theta_{\text{eq}})$	deg
ϵ_{ijk}	Levi-Civita permutation symbol	—
$\boldsymbol{\zeta}$	Vorticity vector, $\nabla \times \mathbf{u}$	s^{-1}
ζ	Out-of-plane component of the vorticity, $\partial v/\partial x - \partial u/\partial y$	s^{-1}
η_1	Scalar (nematic) orientation order parameter, $\lvert\langle e^{i2\theta}\rangle\rvert$	—
$\hat{\eta}_1$	Debiased orientation order parameter	—
$\eta_{1,\text{eq}}$	Equilibrium alignment, $a/(a+b)$	—
$\eta_{1,\max}$	Ceiling of the saturation law (attainable alignment)	—
η^2	Partial eta squared, effect size of the ANCOVA factor	—
θ	In-plane angle of a fiber; θ_i that of fiber i	deg
θ_{dir}	Local director, direction about which the population is ordered	deg
$\bar{\theta}$	Mean director reported for a cell	deg
θ_{eq}	Offset of the population director from the beam axis	deg
λ_1, λ_2	Eigenvalues of the orientation (and nematic) tensor	—
Λ	Spatial dispersion index, R_g/R_g^{unif} (with finite-sample correction)	—
μ_p	Plastic viscosity (Bingham)	Pa s
ξ	Shape factor, $(r_f^2-1)/(r_f^2+1)$	—

Symbol	Description	Units
τ	Shear stress	Pa
τ_0	Yield stress	Pa
ϕ	Fiber volume fraction	—
ϕ_{sd}	Dilute-to-semi-dilute boundary, r_f^{-2}	—
ϕ_c	Semi-dilute-to-concentrated boundary, r_f^{-1}	—
ω	Angular velocity of the fiber axis	$\mathrm{deg\,s^{-1}}$
Operators and notation		
$\langle \cdot \rangle$	Average over the fibers detected within a region	—
$\otimes$	Outer (tensor) product	—
$\mathbf{A}:\mathbf{B}$	Double contraction of two tensors	—
$\lVert \cdot \rVert$	Euclidean norm	—
tr	Trace of a tensor	—
$\nabla \mathbf{u}$	Velocity gradient tensor	$\mathrm{s^{-1}}$
$\mathrm{D}/\mathrm{D}t$	Material derivative, $\partial/\partial t + \mathbf{u} \cdot \nabla$	$\mathrm{s^{-1}}$
Δt	Time step between correlated PIV frames, block/fps	s
Δx	Pair displacement of a fiber between analyzed PTV frames	mm
$\Delta x_{\max}$	Maximum linking distance of the PTV association	mm
$\Delta\hat{\eta}_1$	Difference in debiased order between the two matrices	—
α	Significance level of the statistical tests ($\alpha = 0.05$)	—
Domain labels		
Z1, Z2, Z3	Vertical arm, elbow and outlet zones of the L-shape	—
RiCj	Cell of row i, column j of the beam grid	—
Abbreviations		

Symbol	Description	Units
UHPC	Ultra-High-Performance Concrete	
PIV	Particle Image Velocimetry	
PTV	Particle Tracking Velocimetry	
CFD–DEM	Computational Fluid Dynamics – Discrete Element Method	
ANCOVA	Analysis of Covariance	
RMSE	Root-mean-square error	
IoU	Intersection over Union	
UOD	Universal Outlier Detection	
YOLOv11	You Only Look Once, version 11 (instance segmentation network)	
YOLO26	You Only Look Once, 2026 release (instance segmentation network)	
SAM2	Segment Anything Model 2	
α–β–γ filter	Predictive filter estimating position, velocity and acceleration	
LED	Light-emitting diode	
PLA	Polylactic acid	
NaOH	Sodium hydroxide	
fps	Frames per second	

Bibliography

Abrishambaf, A., Pimentel, M., & Nunes, S. (2017). Influence of fibre orientation on the tensile behaviour of ultra-high performance fibre reinforced cementitious composites. *Cement and Concrete Research*, *97*, 28–40.

Adrian, R. J. (1991). Particle-imaging techniques for experimental fluid mechanics. *Annual Review of Fluid Mechanics*, *23*, 261–304.

Advani, S. G. & Tucker, C. L. (1987). The use of tensors to describe and predict fiber orientation in short fiber composites. *Journal of Rheology*, *31*(8), 751–784.

Alberti, M. G., Enfedaque, A., & Gálvez, J. C. (2017). On the prediction of the orientation factor and fibre distribution of steel and macro-synthetic fibres for fibre-reinforced concrete. *Cement and Concrete Composites*, *77*, 29–48.

Angst, U., Elsener, B., Larsen, C. K., & Vennesland, Ø. (2009). Critical chloride content in reinforced concrete: A review. *Cement and Concrete Research*, *39*(12), 1122–1138.

Aris, R. (2012). *Vectors, tensors and the basic equations of fluid mechanics*. Courier Corporation.

Auernhammer, G. K., Fataei, S., Haustein, M. A., Patel, H. P., Schwarze, R., Secrieru, E., & Mechtcherine, V. (2020). Transparent model concrete with tunable rheology for investigating flow and particle-migration during transport in pipes. *Materials & Design*, *193*, 108673.

Baddeley, A., Rubak, E., Turner, R., et al. (2016). *Spatial point patterns: methodology and applications with R*, volume 1. CRC press Boca Raton.

Bahmani, H., Mostafaei, H., Santos, P., & Fallah Chamasemani, N. (2024). Enhancing the mechanical properties of ultra-high-performance concrete (UHPC) through silica sand replacement with steel slag. *Buildings*, *14*(11), 3520.

Balmforth, N. J., Craster, R. V., Perona, P., Rust, A. C., & Sassi, R. (2007). Viscoplastic dam breaks and the Bostwick consistometer. *Journal of Non-Newtonian Fluid Mechanics*, *142*(1–3), 63–78.

Balmforth, N. J., Frigaard, I. A., & Ovarlez, G. (2014). Yielding to stress: recent developments in viscoplastic fluid mechanics. *Annual Review of Fluid Mechanics*, *46*, 121–146.

Bay, R. S. & Tucker, C. L. (1992). Fiber orientation in simple injection moldings. Part I: Theory and numerical methods. *Polymer Composites*, *13*(4), 317–331.

Bingham, E. C. (1922). *Fluidity and Plasticity*. New York: McGraw-Hill.

Bonn, D., Denn, M. M., Berthier, L., Divoux, T., & Manneville, S. (2017). Yield stress materials in soft condensed matter. *Reviews of Modern Physics*, *89*(3), 035005.

Chang, C. M. & Hossain, A. (2021). Life-cycle cost analysis of ultra high-performance concrete (uhpc) in retrofitting techniques for abc projects. *Accelerated Bridge Construction University Transportation Center Florida International University: Miami, FL, USA*, 3–5.

Chung, D. H. & Kwon, T. H. (2002). Invariant-based optimal fitting closure approximation for the numerical prediction of flow-induced fiber orientation. *Journal of Rheology*, *46*(1), 169–194.

Cieslinski, M. J., Wapperom, P., & Baird, D. G. (2016). Fiber orientation evolution in simple shear flow from a repeatable initial fiber orientation. *Journal of Non-Newtonian Fluid Mechanics*, *237*, 65–75.

Cintra, J. S. & Tucker, C. L. (1995). Orthotropic closure approximations for flow-induced fiber orientation. *Journal of Rheology*, *39*(6), 1095–1122.

Cohen, J. (2013). *Statistical power analysis for the behavioral sciences*. routledge.

Coleman, H. W., Steele, W. G., & Coleman, H. W. (2009). *Experimentation, validation, and uncertainty analysis for engineers*, volume 3. Wiley Online Library.

Coussot, P. (2014). Yield stress fluid flows: a review of experimental data. *Journal of Non-Newtonian Fluid Mechanics*, *211*, 31–49.

Coussot, P., Proust, S., & Ancey, C. (1996). Rheological interpretation of deposits of yield stress fluids. *Journal of Non-Newtonian Fluid Mechanics*, *66*(1), 55–70.

Crocker, J. C. & Grier, D. G. (1996). Methods of digital video microscopy for colloidal studies. *Journal of Colloid and Interface Science*, *179*(1), 298–310.

de Gennes, P.-G. & Prost, J. (1993). *The Physics of Liquid Crystals* (2 ed.). Oxford: Oxford University Press.

de Larrard, F., Ferraris, C. F., & Sedran, T. (1998). Fresh concrete: A Herschel-Bulkley material. *Materials and Structures*, *31*(7), 494–498.

Diggle, P. J. (2013). *Statistical Analysis of Spatial and Spatio-Temporal Point Patterns* (3 ed.). Boca Raton, FL: CRC Press.

Dinkgreve, M., Paredes, J., Denn, M. M., & Bonn, D. (2016). On different ways of measuring "the" yield stress. *Journal of Non-Newtonian Fluid Mechanics*, *238*, 233–241.

Doi, M. & Edwards, S. F. (1988). *The Theory of Polymer Dynamics*, volume 73 of *International Series of Monographs on Physics*. Oxford: Oxford University Press.

Eberle, A. P. R., Vélez-García, G. M., Baird, D. G., & Wapperom, P. (2010). Fiber orientation kinetics of a concentrated short glass fiber suspension in startup of simple shear flow. *Journal of Non-Newtonian Fluid Mechanics*, *165*(3–4), 110–119.

Egelmeers, T. R. N., Jaensson, N. O., Anderson, P. D., & Cardinaels, R. (2025). In situ experimental investigation of fiber orientation kinetics and rheology of polymer composites in shear flow. *Journal of Rheology*, *69*(2), 139–157.

Ferrara, L., Ozyurt, N., & di Prisco, M. (2011). High mechanical performance of fibre reinforced cementitious composites: the role of casting-flow induced fibre orientation. *Materials and Structures*, *44*(1), 109–128.

Ferrara, L., Park, Y.-D., & Shah, S. P. (2008). Correlation among fresh state behavior, fiber dispersion, and toughness properties of steel fiber reinforced concretes. *Journal of Materials in Civil Engineering*, *20*(7), 493–501.

Feys, D., Verhoeven, R., & De Schutter, G. (2009). Why is fresh self-compacting concrete shear thickening? *Cement and Concrete Research*, *39*(6), 510–523.

Folgar, F. & Tucker, C. L. (1984). Orientation behavior of fibers in concentrated suspensions. *Journal of Reinforced Plastics and Composites*, *3*(2), 98–119.

Gerland, F., Vaupel, T., Schomberg, T., & Wünsch, O. (2024). Analysing the influence of fibers on fresh concrete rheometry by the use of numerical simulation. *Construction Materials*, *4*(1), 128–153.

Graybeal, B. A. (2014). Design and construction of field-cast UHPC connections. Technical Report FHWA-HRT-14-084, Federal Highway Administration, McLean, VA.

Gudžulić, V., Dang, T. S., & Meschke, G. (2019). Computational modeling of fiber flow during casting of fresh concrete. *Computational Mechanics*, *63*, 1111–1129.

Guo, Z. (2014). *Principles of Reinforced Concrete*. Butterworth-Heinemann.

Habert, G., Miller, S. A., John, V. M., Provis, J. L., Favier, A., Horvath, A., & Scrivener, K. L. (2020). Environmental impacts and decarbonization strategies in the cement and concrete industries. *Nature Reviews Earth & Environment*, *1*(11), 559–573.

Herschel, W. H. & Bulkley, R. (1926). Konsistenzmessungen von gummi-benzollösungen. *Kolloid-Zeitschrift*, *39*(4), 291–300.

Hinch, E. J. & Leal, L. G. (1976). Constitutive equations in suspension mechanics. Part 2. Approximate forms for a suspension of rigid particles affected by Brownian rotations. *Journal of Fluid Mechanics*, *76*(1), 187–208.

Huang, H., Gao, X., Li, L., & Wang, H. (2018). Improvement effect of steel fiber orientation control on mechanical performance of UHPC. *Construction and Building Materials*, *188*, 709–721.

International Energy Agency (2018). Technology Roadmap: Low-Carbon Transition in the Cement Industry. Technology roadmap, International Energy Agency, Paris. Licence: CC BY 4.0.

Jaqaman, K., Loerke, D., Mettlen, M., Kuwata, H., Grinstein, S., Schmid, S. L., & Danuser, G. (2008). Robust single-particle tracking in live-cell time-lapse sequences. *Nature Methods*, *5*(8), 695–702.

Jeffery, G. B. (1922). The motion of ellipsoidal particles immersed in a viscous fluid. *Proceedings of the Royal Society of London. Series A*, *102*(715), 161–179.

Kang, S.-T. & Kim, J.-K. (2011). The relation between fiber orientation and tensile behavior in an ultra high performance fiber reinforced cementitious composite (UHPFRCC). *Cement and Concrete Research*, *41*(10), 1001–1014.

Kang, S.-T., Lee, Y., Park, Y.-D., & Kim, J.-K. (2010). Tensile fracture properties of an ultra high performance fiber reinforced concrete (UHPFRC) with steel fiber. *Composite Structures*, *92*(1), 61–71.

Kang, W., Mei, K., Liu, W., & Sun, S. (2025). Study on chloride diffusion performance and structural durability design of UHPC under chloride salt erosion. *Buildings*, *15*(19), 3569.

Keane, R. D. & Adrian, R. J. (1992). Theory of cross-correlation analysis of PIV images. *Applied Scientific Research*, *49*(3), 191–215.

Ken-Ichi, K. (1984). Distribution of directional data and fabric tensors. *International journal of engineering science*, *22*(2), 149–164.

Khanam, R. & Hussain, M. (2024). YOLOv11: An overview of the key architectural enhancements. *arXiv preprint arXiv:2410.17725*.

Konrád, P., Künzel, K., Kheml, P., Mára, M., Carrera, K., Beránek, L., Hlavunková, L., Fornusek, J., Konvalinka, P., & Sovják, R. (2025). Analysis of fiber content and orientation in prefabricated slab elements made of UHPFRC: Non-destructive, destructive, and CT scanning methods. *Materials*, *18*(21), 4843.

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, *2*(1–2), 83–97.

Li, J., Wu, Z., Shi, C., Yuan, Q., & Zhang, Z. (2020). Durability of ultra-high performance concrete: A review. *Construction and Building Materials*, *255*, 119296.

Liu, J., Jiang, Z., Zhao, Y., Zhou, H., Wang, X., Zhou, H., Xing, F., Li, S., Zhu, J., & Liu, W. (2020). Chloride distribution and steel corrosion in a concrete bridge after long-term exposure to natural marine environment. *Materials*, *13*(17), 3900.

Luo, X., Zhang, S., Li, A., & Zhang, D. (2023). Fiber content optimization of uhpc and r-uhpc oriented to tensile behavior and cost reduction. *Construction and Building Materials*, *395*, 132314.

Maas, H.-G., Gruen, A., & Papantoniou, D. (1993). Particle tracking velocimetry in three-dimensional flows: Part 1. Photogrammetric determination of particle coordinates. *Experiments in Fluids*, *15*(2), 133–146.

Macosko, C. W. (1994). *Rheology: Principles, Measurements, and Applications*. New York: Wiley-VCH.

Maggi, C. A. (2023). Sistema de velocimetría de imágenes de partículas. Trabajo de titulación para optar al título de ingeniero civil eléctrico, Universidad de los Andes, Santiago, Chile. Profesor Guía: Gustavo Luis Funes.

Mardia, K. V. & Jupp, P. E. (2000). *Directional Statistics*. Chichester: John Wiley & Sons.

Martinie, L., Rossi, P., & Roussel, N. (2010). Rheology of fiber reinforced cementitious materials: classification and prediction. *Cement and Concrete Research*, *40*(2), 226–234.

Martinie, L. & Roussel, N. (2011). Simple tools for fiber orientation prediction in industrial practice. *Cement and Concrete Research*, *41*(10), 993–1000.

Meeker, S. P., Bonnecaze, R. T., & Cloitre, M. (2004). Slip and flow in soft particle pastes. *Physical Review Letters*, *92*(19), 198302.

Munkres, J. (1957). Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, *5*(1), 32–38.

Ouellette, N. T., Xu, H., & Bodenschatz, E. (2006). A quantitative study of three-dimensional Lagrangian particle tracking algorithms. *Experiments in Fluids*, *40*(2), 301–313.

Petrie, C. J. S. (1999). The rheology of fibre suspensions. *Journal of Non-Newtonian Fluid Mechanics*, *87*(2–3), 369–402.

Phelps, J. H. & Tucker, C. L. (2009). An anisotropic rotary diffusion model for fiber orientation in short- and long-fiber thermoplastics. *Journal of Non-Newtonian Fluid Mechanics*, *156*(3), 165–176.

Piau, J.-M. (2007). Carbopol gels: Elastoviscoplastic and slippery glasses made of individual swollen sponges. *Journal of Non-Newtonian Fluid Mechanics*, *144*(1), 1–29.

Plaksyvyi, A., Skublewska-Paszkowska, M., & Powroźnik, P. (2023). A comparative analysis of image segmentation using classical and deep learning approach. *Advances in Science and Technology Research Journal*, *17*(6), 127–139.

Qamar, S., Baba, A. I., Verger, S., & Andersson, M. (2024). Segmentation and characterization of macerated fibers and vessels using deep learning. *arXiv preprint arXiv:2401.16937*.

Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., & Kompenhans, J. (2018). *Particle image velocimetry: a practical guide*. springer.

Ravi, N., Gabeur, V., Hu, Y.-T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al. (2025). Sam 2: Segment anything in images and videos. In *International Conference on Learning Representations*, volume 2025, (pp. 28085–28128).

Ravichandran, D., Prem, P. R., Kaliyavaradhan, S. K., & Ambily, P. S. (2022). Influence of fibers on fresh and hardened properties of ultra high performance concrete (UHPC): A review. *Journal of Building Engineering*, *57*, 104922.

Richard, P. & Cheyrezy, M. (1995). Composition of reactive powder concretes. *Cement and Concrete Research*, *25*(7), 1501–1511.

Roussel, N. (2006a). Correlation between yield stress and slump: comparison between numerical simulations and concrete rheometers results. *Materials and Structures*, *39*(4), 501–509.

Roussel, N. (2006b). A thixotropy model for fresh fluid concretes: theory, validation and applications. *Cement and Concrete Research*, *36*(10), 1797–1806.

Scarano, F. (2002). Iterative image deformation methods in PIV. *Measurement Science and Technology*, *13*(1), R1–R19.

Shi, C., Wu, Z., & Banthia, N. (2024). *Ultra-High Performance Concrete: Design, Performance, and Application*. CRC Press.

Svec, O., Skocek, J., Olesen, J. F., & Stang, H. (2012). Fibre reinforced self-compacting concrete flow simulations in comparison with l-box experiments using carbopol. In Barros, J. A. O. (Ed.), *8th RILEM International Symposium on Fibre Reinforced Concrete (BEFIB 2012)*, (pp. 897–905)., Guimarães, Portugal.

Tapia, V. A. (2026). Experimental analysis of fiber-fluid dynamics in viscoplastic media using PIV. Trabajo de titulación para optar al título de magíster en ciencias de la ingeniería, Universidad de los Andes, Santiago, Chile. Profesor Guía: Patricio Moreno.

Tenne, D. & Singh, T. (2000). Optimal design of α-β-γ filters. In *Proceedings of the American Control Conference*, volume 6, (pp. 4348–4352).

Vaupel, T., Gerland, F., Schomberg, T., & Wünsch, O. (2025). Advanced optical measurement techniques for simultaneous fibre orientation and flow analysis complemented by machine learning. *Scientific Reports*, *15*, 38998.

Wang, D., Shi, C., Wu, Z., Xiao, J., Huang, Z., & Fang, Z. (2015). A review on ultra high performance concrete: Part II. Hydration, microstructure and properties. *Construction and Building Materials*, *96*, 368–377.

Wang, J., O'Gara, J. F., & Tucker, C. L. (2008). An objective model for slow orientation kinetics in concentrated fiber suspensions: Theory and rheological evidence. *Journal of Rheology*, *52*(5), 1179–1200.

Welch, G. & Bishop, G. (1995). An introduction to the Kalman filter. Technical Report TR 95-041, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC.

Westerweel, J. (1997). Fundamentals of digital particle image velocimetry. *Measurement Science and Technology*, *8*(12), 1379–1392.

Westerweel, J. (2008). On velocity gradients in PIV interrogation. *Experiments in Fluids*, *44*(5), 831–842.

Westerweel, J. & Scarano, F. (2005). Universal outlier detection for PIV data. *Experiments in Fluids*, *39*(6), 1096–1100.

Wille, K., Naaman, A. E., & Parra-Montesinos, G. J. (2011). Ultra-high performance concrete with compressive strength exceeding 150 MPa (22 ksi): a simpler way. *ACI Materials Journal*, *108*(1), 46–54.

Wille, K., Tue, N. V., & Parra-Montesinos, G. J. (2014). Fiber distribution and orientation in UHP-FRC beams and their effect on backward analysis. *Materials and Structures*, *47*(11), 1825–1838.

Wolff, L., Maggi, C., Toro Labbe, J. P., Delpiano, J., Gómez, J., Moreno-Casas, P. A., & Paul, Á. (2026). A deep learning-enhanced PTV framework for simultaneous translational and rotational tracking of rod-like particles in non-Newtonian fluids. *Engineering Research Express*.

Younes, E., Castelain, C., Burghelea, T., & Bertola, V. (2020). Slippery flows of a Carbopol gel in a micro-channel. *Physical Review Fluids*, *5*(1), 013301.

Zhang, Q. & Islam, M. M. (2021). Synthesis study quantifying the effect of UHPC fiber dispersion and orientation in structural members. Technical report, FAMU-FSU College of Engineering, Department of Civil and Environmental Engineering. Sponsored by Florida Department of Transportation.

Zhou, B. & Uchida, Y. (2017). Influence of flowability, casting time and formwork geometry on fiber orientation and mechanical properties of UHPFRC. *Cement and Concrete Research*, *95*, 164–177.

Chapter A

Derivation of the Orientation Saturation Law

This appendix derives, in self-contained form, the scalar orientation saturation law used in this work,

$$\eta_1(\Gamma) = \eta_{1,\max}\left(1 - e^{-\Gamma/\Gamma^*}\right), \tag{A.1}$$

together with the planar reduction under which it is confronted with the measurements. The law is not a new constitutive model. It is a compact two-parameter form obtained by projecting the alignment-diffusion balance of the Folgar-Tucker-Advani-Tucker equation onto a scalar order parameter and integrating the resulting relaxation with respect to the accumulated shear strain. The purpose of the appendix is to make every object and every step explicit, including the tensor definitions, the index algebra, and the two approximations on which the reduction rests, all of which are abbreviated in the main text. It is written to be read without reference to Chapter 5, and it consequently repeats, in fuller form, the results that chapter states without proof.

The presentation is organized as follows. Section A.1 defines the tensor objects and fixes the notation, and establishes that the order parameter of the model is identical to the quantity η_1 recovered from the images. Section A.2 derives the evolution equation for the orientation tensor from the single-fiber Jeffery equation and adds the Folgar-Tucker diffusion, in a general d-dimensional form. Section A.3 projects the three terms of that equation onto the scalar order parameter, establishing which projections are exact and which require approximation. Section A.4 shows why the accumulated shear strain, and not time, is the natural independent variable. Section A.5 solves the resulting relaxation equation. Section A.6 interprets the two parameters and examines their identifiability. Section A.7 carries out the planar reduction and states what it costs. Section A.8 collects the assumptions and states the domain within which the law is applicable.

A.1 Tensor Objects and Operations

The derivation uses a small set of tensor quantities, defined here in component form so that the subsequent manipulations are unambiguous. Latin indices run over the spatial directions, taking the values $1, 2, 3$ in three dimensions and $1, 2$ in the planar case, and repeated indices are summed unless stated otherwise. Where a result holds in both cases the dimension is written as d, with $d = 3$ or $d = 2$. The flow is assumed incompressible throughout, so that $\operatorname{tr} \boldsymbol{D} = 0$.

Orientation vector. A fiber is represented by a unit vector $\boldsymbol{p}$ directed along its axis, with $\boldsymbol{p} = (p_1, p_2, p_3)$ and $p_i p_i = 1$. A rod has no distinguished end, so $\boldsymbol{p}$ and $-\boldsymbol{p}$ denote the same physical orientation, and every descriptor built from $\boldsymbol{p}$ must be invariant under this reversal. The outer product $\boldsymbol{p} \otimes \boldsymbol{p}$ is the simplest such descriptor, since it is quadratic in $\boldsymbol{p}$ and therefore blind to the sign, which is the reason it, rather than $\boldsymbol{p}$ itself, is the object averaged below. The point is not a formality: averaging $\boldsymbol{p}$ over a population of fibers lying along a common axis but pointing in opposite senses would return zero, which is a perfectly aligned population reported as a disordered one.

Velocity gradient and its decomposition. The local flow acts on the fiber only through the velocity-gradient tensor, whose components are the derivatives of each velocity component with respect to each coordinate,

$$(\boldsymbol{\nabla u})_{ij} = \frac{\partial u_i}{\partial x_j}. \tag{A.2}$$

Any second-order tensor splits uniquely into a symmetric and an antisymmetric part, and this decomposition separates the two physically distinct actions that the flow exerts on a suspended rod,

$$\boldsymbol{D} = \tfrac{1}{2}\big(\boldsymbol{\nabla u} + \boldsymbol{\nabla u}^{\mathsf{T}}\big), \qquad \boldsymbol{W} = \tfrac{1}{2}\big(\boldsymbol{\nabla u} - \boldsymbol{\nabla u}^{\mathsf{T}}\big), \tag{A.3}$$

with components

$$D_{ij} = \tfrac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right), \qquad W_{ij} = \tfrac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right). \tag{A.4}$$

The rate-of-strain tensor $\boldsymbol{D}$ is symmetric, $D_{ij} = D_{ji}$; its diagonal entries measure stretching along each axis and its off-diagonal entries measure shear between pairs of directions. Because it is symmetric it possesses an orthonormal eigenbasis, and its largest eigenvalue D_1, with eigenvector $\boldsymbol{e}_1$, gives the direction along which the fluid is being stretched most strongly. This direction is the target toward which the deformation drives the fibers, and it reappears at every stage of the derivation. The rotation tensor $\boldsymbol{W}$ is antisymmetric, $W_{ij} = -W_{ji}$, so its diagonal vanishes identically and, in three dimensions, its three independent components are those of the vorticity vector $\boldsymbol{\zeta} = \boldsymbol{\nabla} \times \boldsymbol{u}$, through $W_{ij} = -\frac{1}{2}\,\epsilon_{ijk}\,\zeta_k$, so that $\boldsymbol{W} \cdot \boldsymbol{a} = \frac{1}{2}\,\boldsymbol{\zeta} \times \boldsymbol{a}$ for any vector $\boldsymbol{a}$. Its only real eigenvector is the vorticity itself, with eigenvalue zero, since $\boldsymbol{W} \cdot \boldsymbol{\zeta} = \frac{1}{2}\,\boldsymbol{\zeta} \times \boldsymbol{\zeta} = \boldsymbol{0}$; in the plane a single component remains, $W_{12} = -W_{21} = -\zeta/2$, and $\boldsymbol{W}$ has no real eigenvector. It encodes a local rigid rotation at half the vorticity, which moves the fibers without selecting any direction among them. The scalar magnitude of the deformation, which sets the strength of the diffusion and defines the accumulated strain, is the second invariant of $\boldsymbol{D}$,

$$\dot{\gamma} = \sqrt{2\,\boldsymbol{D}:\boldsymbol{D}} = \sqrt{2\,D_{ij}D_{ij}}, \tag{A.5}$$

normalized so that it reduces to the applied shear rate in a simple shear flow.

Second-order orientation tensor. The orientation state of a population of fibers at a point is described statistically by the second-order orientation tensor, defined as the average of the outer product of $\boldsymbol{p}$ with itself,

$$\boldsymbol{A} = \langle \boldsymbol{p} \otimes \boldsymbol{p} \rangle\,, \qquad A_{ij} = \langle p_i\, p_j \rangle\,, \tag{A.6}$$

where $\langle \cdot \rangle$ denotes the average over the orientation distribution at that point. Written out in three dimensions,

$$\boldsymbol{A} = \begin{pmatrix} \langle p_1^2 \rangle & \langle p_1 p_2 \rangle & \langle p_1 p_3 \rangle \\ \langle p_1 p_2 \rangle & \langle p_2^2 \rangle & \langle p_2 p_3 \rangle \\ \langle p_1 p_3 \rangle & \langle p_2 p_3 \rangle & \langle p_3^2 \rangle \end{pmatrix}, \tag{A.7}$$

each entry admits a direct reading. The diagonal entries are the fractions of orientation aligned with each coordinate direction, and the off-diagonal entries measure the corre-

lation, or tilt, between pairs of directions. The tensor is symmetric by construction, since $\langle p_i p_j \rangle = \langle p_j p_i \rangle$, so it carries six independent components in three dimensions and three in the plane.

Because every fiber is a unit vector, the trace of $\boldsymbol{A}$ is fixed,

$$\operatorname{tr} \boldsymbol{A} = A_{ii} = \langle p_i \, p_i \rangle = \langle 1 \rangle = 1, \tag{A.8}$$

so the diagonal entries behave as a set of fractions that always sum to one, a fixed budget of orientation distributed among the available directions. This constraint is not a modeling choice but an algebraic identity, and it is what makes the isotropic state a definite point rather than an adjustable one. Being symmetric, $\boldsymbol{A}$ is diagonalizable: its eigenvectors are the principal directions of orientation and its eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$, which sum to one by (A.8), measure the amount of orientation carried along each. Perfect alignment along a single direction corresponds to $\lambda_1 = 1$ and all others zero, while the isotropic state, in which no direction is preferred and the correlations vanish by symmetry, distributes the unit budget equally,

$$\boldsymbol{A}_{\text{iso}} = \frac{1}{d} \boldsymbol{I}, \tag{A.9}$$

so that $\lambda_1 = \frac{1}{3}$ in three dimensions and $\lambda_1 = \frac{1}{2}$ in the plane. The leading eigenvalue therefore ranges over $[1/d, 1]$ and is the natural raw measure of alignment. It is not, however, a suitable quantity to report, since it neither vanishes at isotropy nor takes the same isotropic value in the two dimensions.

Fourth-order orientation tensor. The evolution of $\boldsymbol{A}$, derived in Section A.2, involves the fourth-order orientation tensor,

$$\mathbb{A} = \langle \boldsymbol{p} \otimes \boldsymbol{p} \otimes \boldsymbol{p} \otimes \boldsymbol{p} \rangle \,, \qquad A_{ijkl} = \langle p_i \, p_j \, p_k \, p_l \rangle \,, \tag{A.10}$$

which enters through its double contraction with $\boldsymbol{D}$, a second-order tensor with components

$$(\mathbb{A} : \boldsymbol{D})_{ij} = A_{ijkl} \, D_{kl}. \tag{A.11}$$

The appearance of $\mathbb{A}$ in the equation for $\boldsymbol{A}$ is the origin of the closure problem: the evolution of each orientation moment depends on the next higher moment, so the hierarchy does not terminate and $\mathbb{A}$ must be approximated as a function of $\boldsymbol{A}$ through a closure approximation (Hinch et al., 1976; Cintra et al., 1995; Chung et al., 2002). The role of the closure is examined in Section A.3, where it is shown that it affects the deformation term alone and leaves the rotation and interaction terms untouched. One value of $\mathbb{A}$ is needed explicitly below, namely its isotropic value. Any isotropic fourth-order tensor is a linear combination of the three products $\delta_{ij}\delta_{kl}$, $\delta_{ik}\delta_{jl}$ and $\delta_{il}\delta_{jk}$ (Aris, 2012), and since A_{ijkl} is symmetric under any permutation of its indices the three must carry the same coefficient, which gives the isotropic orientation state (Ken-Ichi, 1984),

$$A^{\mathrm{iso}}_{ijkl} = c_d \left(\delta_{ij}\delta_{kl} + \delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk}\right), \qquad c_d = \frac{1}{d(d+2)}, \tag{A.12}$$

where the constant c_d is fixed by requiring that the contraction of the first two indices return the isotropic second-order tensor, $A^{\mathrm{iso}}_{iikl} = c_d(d+2)\,\delta_{kl} = \frac{1}{d}\delta_{kl}$. This gives $c_3 = \frac{1}{15}$ in three dimensions and $c_2 = \frac{1}{8}$ in the plane, the latter reproducing $\langle \cos^4\theta \rangle = \frac{3}{8}$ and $\langle \cos^2\theta \sin^2\theta \rangle = \frac{1}{8}$ for a uniform angular distribution.

Nematic tensor and the scalar order parameter. The tensor $\boldsymbol{A}$ carries the full orientation state, but the analysis requires a single number that is zero at isotropy and one at perfect alignment, in any dimension. The natural object is the nematic tensor, borrowed from the theory of liquid crystals (de Gennes et al., 1993; Doi et al., 1988),

$$\boldsymbol{Q} = \left\langle \frac{d\,\boldsymbol{p}\otimes\boldsymbol{p} - \boldsymbol{I}}{d-1} \right\rangle = \frac{d\,\boldsymbol{A} - \boldsymbol{I}}{d-1}, \tag{A.13}$$

which is the affine rescaling of $\boldsymbol{A}$ that subtracts the isotropic state and normalizes the remainder. Its trace vanishes identically, $\operatorname{tr}\boldsymbol{Q} = (d\cdot 1 - d)/(d-1) = 0$, so $\boldsymbol{Q}$ measures departure from isotropy and nothing else. Because $\boldsymbol{Q}$ is an affine function of $\boldsymbol{A}$, the two tensors share the same eigenvectors, and their eigenvalues are related by the same rescaling. The scalar order parameter is defined as the leading eigenvalue of $\boldsymbol{Q}$,

$$\eta_1 = \frac{d\,\lambda_1 - 1}{d-1}, \tag{A.14}$$

which vanishes at $\lambda_1 = 1/d$ and equals unity at $\lambda_1 = 1$, as required. Specializing (A.14) recovers the two forms used in this work,

$$\eta_1 = \frac{3\lambda_1 - 1}{2} \quad (d = 3), \qquad\qquad \eta_1 = 2\lambda_1 - 1 \quad (d = 2), \tag{A.15}$$

so that the different numerical coefficients are not two separate calibrations but a single definition evaluated in two dimensions. The inverse relation, needed repeatedly below, is $\lambda_1 = [(d-1)\eta_1 + 1]/d$.

A.1.1 Identity With the Measured Order Parameter

The symbol η_1 has been used in (A.14) for a tensor invariant, whereas in the measurement chapters it denotes a circular average over the detected fiber angles,

$$\eta_1 = \left|\left\langle e^{\,i2\theta}\right\rangle\right|. \tag{A.16}$$

The two are the same quantity, and the identification is exact. It is established here at the outset, because it is what licenses the whole of the derivation that follows to be written in a variable the experiment can report.

The relevant case is $d = 2$, since the PTV recovers the fibers projected onto the imaging plane. Let θ be the angle a fiber makes with the first coordinate axis, so that $\boldsymbol{p} = (\cos\theta,\ \sin\theta)$, for which $|\boldsymbol{p}|^2 = 1$ identically. The planar orientation tensor is

$$\boldsymbol{A} = \begin{pmatrix} \langle\cos^2\theta\rangle & \langle\cos\theta\,\sin\theta\rangle \\ \langle\cos\theta\,\sin\theta\rangle & \langle\sin^2\theta\rangle \end{pmatrix}, \tag{A.17}$$

a symmetric tensor whose trace is unity, so its two eigenvalues satisfy $\lambda_1 + \lambda_2 = 1$. Applying the double-angle identities $\cos^2\theta = \frac{1}{2}(1 + \cos 2\theta)$, $\sin^2\theta = \frac{1}{2}(1 - \cos 2\theta)$ and $\cos\theta\sin\theta = \frac{1}{2}\sin 2\theta$ separates it into an isotropic part and a traceless remainder,

$$\boldsymbol{A} = \tfrac{1}{2}\,\boldsymbol{I} + \tfrac{1}{2}\,\boldsymbol{M}, \qquad \boldsymbol{M} = \begin{pmatrix} \langle\cos 2\theta\rangle & \langle\sin 2\theta\rangle \\ \langle\sin 2\theta\rangle & -\langle\cos 2\theta\rangle \end{pmatrix}, \tag{A.18}$$

where $\boldsymbol{M}$ is symmetric with $\operatorname{tr}\boldsymbol{M} = 0$. The doubled angle is not introduced as a

convenience: it is produced by the algebra, and it is the same doubling that (A.16) contains. Its origin is the head-to-tail symmetry noted below (A.6). A rod at θ and a rod at $\theta+\pi$ are the same fiber; the map $\theta \mapsto 2\theta$ sends both onto the same point of the unit circle; the tensor, being quadratic in $\boldsymbol{p}$, is invariant under the reversal automatically, while the circular average is made invariant by doubling the angle explicitly. The two constructions enforce the same symmetry, which is why they agree.

A symmetric traceless 2×2 matrix of the form $\left(\begin{smallmatrix} \alpha & \beta \\ \beta & -\alpha \end{smallmatrix}\right)$ has characteristic polynomial $\mu^2 - (\alpha^2+\beta^2) = 0$, and therefore eigenvalues $\mu_\pm = \pm\sqrt{\alpha^2+\beta^2}$. With $\alpha = \langle \cos 2\theta \rangle$ and $\beta = \langle \sin 2\theta \rangle$,

$$\mu_\pm = \pm\sqrt{\langle \cos 2\theta \rangle^2 + \langle \sin 2\theta \rangle^2} = \pm \left| \left\langle e^{\,i2\theta} \right\rangle \right|, \tag{A.19}$$

the last equality being the definition of the modulus of the complex average $\langle e^{i2\theta} \rangle = \langle \cos 2\theta \rangle + i \langle \sin 2\theta \rangle$. The eigenvalues of $\boldsymbol{A}$ follow from (A.18), remembering that the traceless part enters multiplied by $\frac{1}{2}$,

$$\lambda_{1,2} = \tfrac{1}{2} + \tfrac{1}{2}\,\mu_\pm = \tfrac{1}{2}\left(1 \pm \left| \left\langle e^{\,i2\theta} \right\rangle \right|\right), \tag{A.20}$$

and substituting the leading one into the planar order parameter $\eta_1 = 2\lambda_1 - 1$ of (A.15),

$$\eta_1 = 2 \cdot \tfrac{1}{2}\left(1 + \left| \left\langle e^{\,i2\theta} \right\rangle \right|\right) - 1 = \left| \left\langle e^{\,i2\theta} \right\rangle \right|, \tag{A.21}$$

which is (A.16). The order parameter of the model is therefore *identical* to the quantity extracted from the images, and not merely proportional to it, so that no calibration constant intervenes between the theory and the measurement.

Two consequences follow. The first is that η_1 inherits the properties of a tensor invariant: it does not depend on the coordinate frame in which the angles happen to be measured, so the alignment reported in a cell of the mold is a property of the fiber population and not of the axes chosen to describe it. The second is that the *argument* of the same complex average carries the direction of the alignment, since $\frac{1}{2}\arg\langle e^{i2\theta} \rangle$ is the leading eigenvector of $\boldsymbol{M}$, and hence of $\boldsymbol{A}$, and is therefore the principal orientation direction, the director against which the angular deviation δ of the main text is defined. The modulus measures how ordered the population is, the argument measures the direction of its order, and both are recovered from a single average over the detected

angles.

The derivation that follows is written in η_1 throughout, without further distinction between the measured quantity and the theoretical one.

A.2 From the Single Fiber to the Population

The Jeffery equation. The rotation of a single rigid ellipsoid suspended in a Newtonian fluid at vanishing Reynolds number was obtained by Jeffery (Jeffery, 1922). In terms of the decomposition (A.3), the axis of the particle obeys

$$\frac{\mathrm{d}p_i}{\mathrm{d}t} = W_{ik}\,p_k + \xi\,(D_{ik}\,p_k - D_{kl}\,p_k\,p_l\,p_i)\,, \qquad \xi = \frac{r_f^2 - 1}{r_f^2 + 1}, \tag{A.22}$$

where r_f is the fiber aspect ratio and $\xi \to 1$ for slender fibers. The structure of the two contributions is worth reading carefully, because the whole derivation rests on it. The rotation term $W_{ik}p_k = \frac{1}{2}(\boldsymbol{\zeta} \times \boldsymbol{p})_i$ is a rigid rotation of the axis at half the local vorticity. The deformation term consists of the stretching $D_{ik}p_k$ minus its projection onto $\boldsymbol{p}$ itself, the subtraction being what preserves the constraint $\boldsymbol{p}\cdot\boldsymbol{p} = 1$: contracting (A.22) with p_i gives $p_i\dot{p}_i = W_{ik}p_ip_k + \xi(D_{ik}p_ip_k - D_{kl}p_kp_lp_l) = 0$, since the first term vanishes by the antisymmetry of $\boldsymbol{W}$ and the last two cancel. The rod therefore rotates and never stretches, as a rigid body must.

Evolution of the second-order tensor. The orientation tensor evolves through the outer product. Here and below, $\mathrm{d}/\mathrm{d}t$ denotes the rate of change following a fiber, which, since the fibers move with the fluid, coincides with the material derivative $\mathrm{D}/\mathrm{D}t$ defined after Equation (5.7) and recalled after Equation (A.30). The rate of change of the outer product then follows from the product rule,

$$\frac{\mathrm{d}}{\mathrm{d}t}\,(p_i\,p_j) = \dot{p}_i\,p_j + p_i\,\dot{p}_j. \tag{A.23}$$

Substituting (A.22) into each of the two contributions,

$$\dot{p}_i\,p_j = W_{ik}\,p_k\,p_j + \xi\,(D_{ik}\,p_k\,p_j - D_{kl}\,p_i\,p_j\,p_k\,p_l)\,, \tag{A.24}$$

$$p_i\,\dot{p}_j = W_{jk}\,p_i\,p_k + \xi\,(D_{jk}\,p_i\,p_k - D_{kl}\,p_i\,p_j\,p_k\,p_l)\,, \tag{A.25}$$

and adding them,

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(p_i\,p_j\right) = W_{ik}\,p_k\,p_j + W_{jk}\,p_i\,p_k + \xi\left(D_{ik}\,p_k\,p_j + D_{jk}\,p_i\,p_k - 2\,D_{kl}\,p_i\,p_j\,p_k\,p_l\right). \quad \text{(A.26)}$$

The terms are now rewritten in tensor form. Using the antisymmetry $W_{jk} = -W_{kj}$, the second term becomes $W_{jk}\,p_i\,p_k = -\,p_i\,p_k\,W_{kj} = -\left[(\boldsymbol{p}\otimes\boldsymbol{p})\cdot\boldsymbol{W}\right]_{ij}$, while the symmetry $D_{jk} = D_{kj}$ leaves the corresponding deformation term as $D_{jk}\,p_i\,p_k = \left[(\boldsymbol{p}\otimes\boldsymbol{p})\cdot\boldsymbol{D}\right]_{ij}$. It is here that the two parts of the velocity gradient part company: the antisymmetry converts the rotation contribution into a commutator, whereas the symmetry leaves the deformation contribution as an anticommutator. Equation (A.26) therefore reads

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\boldsymbol{p}\otimes\boldsymbol{p}\right) = \boldsymbol{W}\cdot(\boldsymbol{p}\otimes\boldsymbol{p}) - (\boldsymbol{p}\otimes\boldsymbol{p})\cdot\boldsymbol{W} + \xi\left[\boldsymbol{D}\cdot(\boldsymbol{p}\otimes\boldsymbol{p}) + (\boldsymbol{p}\otimes\boldsymbol{p})\cdot\boldsymbol{D} - 2\,(\boldsymbol{p}\otimes\boldsymbol{p}\otimes\boldsymbol{p}\otimes\boldsymbol{p}):\boldsymbol{D}\right]. \quad \text{(A.27)}$$

Taking the population average, and noting that $\boldsymbol{W}$ and $\boldsymbol{D}$ are properties of the flow at the point and pass outside the average, gives the Jeffery contribution to the evolution of the orientation tensor,

$$\left(\frac{\mathrm{d}\boldsymbol{A}}{\mathrm{d}t}\right)_{\text{Jeffery}} = \boldsymbol{W}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\boldsymbol{W} + \xi\left(\boldsymbol{D}\cdot\boldsymbol{A} + \boldsymbol{A}\cdot\boldsymbol{D} - 2\,\mathbb{A}:\boldsymbol{D}\right). \quad \text{(A.28)}$$

The fourth-order tensor enters through the single term $D_{kl}\,p_i\,p_j\,p_k\,p_l$, whose average is $A_{ijkl}D_{kl}$. The closure problem is therefore not an artifact of the tensorial formulation but a direct consequence of the cubic nonlinearity in $\boldsymbol{p}$ that the rigidity constraint imposes on (A.22).

The Folgar-Tucker diffusion. Equation (A.28) describes a dilute suspension, in which every fiber follows the deterministic trajectory dictated by the flow and no fiber is aware of any other. Folgar and Tucker (Folgar et al., 1984) extended it to concentrated suspensions by superimposing an isotropic rotary diffusion, which represents the collective effect of the hydrodynamic encounters between fibers as a randomizing agency acting on the population, in direct analogy with Brownian rotary diffusion but driven by the flow rather than by temperature. The requirement that fixes its form is that the diffusion must relax the orientation toward the isotropic state (A.9), and must

therefore be proportional to the distance from it,

$$\left(\frac{\mathrm{d}\boldsymbol{A}}{\mathrm{d}t}\right)_{\text{interaction}} \propto \left(\frac{1}{d}\boldsymbol{I} - \boldsymbol{A}\right) = \frac{1}{d}\left(\boldsymbol{I} - d\,\boldsymbol{A}\right), \tag{A.29}$$

where the constant $1/d$ is absorbed into the coefficient. Writing the coefficient as $2C_I\dot{\gamma}$, with C_I the dimensionless interaction coefficient and $\dot{\gamma}$ the scalar shear rate of (A.5), the interaction term takes the general form $2C_I\dot{\gamma}\,(\boldsymbol{I} - d\,\boldsymbol{A})$, which reduces to $2C_I\dot{\gamma}\,(\boldsymbol{I} - 3\boldsymbol{A})$ in three dimensions and to $2C_I\dot{\gamma}\,(\boldsymbol{I} - 2\boldsymbol{A})$ in the plane. Substituting the isotropic state confirms that it vanishes there, $\boldsymbol{I} - d\cdot\frac{1}{d}\boldsymbol{I} = \boldsymbol{0}$, as a physical diffusion must. Two properties deserve emphasis: its strength is proportional to $\dot{\gamma}$, so a suspension at rest neither aligns nor randomizes, and it is independent of $\boldsymbol{W}$, so the rotation of the fluid plays no part in the randomization.

The complete equation. Adding (A.28) and the interaction term gives the Folgar-Tucker-Advani-Tucker equation (Advani et al., 1987),

$$\frac{\mathrm{D}\boldsymbol{A}}{\mathrm{D}t} = \underbrace{(\boldsymbol{W}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\boldsymbol{W})}_{\text{rotation}} + \underbrace{\xi\,(\boldsymbol{D}\cdot\boldsymbol{A} + \boldsymbol{A}\cdot\boldsymbol{D} - 2\,\mathbb{A}\!:\!\boldsymbol{D})}_{\text{strain rate}} + \underbrace{2\,C_I\,\dot{\gamma}\,(\boldsymbol{I} - d\,\boldsymbol{A})}_{\text{interaction}}, \tag{A.30}$$

in which $\mathrm{D}\boldsymbol{A}/\mathrm{D}t = \partial\boldsymbol{A}/\partial t + (\boldsymbol{u}\cdot\boldsymbol{\nabla})\boldsymbol{A}$ is the material derivative, the rate of change following a fluid element. Because the fibers are assumed to move with the fluid, along the trajectory of that element it coincides with the ordinary time derivative, and the projections that follow are written with $\mathrm{d}/\mathrm{d}t$ in that sense. The interaction term acquires a particularly transparent form when written in terms of the nematic tensor (A.13). Since $\boldsymbol{I} - d\boldsymbol{A} = -\,(d\boldsymbol{A} - \boldsymbol{I}) = -(d-1)\,\boldsymbol{Q}$, the interaction term is exactly

$$2\,C_I\,\dot{\gamma}\,(\boldsymbol{I} - d\,\boldsymbol{A}) = -\,2\,(d-1)\,C_I\,\dot{\gamma}\,\boldsymbol{Q}, \tag{A.31}$$

a linear relaxation of $\boldsymbol{Q}$ toward zero. The randomizing mechanism is thus linear in the departure from isotropy, without approximation, a fact used directly in the projection of the next section.

A.3 Projection Onto the Scalar Order Parameter

The eigenvalue evolution formula. The order parameter (A.14) is built from the leading eigenvalue of $\boldsymbol{A}$, so its evolution requires the rate of change of that eigenvalue. Let $\lambda_1(t)$ be the leading eigenvalue and $\boldsymbol{v}_1(t)$ its unit eigenvector, assumed simple, so that $\boldsymbol{A}\boldsymbol{v}_1 = \lambda_1\boldsymbol{v}_1$ and $\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{v}_1 = 1$. Writing the eigenvalue as a Rayleigh quotient, $\lambda_1 = \boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{A}\,\boldsymbol{v}_1$, and differentiating,

$$\frac{\mathrm{d}\lambda_1}{\mathrm{d}t} = \dot{\boldsymbol{v}}_1^{\mathsf{T}}\boldsymbol{A}\,\boldsymbol{v}_1 + \boldsymbol{v}_1^{\mathsf{T}}\dot{\boldsymbol{A}}\,\boldsymbol{v}_1 + \boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{A}\,\dot{\boldsymbol{v}}_1. \tag{A.32}$$

The first and third terms are evaluated using the eigenvalue relation and the symmetry of $\boldsymbol{A}$, which give $\boldsymbol{A}\boldsymbol{v}_1 = \lambda_1\boldsymbol{v}_1$ and $\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{A} = \lambda_1\boldsymbol{v}_1^{\mathsf{T}}$, so that they combine into $\lambda_1\left(\dot{\boldsymbol{v}}_1^{\mathsf{T}}\boldsymbol{v}_1 + \boldsymbol{v}_1^{\mathsf{T}}\dot{\boldsymbol{v}}_1\right) = \lambda_1\,\mathrm{d}\big(\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{v}_1\big)/\mathrm{d}t = 0$, the normalization of the eigenvector being constant. The rotation of the eigenvector therefore drops out, and

$$\boxed{\frac{\mathrm{d}\lambda_1}{\mathrm{d}t} = \boldsymbol{v}_1^{\mathsf{T}}\,\frac{\mathrm{d}\boldsymbol{A}}{\mathrm{d}t}\,\boldsymbol{v}_1} \tag{A.33}$$

which states that the rate of change of the leading eigenvalue is the projection of the full tensor evolution onto the current principal direction. Applying (A.33) term by term to (A.30) decomposes the scalar balance into the three contributions examined next. Since $\boldsymbol{Q}$ and $\boldsymbol{A}$ share eigenvectors, the same formula applies to η_1 through $\mathrm{d}\eta_1/\mathrm{d}t = [d/(d-1)]\,\mathrm{d}\lambda_1/\mathrm{d}t$.

The rotation term contributes nothing. Projecting the commutator onto $\boldsymbol{v}_1$ and using both the eigenvalue relation and the symmetry of $\boldsymbol{A}$,

$$\boldsymbol{v}_1^{\mathsf{T}}\left(\boldsymbol{W}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\boldsymbol{W}\right)\boldsymbol{v}_1 = \boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{W}\left(\boldsymbol{A}\boldsymbol{v}_1\right) - \left(\boldsymbol{A}\boldsymbol{v}_1\right)^{\mathsf{T}}\boldsymbol{W}\boldsymbol{v}_1 = \lambda_1\,\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{W}\boldsymbol{v}_1 - \lambda_1\,\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{W}\boldsymbol{v}_1 = 0. \tag{A.34}$$

The two contributions cancel identically, and each vanishes on its own account as well, since $\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{W}\boldsymbol{v}_1 = 0$ for any antisymmetric $\boldsymbol{W}$ and any vector. The result is exact and requires no approximation of any kind.

The same conclusion holds over finite times, and in a stronger form. Consider the rotation term acting alone, $\mathrm{d}\boldsymbol{A}/\mathrm{d}t = \boldsymbol{W}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\boldsymbol{W}$, and let $\boldsymbol{R}(t)$ solve $\dot{\boldsymbol{R}} = \boldsymbol{W}\boldsymbol{R}$ with

$\boldsymbol{R}(0) = \boldsymbol{I}$. Then $\boldsymbol{R}$ is orthogonal at all times, since

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\boldsymbol{R}^{\mathsf{T}}\boldsymbol{R}\right) = \dot{\boldsymbol{R}}^{\mathsf{T}}\boldsymbol{R} + \boldsymbol{R}^{\mathsf{T}}\dot{\boldsymbol{R}} = \boldsymbol{R}^{\mathsf{T}}\left(\boldsymbol{W}^{\mathsf{T}} + \boldsymbol{W}\right)\boldsymbol{R} = \boldsymbol{0}, \tag{A.35}$$

by the antisymmetry of $\boldsymbol{W}$, and the solution is $\boldsymbol{A}(t) = \boldsymbol{R}(t)\,\boldsymbol{A}(0)\,\boldsymbol{R}(t)^{\mathsf{T}}$, as may be verified by direct substitution. The rotation term therefore acts on the orientation tensor by orthogonal similarity, which is to say by a rigid rotation of the principal axes, and orthogonal similarity leaves the entire spectrum of $\boldsymbol{A}$ invariant. The rotation of the fluid turns the orientation pattern; it does not alter its degree. This is the precise sense in which the alignment of the fibers is driven by the deformation and not by the rotation of the fluid, and it is the reason $\boldsymbol{W}$ does not appear anywhere in the scalar law derived below.

The interaction term projects exactly. Using the form (A.31), the interaction acts on $\boldsymbol{Q}$ as $(\mathrm{d}\boldsymbol{A}/\mathrm{d}t)_{\mathrm{int}} = -2(d-1)C_I\dot{\gamma}\,\boldsymbol{Q}$, and converting to the evolution of $\boldsymbol{Q}$ itself through $\mathrm{d}\boldsymbol{Q}/\mathrm{d}t = [d/(d-1)]\,\mathrm{d}\boldsymbol{A}/\mathrm{d}t$,

$$\left(\frac{\mathrm{d}\boldsymbol{Q}}{\mathrm{d}t}\right)_{\text{interaction}} = \frac{d}{d-1}\left[-2(d-1)\,C_I\,\dot{\gamma}\,\boldsymbol{Q}\right] = -\,2\,d\,C_I\,\dot{\gamma}\,\boldsymbol{Q}. \tag{A.36}$$

Projecting onto the leading eigenvector through (A.33), and recalling that η_1 is the leading eigenvalue of $\boldsymbol{Q}$, so that $\boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{Q}\,\boldsymbol{v}_1 = \eta_1$,

$$\left(\frac{\mathrm{d}\eta_1}{\mathrm{d}t}\right)_{\text{interaction}} = -\,2\,d\,C_I\,\dot{\gamma}\,\eta_1. \tag{A.37}$$

The randomizing contribution is therefore exactly proportional to the order that is present, with no approximation and no dependence on the closure. Anticipating the change of variable of Section A.4, dividing by $\dot{\gamma}$ converts (A.37) into

$$\left(\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma}\right)_{\text{interaction}} = -\,b\,\eta_1, \qquad b = 2\,d\,C_I, \tag{A.38}$$

so that $b = 6\,C_I$ in three dimensions and $b = 4\,C_I$ in the plane. The randomizing rate is thus not merely an increasing function of the interaction coefficient, as asserted in the main text, but is exactly proportional to it. This is the first of the two rates of the saturation law, and it is a material quantity: it depends on the fiber content and on

nothing else.

The deformation term does not project exactly. Projecting the deformation term onto $\boldsymbol{v}_1$ requires the three contractions

$$\boldsymbol{v}_1^{\mathsf{T}} \left(\boldsymbol{D}\cdot\boldsymbol{A}\right) \boldsymbol{v}_1 = \boldsymbol{v}_1^{\mathsf{T}} \left(\boldsymbol{A}\cdot\boldsymbol{D}\right) \boldsymbol{v}_1 = \lambda_1 \left(\boldsymbol{v}_1\cdot\boldsymbol{D}\cdot\boldsymbol{v}_1\right), \tag{A.39}$$

both following from $\boldsymbol{A}\boldsymbol{v}_1 = \lambda_1\boldsymbol{v}_1$ and the symmetry of $\boldsymbol{A}$, together with the contraction of the fourth-order term,

$$\boldsymbol{v}_1^{\mathsf{T}} \left(\mathbb{A}\!:\!\boldsymbol{D}\right) \boldsymbol{v}_1 = v_{1i}\, A_{ijkl}\, D_{kl}\, v_{1j} = \left\langle \left(\boldsymbol{p}\cdot\boldsymbol{v}_1\right)^2 \left(\boldsymbol{p}\cdot\boldsymbol{D}\cdot\boldsymbol{p}\right) \right\rangle. \tag{A.40}$$

Noting further that $\lambda_1 = \boldsymbol{v}_1^{\mathsf{T}}\boldsymbol{A}\boldsymbol{v}_1 = \langle(\boldsymbol{p}\cdot\boldsymbol{v}_1)^2\rangle$, and that $\boldsymbol{v}_1\cdot\boldsymbol{D}\cdot\boldsymbol{v}_1$ is a constant with respect to the population average, the three contractions combine into a single compact expression,

$$\left(\frac{\mathrm{d}\lambda_1}{\mathrm{d}t}\right)_{\text{strain rate}} = 2\,\xi \left\langle \left(\boldsymbol{p}\cdot\boldsymbol{v}_1\right)^2 \left[\left(\boldsymbol{v}_1\cdot\boldsymbol{D}\cdot\boldsymbol{v}_1\right) - \left(\boldsymbol{p}\cdot\boldsymbol{D}\cdot\boldsymbol{p}\right)\right] \right\rangle. \tag{A.41}$$

Equation (A.41) is exact, and it repays reading. The leading eigenvalue grows when the fibers that already lie near the principal orientation direction, weighted by $(\boldsymbol{p}\cdot\boldsymbol{v}_1)^2$, are being stretched less strongly than that direction itself, and it is precisely the excess stretching available to $\boldsymbol{v}_1$ over the individual fibers that drives them toward it. The expression vanishes identically at perfect alignment, since all fibers then satisfy $\boldsymbol{p} = \pm\boldsymbol{v}_1$, the bracket closes, and $\mathrm{d}\lambda_1/\mathrm{d}t = 0$: a fully aligned population cannot be aligned further, and the deformation term switches itself off, as required.

Two features of (A.41) prevent it from being written as a function of η_1 alone. The first is the fourth moment: the average $\langle(\boldsymbol{p}\cdot\boldsymbol{v}_1)^2(\boldsymbol{p}\cdot\boldsymbol{D}\cdot\boldsymbol{p})\rangle$ is a contraction of $\mathbb{A}$, and under any closure other than the linear one it is a nonlinear function of $\boldsymbol{A}$. The second is the orientation of the principal axes: the expression depends on the angle between $\boldsymbol{v}_1$, the principal direction of the fibers, and $\boldsymbol{e}_1$, the principal stretching direction of the flow, and that angle is not encoded in η_1. It follows that the deformation term is not linear in the orientation state, and that the statement, sometimes made in passing, that the scalar reduction follows from the linearity of both mechanisms is incorrect. The reduction of the deformation term to a function of η_1 is an approximation, and it is stated as such below.

The two approximations of the scalar reduction. The reduction rests on two assumptions, which are here made explicit.

Assumption A1 (axis locking). The leading eigenvector of $\boldsymbol{A}$ remains close to the principal stretching direction of $\boldsymbol{D}$, so that $\boldsymbol{v}_1 \approx \boldsymbol{e}_1$ and $\boldsymbol{v}_1 \cdot \boldsymbol{D} \cdot \boldsymbol{v}_1 \approx D_1$. This is the condition under which the deformation term becomes a function of the degree of alignment alone, the direction of alignment having been fixed by the flow. It holds after the initial transient in a flow whose principal directions vary slowly along the fiber trajectory, and it fails whenever the flow reverses or turns sharply, in which case $\boldsymbol{v}_1$ lags behind $\boldsymbol{e}_1$ and the deformation can transiently reduce λ_1 rather than increase it.

Assumption A2 (linear interpolation). Under A1, and after closure, the deformation projection becomes a function $f(\eta_1)$ satisfying two exact constraints: it vanishes at $\eta_1 = 1$, as shown above, and it is maximal at $\eta_1 = 0$, where the population offers the largest room for alignment. The simplest function compatible with both is the straight line joining them,

$$\left(\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma}\right)_{\text{strain rate}} = a\,(1 - \eta_1), \tag{A.42}$$

which is the first-order truncation of $f(\eta_1)$ about the aligned fixed point, extended to the whole interval. The constant a is therefore an *effective* ordering rate, not a fundamental one: it absorbs the shape factor, the closure, the residual misalignment tolerated by A1, and the variation of the flow kinematics along the fiber trajectory. The consequences for the identifiability of the parameters are examined in Section A.6.

Magnitude of the ordering rate. Although a is an effective constant, its order of magnitude is recoverable by evaluating (A.41) at the isotropic state, where $\eta_1 = 0$ and the linear form (A.42) returns a itself. Contracting the isotropic fourth moment (A.12) with $\boldsymbol{v}_1$ and $\boldsymbol{D}$, and using the incompressibility condition $\operatorname{tr}\boldsymbol{D} = 0$,

$$v_{1i}\, A^{\text{iso}}_{ijkl}\, D_{kl}\, v_{1j} = c_d\,[(\boldsymbol{v}_1 \cdot \boldsymbol{v}_1)\operatorname{tr}\boldsymbol{D} + 2\,(\boldsymbol{v}_1 \cdot \boldsymbol{D} \cdot \boldsymbol{v}_1)] = 2\,c_d\,(\boldsymbol{v}_1 \cdot \boldsymbol{D} \cdot \boldsymbol{v}_1)\,, \tag{A.43}$$

and substituting into (A.41) with $\lambda_1 = 1/d$ and, by A1, $\boldsymbol{v}_1 \cdot \boldsymbol{D} \cdot \boldsymbol{v}_1 = D_1$,

$$\left(\frac{\mathrm{d}\lambda_1}{\mathrm{d}t}\right)^{\text{iso}}_{\text{strain rate}} = 2\,\xi\,D_1\left[\frac{1}{d} - \frac{2}{d(d+2)}\right] = \frac{2\,\xi\,D_1}{d+2}. \tag{A.44}$$

Converting to the order parameter through $\mathrm{d}\eta_1 = [d/(d-1)]\,\mathrm{d}\lambda_1$ and dividing by $\dot{\gamma}$ gives the ordering rate,

$$a = \frac{2\,d\,\xi}{(d-1)(d+2)} \cdot \frac{D_1}{\dot{\gamma}}, \tag{A.45}$$

which is $a = \frac{3}{5}\,\xi\,(D_1/\dot{\gamma})$ in three dimensions and $a = \xi\,(D_1/\dot{\gamma})$ in the plane. In a simple shear flow the rate-of-strain tensor has eigenvalues $\pm\dot{\gamma}/2$, so $D_1/\dot{\gamma} = \frac{1}{2}$ and the two rates become $a \simeq 0.30\,\xi$ and $a \simeq 0.50\,\xi$ respectively, approaching 0.30 and 0.50 for slender fibers.

Equation (A.45) exposes a structural asymmetry between the two rates of the model, and it is worth stating plainly. The randomizing rate $b = 2dC_I$ is a *material* quantity, depending on the fiber content alone. The ordering rate a is a *kinematic* quantity, depending on the fiber slenderness through ξ and on the character of the flow through the ratio $D_1/\dot{\gamma}$. With $\dot{\gamma} = \sqrt{2\,\boldsymbol{D}:\boldsymbol{D}}$ this ratio equals $\frac{1}{2}$ both in simple shear and in planar elongation, and reaches its largest value, $1/\sqrt{3}$, in uniaxial elongation; in the plane of the measurement it equals $\frac{1}{2}$ for any flow, since a traceless 2×2 tensor has eigenvalues $\pm D_1$. Per unit of accumulated strain, stretching and shearing therefore supply nearly the same ordering rate, and the difference between the two flows lies in the rotation term, which in shear turns the director away from the direction of stretching and is set aside by the axis-locking assumption.

A.4 The Accumulated Shear Strain as the Independent Variable

The two contributions (A.38) and (A.42) have already been written with respect to the accumulated shear strain,

$$\Gamma = \int_0^t \dot{\gamma}\,(t')\;\mathrm{d}t', \tag{A.46}$$

rather than time, and the reason lies in the structure of (A.30). Introducing the normalized kinematic tensors $\hat{\boldsymbol{D}} = \boldsymbol{D}/\dot{\gamma}$ and $\hat{\boldsymbol{W}} = \boldsymbol{W}/\dot{\gamma}$, which describe the *shape* of the local flow with its intensity divided out, the full equation can be written as

$$\frac{\mathrm{d}\boldsymbol{A}}{\mathrm{d}t} = \dot{\gamma}\left[\left(\hat{\boldsymbol{W}}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\hat{\boldsymbol{W}}\right) + \xi\left(\hat{\boldsymbol{D}}\cdot\boldsymbol{A} + \boldsymbol{A}\cdot\hat{\boldsymbol{D}} - 2\,\mathbb{A}\!:\!\hat{\boldsymbol{D}}\right) + 2\,C_I\,(\boldsymbol{I} - d\,\boldsymbol{A})\right], \tag{A.47}$$

in which the shear rate multiplies the whole of the right-hand side, the deterministic terms and the diffusive one alike. It follows immediately from (A.46) that $\mathrm{d}\Gamma = \dot{\gamma}\,\mathrm{d}t$, and dividing (A.47) by $\dot{\gamma}$ eliminates it entirely,

$$\frac{\mathrm{d}\boldsymbol{A}}{\mathrm{d}\Gamma} = \left(\hat{\boldsymbol{W}}\cdot\boldsymbol{A} - \boldsymbol{A}\cdot\hat{\boldsymbol{W}}\right) + \xi\left(\hat{\boldsymbol{D}}\cdot\boldsymbol{A} + \boldsymbol{A}\cdot\hat{\boldsymbol{D}} - 2\,\mathbb{A}\!:\!\hat{\boldsymbol{D}}\right) + 2\,C_I\left(\boldsymbol{I} - d\,\boldsymbol{A}\right). \qquad \text{(A.48)}$$

The shear rate has disappeared from the problem. It sets the pace at which the orientation evolves, but not the state toward which it evolves, and the orientation acquired by the fibers depends on how much deformation has been accumulated along their path and not on how quickly it was delivered. Two flows that accumulate the same Γ through the same sequence of flow shapes reach the same orientation, whatever their speed. Pouring a mix slowly or quickly through the same geometry does not change the alignment it acquires; changing the geometry does.

Three remarks qualify this conclusion in the present application. First, it presupposes that C_I and ξ are constants, which is the assumption of the model. Second, Γ is a Lagrangian quantity, accumulated along the trajectory of a fluid element, and its value at a point of the hardened element is the integral of the shear rate experienced by the material that came to rest there, not a property of the point itself; its evaluation therefore requires the flow history, which is what the PIV of this work supplies and what a validated simulation would supply beyond the reach of the experiment. Third, the rate independence of (A.48) concerns the *orientation* equation and not the momentum equation: the viscoplastic rheology of fresh UHPC — shear-thinning or shear-thickening depending on the mix, as Section 2.2.1 discusses — enters the problem by determining the velocity field, hence $\boldsymbol{D}$, $\boldsymbol{W}$ and $\dot{\gamma}$, and it is through this dependence, and not through (A.48), that the rheology of the mix governs the final orientation (de Larrard et al., 1998; Martinie et al., 2011).

A.5 The Relaxation Equation and its Explicit Solution

Adding the ordering contribution (A.42) and the randomizing contribution (A.38), and recalling that the rotation term contributes nothing by (A.34), the orientation of the

population obeys the scalar balance

$$\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} = a\,(1-\eta_1) - b\,\eta_1, \tag{A.49}$$

in which the first term pushes η_1 upward in proportion to the disorder that remains and the second pulls it downward in proportion to the order that is present. Expanding and grouping the terms in η_1,

$$\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} = a - (a+b)\,\eta_1, \tag{A.50}$$

and factoring the sum of the two rates,

$$\frac{\mathrm{d}\eta_1}{\mathrm{d}\Gamma} = (a+b)\left[\frac{a}{a+b} - \eta_1\right], \tag{A.51}$$

which exposes the structure of the balance. The orientation changes at a rate proportional to the distance that still separates it from the value $a/(a+b)$, and it changes more slowly the closer it comes to it. This is a first-order relaxation, of the same form as the charging of a capacitor through a resistor, the cooling of a body toward the temperature of its surroundings, or the approach of a reversible chemical reaction to its equilibrium concentration. In every case the equilibrium is fixed by the ratio of two competing rates and the pace of the approach by their sum, and in every case it is reached asymptotically, not in finite time.

The equilibrium is obtained by setting $\mathrm{d}\eta_1/\mathrm{d}\Gamma = 0$ in (A.49), which gives $a\,(1-\eta_{1,\mathrm{eq}}) = b\,\eta_{1,\mathrm{eq}}$, and solving,

$$\eta_{1,\mathrm{eq}} = \frac{a}{a+b}. \tag{A.52}$$

It is a partially aligned state, attained not because either mechanism has ceased to act but because the two have come into balance, and it lies strictly between complete disorder and complete alignment for any finite $b > 0$. This is the central prediction of the model: no casting flow, however long, can align the fibers completely, so long as they interact.

Equation (A.51) is separable. Writing $k = a + b$ and $\eta_\infty = a/(a+b)$, it reads $\mathrm{d}\eta_1/\mathrm{d}\Gamma =$

$k\,(\eta_\infty - \eta_1)$, and separating the variables,

$$\frac{\mathrm{d}\eta_1}{\eta_\infty - \eta_1} = k\,\mathrm{d}\Gamma. \tag{A.53}$$

The left-hand side is integrated with the substitution $u = \eta_\infty - \eta_1$, for which $\mathrm{d}u = -\,\mathrm{d}\eta_1$ and therefore $\int \mathrm{d}\eta_1/(\eta_\infty - \eta_1) = -\int \mathrm{d}u/u = -\ln|u|$, so that

$$-\ln|\eta_\infty - \eta_1| = k\,\Gamma + C_1, \tag{A.54}$$

with C_1 a constant of integration. Multiplying by -1 and exponentiating both sides,

$$|\eta_\infty - \eta_1| = e^{-C_1}\,e^{-k\Gamma} \equiv C\,e^{-k\Gamma}, \tag{A.55}$$

where $C = e^{-C_1}$ absorbs the constant together with the sign of the difference, which is fixed by the initial condition and does not change, since η_1 approaches η_∞ monotonically and cannot cross it. Solving,

$$\eta_1(\Gamma) = \eta_\infty - C\,e^{-k\Gamma}, \tag{A.56}$$

and imposing the entry condition $\eta_1(\Gamma = 0) = \eta_{1,0}$, which gives $C = \eta_\infty - \eta_{1,0}$, the general solution is

$$\eta_1(\Gamma) = \eta_\infty + (\eta_{1,0} - \eta_\infty)\,e^{-k\Gamma}. \tag{A.57}$$

The population departs from whatever order it possessed on entry and approaches the equilibrium exponentially, the initial condition being forgotten at the rate k. It is worth noting that (A.57) does not require $\eta_{1,0} < \eta_\infty$: a population entering the mold *more* aligned than the equilibrium would be disordered by the flow down to the same value, the equilibrium being an attractor from both sides.

When the fibers enter from a disordered state, $\eta_{1,0} = 0$, and (A.57) collapses to the two-parameter form

$$\eta_1(\Gamma) = \eta_\infty\left(1 - e^{-k\Gamma}\right). \tag{A.58}$$

The two parameters are now read directly from the solution rather than postulated.

The asymptotic value approached as $\Gamma \to \infty$ is the maximum attainable order,

$$\eta_{1,\max} \equiv \eta_\infty = \frac{a}{a+b}, \tag{A.59}$$

which is the equilibrium (A.52), and the characteristic strain, defined so that the exponent takes the standard form $e^{-\Gamma/\Gamma^*}$, is the reciprocal of the total rate constant,

$$\Gamma^* \equiv \frac{1}{k} = \frac{1}{a+b}. \tag{A.60}$$

Substituting (A.59) and (A.60) into (A.58) recovers the saturation law (A.1), completing the derivation. Neither parameter was assumed: $a/(a+b)$ arises from factoring $(a+b)$ in (A.51) and is the equilibrium value of the balance, while $1/(a+b)$ is the reciprocal of the rate constant multiplying Γ in the exponent.

A.5.1 The Predicted Magnitude of the Characteristic Strain

Substituting the rates obtained in Section A.3, $b = 2dC_I$ and a given by (A.45), places the two parameters in terms of the physical variables of the problem,

$$\eta_{1,\max} = \left[1 + \frac{(d-1)(d+2)}{\xi} \cdot \frac{\dot\gamma}{D_1} C_I\right]^{-1}, \qquad \Gamma^* = \left[\frac{2d\,\xi}{(d-1)(d+2)} \cdot \frac{D_1}{\dot\gamma} + 2\,d\,C_I\right]^{-1}, \tag{A.61}$$

which shows explicitly that the fiber content enters both parameters, and enters both through C_I alone. For slender fibers in a planar shear flow, for which $\xi \approx 1$, $d = 2$ and $D_1/\dot\gamma = \frac{1}{2}$, these reduce to $\eta_{1,\max} = (1 + 8\,C_I)^{-1}$ and $\Gamma^* = (0.5 + 4\,C_I)^{-1}$, which places Γ^* at order unity for the interaction coefficients reported for fiber suspensions (Bay et al., 1992).

The characteristic strains fitted in this work are one to two orders of magnitude larger than that, and the discrepancy is not a numerical accident but the clearest evidence the data offer against the idealization from which (A.61) descends. Since $\Gamma^* = 1/(a+b)$ and b can only grow with the dosage, a characteristic strain far in excess of unity requires an ordering rate a far below the value (A.45) predicts. Three of the assumptions collected in Section A.8 would produce exactly that. Assumption A1 fails wherever the flow turns, so the principal direction of the fibers lags behind the principal stretching direction of

the flow and the deformation orders the population less efficiently than the aligned-axis estimate assumes. The estimate (A.45) is evaluated at the isotropic state, where the ordering is fastest, and $f(\eta_1)$ is convex rather than linear in a real suspension, so A2 overestimates the ordering supplied at every intermediate degree of alignment. And the deformation delivered in a yield-stress fluid is confined to thin sheared layers, so a fiber travelling within the plug accumulates strain that is counted in Γ only to the extent that it passes through those layers, and the strain that reaches the fiber is a fraction of the strain the flow contains. The prediction of (A.61) is therefore an upper bound on the ordering rate and a lower bound on the characteristic strain, and it is reported here as such: it establishes the scale on which the idealized model would operate, and the distance between that scale and the measured one is the measure of how far a casting departs from the flows for which the model was built.

A.6 Interpretation and Identifiability of the Parameters

Meaning of the two parameters. The maximum order $\eta_{1,\max} = a/(a+b)$ is the share of the ordering rate within the total of the ordering and randomizing rates, and it is the ceiling of the curve: the highest degree of alignment the population can reach, however long the mix continues to flow. When the flow dominates, at low fiber content, $a \gg b$ and $\eta_{1,\max}$ approaches unity; when the interactions dominate, at high fiber content, b grows and $\eta_{1,\max}$ falls. The characteristic strain $\Gamma^* = 1/(a+b)$ is the strain scale of the saturation. At $\Gamma = \Gamma^*$ the population has covered a fraction $1 - e^{-1} \approx 0.63$ of the distance to $\eta_{1,\max}$, at $\Gamma = 2\Gamma^*$ a fraction $1 - e^{-2} \approx 0.86$, and by $\Gamma \approx 3\Gamma^*$ it is saturated for practical purposes. It is not a threshold below which nothing happens: the alignment grows fastest at $\Gamma = 0$ and decelerates thereafter, and Γ^* marks the strain at which the approach is essentially complete rather than the strain at which it begins.

Recovering the individual rates. The two rates can be recovered algebraically from the fitted parameters. Inverting (A.59) and (A.60) gives $a+b = 1/\Gamma^*$ and $\eta_{1,\max} = a\,\Gamma^*$, and therefore

$$a = \frac{\eta_{1,\max}}{\Gamma^*}, \qquad b = \frac{1 - \eta_{1,\max}}{\Gamma^*}. \tag{A.62}$$

These values are not reported as physical results, for three reasons. First, a is an effective constant, as established in Section A.3: it lumps together the shape factor, the closure, the residual violation of the axis-locking assumption, and the variation of the flow kinematics along the fiber trajectory, and it therefore has no universal interpretation. Second, (A.62) inherits and amplifies the uncertainty of $\eta_{1,\max}$ and Γ^* through the division, and is undefined wherever $\eta_{1,\max}$ is itself unidentifiable, which occurs whenever the accumulated strain available in the experiment does not reach the saturation plateau. Third, and most importantly, the two-constant form imposes a coupling between $\eta_{1,\max}$ and Γ^* that the measurements do not support, as shown next.

The two-constant form is structurally incomplete. Both parameters share the denominator $(a + b)$, and this couples them. Suppose, as the model assumes, that a is independent of the fiber content and that the interactions act only by raising $b = 2dC_I$. Then increasing the dosage increases b, which increases $(a+b)$, which lowers $\Gamma^* = 1/(a+b)$ and lowers $\eta_{1,\max} = a/(a+b)$ together. A richer mix would align less, but would saturate *sooner*. The experimental observation that Γ^* *increases* with the fiber content while $\eta_{1,\max}$ decreases is therefore not merely a quantitative discrepancy but a qualitative contradiction of the constant-a assumption, and it can be turned into a statement about a. Writing ϕ for the fiber volume fraction, the observation $\mathrm{d}\Gamma^*/\mathrm{d}\phi > 0$ requires

$$\frac{\mathrm{d}}{\mathrm{d}\phi}\,(a+b) < 0 \qquad \Longrightarrow \qquad \frac{\mathrm{d}a}{\mathrm{d}\phi} < -\,\frac{\mathrm{d}b}{\mathrm{d}\phi} = -\,2\,d\,\frac{\mathrm{d}C_I}{\mathrm{d}\phi} < 0, \tag{A.63}$$

since C_I increases with ϕ. The same conclusion follows more directly from (A.62): $a = \eta_{1,\max}/\Gamma^*$ has a decreasing numerator and an increasing denominator, so a must decrease with the dosage, and it must decrease faster than b increases. The interactions therefore do not merely randomize the population; they also *degrade the effective ordering rate supplied by the flow*, so that a concentrated suspension aligns not only less but more slowly, with respect to accumulated strain, than a dilute one. This is the slow-kinetics behavior documented for concentrated fiber suspensions, and it motivated the reduced-strain-closure and anisotropic-rotary-diffusion extensions of the Folgar-Tucker model (Wang et al., 2008; Phelps et al., 2009). The present work does not adopt those extensions, since the object here is a reduced description rather than a full constitutive model, but it inherits their diagnosis: $\eta_{1,\max}$ and Γ^* are treated as independent phe-

nomenological parameters, fitted directly to the measured orientation, and not as two functions of a single interaction coefficient.

The dimensionless ratio. Where a connection to the interaction coefficient is nevertheless required, it is made through the ratio of the two rates, which is free of the division by Γ^* and therefore of its uncertainty,

$$\frac{b}{a} = \frac{1 - \eta_{1,\max}}{\eta_{1,\max}}, \tag{A.64}$$

a quantity obtained from $\eta_{1,\max}$ alone. Substituting $b = 2dC_I$ and (A.45),

$$\frac{b}{a} = \frac{(d-1)(d+2)}{\xi} \cdot \frac{\dot{\gamma}}{D_1} \cdot C_I, \tag{A.65}$$

so that b/a is proportional to C_I *if and only if* a is independent of the fiber content. Since (A.63) establishes that it is not, (A.65) must be read with care. The ratio remains a monotonically increasing measure of the interaction strength, and it is the quantity the data determine most robustly, but inverting it for C_I under the constant-a assumption overestimates the interaction coefficient, since the fitted $b/a = 2dC_I/a(\phi)$ carries the decrease of a as well as the increase of C_I. Accordingly, b/a is reported here as a dimensionless index of the disordering strength of the mix, and not as a determination of C_I.

The entry condition. The reduction to the two-parameter form (A.58) rests on the assumption $\eta_{1,0} = 0$. The general solution (A.57) retains a nonzero entry order, and collapses to the two-parameter form only when $\eta_{1,0} \approx 0$. In the experiments reported here the condition is supplied by the casting procedure itself: the material falls from the outlet of the L-shape into the beam, and the impact is a disordered deformation event that the alignment carried out of the upstream geometry does not survive, so the fibers arrive in a near-isotropic state independently of the strain they accumulated before it. The consequence of the assumption failing is bounded rather than open. A nonzero entry order absorbed into a two-parameter fit is not simply an error of scale: expanding (A.57) for small Γ shows that an unmeasured $\eta_{1,0} > 0$ raises the intercept of the curve, which the fit can only accommodate by inflating the ceiling and shortening

the characteristic strain, so that the fitted $\eta_{1,\max}$ is biased upward and the fitted Γ^* downward. Where the reset cannot be assumed, the three-parameter form must be fitted instead, at the cost of a parameter that the data are rarely able to constrain, since $\eta_{1,0}$ and Γ^* are strongly correlated in the early part of the curve, which is precisely the part a casting delivers.

A.7 Planar Reduction and What it Costs

The orientation measurements are obtained on a single imaging plane, so the model is applied in its two-dimensional form. The identity between the planar order parameter and the measured quantity was established in Section A.1.1, and is not repeated. What the reduction changes, and what it costs, are collected here.

What changes. Three numerical factors take their planar values, and no other element of the derivation is affected. The isotropic reference becomes $\boldsymbol{A}_{\text{iso}} = \frac{1}{2}\boldsymbol{I}$ rather than $\frac{1}{3}\boldsymbol{I}$, the unit budget of orientation being shared among two directions rather than three, and the interaction term becomes $2C_I\dot{\gamma}\,(\boldsymbol{I} - 2\boldsymbol{A})$ accordingly. The order parameter is $\eta_1 = 2\lambda_1 - 1$ rather than $(3\lambda_1 - 1)/2$, from (A.15). And the two rates take the planar values $b = 4\,C_I$ and $a = \xi\,D_1/\dot{\gamma}$, from (A.38) and (A.45), the planar ordering rate being larger than the spatial one by a factor of $\frac{5}{3}$, since a fiber confined to a plane has one fewer direction into which the deformation can fail to drive it. The isotropic fourth moment likewise takes $c_2 = \frac{1}{8}$ in place of $c_3 = \frac{1}{15}$.
The saturation law (A.1), the first-order relaxation from which it follows, and the interpretation of $\eta_{1,\max}$ and Γ^* are unchanged. The form of the curve is dimension-independent.

Consistency of the two variables. The accumulated deformation is evaluated from the in-plane velocity gradient, with the planar shear rate $\dot{\gamma} = \sqrt{2\,D_{ij}D_{ij}}$ and the indices restricted to the measurement plane, so that the dependent variable η_1 and the independent variable Γ are defined within the same plane and the fit is internally consistent. This is not automatic and it is worth stating: a planar order parameter read against a strain computed from a three-dimensional velocity field would mix two descriptions and produce a fit whose parameters belong to neither.

The cost: out-of-plane projection. The reduction assumes that the out-of-plane component of the fiber orientation is small, which holds when the imaging plane contains the principal flow direction, and it introduces a bias whenever it does not. A fiber inclined out of the plane by an angle ψ projects onto it as a shorter rod, and its projected in-plane angle is not the angle it holds in space. The bias is systematic and it has a definite sign. A rod inclined out of the plane projects toward the direction of the projection axis, which is to say that out-of-plane fibers appear more aligned with the imaging plane than they are, and the measured η_1 therefore *overestimates* the alignment of the three-dimensional population. The magnitude of the effect grows with ψ, and it is smallest exactly where the reduction is most nearly valid.

Two circumstances of the present experiment bound this. The flow is confined between two walls separated by a distance small compared with the length of the mold, so the out-of-plane velocity component is small and the fibers have little occasion to acquire a large ψ. And the bias, being systematic, acts in the same direction at every dosage, so it displaces the three curves of the saturation law together and does not manufacture a separation between them. It inflates $\eta_{1,\max}$; it does not produce the ordering of the dosages that the analysis reports. The reservation therefore attaches to the absolute values of the fitted ceiling and not to the comparison across fiber contents, which is the conclusion the measurements are asked to support.

A.8 Assumptions and Domain of Validity

The derivation has introduced its assumptions at the points where they are needed, and they are collected here so that the range within which the saturation law may be applied is unambiguous. Table A.1 lists them, together with the step at which each enters and the consequence of its failure.

Two of these deserve a closing remark, because they bound the application most tightly. The first is the requirement of a monotone flow. The independent variable Γ accumulates the magnitude of the deformation and is insensitive to its direction, so a flow that reverses, or that turns sharply around a corner of the mold, continues to increase Γ while the principal stretching direction rotates away from the direction in which the fibers have already been aligned. Under such conditions the deformation term ceases to

Table A.1: Assumptions underlying the orientation saturation law.

Assumption	**Enters at**	**Consequence of failure**
Rigid fibers of uniform aspect ratio	Eq. (A.22)	The shape factor ξ is not a single constant; a becomes a distribution average.
Fibers follow the fluid, without migration or settling	Eq. (A.6)	The local fiber content, and therefore C_I, varies with position and b is no longer uniform.
Semi-dilute regime: hydrodynamic interactions, no contact network	Eq. (A.29)	At high dosage, direct contact and jamming introduce mechanisms that a rotary diffusion cannot represent.
Axis locking, $\boldsymbol{v}_1 \approx \boldsymbol{e}_1$ (A1)	Eq. (A.42)	The deformation term is no longer a function of η_1 alone and may transiently reduce the alignment.
Linear interpolation of the ordering term (A2)	Eq. (A.42)	a acquires a dependence on η_1; the exponential form is only approximate.
Monotone flow, no reversal of the principal stretching direction	Eq. (A.46)	Γ continues to accumulate while the alignment falls, and the law loses its monotonicity.
Isotropic entry, $\eta_{1,0} = 0$	Eq. (A.58)	The fitted ceiling is biased upward and the fitted characteristic strain downward.
Orientation confined to the imaging plane	Sec. A.7	The measured η_1 overestimates the three-dimensional alignment.
Constant C_I, independent of the strain history	Eq. (A.48)	The strain independence of (A.48) is lost and the law is no longer rate-independent.

order the population and begins to disorder it, assumption A1 fails, and the monotone saturation predicted by (A.1) is not to be expected. This is not a hypothetical concern in the present geometry, whose defining feature is a ninety-degree bend, and it is the reason the law is traced in the beam, downstream of the bend, rather than across the whole of the mold.

The second is the semi-dilute requirement. The Folgar-Tucker diffusion represents fiber-fiber interactions as a hydrodynamic randomization, an idealization that holds while the fibers disturb one another through the fluid but do not touch. At the dosages employed in a real UHPC mix the suspension exceeds the concentration at which direct contact begins to contribute, and the alignment is limited not only by the randomizing agency represented in the model but also by the geometric obstruction of one fiber by another, a mechanism absent from (A.30). The dosages of the present experiment lie

below that threshold, at the lower edge of the semi-dilute range, so the model is applied within its domain; but it is also applied where the mechanism it represents is weakest, which is a limitation of the experiment rather than of the model, and it is discussed as such in the main text.

Chapter B

PIV and PTV Block-Structure Sensitivity Analysis

The two optical techniques process the recorded frames in fundamentally different ways, and the block structure that governs each follows from that difference. PIV correlates a *pair* of frames and returns the displacement of the tracer pattern between them, so what it must control is the separation between the two frames of a pair: too small and the displacement falls below the correlation's resolution, too large and the pattern decorrelates. PTV links a *sequence* of detections and follows each fiber from frame to frame, so what it must control is the interval between consecutive analyzed frames. Both intervals must widen as the flow decelerates, because a spacing that resolves the fast onset produces motions too small to measure once the material has nearly arrested, and this appendix reports the diagnostic that fixes where each widening is placed, demonstrates that the resulting fields do not depend on the particular choice, and states the block structures adopted.

The two techniques are therefore parameterized differently, and the distinction is carried through the whole appendix.

For PIV, each block is defined by two quantities. The *inter-pair skip* is the number of frames passed over between the two frames of a correlated pair, and it sets the time separation Δt of the pair, and hence the tracer displacement the correlation resolves; it is raised as the flow slows so that the tracers move a measurable distance between the two frames. The *final skip* is the number of frames passed over after a pair is closed and before the next pair begins, and it sets how often a velocity field is produced, that is, the temporal resolution of the series. Together with the two frames of the pair itself, these fix the block length, $\text{block} = \text{skip}_{\text{inter}} + \text{skip}_{\text{final}} + 2$, which is the spacing between the start of one pair and the start of the next, and which is chosen so that the time step block/fps stays approximately constant within each velocity regime.

For PTV, each block is defined by a single *skip*, the number of frames passed over

between two analyzed frames, since the tracking is continuous and Lagrangian and simply follows each fiber through the frames it retains.

B.1 PIV

The PIV diagnostic is the rate of spurious vectors returned by the Universal Outlier Detection test of Equation (3.2). When the inter-pair separation is too small for the flow, the tracer displacement approaches the sub-pixel noise, the correlation returns incoherent vectors, and the spurious rate rises; the appropriate inter-pair skip for a given phase of the filling is the one that minimizes that rate. Figures B.1 and B.2 show, for camera 1 of each matrix, the average velocity recovered under four inter-pair skips together with the corresponding spurious rate. The velocity fields converge as the flow slows, so the choice of skip does not alter the measured field once the flow is in the regime for which the skip is intended, and the spurious rate of each skip is lowest over a distinct interval, moving to later times as the skip grows. The block boundaries are placed where the spurious curve of one skip crosses below that of the next.

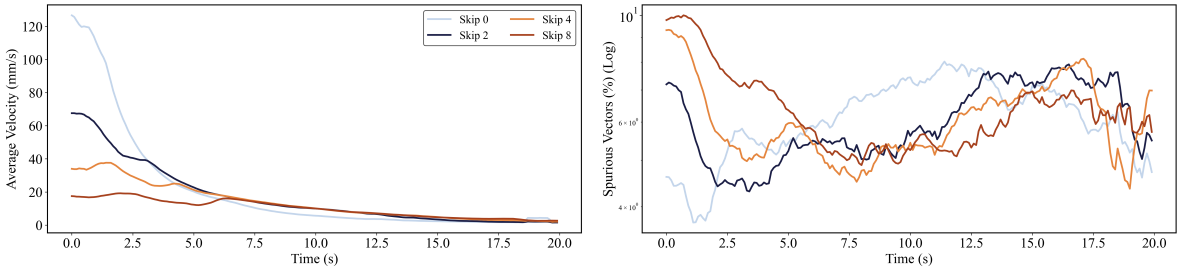


Figure B.1: PIV inter-pair skip sensitivity for Carbopol 0.2%, camera 1. Left: average velocity under four inter-pair skips, showing convergence as the flow stabilizes. Right: percentage of spurious vectors identified by the UOD test, on a logarithmic axis. Each skip minimizes the spurious rate over its own interval of the filling, which fixes the transition to the next.

The onset is the exception, and it is treated separately. There the velocity is highest, so decorrelation rather than under-resolution is the risk, and the smallest inter-pair skip is retained even where a larger one returns a marginally lower spurious rate, because it is the smallest separation that keeps the tracer pattern within the interrogation window between the two frames.

The block structures adopted for PIV are given in Table B.1. Cameras 1 to 3 record

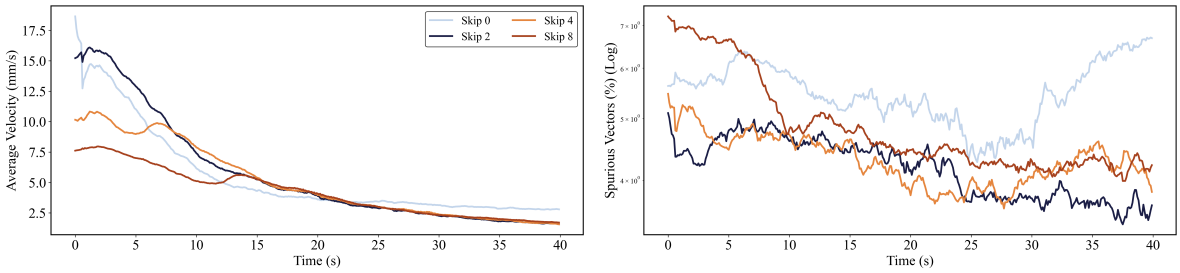


Figure B.2: PIV inter-pair skip sensitivity for Carbopol 0.5%, camera 1, over the longer window the stiffer matrix requires. Axes and conventions as in Figure B.1.

at 220 fps and camera 4, which resolves the fast reorientation at the outlet, at 660 fps. The near-constant time step of each regime, in the final column, is the quantity the block length is chosen to hold fixed.

The first block of camera 3 in the stiffer matrix carries no analysis: the beam it views receives no material for the first eight seconds, so the block is a placeholder spanning that dead interval, and its wide spacing reflects the absence of anything to measure rather than a resolved velocity field.

B.2 PTV

The Lagrangian analysis uses cameras 1 to 3, which view the L-shape and the beam; camera 4, dedicated to the outlet, does not enter the fiber tracking. The PTV diagnostic is the pair displacement Δx, the distance a fiber travels between two analyzed frames, against a floor of 0.25 mm below which the displacement can no longer be separated from the localization noise in the fiber position and the link becomes unreliable. Figures B.3 and B.4 show Δx under a sequence of skips: each skip holds the displacement above the floor over its own interval, and the point at which a given skip approaches the floor is the point at which the next, larger skip takes over.

The block structures adopted for PTV are given in Table B.2, together with the search radius $\Delta x_{\max}$ within which a fiber is sought in the next analyzed frame, tightened as the flow slows and the displacements shorten.

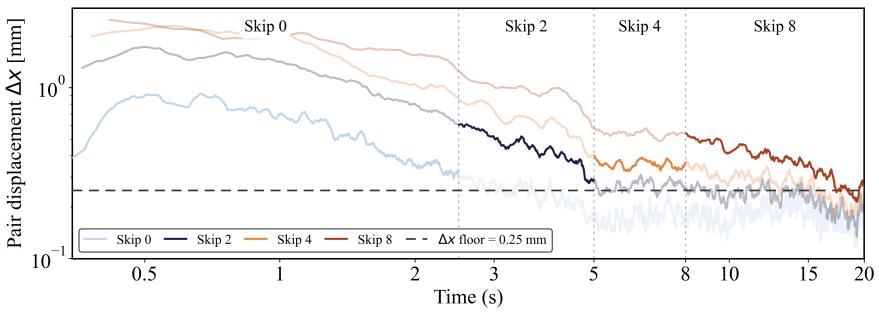


Figure B.3: PTV skip sensitivity for Carbopol 0.2%. Pair displacement Δx, the distance a fiber travels between analyzed frames, under four skips, against the floor of 0.25 mm (dashed). Each skip is drawn bold over the interval in which it is adopted; below the floor the displacement is indistinguishable from the localization noise.

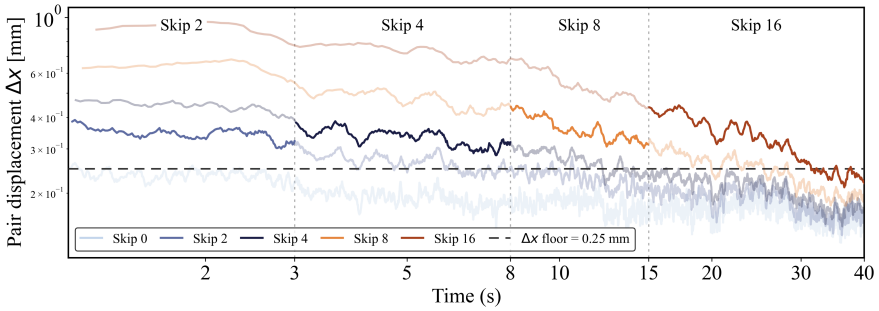


Figure B.4: PTV skip sensitivity for Carbopol 0.5%Pair displacement Δx, the distance a fiber travels between analyzed frames, under four skips, against the floor of 0.25 mm (dashed). Each skip is drawn bold over the interval in which it is adopted; below the floor the displacement is indistinguishable from the localization noise.

B.3 Computational Cost

The block structure reduces the computational load, but the quantity reduced differs between the two techniques, and the two are reported on their own terms rather than on a common one.

For PIV the cost is the number of correlated pairs, since each pair is one cross-correlation regardless of how the frames within it are spaced. A block of length $b = s_\mathrm{i} + s_\mathrm{f} + 2$ produces one pair, and the next pair begins b frames later, so a phase spanning F frames yields F/b pairs, against the F that a correlation of every consecutive pair would require. It is the final skip, not the inter-pair skip, that governs this saving: widening the separation within a pair changes the displacement resolved but not the number of pairs, whereas widening the gap between pairs is what reduces their count.

For PTV the cost is the number of frames analyzed, since the tracking processes each retained frame once. A block of skip s over F frames analyzes $F/(s+1)$ of them. Table B.3 reports both.

The PIV analysis processes about one pair in twenty of what an exhaustive correlation would, a reduction above 94% for every camera and both matrices, because even the fastest phase spaces its pairs widely and the long arrested phase spaces them wider still. The PTV reduction is smaller and more variable, between 72 and 93%, and it is smallest in the softer matrix, whose filling is brief and largely fast enough to demand the smallest skips, and largest in the stiffer matrix, whose long slow arrest is tracked at the widest spacing. In both techniques and both matrices the saving is what makes the full campaign tractable without sacrificing the temporal resolution the fast early phase requires.

Table B.1: PIV block structure by camera and matrix. Each block is defined by its inter-pair skip s_i (frames between the two frames of a correlated pair) and its final skip s_f (frames between one pair and the next); the block length $s_\mathrm{i} + s_\mathrm{f} + 2$ sets the time step $\Delta t = \text{block/fps}$. Cameras 1–3 record at 220 fps, camera 4 at 660 fps.

Matrix	**Cam**	t_0 [s]	t_1 [s]	$s_\mathbf{i}$	$s_\mathbf{f}$	**fps**	Δt [ms]
0.2%	1	0.0	6.5	0	9	220	50.0
		6.5	10.0	2	18	220	100.0
		10.0	20.0	8	12	220	100.0
	2	0.0	1.5	0	9	220	50.0
		1.5	3.0	2	18	220	100.0
		3.0	6.0	4	16	220	100.0
		6.0	20.0	8	12	220	100.0
	3	0.0	2.0	0	9	220	50.0
		2.0	5.0	2	18	220	100.0
		5.0	7.0	4	16	220	100.0
		7.0	20.0	8	12	220	100.0
	4	0.0	4.5	0	20	660	33.3
		4.5	8.0	2	18	660	33.3
		8.0	20.0	8	12	660	33.3
0.5%	1	0.0	0.2	0	9	220	50.0
		0.2	5.0	1	19	220	100.0
		5.0	8.0	2	18	220	100.0
		8.0	15.0	4	16	220	100.0
		15.0	40.0	8	34	220	200.0
	2	0.0	15.0	8	12	220	100.0
		15.0	40.0	16	4	220	100.0
	3	0.0	8.0	0	218	220	1000.0
		8.0	40.0	16	4	220	100.0
	4	0.0	1.5	0	9	660	16.7
		1.5	7.5	2	18	660	33.3
		7.5	20.0	4	16	660	33.3
		20.0	40.0	8	12	660	33.3

Table B.2: PTV block structure by camera and matrix, for the three cameras used in the fiber tracking. Each block is defined by its skip (frames between analyzed frames) and the maximum linking distance $\Delta x_{\max}$. All three cameras record at 220 fps.

Matrix	**Cam**	t_0 **[s]**	t_1 **[s]**	**Skip**	$\Delta x_{\max}$ **[mm]**
0.2%	1	0.0	2.5	0	3.0
		2.5	5.0	2	2.5
		5.0	8.0	4	2.0
		8.0	20.0	8	2.0
	2	0.0	3.0	0	3.0
		3.0	5.0	2	2.5
		5.0	8.0	4	2.0
		8.0	20.0	8	2.0
	3	0.0	0.5	10	3.0
		0.5	1.5	0	3.0
		1.5	3.0	2	2.5
		3.0	6.0	4	2.0
		6.0	20.0	8	2.0
0.5%	1	0.0	0.2	0	3.0
		0.2	3.0	2	2.5
		3.0	8.0	4	2.5
		8.0	15.0	8	2.5
		15.0	40.0	16	2.5
	2	0.0	0.2	0	3.0
		0.2	3.0	4	2.5
		3.0	15.0	8	2.5
		15.0	40.0	16	2.5
	3	0.0	10.0	10	3.0
		10.0	15.0	8	2.5
		15.0	40.0	16	2.5

Table B.3: Reduction in computational load relative to processing every frame: for PIV, the reduction in the number of correlated pairs; for PTV, the reduction in the number of frames analyzed. PTV uses cameras 1–3 only.

	PIV (pairs)		**PTV (frames)**	
Camera	0.2% [%]	0.5% [%]	0.2% [%]	0.5% [%]
1	94.0	96.9	73.7	89.0
2	95.1	95.5	72.0	91.1
3	95.0	96.3	81.5	92.7
4	95.5	95.3	–	–
Overall	**95.1**	**95.7**	**75.7**	**90.9**

Chapter C

Carbopol Mixtures

The yield stress of a Carbopol gel is not a property of its polymer content alone. The polymer is a polyacrylic acid, and it develops its network only once neutralized: the microgel particles swell as the carboxyl groups ionize, and the yield stress rises with them, reaching its maximum in a band around neutrality and falling away on either side of it (Piau, 2007). A batch left acidic is a weak gel and a batch overshot is a weak gel, and neither is the material the experiment intends to pour. The pH is therefore not a quality-control formality but the variable through which the rheology of Table 3.4 is actually set, and a drift in it across batches would be a drift in τ_0, which is the one parameter the two formulations were chosen to contrast.

This appendix reports the pH measured at each of the three stages of the protocol of Section 3.3, for every batch prepared during the campaign. Each value is the average of five points sampled across the mixture, so a batch that had gelled unevenly would be exposed by the spread among them as well as by the mean.

Table C.1 gives the three measurements for the sixteen batches, eight in each formulation, which is the number the campaign requires: four dosages, replicated twice, in each matrix.

Table C.1: Average pH measured during the three stages of the Carbopol preparation, for every batch of the campaign. pH_1 is measured after dispersion of the polymer and before neutralization; pH_2 after the first dose of 1M NaOH; pH_3 after the second dose, and is the value against which the target range of 6.8–7.2 is checked. Each entry is the average of five points sampled across the mixture.

Carbopol	**Mixture No.**	**Avg. pH_1**	**Avg. pH_2**	**Avg. pH_3**
0.2%	M70	3.56	6.89	7.07
	M71	3.59	6.43	6.80
	M73	3.57	6.76	6.98
	M74	3.56	6.75	6.90
	M75	3.57	6.79	6.93
	M77	3.67	6.68	6.81
	M80	3.59	6.93	7.08
	M82	3.53	6.60	6.90
Mean ± s.d.		3.58±0.04	6.73±0.16	6.93±0.11
0.5%	M83	3.37	6.39	7.10
	M86	3.27	6.30	7.01
	M87	3.24	6.37	6.99
	M91	3.28	6.33	6.90
	M92	3.24	6.23	6.96
	M93	3.33	6.28	6.99
	M94	3.28	6.28	6.91
	M95	3.31	6.30	6.95
Mean ± s.d.		3.29±0.04	6.31±0.05	6.98±0.06

Source: Research group.

Chapter D

Set-Up Information

This appendix collects the dimensional and geometric information of the experimental facility that the methodology of Chapter 3 refers to but does not reproduce. Three sets of measurements are given, and each bears on a different part of the analysis. The dimensions of the mold fix the geometry within which the flow develops and the spatial partition of Section 4.2 is defined; the positions of the cameras fix the scale factors of Table 3.1 and the co-registration of the fields recovered by the two techniques; and the renders of the assembled setup show the arrangement of the optical components that the text describes in words.

D.1 Domain Dimensions

Figure D.1 gives the dimensions of the L-shaped device and of the beam into which it discharges. The proportions of the two are what set the range of accumulated strain the casting can deliver: the material must traverse the whole of the L-shape before it reaches the beam, and it is the length of the beam that determines how far the flow can continue to shear the fibers before it arrests. The fill height of 22.4 mm within the L-shape, which fixes the number of fibers present in the measurement domain in Table 3.6, is measured from the floor of the device.

D.2 Camera Positions

Figure D.2 gives the positions and orientations of the four cameras relative to the mold. The distances shown are what determine the scale factors reported in Table 3.1, and through them the resolution at which a fiber is seen, from which the uncertainty of every quantity reported in this work descends. The rigidity of the mounting matters as much as the positions themselves: the analysis requires that a fiber tracked by the PTV be located within the velocity field measured by the PIV, and any relative motion

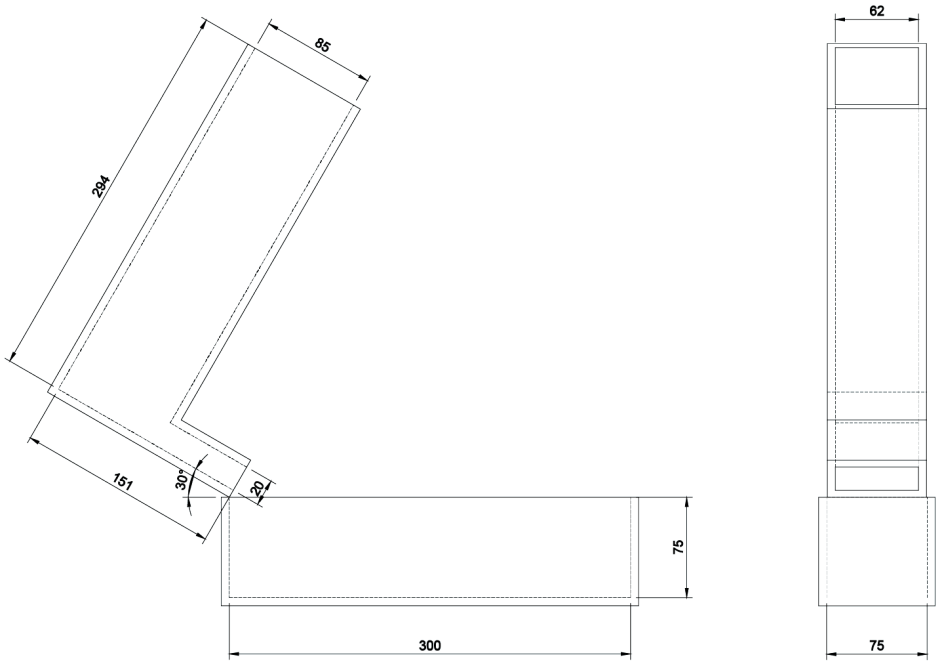


Figure D.1: Dimensions of the L-shaped device and of the beam. All measurements are in millimetres.

of the cameras between the two acquisitions of the same batch would place the fiber in the wrong part of the field.

D.3 Renders of the Experimental Setup

The following renders show the assembled facility from views that the single figure of Chapter 3 does not cover.

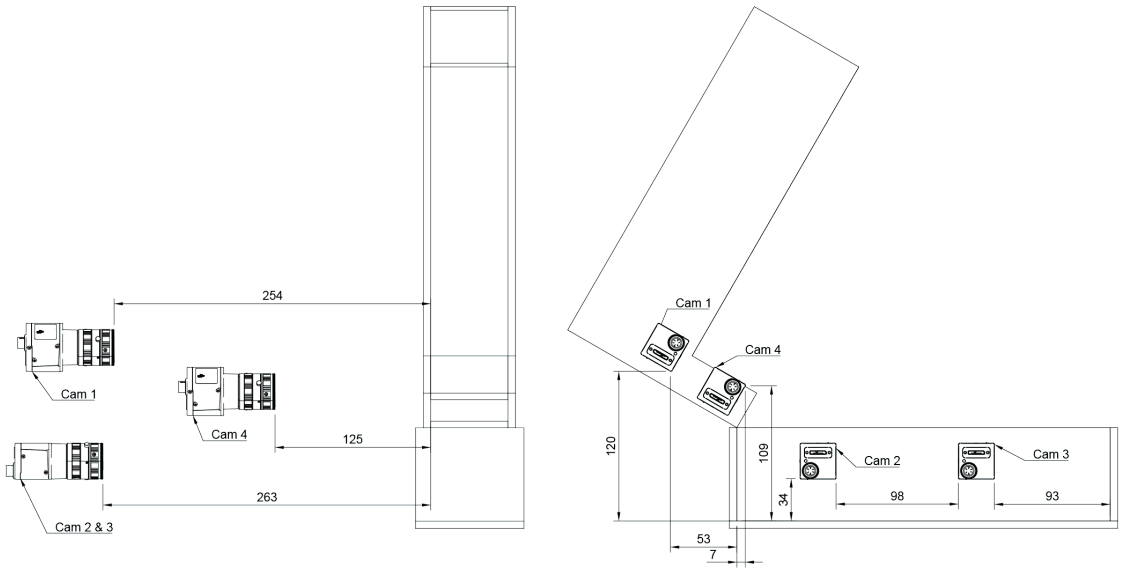


Figure D.2: Positions and viewing directions of the four cameras relative to the mold. Distances are in millimetres.

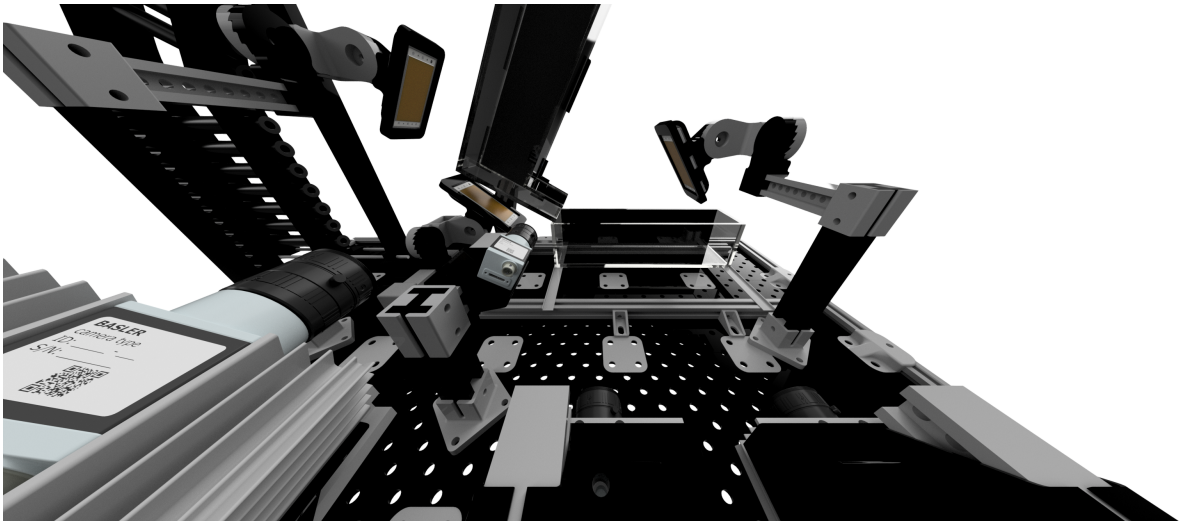


Figure D.3: Render of the assembled setup, showing the mold, the cameras, and the illumination mounted on the optical table.

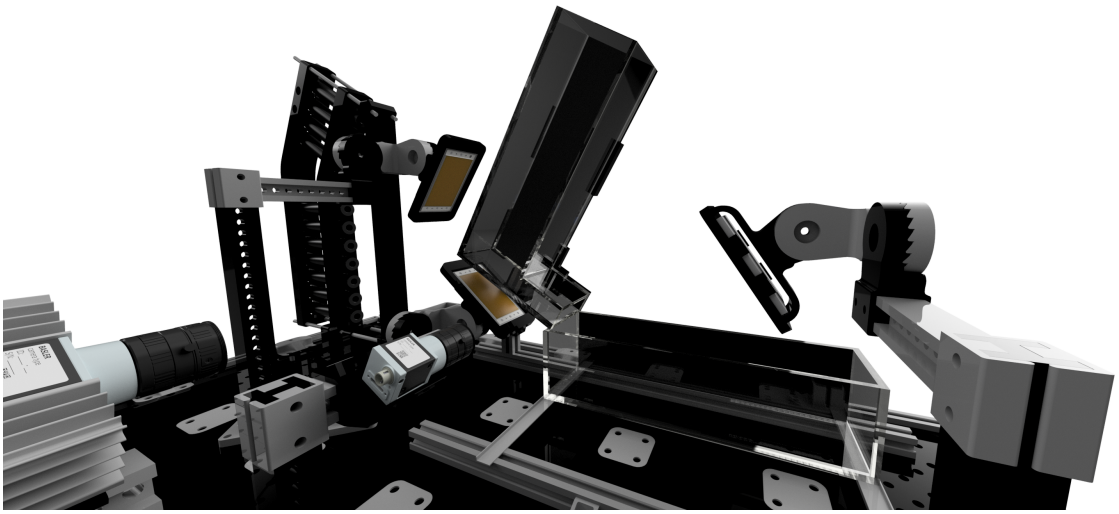


Figure D.4: Render of the assembled setup from a second viewpoint.

Chapter E

Zone-Resolved Statistical Analysis

Chapter 4 reports, for each variable and each flow stage, the largest effect the fiber content exerts anywhere in the domain, together with the zone in which it is found. This appendix gives the effect in every zone. These tables are the source from which the summary tables of the main text are drawn, collected here because the argument of that chapter rests on a claim about the *distribution* of the effect and not merely its peak: that the influence of the fiber content is regional, concentrated where the geometry forces a change in the flow, and that a domain-averaged measurement would report its absence. A table of maxima cannot test that claim; these can.

E.1 Protocol

Each cell of each table reports one Analysis of Covariance, run within a single zone, flow stage, and rheological matrix, with time as a continuous covariate and the fiber content as the factor of interest. The covariate is not optional: the filling is transient, so a zone passes through very different flow conditions as the mold fills, and comparing raw means across dosages would confound the effect of the fibers with the shared temporal trend of the flow. The two numbers reported are the partial eta squared of Equation (3.13), the share of variance the fiber content explains once that trend is removed, and the p-value of the same term.

The two answer different questions, and the tables mark them separately. An entry carries an asterisk when its effect size reaches the relevance threshold of $\eta^2 = 0.02$ fixed in Section 3.7, with further asterisks at 0.10 and 0.30, and is set in bold when that relevant effect is also significant at $p < 0.05$. Bold is the case in which the fiber content is reported as acting. The distinction matters because the samples run to thousands of observations, so the test can return a vanishingly small p for effects far too small to carry any physical consequence; such entries are read, correctly, as the absence of a relevant effect. Significance alone is never taken as evidence that the fibers act.

The zones are those of Figure 4.2, and they are not equally populated in the two matrices. Carbopol 0.5% arrests while the beam is still half empty, so the cells of its distal columns receive little or no material. The tables are therefore restricted, in each matrix, to the zones the flow actually reaches: a zone is retained when the carrier establishes a resolved velocity field within it and the final state deposits a countable population of fibers. The two criteria agree, and are applied identically to the Eulerian and Lagrangian tables, so both cover six zones in the stiffer matrix against eight in the softer one. The distal cells the stiff flow never fills would otherwise yield large but meaningless effect sizes from samples of a few dozen tracks, and it is precisely those cells that this restriction removes.

E.2 Eulerian Variables

Tables E.1 and E.2 report the effect of the fiber content on the speed $|\mathbf{u}|$ and the vorticity magnitude $|\zeta|$ of the carrier, measured by PIV.

Two features carry the argument of Section 4.3.2. The first is that the fiber content acts on the speed almost everywhere the carrier is moving: in Carbopol 0.5% the effect on $|\mathbf{u}|$ is relevant in sixteen of the eighteen zone-stage combinations, and it grows as the flow decelerates, reaching $\eta^2 = 0.525$ at the beam inlet in the quasi-steady stage. The fibers slow the carrier, and do so systematically. The second is that the vorticity does not follow: where the effect on the speed is strongest, in the stiff matrix, the effect on the vorticity never exceeds 0.160. That separation is the observation on which the decoupling of Section 4.3.5 rests.

E.3 Lagrangian Variables

Tables E.3 and E.4 report the effect of the fiber content on the fiber speed $|\mathbf{v}|$, the angular velocity $|\omega|$, and the orientation order η_1, measured by PTV, over the same zones the Eulerian analysis resolves.

The angular velocity is the variable these tables single out. It carries a relevant and significant effect in most zones and stages of both matrices, and its peaks are among the largest anywhere in the dataset, reaching $\eta^2 = 0.502$ in the vertical arm of the softer

Table E.1: Effect of fiber content N on the Eulerian variables, Carbopol 0.2%, by zone and flow stage. Each variable reports the partial eta squared and the covariate-adjusted p-value of the fiber-content term. Bold marks a relevant and significant effect ($\eta^2 \geq 0.02$ and $p < 0.05$); asterisks mark the effect-size bands alone. $^{*}\,\eta^2 \geq 0.02$; $^{**}\,\eta^2 \geq 0.10$; $^{***}\,\eta^2 \geq 0.30$.

Stage	**Zone**	$\lvert\mathbf{u}\rvert$		$\lvert\boldsymbol{\zeta}\rvert$	
		η^2	p	η^2	p
Onset	Z1	0.004	0.191	**0.283****	9×10^{-47}
	Z2	**0.049***	5×10^{-23}	0.015	6×10^{-6}
	Z3	0.015	0.010	0.001	0.717
	r1c1	**0.165****	4×10^{-26}	0.005	0.353
	r1c2	0.012	5×10^{-7}	**0.032***	0.044
	r2c1	**0.057***	4×10^{-16}	0.003	0.211
	r2c2	0.003	0.222	0.002	0.405
	r2c3	**0.022***	1×10^{-7}	0.015	1×10^{-5}
Transition	Z1	**0.057***	2×10^{-146}	**0.228****	9×10^{-42}
	Z2	**0.102****	3×10^{-146}	**0.036***	2×10^{-31}
	Z3	0.010	9×10^{-28}	0.004	<0.001
	r1c1	**0.021***	2×10^{-71}	0.004	3×10^{-6}
	r1c2	**0.022***	2×10^{-73}	0.003	0.040
	r2c1	**0.031***	2×10^{-82}	0.019	4×10^{-18}
	r2c2	0.017	6×10^{-38}	0.005	0.001
	r2c3	**0.045***	6×10^{-105}	**0.065***	2×10^{-32}
Quasi-steady	Z1	**0.021***	2×10^{-54}	**0.020***	2×10^{-64}
	Z2	**0.049***	5×10^{-143}	0.004	6×10^{-22}
	Z3	0.010	9×10^{-50}	0.000	0.424
	r1c1	0.020	1×10^{-117}	0.001	0.145
	r1c2	**0.023***	2×10^{-20}	**0.042***	6×10^{-41}
	r2c1	0.014	1×10^{-83}	0.003	1×10^{-9}
	r2c2	0.015	3×10^{-71}	0.002	0.026
	r2c3	**0.048***	1×10^{-28}	**0.061***	8×10^{-44}

matrix during the transition and $\eta^2 = 0.430$ in the first beam column of the stiffer one once the flow has settled. The orientation order climbs higher still, to $\eta^2 = 0.713$ at the beam inlet of the softer matrix in the quasi-steady stage, the single largest response of the whole dataset. Read against the Eulerian tables, where the vorticity of the carrier responds far more weakly to how many fibers it carries, this is the decoupling of Section 4.3.5: the fibers turn faster as their number grows, and the fluid carrying them does so far less.

Table E.2: Effect of fiber content N on the Eulerian variables, Carbopol 0.5%, by zone and flow stage. The stiffer matrix arrests before filling the third column of the beam, so its distal cells receive no material and are absent from the table. Conventions as in Table E.1.

Stage	**Zone**	$\lvert\mathbf{u}\rvert$		$\lvert\boldsymbol{\zeta}\rvert$	
		η^2	p	η^2	p
Onset	Z1	**0.070***	2×10^{-21}	**0.084***	2×10^{-25}
	Z2	**0.292****	4×10^{-245}	0.009	1×10^{-7}
	Z3	**0.097***	3×10^{-72}	**0.094***	2×10^{-83}
	r1c1	**0.037***	3×10^{-35}	0.005	4×10^{-8}
	r2c1	**0.027***	7×10^{-25}	0.005	2×10^{-9}
	r2c2	**0.050***	2×10^{-27}	**0.052***	7×10^{-25}
Transition	Z1	**0.116****	2×10^{-100}	**0.047***	2×10^{-74}
	Z2	**0.434*****	$<10^{-300}$	**0.062***	4×10^{-64}
	Z3	**0.044***	2×10^{-146}	**0.095***	7×10^{-66}
	r1c1	**0.131****	$<10^{-300}$	0.005	4×10^{-13}
	r2c1	**0.070***	$<10^{-300}$	**0.034***	1×10^{-37}
	r2c2	0.017	6×10^{-91}	**0.030***	1×10^{-34}
Quasi-steady	Z1	0.013	1×10^{-90}	**0.077***	$<10^{-300}$
	Z2	**0.150****	$<10^{-300}$	0.011	3×10^{-75}
	Z3	**0.024***	6×10^{-319}	**0.070***	3×10^{-286}
	r1c1	**0.525*****	1×10^{-75}	**0.098***	2×10^{-10}
	r2c1	**0.142****	3×10^{-269}	**0.108****	2×10^{-125}
	r2c2	**0.067***	6×10^{-104}	**0.160****	8×10^{-231}

E.4 Reading the Distribution

Two properties of the whole set of tables bear on the interpretation of the main text, and neither is visible in a table of maxima.

The first is that the effect is regional. In every matrix and for every variable, zones in which the fiber content acts strongly sit alongside zones in which it does not act at all, and the strong ones are those where the geometry forces a change in the flow: the elbow, where the material turns; the beam inlet, where it lands; and the vertical arm, where the descent first develops. This is why the effect nearly vanishes when the domain is pooled, and why the pooled figure is the less informative of the two. Pooling places the variance between a plug descending the arm and a nearly arrested beam cell into the error term of Equation (3.13), and that variance dominates any effect the fibers could exert. The pooled value therefore establishes only that the fiber content is small relative to the spatial heterogeneity of the flow, which holds for any flow through this geometry and is not a statement about the fibers.

Table E.3: Effect of fiber content N on the Lagrangian variables, Carbopol 0.2%, by zone and flow stage. The table covers the same zones as the Eulerian analysis of Table E.1; the distal cell the flow reaches only in sparse transient passages, unable to support the test, is omitted. Conventions as in Table E.1.

Stage	**Zone**	$\lvert\mathbf{v}\rvert$		$\lvert\boldsymbol{\omega}\rvert$		$\boldsymbol{\eta_1}$	
		η^2	p	η^2	p	η^2	p
Onset	Z1	**0.043***	6×10^{-9}	**0.426*****	3×10^{-102}	**0.180****	7×10^{-43}
	Z2	0.018	4×10^{-6}	**0.143****	1×10^{-52}	**0.284****	2×10^{-101}
	Z3	**0.037***	3×10^{-12}	**0.093***	4×10^{-31}	**0.140****	5×10^{-49}
	r1c1	**0.072***	1×10^{-26}	0.003	0.191	0.015	0.002
	r1c2	**0.146****	4×10^{-65}	0.019	0.168	**0.173****	0.009
	r2c1	0.020	2×10^{-6}	**0.086***	1×10^{-23}	**0.076***	1×10^{-19}
	r2c2	**0.165****	2×10^{-92}	**0.183****	1×10^{-88}	**0.046***	7×10^{-20}
	r2c3	0.020	2×10^{-8}	**0.026***	4×10^{-5}	**0.234****	3×10^{-25}
Transition	Z1	**0.279****	2×10^{-249}	**0.502*****	$<10^{-300}$	**0.141****	2×10^{-106}
	Z2	**0.145****	4×10^{-207}	**0.178****	1×10^{-211}	**0.082***	2×10^{-48}
	Z3	**0.044***	6×10^{-45}	0.007	8×10^{-8}	0.006	<0.001
	r1c1	**0.093***	1×10^{-145}	**0.090***	9×10^{-131}	0.002	0.076
	r1c2	0.017	6×10^{-12}	0.015	5×10^{-9}	0.007	0.002
	r2c1	**0.221****	$<10^{-300}$	**0.201****	3×10^{-274}	**0.161****	1×10^{-113}
	r2c2	**0.086***	7×10^{-113}	**0.089***	8×10^{-102}	0.015	4×10^{-10}
	r2c3	**0.075***	3×10^{-33}	**0.226****	5×10^{-80}	**0.057***	2×10^{-17}
Quasi-steady	Z1	**0.031***	6×10^{-43}	**0.054***	9×10^{-77}	**0.069***	7×10^{-57}
	Z2	**0.020***	1×10^{-34}	**0.054***	7×10^{-96}	0.015	6×10^{-15}
	Z3	0.012	5×10^{-19}	**0.031***	2×10^{-26}	**0.053***	2×10^{-37}
	r1c1	**0.068***	3×10^{-82}	**0.239****	2×10^{-247}	**0.278****	8×10^{-176}
	r1c2	**0.133****	4×10^{-78}	**0.203****	1×10^{-117}	**0.105****	1×10^{-53}
	r2c1	**0.046***	5×10^{-69}	**0.071***	8×10^{-107}	**0.713*****	$<10^{-300}$
	r2c2	**0.173****	6×10^{-302}	**0.358*****	$<10^{-300}$	**0.359*****	$<10^{-300}$
	r2c3	**0.396*****	3×10^{-136}	**0.323*****	4×10^{-97}	**0.187****	1×10^{-50}

The second is that the effect broadly strengthens as the flow decelerates. Reading down the three stages of a column, the entries tend to rise from the onset toward the quasi-steady stage, and the strongest responses of the dataset are recorded in these later stages, once the flow has settled or is settling into arrest. This is what would be expected if the fibers act by adding hydrodynamic resistance to a carrier already near its yield threshold, since the same increment of effective consistency matters more where the material is barely moving than where it descends under gravity, and it is one reason the analysis of the fiber final state in Section 4.4 is conducted on the arrested configuration rather than on the flow that produced it.

Table E.4: Effect of fiber content N on the Lagrangian variables, Carbopol 0.5%, by zone and flow stage. As in the Eulerian analysis, the distal columns the arresting flow never fills are absent from the table; the six zones listed are those in which the carrier establishes a resolved field and a countable fiber population is tracked. Conventions as in Table E.1.

Stage	**Zone**	$\lvert\mathbf{v}\rvert$		$\lvert\boldsymbol{\omega}\rvert$		$\boldsymbol{\eta_1}$	
		η^2	p	η^2	p	η^2	p
Onset	Z1	$\mathbf{0.161}^{**}$	8×10^{-68}	$\mathbf{0.333}^{***}$	5×10^{-134}	$\mathbf{0.407}^{***}$	1×10^{-120}
	Z2	$\mathbf{0.050}^{*}$	1×10^{-61}	$\mathbf{0.219}^{**}$	8×10^{-262}	$\mathbf{0.186}^{**}$	1×10^{-113}
	Z3	0.018	2×10^{-11}	$\mathbf{0.110}^{**}$	3×10^{-95}	$\mathbf{0.095}^{*}$	2×10^{-56}
	r1c1	0.003	0.024	0.015	1×10^{-6}	$\mathbf{0.039}^{*}$	7×10^{-16}
	r2c1	0.000	0.973	0.009	3×10^{-5}	$\mathbf{0.048}^{*}$	3×10^{-26}
	r2c2	$\mathbf{0.029}^{*}$	4×10^{-11}	$\mathbf{0.025}^{*}$	4×10^{-5}	0.010	0.094
Transition	Z1	$\mathbf{0.091}^{*}$	2×10^{-123}	$\mathbf{0.327}^{***}$	$<10^{-300}$	$\mathbf{0.354}^{***}$	7×10^{-232}
	Z2	0.005	0.009	$\mathbf{0.126}^{**}$	7×10^{-56}	$\mathbf{0.378}^{***}$	2×10^{-123}
	Z3	0.004	0.033	$\mathbf{0.025}^{*}$	9×10^{-10}	$\mathbf{0.105}^{**}$	2×10^{-32}
	r1c1	$\mathbf{0.030}^{*}$	1×10^{-12}	$\mathbf{0.230}^{**}$	2×10^{-93}	$\mathbf{0.109}^{**}$	2×10^{-40}
	r2c1	$\mathbf{0.057}^{*}$	1×10^{-18}	$\mathbf{0.138}^{**}$	8×10^{-39}	$\mathbf{0.523}^{***}$	2×10^{-165}
	r2c2	0.007	0.008	0.018	2×10^{-5}	0.009	0.006
Quasi-steady	Z1	$\mathbf{0.083}^{*}$	6×10^{-102}	$\mathbf{0.285}^{**}$	1×10^{-273}	0.012	1×10^{-6}
	Z2	$\mathbf{0.022}^{*}$	2×10^{-17}	$\mathbf{0.234}^{**}$	7×10^{-161}	$\mathbf{0.384}^{***}$	3×10^{-237}
	Z3	0.000	0.559	$\mathbf{0.060}^{*}$	3×10^{-31}	0.000	0.694
	r1c1	$\mathbf{0.280}^{**}$	3×10^{-15}	$\mathbf{0.247}^{**}$	2×10^{-13}	$\mathbf{0.306}^{***}$	3×10^{-16}
	r2c1	$\mathbf{0.368}^{***}$	1×10^{-95}	$\mathbf{0.430}^{***}$	2×10^{-121}	$\mathbf{0.340}^{***}$	1×10^{-84}
	r2c2	$\mathbf{0.029}^{*}$	1×10^{-5}	$\mathbf{0.026}^{*}$	5×10^{-5}	$\mathbf{0.052}^{*}$	1×10^{-10}

Chapter F

Final State of the Fiber Population

Section 4.4 reports which factors govern the state in which the fibers come to rest, through effect sizes aggregated over the domain. This appendix gives the underlying measurements: the descriptors of the deposited population in every beam cell, at every dosage and in both matrices; the tests that isolate the effect of the rheology from that of the fiber content; and the photographs of the arrested elements from which the descriptors are computed.

Three quantities describe the final state, defined in Section 2.3.3. The orientation order $\hat{\eta}_1$ measures how aligned the fibers of a cell are, from zero at complete disorder to one at perfect alignment; it is reported debiased, because the number of tracked fibers varies across cells and the raw order parameter carries a finite-sample bias that falls as that number rises. The mean director $\bar{\theta}$ gives the direction of that alignment, reported alongside the order because a cell may be strongly ordered about a direction other than the one the flow would impose. The dispersion index Λ compares the radius of gyration of the fiber centroids with that of a uniform arrangement over the same cell, so that $\Lambda < 1$ denotes clustering, $\Lambda = 1$ uniformity, and $\Lambda > 1$ an arrangement more spread out than uniform.

The dosages are given throughout by their nominal fiber count, 750, 1500 and 3000; the corresponding counts within the measurement domain are 547, 1094 and 2188, from Table 3.6.

F.1 Descriptors of the Deposited Population

Table F.1 gives the three descriptors in every populated cell. Two features of it govern the reading of everything that follows.

The first is the asymmetry between the matrices. Carbopol 0.2% populates five cells of the beam and reaches the far column, while Carbopol 0.5% populates three, all clustered near the entry. The stiffer gel arrests while the beam is still half empty, which is why

its rows are fewer and why the accumulated-strain analysis of Chapter 6 is conducted on the softer matrix alone.

The second is the tracked count n. It falls as low as 19 in the upper row and reaches 239 in the lower one, so the descriptors are not computed on samples of comparable size, and the cells holding few fibers carry correspondingly wide uncertainty. This is why $\hat{\eta}_1$ is debiased rather than raw, and why the entries with the smallest n are not read as measurements in their own right.

Table F.1: Final-state descriptors of the deposited fiber population in each beam cell, by matrix and nominal fiber content. n: tracked fibers in the cell; $\hat{\eta}_1$: debiased orientation order; $\bar{\theta}$: mean director; R_g: radius of gyration of the centroids, against a uniform reference of $R_g^{\text{unif}} = 30.8$ mm; $\Lambda = R_g/R_g^{\text{unif}}$: spatial dispersion index, with $\Lambda < 1$ denoting clustering.

Matrix	**Cell**	N	n	$\hat{\eta}_1$	$\bar{\theta}$ [°]	Λ
0.2%	r1c1	750	97	0.501	161	0.845
	r1c1	1500	115	0.454	169	0.853
	r1c1	3000	120	0.357	166	0.923
	r1c2	750	19	0.759	175	0.700
	r1c2	1500	37	0.598	178	0.827
	r1c2	3000	31	0.493	175	0.701
	r2c1	750	162	0.524	173	0.999
	r2c1	1500	199	0.310	4	0.968
	r2c1	3000	149	0.323	170	1.026
	r2c2	750	202	0.685	1	1.031
	r2c2	1500	239	0.550	5	1.022
	r2c2	3000	200	0.496	6	1.006
	r2c3	750	89	0.637	1	0.476
	r2c3	1500	106	0.524	176	0.480
	r2c3	3000	67	0.697	2	0.498
0.5%	r1c1	750	44	0.147	90	0.479
	r1c1	1500	72	0.116	125	0.519
	r1c1	3000	63	0.421	92	0.532
	r2c1	750	133	0.479	23	0.906
	r2c1	1500	160	0.256	26	0.911
	r2c1	3000	102	0.192	27	0.888
	r2c2	750	51	0.368	26	0.626
	r2c2	1500	34	0.219	11	0.821
	r2c2	3000	29	0.226	11	0.831

F.1.1 The Orientation

Within Carbopol 0.2% the order falls monotonically with the dosage in three of the five cells, r1c1, r1c2 and r2c2, and the fall is substantial: at the beam inlet it runs from

0.501 at the lowest loading to 0.357 at the highest, and in the cell above it from 0.759 to 0.493. This is the trend the saturation law predicts, and the one Chapter 6 recovers when the same orientations are read against the strain that produced them rather than the cell in which they were measured.

Two cells do not follow it. r2c1 falls from 0.524 to 0.310 and rises marginally to 0.323, and r2c3 falls from 0.637 to 0.524 and rises to 0.697, so the richest mix carries the highest order recorded in the far column. Both hold the fewest tracked fibers at the highest dosage, 149 and 67, and r2c3 is the most distal cell, reached last and by the fewest fibers. The departure is therefore attributed to sampling rather than physics, and r2c3 is the entry the fit of Chapter 6 identifies as the least reliable of its arrest points.

The director is nearly uniform across the softer matrix, within a few degrees of the horizontal in every cell, the values near 180° and near 0° describing the same direction since a rod has no distinguished end. The fibers of the softer matrix align with the flow along the beam, whatever the degree of that alignment. The stiffer matrix does not: its directors sit near 25° in the lower row and near 90° at the inlet, so the fibers arriving into the stiff gel arrest inclined to the flow, and at the inlet transverse to it. This is a difference in the direction of the order, not its degree, and a further sign that the material entering the stiff beam does not flow along it but lands and stops.

F.1.2 The Spatial Arrangement

The dispersion separates the two matrices more sharply than the orientation does. Carbopol 0.2% reaches Λ at or slightly above unity in the lower row, r2c1 and r2c2, so its fibers there are arranged about as uniformly as a random placement would arrange them, and marginally more spread out than that in r2c2. Carbopol 0.5% stays below unity in every populated cell, reaching 0.479 at the inlet, so its fibers arrive clustered and stay clustered.

The contrast follows from how each matrix arrests. The stiffer gel stops flowing while still advancing and locks the fibers wherever they happen to be, preserving whatever clustering they carried out of the L-shape. The softer gel keeps flowing for longer, so the fibers are carried and rearranged more completely before the material sets, and at the inlet the widening flow spreads them slightly past uniform, pushing Λ above one.

That value is not an artifact: it indicates fibers sitting closer to the cell boundaries than a uniform set would.

One cell of the softer matrix runs against this. r2c3 holds $\Lambda \approx 0.48$ at every dosage, the most strongly clustered arrangement in the dataset, and it is the far end of the beam, where the flow arrives last and arrests soonest after arriving. The fibers reaching it are those the flow carried furthest, and they arrest before the material has time to redistribute them.

F.2 The Effect of the Rheology

Table F.2 compares the two matrices at each zone and dosage, testing three quantities separately: the difference in orientation order, by permutation; the difference in dispersion, by Mann-Whitney, as a rank-biserial correlation; and the difference in director, by the Watson-Williams test, the circular analogue of a comparison of means.

Table F.2: Effect of the rheological matrix on the final state, by zone and nominal fiber content. $\Delta\hat{\eta}_1 = \hat{\eta}_1^{0.2} - \hat{\eta}_1^{0.5}$, with its permutation p-value; r_Λ is the rank-biserial effect on the dispersion, with its Mann-Whitney p-value; p_{WW} is the Watson-Williams p-value for the difference in director. Bold marks $p < 0.05$.

Zone	N	$\hat{\eta}_1^{0.2}$	$\hat{\eta}_1^{0.5}$	$\Delta\hat{\eta}_1$ (p)	r_Λ (p)	p_{WW}
Z1	750	0.770	0.493	0.276 (0.153)	**0.581** (0.004)	**0.012**
	1500	0.412	0.457	0.045 (0.827)	0.085 (0.629)	0.182
	3000	0.747	0.443	0.304 (0.067)	**0.372** (0.018)	0.842
Z2	750	0.705	0.677	0.028 (0.821)	0.142 (0.282)	4×10^{-4}
	1500	0.531	0.564	0.033 (0.784)	0.021 (0.886)	0.138
	3000	0.142	0.455	**0.313** (0.036)	0.259 (0.137)	0.506
Z3	750	0.722	0.677	0.046 (0.748)	0.077 (0.701)	0.134
	1500	0.792	0.750	0.042 (0.715)	0.127 (0.478)	0.810
	3000	0.730	0.760	0.030 (0.790)	0.085 (0.632)	0.142
r1c1	750	0.501	0.147	**0.354** (1×10^{-3})	**0.321** (0.002)	8×10^{-7}
	1500	0.454	0.116	**0.338** (1×10^{-3})	**0.281** (0.001)	5×10^{-5}
	3000	0.357	0.421	0.063 (0.531)	0.018 (0.843)	2×10^{-19}
r2c1	750	0.524	0.479	0.044 (0.485)	0.095 (0.159)	2×10^{-12}
	1500	0.310	0.256	0.054 (0.455)	0.092 (0.136)	6×10^{-5}
	3000	0.323	0.192	0.131 (0.122)	0.105 (0.158)	8×10^{-7}
r2c2	750	0.685	0.368	**0.317** (1×10^{-3})	**0.361** (7×10^{-5})	1×10^{-5}
	1500	0.550	0.219	**0.331** (0.002)	**0.357** (8×10^{-4})	0.513
	3000	0.496	0.226	**0.270** (0.030)	**0.263** (0.022)	0.640

The three quantities do not respond alike.

The director is the quantity the rheology separates most consistently. It differs significantly between matrices in every beam cell and at nearly every dosage, with the smallest p-values in the table, reaching 2×10^{-19} at the inlet. This is the result already read from the descriptors: the fibers of the softer matrix arrest along the beam and those of the stiffer one inclined to it, and the difference is not marginal.

The order separates the matrices in the beam but not in the L-shape. Within the L-shape the difference is not significant, except at the highest dosage in the elbow, and at the outlet Z3 the two are indistinguishable at every loading. Downstream, in r1c1 and r2c2, the softer matrix carries an order some 0.3 higher than the stiffer one, significant at nearly every dosage. What separates the matrices is therefore not the state in which the fibers leave the L-shape but the state they reach once inside the beam, consistent with the softer matrix continuing to flow and order its fibers while the stiffer one arrests.

The dispersion follows the order. It is separated significantly at the inlet and in the second cell of the lower row, and is indistinguishable at the outlet Z3, so the clustering the two matrices show in the beam is acquired within it and not carried in.

F.3 The Effect of the Fiber Content on the Spatial Arrangement

Table F.3 tests the dispersion of the centroids against the two factors separately: the fiber content within each matrix, and the rheology at each dosage.

The rheology governs the dispersion throughout the beam. Its effect is significant at the inlet and in both cells of the lower row, at nearly every dosage, and absent in the L-shape, reinforcing the preceding section: the clustering is acquired in the beam, not carried into it.

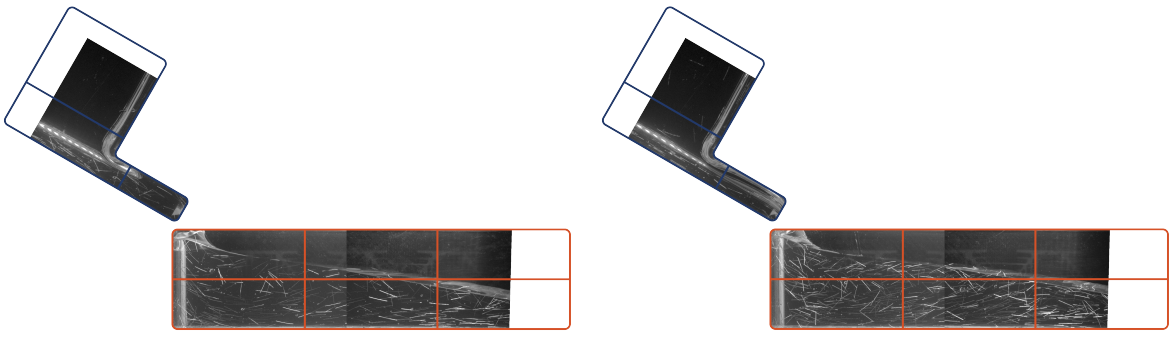
The fiber content, by contrast, is inert in most of the domain. It carries no significant effect in the elbow, the outlet, or any beam cell but one, and its effect sizes there are indistinguishable from zero. Two zones depart from this, and must be stated rather than folded into the general conclusion.

The first is the vertical arm, Z1, where the fiber content of the softer matrix carries an effect of $\eta^2 = 0.302$ on the dispersion, with $p = 0.003$. This is a large effect, larger than

any the fiber content exerts on the orientation anywhere in the dataset, and it is not compatible with the arrangement of the centroids being a matter of rheology alone. Z1 is also where the rheological effect reverses sign with the dosage, from $r = +0.469$ at the lowest loading to $r = -0.438$ at the highest, which is an interaction between the two factors rather than an effect of either. The arm is where the material descends before it has begun to spread, and where the fibers are still packed as they were introduced, so it is the zone in which the initial distribution is least modified by the flow. Whether the effect there describes the deposition or the introduction is not settled by these data. The second is the beam inlet, r1c1, where the effect of the fiber content is significant in the softer matrix but of size $\eta^2 = 0.013$, below the relevance threshold of Section 3.7. It is reported and not interpreted.

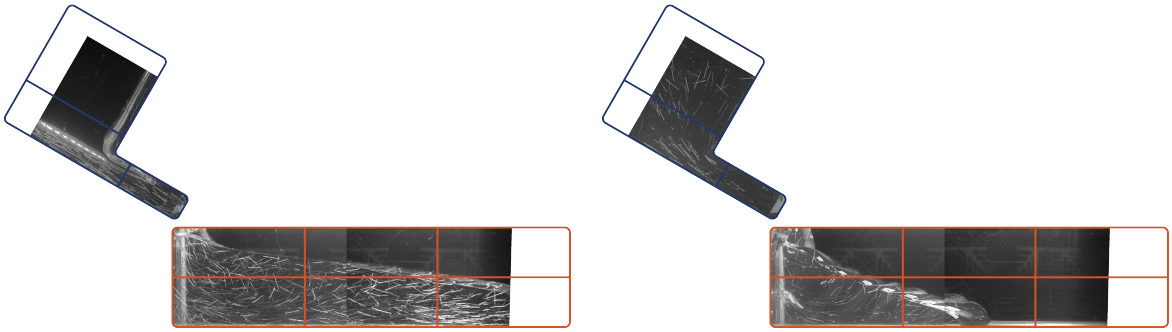
F.4 Photographs of the Arrested Elements

The photographs below show the beam once the material has come to rest, for each dosage and both matrices. They are the images from which the final-state descriptors of Table F.1 are computed, and they carry directly the two contrasts the tables report: the softer matrix reaches the far end of the beam while the stiffer one does not, and the fibers of the softer matrix lie along the beam while those of the stiffer one lie inclined to it.



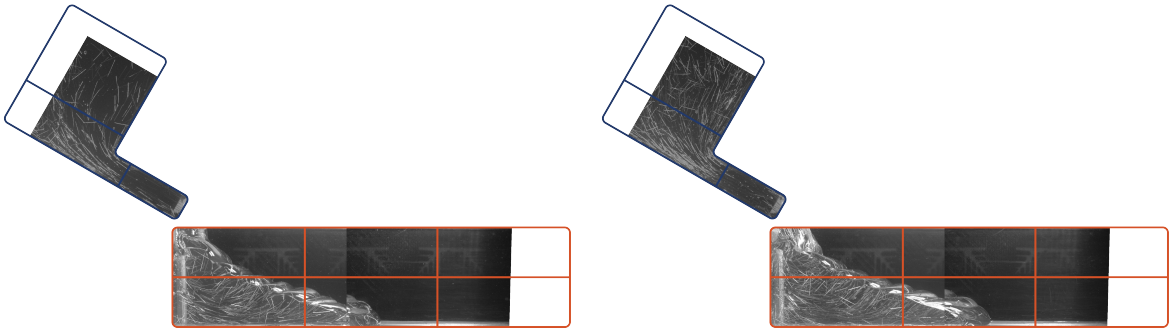
(a) (b)

Figure F.1: Arrested state of the beam in Carbopol 0.2%. (a) Run M74, 750 fibers. (b) Run M73, 1500 fibers.



(a) (b)

Figure F.2: (a) Run M70, 3000 fibers in Carbopol 0.2%, which completes the series of the softer matrix. (b) Run M86, 750 fibers in Carbopol 0.5%, which begins the series of the stiffer one. The two are shown together so that the contrast in filling extent is immediate: the softer matrix reaches the far end of the beam at its highest dosage, while the stiffer one does not reach it at its lowest.



(a) (b)

Figure F.3: Arrested state of the beam in Carbopol 0.5%. (a) Run M94, 1500 fibers. (b) Run M83, 3000 fibers. The material arrests before filling the third column of the beam at every dosage of this matrix.

Table F.3: Tests on the final centroid arrangement, by zone. The effect of the fiber content is tested within each matrix and reported as η^2; the effect of the rheology is tested at each nominal dosage and reported as a rank-biserial correlation r. Bold marks $p < 0.05$. Cells absent from a matrix are omitted.

Zone	**Comparison**	**Test**	Effect	p
Z1	N (Carbopol 0.2%)	One-way ANOVA	**0.302**	0.003
	N (Carbopol 0.5%)	Kruskal-Wallis	−0.003	0.635
	Rheology ($N = 750$)	Mann-Whitney	**0.469**	0.021
	Rheology ($N = 1500$)	Mann-Whitney	0.062	0.727
	Rheology ($N = 3000$)	Mann-Whitney	−0.438	0.005
Z2	N (Carbopol 0.2%)	Kruskal-Wallis	−0.032	0.879
	N (Carbopol 0.5%)	Kruskal-Wallis	−0.001	0.411
	Rheology ($N = 750$)	Mann-Whitney	0.073	0.579
	Rheology ($N = 1500$)	Mann-Whitney	0.084	0.554
	Rheology ($N = 3000$)	Mann-Whitney	0.135	0.439
Z3	N (Carbopol 0.2%)	Kruskal-Wallis	−0.001	0.378
	N (Carbopol 0.5%)	Kruskal-Wallis	−0.031	0.704
	Rheology ($N = 750$)	Mann-Whitney	0.194	0.323
	Rheology ($N = 1500$)	Mann-Whitney	−0.107	0.551
	Rheology ($N = 3000$)	Mann-Whitney	−0.091	0.608
r1c1	N (Carbopol 0.2%)	Kruskal-Wallis	**0.013**	0.046
	N (Carbopol 0.5%)	Kruskal-Wallis	−0.011	0.987
	Rheology ($N = 750$)	Mann-Whitney	−0.326	0.002
	Rheology ($N = 1500$)	Mann-Whitney	−0.376	2×10^{-5}
	Rheology ($N = 3000$)	Mann-Whitney	−0.533	3×10^{-9}
r1c2	N (Carbopol 0.2%)	Kruskal-Wallis	−0.015	0.688
r2c1	N (Carbopol 0.2%)	Kruskal-Wallis	−0.001	0.447
	N (Carbopol 0.5%)	Kruskal-Wallis	−0.003	0.623
	Rheology ($N = 750$)	Mann-Whitney	−0.138	0.041
	Rheology ($N = 1500$)	Mann-Whitney	−0.093	0.131
	Rheology ($N = 3000$)	Mann-Whitney	−0.225	0.003
r2c2	N (Carbopol 0.2%)	Kruskal-Wallis	−0.002	0.753
	N (Carbopol 0.5%)	Kruskal-Wallis	0.009	0.222
	Rheology ($N = 750$)	Mann-Whitney	−0.499	4×10^{-8}
	Rheology ($N = 1500$)	Mann-Whitney	−0.353	9×10^{-4}
	Rheology ($N = 3000$)	Mann-Whitney	−0.305	0.008
r2c3	N (Carbopol 0.2%)	Kruskal-Wallis	−0.006	0.791

Chapter G
Design

Every figure in this work reports one of two rheological matrices, and most report three fiber dosages within them, so the colour of a curve is not decoration but the primary means by which the reader identifies what it describes. The palette is therefore fixed once, in Figure G.1, and applied without exception across the whole document.

The rheology sets the hue. Carbopol 0.2% is always drawn in blues and Carbopol 0.5% always in oranges, and no other pairing occurs anywhere in this work, so the matrix a curve belongs to is legible at a glance and without recourse to the legend. The two hues sit opposite one another on the colour wheel, which is what allows a figure to superimpose the two matrices without their curves being confused, and the choice is not arbitrary in the other direction either: the comparison the figures most often ask the reader to make is between matrices, and the strongest visual contrast is spent on the strongest distinction.

The dosage sets the tone. Where three fiber contents are reported within a matrix, they are drawn light, medium and dark within its hue, and the ordering is monotonic, the darkest tone always denoting the highest dosage. The reader can therefore read a trend with concentration from the shading alone, without consulting the legend for each curve, and a figure carrying six curves resolves into two families of three rather than into six unrelated lines. The tonal separation is set wide enough that the three levels remain distinguishable when a figure is printed in greyscale, which the hue distinction alone would not survive.

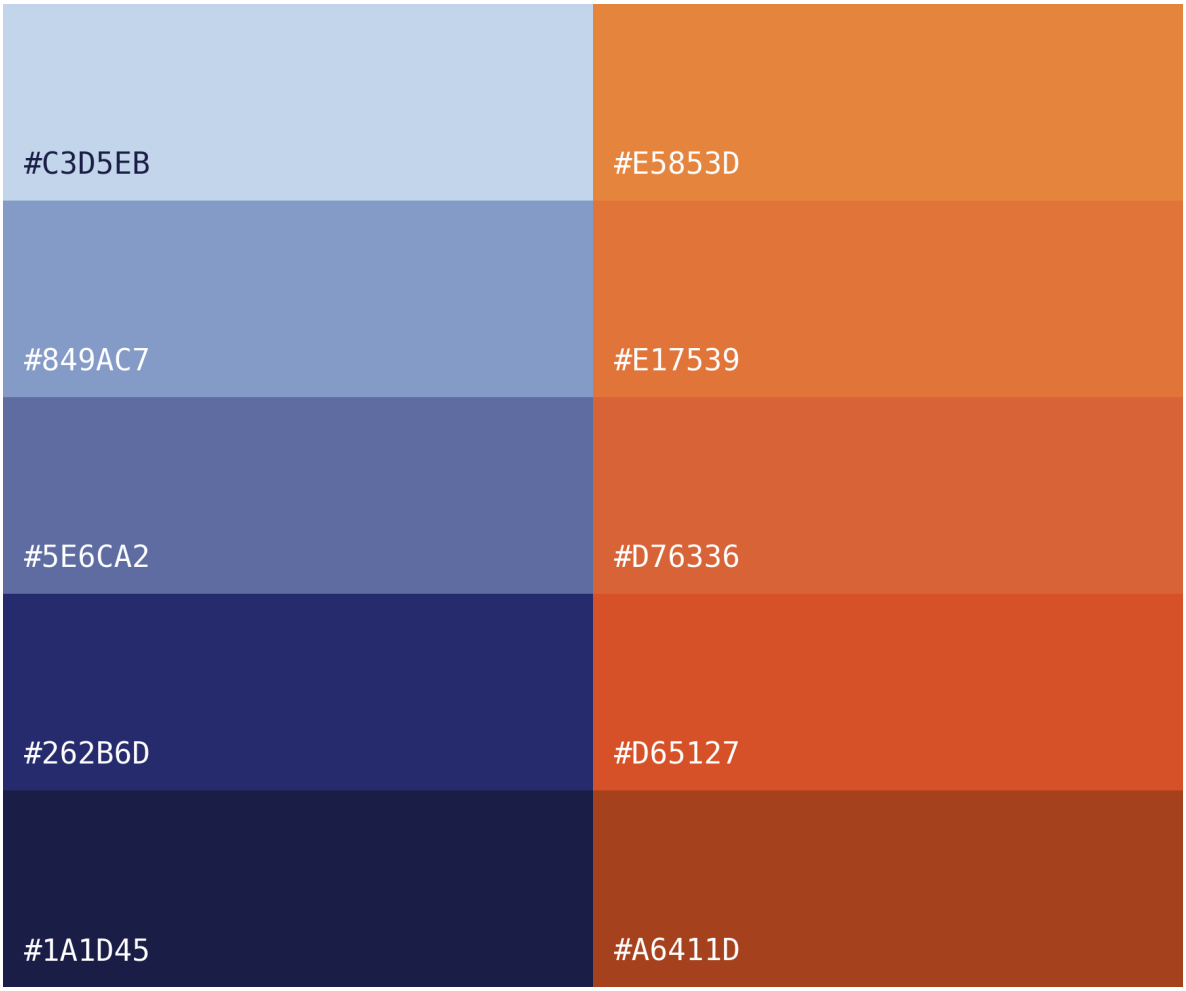


Figure G.1: Colour palette used throughout this work. Carbopol 0.2% is drawn in blues and Carbopol 0.5% in oranges; within each matrix, the three fiber dosages are drawn in a light, medium and dark tone, the darkest denoting the highest fiber content.

