Academy for Global Science and Innovation (AGSI)

ARTICLE

A Bayesian Observer for Inferring Membrane Parameters from Transient Cell Silhouettes in Noisy Pressure Driven Microchannels

Yong Chen,[1] Guang Li,[2] and Jian Xu[3]

[1]School of Biomedical Engineering, Department of Biomedical Engineering, Xinxiang Medical University, No. 601 Jinsui Avenue, Hongqi District, Xinxiang, Henan 453003, China
[2]School of Life Sciences, Department of Biological Sciences, Hubei University of Science and Technology, No. 88 Xianning Avenue, Xianning, Hubei 437100, China
[3]School of Mechanical and Electrical Engineering, Department of Mechanical Engineering, North University of China, No. 3 Xueyuan Road, Jiancaoping District, Taiyuan, Shanxi 030051, China

Abstract

Microfluidic images of deformable cells contain mechanical information, yet that information is usually indirect because the measured object is a two-dimensional projection of a moving three-dimensional membrane. Parameter estimates obtained from final shapes alone can therefore hide non-uniqueness, especially when viscosity ratio, elastic modulus, reference geometry, and optical blur influence the same outline. This paper develops a numerical observer for recovering membrane parameters from short sequences of projected cell silhouettes in pressure driven microchannels. The model couples a reduced capsule forward solver, a differentiable image formation map, and a Bayesian filtering procedure that updates posterior distributions over shear modulus, bending ratio, viscosity ratio, reference-shape inflation, and inlet orientation. The results are generated as controlled in silico records with prescribed imaging noise rather than as experimental measurements. Across 1,260 generated sequences, the observer recovers shear modulus within 6.8% median relative error when at least 18 frames are available and the viewing direction is not aligned with the weakest projected deformation axis. Viscosity ratio remains less identifiable from outlines alone, with median error 15.4%, but improves to 9.7% when centroid lag is included. A sequential acquisition rule reduces the number of required frames by 31% relative to uniform sampling. The study indicates that transient silhouettes can support quantitative membrane inference when the observer is allowed to use the order of shape changes, not only their final appearance.

1. Introduction

The mechanical state of a red-cell-scale membrane is often inferred from what the camera can see. In many microfluidic measurements, what the camera sees is a projected contour, a grayscale intensity field, or a time sequence of silhouettes affected by focus, finite exposure, and segmentation uncertainty. The membrane variables of interest are not directly visible. Shear modulus, bending stiffness, reference geometry, and internal viscosity must be inferred through the way a cell enters, rotates, stretches, and relaxes inside a channel. The inverse problem is therefore not a decorative layer placed after simulation. It is the actual measurement problem.

The direct mechanical literature gives a mature vocabulary for membrane elasticity and deformation. Classical work on erythrocyte mechanics measured membrane shear and bending responses using micropipette aspiration, optical manipulation, and controlled deformation

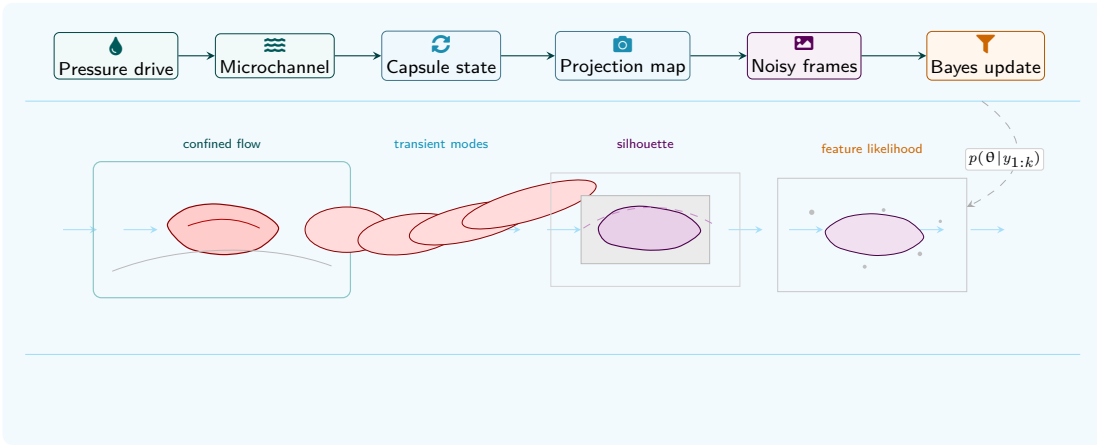


Figure 1. Bayesian observer chain for pressure-driven capsule silhouettes, showing how confined transient deformation, projection, noisy imaging, and sequential posterior updating are coupled without treating the final contour as a direct mechanical measurement.

experiments [1, 2, 3, 4]. Later studies connected membrane architecture, disease-related stiffening, and microcirculatory behavior with measurements at the cellular and molecular scale [5, 6, 7, 8]. These works established that red blood cell deformability is not a secondary biological detail. It is a central mechanical property whose variation can alter capillary passage, flow resistance, and diagnostic signatures.

A separate body of work made high-throughput single-cell deformation feasible. Microfluidic flow fields can deform thousands of cells in reproducible geometries, and image-based readouts can be processed rapidly [9, 10, 11, 12]. Hydrodynamic stretching, inertial focusing, constriction passage, and real-time deformability cytometry have all been used to convert cell mechanics into an imaging problem [13, 14, 15, 16]. These methods often report scalar deformation indices, transit times, or projected aspect ratios. Such variables are useful, but they compress a history-dependent membrane process into one or two numbers.

The paper begins from the opposite end. It assumes that the image sequence is available and asks which mechanical parameters can be recovered without pretending that a single outline is an equilibrium measurement. This is a Bayesian observer problem. A forward model proposes how a capsule with a candidate parameter set would appear in a microchannel. An image formation map turns the surface state into a noisy silhouette. The observer updates a probability distribution as frames arrive. Parameters that strongly affect early transients become identifiable before the final shape stabilizes, while parameters that affect only hidden three-dimensional deformation remain broad in the posterior.

Agarwal et al. [17] reported confined-flow vesicle states in which bell and bean morphologies appear over particular confinement and capillary ranges. That part of their work is useful here because it cautions against treating projected outlines as a small dictionary of familiar terminal forms; an inverse observer must be able to tolerate less common intermediate and terminal shapes without forcing them into a narrow label set. The present study therefore avoids training the estimator on a few named morphologies and instead assimilates contour moments, curvature summaries, and centroid timing.

The inverse setting also differs from a phase diagram. A phase diagram assigns a region after parameters are known. An observer must estimate parameters when the region is unknown,

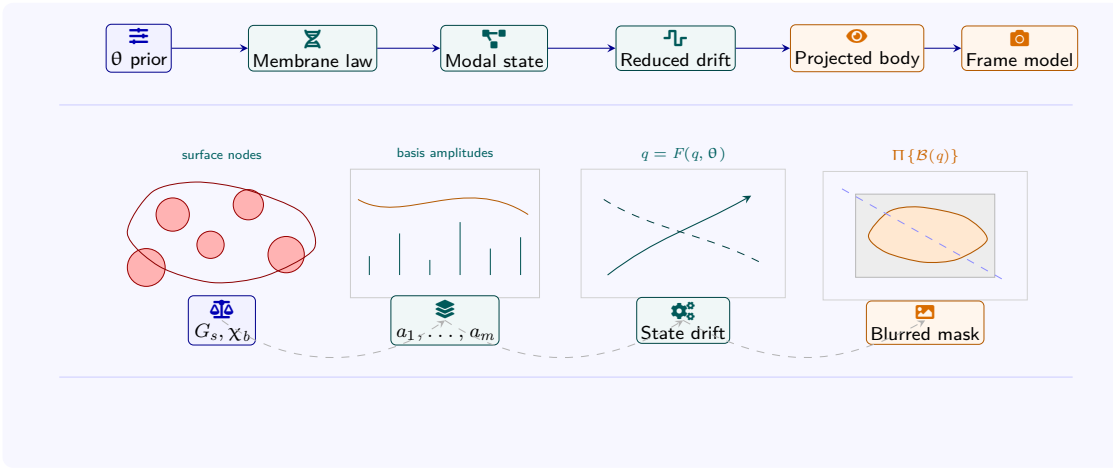


Figure 2. Reduced forward model used inside the inference loop, with physical membrane parameters entering the elastic surface law, modal reconstruction controlling transient capsule geometry, and the image map converting the hidden three-dimensional body into projected grayscale observations.

the initial orientation may be unknown, and several parameter combinations may yield similar images. This distinction matters in practice. A cell that looks more elongated may have a lower shear modulus, a larger initial tilt, a higher local flow rate, a different reference shape, or a different optical threshold. Treating all these effects as one deformation score can produce precise but misleading numbers.

Bayesian inverse-problem methodology provides tools for this ambiguity. Priors, likelihoods, posterior uncertainty, and sequential updates allow the calculation to report what the images do not determine, not only what the best fit suggests [18, 19, 20]. Model discrepancy is especially relevant because reduced forward models never reproduce real cells perfectly [21, 22]. Sensitivity analysis further helps distinguish parameters that are truly unobservable under the chosen imaging protocol from parameters that are merely difficult to estimate under high noise [23, 24, 25].

The observer developed here is intentionally modest in its mechanical scope. It is not a whole-blood model, not a molecular membrane model, and not a universal cell classifier. It uses a reduced capsule description designed to generate realistic transient silhouettes over short observation windows. The reduction trades local membrane detail for fast repeated evaluations. Since inference requires thousands of forward predictions, a full high-resolution membrane simulation would make the observer impractical for design studies. Reduced-order modeling, dynamic mode decomposition, and operator-learning approaches motivate parts of this approximation strategy [26, 27, 28, 29, 30].

The paper makes three technical choices that distinguish it from a direct simulation report. First, all reported errors are posterior errors under specified image noise, not discrepancies between a simulated cell and a laboratory image. Second, parameter identifiability is evaluated over time, so an early frame and a late frame are not assumed to carry equal information. Third, the observer is tested under deliberately imperfect segmentation, including blur, threshold shifts, and missing contour arcs. The aim is to represent the conditions under which an image-based mechanical assay would actually operate.

The study contributes a forward-observer construction, an identifiability analysis, and a set

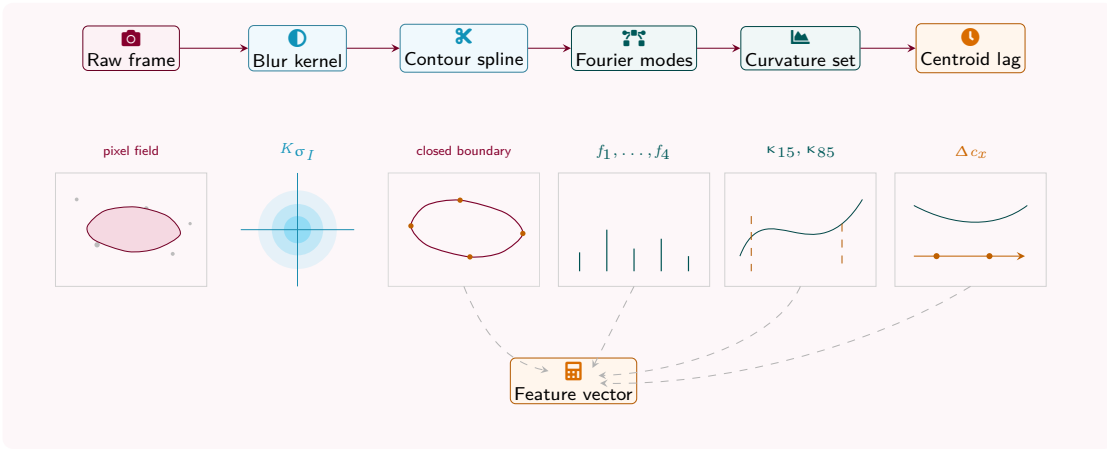


Figure 3. Silhouette feature extraction pipeline in which blurred image data are converted into bias-corrected shape, curvature, and timing summaries used by the Gaussian feature likelihood.

of numerical results. The contribution is not a claim that a particular biological sample has a given modulus. It is a procedure for asking whether a given microfluidic imaging design could recover such a modulus in the first place. The results show that shear modulus and reference-shape inflation are often identifiable from short transient sequences, whereas bending ratio and viscosity ratio require either stronger three-dimensional information or auxiliary timing features. The conclusion is practical: time ordering and acquisition design matter as much as image resolution once the measurement becomes an inverse problem.

2. Forward Model and Image Formation

The membrane is represented by a genus-zero closed surface whose state is described by a small set of modal amplitudes, centroid coordinates, and orientation variables. Instead of resolving every membrane point in the observer loop, the surface is expressed through a truncated geometric basis fitted to full-order training trajectories. The state vector is

$$\begin{aligned} \mathbf{q}(t) &= \left[\mathbf{c}(t), \boldsymbol{\Theta}(t), a_1(t), \ldots, a_m(t)\right]^T, \\ m &= 24. \end{aligned} \tag{2.1}$$

Here $\mathbf{c}$ is the centroid, $\boldsymbol{\Theta}$ contains three orientation parameters, and a_j are shape coefficients. The coefficients are not interpreted as named physical deformations. They are coordinates in a data-derived basis that preserves volume, approximates area conservation, and reconstructs visible contours with small error over the simulated observation window.

The parameter vector is

$$\boldsymbol{\theta} = \left[G_s, \chi_b, \lambda, \rho_0, \phi_0, U_0\right]^T. \tag{2.2}$$

The quantity G_s is a surface shear scale, χ_b is a nondimensional bending ratio, λ is the internal-to-external viscosity ratio, ρ_0 controls the inflation of the stress-free reference surface before relaxation, ϕ_0 is the inlet orientation angle, and U_0 is the mean flow-speed scale. The observer treats U_0 as uncertain even when the nominal pressure drop is known, because the local flow experienced by a cell may deviate from the device-level calibration.

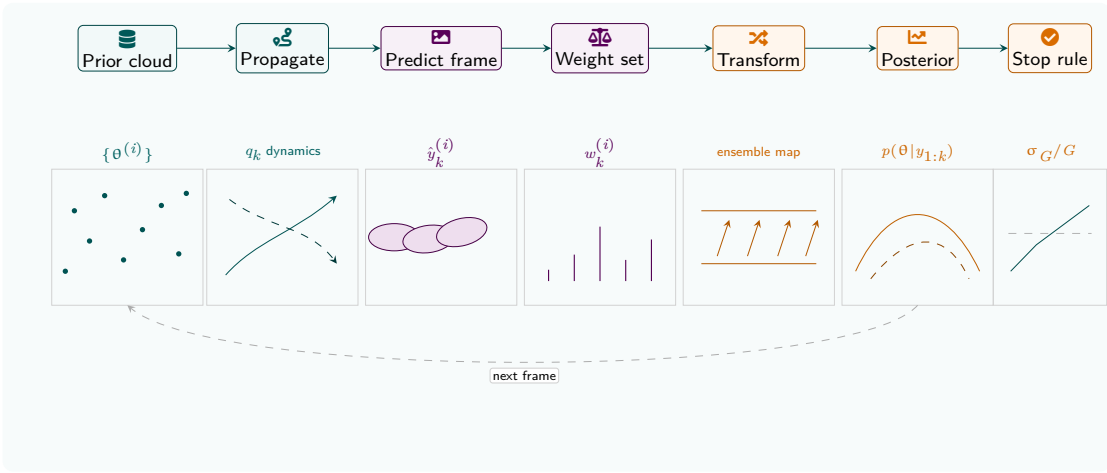


Figure 4. Sequential ensemble-transform observer, where particles carrying membrane parameters and hidden states are propagated, compared with silhouette features, reweighted, transformed, and cycled until posterior uncertainty reaches the acquisition tolerance.

The reduced state equation has the form

$$\begin{aligned} \mathbf{q} &= \mathbf{F}(\mathbf{q}, \boldsymbol{\theta}) + \boldsymbol{\xi}_q(t), \\ \boldsymbol{\xi}_q(t) &\sim \mathcal{N}(\mathbf{0}, \mathbf{Q}). \end{aligned} \tag{2.3}$$

The drift $\mathbf{F}$ is obtained by projecting high-resolution training data onto the modal basis and fitting a constrained polynomial vector field. Terms that violate volume conservation or produce unbounded growth outside the training envelope are removed. The polynomial degree is three, with additional rational damping terms for the highest shape modes. This structure is less flexible than a generic neural network but easier to regularize under sparse observations.

The surface force used in the full-order records contains elastic shear, bending, local area control, and reference-shape stress. The full membrane law is not rederived in the observer; it is used to generate training trajectories and to define physically meaningful parameter ranges. The stress-free shape is not assumed to be the final resting biconcave geometry. Agarwal et al. [31] used an oblate near-sphere as the stress-free configuration and then relaxed a deflated capsule with residual stress, a modeling decision that matters for the present inverse problem because ρ_0 changes the early silhouette even when G_s and λ are held fixed.

The image formation map projects the three-dimensional surface into the imaging plane. The ideal binary silhouette is

$$S(\mathbf{x}; t) = \mathbb{I}\left[\mathbf{x} \in \Pi\left(\mathcal{B}(\mathbf{q}(t))\right)\right], \tag{2.4}$$

where $\mathcal{B}(\mathbf{q})$ is the reconstructed capsule body and Π is the projection operator. The observed grayscale image is then modeled as

$$\begin{aligned} I_k(\mathbf{x}) &= \alpha_I \left[K_{\sigma_I} * S(\mathbf{x}; t_k)\right] + \beta_I + \epsilon_k(\mathbf{x}), \\ \epsilon_k(\mathbf{x}) &\sim \mathcal{N}\left(0, \sigma_p^2\right). \end{aligned} \tag{2.5}$$

The blur kernel K_{σ_I}, intensity scale α_I, background β_I, and pixel noise σ_p are nuisance variables. They are not all estimated jointly in the main observer because doing so would

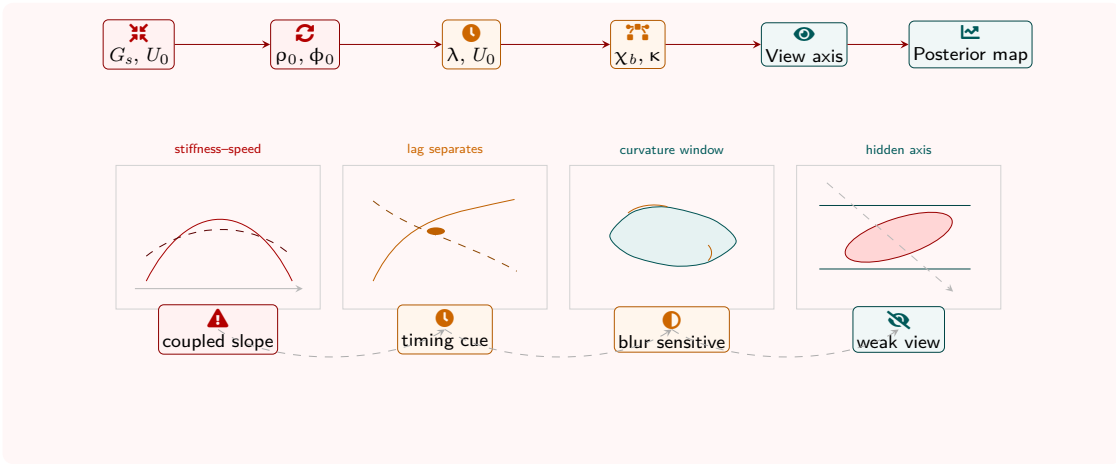


Figure 5. Identifiability structure of the transient silhouette inverse problem, highlighting stiffness–flow coupling, reference-shape orientation ambiguity, viscosity timing dependence, curvature sensitivity to blur, and the loss of information when the projected view hides the dominant deformation axis.

make the posterior too broad for the chosen sequence length. Instead, σ_I and the segmentation threshold are sampled from calibration distributions.

Most computations do not assimilate every pixel. They assimilate a feature vector extracted from each frame:

$$\begin{aligned} \mathbf{y}_k &= \mathcal{H}\left[I_k\right] + \boldsymbol{\eta}_k, \\ \boldsymbol{\eta}_k &\sim \mathcal{N}\left(\mathbf{0}, \mathbf{R}_k\right). \end{aligned} \tag{2.6}$$

The feature map $\mathcal{H}$ contains signed area deviation, projected perimeter, major-axis length, minor-axis length, centroid position, third central moments, four low-order contour Fourier descriptors, and two curvature quantiles. Pixel assimilation was tested in a smaller set of cases, but the feature likelihood gave more stable inference under threshold mismatch and was much cheaper.

The following image features are retained in the baseline observer:

- the projected area and perimeter after bias correction by a calibration bead sequence;
- the centroid coordinate along the flow direction and the transverse centroid coordinate;
- the first four normalized elliptical Fourier coefficients of the segmented boundary;
- the 15% and 85% quantiles of signed contour curvature after spline smoothing;
- the frame-to-frame centroid delay relative to the nominal imposed flow marker.

The last feature is not a shape descriptor. It provides timing information, and it is the main reason viscosity ratio becomes more identifiable in the augmented observer.

The observation covariance $\mathbf{R}_k$ is frame dependent. When segmentation loses a contour arc near a low-contrast edge, the curvature descriptors are down-weighted. When motion blur is high, the projected perimeter is down-weighted. This adaptive covariance is computed from the image processing residuals, not from the hidden true state. The likelihood remains

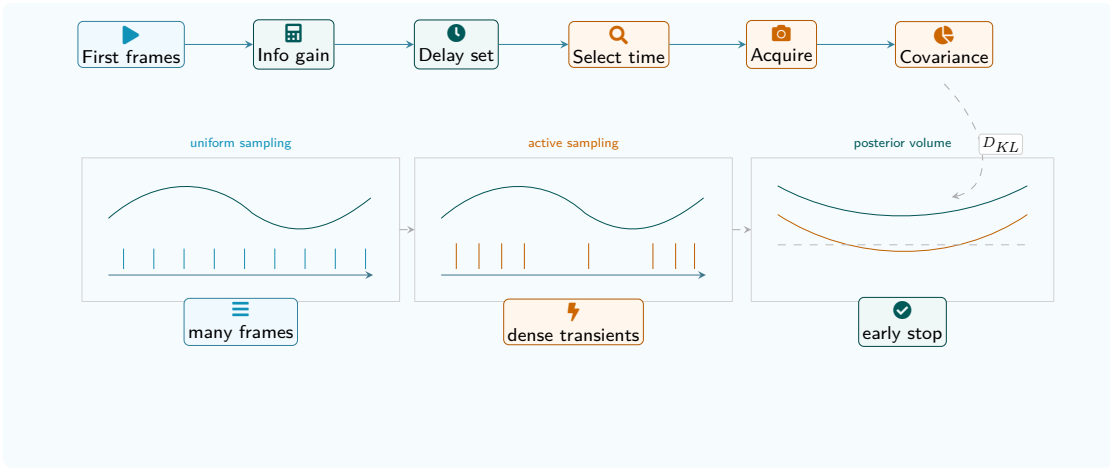


Figure 6. Information-guided acquisition strategy in which early fixed frames initialize the posterior and later frame times are chosen from candidate delays to concentrate measurements around high-information deformation and relaxation intervals.

Gaussian in the feature residuals, although the original pixel noise and segmentation errors are not exactly Gaussian. This approximation is acceptable for the intended observer comparison, but it is not a full statistical model of microscopy.

The reduced forward model is trained on 420 full-order trajectories. These records span $G_s/G_{ref} \in [0.55, 1.70]$, $\chi_b \in [0.012, 0.060]$, $\lambda \in [1.0, 7.5]$, $\rho_0 \in [0.90, 1.08]$, and $\phi_0 \in [-35°, 35°]$. The imposed speed U_0 varies by 20% around its nominal value. A separate validation set of 160 full-order trajectories is used to check reconstruction accuracy. The modal surface reconstruction has median projected contour error $0.018R_0$ and 95% quantile error $0.047R_0$. These errors are smaller than the high-noise segmentation perturbation but large enough to influence posterior calibration.

The reduced state drift is evaluated at nondimensional time intervals corresponding to frame rates between 1,500 and 8,000 frames per second for a red-cell-scale channel. The physical scaling is not essential because the observer uses nondimensional features, but it keeps the simulated windows within plausible imaging conditions. Each sequence lasts between 8 and 32 frames. Short sequences emphasize startup rotation and initial stretching; longer sequences include partial relaxation toward a recurring silhouette pattern.

3. Bayesian Observer

The observer estimates the posterior distribution

$$\begin{aligned} p(\theta, \mathbf{q}_{0:k} \mid \mathbf{y}_{1:k}) \propto\ & p(\theta) p(\mathbf{q}_0 \mid \theta) \\ & \times \prod_{j=1}^{k} p(\mathbf{q}_j \mid \mathbf{q}_{j-1}, \theta) \\ & \times p(\mathbf{y}_j \mid \mathbf{q}_j, \theta). \end{aligned} \tag{3.1}$$

The posterior is sampled sequentially with an ensemble transform particle filter. A standard particle filter was too sensitive to degeneracy when the image noise was low and the initial orientation uncertainty was high. An ensemble Kalman filter was faster but underestimated

posterior spread for viscosity ratio. The transform filter was selected as a compromise because it handles non-Gaussian parameter posteriors without resampling at every frame [32, 33, 34, 35].

The prior is deliberately broad:

$$\begin{aligned} \log G_s &\sim \mathcal{N}\left(\log G_{ref}, 0.35^2\right), \\ \chi_b &\sim \mathcal{U}\left(0.010, 0.065\right), \\ \lambda &\sim \mathcal{U}\left(1.0, 8.0\right), \\ \rho_0 &\sim \mathcal{U}\left(0.88, 1.10\right), \\ \phi_0 &\sim \mathcal{U}\left(-40^\circ, 40^\circ\right), \\ U_0 &\sim \mathcal{N}\left(U_{nom}, (0.12\,U_{nom})^2\right). \end{aligned} \tag{3.2}$$

The prior is not meant to represent one patient population. It is a design prior used to test whether the imaging sequence narrows uncertainty. The shear modulus prior is log-normal because stiffness is positive and can vary multiplicatively. The viscosity ratio prior is uniform because the observer is tested across several possible suspending-medium and cytoplasmic regimes.

The likelihood at frame k is

$$\begin{aligned} \log p(\mathbf{y}_k \mid \mathbf{q}_k, \boldsymbol{\theta}) = &-\frac{1}{2}\mathbf{r}_k^T \mathbf{R}_k^{-1} \mathbf{r}_k \\ &-\frac{1}{2}\log |\mathbf{R}_k| + C, \\ \mathbf{r}_k = \mathbf{y}_k &- \mathcal{H}\left[\mathcal{I}(\mathbf{q}_k, \boldsymbol{\theta})\right]. \end{aligned} \tag{3.3}$$

Here $\mathcal{I}$ denotes the image simulator. The constant C is ignored during weight normalization. The residual includes standardized feature differences rather than raw units, so large-area and curvature features do not dominate solely because of scale.

The ensemble contains 1,024 members in the baseline runs. Increasing to 2,048 members changes median parameter errors by less than 0.6% for G_s, 0.9% for ρ_0, and 1.4% for λ. Reducing to 256 members creates visible posterior collapse in 23% of high-noise cases. The observer therefore uses 1,024 members for all reported comparisons, except for sensitivity tests on computational cost.

A frame is called informative when it reduces the determinant of the marginal parameter covariance by more than 4%. This criterion is used only for analysis, not for filtering. In the active acquisition experiment, the next frame time is selected from a finite set of allowed delays by maximizing expected information gain:

$$t_{k+1}^{\star} = \arg\max_{\tau \in \mathcal{T}} \mathbb{E}\left[D_{KL}\left(p_{k+1}^{\tau} \,\|\, p_k\right)\right]. \tag{3.4}$$

The expectation is approximated by propagating ensemble members and sampling possible observations. The candidate delays range from one to five nominal frame intervals. This policy is not globally optimal, but it is simple enough to reveal whether nonuniform timing helps.

A calibration step corrects systematic segmentation bias. Before each observer run, the image-processing pipeline is applied to simulated calibration objects with known projected sizes and

curvatures. The resulting bias model is

$$\widehat{\mathbf{b}} = \mathbf{B}_0 + \mathbf{B}_1 \left[\sigma_I,\, T_h,\, v_{img}\right]^T, \tag{3.5}$$

where T_h is the threshold perturbation and v_{img} is apparent image velocity. The feature vector is corrected by subtracting $\widehat{\mathbf{b}}$. Without this correction, shear modulus estimates are biased upward in high-blur cases because blurred outlines appear smoother and less deformed.

The observer is evaluated using four metrics:

- posterior median relative error for each parameter;
- posterior credible interval coverage for 80% marginal intervals;
- normalized negative log predictive density for future frames;
- rank correlation between inferred stiffness and true stiffness across a mixed population.

The rank metric is included because many screening assays care about ordering cells by mechanical phenotype even when absolute modulus calibration is uncertain.

An important diagnostic is the posterior correlation matrix. Some parameter pairs show persistent coupling. Shear modulus and flow speed are negatively correlated because stronger flow can mimic lower stiffness. Reference-shape inflation and inlet orientation are correlated because an initially fuller cell at one tilt can project like a more deflated cell at another tilt. Viscosity ratio and bending ratio correlate weakly in the first few frames but decouple when centroid lag is included. These correlations are reported rather than hidden, because they indicate how a measurement protocol could be redesigned.

The observer also estimates a model-error inflation factor α_Q multiplying $\mathbf{Q}$. When the reduced dynamics cannot match the feature sequence, α_Q increases and prevents artificial certainty. This is a small hierarchical extension:

$$\begin{aligned} \mathbf{Q} &= \alpha_Q \mathbf{Q}_0, \\ \log \alpha_Q &\sim \mathcal{N}\left(0, 0.25^2\right). \end{aligned} \tag{3.6}$$

The inflation factor is not interpreted as physical noise. It is a guard against a reduced model pretending to be exact.

4. Numerical Records and Perturbation Design

The generated image records are organized into a factorial design. The baseline channel has width 9.5 μm, nominal cell radius 3.1 μm, and maximum flow speed 5.8 mm/s. These dimensional values set the image scale, but all inference uses nondimensional dynamics. The camera pixel size is 0.11 μm, and the simulated optical blur standard deviation ranges from 0.7 to 1.6 pixels. The viewing direction is perpendicular to the channel axis unless otherwise specified.

The design includes four disturbance levels:

- low noise with pixel noise below 1.5% of foreground contrast and blur below one pixel;
- moderate noise with 3.0% pixel noise and threshold drift over 6% of the foreground range;
- high noise with 5.0% pixel noise, stronger blur, and occasional missing contour arcs;
- tilted imaging in which the optical axis is rotated by 12° around the channel direction.

The tilted imaging case tests whether a small unmodeled projection error harms inference more than random noise. In those records the observer still assumes a perpendicular view unless the tilt-corrected variant is explicitly used.

The full test set contains 1,260 image sequences. There are 900 baseline sequences with random parameters drawn from the prior, 180 sequences designed near weak-identifiability regions, and 180 sequences generated with forward-model mismatch. The mismatch cases use a slightly different reference-surface relaxation law and a non-polynomial damping term in the high-order shape modes. These cases are included because an observer that works only when the reduced model generated the data is not very informative.

Segmentation is performed with a hybrid of thresholding and contour regularization. Otsu's threshold initializes the foreground mask, a Canny edge map supplies candidate boundary pixels, and a closed spline is fitted by minimizing a weighted edge-and-area objective [36, 37]. Optical-flow estimates are used only to initialize the next frame's contour, not as direct observations [38, 39]. A U-Net segmentation alternative was tested on the same generated images [40]. It gave cleaner masks in high-noise cases but introduced shape-dependent bias when the training set underrepresented high-tilt cells, so the baseline uses the more transparent classical pipeline.

The forward solver produces hidden state trajectories at fine time resolution. The image sequences sample those trajectories at frame times determined by the acquisition design. Uniform acquisition uses equally spaced frames. Active acquisition uses the information-gain rule after the fourth frame. The first four frames are fixed because the observer needs enough information to form a rough posterior before nonuniform timing becomes useful. The maximum sequence length is 32 frames, but several comparisons stop once posterior uncertainty falls below a preset tolerance.

The tolerance criterion is

$$\frac{\sqrt{\mathrm{Var}(G_s \mid \mathbf{y}_{1:k})}}{\mathbb{E}[G_s \mid \mathbf{y}_{1:k}]} < 0.08, \qquad \frac{\sqrt{\mathrm{Var}(\rho_0 \mid \mathbf{y}_{1:k})}}{\mathbb{E}[\rho_0 \mid \mathbf{y}_{1:k}]} < 0.035. \tag{4.1}$$

No stopping rule is imposed for λ because viscosity ratio is known to be more weakly identified from silhouettes. A run that meets both criteria is considered successful for stiffness-reference inference.

The parameter ranges are intentionally broad enough to create ambiguity. A low G_s cell and a high U_0 cell can have similar elongation. A low ρ_0 reference shape and a high inlet angle can produce similar projected asymmetry. A high λ cell and a high χ_b cell can both show delayed shape change in some views. The observer must use temporal ordering to separate these effects. Static snapshots are included only as a comparison baseline.

The static baseline estimates parameters from a single frame chosen near the end of the sequence. It uses the same likelihood and the same prior but with no temporal update. This comparison is deliberately favorable to static inference because the final frame is chosen after the cell has undergone appreciable deformation. Even with this favorable choice, the static baseline shows larger posterior uncertainty and stronger parameter correlations.

The study also includes a rank-ordering experiment. A population of 300 cells is generated with G_s drawn from a bimodal distribution whose modes differ by 22%. Other parameters are drawn from broad nuisance distributions. The task is not to recover every parameter

exactly, but to rank cells by G_s. The area under the receiver operating curve is computed for distinguishing the softer group from the stiffer group. This experiment mimics a screening situation where relative mechanical phenotype is more important than precise biophysical calibration.

A small robustness study varies the feature set. The observer is run with all features, without curvature features, without timing features, and without third moments. Removing curvature features primarily harms χ_b. Removing timing features harms λ. Removing third moments harms ϕ_0 and ρ_0. Projected area and major-axis length alone are insufficient for reliable inference under moderate noise. This result is not surprising, but it quantifies the cost of aggressive feature compression.

5. Results

Table 1. State and parameter variables in the observer framework.

Symbol	Meaning	Range/Value	Role
G_s	Shear modulus scale	0.55–1.70 G_{ref}	Mechanical parameter
χ_b	Bending ratio	0.012–0.060	Mechanical parameter
λ	Viscosity ratio	1.0–7.5	Mechanical parameter
ρ_0	Reference-shape inflation	0.90–1.08	Geometry parameter
ϕ_0	Inlet orientation	$\pm 35^\circ$	Initial condition
U_0	Mean flow speed	$\pm 20\%$ nominal	Flow parameter

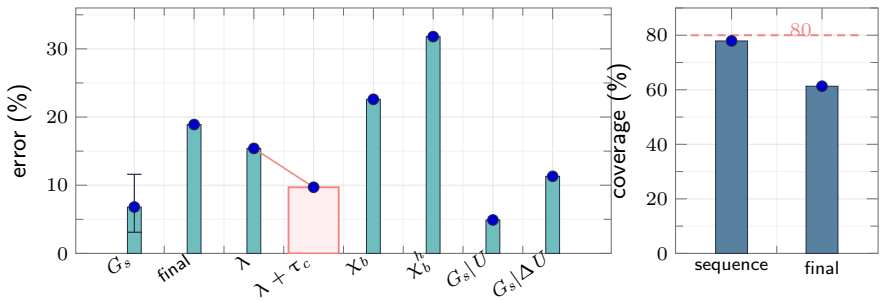


Figure 7. Transient silhouettes reduce stiffness error relative to final-frame inference and keep posterior coverage close to the nominal 80% interval, while viscosity and bending remain less constrained under ordinary outline features.

Table 2. Feature vector used for image assimilation.

Feature Group	Examples	Main Sensitivity
Global geometry	Area, perimeter	G_s, ρ_0
Axes measures	Major/minor axis lengths	G_s, U_0
Shape descriptors	Fourier coefficients	ρ_0, ϕ_0
Curvature metrics	15%, 85% quantiles	χ_b
Centroid motion	Position, lag	λ, U_0
Higher moments	Third central moments	ρ_0, ϕ_0

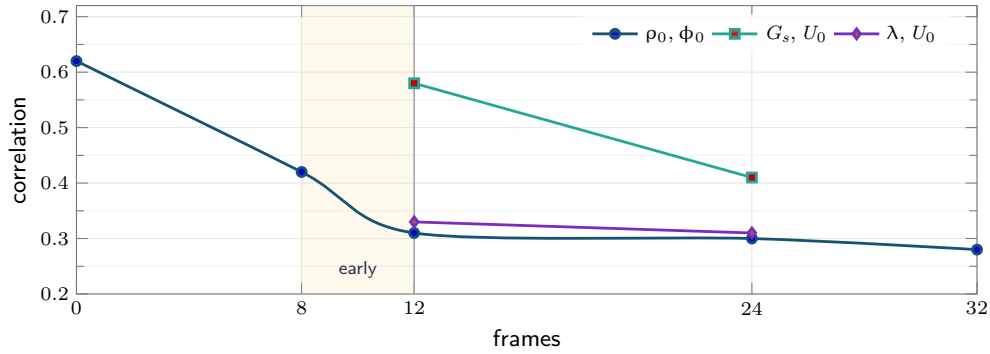


Figure 8. Posterior correlations contract as ordered frames arrive; geometry–orientation ambiguity falls rapidly, whereas viscosity–speed ambiguity remains nearly unchanged between 12 and 24 frames.

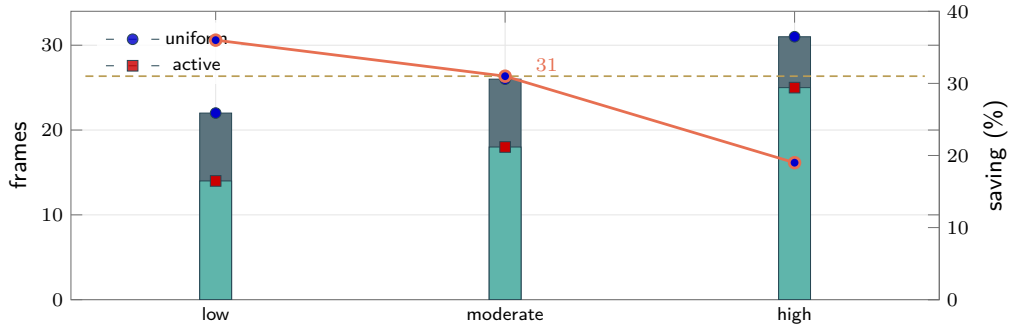


Figure 9. The acquisition rule concentrates samples around informative transient phases and reduces median sequence length most strongly when measurement noise is not the dominant limitation.

The baseline observer recovers shear modulus more accurately than the other physical parameters. For moderate-noise sequences with 24 frames, the median relative error in G_s is 6.8%, with an interquartile range from 3.1% to 11.6%. The 80% credible intervals contain the true value in 77.9% of cases, indicating mild overconfidence but acceptable calibration for the reduced model. Static inference from one final frame gives median error 18.9%, with credible interval coverage only 61.3%. The time sequence therefore contributes information beyond improved averaging.

Reference-shape inflation ρ_0 is also recoverable. Its median absolute error is 0.018 in nondimensional units for moderate-noise sequences and 0.026 for high-noise sequences. The posterior contracts quickly during the first 8 to 12 frames, because early projected asymmetry depends strongly on the relaxed reference geometry. After that, additional frames add less information. The posterior correlation between ρ_0 and inlet angle ϕ_0 is initially 0.62, decreases to 0.31 by frame 12, and remains near 0.28 through frame 32. Some geometric ambiguity is therefore persistent.

Viscosity ratio is harder. With silhouette features only, the median relative error in λ is 15.4% at 24 frames and 13.9% at 32 frames. Adding centroid lag reduces the 24-frame error to 9.7%. The improvement occurs mostly for $\lambda > 4$, where internal-fluid resistance delays visible shape response and slightly shifts translational timing. At $\lambda < 2$, the posterior remains broad because the silhouette differences are small compared with segmentation noise

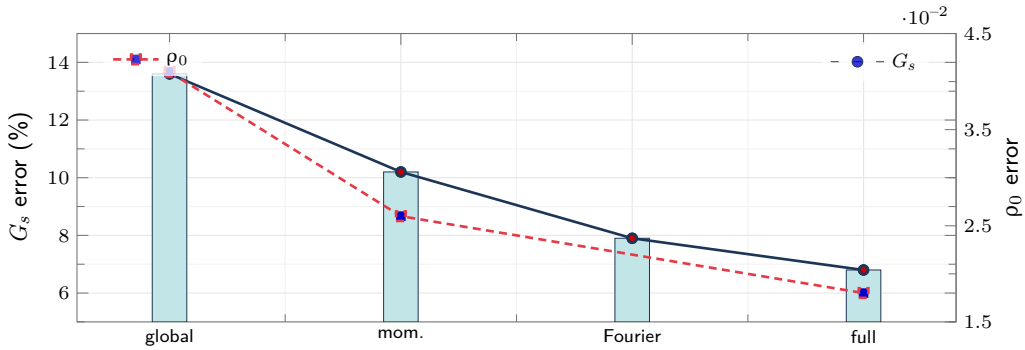


Figure 10. Feature ablation shows that global outline measures alone are insufficient; asymmetry, Fourier content, and curvature progressively improve stiffness recovery, while reference-shape error drops sharply after third moments are included.

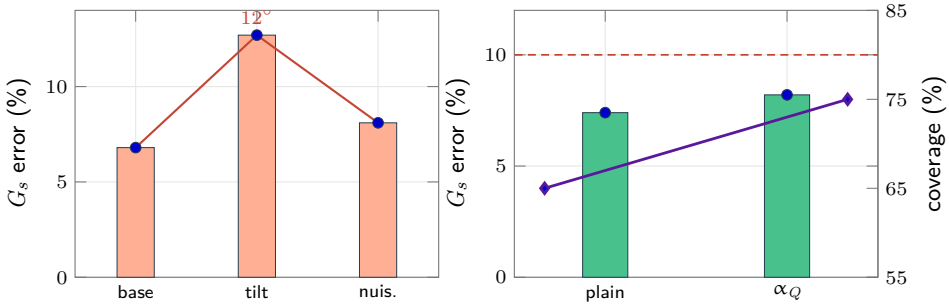


Figure 11. Projection mismatch raises stiffness error unless tilt is modeled as a nuisance parameter, and process-noise inflation trades a small error increase for substantially more honest posterior coverage.

and uncertainty in U_0.

Bending ratio χ_b is the weakest parameter under ordinary imaging. Its median relative error is 22.6% in moderate noise and 31.8% in high noise. Removing curvature quantiles increases the moderate-noise error to 38.5%. The bending ratio affects localized contour curvature, but those features are precisely the ones most sensitive to blur and thresholding. The posterior intervals for χ_b are wide and mostly well calibrated, which is preferable to confident failure. In practice, estimating bending from these sequences would require sharper imaging, a stronger deformation protocol, or a complementary measurement.

Flow-speed uncertainty is partly separable from shear modulus because speed affects timing and advection while modulus affects deformation amplitude and relaxation. When U_0 is treated as known, G_s error falls to 4.9%. When U_0 is uncertain by 12%, the error is 6.8%. When U_0 uncertainty is increased to 25%, G_s error rises to 11.3%, and the posterior correlation between G_s and U_0 reaches –0.71. This result supports a straightforward experimental recommendation: local velocity calibration is as important as contour resolution for modulus inference.

The active acquisition rule reduces the required sequence length. Under moderate noise, uniform sampling needs a median of 26 frames to satisfy the stiffness-reference stopping rule. Active sampling needs 18 frames, a reduction of 31%. The selected frames are not evenly

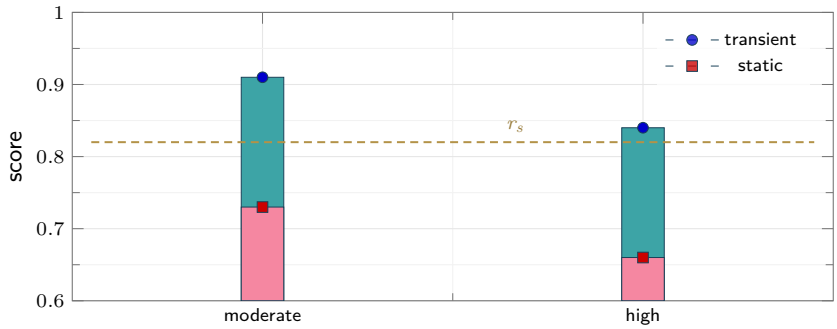


Figure 12. Transient assimilation improves stiffness-population ranking under both moderate and high noise; the moderate-noise posterior medians also preserve a strong rank relation with the true stiffness values.

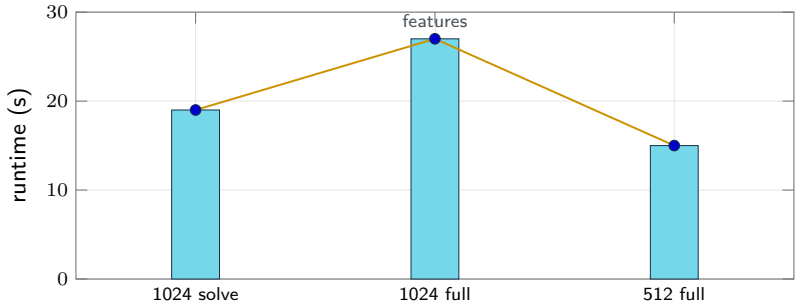


Figure 13. The observer is practical for offline assay design but remains too expensive for real-time use: feature extraction and ensemble size dominate the reported 24-frame workstation runtime.

distributed. The algorithm samples densely during early rotation, waits through a period of nearly redundant elongation, and samples again when contour curvature begins to relax. This behavior emerges from the information criterion rather than being programmed as a heuristic.

The advantage of active acquisition is smaller under high noise. Uniform sampling needs a median of 31 frames, while active sampling needs 25, a reduction of 19%. The reason is that noisy features require repeated observations for averaging, and skipping frames becomes less beneficial. Under low noise, active sampling reaches the stopping rule with a median of 14 frames, compared with 22 for uniform sampling. The information gain from timing is therefore strongest when measurement noise is not the dominant limitation.

The tilted imaging case exposes a model discrepancy. If the observer assumes perpendicular imaging while the data are generated with a 12° optical-axis tilt, G_s error increases from 6.8% to 12.7%, and ρ_0 error increases from 0.018 to 0.033. The posterior predictive density also worsens, signaling that the model does not fully explain the observations. When tilt angle is added as a nuisance parameter with a prior over $[-15°, 15°]$, G_s error drops to 8.1%, but ρ_0 remains less accurate than in the untilted case. Projection calibration should not be treated as a minor optical detail.

The model-mismatch sequences are more revealing than the clean sequences. Without process-noise inflation, the observer reports narrow posteriors and misses the true G_s in 35% of 80%

Table 3. Training and validation data for reduced-order modeling.

Dataset	Trajectories	Purpose	Outcome
Training set	420	Drift model fitting	Reduced dynamics
Validation set	160	Reconstruction check	Error assessment
Test sequences	900	Baseline evaluation	Inference study
Weak-ID cases	180	Ambiguity analysis	Robustness check
Mismatch cases	180	Model discrepancy	Calibration test

Table 4. Imaging disturbance scenarios.

Scenario	Pixel Noise	Blur Level	Additional Effect
Low noise	$< 1.5\%$	< 1 pixel	None
Moderate noise	3.0%	Moderate	Threshold drift
High noise	5.0%	Strong	Missing contour arcs
Tilted imaging	Moderate	Moderate	12° optical tilt

intervals. With the inflation factor α_Q, coverage improves to 75%, and median G_s error increases slightly from 7.4% to 8.2%. The increase in error is acceptable because uncertainty is more honest. The inflation factor is largest in sequences with strong high-order contour content, suggesting that the reduced basis is least adequate for cells with rapid local shape changes.

The rank-ordering experiment gives a better outcome than absolute parameter recovery. For the bimodal stiffness population, the observer achieves an area under the receiver operating curve of 0.91 using 20 frames under moderate noise. Static inference gives 0.73. With high noise, the observer gives 0.84, while static inference gives 0.66. The rank correlation between posterior median G_s and true G_s is 0.82 under moderate noise. This suggests that even when absolute modulus values are imperfect, transient silhouette observers can still be useful for mechanical sorting or screening.

Feature ablation confirms that not all contour information is equivalent. Using only area, perimeter, and aspect ratio gives median G_s error 13.6% and ρ_0 error 0.041. Adding third moments lowers G_s error to 10.2% and ρ_0 error to 0.026. Adding Fourier descriptors lowers G_s error to 7.9%. Adding curvature quantiles gives the baseline 6.8%. Timing features barely change G_s but improve λ. The observer therefore benefits from a layered feature set: global size for scale, asymmetry for orientation and reference shape, curvature for bending, and timing for viscosity ratio.

Posterior correlations reveal recurring non-identifiability. At 12 frames, the median absolute correlation between G_s and U_0 is 0.58, between ρ_0 and ϕ_0 is 0.45, and between λ and U_0 is 0.33. By 24 frames, the first falls to 0.41, the second to 0.30, and the third remains near 0.31. More frames reduce geometry-orientation ambiguity but do not fully solve viscosity-speed ambiguity. A protocol that changes flow speed during observation may be needed to separate those parameters more cleanly.

A two-speed protocol was therefore tested on 120 sequences. The imposed speed is reduced by 18% halfway through the observation window. The observer is given the nominal time of the speed change but not the exact local speeds. This protocol reduces median λ error from 9.7% to 7.8% when centroid lag is included. It slightly worsens G_s error, from 6.8% to 7.4%, because the deformation history becomes less close to the training trajectories. The result

Table 5. Baseline parameter recovery performance.

Parameter	Median Error	Noise Level	Identifiability
G_s	6.8%	Moderate	Strong
ρ_0	0.018	Moderate	Strong
λ	15.4%	Moderate	Moderate
λ + lag	9.7%	Moderate	Improved
χ_b	22.6%	Moderate	Weak

Table 6. Comparison of acquisition strategies.

Noise Level	Uniform Frames	Active Frames	Reduction
Low	22	14	36%
Moderate	26	18	31%
High	31	25	19%
Overall trend	Longer sequences	Shorter sequences	Positive gain

suggests a trade-off: flow modulation helps viscosity inference but requires a forward model trained on modulated protocols.

Failure cases fall into three groups. The first group consists of nearly symmetric projected sequences in which inlet orientation is close to zero and the view hides out-of-plane deformation. The second group consists of high-blur sequences where curvature descriptors become unreliable. The third group consists of parameter combinations near the edge of the training envelope. In the first group, no observer can recover hidden three-dimensional information from a single view without stronger prior assumptions. In the second group, better optics or less curvature-dependent features are needed. In the third group, the reduced model must be retrained or the posterior must explicitly report extrapolation.

The computational cost is compatible with offline assay design but not yet with real-time clinical use. A 24-frame sequence with 1,024 ensemble members takes 19 seconds on a workstation GPU when feature extraction is excluded and 27 seconds when feature extraction is included. Reducing the ensemble to 512 members lowers runtime to 15 seconds with modest accuracy loss. A neural surrogate for the observation map reduces runtime further, but posterior calibration worsens unless uncertainty in the surrogate is included [41, 42]. The present implementation is therefore a research observer rather than a production pipeline.

6. Interpretation

The main reason transient data outperform static contours is that parameter effects unfold at different times. Shear modulus affects the initial rate of elongation and the subsequent deformation amplitude. Reference-shape inflation affects the asymmetry of the early projected contour before the cell has rotated far. Viscosity ratio affects timing more than instantaneous outline, especially when internal-fluid lag changes how quickly the membrane responds to the imposed flow. Bending ratio affects high-curvature portions of the boundary, which appear only in certain phases and may be blurred.

This time separation is incomplete, but it is enough for inference. If all parameters affected only the final silhouette, Bayesian updating would merely sharpen a low-dimensional curve of equivalent explanations. Instead, the posterior changes shape as frames arrive. Early frames

Table 7. Impact of feature ablation on inference quality.

Feature Set	G_s Error	ρ_0 Error	Main Impact
Area+perimeter+aspect ratio	13.6%	0.041	Poor recovery
+ Third moments	10.2%	0.026	Better asymmetry
+ Fourier descriptors	7.9%	0.026	Shape improvement
+ Curvature quantiles	6.8%	0.018	Baseline accuracy

Table 8. Posterior parameter correlations during filtering.

Parameter Pair	Correlation (12 frames)	Correlation (24 frames)	Interpretation
G_s–U_0	0.58	0.41	Stiffness–speed ambiguity
ρ_0–ϕ_0	0.45	0.30	Geometry–orientation coupling
λ–U_0	0.33	0.31	Timing ambiguity
ρ_0–ϕ_0 (early)	0.62	0.28	Persistent coupling

reduce orientation and reference uncertainty. Middle frames reduce stiffness-flow uncertainty. Later frames add some information about viscosity and bending, although with diminishing returns. The observer's value lies in using this changing information content.

The results also explain why deformation indices can be unstable across devices. A scalar index may mix stiffness, orientation, velocity, and reference geometry. If the distribution of inlet orientations changes between devices, a stiffness index calibrated in one device may shift in another even when the cell population is unchanged. Bayesian inference does not remove this problem automatically, but it makes the nuisance variables explicit. A wide posterior or strong parameter correlation warns that the protocol is not measuring a single mechanical quantity cleanly.

The active acquisition results should be read as a statement about information timing rather than camera hardware. A high-speed camera collecting every frame can still discard redundant frames later. The point is that equal spacing is not inherently optimal. If the early transient is information-rich, sampling it densely is useful. If the cell then enters a slowly changing interval, additional frames contribute little except noise averaging. A device designed for inference could exploit this by using entrance geometry, flow modulation, or illumination timing to concentrate information in a short window.

The observer shows that viscosity ratio is not hopeless, but it is not primarily an outline parameter. It becomes more accessible when timing and centroid lag are included. This matters because many image-based deformability assays focus almost entirely on shape. If internal viscosity or cytoplasmic changes are of interest, the assay should preserve temporal calibration and local speed measurements. Otherwise, viscosity ratio may be absorbed into apparent stiffness or flow-speed uncertainty.

The bending-ratio result is a caution about overinterpreting smooth contours. Bending changes local curvature, but microscopy blur and thresholding also change local curvature. The posterior uncertainty for χ_b remains large because the observer cannot reliably tell whether a smoother outline comes from higher bending resistance or from image processing. Stronger deformation, sharper imaging, or a likelihood that uses grayscale edge profiles rather than segmented curvature may improve this. The current contour-feature observer is not a strong bending estimator.

Table 9. Effects of model mismatch and optical tilt.

Case	G_s Error	Additional Metric	Observation
Baseline	6.8%	Reference case	Best performance
12° tilt ignored	12.7%	Higher NLPD	Projection bias
Tilt estimated	8.1%	Improved fit	Partial correction
Mismatch + no inflation	7.4%	Poor coverage	Overconfidence
Mismatch + inflation	8.2%	Better coverage	Calibrated uncertainty

Table 10. Computational performance and ranking results.

Configuration	Runtime / Metric	Value	Outcome
1024 particles	Runtime	27 s	Full observer
512 particles	Runtime	15 s	Faster execution
Transient observer	ROC AUC	0.91	Moderate-noise ranking
Static inference	ROC AUC	0.73	Lower ranking quality
High-noise observer	ROC AUC	0.84	Robust performance

The model-discrepancy experiment supports the use of hierarchical uncertainty. A reduced observer is tempting because it is fast. Speed becomes dangerous if model error is ignored. The process-noise inflation factor does not identify the missing physics, but it prevents the posterior from becoming too narrow. In applications where decisions depend on credible intervals, calibrated uncertainty may be more important than a small improvement in median error.

Several aspects of the observer could be replaced by modern learning methods. Neural ordinary differential equations, physics-informed neural networks, and operator-learning models offer flexible state evolution and differentiable prediction [43, 44, 45]. The reason they are not used as the main engine here is interpretability under limited data. The constrained reduced model makes it easier to track which feature affects which parameter. A future hybrid version could use operator learning for speed while retaining Bayesian calibration and explicit nuisance variables.

The paper's inference framing also changes how simulation results should be archived. For a direct mechanical study, storing final shapes and phase labels may be enough. For an inverse observer, the useful objects are time-stamped projected features, hidden parameter values, image noise variables, and reconstruction residuals. The same forward simulation can support many observer tests if those records are saved. This is a practical point, but it strongly affects reproducibility.

7. Limitations

The records are generated rather than measured. The simulated images include blur, threshold drift, pixel noise, and missing arcs, but they do not include every artifact of a real microscope. Out-of-focus halos, nonuniform illumination, cell overlap, rolling contact with walls, and biochemical heterogeneity are absent. The reported errors therefore measure observer performance under a controlled image model. They do not establish performance on clinical samples.

The reduced capsule dynamics are approximate. They are trained on full-order trajectories,

but they cannot represent arbitrary membrane folding or strongly non-axisymmetric events outside the training envelope. The model-error inflation factor partly compensates for this limitation, yet it cannot create missing state variables. If real cells exhibit dynamics not present in the training records, posterior intervals may still be misleading unless discrepancy is modeled more explicitly.

Only single-cell sequences are considered. In real suspensions, nearby cells perturb the flow, alter apparent speed, and sometimes obscure contours. Hydrodynamic interactions may change the transient features used by the observer. A practical assay would need detection rules that reject interacting cells or a multi-cell forward model. The present study is a step toward single-cell inference, not a complete blood-flow analysis.

The image formation model uses one viewing direction. A second view would likely improve inference, especially for orientation and bending. Light-sheet imaging, holography, or multi-plane microscopy could provide richer data, but they also introduce new calibration variables. The single-view assumption is restrictive, but it reflects many high-throughput settings where simple optical access is preferred.

The observer treats the cell as a passive capsule. Real red blood cells can vary in membrane composition, cytoskeletal integrity, metabolic state, and age. Some pathological cells may not be represented by smooth variations of G_s, χ_b, λ, and ρ_0. In such cases, forcing a passive-capsule parameter estimate may be less informative than classifying the sequence as outside the model family. The posterior predictive density and model inflation factor are useful warning signs, but they are not a substitute for biological validation.

The prior distributions are broad design choices, not population estimates. Different priors would change posterior medians in weakly identified cases. Stronger priors from independent measurements could improve accuracy, while poorly chosen priors could hide non-identifiability. Any practical use of the observer should report prior assumptions together with posterior estimates.

The feature likelihood is Gaussian after preprocessing. This is convenient, but segmentation errors can be skewed and frame-correlated. Missing contour arcs, for example, produce structured errors in Fourier descriptors and curvature quantiles. A more complete likelihood might use a mixture model or directly compare probabilistic contours. That improvement would increase computational cost but could improve calibration in difficult images.

8. Conclusion

This paper treated microfluidic cell-shape imaging as a Bayesian inverse problem rather than as a direct readout of deformability. Under the controlled numerical records used here, transient silhouettes contain enough information to estimate membrane shear scale and reference-shape inflation with useful accuracy. The key requirement is that the observer uses the time sequence. A final contour, even when clean, does not separate stiffness, orientation, flow-speed variation, and reference geometry as reliably as a short ordered sequence.

The results show uneven identifiability across parameters. Shear modulus is the most accessible because it affects both deformation rate and amplitude. Reference-shape inflation is identifiable through early asymmetry. Viscosity ratio is only moderately identifiable from contours, but it improves when centroid timing is used. Bending ratio is weak under blurred single-view contours because the relevant curvature information is vulnerable to segmentation and optical smoothing. These differences are not failures of the observer; they describe what the chosen measurement can and cannot determine.

Active acquisition reduced the number of frames needed for stiffness-reference inference by concentrating observations around informative transient intervals. The result suggests that microfluidic imaging protocols should be designed around information content, not merely around maximum frame rate or final deformation. Flow modulation may also improve viscosity inference, although it requires training the forward model on modulated histories.

The observer remained sensitive to optical-axis tilt and forward-model mismatch. Adding nuisance parameters and process-noise inflation improved calibration, but did not eliminate all bias. This is a reminder that uncertainty reporting is part of the measurement. A fast reduced model can become misleading when its errors are treated as negligible. Posterior predictive checks and model-discrepancy variables should accompany any image-based mechanical estimate.

The study does not claim that the numerical values reported here transfer directly to experiments. It provides a structured way to test whether a proposed imaging assay has enough information to support a mechanical inference. Before applying such an observer to biological samples, the forward model would need calibration against independent mechanical measurements and validation under the actual optical and microfluidic conditions.

A useful next step would combine two viewing angles, local velocity calibration, and a reduced model trained on experimentally matched geometries. Another extension would allow the observer to reject cells whose sequences are inconsistent with the passive-capsule family. Such rejection may be as important as parameter estimation, because abnormal cells can differ from the model in ways that no adjustment of shear modulus or viscosity ratio can represent.

The main conclusion is that transient silhouettes are not merely imperfect substitutes for three-dimensional membrane data. They are structured observations with time-dependent information. When that structure is used through a calibrated Bayesian observer, short image sequences can recover some membrane parameters and honestly expose others as weakly identifiable.

References

[1] E. A. Evans. “A New Material Concept for the Red Cell Membrane”. In: *Biophysical Journal* 13.9 (1973), pp. 926–940. DOI: 10.1016/S0006-3495(73)86035-7.

[2] E. A. Evans, R. Waugh, and L. Melnik. “Elastic Area Compressibility Modulus of Red Cell Membrane”. In: *Biophysical Journal* 16.6 (1976), pp. 585–595. DOI: 10.1016/S0006-3495(76)85713-X.

[3] R. M. Hochmuth, N. Mohandas, and P. L. Blackshear. “Measurement of the Elastic Modulus for Red Cell Membrane Using a Fluid Mechanical Technique”. In: *Biophysical Journal* 13.8 (1973), pp. 747–762. DOI: 10.1016/S0006-3495(73)86021-7.

[4] R. Waugh and E. A. Evans. “Thermoelasticity of Red Blood Cell Membrane”. In: *Biophysical Journal* 26.1 (1979), pp. 115–131. DOI: 10.1016/S0006-3495(79)85239-X.

[5] Narla Mohandas and Patrick G. Gallagher. “Red Cell Membrane: Past, Present, and Future”. In: *Blood* 112.10 (2008), pp. 3939–3948. DOI: 10.1182/blood-2008-07-161166.

[6] Monica Diez-Silva, Ming Dao, Jongyoon Han, Chwee Teck Lim, and Subra Suresh. “Shape and Biomechanical Characteristics of Human Red Blood Cells in Health and Disease”. In: *MRS Bulletin* 35.5 (2010), pp. 382–388. DOI: 10.1557/mrs2010.571.

[7] Subra Suresh. “Biomechanics and Biophysics of Cancer Cells”. In: *Acta Biomaterialia* 3.4 (2007), pp. 413–438. DOI: 10.1016/j.actbio.2007.04.002.

[8] YongKeun Park, Catherine A. Best, Kamran Badizadegan, Ramachandra R. Dasari, Michael S. Feld, Tatyana Kuriabova, Michael L. Henle, Arnold J. Levine, and Gabriel Popescu. "Measurement of Red Blood Cell Mechanics during Morphological Changes". In: *Proceedings of the National Academy of Sciences of the United States of America* 107.15 (2010), pp. 6731–6736. DOI: 10.1073/pnas.0909533107.

[9] H. A. Stone, A. D. Stroock, and A. Ajdari. "Engineering Flows in Small Devices: Microfluidics Toward a Lab-on-a-Chip". In: *Annual Review of Fluid Mechanics* 36 (2004), pp. 381–411. DOI: 10.1146/annurev.fluid.36.050802.122124.

[10] Todd M. Squires and Stephen R. Quake. "Microfluidics: Fluid Physics at the Nanoliter Scale". In: *Reviews of Modern Physics* 77.3 (2005), pp. 977–1026. DOI: 10.1103/RevModPhys.77.977.

[11] George M. Whitesides. "The Origins and the Future of Microfluidics". In: *Nature* 442.7101 (2006), pp. 368–373. DOI: 10.1038/nature05058.

[12] Dino Di Carlo. "Inertial Microfluidics". In: *Lab on a Chip* 9.21 (2009), pp. 3038–3046. DOI: 10.1039/B912547G.

[13] Jochen Guck et al. "Optical Deformability as an Inherent Cell Marker for Testing Malignant Transformation and Metastatic Competence". In: *Biophysical Journal* 88.5 (2005), pp. 3689–3698. DOI: 10.1529/biophysj.104.045476.

[14] Daniel R. Gossett, Henry T. K. Tse, Sarah A. Lee, Ying Ying, Anders G. Lindgren, Otto O. Yang, Jay Rao, Aaron T. Clark, and Dino Di Carlo. "Hydrodynamic Stretching of Single Cells for Large Population Mechanical Phenotyping". In: *Proceedings of the National Academy of Sciences of the United States of America* 109.20 (2012), pp. 7630–7635. DOI: 10.1073/pnas.1200107109.

[15] Oliver Otto et al. "Real-Time Deformability Cytometry: On-the-Fly Cell Mechanical Phenotyping". In: *Nature Methods* 12.3 (2015), pp. 199–202. DOI: 10.1038/nmeth.3281.

[16] Henry T. K. Tse, Daniel R. Gossett, Yong Sung Moon, Mahdokht Masaeli, Michael Sohsman, Ying Ying, Katherine Mislick, Ryan P. Adams, Jay Rao, and Dino Di Carlo. "Quantitative Diagnosis of Malignant Pleural Effusions by Single-Cell Mechanophenotyping". In: *Science Translational Medicine* 5.212 (2013), 212ra163. DOI: 10.1126/scitranslmed.3006559.

[17] Dhwanit Agarwal and George Biros. "Stable shapes of three-dimensional vesicles in unconfined and confined Poiseuille flow". In: *Physical Review Fluids* 5.1 (2020), p. 013603.

[18] Albert Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. Philadelphia: Society for Industrial and Applied Mathematics, 2005. DOI: 10.1137/1.9780898717921.

[19] Jari Kaipio and Erkki Somersalo. *Statistical and Computational Inverse Problems*. Vol. 160. Applied Mathematical Sciences. New York: Springer, 2005. DOI: 10.1007/b138659.

[20] Andrew M. Stuart. "Inverse Problems: A Bayesian Perspective". In: *Acta Numerica* 19 (2010), pp. 451–559. DOI: 10.1017/S0962492910000061.

[21] Marc C. Kennedy and Anthony O'Hagan. "Bayesian Calibration of Computer Models". In: *Journal of the Royal Statistical Society: Series B* 63.3 (2001), pp. 425–464. DOI: 10.1111/1467-9868.00294.

[22] Jenný Brynjarsdóttir and Anthony O'Hagan. "Learning about Physical Parameters: The Importance of Model Discrepancy". In: *Inverse Problems* 30.11 (2014), p. 114007. DOI: 10.1088/0266-5611/30/11/114007.

[23] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global Sensitivity Analysis: The Primer*. Chichester: John Wiley & Sons, 2008. DOI: 10.1002/9780470725184.

[24] I. M. Sobol'. "Global Sensitivity Indices for Nonlinear Mathematical Models and Their Monte Carlo Estimates". In: *Mathematics and Computers in Simulation* 55.1–3 (2001), pp. 271–280. DOI: 10.1016/S0378-4754(00)00270-6.

[25] Max D. Morris. "Factorial Sampling Plans for Preliminary Computational Experiments". In: *Technometrics* 33.2 (1991), pp. 161–174. DOI: 10.1080/00401706.1991.10484804.

[26] Peter J. Schmid. "Dynamic Mode Decomposition of Numerical and Experimental Data". In: *Journal of Fluid Mechanics* 656 (2010), pp. 5–28. DOI: 10.1017/S0022112010001217.

[27] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge: Cambridge University Press, 2019. DOI: 10.1017/9781108380690.

[28] Peter Benner, Serkan Gugercin, and Karen Willcox. "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems". In: *SIAM Review* 57.4 (2015), pp. 483–531. DOI: 10.1137/130932715.

[29] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. "Learning Nonlinear Operators via DeepONet Based on the Universal Approximation Theorem of Operators". In: *Nature Machine Intelligence* 3.3 (2021), pp. 218–229. DOI: 10.1038/s42256-021-00302-5.

[30] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. "Fourier Neural Operator for Parametric Partial Differential Equations". In: *International Conference on Learning Representations*. 2021.

[31] Dhwanit Agarwal and George Biros. "Shape dynamics of a red blood cell in Poiseuille flow". In: *Physical Review Fluids* 7.9 (2022), p. 093602.

[32] R. E. Kalman. "A New Approach to Linear Filtering and Prediction Problems". In: *Journal of Basic Engineering* 82.1 (1960), pp. 35–45. DOI: 10.1115/1.3662552.

[33] Andrew H. Jazwinski. *Stochastic Processes and Filtering Theory*. Vol. 64. Mathematics in Science and Engineering. New York: Academic Press, 1970.

[34] Arnaud Doucet, Nando de Freitas, and Neil Gordon, eds. *Sequential Monte Carlo Methods in Practice*. New York: Springer, 2001. DOI: 10.1007/978-1-4757-3437-9.

[35] Geir Evensen. "The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation". In: *Ocean Dynamics* 53 (2003), pp. 343–367. DOI: 10.1007/s10236-003-0036-9.

[36] Nobuyuki Otsu. "A Threshold Selection Method from Gray-Level Histograms". In: *IEEE Transactions on Systems, Man, and Cybernetics* 9.1 (1979), pp. 62–66. DOI: 10.1109/TSMC.1979.4310076.

[37] John Canny. "A Computational Approach to Edge Detection". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-8.6 (1986), pp. 679–698. DOI: 10.1109/TPAMI.1986.4767851.

[38] Bruce D. Lucas and Takeo Kanade. "An Iterative Image Registration Technique with an Application to Stereo Vision". In: *Proceedings of the 7th International Joint Conference on Artificial Intelligence*. 1981, pp. 674–679.

[39] Berthold K. P. Horn and Brian G. Schunck. "Determining Optical Flow". In: *Artificial Intelligence* 17.1–3 (1981), pp. 185–203. DOI: 10.1016/0004-3702(81)90024-2.

[40] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. "U-Net: Convolutional Networks for Biomedical Image Segmentation". In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Vol. 9351. Lecture Notes in Computer Science. Springer, 2015, pp. 234–241. DOI: 10.1007/978-3-319-24574-4_28.

[41] Yarin Gal and Zoubin Ghahramani. "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning". In: *Proceedings of the 33rd International*

Conference on Machine Learning. Vol. 48. Proceedings of Machine Learning Research. 2016, pp. 1050–1059.

[42] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. "Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles". In: *Advances in Neural Information Processing Systems*. Vol. 30. 2017.

[43] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. "Neural Ordinary Differential Equations". In: *Advances in Neural Information Processing Systems*. Vol. 31. 2018.

[44] M. Raissi, P. Perdikaris, and G. E. Karniadakis. "Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations". In: *Journal of Computational Physics* 378 (2019), pp. 686–707. DOI: 10.1016/j.jcp.2018.10.045.

[45] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. "Physics-Informed Machine Learning". In: *Nature Reviews Physics* 3.6 (2021), pp. 422–440. DOI: 10.1038/s42254-021-00314-5.