Archived source document. Publication links are updated; scientific text and reported values are retained. Raw and complete data remain in the unpublished Zenodo preparation package.

Back to Model

Complete supporting record

Reader simulation record

- Optical methods and filter comparisons
- Gas and liquid transport methods
- Validation plan and simulation files

Optical methods and filter comparisons

Full optical model, maps and numerical checks

Geometry and transport assumptions

The photon-packet formulation uses absorption weights, Fresnel interfaces and Henyey–Greenstein scattering, established transport ingredients described by Wang et al. The AeroSense solver implements its own rectilinear geometry and scoring rather than running MCML. [2]

The solver follows packets across **rectilinear material boundaries**, with Fresnel reflection and total internal reflection. It does not trace arbitrary STL triangles. The four well cavities are 4 × 2.8 mm; the floor is at Z = 43 mm and the 100 µL meniscus at Z ≈ 51.93 mm. LED centers are (±3.3, ±5.7) mm; each has a 1 × 0.3 mm Lambertian emitting area and no lens.

The PD is a 3 × 2.5 mm equal-area proxy for a 7.5 mm² sensitive area, centered at (0, −0.4, 40.5) mm. The 5.2 × 4.2 × 1 mm filter is a homogeneous absorptive proxy. Refractive indices are 1.0 for air, 1.333 for liquid and 1.5 for resin; resin scattering and parasitic absorption are omitted. Actual detector geometry, filter angle response and printed-resin optics need measurement.

Collected power is scored as **LED power × absorbed fraction × quantum yield × (470 / 520) × collection fraction**. This separates excitation absorption from the probability that a fluorescence packet reaches the finite detector.

Table S1. Monte Carlo sampling budgets for the optical sweep, floor irradiance, detector collection map and additional filter study. Variation among seeds estimates numerical sampling uncertainty; it does not calibrate the physical instrument.

Calculation	Sampling	Interpretation
Main optical sweep	24 conditions × 3 seeds; 24,000 excitation + 60,000 emission packets / seed	6,048,000 packets; seed variation estimates Monte Carlo sampling uncertainty
Floor irradiance	600,000 excitation packets	Spatial illumination, not a measured light field
Collection map	224 positions × 65,536 isotropic-source packets	Finite detector acceptance, not a detector calibration
Additional LED / filter study	12 conditions × 3 seeds; 3,024,000 packets	Per-LED excitation rejection assumptions

Main result figure: The optical parameter sweep identifies which assumed inputs dominate the conditional collection estimate.

Spatial illumination and source-well mixing

The floor maps in Main result figure and the transfer matrix in Main result figure resolve different questions: where the modeled light reaches the cell layer, and how fluorescence from each well contributes to an LED-tagged channel.

Interpret the power and spatial results separately

For the same reference assumptions, LED-on collected powers are NE 0.593, NW 0.723, SW 0.829 and SE 1.008 nW. These are not the recovered AC amplitudes: at 25% duty, the time average is 0.25 times the on-state power and the fundamental amplitude is approximately 0.450 times it. The NE floor mean upward irradiance is 4.366 μW/mm² with CV 46.2%; this is neither all-direction fluence nor a prediction of cell distribution.

About 32% of collected NE fluorescence reflects after emission and remains useful signal. Artificial detector acceptance cones retain only 19.3% within 45° and 64.8% within 65°. The photodiode's quoted half-sensitivity angle is not a hard cutoff, so material and angular characterization are needed before converting this model to a detector specification.

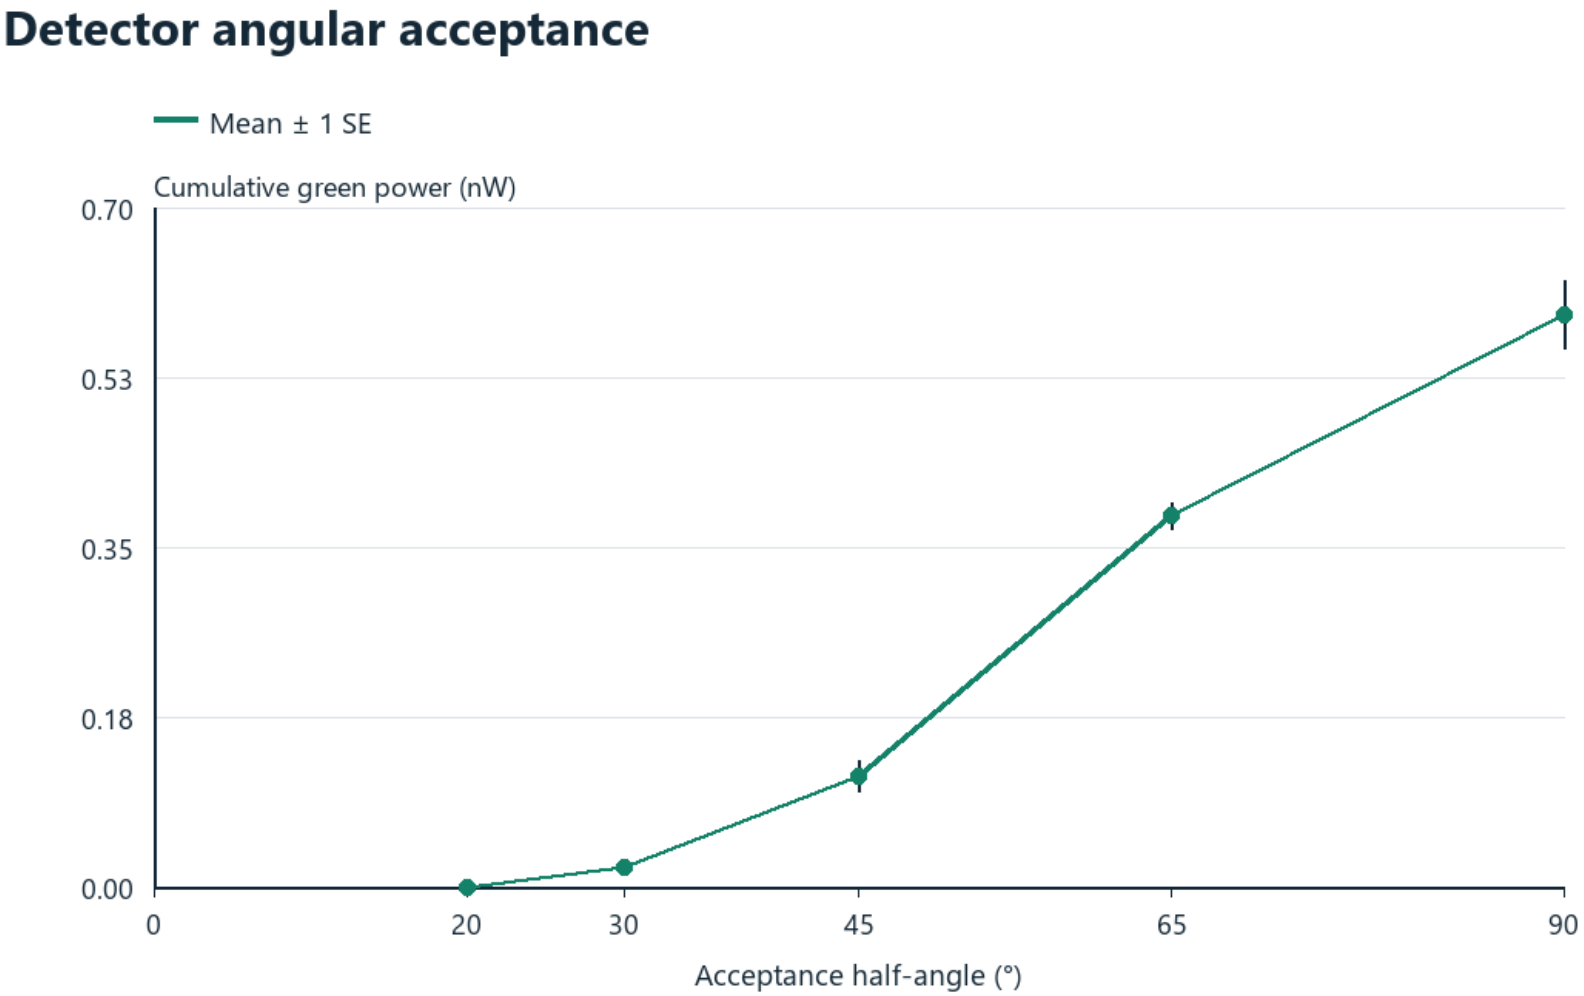

Figure S1. Detector-angle and geometry sensitivity under the idealized transport model. Do not interpret a datasheet half-sensitivity angle as an acceptance cutoff.

Numerical checks · conservation, uncertainty and incomplete paths

The main sweep contains 6,048,000 packets. Reported uncertainty is the three-seed mean ± one Monte Carlo SE, not a 95% confidence interval or a bound on real-device uncertainty. The 224-position collection map has a median relative SE of 17.3%.

Analytic solid-angle, normal-incidence Fresnel, total-internal-reflection, scattering and material-location checks accompany the implementation. Maximum energy-ledger discrepancy is 2.9×10^{-15}. At the 2,500-event limit, up to 0.0117% of paths remain unfinished; the worst conservative unresolved green-power bound is approximately 0.121 nW in an individual run. Extra NE runs compare 2,500 and 10,000 events; they do not establish global convergence of every geometry.

Validation record — Zenodo (publication pending) · Publication checks · Complete numerical interpretation

Why a fluorescence standard uses a different model

The homogeneous-volume calculation represents sodium fluorescein distributed throughout the well. Its row-normalized matrix differs from the thin bottom-layer cell model because the emitting volume differs.

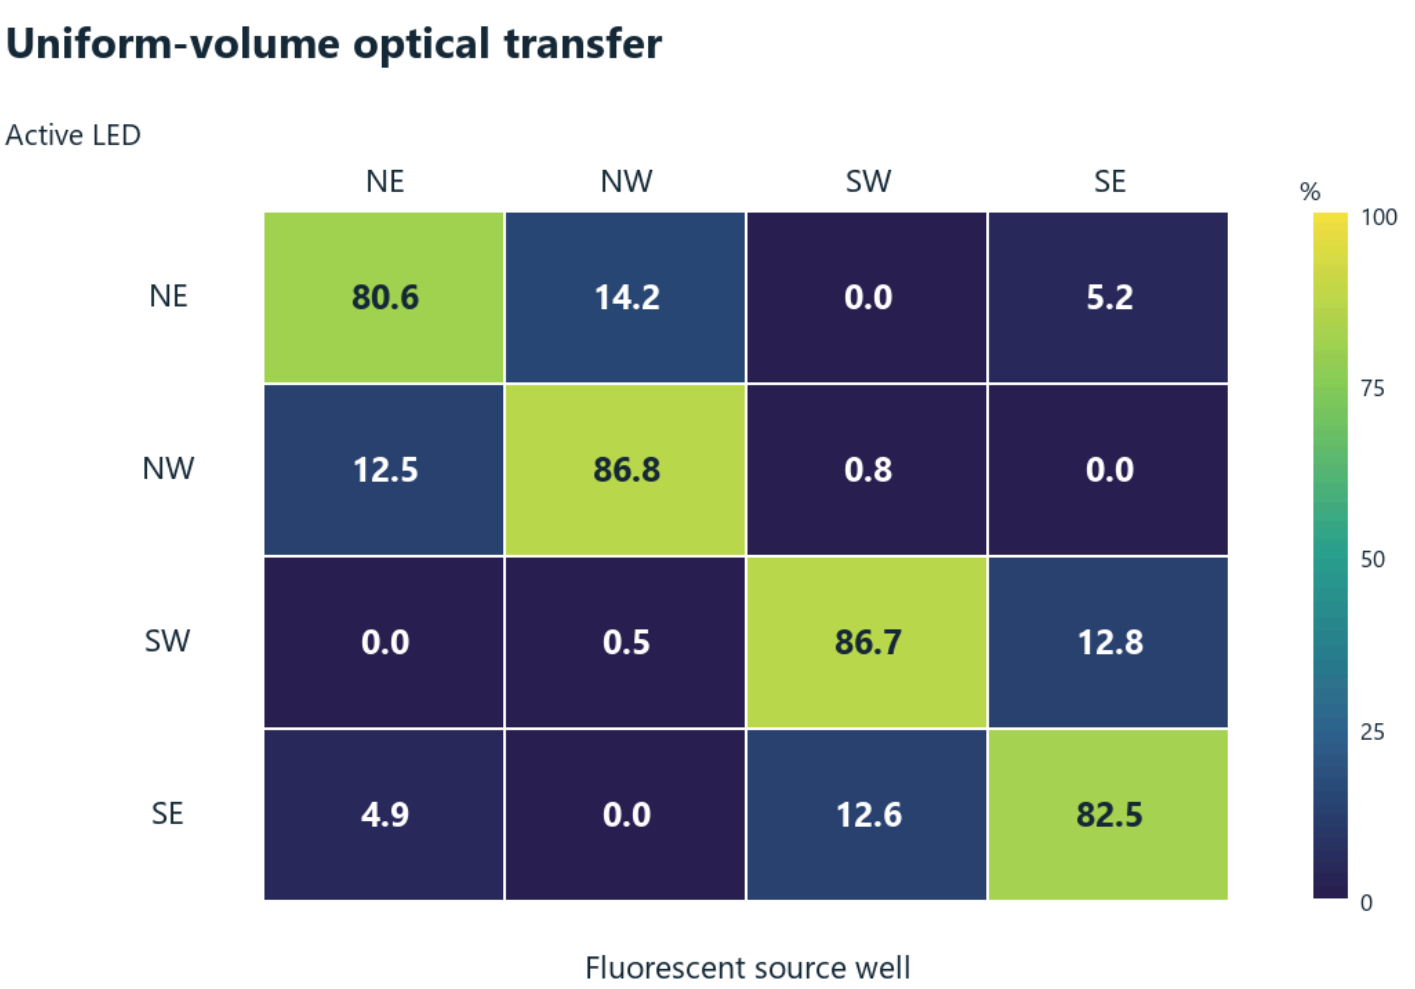

Figure S2. Separate uniform-volume model for the planned cell-free standard experiment.

Standard and later cell-comparison guide

Complete methods and assumptions · Primary numerical summary — Zenodo (publication pending) · Figure-by-figure notes · Hardware decision →

Full spectral assumptions, noise model and filter comparisons

Separate supplied spectra from constructed assumptions

The LEE curve is interpolated from the vendor's sampled transmission data [3]; two layers use $T(\lambda)^2$ without an added interlayer reflection model. The source LED is an assumed Gaussian centered at 468 nm with 25 nm FWHM. The green proxy has its mode at 515 nm and half-Gaussian widths of 10 and 25 nm; it is not a measured GCaMP6f spectrum. The detector uses a digitized typical responsivity curve, not a calibration of our device.

For Edmund #29-190, the manufacturer gives band specifications, not the pointwise curve used here [4]. The explicit scenario is 10^{-4} through 482.7 nm, rises logarithmically to 50% at 500 nm, then linearly to 90% at 517.7 nm. Average OD ≥ 4 in a specified rejection band must not be read as uniform pointwise blocking across the LED tail. Varying the assumed edge gives retained LED current of 0.0104–9.36% and green current of 71.8–90.3%; these are scenario ranges, not confidence intervals.

Main result figure makes the trade-off explicit: blue leakage and useful green light are weighted by different source spectra and the same detector response.

Use one geometry and noise protocol across the four cases

The unfiltered Monte Carlo input comprises four LED cases × three seeds × (24,000 excitation + 60,000 emission packets), or 1,008,000 primary packets. The current unfiltered green powers are 0.7848 ± 0.0723, 0.6559 ± 0.0685, 0.8228 ± 0.0501 and 0.9558 ± 0.0772 nW for NE/NW/SW/SE, respectively (mean ± one Monte Carlo SE). Blue powers are 110.49 ± 13.81, 62.26 ± 23.85, 123.77 ± 20.89 and 124.20 ± 12.03 nW. These finite-sampling uncertainties are appreciable and are separate from the old filtered 0.593 nW reference.

Each 2 s trace contains 8,000 samples at 4 kHz and 25% duty, with NE/NW/SW/SE tags of 89/113/41/67 Hz. Assumptions include electronic noise of 1.5 pA RMS per sample, 0.2 nA dark current, and ambient power of 0.15 nW plus a 0.05 nW sinusoid at 100.3 Hz, distributed uniformly across 500–650 nm. Poisson counts use 250 µs bins. Signal seed 20261011 and blank seed 20262011 are recorded; five blank records are averaged. Cases share electronic-noise draws and Poisson seeds, but different expected photoelectron counts give different shot-noise draws.

The offline joint fit keeps only the four AC fundamentals. It omits DC, PWM harmonics, analog transfer, ADC clipping and timing jitter. Sixty-four noise realizations test the analytic noise calculation; they do not propagate Monte Carlo geometry, spectrum, material or biological uncertainty. The displayed corrected trace is not a reconstruction of every raw frequency component.

Figure S3 shows the raw signal and corrected AC on their own scales; overlapping curves under an ideal blank are not evidence of instrument validation.

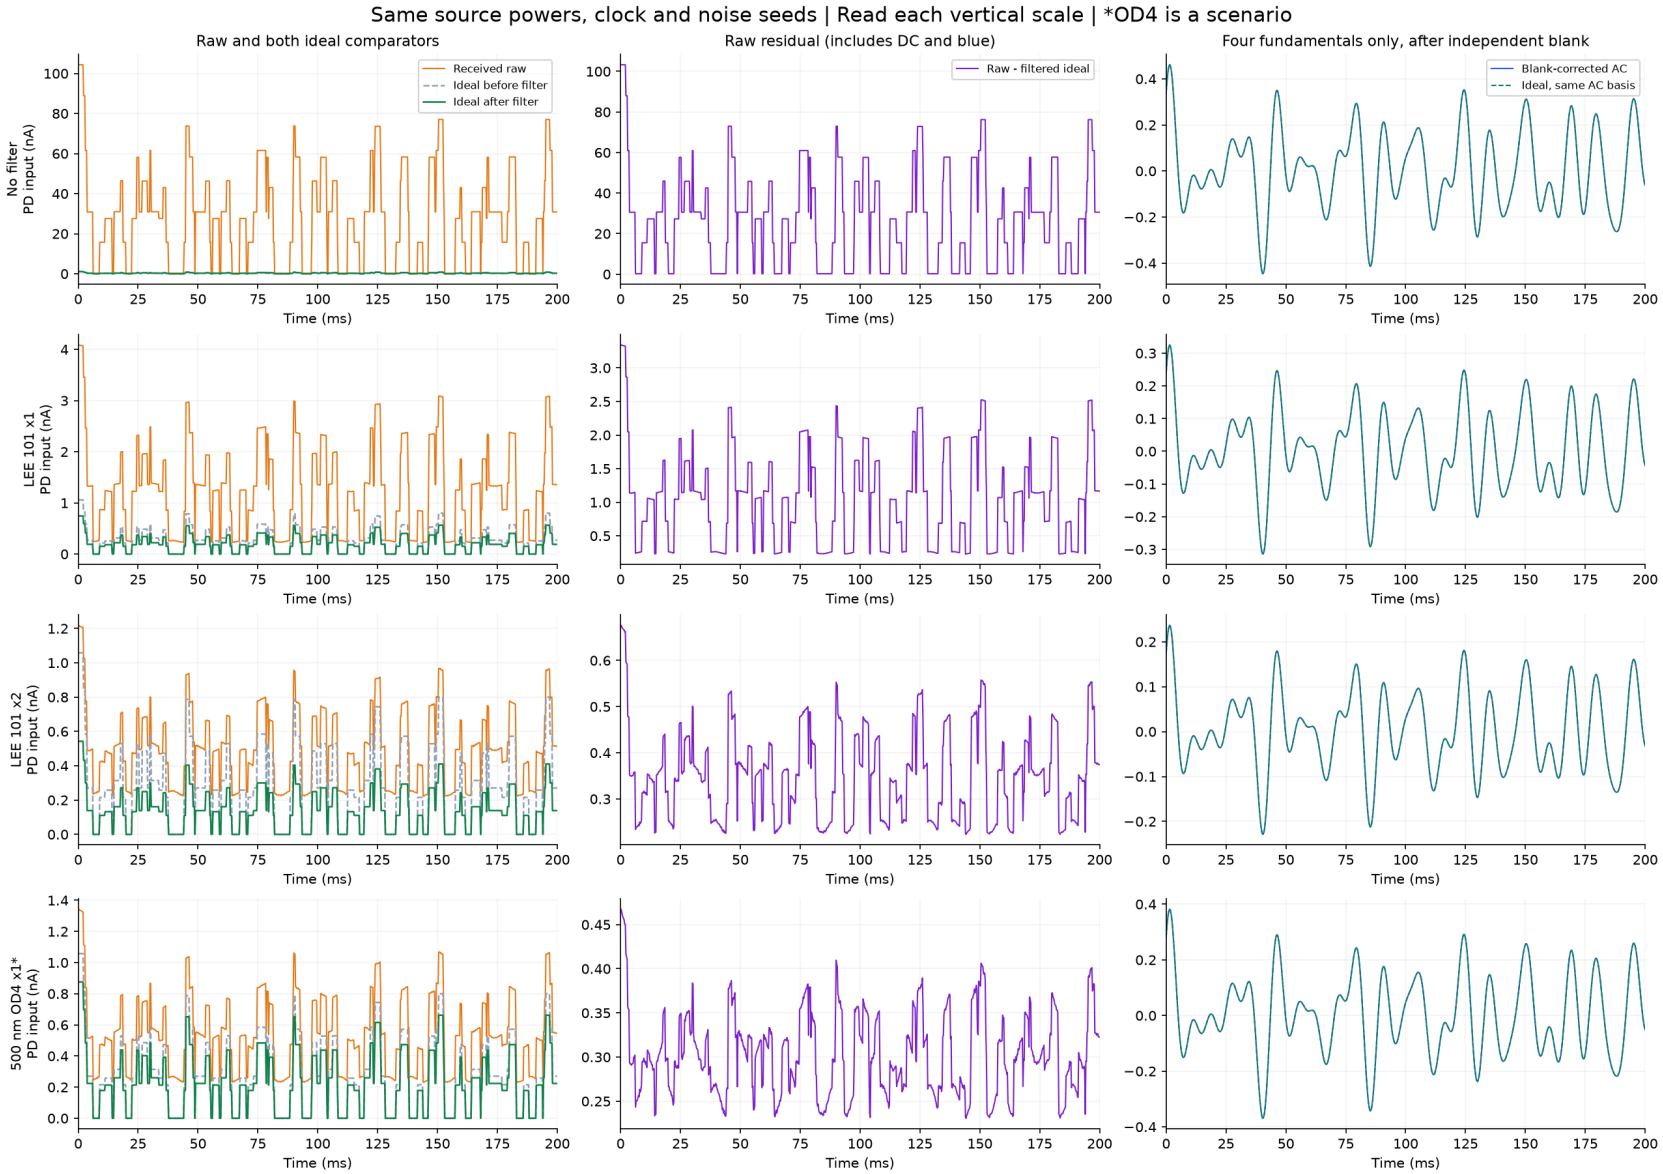

Figure S3. Simulation · first 200 ms of each 2 s record. Columns are raw/ideal current, raw residual including DC and excitation leakage, and four-fundamental blank-corrected AC. The blank shares the deterministic leakage with the sample. Fixed seeds are recorded above; no uncertainty bands or measured hardware traces are plotted. Source methods and uncertainty.

Ask how much the blank can change

Worst raw complex leakage relative to green is 11,210.32%, 475.46%, 89.55% and 26.20% for no filter / LEE×1 / LEE×2 / OD4 scenario. Under the ideal stable blank, residual complex errors are 0.1341%, 0.0704%, 0.0849% and 0.0534%. This favorable correction depends on preserving the deterministic background; it does not remove its shot noise.

In Main result figure, a 0.1% blue-gain change after blank acquisition gives worst residual biases of about 11.2105%, 0.47553%, 0.08963% and 0.02625% of the green signal. Compared with a 1% fluorescence change, those are 11.21×, 0.476×, 0.0896× and 0.0263× the change. This is a sensitivity calculation; no real-device drift was measured.

Earlier monochromatic photon-budget target

This separate reference calculation predates the spectral filter comparison. At 500 µW LED-on power, quantum yield 0.6 and a 10% fluorescence change, the stated photon-budget model predicts a **16.03 pW lock-in amplitude change**. A chosen SNR ≥ 5 target would require repeat lock-in RMS below **3.21 pW**. This is an engineering target under those assumptions, not a measured detection limit or an ADC sample-noise requirement.

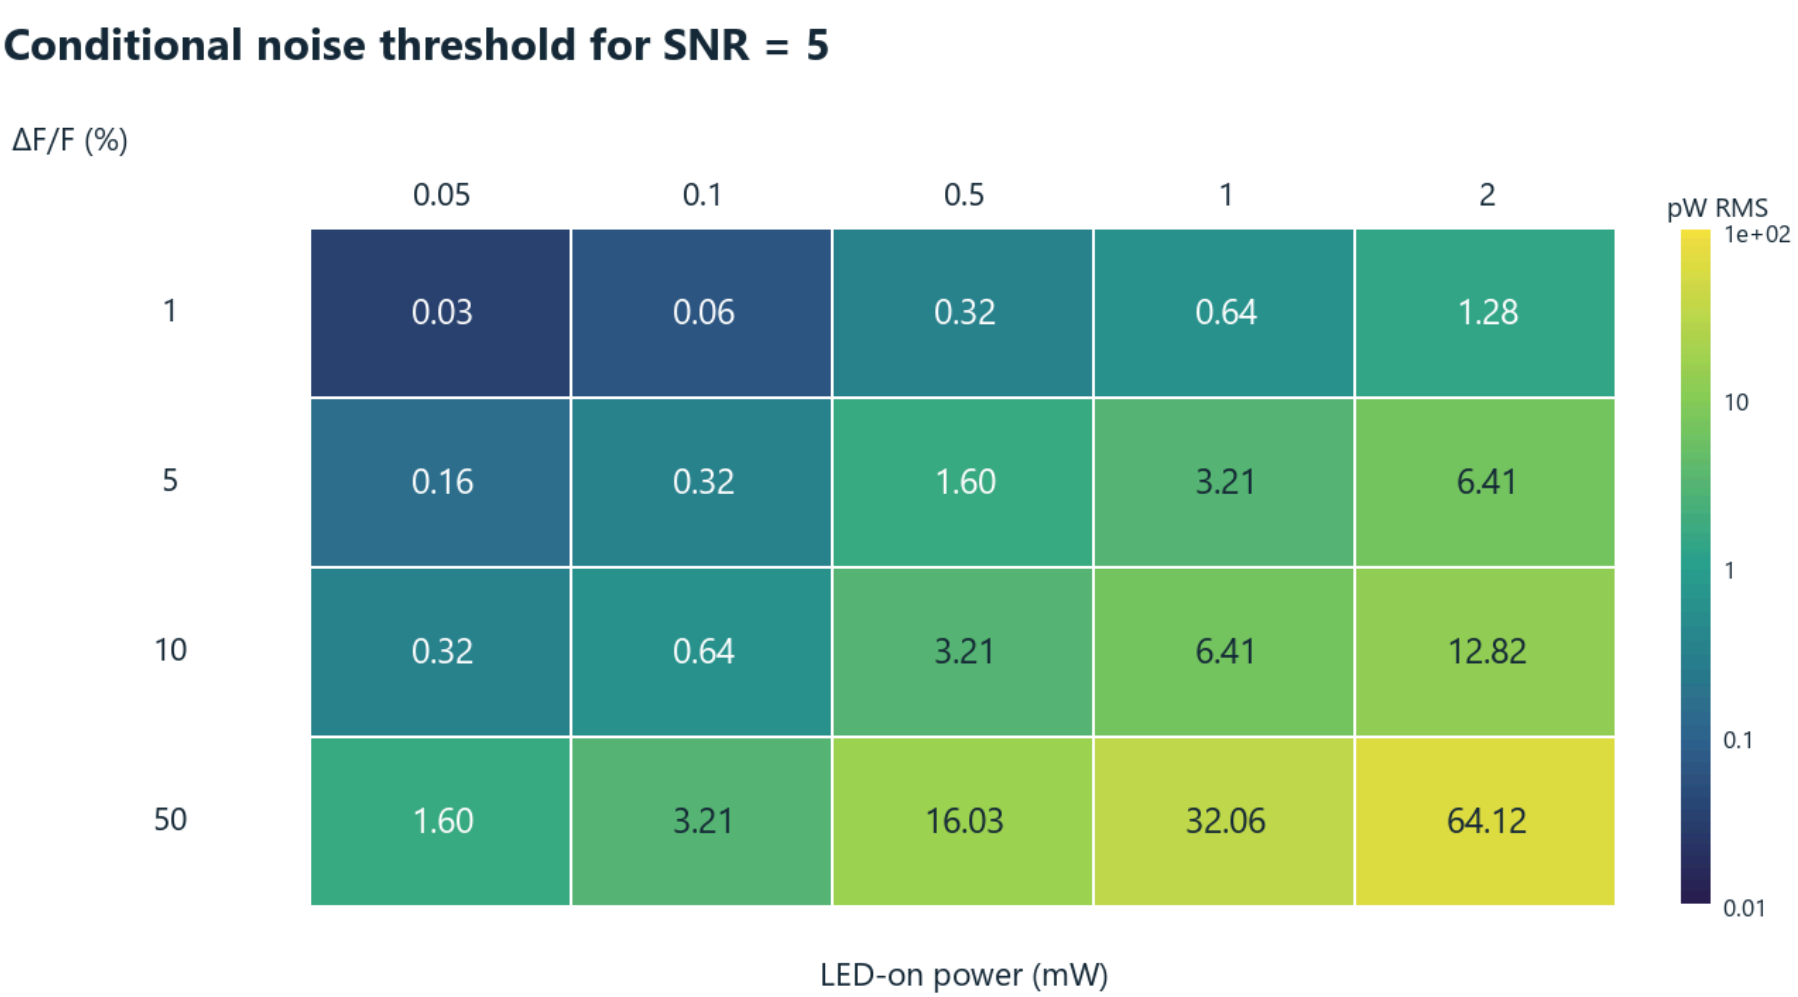

Figure S4. Photon-budget sensitivity. Changing quantum yield, optical power or fluorescence change changes the target; actual cells and electronics remain to be measured.

English methods and reproduction guide · Numerical summary — Zenodo (publication pending) · Case-level CSV — Zenodo (publication pending) · Numerical checks — Zenodo (publication pending)

Gas and liquid transport methods

Surface calculation and checks

The confirmed working fill is 100 µL per well. The team estimates a printed brim capacity of approximately 120 µL; the nominal CAD brim remains 140 µL. That capacity estimate does not define a new cross-section, so the calculation retains the existing CAD geometry. On the main 0.4 mm grid, equal-area gas cells centered 0.2 mm above the nominal liquid surface supply each surface mean.

The saved D3Q19 lattice-Boltzmann velocity fields supply a conservative finite-volume advection–diffusion calculation. Velocity is scaled to the sampled flow; the scalar cases are not independent velocity solves at each flow. The liquid boundary has zero uptake. Sample-gas flow is independent of the PCB cooling fan.

Table S2. Nominal first-gas-layer arrival times: Q = 3 mL/min, D = 10 mm²/s, 100 µL per well, 0.4 mm grid and 0.25 s time step. Symmetric pairs follow the nominal geometry.

Well	t10 (s)	t50 (s)	t90 (s)
NE	6.29	18.24	49.42
NW	5.30	16.66	47.80
SW	5.30	16.66	47.80
SE	6.29	18.24	49.42

The current first-layer calculation contains 26 cases: 21 selected-shelf combinations of seven flows and three gas diffusivities (5, 10 and 20 mm²/s), two geometry controls and three numerical checks. Time zero is an inlet concentration step. t10 indicates the rise of the surface mean, not first molecular contact; external tubing delay is excluded.

The shelf is retained for directional flow control and a simple integral structure, rather than the fastest modeled arrival. At 0.4 mm, the downward momentum ratios for open, original plate and shelf are 1.00, 0.57 and 0.69. This approximately 31% shelf reduction is a gas-flow proxy, not cell shear or biological safety.

All 26 cases pass concentration-bound, uniform-field and global mass-balance checks; the largest relative mass residual is 9.13×10^{-12}. At 2.5 mL/min, the worst t90 is 60.06 s, within the tested numerical sensitivity of a 60 s target. The 3 mL/min choice provides timing margin, not a unique optimum.

The surface movie interpolates saved one-second concentration maps without extrapolation. Geometry and flow comparisons each span 0–180 physical seconds. Its two chapters use the same color scale and reset the physical clock. The three-dimensional movie uses these scalar fields with saved velocity fields; marker density does not encode concentration.

Surface histories, thresholds and numerical checks (CSV) — Zenodo (publication pending) · Full airflow methods and reproduction guide

Full gas-flow methods, geometry sweeps and wash-in checks

Velocity model: geometry is the variable

A double-precision D3Q19 BGK lattice-Boltzmann solver screens 18 geometries in 28 original velocity solves: the open plenum, original 12-hole plate, vertical walls, L-shaped hoods and horizontal shelves. Air density is 1.2 kg/m³ and viscosity 1.8×10^{-5} Pa·s. Solid walls and the stationary liquid surface use halfway bounce-back; pressure terminals use non-equilibrium extrapolation. [5] [6]

The BGK collision model follows Qian et al.; the pressure boundary reconstructs equilibrium at the imposed density and copies the adjacent fluid node's non-equilibrium part, as in Guo et al. The scalar advection–diffusion calculation below is a separate finite-volume solver.

Saved fields are Stokes-normalized to 1 mL/min. Convergence requires velocity change below 2×10^{-5} over 200 steps and mass imbalance below 0.2% for three consecutive checks. All 28 runs meet the checks; maximum reported imbalance is 0.0444%. A separate near-1 mL/min solve confirms the directional trend.

The momentum proxy is evaluated 1 mm below the well rim; gas-speed P95 is sampled 0.6–1.4 mm above the liquid. These are **gas-side quantities**, not submerged-cell shear, biological response or tissue-injury criteria.

The shelf improves one metric, with a tradeoff

The original geometry screen in Figure S5 retains the rejected alternatives so the shelf choice can be inspected alongside its tradeoffs.

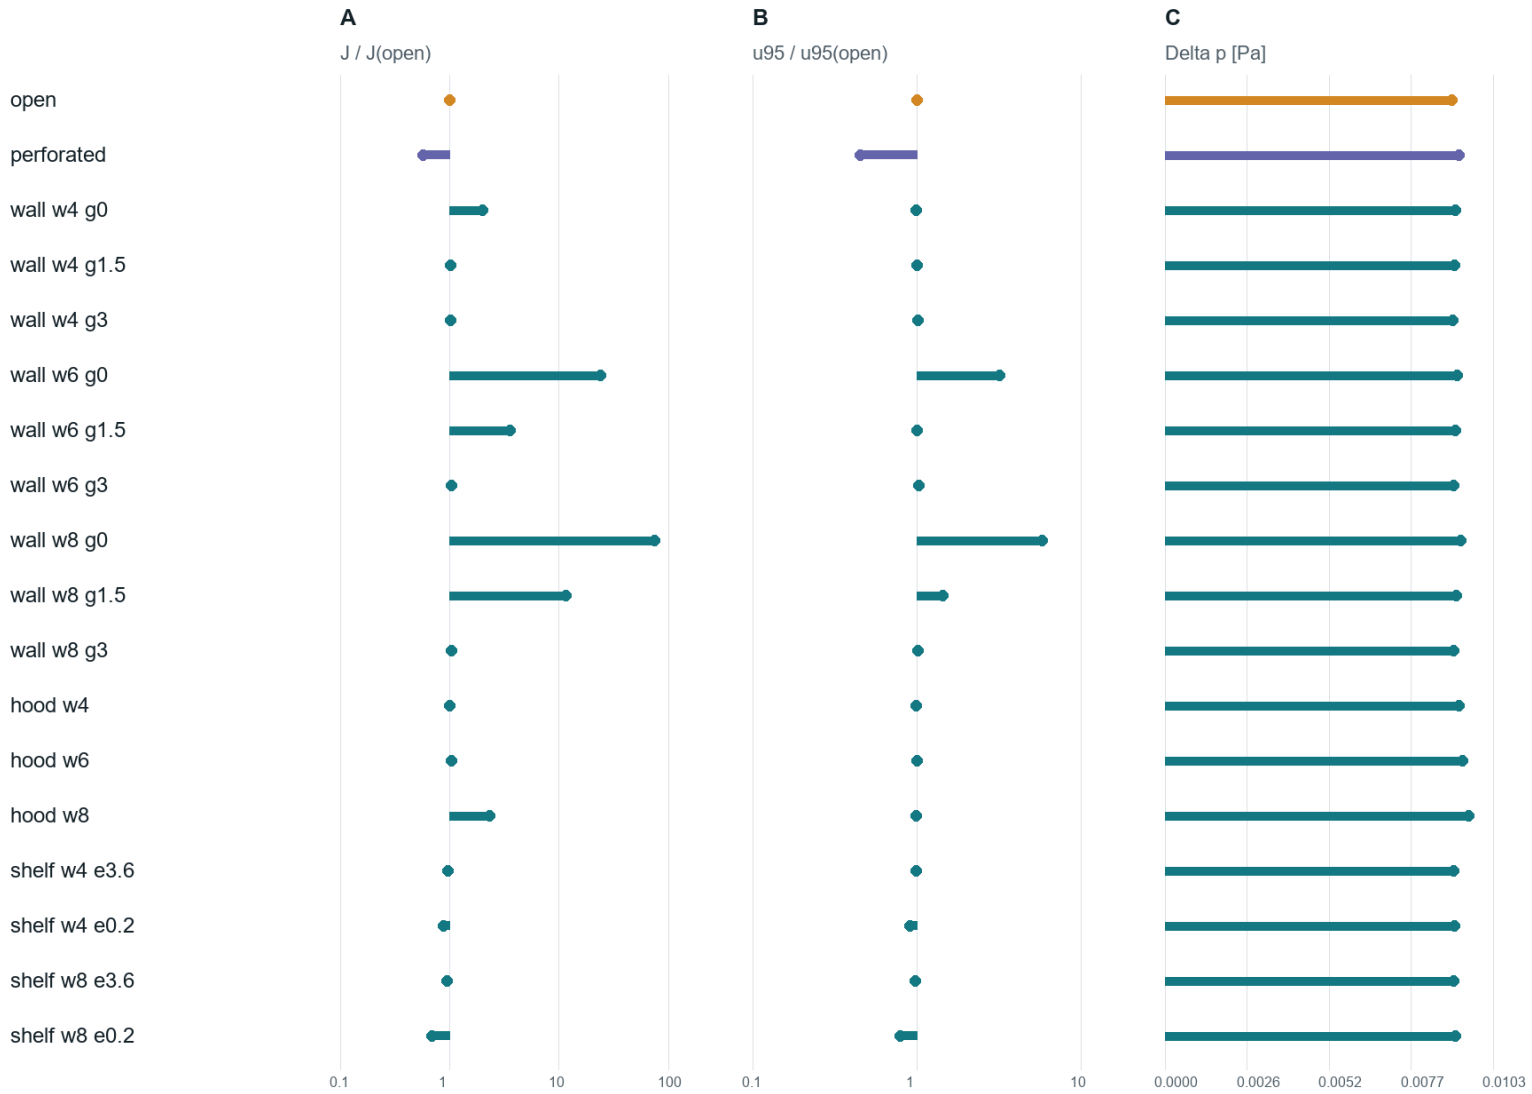

Figure S5. Simulation · gas-lid geometry screen. This figure compares 18 deterministic geometry cases on the original 0.4 mm grid: worst downward-momentum proxy, near-liquid P95 speed normalized to the open design, and pressure loss. Cases are design alternatives, not statistical replicates; no Monte Carlo seeds or error bars apply. The separate shelf/open refinement gives momentum reductions of 30.9%, 30.1% and 36.1% at 0.4, 0.3 and 0.2 mm respectively, so grid convergence is not established. All cases use the nominal CAD gas space with a 100 µL fill. Source methods and uncertainty.

Hardware: the integral shelf and the provisional flow setting →

Scalar model: gas concentration is a separate calculation

A conservative cell-centered finite-volume advection–diffusion solver uses backward Euler in time, upwind advection and central diffusion. The inlet steps from zero to normalized concentration 1; the outlet is convective with no diffusive flux. Walls and liquid are impermeable in this gas-only calculation. Projecting face fluxes to a divergence-free field changes the baseline velocity flux by 1.54%.

Three assumed gas diffusivities, 5/10/20 mm²/s, test molecular-transport uncertainty. They are not pump settings. Near-liquid concentration is averaged separately for each well; the four-well worst t90 is used in the operating comparison.

Advection–diffusion then estimates gas concentration near the liquid. The 1 mL/min, D_{gas} = 10 mm²/s NE baseline has t90 = **149.55 s**. Velocity markers are not molecular concentration, and neither observable measures bottom-cell exposure.

Figure S6: The gas-delivery curves quantify arrival in the nominal gas domain; they do not establish delivery to cells.

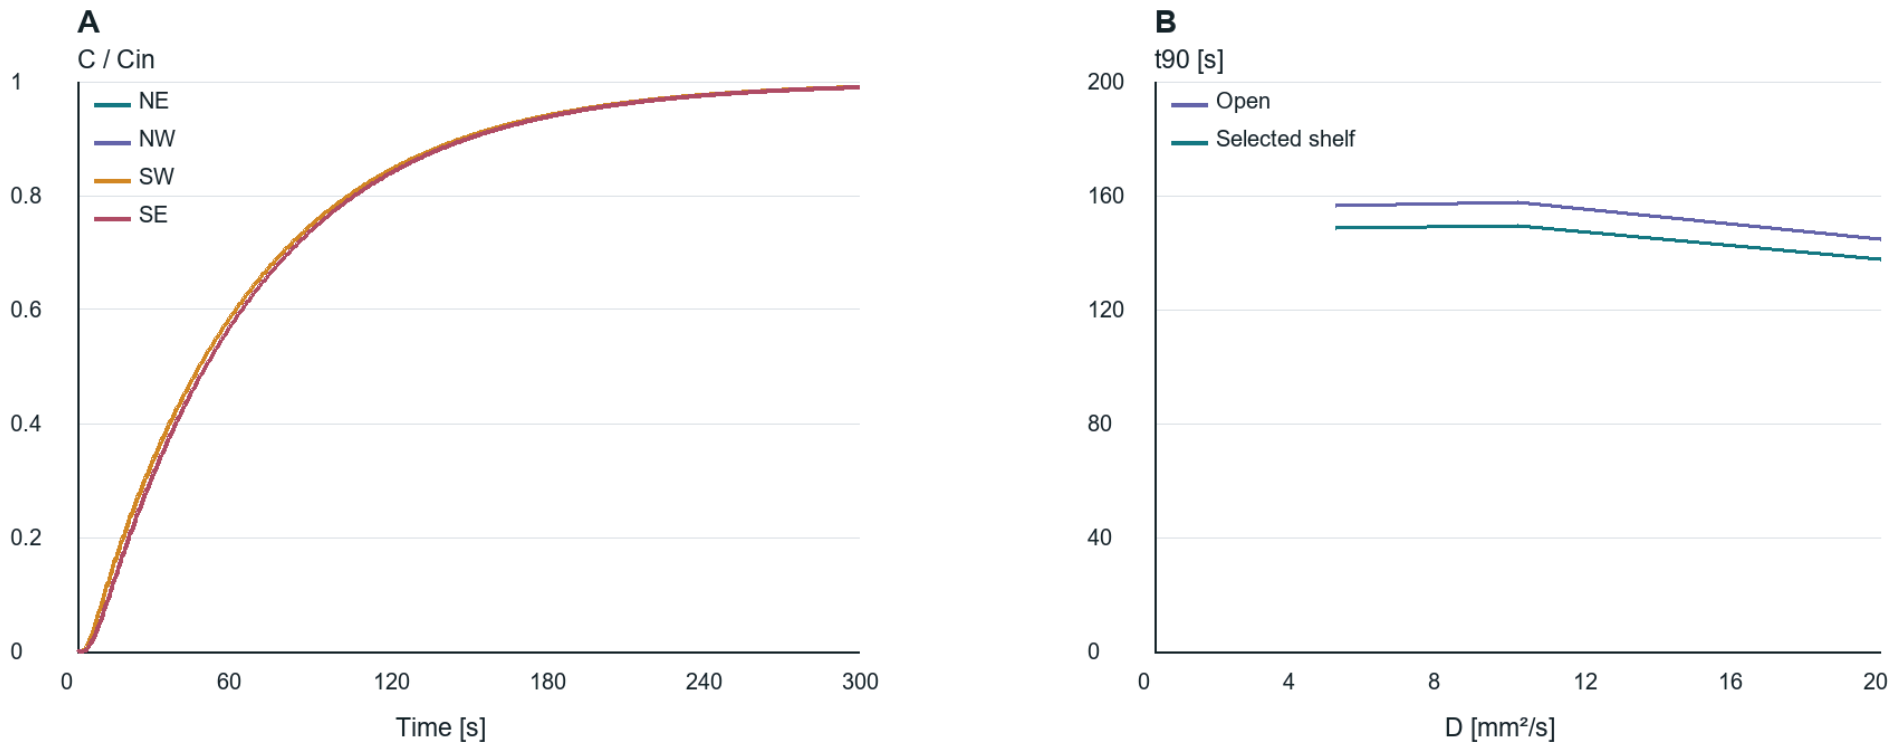

Figure S6. At 1 mL/min and Dgas = 10 mm²/s, the NE baseline t90 is 149.55 s. Gas concentration at the liquid interface is distinct from concentration at bottom-attached cells.

Choose a sampled operating point against declared criteria

The operating-band study evaluates 28 combinations: 0.25/0.5/1/2/2.5/3/5 mL/min at each of four diffusivity assumptions (5/10/15/20 mm²/s). It samples gas 0.6–1.4 mm above the liquid, separately from the current first-layer calculation. The comparison requires worst-well gas t90 ≤ 60 s, port Reynolds number ≤ 3 and near-liquid P95 gas speed ≤ 0.01 mm/s. The speed cap is an engineering comparison criterion, not a measured cell-injury threshold.

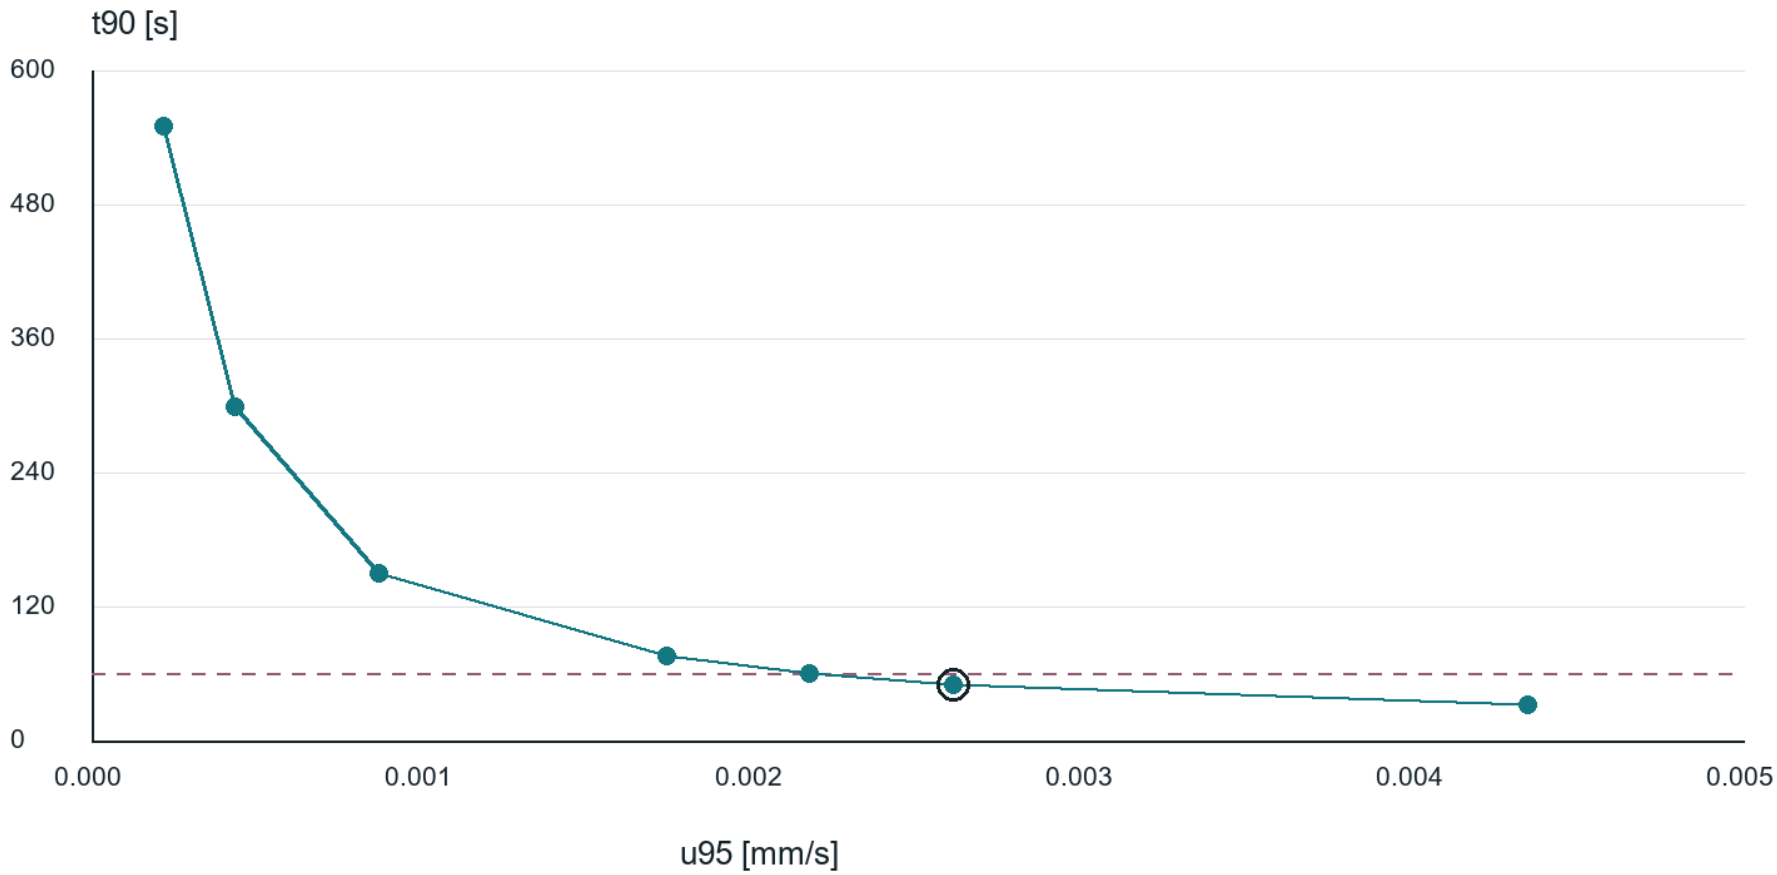

Figure S7. Fixed-CAD flow/diffusivity study. Gas diffusivities 5/10/15/20 mm²/s are uncertain material assumptions; the candidate is the lowest qualifying sampled flow, not a global optimum.

The full time-to-arrival screen in Figure S9 shows why the operating candidate is conditional on geometry and diffusivity.

The ledger includes 28 original velocity solves (18 geometries), 18 hole conditions plus 2 refinements, 9 baseline scalar conditions plus an analytic benchmark, and 28 operating points plus 6 scalar checks and 1 direct velocity confirmation. Solver counts are not experimental replicates. The supplemental ledger — Zenodo (publication pending) and operating-point record — Zenodo (publication pending) give the exact accounting.

The hardware gas-path specification uses one regulated downstream pump. Vergara et al. (2012) provides the separate metal-oxide dataset context for recording delivery conditions and time/batch metadata. Vergara et al. (2012) — Zenodo (publication pending) Hardware implementation → Operating method and assumptions

Explore the 28 computed operating points

Rotate the saved flow-rate/diffusivity study and inspect its values. Gas diffusivity is an uncertain material assumption, not an operator control. The speed cap is an engineering comparison value, not a cell-injury threshold.

Open document in a new tab ↗ — Zenodo (publication pending)
Explore the physical simulations on the Model page.

The document link above opens the complete preview.

Geometry sweeps and numerical verification

The shelf/open momentum reduction is 30.9%, 30.1% and 36.1% on 0.4, 0.3 and 0.2 mm grids. Local flow metrics are therefore not grid independent. A separate 18-condition six-hole sweep varies diameter (1.2/1.8/2.4 mm), pitch (3.6/4.2/4.8 mm) and plate height (57.9/58.2 mm). Its smallest momentum ratio changes from 0.0246 to 0.0351 on refinement. Narrow holes also introduce fabrication and cleaning tradeoffs; the current STLs retain the shelf.

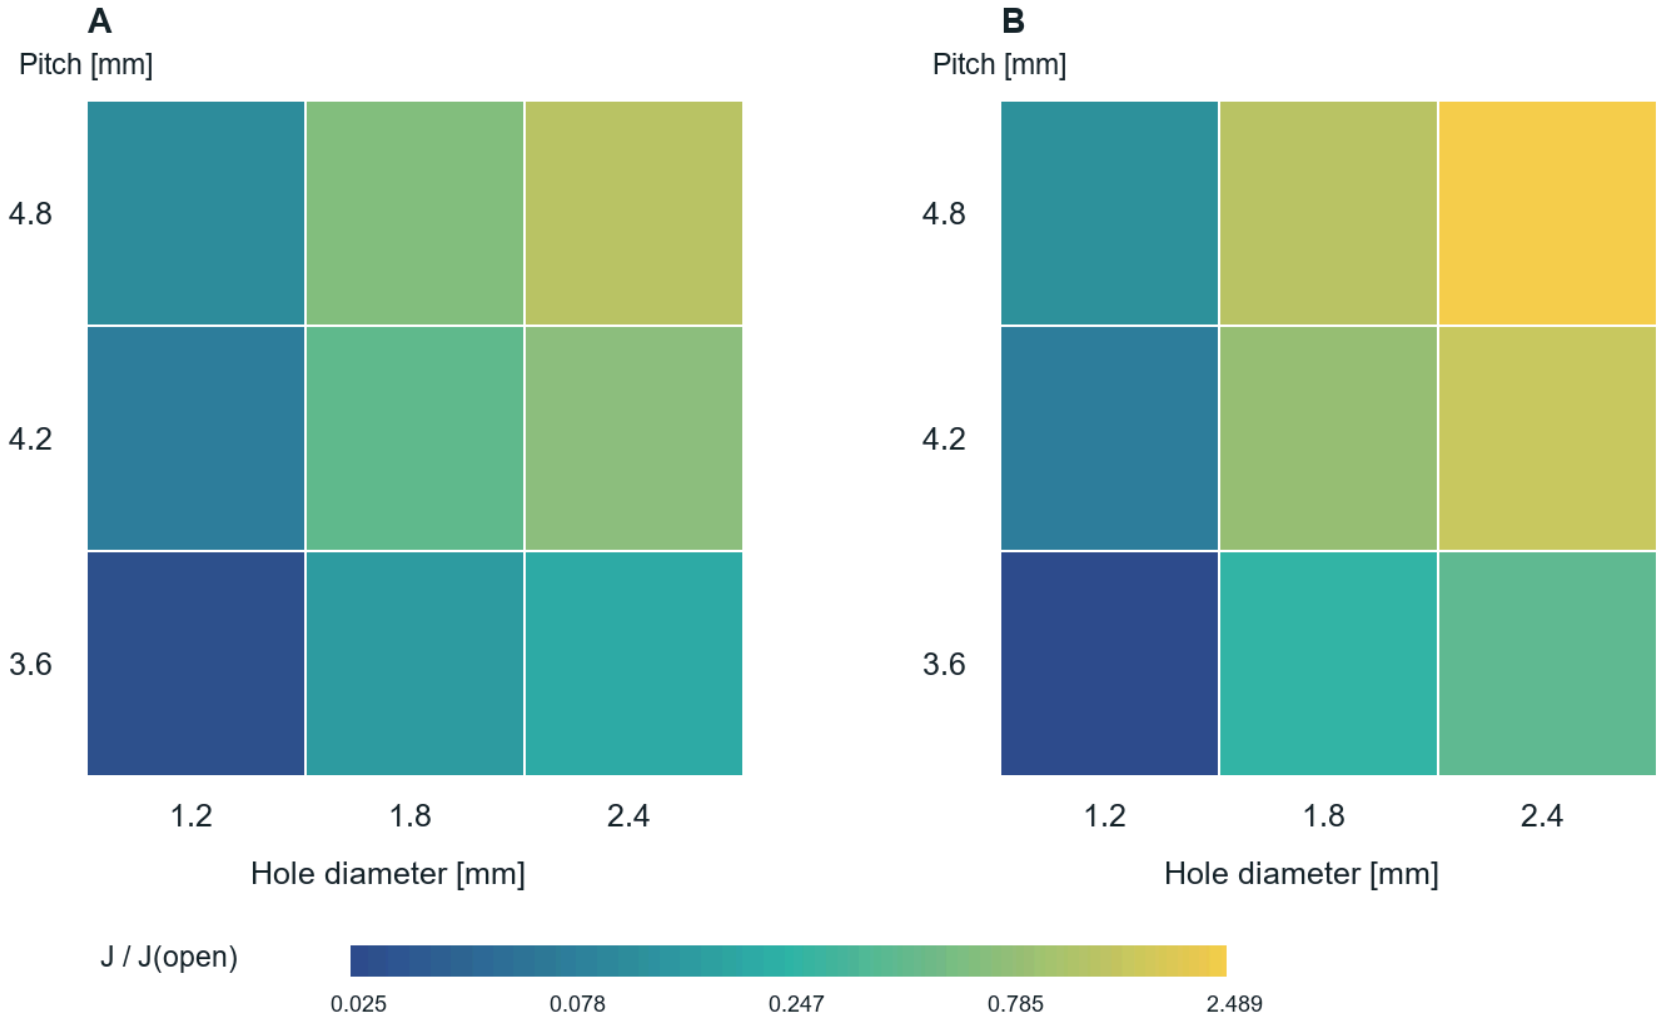

Figure S8. Exploratory six-hole sweep. This is a comparison set, not an alternate selected print or proof of a global optimum.

The C# and NumPy velocity implementations differ by a relative L2 norm of 3.09×10^{-12}; the Hagen–Poiseuille benchmark errors are 2.54% at $\tau = 0.8$ and −1.13% at $\tau = 0.55$. For scalar transport, the mass-ledger discrepancy is 8.65×10^{-12} and the one-dimensional analytic mean-concentration error is 0.000246. Baseline gas t90 changes by about +1.7% with grid refinement and −0.2% with a halved time step. These verify numerical implementation, not the actual apparatus.

Full solver methods · Explore the saved hole sweep ↗ — Zenodo (publication pending)

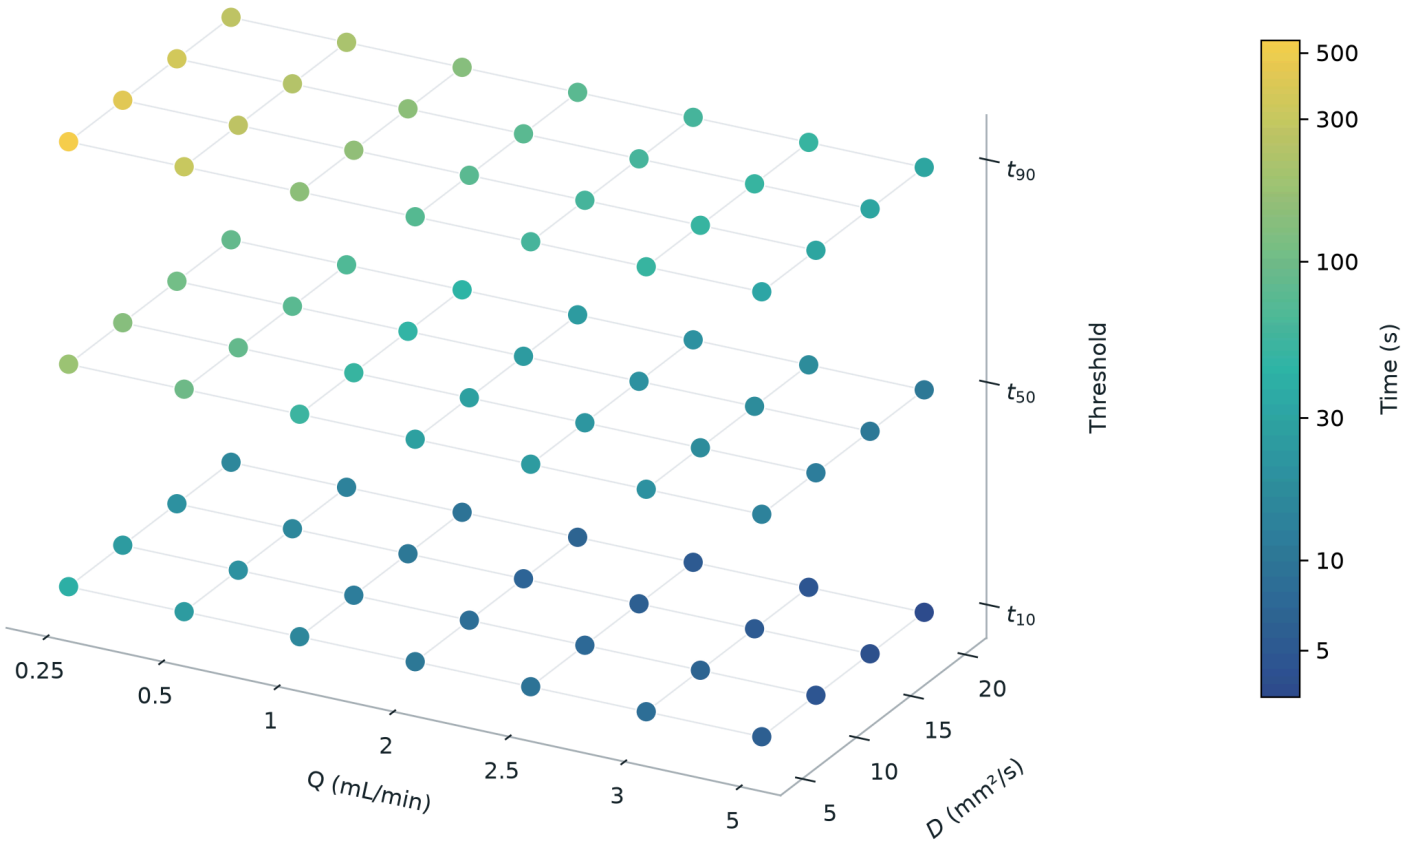

Figure S9. Twenty-eight computed operating-band conditions across seven flows and four assumed diffusivities (5, 10, 15 and 20 mm²/s). The three-dimensional scatter shows worst-well gas t90 in the 0.6–1.4 mm sampling band above the liquid. It is separate from the current first-gas-layer surface calculation.

Open plenum, perforated plate and integral shelfOriginal simulation · 20 s · 2.81 MB

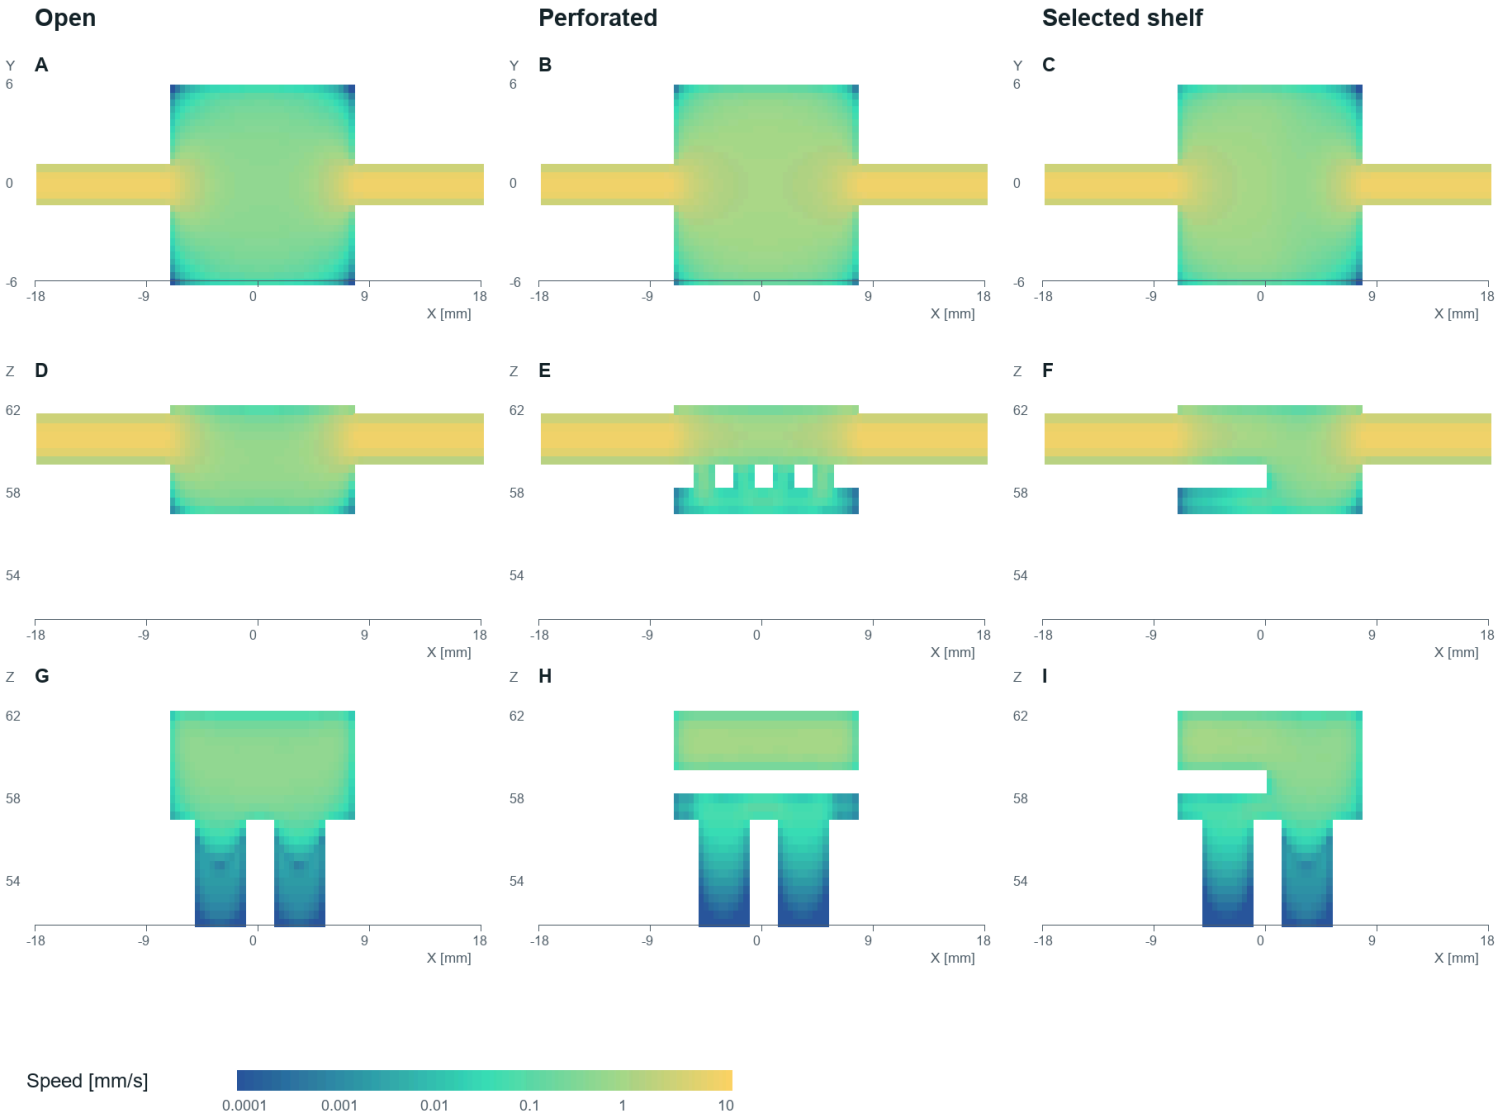

Open the original MP4 — Zenodo (publication pending)

Watch the corresponding airflow animation — VideoUniverse

Figure S10. 60 physical seconds in 20 seconds of playback. Panels A–C advect saved velocity markers in the open, 12-hole and selected-shelf gas geometries at the **1 mL/min baseline**. Marker density is not concentration. Panel D independently plots finite-volume gas concentration for the shelf with D_{gas} = 10 mm²/s. The liquid remains a stationary boundary; no liquid mixing or cell motion is simulated. This is not a movie of the later 3 mL/min candidate. Original MP4 ↗ — Zenodo (publication pending) · Source timing and scope — Zenodo (publication pending).

Open this simulation movie — Zenodo (publication pending).

Table S3. Operating-band comparison at the 3 mL/min candidate. These values sample 0.6–1.4 mm above the liquid; current first-layer results and checks are reported separately above.

Metric	Computed result	Declared criterion / boundary
Worst near-liquid gas t90	50.00 s	≤60 s across four wells and four gas-diffusivity assumptions
Refined grid / half time step	50.98 s / 49.85 s	Numerical checks at the chosen candidate
Near-liquid gas P95 speed	0.00280 mm/s	≤0.01 mm/s engineering comparison cap; not a cell-injury threshold
Port Reynolds number	1.77	≤3 in the declared comparison
Chamber pressure drop	0.02737 Pa	Excludes external filters and tubing
Nearby 2.5 mL/min point	Worst t90 ≈60.04 s	Numerically borderline; not proven physically inadequate

Hardware simulations / 04

Liquid-path diffusion model

Earlier screening model. These idealized liquid and PDMS calculations are retained for comparison; the current gas-arrival calculation ends at the gas side of the liquid surface.

Inputs
: Depth, diffusivity and assumed partition coefficients

Method
: One-dimensional diffusion / coupled layer model

Output
: Cell-plane delivery time and membrane sensitivity

Decision
: Measure cell-plane delivery before adding a film

The three-dimensional gas model stops at the liquid interface. A separate one-dimensional screening model asks how molecules could travel through the gas gap, an optional membrane and the liquid to bottom-attached cells. At the original 100 µL fill, the liquid is 8.929 mm deep and the gas gap below the rim is 3.571 mm. No membrane is selected for the current eight-part build.

Firpo et al. provide the solution–diffusion and surface-resistance context, as well as the oxygen permeability and diffusivity used in the source package's separate oxygen example. Generic VOC diffusivities and partition coefficients in this sweep remain assumptions. [7]

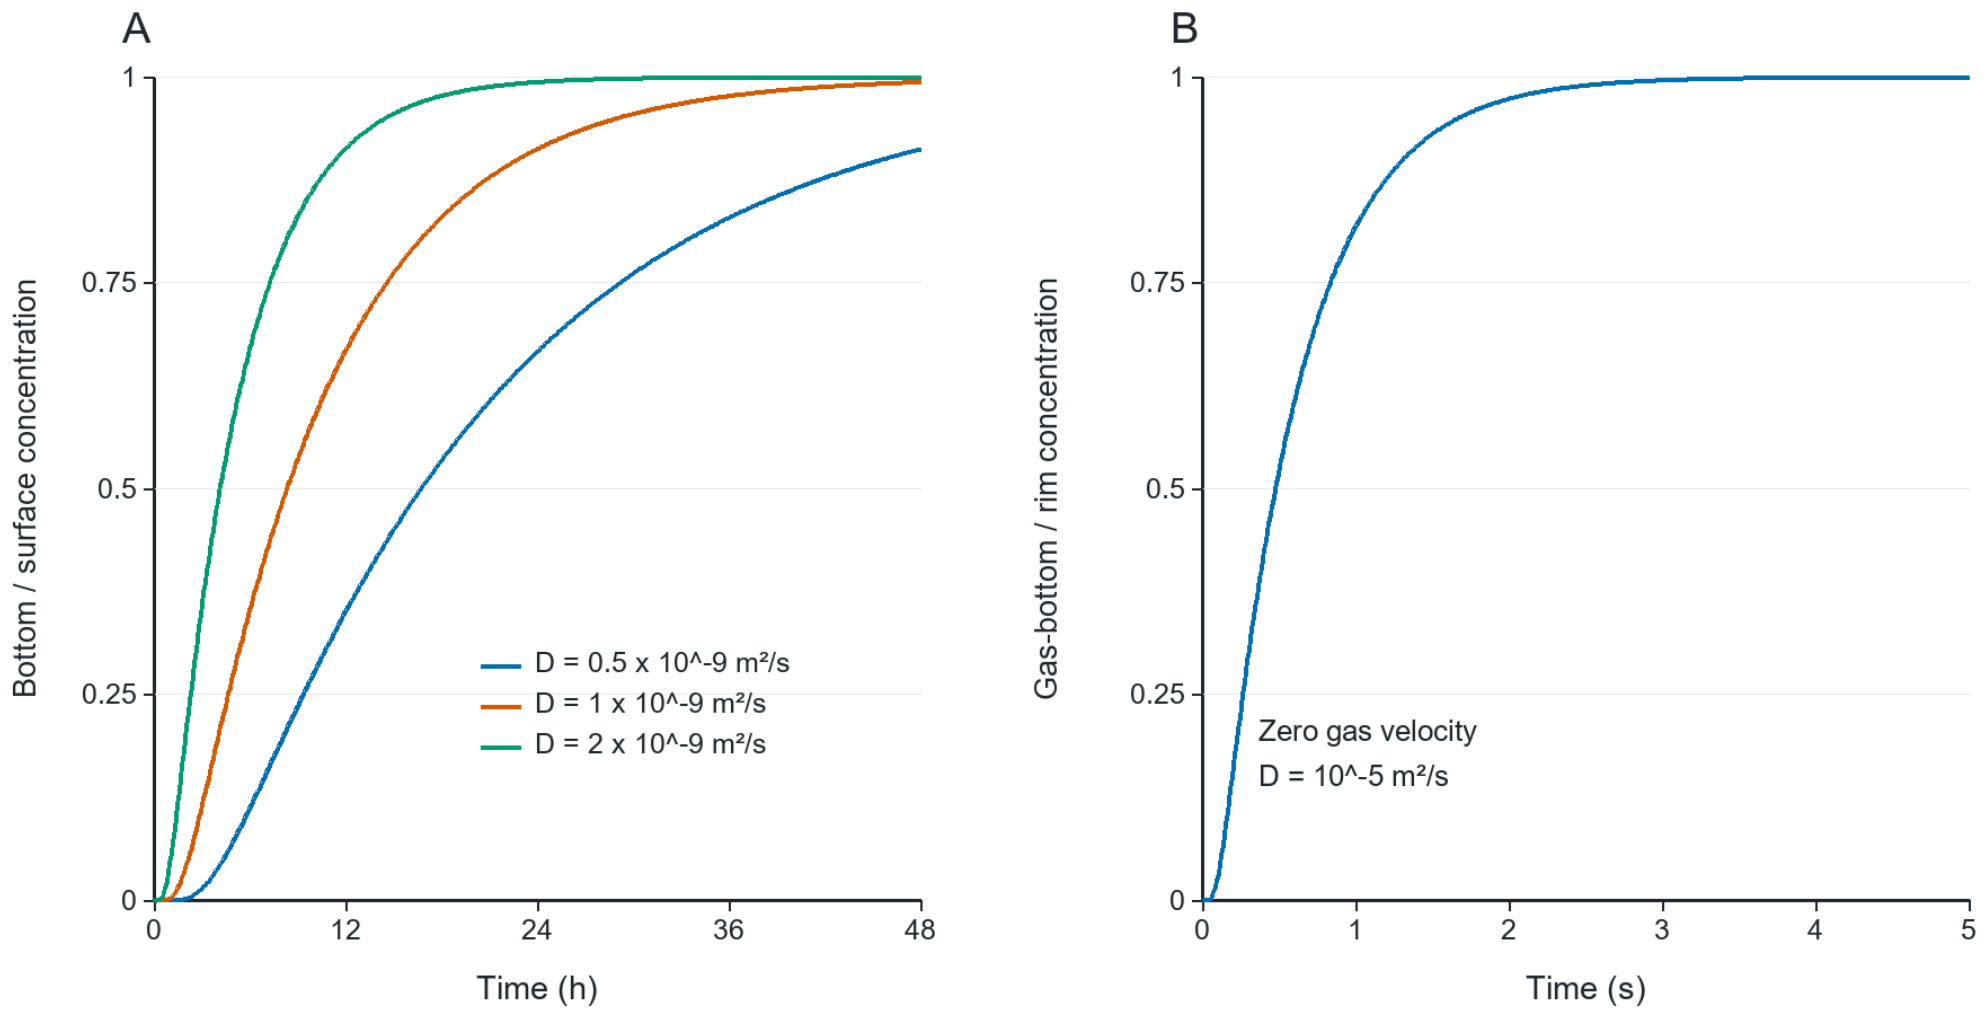

Figure S11. Simulation · gas arrival versus cell exposure. Exact one-dimensional surface-step diffusion is evaluated at three assumed diffusivities, with nominal 8.929 mm liquid depth and a separately modeled 3.571 mm gas gap. The curves are deterministic scenarios, not repeated observations; no sampling seed or error bars apply. At D = 10^{-9} m²/s, the liquid-bottom t90 is 22.83 h. These earlier idealized calculations use the nominal CAD geometry and a 100 µL fill. Rapid gas wash-in therefore does not establish rapid sensing by bottom-attached cells. Source methods and uncertainty.

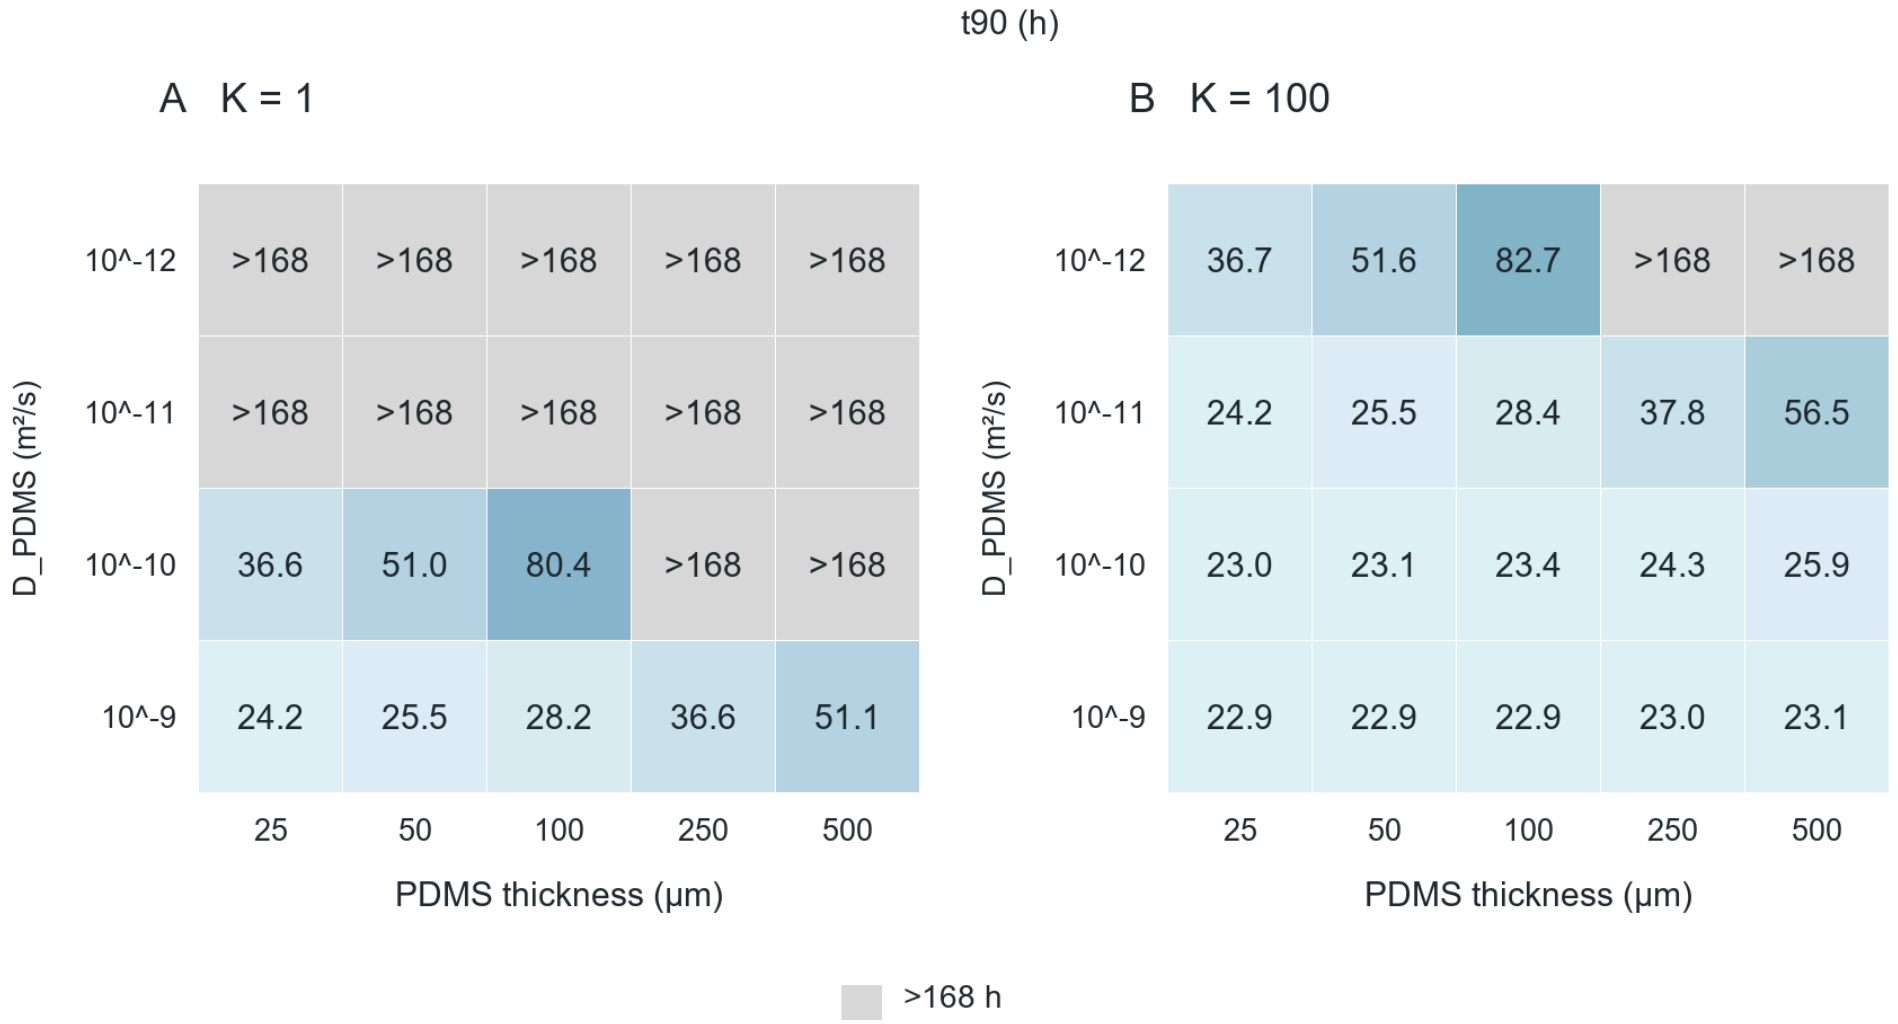

Figure S12. Simulation · film selection depends on target properties. Forty displayed cells are a slice of 243 deterministic parameter combinations, assuming 100 µL fill, liquid diffusivity 10^{-9} m²/s and liquid–gas partition coefficient 10. There are no biological replicates, Monte Carlo seeds or statistical error bars. Gray cells did not reach 90% within the 168-hour calculation; they are censored, not zero-delay cases. Material inputs are illustrative and these long model horizons do not establish cell stability. Source methods and uncertainty.

For an ideal step directly at the liquid surface, no mixing and $D_{liquid} = 10^{-9}$ m²/s, bottom t90 is **22.83 hours**. This follows t90 ≈ 1.031105 L²/D; halving the liquid path would reduce that idealized time by a factor of four. The result is a transport boundary case, not a receptor or GCaMP response time.

Hardware: retain the no-film build and measure delivery at the cell plane →

Full diffusion, membrane and transient-transport analysis

What is solved, and what is omitted

Define gas-equivalent concentration *u*, partition coefficient *K*, and material concentration $c = Ku$. The layer model solves $\mathbf{K\,\partial u/\partial t = \partial(DK\,\partial u/\partial x)/\partial x}$, with instant interfacial partitioning, a concentration step upstream and a non-consuming bottom probe. A finite-volume spatial operator and matrix exponential handle the step-forcing solution.

Diffusivities and partition coefficients are generic sensitivity inputs: target VOCs and their measured properties are not established. The model assumes dilute linear transport, constant properties, uniform area and no convection, reaction, binding or cell consumption. It excludes meniscus deformation and leaks. Faster effective diffusion is a mixing proxy, not a demonstrated mixing mechanism.

Liquid depth can dominate the response time

The two transport stages in Figure S11 must remain separate when interpreting a proposed assay response time.

Across $D_{liquid} = 0.5/1/2 \times 10^{-9}$ m²/s, the ideal bottom t90 is 45.67/22.83/11.42 h. A shallower 25 µL counterfactual gives about 1.43 h at the middle diffusivity; it is not the selected operating fill. These earlier screening calculations retain the nominal 100 µL fill.

Figure S13: The depth and mixing comparison shows why cell-plane delivery needs its own measurement.

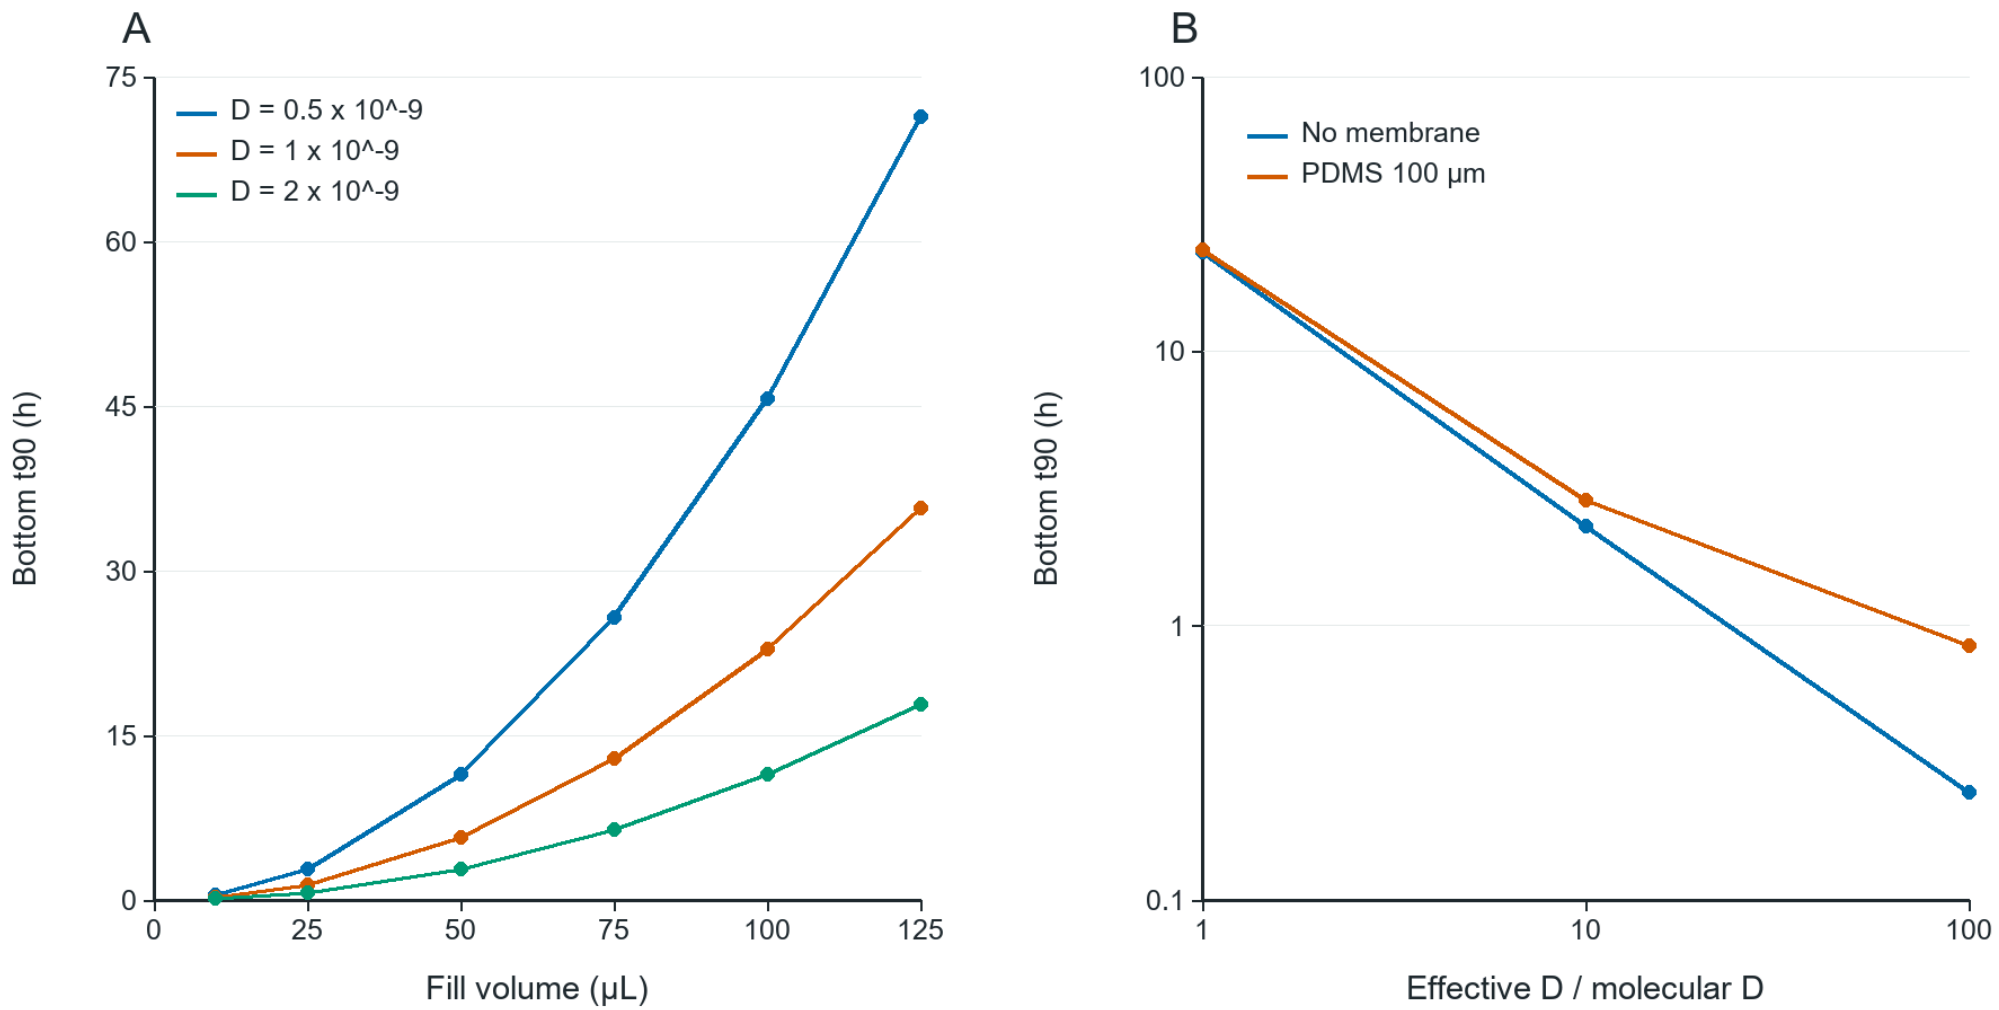

Figure S13. The depth and effective-diffusivity studies locate a transport bottleneck. They do not demonstrate viable cell operation or an implemented mixing mechanism.

PDMS is a research branch, not an installed solution

The 243-condition study includes 25/50/100/250/500 µm films and no-film controls, membrane diffusivities from 10^{-12} to 10^{-9} m²/s, and a broad partition-coefficient grid. Seventy-four conditions do not reach t90 within seven days; missing times must not be read as zero.

Under a representative stagnant-liquid condition, adding a 100 µm film (D = 10^{-10} m²/s, K = 100) changes the coupled t90 from about 22.85 to 23.39 h. If effective liquid diffusion is increased 100-fold, the respective times are approximately 0.25 and 0.85 h, and the membrane can become the bottleneck. A film can alter delay and recovery; hydrophobicity alone does not establish a water-vapor barrier. [8]

The sensitivity slice in Figure S12 illustrates why the film is not a required part of the selected enclosure.

Transient forcing and model checks

The saved CFD concentration history can drive the layer model in one direction. The source runs to 300 s; at that time the simulated bottom fraction remains below the 10^{-10} reporting floor. Extending the boundary thereafter as either 1 or its final 0.991144 value gives no-film t90 between 22.87 and 23.62 h. This is not a 23-hour CFD simulation or two-way mass-conserving coupling. The source rim t50/t90 is 48.78/149.02 s. The same continuation gives 23.41–24.18 h with the nominal film. These outcomes describe the stated boundary forcing, not measured delivery.

A 20-minute input pulse produces a delayed peak of about 2.78% after 3.92 h in the illustrated boundary case. It shows transport memory under the assumptions, not adsorption, receptor activation or a measured recovery curve.

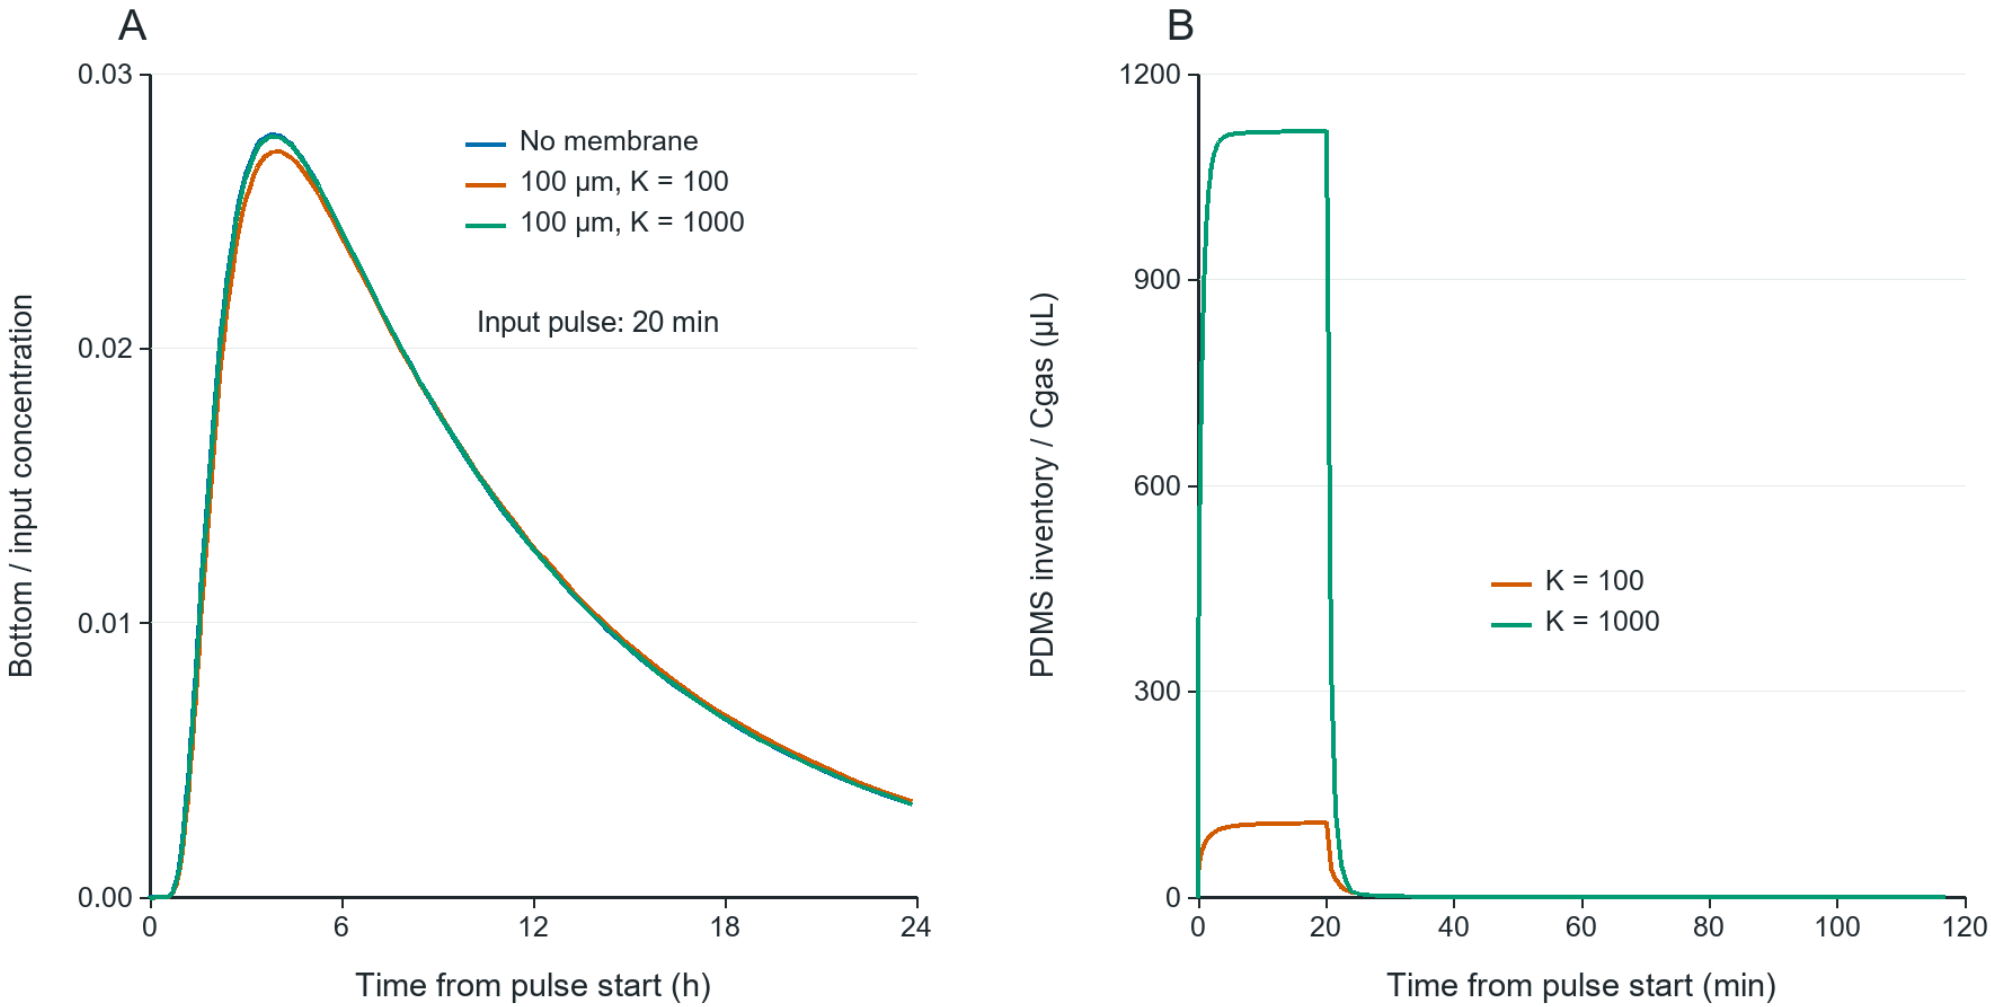

Figure S14. Transient storage and delayed delivery are consequences of the stated transport model; they are not biological measurements.

The standalone 128-cell membrane check reproduces the analytical time lag within 0.00305%; liquid-bottom t90 differs from the exact series by 0.000762%. Refining the coupled liquid mesh from 48 to 96 cells changes t90 by 0.00405%; worst normalized mass-balance residual is 2.17×10^{-5}.

Before selecting a film or claiming rapid sensing, define a target molecule, measure no-film delivery and recovery, then compare a characterized film using an independent concentration assay. GCaMP alone cannot separate transport from receptor sensitivity or optical artifacts.

Full transport methods and limits · Parameter and result summary — Zenodo (publication pending) · Back to the build decision →

Validation plan and simulation files

Hardware simulations / Validation boundary

Predictions await measurement

No calibrated detector power, dark-noise trace, standard-response curve or assembled-path gas measurement is archived. A numerical agreement plot would therefore be premature. The following pairs define the first tests; they contain no fitted experimental values.

Table S4. Numerical optical, demodulation and transport predictions are paired with the physical measurements needed for validation. The table contains proposed validation pairs, not fitted experimental results.

Prediction	Available evidence	Required paired observation
Optical collection: NE 0.593 ± 0.035 nW under the stated reference	Monte Carlo outputs and convergence checks	Traceable fluorescence standard, actual filter spectrum and photodiode calibration in the same geometry
Synthetic demodulation: residual amplitude bias 2.63–3.52%	Generated interference and host-tested analysis	Independent illuminated blank and reference recordings, including saturation and drift
Gas candidate: 3 mL/min; worst first-layer t90 50.024 s	Nominal CAD geometry and confirmed 100 µL fill	Measured well capacity, assembled-path flow and gas-interface time course
Stagnant liquid: bottom t90 22.83 h at 100 µL, D = 10^{-9} m²/s	Idealized diffusion calculation	Concentration time course at the cell plane under verified depth and mixing

Physical test sequence and blank recording forms →

Hardware simulations / Hypothesis

Capture and docking are exploratory

The reviewed local and repository sources contain no reproducible project capture-protein calculation in that release. Source audit: the docking / OBP evidence note — Zenodo (publication pending)

Docking can propose poses or rank candidates within one scoring model. A score is not measured Kd, an absorption rate, receptor activation or proof that liquid-delivery limits have been resolved. Keep this mechanism separate from the installed apparatus until its identity, assumptions and independent tests are available.

Hardware simulations / Data availability

Reproduce the physical simulations

Physical simulation index · Source archives and environments · Parameter / figure map. The optics, airflow and membrane implementations are available.

The independent physical calculations retain their source and figure index and reproducibility and assembly packages. The current gas-surface study uses 100 µL per well in the nominal CAD geometry.

Explore the separate odor-decoding study →

Liquid and capture hypotheses

Table S5. Supplementary decisions from the earlier liquid and capture screening.

Question	Evidence	Decision	Remaining measurement
Should a membrane be part of the current build?	Gas–liquid transport and conditional membrane sensitivity; no measured target-specific permeability	Keep the current enclosure without a film; establish the no-film baseline first	Defined-molecule arrival and recovery at the cell plane, then measured film selectivity and carryover
Does a capture protein speed cell exposure?	No target-specific uptake or binding records supplied in this package	Retain a separately labeled hypothesis	Identify protein and ligand; measure uptake and release

Hardware simulation sources

1. Burton, A., Obaid, S. N., Vázquez-Guardado, A., Schmit, M. B., Stuart, T., Cai, L., Chen, Z., Kandela, I., Haney, C. R., Waters, E. A., Cai, H., Rogers, J. A., Lu, L., & Gutruf, P. (2020). Wireless, battery-free subdermally implantable photometry systems for chronic recording of neural dynamics. Proceedings of the National Academy of Sciences, 117(6), 2835–2845. Source ↗

Source highlights

Figure 2H–J separates excitation fluence, detector acceptance and their overlap. It supports the optical-map method and excitation-rejection rationale; its implant geometry, filter and measured performance are not AeroSense inputs.

Back to citation

2. Wang, L., Jacques, S. L., & Zheng, L. (1995). MCML—Monte Carlo modeling of light transport in multi-layered tissues. Computer Methods and Programs in Biomedicine, 47(2), 131–146. Source ↗

Source highlights

Primary description of weighted-photon Monte Carlo transport through absorbing and scattering media with refractive interfaces. AeroSense uses these transport ingredients in independent code, not the MCML program or its tissue parameters.

Back to citation

3. LEE Filters. 101 Yellow. Accessed 11 October 2026. Source ↗

Source highlights

Manufacturer transmission information for the sampled LEE101 curve. Layer stacking, source spectra and complete-reader performance are separate modeling assumptions.

Back to citation

4. Edmund Optics. 500 nm, 12.5 mm Ultra-Thin OD4 Longpass Filter (#29-190). Accessed 11 October 2026. Source ↗

Source highlights

Manufacturer band specifications and dimensions: 12.45 ± 0.15 mm diameter, 0.30 mm thickness and 90% clear aperture. The wiki pointwise transmission is an explicit constructed scenario; the product page does not establish fit in AeroSense or measured device sensitivity.

Back to citation

5. Qian, Y. H., d’Humières, D., & Lallemand, P. (1992). Lattice BGK Models for Navier-Stokes Equation. Europhysics Letters, 17(6), 479–484. Source ↗

Source highlights

Foundational lattice-BGK collision and equilibrium formulation used by the D3Q19 airflow solver. It supports the numerical method, not the chosen geometry, pump flow or biological safety limits.

Back to citation

6. Guo, Z.-L., Zheng, C.-G., & Shi, B.-C. (2002). Non-equilibrium extrapolation method for velocity and pressure boundary conditions in the lattice Boltzmann method. Chinese Physics, 11(4), 366–374. Source ↗

Source highlights

Boundary distributions combine imposed equilibrium values with an extrapolated non-equilibrium part from a neighboring fluid node. This matches the pressure-terminal reconstruction in the airflow source code; it is not a citation for the separate scalar solver.

Back to citation

7. Firpo, G., Angeli, E., Guida, P., Pezzuoli, D., Repetto, D., Repetto, L., & Valbusa, U. (2019). The Role of Surfaces in Gas Transport Through Polymer Membranes. Polymers, 11(5), 910. Source ↗

Source highlights

Measures PDMS gas transport and distinguishes bulk diffusion from surface resistance. Table 1 supplies the source package’s oxygen example at 293 K; it does not supply the generic VOC grid or characterize the proposed film.

Back to citation

8. Bian, P., Wang, Y., & McCarthy, T. J. (2021). Rediscovering Silicones: The Anomalous Water Permeability of “Hydrophobic” PDMS Suggests Nanostructure and Applications in Water Purification and Anti-Icing. Macromolecular Rapid Communications, 42(5), 2000682. Source ↗

Source highlights

Experiments on unfilled, silica-free PDMS show that water repellency does not imply low water-vapor transmission; thin-film surface kinetics can matter. No evaporation rate or culture lifetime is transferred to AeroSense.

Back to citation