An open implementation of slot-die and blade coating operability limits for scalable perovskite photovoltaics

Abdulai Rabiu · ORCID 0000-0002-3739-140X

Wright Center for Photovoltaics Innovation and Commercialization, Department of Physics & Astronomy, University of Toledo, Toledo, Ohio 43606, USA

arabiu@rockets.utoledo.edu

Slot-die and blade coating deposit the functional layers of printed optoelectronic devices, and both fail outside a bounded region of process parameters. The boundaries of that region have been derived and measured since the 1940s, but the results are scattered across the coating-science literature in notation that differs between papers, and they are rarely applied by the materials groups who most need them. In practice, a laboratory establishes its coating window by trial and error, one substrate at a time, and the resulting knowledge stays tacit and local. We present pvcoat, an open-source Python package that computes wet and dry film thickness and the operability window for both coating heads directly from process parameters, implementing the low-flow limit, the coating-bead vacuum window, the Landau-Levich film thickness, and the air-entrainment speed together with the validity range each model was derived for. The package has no runtime dependencies, accepts quantities in the units coating practitioners use, and attaches an explicit warning to any result computed outside a model's range of derivation. We verify the implementation against seven closed-form identities, all satisfied to machine precision, and exercise it on a worked case representative of a formamidinium lead triiodide precursor ink. Two results of the framework deserve emphasis for practitioners. First, the low-flow limit rises as the two-thirds power of the capillary number, so increasing line speed makes thin films harder rather than easier: for the worked ink at a 100 µm gap, a 700 nm dry layer remains coatable only up to 106 mm s^{-1}, a factor of 31 below the air-entrainment ceiling of 3.31 m s^{-1}. Second, the speed at which a blade-coated film is thinnest is independent of solids loading, because the solids volume fraction cancels from both competing branches; diluting an ink lowers the whole thickness-speed curve without moving the crossover. We emphasise that pvcoat reproduces published equations and has not been validated against coating experiments, and we set out the validation programme this requires. The package, its test suite and the scripts reproducing every figure and table in this work are openly available under the MIT licence.

Keywords: Slot-die coating, blade coating, operability window, Landau-Levich, air entrainment, perovskite photovoltaics, scalable manufacturing, open-source scientific software

I. INTRODUCTION

Perovskite photovoltaics has moved from a certified power conversion efficiency of 3.8% in 2009 to above 26% on laboratory-scale single-junction devices, a trajectory no other photovoltaic technology has matched [1-4]. Almost the whole of that record was set on spin-coated films of a few square centimetres. Spin coating cannot be scaled: it wastes the great majority of the precursor, it does not translate to continuous substrates, and the hydrodynamics that make it reproducible on a 25 mm square do not survive the transition to a 300 mm web [5-8]. The processes that will carry perovskite photovoltaics into manufacturing are the pre-metered and self-metered coating methods already used for batteries, barrier films and organic photovoltaics: slot-die coating and blade coating [9-14]. Both have been demonstrated on perovskite cells and minimodules, in sheet-to-sheet and roll-to-roll form [62-66].

Both methods fail outside a bounded region of the process-parameter space. Below a minimum wet thickness, the coating bead cannot sustain itself, and the film breaks into ribs, rivulets, or breaklines. Above a critical speed, the dynamic wetting line becomes unstable and entrains air, which dries into pinholes and craters. Between these bounds, the vacuum applied behind the bead must lie inside a window set by the curvature the two menisci can sustain. The existence of these boundaries, and the physics that sets them, has been understood for decades: the Landau-Levich analysis of film entrainment dates from 1942 [15,16], Ruschak's limiting-flow treatment of pre-metered coating from 1976 [17], and the Higgins-Scriven bounds on coating-bead operability from 1980 [18]. A mature literature has since mapped the minimum wet thickness experimentally [19-23], characterised dynamic wetting failure [24-30], and reviewed the operating limits as a coherent set [31-33].

That literature is not reaching the laboratories that need it. The notation differs between papers; the key reviews sit behind subscription paywalls; the expressions are stated in forms that assume familiarity with lubrication theory; and the validity ranges, where stated at all, are given in words rather than enforced. The practical consequence is visible in the perovskite coating literature, where coating gaps, speeds and flow rates are frequently reported as the outcome of an empirical optimisation with no reference to the operability limits that bound them, and where the same optimisation is then repeated from scratch by the next group. A calculation that would take a minute is instead performed as an experiment that takes a week.

This paper presents pvcoat, an open-source Python package that closes that gap for the two coating methods most relevant to scalable perovskite photovoltaics. The package takes an ink and a coating head, both described in the units a practitioner actually uses, and returns the wet and dry film thickness, the dimensionless groups, the margin to each operability limit, the admissible vacuum window, and an explicit warning whenever a result has been computed outside the range its underlying model was derived for. It is small, has no runtime dependencies, and is designed so that every number it produces can be traced to a published equation.

Our contribution is one of synthesis, implementation and verification rather than new physics. Every model implemented here is published, and we have added nothing to it. What the package adds is a single consistent notation, a programmatic interface, an executable statement of each model's validity range, and a test suite that checks the implementation against closed-form identities rather than against recorded outputs. We state at the outset, and return to it in Section VI, that pvcoat has not been validated against coating experiments; it is a calculator for published models, and should be described as nothing more than that until the validation programme of Section VI C has been carried out.

The remainder of this paper is organised as follows. Section II sets out the theoretical framework, deriving each implemented limit from its source and stating the assumptions and the validity range. Section III describes the implementation, the verification strategy and the provenance discipline. Section IV applies the package to a worked case representative of a formamidinium lead triiodide precursor ink and reports the resulting operability maps. Section V discusses what the maps imply for the design of perovskite coating processes. Section VI states the limitations candidly and sets out the work required to close them.

II. THEORETICAL FRAMEWORK

This section states each implemented model, its assumptions and its range of validity. Table I defines the symbols. Table II summarises the models, their sources and their validity ranges in one place. Throughout, the sign and naming conventions follow the review of Ding and co-workers [31]: *upstream* is the side the substrate arrives from and where a vacuum box is attached; *downstream* is the side the wet film leaves on.

TABLE I. Nomenclature. pvcoat computes in SI throughout; the final column gives the unit the package expects internally.

Symbol	Meaning	SI unit
U	substrate or coating speed	m s^{-1}
Q	volumetric flow rate through the die	m^3 s^{-1}
q	flow rate per unit width, Q/W	m^2 s^{-1}
W	coated width	m
h	wet film thickness	m
h_{dry}	dry film thickness	m
h_{min}	low-flow limit (minimum wet thickness)	m
t_0	minimum wet thickness with no applied vacuum	m
H_d, H_u	downstream, upstream coating gap	m
L_d, L_u	downstream, upstream lip land length	m
R	meniscus radius of curvature	m
μ	dynamic viscosity	Pa s
σ	surface tension	N m^{-1}
ρ, ρ_f	solution, dried-film density	kg m^{-3}
c	solids mass per volume of solution	kg m^{-3}
φ	solids volume fraction, c/ρ_f	–
E	evaporation flux per unit blade length	m^2 s^{-1}
ΔP	applied vacuum, $P_{atm} - P_{vac}$	Pa
Ca	capillary number, $\mu U/\sigma$	–
Re	Reynolds number, $\rho q/\mu$	–
U_{ae}	air-entrainment (wetting failure) speed	m s^{-1}
U^*	blade-coating regime crossover speed	m s^{-1}

A. Metering and the dry-film closure

Slot-die coating is *pre-metered*: the pump, not the head, sets the thickness. For an incompressible ink delivered at volumetric flow rate Q over a coated width W on a substrate moving at speed U, mass conservation fixes the wet thickness exactly,

$$h = Q / (W U) , \tag{1}$$

with no adjustable parameter. Equation (1) says nothing about whether the resulting coating is *stable*; that is the subject of the remaining subsections. Drying removes the solvent, and for a fully dense film with no loss of solid the dry thickness follows from the solids volume fraction φ,

$$h_{dry} = \varphi\, h , \quad \varphi = c / \rho_f , \tag{2}$$

where c is the mass of solid per unit volume of solution and ρ_f the density of the dried film. Equation (2) is a *floor* on the thickness rather than an estimate: residual porosity, retained solvent and volume change on crystallisation all make the real film thicker. For the worked ink of Section IV, φ = 0.2160, so the film shrinks by a factor of 4.63 on drying and a 700 nm perovskite layer requires 3.241 μm of wet film.

Blade coating, by contrast, is *self-metered*: the speed and the meniscus set the thickness, and the operative question runs in the opposite direction. That case is treated in Section II E.

B. The low-flow limit

The thinnest wet film a slot-die bead can sustain follows from the Landau-Levich boundary condition applied to the downstream meniscus. A meniscus of radius of curvature R dragging film from a reservoir onto a substrate deposits a layer of thickness [15-17]

$$h = 1.34\, R\, Ca^{2/3} , \quad Ca = \mu U / \sigma , \tag{3}$$

so a thinner film demands a more sharply curved meniscus. The Landau-Levich problem and its extensions are reviewed in Refs. [32,33,52,55,56]. The curvature cannot increase without bound. The downstream meniscus is pinned at the die-lip corner and must remain tangent to the moving substrate; taking a circle of radius R with its centre at height R above the substrate and requiring it to pass through the corner at height H_d gives

$$x_c^2 = 2 R H_d - H_d^2 \;\Rightarrow\; R \geq H_d / 2 , \tag{4}$$

so that the tightest admissible arc has $R_{min} = H_d/2$. Substituting into Eq. (3),

$$h_{min} = 1.34\, (H_d/2)\, Ca^{2/3} = 0.67\, H_d\, Ca^{2/3} . \tag{5}$$

Below h_{min} the downstream meniscus recedes into the gap, air fingers invade the bead, and the coating breaks up. Because the limit is set by the downstream meniscus alone, **applying more vacuum does not move it** - a point made explicitly by Ding and co-workers in distinguishing the low-flow limit from the zero-vacuum minimum thickness [31], and one that is routinely missed in practice.

Equation (5) carries a consequence that runs against intuition. Because h_{min} grows as $Ca^{2/3}$ and Ca is proportional to U, the thinnest coatable film *increases* with line speed: doubling the speed raises the floor by a factor of $2^{2/3} \approx 1.59$. Thin films are a low-speed proposition, and the window narrows from below as the line is sped up. Section IV C quantifies this for the worked case.

Equation (5) is an asymptotic low-Ca result. Experiment departs from it in two stages. Chang and co-workers identify three regions in the measured minimum wet thickness [22]: a low-Ca branch following the two-thirds scaling, an intermediate plateau, and a high-Ca branch in which h_{min} falls again as inertia lengthens the bead. Khandavalli and Rothstein measure the plateau at h/H between roughly 0.44 and 0.55 [23], and Carvalho and Kheshgi document the high-Ca decrease [21]. pvcoat implements the first branch only and warns outside it; Section II F describes the mechanism.

C. The coating-bead vacuum window

The admissible range of applied vacuum follows from a pressure balance across the bead. Writing $\Delta P = P_{atm} - P_{vac}$ so that a positive value is a suction, the balance sums the capillary jump at each meniscus and the lubrication pressure drop under each lip,

$$\Delta P = \sigma/R_d - \sigma/R_u + 6\mu\, U\, L_u/H_u^2 + (6\mu\, U\, L_d/H_d^2)(1 - 2h/H_d)\ . \tag{6}$$

Term by term: the downstream capillary term σ/R_d is fixed by the Landau-Levich condition of Eq. (3) to $1.34\ \sigma\ Ca^{2/3}/h$, so thin films demand more vacuum; the upstream term σ/R_u is free to swing about its pinning corner over $\pm 2\sigma/H_u$, and those two extremes are the two ends of the window; and the viscous terms follow from lubrication flow between lip and substrate, with zero net flow upstream of the slot and $q = hU$ downstream, the latter changing sign at $h = H_d/2$.

An immediate and useful consequence is that the window always has width

$$\Delta P_{max} - \Delta P_{min} = 4\sigma / H_u\ , \tag{7}$$

independent of speed, thickness and viscosity. The width of a slot-die process window in vacuum is therefore set by the upstream gap and the surface tension alone, which gives a designer a direct handle on process latitude. For the worked case this is 12.0 mbar. Equation (7) is one of the identities the test suite asserts (Section III B).

With both lip lengths set to zero the viscous terms vanish and Eq. (6) reduces to Ruschak's capillary model [17]; with them retained it is the viscocapillary model of Higgins and Scriven [18]. Setting $\Delta P_{min} = 0$ and solving for h gives the thinnest film coatable with no vacuum box, denoted t_0. The left-hand side falls monotonically with h - both the downstream capillary term and the adverse viscous gradient weaken as the film thickens - so the root is unique and pvcoat locates it by bisection. In the capillary limit it reduces to the closed form $0.67\ H_u\ Ca^{2/3}$. The quantities t_0 and h_{min} are distinct and have different physics: t_0 is set by how far the *upstream* meniscus can swing, h_{min} by the *downstream* meniscus.

D. Dynamic wetting failure

Above a critical speed the dynamic wetting line breaks into a sawtooth and entrains air, which dries into pinholes and craters [24-26]. This is a property of the contact line rather than of the coating head, so the same bound applies to slot-die and blade coating alike, and it is the *upper* speed bound of every window; the underlying moving-contact-line physics is reviewed in Refs. [53,54]. pvcoat implements two models. The first is the dimensional correlation of Gutoff and Kendrick [27],

$$U_{ae} = 6.9\ \mu^{-0.67} \quad (\mu \text{ in mPa s},\ U_{ae} \text{ in m s}^{-1})\ , \tag{8}$$

in which viscosity dominates and surface tension plays only a secondary role, which is why no σ appears. The second expresses the bound through a critical capillary number, $U_{ae} = Ca_{crit}\ \sigma/\mu$, for use when Ca_{crit} has been measured on the actual ink and substrate.

The two models disagree, often by a factor of several, because air entrainment depends on substrate roughness, ambient pressure and contact-line chemistry that neither captures [28-30]. Section IV G quantifies the disagreement. For polymer solutions, Cohu and Benkreira report that the *solvent* viscosity predicts the failure speed better than the solution viscosity [28], which suggests the failure is set by small molecules near the contact line rather than by bulk rheology; where an ink carries polymer, the solvent viscosity should be passed.

E. Blade coating: two competing deposition regimes

Blade coating is self-metered, and two mechanisms compete to set the deposited thickness [34-37]. At low speed the ink dries at the meniscus as fast as it is dragged out, so solid deposits directly and mass conservation gives the evaporation regime,

$$h_{dry} = c\ E / (\rho_f\, U)\ , \tag{9}$$

where E is the solvent volume leaving the meniscus per unit blade length per second. Thickness falls as $1/U$, the signature of the regime. We stress that E is an *apparatus* property and not a material constant: it depends on the blade slit width, the substrate temperature, the gas flow over the meniscus and the solvent volatility, and must be obtained by fitting the low-speed branch of a thickness-speed sweep on the coater in question.

At high speed, viscous drag pulls out a wet film which dries afterwards, and the thickness follows the Landau-Levich result of Eq. (3) with the meniscus radius R, rising as $U^{2/3}$. Setting the two dry thicknesses equal, the solids volume fraction cancels identically from both sides,

$$c\ E / (\rho_f\, U) = \varphi \cdot 1.34\ R\ (\mu\ U/\sigma)^{2/3}\ , \quad \varphi = c/\rho_f \Rightarrow E/U = 1.34\ R\ (\mu\ U/\sigma)^{2/3}\ , \tag{10}$$

leaving a crossover speed that depends only on the evaporation flux, the meniscus radius and the ink's viscosity and surface tension,

$$U^* = [\ E\ \sigma^{2/3} / (1.34\ R\ \mu^{2/3})\ ]^{3/5}\ . \tag{11}$$

The crossover speed is independent of solids loading. Diluting an ink moves the whole thickness-speed curve downward without moving U^*. We have not found this stated explicitly in the meniscus-guided coating literature, though it follows in two lines from the standard two-regime framework; Section IV F reports the numerical confirmation and Section V B the practical consequence. U^* is also where the deposited film is thinnest, which makes it attractive for thin layers and unattractive for reproducibility: dh/dU passes through zero, so speed error enters only at second order, but the film morphology on either side of the crossover differs, and small drifts in E change which branch the process is on.

F. Validity ranges as a first-class output

Every model above has a range it was derived and tested in, and a value computed outside that range is still a number, which is precisely the problem. pvcoat therefore attaches explicit warnings to any result computed outside the ranges of Table II. The thresholds are deliberate judgements rather than sharp physical boundaries: Carvalho and Kheshgi report the viscocapillary model holding to about Ca = 1 [21], which fixes the capillary-number ceiling; Chang and co-workers report inertia mattering near Re = 20 [22], and pvcoat warns an order of magnitude earlier, at Re = 1, on the grounds that a warning is cheap and a silently invalid number is not. All thresholds are module constants and are intended to be moved by any user whose own evidence puts them elsewhere.

TABLE II. Models implemented in pvcoat v0.1.0, with the equation in this paper, the primary source and the range of validity enforced by the package.

Quantity	Eq.	Source	Validity enforced
Wet thickness $h = Q/(WU)$	(1)	mass conservation	exact
Dry thickness $h_{dry} = \varphi h$	(2)	mass conservation	dense-film floor
Low-flow limit $0.67\, H_d\, Ca^{2/3}$	(5)	[15,17,18,21,31]	$Ca \le 1$, $Re \le 1$
Vacuum window, width $4\sigma/H_u$	(6),(7)	[17,18,31]	$Ca \le 1$, $Re \le 1$
Zero-vacuum minimum t_0	(6)	[17,18,31]	$Ca \le 1$, $Re \le 1$
Air entrainment $6.9\mu^{-0.67}$	(8)	[27]	Newtonian, smooth substrate
Critical-Ca wetting failure	(8)	[25,26,30]	Ca_{crit} must be measured
Landau-Levich $1.34\, R\, Ca^{2/3}$	(3)	[15,16,34]	$Ca \ll 1$
Blade evaporation $cE/(\rho_f U)$	(9)	[34,67]	E must be measured
Regime crossover U^*	(11)	[34,67]	$Ca \le 1$
Flooding check $h > H_d$	–	geometric	hard error

III. IMPLEMENTATION

A. Architecture and design decisions

pvcoat is written in pure Python and requires version 3.9 or newer. Four design decisions shape it, each made to serve the stated purpose of a tool whose numbers can be traced to published equations.

No runtime dependencies. The package imports only the standard library. Numerical root-finding, where it is needed for the zero-vacuum thickness of Section II C, uses a bisection routine of twenty lines rather than a SciPy call [69]. The cost is a few lines of code; the benefit is that the package installs anywhere, cannot break when a numerical library changes its API, and can be read end to end by a reviewer. Matplotlib is required only by the example scripts and lives in an optional extra; the analysis reported in Section IV uses NumPy, Matplotlib and seaborn [68,70,71].

Units at the boundary, SI inside. Coating practice does not use SI. Gaps are quoted in micrometres, viscosities in centipoise, surface tensions in dyne per centimetre, flow rates in millilitres per minute and vacuum in millibar or inches of water. Every public constructor therefore accepts either a bare SI number or a string carrying a unit, so that Fluid(viscosity="3 cP") and Fluid(viscosity=0.003) are the same object. Internally everything is SI. This removes the single most common source of error in process calculations without asking the user to change how they think.

Results, not numbers. An evaluation returns an object carrying the thicknesses, the dimensionless groups, the margin to each limit, the vacuum window, the governing regime, any validity warnings and the bibliographic keys of the models used. A bare float would discard exactly the context that makes the number safe to act on.

Traceability by construction. Every model function carries a .references attribute holding the keys of the works it implements, resolvable through a bibliography module that ships with the package and is printable from the command line. No coefficient is fitted or tuned anywhere in the package; where a model requires a coefficient that its source reports empirically, that coefficient is a module-level constant so a user can replace it with their own measurement.

The public interface comprises objects for the ink and the two coating heads, modules for the dimensionless groups, the metering arithmetic, the operability limits and the validity checks, evaluation entry points for each head, and a command-line interface exposing single-point evaluation, window sweeps and the bibliography.

B. Verification strategy

The package ships 78 tests. The governing principle is that a test should check a closed form, a limiting case or an analytic identity, not a recorded output: a test that merely pins today's number protects nothing, because it will be updated to match whatever the code produces after a change. Table III lists the seven identities that carry the most weight, together with the residuals measured at the working point of Section IV.

TABLE III. Closed-form identities asserted by the test suite, with residuals measured at the working point of Table V. Residuals are relative except C3, which is an absolute pressure residual, and C5, an absolute error on an exponent.

ID	Identity	Residual
C1	Vacuum-window width equals 4 sigma / H_u	0 (exact)
C2	Capillary-limit t_0 equals 0.67 H_u Ca^(2/3)	4.5e-10
C3	Viscocapillary t_0 is a root of dP_min(h)=0	2.9e-07
C4	Low-flow limit equals 0.67 H_d Ca^(2/3)	0 (exact)
C5	Landau-Levich exponent recovered by regression	6.7e-16
C6	Blade branches coincide at the crossover U*	7.7e-16
C7	Metering closure h_wet x phi equals the dry target	0 (exact)

All seven are satisfied to machine precision. Three deserve comment. Identity C1 checks the analytic window width of Eq. (7), which follows from the structure of the pressure balance and would be violated by almost any sign error in Eq. (6). Identity C3 confirms that the bisection root genuinely satisfies the defining equation rather than merely terminating; the residual of 287.3 nPa is the bisection tolerance, not a physical quantity. Identity C5 recovers the Landau-Levich exponent of 2/3 by linear regression of log h on log U across three decades of speed, which tests the composition of the capillary-number definition with the thickness law rather than either in isolation.

Beyond these identities the suite covers the unit parser against hand-computed conversions, the rejection of unphysical inputs, the equivalence of specifying a flow rate and specifying a target thickness, the monotonicity of the limits in speed and gap, the switching of the governing blade regime at the crossover, the invariance reported in Section IV F, and the command-line interface. Continuous integration runs the suite

on Python 3.9 through 3.13 on Linux, macOS and Windows, and separately verifies that the package imports and runs with no optional dependency installed.

We distinguish this verification - the code solves the equations correctly - from validation, which would ask whether the equations describe a real coater. Only the former has been done. Section VI returns to this.

C. Reproducibility and provenance

Every number reported in this paper, in the package documentation and in the figures is computed by a script and written to a machine-readable statistics file, from which the text is then interpolated. No value is typed by hand into prose. A continuous integration job regenerates the statistics file and fails the build if any reported number has drifted from the code that produced it. The analysis script that generates every figure and table here is included in the archived release. This discipline follows established guidance on reproducible computational practice and software citation [72-75], and responds to the documented fragility of research code over time [76,77].

The release carries a verification record stating what the author checked before publication, including the provenance of each empirical coefficient, and a limitations document enumerating what the package does not do. We regard both as part of the software rather than as commentary on it.

IV. RESULTS

A. A worked case

We exercise the package on a case representative of a formamidinium lead triiodide precursor ink in a dimethylformamide/dimethyl sulfoxide solvent system, of the kind used for blade- and slot-die-coated p-i-n perovskite minimodules. Table IV lists the parameters. We stress that the ink properties are plausible literature-scale values chosen to exercise the models, not measurements; the evaporation flux in particular is an order-of-magnitude placeholder. Any application to a real process must substitute measured rheology, tensiometry, pycnometry and a fitted low-speed thickness sweep.

TABLE IV. Worked-case parameters. Ink properties are illustrative values chosen to exercise the models over a realistic range; they are not measurements, and the evaporation flux in particular is an order-of-magnitude placeholder.

Parameter	Symbol	Value
Viscosity	μ	3 mPa s
Surface tension	σ	30 mN m^{-1}
Solution density	ρ	1450 kg m^{-3}
Solids concentration (1.4 M $FAPbI_3$)	c	885.6 kg m^{-3}
Dried-film density	ρ_f	4100 kg m^{-3}
Solids volume fraction	φ	0.2160
Capillary length	ℓ_c	1.45 mm
Evaporation flux (assumed)	E	1.0e-10 m^2 s^{-1}
Slot-die gap (both lips)	H_d, H_u	100 µm
Lip land length (both)	L_d, L_u	1.0 mm
Coated width	W	50 mm
Blade gap	G	150 µm
Blade meniscus radius	R	75 µm
Target dry thickness	h_{dry}	700 nm

At the nominal working point - a 100 µm gap, a 50 mm stripe, 10 mm s^{-1} and a 700 nm dry target - the metering arithmetic of Eqs. (1) and (2) gives a wet thickness of 3.241 µm and a pump setting of 0.0972 mL min^{-1}. The capillary number is 1.00×10^{-3} and the Reynolds number 0.0157, both two to three orders of magnitude inside the viscocapillary validity range, so the models are being used where they were derived. Table V reports the full evaluation.

TABLE V. Evaluation of the slot-die working point: a 700 nm dry film at 10 mm s^{-1} over a 50 mm stripe.

Quantity	Value
Wet thickness for the target	3.2406 µm
Pump setting	0.0972 mL min^{-1}
Capillary number Ca	1.000e-03
Reynolds number Re	0.0157
Low-flow limit h_{min}	0.6700 µm
Margin h/h_{min}	4.84
Zero-vacuum minimum t_0	0.7124 µm
Vacuum window	-4.41 to 7.59 mbar
Window width $4\sigma/H_u$	12.00 mbar
Air-entrainment speed U_{ae}	3.305 m s^{-1}
Speed margin U_{ae}/U	331
Verdict	operable

The low-flow limit is 0.670 µm, giving a margin of 4.84 between the metered film and the thinnest the bead can sustain. The vacuum window runs from -4.41 to 7.59 mbar; because the lower bound is negative, no vacuum box is required at this point, and the zero-vacuum minimum thickness $t_0 = 0.712$ µm sits just above h_{min}, as expected for a parallel die. The air-entrainment speed is 3.31 m s^{-1}, a margin of 331 on the operating speed. The point is comfortably operable, and the binding constraint is nowhere near the contact line.

B. The operability window

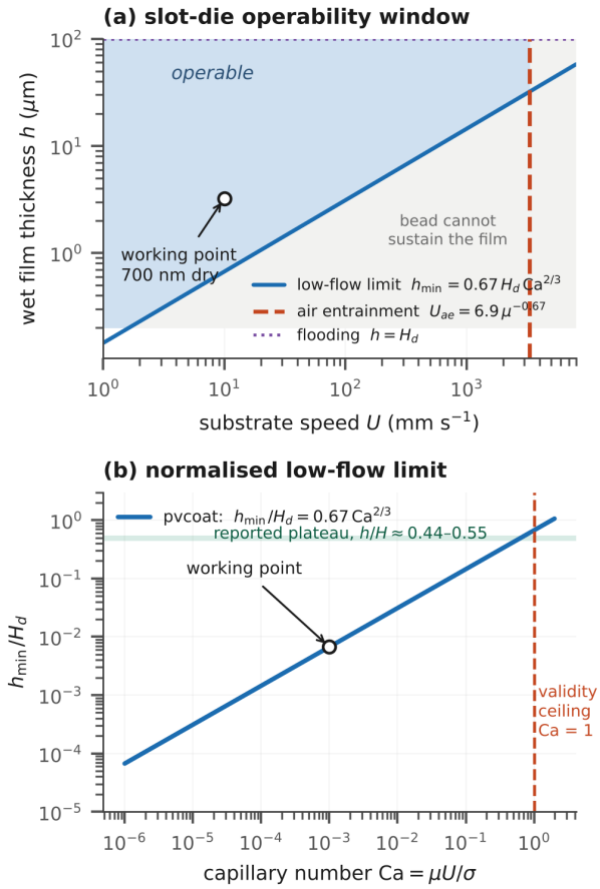


FIG. 1. (a) The slot-die operability window for the worked ink at a 100 µm gap. The shaded region is operable; it is bounded below by the low-flow limit of Eq. (5), to the right by the air-entrainment speed of Eq. (8), and above by flooding at h = H_d. The working point of Table V is marked. (b) The same limit normalised, showing the $Ca^{2/3}$ scaling of Eq. (5) against the plateau at h/H ≈ 0.44–0.55 reported experimentally [23] and the Ca = 1 validity ceiling beyond which pvcoat warns.

Figure 1(a) shows the window. Its shape is characteristic and worth reading carefully. The lower boundary is not horizontal: it rises with speed as $Ca^{2/3}$, so the window closes from below as the line is sped up, and the operable region is a wedge rather than a band. The right-hand boundary, air entrainment, is vertical because it is a property of the contact line and does not depend on the metered thickness. The two boundaries intersect, and the intersection defines the thinnest film the system can ever coat at any speed.

Figure 1(b) places the implemented limit against the experimental picture. The two-thirds branch is valid over the five decades of capillary number where the working point sits, but it crosses the measured plateau band near Ca ≈ 0.5 and the validity ceiling at Ca = 1. For the worked ink at 10 mm s^{-1} the capillary number is 1.00 × 10^{-3}, so the model is used three decades inside its range; reaching the ceiling would require a line speed of several metres per second or a viscosity above about 3 Pa s.

C. Speed dependence and the thin-film constraint

TABLE VI. Speed dependence at fixed metered thickness h = 3.241 µm (700 nm dry). The low-flow limit rises as $Ca^{2/3}$ and eventually overtakes the metered film.

U (mm s^{-1})	Ca	Re	h_{min} (µm)	h/h_{min}	ΔP_{min} (mbar)	Verdict
2	2.00e-4	0.003	0.229	14.14	-5.5	operable
5	5.00e-4	0.008	0.422	7.68	-5.0	operable
10	1.00e-3	0.016	0.670	4.84	-4.4	operable
20	2.00e-3	0.031	1.064	3.05	-3.3	operable
50	5.00e-3	0.078	1.959	1.65	-0.6	operable
100	1.00e-2	0.157	3.110	1.04	3.2	marginal
200	2.00e-2	0.313	4.937	0.66	10.1	not operable
500	5.00e-2	0.783	9.093	0.36	28.3	not operable

Table VI sweeps the speed at fixed target thickness and makes the consequence of the $Ca^{2/3}$ scaling concrete. Raising the speed from 2 to 500 mm s^{-1} - a factor of 250 - raises the low-flow limit by a factor of $250^{2/3} \approx 40$, from 0.229 to 9.09 µm, while the metered film stays at 3.24 µm. The margin falls from 14.1 to 0.36, and the process crosses from comfortably operable to not operable at all. For this ink and gap the 700 nm target survives only to 106 mm s^{-1}.

This is the central practical result for perovskite coating. The air-entrainment ceiling of 3.31 m s^{-1} is often quoted as the speed limit of a coating process, and for thick films it is. For the sub-micrometre wet films that perovskite absorbers require, it is not the operative bound: the binding constraint arrives a factor of 31 earlier, and it arrives from the bead rather than the contact line. A throughput calculation that takes the air-entrainment speed as the ceiling will overestimate the attainable line speed by roughly that factor, and the error grows as the target film is made thinner, since h_{min} is unchanged while the metered thickness falls.

D. The vacuum window

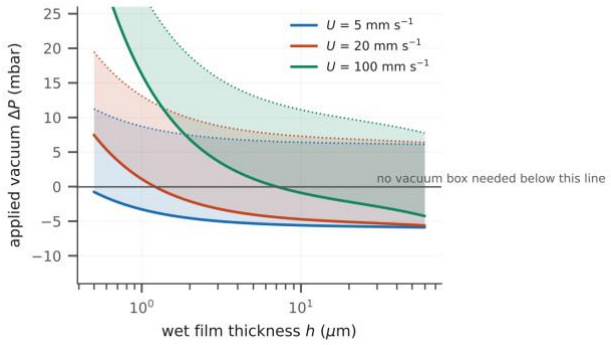


FIG. 2. Admissible applied vacuum as a function of wet film thickness at three speeds, from Eq. (6). Shaded bands span ΔP_{min} to ΔP_{max}; solid lines mark the lower bound, dotted lines the upper. The band width is $4\sigma/H_u$ = 12.0 mbar everywhere, independent of speed and thickness, as Eq. (7) requires. Where the lower bound falls below zero the bead is operable with no vacuum box.

Figure 2 shows how the vacuum requirement varies. Two features matter operationally. First, the band width is constant - the bands at 5, 20 and 100 mm s^{-1} are vertically offset but identical in height, as Eq. (7) requires. Process latitude in vacuum is a property of the die and the ink, not of the operating point, and a designer who wants more latitude must change the upstream gap or the surface tension. Second, the whole band rises steeply as the film thins, because the downstream capillary term scales as 1/h. Thin films need vacuum, thick films do not, and the transition is sharp. At the worked point the lower bound is -4.4 mbar, comfortably negative; at a tenth of that thickness it would be strongly positive and a vacuum box would become mandatory.

E. Blade coating and the regime crossover

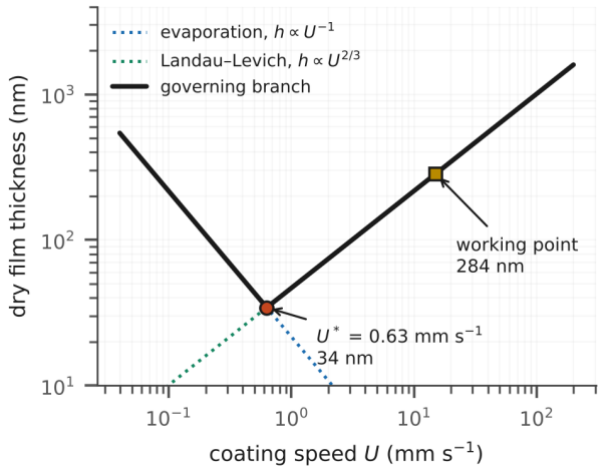


FIG. 3. Dry film thickness against coating speed for the worked ink at a 75 µm meniscus radius. The evaporation branch of Eq. (9) falls as U^{-1}, the Landau-Levich branch of Eq. (3) rises as $U^{2/3}$, and the governing branch is the upper envelope. The crossover $U^* = 0.63$ mm s^{-1} is the thinnest attainable film, 34 nm. The square marks the working point of Table IV.

Figure 3 shows the two-regime structure. At 15 mm s^{-1} the worked case sits on the Landau-Levich branch and deposits 284 nm, well above the crossover at 0.63 mm s^{-1} where the thinnest attainable film is 34 nm. Table VII tabulates the sweep. The V shape has a practical reading: a target thickness is generally reachable at *two* speeds, one on each branch, and the two will not give the same film. The evaporation branch dries at the meniscus and tends to produce different crystallisation behaviour from the Landau-Levich branch, which deposits a wet film that dries afterwards. Reporting a blade speed without stating which branch it lies on leaves the process under-specified.

TABLE VII. Blade coating: the two branches and the governing dry thickness. The crossover lies at $U^* = 0.629$ mm s^{-1}.

U (mm s^{-1})	Ca	Evaporation (nm)	Landau-Levich (nm)	Governing (nm)	Regime
0.100	1.00e-5	216.0	10.1	216.0	evaporation
0.300	3.00e-5	72.0	21.0	72.0	evaporation
0.629	6.29e-5	34.3	34.3	34.3	evaporation
1.000	1.00e-4	21.6	46.8	46.8	landau-levich
3.000	3.00e-4	7.2	97.3	97.3	landau-levich
10.000	1.00e-3	2.2	217.1	217.1	landau-levich
15.000	1.50e-3	1.4	284.5	284.5	landau-levich
30.000	3.00e-3	0.7	451.6	451.6	landau-levich
100.000	1.00e-2	0.2	1007.6	1007.6	landau-levich

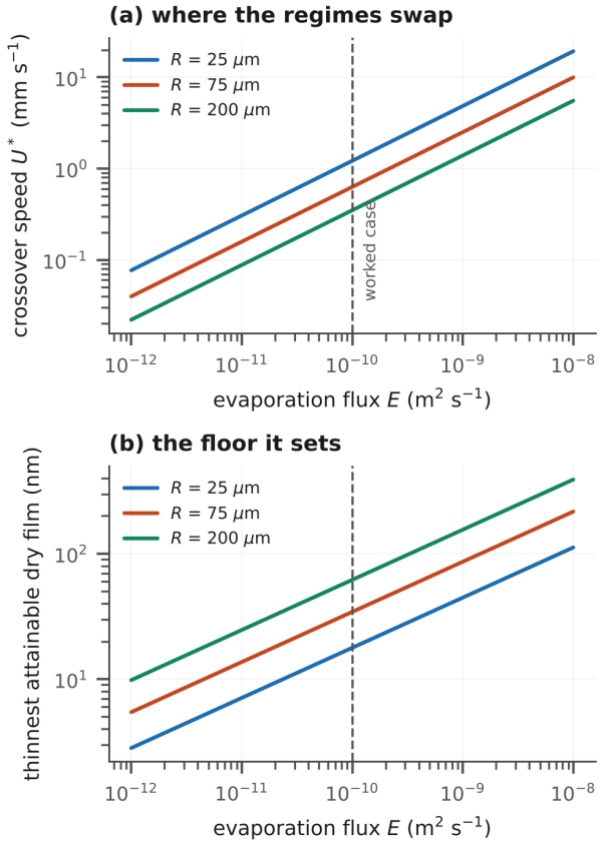


FIG. 4. (a) Crossover speed U^* from Eq. (11) against evaporation flux, for three meniscus radii. (b) The corresponding thinnest attainable dry film. The dashed line marks the worked case. Because the solids volume fraction cancels in Eq. (10), neither panel depends on ink concentration.

Figure 4 maps the crossover over four decades of evaporation flux. Since E is an apparatus property that must be measured, and since laboratory values are seldom known better than an order of magnitude, the map is the honest way to present the result: it tells a user what U^* would be once they have measured E, rather than asserting a value they have no way to check.

F. Concentration invariance of the crossover

Equation (10) predicts that U^* does not depend on solids loading. We confirm this numerically: reducing the solids concentration of the worked ink to one quarter, from 885.6 to 221.4 kg m^{-3}, leaves the crossover speed at 0.6291 mm s^{-1} against 0.6291 mm s^{-1} for the full-strength ink, a relative difference of 0.0. The test suite asserts this identity.

The practical content is that dilution is a clean thickness knob for blade coating. Halving the concentration halves the deposited thickness on both branches and leaves the regime structure - and therefore the speed at which the process is most tolerant to speed error - exactly where it was. A process developer tuning thickness by concentration need not re-map the regime boundary; one tuning thickness by speed must.

G. Disagreement between air-entrainment models

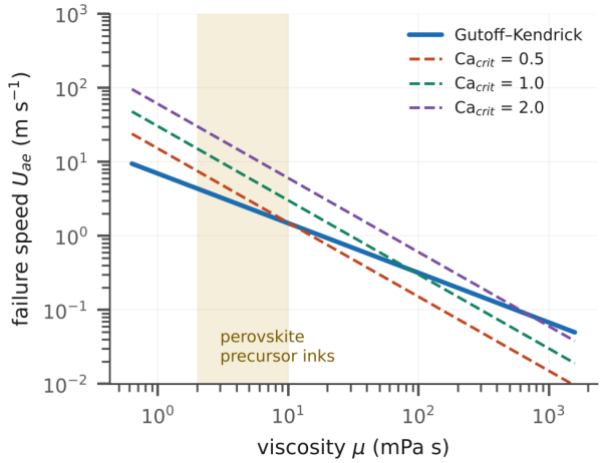


FIG. 5. Dynamic wetting failure speed against viscosity. The Gutoff-Kendrick correlation of Eq. (8) is compared with the critical-capillary-number form at three values of Ca_{crit}, for $\sigma = 30$ mN m^{-1}. The shaded band marks the viscosity range of typical perovskite precursor inks. The two model families differ in both magnitude and slope.

Figure 5 compares the two air-entrainment models across three decades of viscosity and Table VIII tabulates the comparison. The disagreement is structural rather than numerical: the Gutoff-Kendrick correlation scales as $\mu^{-0.67}$ while the critical-capillary form scales as μ^{-1}, so they cannot be reconciled by choosing a value of Ca_{crit}. Reading the implied critical capillary number off the Gutoff-Kendrick curve gives 0.230 at 1 mPa s rising to 2.25 at 1000 mPa s, a span of an order of magnitude, which is simply a restatement of the different exponents.

For the low-viscosity inks of interest here the two models happen to agree to within a factor of a few, and in any case both place the limit far above the speeds at which the low-flow limit binds. The disagreement would matter for a high-viscosity ink or a high-speed line, and in that regime neither model should be trusted without a measurement on the actual ink and substrate.

TABLE VIII. Air-entrainment speed from the two implemented models, at $\sigma = 30$ mN m^{-1}. The final column is the critical capillary number implied by the Gutoff-Kendrick correlation, which is not constant.

μ (mPa s)	Gutoff-Kendrick (m s^{-1})	Ca_{crit}=0.5	Ca_{crit}=1.0	Ca_{crit}=2.0	Implied Ca_{crit}
1	6.900	15.00	30.00	60.00	0.230
3	3.305	5.00	10.00	20.00	0.331
10	1.475	1.50	3.00	6.00	0.492
30	0.707	0.50	1.00	2.00	0.707
100	0.315	0.15	0.30	0.60	1.051
300	0.151	0.05	0.10	0.20	1.511
1000	0.067	0.01	0.03	0.06	2.248

H. Process design maps

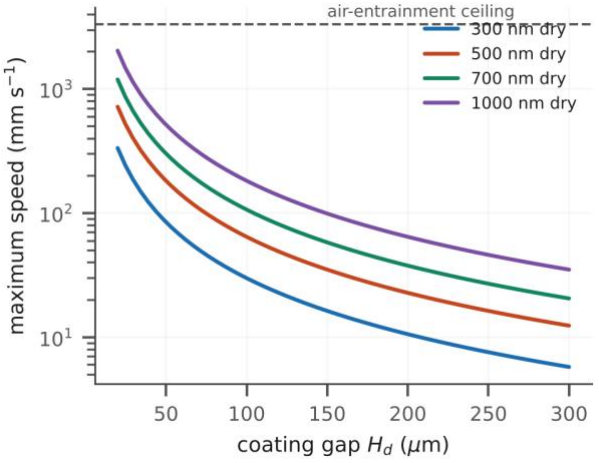


FIG. 6. Maximum coating speed at which a given dry thickness remains above the low-flow limit, as a function of coating gap, for the worked ink. The dashed line is the air-entrainment ceiling of 3.31 m s^{-1}, which binds only for the thickest films at the narrowest gaps.

Figure 6 inverts the window into the form a process engineer needs: given a target dry thickness and a coating gap, how fast can the line run? The curves fall steeply with gap, because h_{min} is proportional to H_d. Closing the gap is therefore the single most effective lever for throughput on thin films - a point that follows directly from Eq. (5) but is easy to miss when the gap is treated as a fixed property of the head.

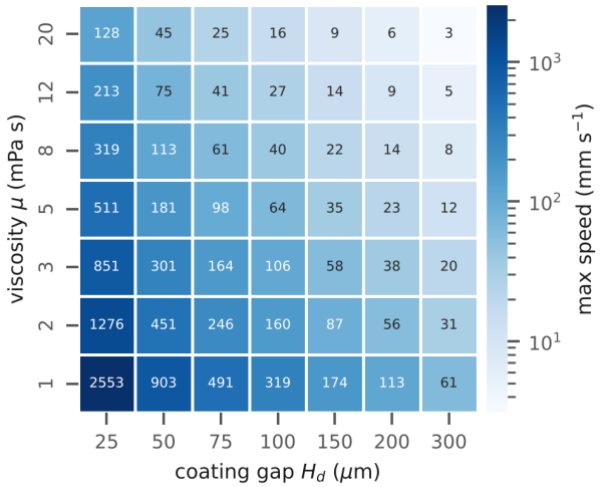


FIG. 7. Maximum line speed in mm s^{-1} for a 700 nm dry film, over coating gap and ink viscosity, bounded by the low-flow limit and the air-entrainment ceiling. Values span 3 to 2,553 mm s^{-1} across the mapped range.

Figure 7 extends this over the two parameters a formulator and a tool designer respectively control. Within the mapped range the attainable speed spans 3 to 2,553 mm s^{-1}, a factor of 831. Both levers act through the capillary number, and the map shows that gap and viscosity trade against one another: a 100 µm gap with a 3 mPa s ink and a 50 µm gap with an 8 mPa s ink permit similar speeds. For a process developer who can change the formulation more readily than the hardware, or vice versa, that equivalence is the useful content of the figure.

TABLE IX. Gap dependence for the 700 nm dry target. Closing the gap raises the attainable speed and narrows the vacuum window, the latter exactly as $4\sigma/H_u$.

H_d (µm)	h_{min} (µm)	h/h_{min}	h/H_d	ΔP window (mbar)	ΔP_{min} (mbar)	U_{max} (mm s^{-1})
25	0.168	19.35	0.130	48.0	-17.7	851
50	0.335	9.67	0.065	24.0	-9.4	301
75	0.503	6.45	0.043	16.0	-6.1	164
100	0.670	4.84	0.032	12.0	-4.4	106
150	1.005	3.22	0.022	8.0	-2.6	58
200	1.340	2.42	0.016	6.0	-1.7	38
300	2.010	1.61	0.011	4.0	-0.7	20

Table IX reports the gap sweep underlying these maps. The attainable speed falls by a factor of roughly 42 between a 25 and a 300 µm gap. The vacuum window width, by contrast, *narrows* as the gap closes, from 4.0 mbar at 300 µm to 48.0 mbar at 25 µm, exactly as the $4\sigma/H_u$ scaling of Eq. (7) requires. Closing the gap buys speed and costs vacuum latitude, and that trade-off is quantitative rather than a matter of judgement.

V. DISCUSSION

A. Reading a window diagram

The practical value of Fig. 1 is not that it marks a point operable, but that it shows which constraint is binding and by how much. A process at a margin of 4.8 to the low-flow limit and 331 to air entrainment is not merely 'in the window'; it is a process whose only nearby boundary is the bead, and whose throughput can be raised until that boundary is reached and no further. That is actionable in a way that a pass/fail verdict is not.

It also reframes a common troubleshooting pattern. A laboratory that sees ribbing or breaklines on thin films frequently responds by increasing the applied vacuum, on the reasonable intuition that vacuum stabilises the bead. Equation (5) says this cannot work when the low-flow limit is the binding constraint, because h_{min} is set by the downstream meniscus and the vacuum acts on the upstream one. The effective responses are to close the gap, slow the line, or thicken the film - and the window diagram says which of the three has the most room.

B. Implications for perovskite minimodules

Three consequences bear directly on scalable perovskite fabrication.

Throughput for sub-micrometre absorbers is bead-limited, not wetting-limited. The perovskite absorber is the thinnest demanding layer in the stack, typically 500-800 nm dry and therefore a few micrometres wet. Section IV C showed that for such a film the low-flow limit binds a factor of 31 below the air-entrainment speed, and that the gap between the two widens as the film is thinned. Roadmaps that project line speeds from contact-line physics alone will be optimistic by a large factor; the relevant lever is the coating gap.

Gap control is a throughput specification, not only a uniformity specification. Because $h_{min} \propto H_d$, the gap enters the attainable speed linearly (Fig. 6). A head capable of a reproducible 50 µm gap permits roughly twice the line speed of one limited to 100 µm at the same thickness. Gap tolerance is usually specified from cross-web uniformity requirements; these results suggest it should be specified from throughput requirements too.

Dilution and speed are not interchangeable thickness knobs in blade coating. Section IV F showed that concentration moves the thickness-speed curve without moving the regime crossover, whereas speed moves along the curve and can cross it. Two processes depositing the same thickness, one by dilution and one by speed, may sit on different branches and dry by different mechanisms. Since the drying pathway governs intermediate-phase formation and ultimately the grain structure of the perovskite [38-41,44,57], this is not a cosmetic distinction, and the branch should be reported alongside the speed.

More broadly, the models implemented here concern what lands on the substrate, not what it becomes. For perovskites the relationship between a wet film and the final polycrystalline absorber is mediated by solvent coordination, intermediate phases, quenching and annealing [38-44], none of which pvcoat addresses. The package should be understood as fixing the boundary conditions within which that chemistry then operates - a necessary input to process design, and a long way from a sufficient one.

C. Relation to existing tools

Open scientific software in photovoltaics has concentrated on device and system modelling - pvlib for system performance [45], SCAPS and similar for device physics - and on data standards for stability testing [46,47]. Process-side tooling has lagged. Computational fluid dynamics solutions of the coating bead exist and are far more capable than the lubrication-level models implemented here [21,48,61], but they require meshing, a licence or substantial setup, and expertise that a materials group typically does not have in-house; they answer a different question, at a different cost, on a different timescale. Related lubrication-level treatments exist for the neighbouring processes of roll, gravure and printing-nip transfer [58,59,60], and the same style of analysis could usefully be extended to them. pvcoat occupies the gap between a paper-and-pencil estimate and a full simulation: it is the calculation a process engineer should do before deciding whether a simulation is warranted.

VI. LIMITATIONS

We state these numbered so they can be cited in review or in an issue report. They are not caveats appended to a result; they are the boundary of what the work supports.

1. No experimental validation. pvcoat reproduces published equations and checks itself analytically. No coating trial has been run against its predictions. Until that is done it is a calculator for published models and must be described as such. This is the most important limitation in this list and the first item of the roadmap below.

2. Newtonian inks only. Every model assumes a single viscosity. Concentrated perovskite precursors, and any ink carrying polymer or colloidal content, are not Newtonian, and the effect is not captured by substituting an apparent viscosity. Khandavalli and Rothstein show that at matched shear-rate-dependent capillary number a shear-thinning fluid requires a wet film roughly twice as thick before the coating is stable, while shear thickening reduces the minimum thickness by of order 10% [23,49].

3. Three failure modes of many. The low-flow limit, air entrainment, and the vacuum bounds are implemented. Ribbing, rivulets, breaklines, dripping, bubble entrainment from the feed, and streaklines from lip defects are not. A point pvcoat calls operable can still fail by a mechanism pvcoat does not model, and in published experiments the observed failure mode changes with gap and ink within a region a three-limit model would pass [22,23].

4. Two-dimensional, steady and isothermal. No cross-web variation, no edge effects, no start-stop transients, no substrate temperature gradient. Edge beads and the first and last centimetres of a stripe are precisely where minimodule yield is lost, and the package is silent on them.

5. No inertial regime. The viscocapillary models are asymptotic in Ca and Re. Above Ca $\approx$ 1, the measured minimum wet thickness departs from the two-thirds scaling and at higher Ca and Re falls again [21,22], so pvcoat is *conservative* there. Carvalho and Kheshgi's high-Ca branch is not implemented.

6. Perfect wetting assumed. Contact angles do not appear; both menisci are pinned at lip corners. Imperfect wetting raises the thickness at which the low-flow limit is reached, so a poorly wetting substrate is worse than predicted. This matters directly for coating onto self-assembled-monolayer-treated electrodes, where substrate energy is the variable being engineered [50,51].

7. The dry thickness is a floor. Equation (2) assumes a fully dense film and no solid lost in drying or annealing. Residual porosity, retained solvent, and volume change on crystallisation all make the real film thicker, and for perovskites the intermediate-phase route is not a volume-conserving dry-down. This is likely the largest single source of discrepancy between a predicted and a measured layer.

8. Drying is not modelled. There is no treatment of solvent removal, gas-knife or vacuum quenching, nucleation, or crystallisation.

9. The blade evaporation flux must be measured. E is an apparatus constant, not a material property, and every evaporation regime and crossover number is only as good as that measurement.

10. No uncertainty propagation. A viscosity known to ±20% gives a low-flow limit known to about ±13% through the two-thirds power. pvcoat reports point estimates and propagates neither input nor output uncertainty.

11. Two empirical coefficients carry external provenance. The Gutoff-Kendrick prefactor and exponent, and the plateau thresholds used in the validity checks, are taken from the literature rather than re-derived. Both are module constants and are documented as such.

A. Roadmap

The priority is validation. A v0.2 release should carry digitised operating-window data from published figures [21-23] with the agreement reported quantitatively, together with the author's own coating trials: predicted against measured thickness across a speed sweep for both heads. Ten well-characterised points would change what the package is. Beyond that, the ordered programme is the high-Ca and plateau branches of the low-flow limit, non-Newtonian rheology through a power-law or Carreau fit evaluated at the gap shear rate, additional failure modes beginning with ribbing, uncertainty propagation, and contact angles in the meniscus pinning conditions. Curtain, slide, gravure and roll coating, device physics and process optimisation are deliberately out of scope.

VII. CONCLUSIONS

We have presented pvcoat, an open-source Python implementation of the published operability limits for slot-die and blade coating, aimed at the scalable deposition of perovskite photovoltaic layers. The package computes wet and dry film thickness, the low-flow limit, the coating-bead vacuum window, the Landau-Levich thickness, the air-entrainment speed and the blade-coating regime crossover from process parameters, and attaches an explicit validity warning to any result computed outside the range its model was derived for. It has no runtime dependencies, accepts practitioner units, and traces every result to a published source. Seven closed-form identities are satisfied to machine precision across a 78-test suite.

Two results of the framework carry directly into practice. The low-flow limit rises as the two-thirds power of the capillary number, so for the sub-micrometre wet films that perovskite absorbers require, line speed is limited by the coating bead at a value a factor of 31 below the air-entrainment ceiling - and the effective lever on throughput is the coating gap, not the contact line. Separately, the speed at which a blade-coated film is thinnest is independent of solids loading, because the solids volume fraction cancels from both competing branches, so dilution and speed are not interchangeable thickness knobs.

The package is verified but not validated. It reproduces published equations correctly; whether those equations describe any particular coater remains an experimental question, and we have set out the programme that would answer it. We release the code, its tests and the scripts that reproduce every figure and table here in the hope that the validation is carried out by more laboratories than ours.

DATA AND CODE AVAILABILITY

The archived release of pvcoat v0.1.0, including the source, the test suite, the documentation, the verification record and the analysis scripts that generate every figure and table in this paper, is deposited at Zenodo under https://doi.org/10.5281/zenodo.22270164 and is released under the MIT licence. The machine-readable statistics file from which every number in this manuscript is interpolated is included in that deposit. No experimental data are reported in this work.

AUTHOR DECLARATIONS

Conflicts of interest. The author declares no competing interests.

AI assistance. Code implementation, documentation drafting, literature triage, and manuscript preparation were assisted by a large language model. Every modelling decision, empirical coefficient, and validity threshold is the author's, and every number reported here was verified by the author against the source publications before release; the verification record is included in the archived deposit.

ACKNOWLEDGMENTS

The author thanks colleagues at the Wright Center for Photovoltaics Innovation and Commercialization for discussions on scalable perovskite deposition.

REFERENCES

[1] A. Kojima, K. Teshima, Y. Shirai, and T. Miyasaka, Organometal halide perovskites as visible-light sensitizers for photovoltaic cells, J. Am. Chem. Soc. 131, 6050 (2009).

[2] M. A. Green, A. Ho-Baillie, and H. J. Snaith, The emergence of perovskite solar cells, Nat. Photonics 8, 506 (2014).

[3] National Renewable Energy Laboratory, Best Research-Cell Efficiency Chart (2026), https://www.nrel.gov/pv/cell-efficiency.html.

[4] J.-P. Correa-Baena, M. Saliba, T. Buonassisi, M. Gratzel, A. Abate, W. Tress, and A. Hagfeldt, Promises and challenges of perovskite solar cells, Science 358, 739 (2017).

[5] N.-G. Park and K. Zhu, Scalable fabrication and coating methods for perovskite solar cells and solar modules, Nat. Rev. Mater. 5, 333 (2020).

[6] Z. Li, T. R. Klein, D. H. Kim, M. Yang, J. J. Berry, M. F. A. M. van Hest, and K. Zhu, Scalable fabrication of perovskite solar cells, Nat. Rev. Mater. 3, 18017 (2018).

[7] Y. Rong, Y. Hu, A. Mei, H. Tan, M. I. Saidaminov, S. I. Seok, M. D. McGehee, E. H. Sargent, and H. Han, Challenges for commercializing perovskite solar cells, Science 361, eaat8235 (2018).

[8] A. Extance, The reality behind solar power's next star material, Nature 570, 429 (2019).

[9] F. C. Krebs, Fabrication and processing of polymer solar cells: A review of printing and coating techniques, Sol. Energy Mater. Sol. Cells 93, 394 (2009).

[10] F. C. Krebs, Polymer solar cell modules prepared using roll-to-roll methods: Knife-over-edge coating, slot-die coating and screen printing, Sol. Energy Mater. Sol. Cells 93, 465 (2009).

[11] R. R. Sondergaard, M. Hosel, and F. C. Krebs, Roll-to-roll fabrication of large area functional organic materials, J. Polym. Sci. B 51, 16 (2013).

[12] R. Sondergaard, M. Hosel, D. Angmo, T. T. Larsen-Olsen, and F. C. Krebs, Roll-to-roll fabrication of polymer solar cells, Mater. Today 15, 36 (2012).

[13] D. Vak, K. Hwang, A. Faulks, Y.-S. Jung, N. Clark, D.-Y. Kim, G. J. Wilson, and S. E. Watkins, 3D printer based slot-die coater as a lab-to-fab translation tool for solution-processed solar cells, Adv. Energy Mater. 5, 1401539 (2015).

[14] K. Hwang, Y.-S. Jung, Y.-J. Heo, F. H. Scholes, S. E. Watkins, J. Subbiah, D. J. Jones, D.-Y. Kim, and D. Vak, Toward large scale roll-to-roll production of fully printed perovskite solar cells, Adv. Mater. 27, 1241 (2015).

[15] L. D. Landau and B. G. Levich, Dragging of a liquid by a moving plate, Acta Physicochim. URSS 17, 42 (1942).

[16] B. V. Derjaguin, On the thickness of the liquid film adhering to the walls of a vessel after emptying, Acta Physicochim. URSS 20, 349 (1943).

[17] K. J. Ruschak, Limiting flow in a pre-metered coating device, Chem. Eng. Sci. 31, 1057 (1976).

[18] B. G. Higgins and L. E. Scriven, Capillary pressure and viscous pressure drop set bounds on coating bead operability, Chem. Eng. Sci. 35, 673 (1980).

[19] K. Y. Lee, L. D. Liu, and T.-J. Liu, Minimum wet thickness in extrusion slot coating, Chem. Eng. Sci. 47, 1703 (1992).

[20] S. F. Kistler and P. M. Schweizer, eds., Liquid Film Coating: Scientific Principles and Their Technological Implications (Chapman and Hall, London, 1997).

[21] M. S. Carvalho and H. S. Kheshgi, Low-flow limit in slot coating: Theory and experiments, AIChE J. 46, 1907 (2000).

[22] Y.-R. Chang, H.-M. Chang, C.-F. Lin, T.-J. Liu, and P.-Y. Wu, Three minimum wet thickness regions of slot die coating, J. Colloid Interface Sci. 308, 222 (2007).

[23] S. Khandavalli and J. P. Rothstein, The effect of shear-thickening on the stability of slot-die coating, AIChE J. 62, 4536 (2016).

[24] R. Burley and B. S. Kennedy, An experimental study of air entrainment at a solid/liquid/gas interface, Chem. Eng. Sci. 31, 901 (1976).

[25] T. D. Blake and K. J. Ruschak, A maximum speed of wetting, Nature 282, 489 (1979).

[26] T. D. Blake, The physics of moving wetting lines, J. Colloid Interface Sci. 299, 1 (2006).

[27] E. B. Gutoff and C. E. Kendrick, Dynamic contact angles, AIChE J. 28, 459 (1982).

[28] O. Cohu and H. Benkreira, Air entrainment in angled dip coating, Chem. Eng. Sci. 53, 533 (1998).

[29] H. Benkreira and M. I. Khan, Air entrainment in dip coating under reduced air pressures, Chem. Eng. Sci. 63, 448 (2008).

[30] E. Vandre, M. S. Carvalho, and S. Kumar, Delaying the onset of dynamic wetting failure through meniscus confinement, J. Fluid Mech. 707, 496 (2012).

[31] X. Ding, J. Liu, and T. A. L. Harris, A review of the operating limits in slot die coating processes, AIChE J. 62, 2508 (2016).

[32] S. J. Weinstein and K. J. Ruschak, Coating flows, Annu. Rev. Fluid Mech. 36, 29 (2004).

[33] K. J. Ruschak, Coating flows, Annu. Rev. Fluid Mech. 17, 65 (1985).

[34] M. Le Berre, Y. Chen, and D. Baigl, From convective assembly to Landau-Levich deposition of multilayered phospholipid films of controlled thickness, Langmuir 25, 2554 (2009).

[35] Y. Diao, L. Shaw, Z. Bao, and S. C. B. Mannsfeld, Morphology control strategies for solution-processed organic semiconductor thin films, Energy Environ. Sci. 7, 2145 (2014).

[36] Y. Diao, B. C.-K. Tee, G. Giri, J. Xu, D. H. Kim, H. A. Becerril, R. M. Stoltenberg, T. H. Lee, G. Xue, S. C. B. Mannsfeld, and Z. Bao, Solution coating of large-area organic semiconductor thin films with aligned single-crystalline domains, Nat. Mater. 12, 665 (2013).

[37] S. Patel and Y. Diao, Multiscale assembly of solution-processed organic electronics: The critical roles of confinement, fluid flow, and interfaces, Nanotechnology 29, 044004 (2018).

[38] Y. Deng, E. Peng, Y. Shao, Z. Xiao, Q. Dong, and J. Huang, Scalable fabrication of efficient organolead trihalide perovskite solar cells with doctor-bladed active layers, Energy Environ. Sci. 8, 1544 (2015).

[39] Y. Deng, X. Zheng, Y. Bai, Q. Wang, J. Zhao, and J. Huang, Surfactant-controlled ink drying enables high-speed deposition of perovskite films for efficient photovoltaic modules, Nat. Energy 3, 560 (2018).

[40] Y. Deng, C. H. Van Brackle, X. Dai, J. Zhao, B. Chen, and J. Huang, Tailoring solvent coordination for high-speed, room-temperature blading of perovskite photovoltaic films, Sci. Adv. 5, eaax7537 (2019).

[41] N. J. Jeon, J. H. Noh, Y. C. Kim, W. S. Yang, S. Ryu, and S. I. Seok, Solvent engineering for high-performance inorganic-organic hybrid perovskite solar cells, Nat. Mater. 13, 897 (2014).

[42] M. Yang, Z. Li, M. O. Reese, O. G. Reid, D. H. Kim, S. Siol, T. R. Klein, Y. Yan, J. J. Berry, M. F. A. M. van Hest, and K. Zhu, Perovskite ink with wide processing window for scalable high-efficiency solar cells, Nat. Energy 2, 17038 (2017).

[43] A. D. Taylor, Q. Sun, K. P. Goetz, Q. An, T. Schramm, Y. Hofstetter, M. Litterst, F. Paulus, and Y. Vaynzof, A general approach to high-efficiency perovskite solar cells by any antisolvent, Nat. Commun. 12, 1878 (2021).

[44] A. F. Routh, Drying of thin colloidal films, Rep. Prog. Phys. 76, 046603 (2013).

[45] W. F. Holmgren, C. W. Hansen, and M. A. Mikofski, pvlib python: A python package for modeling solar energy systems, J. Open Source Softw. 3, 884 (2018).

[46] M. V. Khenkin, E. A. Katz, A. Abate, et al., Consensus statement for stability assessment and reporting for perovskite photovoltaics based on ISOS procedures, Nat. Energy 5, 35 (2020).

[47] M. O. Reese, S. A. Gevorgyan, M. Jorgensen, et al., Consensus stability testing protocols for organic photovoltaic materials and devices, Sol. Energy Mater. Sol. Cells 95, 1253 (2011).

[48] O. J. Romero, W. J. Suszynski, L. E. Scriven, and M. S. Carvalho, Low-flow limit in slot coating of dilute solutions of high molecular weight polymer, J. Non-Newtonian Fluid Mech. 118, 137 (2004).

[49] S. Khandavalli and J. P. Rothstein, Extensional rheology of shear-thickening fumed silica nanoparticles dispersed in an aqueous polyethylene oxide solution, J. Rheol. 58, 411 (2014).

[50] A. Al-Ashouri, E. Kohnen, B. Li, et al., Monolithic perovskite/silicon tandem solar cell with >29% efficiency by enhanced hole extraction, Science 370, 1300 (2020).

[51] Q. Jiang, J. Tong, Y. Xian, et al., Surface reaction for efficient and stable inverted perovskite solar cells, Nature 611, 278 (2022).

[52] D. Quere, Fluid coating on a fiber, Annu. Rev. Fluid Mech. 31, 347 (1999).

[53] D. Bonn, J. Eggers, J. Indekeu, J. Meunier, and E. Rolley, Wetting and spreading, Rev. Mod. Phys. 81, 739 (2009).

[54] J. H. Snoeijer and B. Andreotti, Moving contact lines: Scales, regimes, and dynamical transitions, Annu. Rev. Fluid Mech. 45, 269 (2013).

[55] R. Krechetnikov and G. M. Homsy, Experimental study of substrate roughness and surfactant effects on the Landau-Levich law, Phys. Fluids 17, 102108 (2005).

[56] M. Maleki, M. Reyssat, F. Restagno, D. Quere, and C. Clanet, Landau-Levich menisci, J. Colloid Interface Sci. 354, 359 (2011).

[57] R. D. Deegan, O. Bakajin, T. F. Dupont, G. Huber, S. R. Nagel, and T. A. Witten, Capillary flow as the cause of ring stains from dried liquid drops, Nature 389, 827 (1997).

[58] H. Benkreira and R. Patel, Direct gravure roll coating, Chem. Eng. Sci. 48, 2329 (1993).

[59] S. Kumar, Liquid transfer in printing processes: Liquid bridges with moving contact lines, Annu. Rev. Fluid Mech. 47, 67 (2015).

[60] P. H. Gaskell, G. E. Innes, and M. D. Savage, An experimental investigation of meniscus roll coating, J. Fluid Mech. 355, 17 (1998).

[61] M. S. Carvalho and H. S. Kheshgi, Three-dimensional stability analysis of free surface flows: Application to forward deformable roll coating, J. Comput. Phys. 151, 534 (1999).

[62] T. M. Schmidt, T. T. Larsen-Olsen, J. E. Carle, D. Angmo, and F. C. Krebs, Upscaling of perovskite solar cells: Fully ambient roll processing of flexible perovskite solar cells with printed back electrodes, Adv. Energy Mater. 5, 1500569 (2015).

[63] F. Di Giacomo, S. Shanmugam, H. Fledderus, et al., Up-scalable sheet-to-sheet production of high efficiency perovskite module and solar cells on 6-inch substrate using slot die coating, Sol. Energy Mater. Sol. Cells 181, 53 (2018).

[64] Y. Galagan, F. Di Giacomo, H. Gorter, G. Kirchner, I. de Vries, R. Andriessen, and P. Groen, Roll-to-roll slot die coated perovskite for efficient flexible solar cells, Adv. Energy Mater. 8, 1801935 (2018).

[65] D. Burkitt, J. Searle, and T. Watson, Perovskite solar cells in N-I-P structure with four slot-die-coated layers, R. Soc. Open Sci. 5, 172158 (2018).

[66] G. Cotella, J. Baker, D. Worsley, F. De Rossi, C. Pleydell-Pearce, M. Carnie, and T. Watson, One-step deposition by slot-die coating of mixed lead halide perovskite for photovoltaic applications, Sol. Energy Mater. Sol. Cells 159, 362 (2017).

[67] Y. Huang et al., Meniscus-modulated blade coating enables high-quality alpha-phase formamidinium lead triiodide crystals and efficient perovskite minimodules, Joule 8, 2313 (2024).

[68] C. R. Harris, K. J. Millman, S. J. van der Walt, et al., Array programming with NumPy, Nature 585, 357 (2020).

[69] P. Virtanen, R. Gommers, T. E. Oliphant, et al., SciPy 1.0: Fundamental algorithms for scientific computing in Python, Nat. Methods 17, 261 (2020).

[70] J. D. Hunter, Matplotlib: A 2D graphics environment, Comput. Sci. Eng. 9, 90 (2007).

[71] M. L. Waskom, seaborn: Statistical data visualization, J. Open Source Softw. 6, 3021 (2021).

[72] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, et al., The FAIR Guiding Principles for scientific data management and stewardship, Sci. Data 3, 160018 (2016).

[73] A. M. Smith, D. S. Katz, and K. E. Niemeyer, Software citation principles, PeerJ Comput. Sci. 2, e86 (2016).

[74] G. Wilson, D. A. Aruliah, C. T. Brown, et al., Best practices for scientific computing, PLoS Biol. 12, e1001745 (2014).

[75] H. Anzt, F. Bach, S. Druskat, et al., An environment for sustainable research software in Germany and beyond, F1000Research 9, 295 (2021).

[76] M. Baker, 1,500 scientists lift the lid on reproducibility, Nature 533, 452 (2016).

[77] J. M. Perkel, Challenge to scientists: Does your ten-year-old code still run?, Nature 584, 656 (2020).