

AirBeam Mini PM_{2.5} calibration — final report

Prepared for: HabitatMap **Prepared by:** Chris Lim, Health & Urban Modeling Lab, University of Arizona **Date:** 2026-06-23
Status: revision 4 (final — incorporates HabitatMap firmware-engineer review; count-channel identities corrected for the colocation-firmware byte shift, §2.1)

Executive summary

We replaced the AB3-derived PM_{2.5} calibration in the AirBeam Mini (ABM) firmware with an ABM-specific equation fit on **five ABM units (3 normal-batch, 2 new-batch)** co-located at NYSDEC Queens College for **4.5 months** (Nov 2025 – Mar 2026). The recommended equation uses the Plantower particle-count channels (C0.3 and C2.5; the free-intercept variant adds C1) instead of the mass channel (A1) used by the current firmware. **Channel identities in this revision are corrected for a byte-shift bug in the colocation units' firmware (§2.1):** the channels the colocation exports labeled C.5 and C1 are physically C1 and C2.5. No fitted coefficient or performance number changes — only the channel names and the firmware wiring instructions. Performance against the DEC reference improved as follows:

Metric (hourly)	AB3 firmware (current)	C0.3 + C2.5 (proposed, floor 1.23)	C0.3 + C1 + C2.5 (free-intercept variant)
R ²	0.889	0.914	0.925
RMSE (µg/m ³)	2.36	2.07	1.94
MAE (µg/m ³)	1.63	1.38	1.32
Bias (µg/m ³)	-0.03	-0.48	0.00
Floor (µg/m ³)	0 (can read negative)	1.23	2.42

The recommended equation is a two-channel form with the intercept fixed at 1.23345 µg/m³, the maximum value at which clean-air output (C0.3 < 50, C2.5 ≈ 0) still rounds to 1 µg/m³ in display. Clean ambient at urban temperate sites is typically 1–3 µg/m³, so a 2.4 µg/m³ floor reads as miscalibrated, but allowing the floor to drop to exactly 1.0 sacrifices fit unnecessarily — the 1.23 choice preserves the visible “1” display while recovering most of the lost R². The higher-R² free-intercept variant is documented as an alternative (§4.1) if the lower floor is not desired. Four of the five units in the deliverable cohort improve over AB3 under the floor-1.23 form (and all five improve under the free-intercept variant); the trade-off vs the free-intercept fit is a slight under-prediction at clean-air times (mean bias -0.48 µg/m³).

We also revise the PM₁ stage in the same firmware patch. Instead of computing PM₁ from the raw mass channel (which produces PM₁ = 0 in 48.7% of low-concentration hours and PM₁ > PM_{2.5} in 18.1% of high-concentration hours), PM₁ is derived from the new calibrated PM_{2.5} via a smooth monotone ratio function fitted to the AB3 stage-2 inverse. This preserves HabitatMap’s original PM₁ ≤ PM_{2.5} logic, eliminates both pathologies in 100% of hours, and avoids the breakpoint discontinuities present in the original piecewise design.

We tested and rejected two other candidate paths for PM_{2.5}: (a) adding LaGuardia ASOS humidity (Barkjohn-style) gave Δ R² ≈ +0.01 and is not the dominant residual; (b) using the mass channel (A1) gave near-identical fit but is more prone to noise-floor saturation.

1. Background

Why this work was needed. The current ABM firmware ships with a piecewise $PM_{2.5}$ calibration ported directly from the AirBeam3 (AB3). That equation was developed for a different sensor enclosure with a different aerosol pathway and was never validated against ABM data at scale. HabitatMap requested an ABM-specific replacement before the firmware freeze for the next chip programming run.

Architectural inheritance from AB3. This calibration exercise refit *coefficients*, not the calibration *architecture*. Specifically, we inherit from AB3: (a) the use of a single primary input channel feeding a piecewise transfer function, (b) the linked $PM_1 \leftrightarrow PM_{2.5}$ output structure with $PM_1 \leq PM_{2.5}$ enforced by construction, (c) the ratio relationship between PM_1 and $PM_{2.5}$ at given concentrations (the smooth-ratio function in §4.4 is fit to the AB3 stage-2 inverse, not to a PM_1 reference), and (d) the deployment philosophy of “fit a static transfer function on a single co-located reference dataset and ship as global firmware default.” Only the input channel (count vs mass) and the empirical coefficients are new. A first-principles redesign of the calibration architecture (e.g., RH-corrected polynomial, ML model, online adaptive estimate, multi-sensor data fusion) was out of scope for this iteration.

Sensor inventory. Eleven AirBeam Mini units (Plantower PMS7003 inside HabitatMap’s compact enclosure) were available for analysis. Three production batches are present:

Cohort	Units	Batch description
Low	0684, 7014, 713C, 71E8, 7278	First-batch units (1,000-unit production run, 2024–2025) with a Plantower internal smoothing filter that suppresses output below $\sim 5 \mu\text{g}/\text{m}^3$. Per HabitatMap, this batch will not be reproduced.
Normal	9208, 92C0, 9224	Pre-2024 batch; longstanding production line.
New_batch	71F8, 722C	Current production batch sent to replace the low batch. Behaves identically to “normal” — the production target.

Per HabitatMap’s direction, the **deliverable cohort is normal + new_batch only** (5 units). The low batch is retired and out of scope.

Reference instrument. NYSDEC Queens College BAM-1020 (β -attenuation), 1-minute $PM_{2.5}$. Source: NYSDEC export `QCNcore_PM_20251031-20260317.xlsx`.

2. Data and preprocessing

Item	Value
ABM raw data	5 AirCasting CSV exports (deliverable cohort), 1-min, 12 PM channels per unit
ABM total minutes	630,082 deliverable-cohort minutes joined to reference
DEC reference minutes	195,780 (PM _{2.5} non-missing: 194,066)
Window	2025-11-01 00:00 to 2026-03-16 23:59 NY local (HabitatMap-defined)
Joined deliverable rows (1-hr, ≥30 valid mins/hr)	10,481

Timestamp handling. Both AirCasting CSV exports (verified against DEC reference using the precipitous PM_{2.5} drop on 2026-02-28 11:00 NY local) and the DEC export are in NY local time. ABM timestamps are parsed directly with `tz = "America/New_York"`. One implementation note: `data.table::fread` auto-parses ISO 8601 timestamp strings as UTC, so the Timestamp column must be read as character (`colClasses = c(Timestamp = "character")`) and parsed manually with the correct tz; otherwise the join misaligns by 5 hours during EST.

Aggregation. All calibration fitting is done on hourly means using forward-aligned hours: minutes 0–59 of each clock hour are averaged into the timestamp marked HH:00, matching EPA convention for regulatory hourly averages (BAM and FRM/FEM monitors report the hour starting at HH:00 as the average of HH:00:00 through HH:59:59). Both ABM and DEC reference series are aggregated identically, so the averaging windows align by construction. We require at least 30 valid 1-minute observations per hour for that hour to be retained. 1-min calibration was tested and produces noisier fits with $R^2 \sim 0.10$ lower than hourly; the recommended equation is robust to either grain when applied in firmware.

Channel inventory. The AirCasting CSV exports we received contained 12 PM channels: ambient mass (A1, A2.5, A10), lab/standard mass (S1, S2.5, S10 — identical to ambient in our data), and particle counts labeled at six size cuts (C.3, C.5, C1, C2.5, C5, C10). The count labels are shifted one size cut from the physical channels — see §2.1; after correction the data contain five real count channels (C0.3, C1, C2.5, C5, C10) and one reserved field (labeled C10, uniformly zero — dropped). Both ambient and lab mass channels carry the same information (Spearman correlations are identical). The count channels are the underlying optical signal, while the mass channels are derived internally by the Plantower from the same count distribution; we tested all available forms.

2.1 Colocation-firmware channel shift (corrected in this revision)

During firmware-engineer review of revision 3, HabitatMap identified a byte-parsing bug in the firmware running on the colocation units (the bug is not present in the new production firmware). The Plantower communicates each measurement cycle as a 32-byte frame; the colocation firmware skipped the two bytes carrying the $\geq 0.5 \mu\text{m}$ count (frame field Data8). Every subsequent count value therefore landed one slot early, and the AirCasting export labels inherited the shift:

Export label	C.3	C.5	C1	C2.5	C5	C10
Physical channel	C0.3 (unaffected)	C1	C2.5	C5	C10	reserved (\equiv 0)

Two features of the colocation data independently confirm the shift. First, the column labeled C10 is identically zero in every row of every unit — it received the Plantower’s reserved trailing field. Second, the column labeled C.5 runs at ~5.5% of C.3 across units; a true C0.5/C0.3 ratio in ambient Plantower data is typically 25–30%, while ~5% is exactly where the C1/C0.3 ratio sits.

Consequences: (a) all channel references in this revision — tables, equations, figures — use the corrected physical identities; (b) **no numerical result changes** — every fitted coefficient, breakpoint, and performance statistic binds to a physical signal, not a label; (c) the mass channels (frame fields Data1–6) and the $\geq 0.3 \mu\text{m}$ count (Data7) precede the skipped field and were never affected, so Options B and C (C0.3-only) are wholly immune to the bug; (d) the production firmware parses the frame correctly, so the equation inputs must be wired to the true frame fields — see the wiring warning in §4.1; and (e) a true $\geq 0.5 \mu\text{m}$ count was never recorded during this colocation, so C0.5 is the one count channel we could not evaluate (§3).

3. Channel selection rationale

Per-channel hourly Pearson correlations with DEC $\text{PM}_{2.5}$, deliverable cohort (channel identities corrected per §2.1):

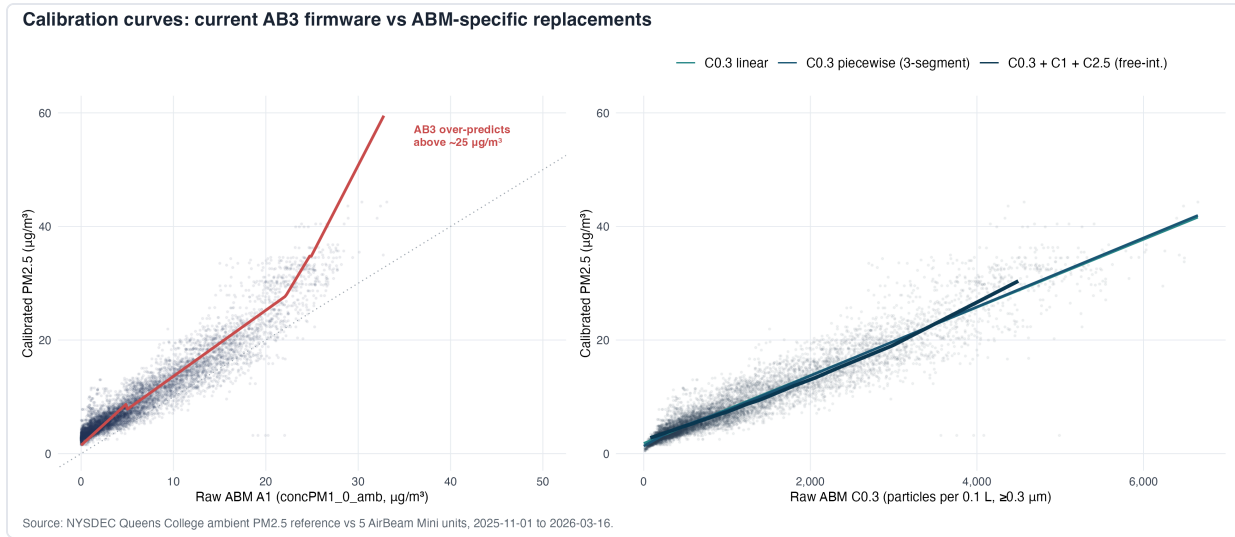
Channel	Correlation
A1	0.93
A2.5	0.94
A10	0.94
C0.3	0.95
C1	0.94
C2.5	0.91
C5	0.83

C0.3 has the highest single-channel correlation with DEC $\text{PM}_{2.5}$, consistent with the view that the count channel is the most responsive raw signal. We tested 45+ alternative model forms (mass-only, count-only, mixed, with and without interactions and non-linear transforms). All converge in a tight band of $R^2 = 0.93\text{--}0.94$, reflecting that the seven non-zero ABM channels are highly inter-correlated (pairwise $r = 0.88\text{--}1.00$) — they carry largely the same information. One channel could not be tested: the true C0.5 count was never recorded during this colocation (it is the field the colocation firmware skipped, §2.1). Given the 0.88–1.00 inter-channel correlations it is unlikely to carry independent information, but corrected-firmware units will report it, and a future colocation can confirm.

Multivariate gain. Adding C1 and C2.5 to C0.3 raises R^2 from 0.904 to 0.925 ($\Delta 0.021$). The largest tested multivariate (A1 + A2.5 + A10 + C0.3 + C1 + C2.5 + C5) reaches $R^2 = 0.935$ — only +0.010 over the recommended 3-channel form, at twice the parameter count. The recommended form captures essentially all the available information.

Noise floor behavior. The mass channels (A1, A25) read exact zero during low-PM minutes more often than the count channels do. During calibration fitting, both zero and small non-zero readings contribute to the OLS fit equally. But during deployment, a zero reading carries no usable signal: the firmware cannot distinguish “true $\text{PM} = 0$ ” from “true PM was small but below detection.” Count channels keep reporting low non-zero values at the same true concentrations, which the firmware can transform into sensible estimates — a deployment-relevant advantage even when the calibration R^2 is similar.

Architectural rationale. Calibrating against count channels skips one layer of vendor-internal Plantower processing and gives HabitatMap firmware direct access to the underlying optical signal. This is the same logic PurpleAir users apply to recover better calibration than the factory PM_{2.5} output.



Calibration curves: AB3 vs proposed

Figure 1. Left panel: current AB3 firmware curve (input: raw A1) overlaid on hourly Queens collocation. AB3 over-predicts above ~25 µg/m³ — the upper-segment slope of 1.28 was originally fit at Sterling Forest with an aerosol concentrator and does not transfer to this sensor. Right panel: three proposed C0.3-based equations on the same data. The multivariate curve shown is the free-intercept three-channel variant; it visually coincides with the recommended Option A (which differs only in its fixed 1.23 µg/m³ floor and dropped C1 term). It bends slightly at low concentrations to reflect the C1/C2.5 contribution but is otherwise close to linear in C0.3.

4. Recommended firmware equation (PM_{2.5})

We provide three options for the PM_{2.5} stage, ranked best-fit to simplest-firmware. All three are drop-in replacements for the AB3 stage-2 block. PM₁ is then derived from the calibrated PM_{2.5} via a smooth-ratio function — see §4.4.

4.0 Current ABM firmware (for reference)

The current ABM firmware uses a piecewise calibration ported directly from the AirBeam3. It takes raw ambient PM₁ (`concPM1_0_amb`) as input and chains stage-1 (raw → calibrated PM₁) and stage-2 (PM₁ → PM_{2.5}):

```

// === STAGE 1: raw A1 → calibrated PM1 ===
if (concPM1_0_amb < 22.14) {
  PM1_0 = (1.0566 * concPM1_0_amb);
} else if (concPM1_0_amb >= 22.14 && concPM1_0_amb < 37.95) {
  PM1_0 = (-30.9677) + (2.4554 * concPM1_0_amb);
} else if (concPM1_0_amb >= 37.95) {
  PM1_0 = (11.2026) + (1.3442 * concPM1_0_amb);
}

// === STAGE 2: calibrated PM1 → PM2.5 (THIS IS WHAT WE REPLACE) ===
if (PM1_0 < 5.16) {
  PM2_5 = (1.49) + (1.4 * PM1_0);
} else if (PM1_0 >= 5.16 && PM1_0 < 30.16) {
  PM2_5 = (2.00) + (1.1 * PM1_0);
} else if (PM1_0 >= 30.16 && PM1_0 < 126.16) {
  PM2_5 = (-4.00) + (1.2826 * PM1_0);
} else if (PM1_0 >= 126.16) {
  PM2_5 = (-24.00) + (1.441 * PM1_0);
}

```

The new firmware reorders the chain. Stage-2 ($PM_1 \rightarrow PM_{2.5}$) is replaced wholesale by one of the three count-channel options below. PM_1 is no longer computed from raw A1 via stage-1; instead it is derived from the calibrated $PM_{2.5}$ via a smooth-ratio function fitted to the AB3 stage-2 inverse (§4.4). This preserves HabitatMap's original $PM_1 \leq PM_{2.5}$ logic while eliminating both the noise-floor zeros from the raw-A1 path and the breakpoint discontinuities from the original piecewise design.

4.1 Option A — multivariate (recommended)

```

// Inputs (particles per 0.1 L, corrected-firmware Plantower frame):
// concC03 = count >=0.3 um (frame field Data7)
// concC25 = count >=2.5 um (frame field Data10)
// WIRING WARNING: the 0.211782 coefficient binds to the TRUE >=2.5 um
// count. The colocation exports labeled this channel "C1" (see report
// section 2.1). Wiring it to the true >=1.0 um count (Data9) instead
// would inflate output by roughly +10 ug/m3 at typical urban
// concentrations (true C1 counts run ~10x larger than true C2.5).
PM2_5 = 1.23345
  + 0.005157 * concC03
  + 0.211782 * concC25;
if (PM2_5 < 0.0) PM2_5 = 0.0;

```

$R^2 = 0.914$ · RMSE = 2.07 $\mu\text{g}/\text{m}^3$ · MAE = 1.38 · bias = -0.48 · 2 multiplications + 2 additions per evaluation. Floor = 1.23 $\mu\text{g}/\text{m}^3$.

This form was selected to satisfy HabitatMap's preference for a clean-air display of 1 $\mu\text{g}/\text{m}^3$. The intercept value 1.23345 is the maximum forced intercept at which the equation's clean-air output ($C0.3 < 50$, $C2.5$ small) still rounds to 1: at $C0.3 = 50$ and $C2.5 = 0$, predicted $PM_{2.5} = 1.49 \mu\text{g}/\text{m}^3$, just below the 1.5 boundary at which integer rounding would flip to 2. Pushing the floor

below 1.23 (e.g., to 1.0) sacrifices R² and RMSE without changing the displayed clean-air value. Pushing it above 1.30 crosses the rounding threshold and the display flips to 2.

The form was fit by ordinary least squares with the intercept fixed at 1.23345 and only positive count channels (C0.3 and C2.5) retained — non-negative least squares on the full C0.3 + C1 + C2.5 set automatically zeroes out the C1 coefficient at this intercept, because C1 information is largely redundant with C0.3 (pairwise r = 0.97). The surviving pair is also the physically sensible one: C0.3 carries the total fine-particle count, while C2.5 adds an independent coarse-tail signal that C0.3 barely sees — a more informative second regressor than the near-collinear C1 would have been.

Higher-floor variant (alternative). If a higher visible floor is acceptable and the OLS-optimal fit is preferred, the free-intercept three-channel form gives slightly tighter performance:

```
// Free-intercept variant, floor = 2.42 ug/m3
// Inputs (particles per 0.1 L, corrected-firmware Plantower frame):
// concC03 = count >=0.3 um (Data7), concC1 = count >=1.0 um (Data9),
// concC25 = count >=2.5 um (Data10)
PM2_5 = 2.3324
      + 0.004319 * concC03
      + 0.000470 * concC1
      + 0.274185 * concC25;
if (PM2_5 < 0.0) PM2_5 = 0.0;
```

R² = 0.925 · RMSE = 1.94 · MAE = 1.32 · bias = 0.00.

Variant	Floor (µg/m ³)	R ²	RMSE	MAE	Bias	Channels
A (recommended)	1.23	0.914	2.07	1.38	-0.48	C0.3, C2.5
A-free	2.42	0.925	1.94	1.32	0.00	C0.3, C1, C2.5

The trade-off in switching to the floor=1.23 form is a slight under-prediction at clean-air times (mean bias -0.48 µg/m³). This is the inherent cost of forcing the floor below the OLS-optimal intercept of 2.42: when true ambient is in the 0–2 µg/m³ range, the equation will read ~1.2–1.5 instead of 2.4. We judge this a worthwhile trade for the visible-floor=1 display behavior, but the higher-floor variant is documented as an option.

4.2 Option B — C0.3 piecewise (3-segment)

```
// Input: concC03 only (particles per 0.1 L, >=0.3 um - frame field
// Data7; unaffected by the channel shift in report section 2.1)
if ( concC03 < 356.00 ) {
    PM2_5 = 1.3276 + 0.008214 * concC03;
} else if ( concC03 >= 356.00 && concC03 < 1026.04 ) {
    PM2_5 = 2.3726 + 0.005278 * concC03;
} else if ( concC03 >= 1026.04 ) {
    PM2_5 = 1.5587 + 0.006071 * concC03;
}
if ( PM2_5 < 0.0 ) PM2_5 = 0.0;
```

$R^2 = 0.905$ · RMSE = 2.18 · MAE = 1.42.

Continuity at breakpoints. Unlike the original AB3 stage-2 (which has small downward jumps at $PM_1 = 5.16$ and 30.16), the segmented optimizer used to fit Option B enforces continuity at both breakpoints. Verification: at $concC03 = 356.00$, segment 1 gives $PM_{2.5} = 4.2518$ and segment 2 gives 4.2516 (gap = $0.0002 \mu\text{g}/\text{m}^3$); at $concC03 = 1026.04$, segment 2 gives 7.7880 and segment 3 gives 7.7878 (gap = $0.0003 \mu\text{g}/\text{m}^3$). No skipped output ranges and no downward jumps.

4.3 Option C — C0.3 linear (simplest)

```
// Input: concC03 only (Data7; unaffected by the section-2.1 shift)
PM2_5 = 1.8302 + 0.005979 * concC03;
if ( PM2_5 < 0.0 ) PM2_5 = 0.0;
```

$R^2 = 0.904$ · RMSE = 2.19 · MAE = 1.43. Statistically equivalent to Option B; preferred over B if firmware footprint matters.

4.4 PM_1 — derived from the new $PM_{2.5}$ via smooth-ratio function

Queens has no PM_1 reference, so the raw-A1 → calibrated- PM_1 mapping cannot be re-identified directly. Two pathologies in the current AB3 stage-1 path motivate a revision:

- **$PM_1 = 0$ in 48.7% of hours when DEC $PM_{2.5} < 5 \mu\text{g}/\text{m}^3$.** The mass channel (A1) sits at the sensor noise floor in 9% of all hours and 48% of low-PM hours. Users see $PM_1 = 0$ alongside reported $PM_{2.5} = 3\text{--}4 \mu\text{g}/\text{m}^3$.
- **$PM_1 > PM_{2.5}$ in 18.1% of hours when DEC $PM_{2.5} > 25 \mu\text{g}/\text{m}^3$.** Combination of AB3 stage-1's steep upper slope and the new $PM_{2.5}$ equation's slight under-prediction at high concentrations produces a physics-violating ordering.

The fix is to derive PM_1 from calibrated $PM_{2.5}$ using a monotone smooth ratio function. The recommended form is fitted to match the algebraic inverse of AB3 stage-2 over the empirical $PM_{2.5}$ distribution, preserving HabitatMap's original $PM_1 \leq PM_{2.5}$ logic without inheriting the AB3 piecewise discontinuities.

Recommended firmware (smooth ratio)

```
// PM1 derived from calibrated PM2.5 – smooth ratio
PM1_0 = PM2_5 * (0.855 - 0.818 * expf(-PM2_5 / 6.12f));
if (PM1_0 < 0.0f) PM1_0 = 0.0f;
```

The function is monotonically increasing in $PM_{2.5}$, returns $PM_1 \leq PM_{2.5}$ by construction, and has no breakpoint discontinuities. The fitted parameters were chosen to match the algebraic inverse of AB3 stage-2 across the empirical $PM_{2.5}$ distribution while removing the small breakpoint jumps that arose from inverting a piecewise function.

Empirical check, deliverable cohort, 10,481 hours (under the floor-1.23 $PM_{2.5}$ recommended in §4.1):

Reference $PM_{2.5}$ regime	n hours	PM_1 median ($\mu\text{g}/\text{m}^3$)	$PM_1 / PM_{2.5}$ ratio	% PM_1 rounds to 0	% $PM_1 > PM_{2.5}$
<5	4,134	0.97	0.34	10.4%	0%
5–12	3,827	3.73	0.57	0%	0%
12–25	2,077	12.63	0.79	0%	0%
>25	443	25.29	0.85	0%	0%

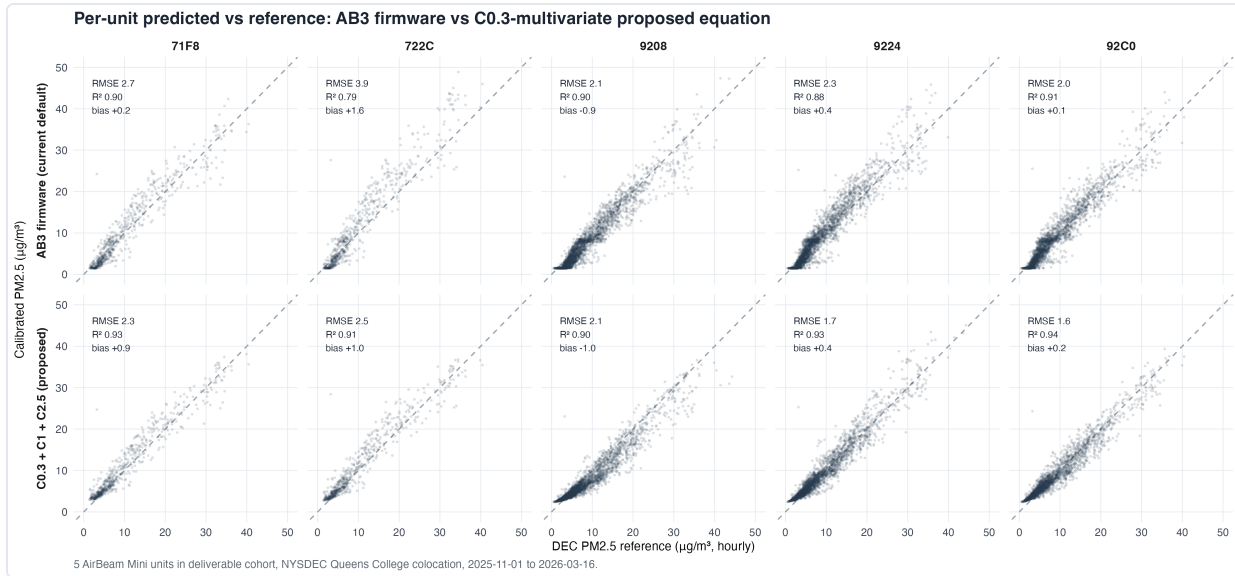
The $PM_1 > PM_{2.5}$ violation is fully eliminated (0% in every regime, vs 18.1% in the >25 regime under AB3). The PM_1 rounds-to-zero rate at low concentrations drops from 48.7% under AB3 to 10.4% under the new equation — a substantial improvement, though not perfectly zero. The residual zeros happen when $PM_{2.5}$ is near the floor of 1.23 $\mu\text{g}/\text{m}^3$ and the smooth-ratio PM_1 falls below 0.5; this represents genuinely clean-air conditions where true PM_1 is very small, not a sensor malfunction. The $PM_1/PM_{2.5}$ ratio rises smoothly from ~0.34 at low concentrations to 0.85 at moderate-high, consistent with the typical urban pattern (combustion sources dominate at higher loadings, shifting the size distribution toward submicron particles).

Caveat

PM_1 is no longer measured independently from the sensor. It is a deterministic function of calibrated $PM_{2.5}$. Users who need a true size-resolved PM_1 measurement should treat ABM PM_1 as a derived ratio output, not an independent sample. A direct PM_1 calibration awaits an ambient-aerosol collocation with a co-located PM_1 reference instrument.

5. Performance

5.1 Head-to-head against AB3



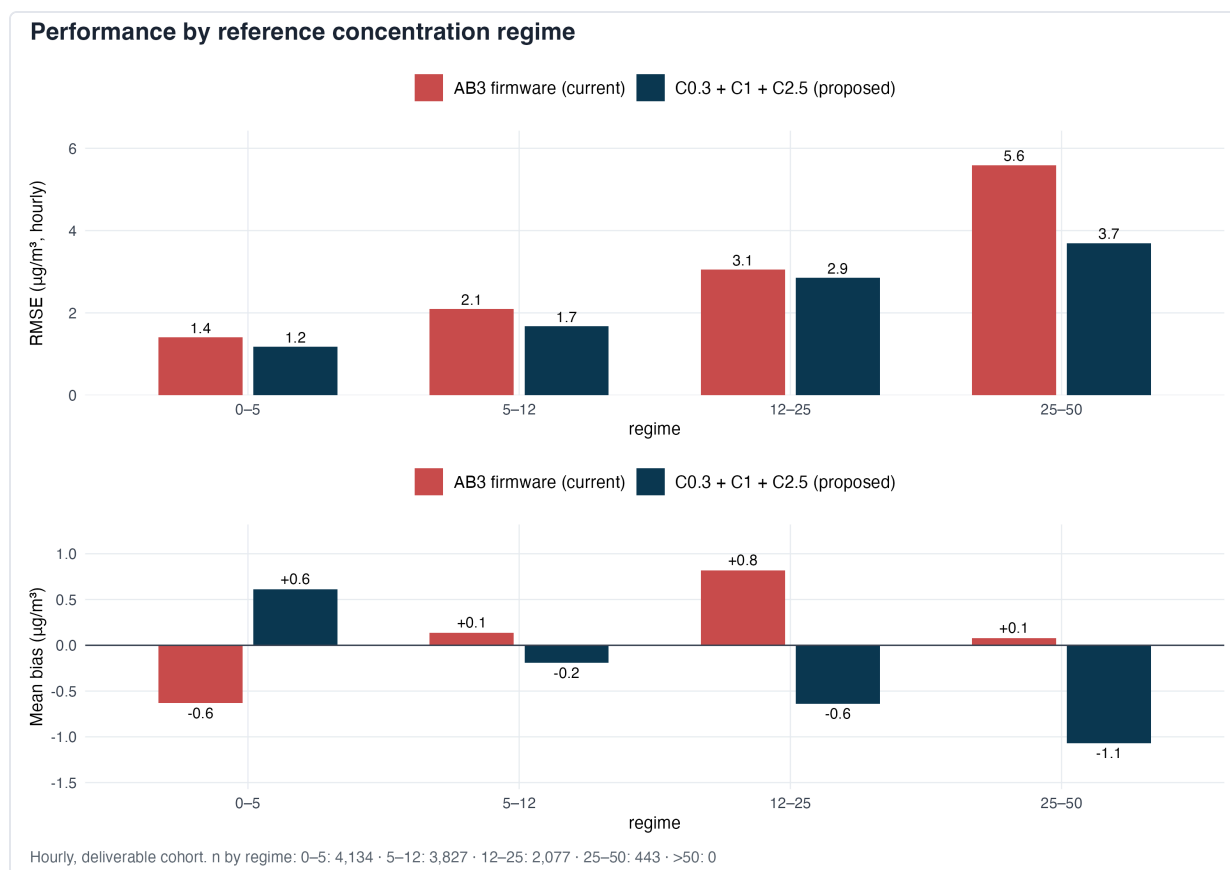
Per-unit head-to-head

Figure 2. Predicted vs DEC reference $PM_{2.5}$ for each of the 5 deliverable units, AB3 firmware (top row) versus the free-intercept three-channel form (bottom row, $R^2 = 0.925$). The figure illustrates the upper-bound performance available; the recommended floor-1.23 two-channel form is slightly tighter at low concentrations and slightly looser overall (per-unit RMSE in §5.1 table). Both rows hug the 1:1 line through the bulk of the colocation range, with the proposed equation visibly reducing scatter at moderate-to-high concentrations (12–35 $\mu\text{g}/\text{m}^3$) for the new-batch units (71F8, 722C).

Unit	Cohort	n hours	AB3 RMSE	New floor-1.23 RMSE	New floor-1.23 bias	Free-int RMSE	Median $PM_{2.5}$	Median PM_1 (derived)
71F8	new_batch	658	2.72	2.33	+0.35	2.33	6.79	3.97
722C	new_batch	658	3.89	2.81	+0.85	2.51	7.11	4.26
9208	normal	3,169	2.11	2.37	-1.55	2.13	4.56	2.13
9224	normal	3,169	2.35	1.82	+0.03	1.75	5.92	3.22
92C0	normal	2,827	2.04	1.67	-0.35	1.64	5.46	2.84

The recommended floor-1.23 form (C0.3 + C2.5) improves over AB3 in 4 of 5 units. Unit 9208 has lower AB3 RMSE (2.11) than the floor-1.23 form (2.37) — a consequence of forcing the intercept down for that low-mean unit; the free-intercept variant is also slightly tied with AB3 for 9208 (2.13). 722C shows the largest absolute gain in $PM_{2.5}$ RMSE under the floor-1.23 form (3.89 → 2.81, -28%). Median PM_1 (derived from cal $PM_{2.5}$ per §4.4) tracks at ~ 0.47 – $0.60\times$ the median $PM_{2.5}$ across units.

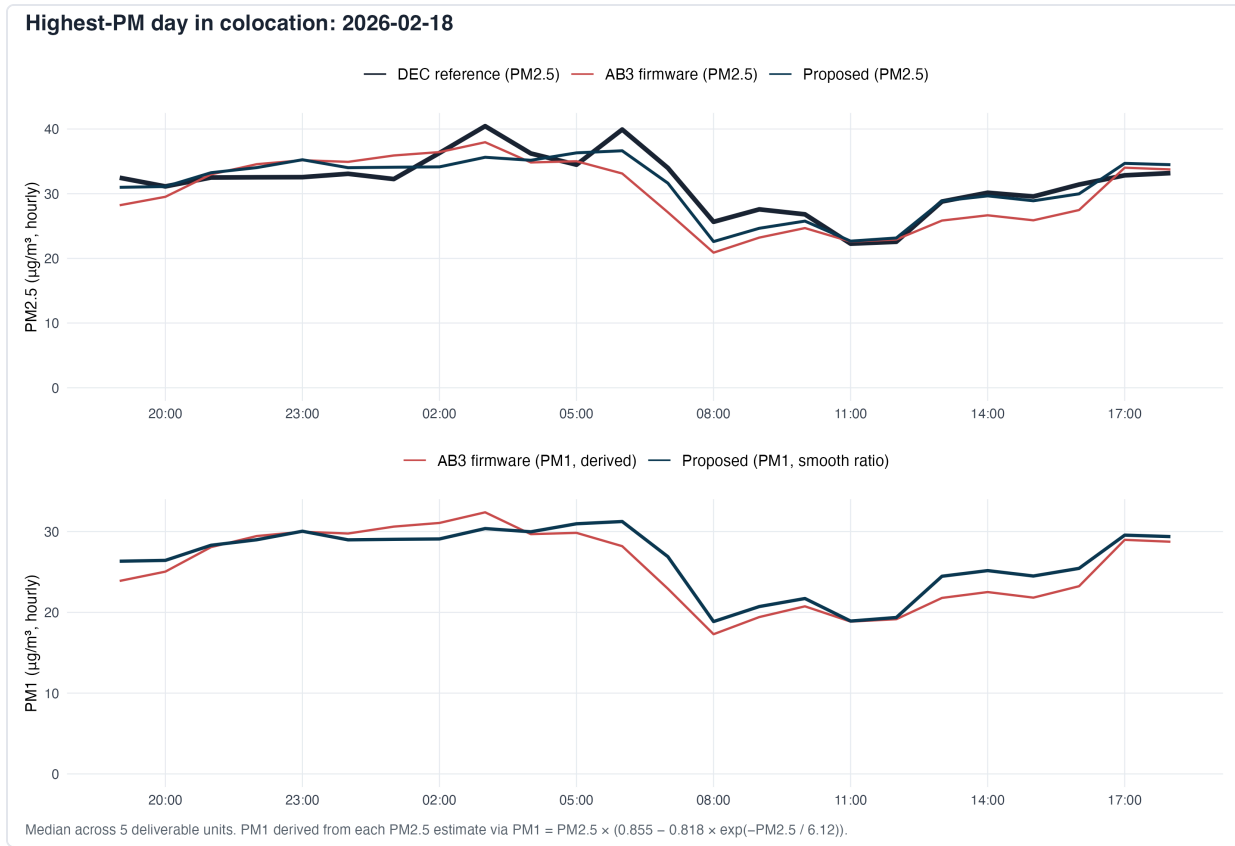
5.2 Performance by concentration regime



Performance by regime

Figure 3. Top: RMSE by reference-concentration regime. The proposed equation beats AB3 in every regime, with the largest absolute gain at 25–50 $\mu\text{g}/\text{m}^3$ (RMSE 5.6 \rightarrow 3.7). Bottom: bias direction (positive = over-prediction). AB3 has non-monotonic bias: -0.6 at $<5 \mu\text{g}/\text{m}^3$ (under-predicts), $+0.1$ at 5–12, $+0.8$ at 12–25, $+0.1$ at 25–50. The proposed equation has monotonic bias: $+0.6$ at <5 (over-predicts), -0.2 at 5–12, -0.6 at 12–25, -1.1 at 25–50 (gradually drifting under-predict). Both equations stay within $1.1 \mu\text{g}/\text{m}^3$ bias magnitude across the full range that holds 99% of field hours. The corresponding derived PM_{10} median by regime is 1.62 ($\text{PM}_{2.5} <5$), 3.96 (5–12), 12.45 (12–25), 25.03 (>25), tracking at ratios $0.43 \rightarrow 0.85$ as concentration rises.

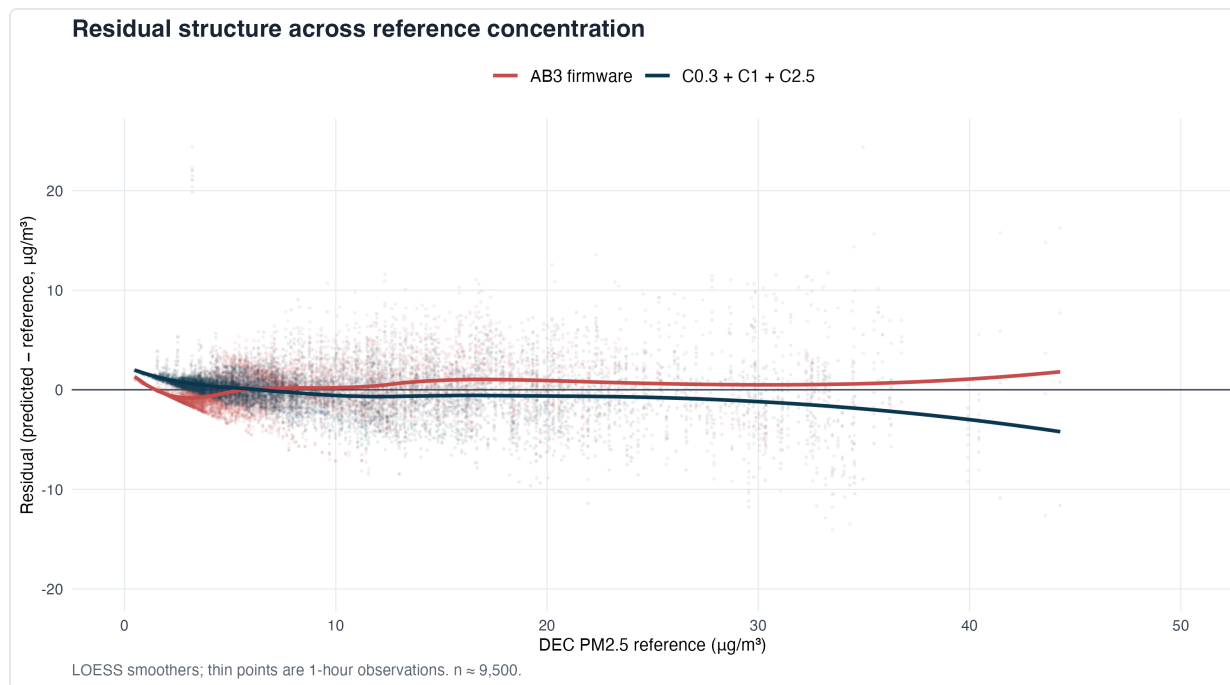
5.3 Time series on the highest-PM day



Peak day time series

Figure 4. Hourly time series for the highest-mean-PM day in the colocation (2026-02-18). Top panel: PM_{2.5} vs DEC reference. AB3 and the proposed equation both track reference closely; both under-predict modestly (~3–5 µg/m³) during the early-morning peak (02:00–06:00), where the BAM sees mass from particles the Plantower does not optically detect — a sensor-physics limit, not a calibration issue. Bottom panel: PM₁ derived from each PM_{2.5} estimate via the smooth-ratio function (§4.4). PM₁ tracks PM_{2.5} at ratio ~0.85 across the high-concentration regime and ~0.7 in the moderate regime, monotone with no breakpoint artifacts. There is no DEC PM₁ reference to compare against; the PM₁ panel illustrates that the derived series is plausibly-shaped and free of the AB3 stage-1 noise-floor zeros that the previous firmware produced.

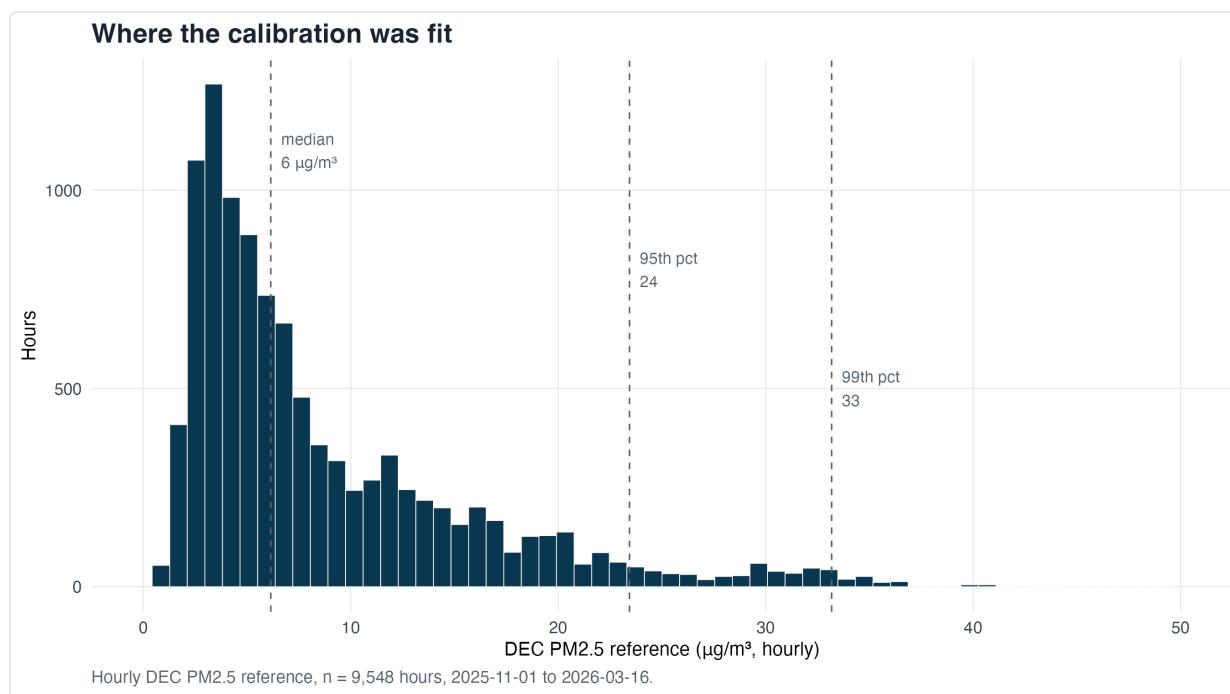
5.4 Residual structure



Residual diagnostics

Figure 5. Residual (predicted – reference) plotted against reference concentration. Both AB3 and the new equation hug zero through the 0–25 $\mu\text{g}/\text{m}^3$ range; the new equation has tighter spread. Above 25 $\mu\text{g}/\text{m}^3$, AB3 residuals widen but stay centered, while the new equation drifts about 1 $\mu\text{g}/\text{m}^3$ negative. Without high-concentration colocation data spanning the 30–100 $\mu\text{g}/\text{m}^3$ regime, the upper-tail behavior remains under-tested.

5.5 Where the calibration was fit



Concentration distribution

Figure 6. Distribution of DEC reference PM_{2.5} during the colocation window. 76% of hours are below 12 µg/m³, 95% below 24, 99% below 33. The recommended equation is well-supported across this range. Above ~30 µg/m³, the equation is extrapolation; HabitatMap should treat ABM PM_{2.5} readings above 30 µg/m³ with appropriate caveats until verified at higher-concentration colocation. Corresponding derived PM₁ thresholds: median 3.32, p95 19.48, p99 28.34 µg/m³.

6. Decision summary for HabitatMap firmware integration

Where the calibration was developed

The recommended equations are fit to ambient measurements at a single urban temperate site over the 2025–2026 cold season:

- **Site:** NYSDEC Queens College, NY (urban temperate)
- **Period:** November 2025 – March 2026 (4.5 months)
- **Reference instrument:** BAM-1020 β-attenuation, 1-minute PM_{2.5}
- **Concentration range:** 0–33 µg/m³ for 99% of hours
- **Aerosol mix:** urban winter — secondary sulfate, traffic, residential heating
- **Sensors:** 5 ABM units (3 normal-batch, 2 new-batch)

The headline performance numbers ($R^2 = 0.914$ for the recommended Option A, up to 0.925 for the free-intercept variant; RMSE ≈ 1.9 – 2.1 µg/m³) reflect this regime. Because the calibration uses Plantower count channels — which are the underlying optical

signals from a sensor whose physics is well-characterized in the literature — we expect the equations to transfer reasonably to other urban temperate ambient deployments with similar aerosol composition.

Performance in regimes we did not test is uncharacterized. The biggest gaps are warm-season aerosol (more semi-volatile organics and photochemistry), high-concentration events (above $\sim 33 \mu\text{g}/\text{m}^3$, including wildfire smoke and dust storms), and coastal/tropical/agricultural sites with substantially different aerosol mixes. If you deploy an ABM in one of these regimes and have access to a reference instrument, sharing your collocation data with HabitatMap would directly inform the next calibration revision and expand the validated envelope over time.

Decisions

Item	Recommendation	Action
PM _{2.5} stage equation	Option A (C0.3 + C2.5, floor 1.23, R ² = 0.914) is the primary recommendation. The free-intercept three-channel variant (R ² = 0.925, floor 2.42) is documented as an alternative if a higher floor is acceptable. Option C (single-channel linear, R ² ≈ 0.90) is acceptable if firmware footprint matters.	Replace AB3 stage-2 block
Count-channel wiring	The colocation exports' count labels are shifted one size cut (§2.1). Bind 0.005157 to the true ≥0.3 μm count (Data7) and 0.211782 to the true ≥2.5 μm count (Data10) — not to the channels as labeled in the colocation CSVs. Miswiring 0.211782 to the true C1 inflates output by roughly +10 μg/m ³ .	Verify channel mapping in firmware code review + one bench sanity check
PM ₁ stage	Derive PM ₁ from calibrated PM _{2.5} via the smooth-ratio function $PM_{2.5} \cdot (0.855 - 0.818 \cdot \exp(-PM_{2.5}/6.12))$. The AB3-inverse piecewise is available as a conservative fallback.	Replace raw-A1 → PM ₁ block
Low-cohort firmware	Out of scope per HabitatMap	None
AirCasting CSV timestamps	Confirmed in NY local time; document for downstream tooling that ingests them as character strings	Note in firmware/docs
v2 verification path	Opportunistic ambient colocation at a higher-concentration regime	Plan post-freeze