HELIOSOLVE

PV Simulation & Yield Analysis Platform

Technical Validation Report

Heliosolve PV Simulation Engine

Post-Release Replication of the Publicly-Released 2021 PVPMC Modeling Intercomparison Dataset

AUTHOR **David Martinez** — Founder, Heliosolve

DATE March 6, 2026

DOCUMENT Technical Validation Whitepaper — Public Release

CONTACT contact@heliosolve.com

© 2026 Heliosolve. All rights reserved. This document may be shared with attribution for non-commercial purposes.

EXECUTIVE SUMMARY

EXECUTIVE SUMMARY AT A GLANCE

HelioSolve's solar-simulation engine reproduced the publicly-released Sandia PV-modelling measurement dataset to within $\pm 2\%$ **on annual DC-power output**, across **two different solar-panel technologies** at the same Albuquerque, New Mexico research site. That is the simulation-side share of the modelling uncertainty contributing to the industry's documented $\sim$**8% energy-yield-forecast gap** [3, 5, 6].

- **What was tested**: HelioSolve replicated two scenarios from the publicly-released 2021 Sandia PV Performance Modeling Collaborative dataset [1, 2] — a Panasonic HIT panel and a Canadian Solar mono-crystalline panel, both fixed-tilt, both at the Sandia research site in Albuquerque.
- **What we found**: Modelled DC power matched the measured reference to -1.93% (Panasonic HIT) and $+0.52\%$ (Canadian Solar) on the year, with measurement-vs-model correlation of $R^2 = 0.9843$ **and** 0.9862. Both biases sit inside the 2021 PVPMC blind intercomparison cohort's median band (-3.3% annual DC energy across 29 submissions) [1].
- **Why it matters**: These results address the *modelling* share of the industry's $\sim$8% yield-forecast gap. Operational, financial and degradation contributors lie outside this report. Tighter modelling on the DC side supports lender P50/P90 inputs as one factor among several [7].

Scope: two scenarios at one research site, one year of measured data. Self-conducted post-release replication (HelioSolve did not submit to the 2021 blind round); not a third-party audit.

Table 1: Annual DC-power validation metrics for the two Albuquerque scenarios. See the methodology appendix for term definitions.

Validation Parameter	Scenario S1 (HIT)	Scenario S2 (mono-c-Si)
Annual DC-power bias	-1.93%	$+0.52\%$
Measurement-vs-model correlation (R^2)	**0.9843**	**0.9862**
Average hourly relative error (nRMSE)	**6.12%**	**5.60%**
Average hourly absolute error (MAPE)	**6.68%**	**5.88%**

Bias = annual integrated difference between modelled and measured DC power, expressed as a percentage of the measured total. AC-side losses (inverter, wiring, availability) and project-level effects (soiling, degradation) are not part of this validation.

Contents

WHY THIS MATTERS

Modern utility-scale solar plants produce roughly **8% less energy** on a weather-adjusted basis than their pre-construction forecasts had predicted [3, 5]. Lawrence Berkeley National Laboratory has separately documented plant-level capacity factors and degradation patterns across more than 31 GW of installed capacity [6]. The gap has multiple drivers — soiling, availability, degradation, EPC quality, weather inputs, and modelling. This report addresses the **modelling** share:

- **The 8% shortfall is real and measured**: In its 2025 update, kWh Analytics reports an 8.6% mean shortfall across 34,000 system-months of operating data [3, 5].
- **It feeds straight into project finance**: Higher modelling uncertainty widens the gap between best-case and worst-case revenue forecasts, which can depress lender debt-service-coverage margins [7].
- **Simulator-vs-measurement gaps are documented in the literature**: A 9-plant close test by NREL of one reference simulator reports annual-energy prediction errors below 3% for fixed-tilt installations and below 8% for tracker installations [9]; the 2021 PVPMC blind intercomparison cohort itself reported a median annual DC-energy bias of -3.3% across 29 submissions [1].

This report focuses on the simulation-model piece, and on a single question: can a simulation engine reproduce what reference instruments actually measured at a research site under real outdoor weather, over a full year? The answer, on these two scenarios, is yes — to within $\pm 2\%$.

WHAT WE VALIDATED

HelioSolve is a physics-based simulation engine that estimates how much sunlight reaches a solar panel and, from that, how much DC electrical power the panel produces under real-world weather and operating conditions. The engine's internals are conventional in their components — standard models for sunlight on the panel surface, panel temperature, and electrical output — and the validation in Section 3 is the empirical evidence that, taken end-to-end, the chain reproduces the measurement reference within the reported error bands.

The validation reference is the **publicly-released 2021 Sandia PV Performance Modeling Collaborative (PVPMC) dataset** — the same dataset that 29 blind-cohort participants from 28 institutions in 12 countries used to test their simulators [1]. The full measurement record was made openly available in April 2023 by Sandia National Laboratories [2]. HelioSolve did *not* submit to the 2021 blind round; the results presented here are a self-conducted post-release replication.

Why this is the right benchmark

- **Fixed, public ground truth**: HelioSolve's local input data was checked against the public Sandia release row by row. The result: **zero mismatches across 26,280 measurement cells per scenario** (front-side sunlight, panel temperature, DC power) to two decimal places. This means simulation tuning to measurement after the fact would be detectable against the cohort's published error band.
- **Two technologies, same site**: The replicated cases cover two distinct solar-panel technologies

(Panasonic HIT and Canadian Solar mono-crystalline) on the same fixed-tilt rack at the Sandia Albuquerque research site — a clean test of whether the simulation chain works across different cell types under identical weather inputs.

- **Peer-reviewed methodology**: All metrics follow the methodology published by the PVPMC in Theristis et al. (2023) [1].

> **The Investor View:** The 2021 PVPMC blind cohort, using the same inputs HelioSolve uses here, came in at a median annual DC-energy bias of -3.3% [1]. HelioSolve's two scenarios came in at -1.93% and $+0.52\%$. That is one piece of evidence that simulation-side modelling uncertainty can be made tighter than the cohort baseline on this dataset; it is not, by itself, a guarantee that the same precision generalises to new sites or new years.

THE RESULT

Over a year of measured data on both Albuquerque scenarios, HelioSolve's modelled DC power matched the measured reference to -1.93% (Panasonic HIT) and $+0.52\%$ (Canadian Solar mono-crystalline) — both well inside the 2021 PVPMC blind intercomparison cohort's own published error band (-3.3% median annual DC-energy bias) [1].

Figure 1 plots modelled vs. measured DC power for each filtered hour of the year, with point density shown by colour. The per-channel breakdown for the Panasonic HIT scenario (Table 2) confirms that the simulation chain is individually well-calibrated at each step — sunlight on the panel, panel temperature, and electrical output — rather than relying on errors cancelling at the aggregate level.

What is, and isn't, in scope: This validation covers the **DC side** of the simulation chain — sunlight on the panel, panel temperature, and the DC electrical power the panel produces before the inverter. AC-side losses (inverter, wiring, transformer), soiling, equipment availability, and grid curtailment are *not* part of this validation. Translating the DC-power agreement into a project-level revenue forecast therefore requires the user's own assumptions on those operational and AC-side losses.

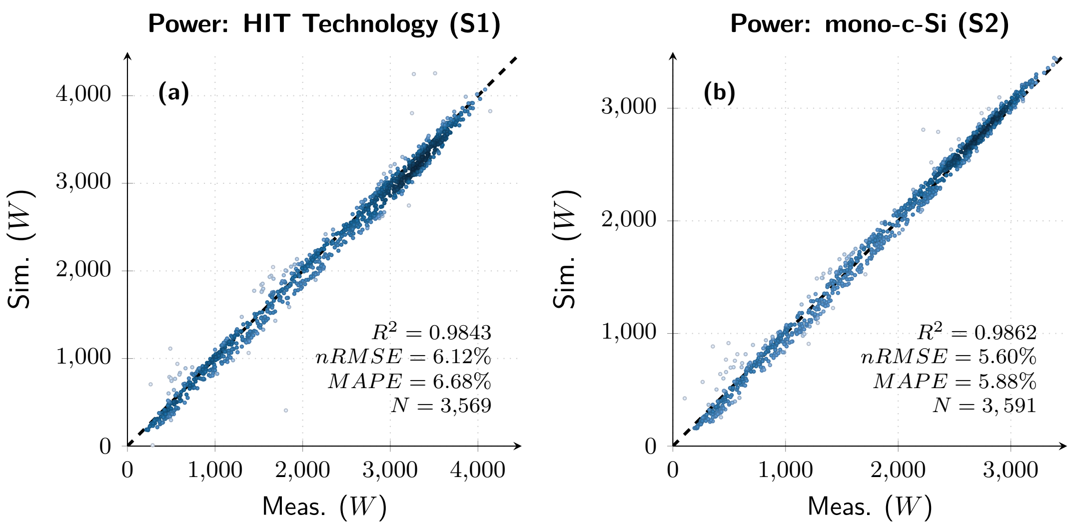

Figure 1: Modelled vs. measured DC power for both Albuquerque scenarios: (a) Panasonic HIT and (b) Canadian Solar mono-crystalline. Each dot is one hour of operation; colour shows how many hours fall at that point. The dashed black line is perfect agreement. Closer to the line is better.

For the Panasonic HIT scenario, the simulation chain is also broken out by step (Table 2): how well the model reproduced the sunlight on the panel, the panel's temperature, and finally the DC electrical power.

Table 2: Step-by-step validation of the Panasonic HIT simulation chain. Each row is one channel of the model. Bias is the annual integrated difference (closer to zero is better); correlation R^2 is the agreement vs. a perfect 1:1 line.

Step in the chain	Hours	R^2	Annual bias	Hourly rel. error
Sunlight on the panel surface	3,569	0.9885	-0.20%	5.18%
Panel temperature	3,548	0.9597	$+0.97\%$	8.59%
DC electrical power	3,569	0.9843	-1.93%	6.12%

Computed on the hours that pass the standard validation filters published with the dataset [2].

The chain holds together at each step: sunlight on the panel reproduced to -0.20%, panel temperature to $+0.97\%$, and DC power to -1.93%. The simulation is therefore well-calibrated at the level of each physical sub-model, not just at the aggregate DC output.

WHAT THIS CHANGES FOR PROJECT FINANCE

Tighter simulation on the DC side addresses one element of the forecast-uncertainty budget that lenders use to size debt and equity returns. Pacudan (2016) documents the link between resource-estimation uncertainty and project-finance metrics — the narrower the modelled distribution of plant output, the better the debt-service-coverage ratio and net present value of the asset [7]. The operational gap between forecast and realized generation reported by kWh Analytics [3, 5] and Bolinger et al. [6] captures the full mix of modelling, soiling, availability, degradation and weather drivers; this report addresses only the modelling component, on two scenarios at one site.

What this validation is, and isn't: It is evidence that the simulation side of a yield forecast can be made tighter than the 2021 PVPMC cohort baseline on a public reference dataset. It is *not* a closed business case for any specific asset. We do not, in this document, quantify the project-level NPV or DSCR impact of any individual project; that requires a project-by-project analysis using the client's own capital-stack assumptions.

SCOPE, LIMITATIONS, AND THE NEXT STEP

This validation covers **two scenarios at one research site** (the Sandia Albuquerque site) over a **twelve-month window** drawn from the publicly-released 2021 PVPMC dataset [1, 2]. Before the engine can be considered generally validated for production-grade forecasting, the same exercise needs to be repeated on additional sites, additional years, and additional cell technologies. That is the next phase of work.

Request a pilot / technical review

Asset owners, lenders, and developers are invited to benchmark their own simulation assumptions against a HelioSolve simulation under matched inputs.

Request a pilot / technical review

Contact the HelioSolve team to schedule an independent audit of your project's yield projections.

contact@heliosolve.com

APPENDIX A — METHODOLOGY AND DEFINITIONS

Glossary

- **Sunlight on the panel surface** (in the technical literature, "plane-of-array irradiance" or POA): the actual sunlight reaching the panel's working face after accounting for the panel's tilt, the sun's position, and any shading.
- **Panel temperature**: the operating temperature of the cells inside a panel under sunlight, wind and ambient air. Higher temperature reduces electrical efficiency, so getting it right matters for the DC-power output.
- **DC power**: the electrical power produced by the panels before it reaches the inverter (which converts DC to grid-quality AC).
- **Fixed-tilt monofacial array**: solar panels mounted at a fixed angle (no daily rotation), with electrical output coming only from the front of the panel. This is the simpler of the two main utility-scale array geometries; the other, more complex case is bifacial single-axis tracking.
- **Annual bias**: the difference between the simulated value and the measured reference, integrated over the validation window, expressed as a percentage of the measured total. A bias of -1.93% means the simulation came in 1.93% *below* the measured total over the year.
- **Correlation (R^2)**: how well the simulated and measured values agree on a 1:1 line. $R^2 = 1$ would be perfect; $R^2 = 0.98$ means the simulation explains 98% of the hour-by-hour variation in the measured signal.
- **Hourly relative error (nRMSE) / hourly absolute error (MAPE)**: average size of the hour-by-hour difference between simulation and measurement, with two different ways of expressing the average. Both are reported in percent of the measurement.
- **P50 / P90 / P99**: probabilistic forecast percentiles. P50 is the median (best-estimate) annual production; P90 is the level of production that is exceeded with 90% probability (used by lenders as a worst-case sizing metric).
- **DSCR**: Debt-Service Coverage Ratio — the ratio of a project's operating cash flow to its debt service. A higher DSCR means more headroom against under-performance.
- **NPV**: Net Present Value — the project's expected cash flows discounted to today.

What was filtered out

The metrics in this report are computed on the hours that pass the standard validation filters published with the Sandia dataset: front-side sunlight $\geq 100\,\mathrm{W/m^2}$, DC power $\geq 50\,\mathrm{W}$, solar elevation above $5°$, and ambient temperature between $-5°$C and $+45°$C. Night-time hours and sensor-outage rows (encoded as zero values) are also excluded. After filtering, $\approx 3{,}569$ hours remain for the Panasonic HIT scenario and $\approx 3{,}591$ for the Canadian Solar scenario, out of the year's $8{,}760$ hours.

What is, and isn't, in scope

This report covers the DC-side simulation chain only — sunlight on the panel, panel temperature, and DC electrical output. AC-side losses (inverter conversion, clipping, MV/LV transformer losses), soiling, availability and grid curtailment are not part of this validation. Translating the DC-power agreement into a project-level revenue forecast requires the user's own assumptions on those operational and AC-side losses.

REFERENCES

[1] Theristis, M., Riedel-Lyngskær, N., Stein, J. S., et al. (2023). Blind photovoltaic modeling intercomparison: A multidimensional data analysis and lessons learned. *Progress in Photovoltaics: Research and Applications*, 31(11), 1144–1157. DOI: 10.1002/pip.3729. Available: https://doi.org/10.1002/pip.3729

[2] Theristis, M., and Stein, J. S. (2023). *The 2021 Blind PVPMC Modeling Intercomparison*. Sandia National Laboratories Technical Report SAND2023-03135. OSTI: https://www.osti.gov/biblio/2293576 (direct PDF: https://www.osti.gov/servlets/purl/2293576/). Accompanying open-access dataset on the DuraMAT Data Hub (DOI: 10.21948/1970772): https://datahub.duramat.org/dataset/pv-performance-modeling-data.

[3] kWh Analytics (December 2022). *2022 Solar Generation Index: Evaluating Actual vs. Expected Production Across the US Fleet*. Industry Research Report (registration required at kwhanalytics.com). Press coverage: pv-magazine-usa.com (Dec. 6, 2022).

[4] Flyvbjerg, B., Garbuio, M., and Lovallo, D. (2009). Delusion and Deception in Large Infrastructure Projects: Two Models for Explaining and Preventing Executive Disaster. *California Management Review*, 51(2), 170–193. DOI: 10.2307/41166485. Available: https://doi.org/10.2307/41166485

[5] kWh Analytics (June 2025). *2025 Solar Risk Assessment: Market Insight & Empirical Performance Data for Project Finance*. Industry Research Report (seventh edition; registration required at kwhanalytics.com). Press coverage: solaranywhere.com (June 2025).

[6] Bolinger, M., Gorman, W., Millstein, D., and Jordan, D. C. (2022). Plant-level performance and degradation of 31 GW-DC of utility-scale PV in the United States. Lawrence Berkeley National Laboratory Report. Available: https://emp.lbl.gov/publications/plant-level-performance-and

[7] Pacudan, R. (2016). Implications of applying solar industry best practice resource estimation on project financing. *Energy Policy*, 95, 489–497. DOI: 10.1016/j.enpol.2016.02.021. Available: https://doi.org/10.1016/j.enpol.2016.02.021

[8] Elnosh, A., Calais, M., and Parlevliet, D. (2026). A systematic literature review of digital twin research for photovoltaic systems: Trends, challenges, and opportunities. *Renewable and Sustainable Energy Reviews*, 226, 116326. DOI: 10.1016/j.rser.2025.116326. Available: https://doi.org/10.1016/j.rser.2025.116326

[9] Freeman, J., Whitmore, J., Blair, N., and Dobos, A. P. (2014). Validation of Multiple Tools for Flat-Plate Photovoltaic Modeling Against Measured Data. *NREL Technical Report NREL/TP-6A20-61497*. Available: https://www.nrel.gov/docs/fy14osti/61497.pdf

HELIOSOLVE — Precision Yield Simulation for the Solar Industry

David Martinez, Founder · contact@heliosolve.com

© 2026 Heliosolve. All rights reserved. This document may be shared with attribution.