

Response to RC2

This paper presents a detailed evaluation of wind and solar energy generation estimates derived from E3SM SCREAM regionally refined simulations over California, with comparisons against HRRR, E3SM-25km NARRM, and EIA-reported capacity factors. The topic is timely and relevant, particularly given the growing importance of high-resolution meteorological modeling for renewable energy assessment. The modeling framework and the scope of the analysis are potentially valuable. However, in its current form, the manuscript requires substantial revision before it can be considered for publication. The primary issues concern clarity, precision, and organization of the writing, as well as conceptual ambiguities in the framing of the research questions and interpretation of results. Detailed comments are provided below.

Thank you for your careful reading of the manuscript and for your supportive comments! In response to your questions and suggestions, we provide detailed replies to each point below. Following your suggestions regarding wording, structure, and presentation, we have revised the manuscript accordingly, particularly in the Introduction and Results sections. In addition, because the EIA data for 2024 have recently become available, we updated all figures related to EIA and observations (reanalysis), except for Figs. 1, 7, and 14. The additional year of observations increases the sample size for EIA plant-level analyses, although the climatological comparison results remain almost unchanged.

Major comments

1. Several abbreviations are introduced without clear or consistent definition, particularly in the abstract.

(1) The abbreviation RRM is used inconsistently. The manuscript refers to “regional mesh refinement” in the context of CARRM and “regionally refined model” in the context of NARRM. A consistent naming convention should be adopted and clearly introduced at first use.

(2) SCREAM is introduced in a convoluted manner in the abstract (“derived from the US Department of Energy’s Simple Cloud-Resolving Energy Exascale Earth System Model (E3SM) Atmosphere Model (SCREAM)”), which makes it unnecessarily difficult to identify what SCREAM stands for. The full name should be stated clearly and directly at first mention.

(3) ERA5 and MERRA-2 are repeatedly referenced without explicitly stating that they are reanalysis datasets, nor is it specified which meteorological variables from these datasets are used. In multiple places, the manuscript refers to the reanalysis data as “meteorological models” or “meteorological model inputs,” which is inaccurate and potentially misleading.

Thank you for pointing out these semantic ambiguities!

1) In the literature, RRM can refer both to the Regionally Refined *Mesh* and to a Regionally Refined *Model*. Strictly speaking, RRM is a technical approach for modifying grid-related configurations within a given host model, and it does not exist independently of an underlying model. However, in the literature it is often simplified to refer to a specific model version that applies a regionally refined mesh. We have chosen to consistently retain the use of Regionally Refined Model throughout the manuscript.

2) By directly referencing E3SM, we have revised the description of SCREAM in the abstract to its more commonly used full name “Simple Cloud-Resolving E3SM Atmosphere Model (SCREAM)”.

3) When ERA5 is first introduced, its full name naturally indicates that it is a reanalysis product, whereas the full name of MERRA2 does not explicitly include this information. We have therefore added explicit statements at first mention clarifying that both are reanalysis datasets “...compared to the coarse-resolution reanalyses such as ERA5 and MERRA2”. We do not calculate CFs using ERA5 or MERRA2 in this study but cite their use in the existing literature. In the subsection “PLUSWIND-derived CFs” we have added which variables are used in PLUSWIND: “They use hub-height wind speeds from HRRR, model-level wind speeds from ERA5, and near-surface wind speeds from MERRA2 to estimate wind power generation”.

We apologize for the misleading word choice. The phrase meteorological model inputs appeared once in the original text: “This suggests that the effect of modeling assumptions between PLUSWIND and PySAM is larger than the effect of internal corrections within PLUSWIND, but smaller than the effect of varying meteorological model inputs”. Here, “inputs” refers to meteorological datasets (i.e., outputs from atmospheric models) used as inputs to the generation model (e.g., PySAM). To avoid ambiguity, we have revised this to “varying meteorological model datasets”. In this paper, unless otherwise specified, each meteorological model corresponds one-to-one with its associated dataset; e.g., HRRR, ERA5, and MERRA2 are distinct models, each paired with the meteorological dataset generated by that model.

2. The research questions listed in the introduction are not well separated conceptually and often mix scientific questions with methodological choices.

(1) First goal (generation modeling assumptions): This appears primarily to be a methodological sensitivity analysis, yet it is framed as a main scientific objective. It is not clear, until later sections, why this comparison is essential for the broader goals of the paper. If this is a key component, the motivation and what the “generation modeling assumptions” represent should be stated more clearly in the introduction.

(2) Second goal (meteorological models): The phrase “what do meteorological models represent” is vague. It is unclear whether this refers to differences in simulated meteorological fields or differences in model physics. The authors should be more precise in their wording. In addition, the comparison between SCREAM-3kmCARRM forced by a 2005–2014 climatology and HRRR simulations from 2018–2022 is not a fair comparison. While the manuscript argues that monthly climatology justifies this approach, interannual variability during 2018–2022 may still strongly affect the results, and this issue is not adequately addressed. There is also an inconsistency in how the “2010climo” period is defined: it is described as 2005–2014 in the text, whereas Table 1 lists 2010–2019. This discrepancy should be clarified and made consistent.

Thank you for raising this point! These two questions are indeed related. In the original structure, generation modeling assumptions were listed as the first item, and meteorological models as the second, but both are components of the first core question we aim to address: Is km-scale E3SM (SCREAM) combined with PySAM an effective framework? The second core question is: Can km-scale E3SM simulations substantially reduce the pronounced wind power biases identified in 25 km E3SM simulations over California, and how important is horizontal resolution for wind and solar energy simulations?

For the first question, assessing “effectiveness” requires multiple reference benchmarks. The benchmark closest to applications is the EIA monthly data; however, discrepancies between simulated and EIA CFs arise from multiple sources of uncertainty: (1) biases in SCREAM-CARRMs simulation data, (2) PySAM energy modeling, and (3) EIA data collection and quality control. We therefore include a second benchmark, HRRR, another km-scale meteorological dataset validated in the wind energy literature. By combining HRRR with the same PySAM settings applied to SCREAM-CARRMs, uncertainties related to horizontal resolution and energy modeling are removed, isolating differences arising solely from the meteorological model (SCREAM-CARRMs vs. HRRR). A third benchmark is the PLUSWIND-derived HRRR wind CFs, which, when compared with HRRR + PySAM, removes meteorological uncertainty

(both use HRRR) and isolates differences only due to the generation model (PySAM vs. PLUSWIND).

To the specific questions:

1) The motivation for introducing generation modeling assumptions in the first goal is as follows. We begin by comparing HRRR + PLUSWIND vs. HRRR + PySAM to first assess how large the variability arising from the use of different generation models can be. This is to validate our use of PySAM by comparing against a well-established dataset. While this comparison resembles a sensitivity analysis (as it involves two different wind power models), the focus of our study is the energy evaluation of SCREAM-CARRMs and the effects of grid resolution. Thus, except for axis-tracking, we adopt the default PySAM settings rather than doing a detailed sensitivity analysis.

2) In the second goal, “meteorological models” refers specifically to different meteorological models at comparable horizontal resolution, namely SCREAM-3kmCARRM vs. HRRR. In a climatological context, we treat meteorological models (including physics) and the resulting simulation datasets as closely linked, as differences among simulated fields are primarily driven by model differences. We fully agree with the concern regarding mismatched simulation periods; this is an unavoidable challenge when comparing climatological simulations with reanalysis data. We also recognize that the original question framing was too broad. As discussed above, HRRR serves as one of several benchmarks for evaluating SCREAM-3kmCARRM, analogous to the role of PLUSWIND in the first question. Both questions arise because EIA-based evaluation alone contains multiple sources of uncertainty, motivating additional comparisons to help gauge the relative importance of different uncertainty sources (horizontal resolution, meteorological model, and generation model).

The 2010climo simulation is forced by present-day observational climatology of the 2010s, with SST and sea ice prescribed from the climatological annual cycle averaged over 2005–2014. The nominal simulation period spans 2010–2019. However, for climate-type simulations, nominal years can be misleading, particularly when other datasets correspond to actual calendar years. Therefore, we have revised Table 1 to report simulation durations rather than nominal simulation years for all E3SM simulations.

=====

Following your suggestion, we have reorganized and clarified the questions listed in the introduction:

“By estimating CFs using PySAM and comparing all modeled results with EIA monthly data, we construct multiple angles of comparison. Our high-level study objectives are to access: 1) whether km-scale SCREAM-CARRM climate data combined with PySAM constitutes an effective energy modeling framework, and 2) whether it can substantially reduce the pronounced wind energy biases identified in 25 km E3SM simulations over California (i.e., how important is horizontal resolution is for wind and solar energy simulations?).

EIA provides the benchmark most directly relevant to applications; however, discrepancies relative to EIA reflect combined uncertainties associated with SCREAM-CARRM datasets, PySAM energy modeling, and EIA data quality. We therefore introduce additional evaluation targets to approximately assess the relative importance of these uncertainties. Specifically, we first assess whether PySAM is applied appropriately, then compare against HRRR at comparable resolution to evaluate whether SCREAM-3kmCARRM has similar performance, and finally examine the role of horizontal resolution. Accordingly, we address the following three questions:

\begin{enumerate}

\item \emph{Does the use of different energy models affect wind CFs? How do different axis-tracking methods affect solar CFs?}

This question focuses on validating our use of PySAM in assessing whether the SCREAM-CARRM + PySAM framework is effective. For wind, we compare HRRR-derived CFs from PLUSWIND with those estimated using PySAM, with both evaluated against EIA monthly CFs. We also reproduce the PLUSWIND comparisons among ERA5, MERRA2, and HRRR, which helps assess the relative importance of horizontal resolution and generation modeling. For solar, we evaluate the impacts of fixed, single-axis, and dual-axis tracking assumptions in PySAM.

\item \emph{How different are wind and solar CFs between SCREAM-3kmCARRM and HRRR at comparable horizontal resolution?}

Using the same PySAM configuration, we assess the performance of SCREAM-CARRM by comparison with HRRR. Although these two meteorological models differ substantially in model structure, physical parameterizations, and simulation periods (HRRR forecasts over 2018–2022 vs. SCREAM-3kmCARRM forced by observed climatology over 2005–2014), they share a very similar horizontal resolution at the kilometer scale. We therefore ask whether, despite

these differences, comparable horizontal resolution can dominate and lead to statistically similar wind and solar CFs.

\item \emph{How much does horizontal resolution affect wind and solar energy calculations?}

We first compare CFs computed with PySAM from the E3SM-25kmNARRM and 3.25 km SCREAM-CARRM simulations to assess whether km-scale resolution substantially reduces the wind generation biases identified at 25 km. Sensitivity to further resolution refinement is then assessed by comparing 3.25 km and 800 m SCREAM-CARRM simulations, which share identical configurations (driven by CMIP6 forcings and evaluated for the present day) and differ only in horizontal resolution.

\end{enumerate}”

3. The manuscript states that SCREAM-RRM requires nudging because it lacks a scale-aware deep convection scheme. However, nudging does not directly address the absence of scale-aware convection. The authors should clarify the actual purpose of nudging in these simulations, for example whether it is primarily used to constrain large-scale circulation. The choice of a 2-day relaxation timescale also appears relatively weak and should be justified. In addition, the explanation of how “reduced nudging strength near the surface and model top” allows “free-running conditions over California” while maintaining large-scale constraints elsewhere is not sufficiently clear.

Yes, nudging is used primarily to constrain the large-scale circulation. It does not directly compensate for the absence of a scale-aware deep convection scheme; rather, it helps maintain realistic large-scale thermodynamic and dynamical conditions, which in turn limits the unrealistic buildup of instability in the absence of a convective parameterization and ensures consistent large-scale forcing into the refined domain.

Regarding the nudging details, we apologize for the lack of clarity in the original description. The statement “this allows free-running conditions over California while nudging the coarser-resolution domain” corresponds to “using the E3SM regional configuration with a tanh-based smooth windowing function to constrain the horizontal domain.” Because SCREAM is a global model, a spatial weighting function is used to vary nudging strength across the domain; in E3SM, this is implemented via a tanh-based smooth windowing function. Vertically, nudging strength is also defined via tanh-based smooth window, with users specifying topmost and bottommost model levels and transition gradients (deltas) in the user namelist. We adopted the default

configuration, in which nudging spans all vertical levels with delta values of 0.1 at both the top and bottom, resulting in weaker nudging near the surface and at the top and full-strength nudging elsewhere. The two-day relaxation timescale follows previous SCREAM-3kmCARRM climate simulations (Zhang et al., 2024). In earlier experiments we tested shorter timescales and found that a 2-day relaxation produced the most consistent precipitation patterns between RRM and the nudging data.

To clarify, we have revised the nudging description accordingly:

“Nudging data are interpolated onto the RRM physics grid using NCO’s linear-in-log-pressure vertical interpolation and TempestRemap’s high-order horizontal interpolation. U, V, T, and Q nudging with a 2-day relaxation timescale is applied following \citep{Zhang2024}. Regional nudging in E3SM employs a tanh-based smooth windowing function that is zero within the refined region and transitions toward one in the surrounding coarse-resolution domain \citep[e.g.,]{Tang2019}, allowing free-running conditions over California and applying nudging in the coarser-resolution domain. Vertically, the default tanh-based smooth window is applied, which reduces the effective nudging strength near the surface and at the top and full strength in between.”

4. Section 3.1 emphasizes substantial uncertainties in the EIA monthly data and explicitly states that model–EIA differences should be interpreted as “discrepancies” rather than biases. Given this, the scientific value of using EIA data as the primary reference for evaluating model performance is unclear. The authors should clarify how conclusions about model skill can be drawn when the reference dataset itself is highly uncertain, and whether the PLUSWIND dataset might provide a more reliable benchmark.

Thank you for raising this issue! As discussed in our response to earlier comments, this is precisely why we supplement EIA-based evaluation with additional datasets. These include comparisons between HRRR + PLUSWIND vs. HRRR + PySAM to confirm our correct use of PySAM, as well as comparisons between HRRR + PySAM and SCREAM-3kmCARRM + PySAM.

Despite its limitations, EIA remains the best publicly available dataset of observed generation and one of the most widely used public energy datasets in the U.S. (\url{https://catalystcoop-pudl.readthedocs.io/en/latest/index.html}, last accessed: February 22, 2026). HRRR CFs from PLUSWIND closely resemble our HRRR + PySAM results and have therefore been used as a benchmark for wind CFs, although PLUSWIND does not include solar. Both HRRR + PLUSWIND and HRRR + PySAM are

subject to uncertainties associated with HRRR model biases, particularly for solar CFs. Indeed, HRRR solar simulations deviate more from reference datasets (EIA for CFs and ERA5 for surface downwelling shortwave radiation) than SCREAM-3kmCARRM, suggesting caution in treating HRRR solar results as closer to reality. Unfortunately, more suitable observational datasets are either unavailable or not publicly accessible. Our conclusions therefore emphasize SCREAM's potential applicability in regions with similar climates, based on the available comparisons.

We have expanded the discussion of potential alternative datasets in the section "Uncertainty in EIA monthly data and interpretation of discrepancies" and added the following interpretation:

"Given these limitations, the use of monthly EIA data in this study reflects a trade-off between observational fidelity and data availability. We therefore caution against attributing discrepancies between modeled generation and EIA monthly CFs solely to deficiencies in either the atmospheric simulations or the power generation modeling. Instead, these discrepancies should be interpreted in light of the combined uncertainties inherent in both the observational data and the meteorological-generation modeling framework. This is why we also include comparisons with PLUSWIND and HRRR to help assess the relative importance of different uncertainty sources."

5. Section 3.2 ("Evaluation group design and key questions") is vague and reads more like informal notes than formal academic writing. This section should be rewritten to clearly and concisely describe the evaluation framework.

Thank you for this suggestion! We have rewritten this section as follows:

"\subsection{Evaluation design and structure}

The Results section is organized to address the three guiding questions introduced in the Introduction through a set of controlled comparisons. All CFs are calculated using a consistent PySAM configuration unless otherwise noted. Because discrepancies relative to EIA reflect combined uncertainties from meteorological datasets, power-generation modeling, and EIA data quality, additional benchmarks are introduced to help contextualize the relative importance of these uncertainty sources.

For both wind and solar energy, the analysis follows a common structure in general:

\begin{enumerate}

\item \emph{Generation model dependence (Q1).}

For wind, we compare CFs derived from HRRR using PLUSWIND and PySAM to assess whether generation model choices materially affect CF estimates. For solar PV, we quantify the sensitivity of modeled CFs to axis-tracking assumptions within PySAM. These tests serve as a validation of PySAM before inter-comparing meteorological datasets.

\item \emph{Meteorological dataset dependence at comparable resolution (Q2).}

We compare SCREAM-3kmCARRM 2010climo against HRRR, both of which provide km-scale meteorological fields to PySAM. This comparison quantifies the magnitude of CF differences attributable to the meteorological model when horizontal resolution is not the primary differentiator.

\item \emph{Horizontal resolution sensitivity (Q3).}

We compare present-day SCREAM-CARRM simulations at 3.25 km with the E3SM-25kmNARRM simulation to assess whether km-scale configurations substantially reduce the pronounced wind CF phase errors identified at 25 km resolution over California. Sensitivity to further refinement of resolution is then examined by comparing the 3.25 km and 800 m CARRM simulations.

\end{enumerate}

Within each comparison group, results are presented separately for wind and solar. For each energy type, we first show multi-year monthly CF climatology aggregated by plant clusters, followed by diagnostics that help interpret CF differences, including diurnal cycles of simulated CFs, statewide spatial patterns of CFs and key meteorological driver variables used by PySAM. Each energy-type section concludes with a summary plot to synthesize the main findings across meteorological datasets.”

6. The statement that “each following subsection presents results by energy type (wind or solar)” is misleading, as the first evaluation (generation modeling assumptions) applies only to wind. A consistent subsection addressing generation modeling assumptions for solar is missing.

Following your earlier comments, we have clarified this issue by explicitly distinguishing between wind and solar. For wind, the question is whether the use of different energy models affect wind CFs; for solar, the question is how do different axis-tracking methods affect biases in modeled solar CFs. We have placed the axis-tracking results in a dedicated subsection.

The content of this subsection has been updated accordingly:

“\subsection{How do different axis-tracking methods affect solar CFs?}

Figure~\ref{SolarMon3kmCARRMF2010vsHRRR} shows the solar CFs in California as estimated by PySAM, using SCREAM-3kmCARRM 2010climo and HRRR simulations. Compared to wind, the seasonal cycle of solar CFs is much simpler, with a clear summer peak and winter minimum. Even for dual-axis systems, despite having only three samples, the seasonal pattern remains smooth. This likely reflects the relatively stable solar resource in California’s sunny climate; in cloudier regions with higher solar variability, such smoothness may not hold.

Across the three tracker classes, the overestimation of simulated solar CFs relative to EIA records increases with tracker complexity. For fixed-axis arrays, both SCREAM-3kmCARRM 2010climo and HRRR reproduce the observations well: on average, 3kmCARRM deviates by less than 3 pp in every month. For single-axis systems, the overestimation in both models remains below 6 pp, while for dual-axis systems it stays below 13 pp.

This monotonic increase suggests that the tracker algorithms implemented in PySAM are somewhat idealized. Solar tracking systems are designed to increase energy production by maintaining the panel orientation close to perpendicular to the incoming solar radiation. In principle, dual-axis trackers should therefore achieve higher efficiency than single-axis systems. However, this pattern is not evident in the EIA generation records. One possible reason is that dual-axis systems are mechanically more complex, and operational constraints associated with the control and driving mechanisms may offset part of the theoretical energy gain. In addition, the benefit of multi-axis tracking may be limited in hot climates, where the increased incident irradiance can raise module temperatures and reduce PV output due to thermal losses \citep{Hammoumi2022}.

In the SAM PV model, tracker orientation follows geometric tracking algorithms that determine panel tilt and azimuth from solar position and array geometry:

for single-axis systems, the panel azimuth tracks the solar azimuth while the tilt remains fixed relative to the rotation axis, whereas for dual-axis systems both tilt and azimuth follow the solar zenith and azimuth angles, respectively \citep{SAMPV2018}. In practice, mechanical constraints, tracking control strategies, and stow conditions can lead to slower or less precise tracking, reducing the realized energy gains relative to these idealized simulations.”

In addition, we revised the grouping logic for tracking types in the Methods section as follows:

“For PV plants, we group samples by both model grid point and axis-tracking type. For example, if a grid point G contains three single-axis plants and one fixed-axis plant, two groups are formed: G_1 (single-axis) and G_0 (fixed). The seasonal cycles and evaluation metrics are then computed by clustering according to axis-tracking type. This approach preserves the plant-level tracking type information reported in the EIA dataset, and the total nameplate capacity within each cluster does not vary with model resolution. As will be shown in the results section, the axis-tracking type has a substantial influence on PV power generation.”

7. Sections 3.3.2–3.3.6 are unevenly structured. Some subsections emphasize resolution effects even though they are nominally focused on meteorological model differences, while others contain only a single paragraph. The organization of these subsections should be improved to better align with the stated evaluation goals.

Section 3.3.2 compares HRRR + PySAM with SCREAM-3kmCARRM 2010climo + PySAM. Although the two simulations have similar horizontal resolutions, the models themselves differ in many aspects. As noted above, the inference regarding which source of uncertainty is more influential emerges only when this comparison is considered together with the subsequent analysis of horizontal resolution sensitivity and the contrasting results between wind and solar CFs. To clarify this point, we moved the discussion related to horizontal resolution from this section to the following subsection.

In addition, we added a subsection discussing the diurnal cycle of solar generation and moved the corresponding figure forward to immediately follow the EIA monthly evaluation, because the diurnal-cycle analysis is also based on EIA plant sampling (although no hourly EIA data are available here).

As explained in our responses to the previous two comments, we have revised the descriptions introducing each subsection accordingly. Because these revisions involve extensive changes throughout the results section, please refer to the tracked-changes version for the detailed modifications.

8. In Figure 4 and subsequent figures, the EIA curves appear different between the top and bottom panels. The reason for this difference is unclear and should be explained.

Thanks for asking this! This choice was made because our primary goal is model evaluation at the model-resolution (i.e., grid-cell) scale, rather than plant-level applications. This approach has several advantages:

1) It avoids artificially reducing the spread of the modeled climatological seasonal cycle. In plant-level evaluations, multiple plants located within the same model grid cell would be represented identically. This would substantially reduce the apparent monthly climatological spread and make modeled variability appear artificially smaller than that from observations.

2) Aggregating to the model grid ensures consistency between local evaluations and assessments of spatial patterns across California, as both are conducted on the same grid.

In addition, we identified a minor issue in the previous version of the code used for the seasonal cycle plots. Small differences in the EIA curves after aggregation to different model grids arose from an incorrect application of nameplate capacity weights during gridcell averaging. This has been corrected in the revised manuscript, and the resulting changes are minor and do not affect the conclusions.

We have added this explanation to the subsection "EIA monthly CF data":

"We adopted a gridcell-based evaluation framework instead of plant-level evaluation for the following reasons: 1) It avoids artificially reducing the spread of the modeled climatological seasonal cycle. In plant-level evaluations, multiple plants located within the same model grid cell would be represented identically. This would substantially reduce the apparent monthly climatological spread and make modeled variability appear artificially smaller than that from observations. 2) Aggregating to the model grid ensures consistency between local evaluations and assessments of spatial patterns across California, as both are conducted on the same grid. In the seasonal-cycle analysis, we applied a weighted average

based on total nameplate capacity within each grid cell, so that the EIA mean line remains consistent across comparisons at different model resolutions.”

9. In some regions (e.g., Kern County and Southern California), the seasonal cycles between datasets do not appear as similar as claimed in the text, for example the statement that “both models captured the July peak well.” These interpretations should be revisited and aligned more closely with what is shown in the figures.

We apologize that the discussion of SoCal and NorCal was inadvertently reversed in several places! The statement in question actually refers to NorCal. The relevant text has been corrected.

10. Lines 475–485 discuss deficiencies in the simulation of large-scale circulation, but then abruptly conclude that large-scale circulation plays a relatively minor role. The logical connection between these statements is unclear and requires clarification.

The initial discussion examines large-scale circulation as a commonly used diagnostic to understand the high winter wind speeds in the 25 km simulation. However, the ridge pattern in SCREAM-3kmCARRM 2010climo closely resembles that in 25kmNARRM. If biases in large-scale circulation were the dominant cause of the winter wind bias, the 3 km simulation would be expected to exhibit similarly strong biases, which it does not. This indicates that large-scale circulation is unlikely to be the primary driver of the high winter wind bias in 25kmNARRM. This may not be surprising for California, where complex topography implies a stronger influence of local, small-scale processes on wind patterns. Nevertheless, examining circulation differences remains a useful diagnostic to support this inference.

11. Section 3.4.1 aims to assess the impact of meteorological models on solar energy, but the results primarily describe model behavior without clearly answering this question.

Thank you for this reminder! We have refined the question to “How different are solar CFs between SCREAM-3kmCARRM and HRRR at comparable horizontal resolution?” Given the many differences between the two models, our evaluation quantifies the magnitude of differences in simulated solar CFs but cannot isolate specific causal mechanisms. Nevertheless, we examined key meteorological drivers of solar CFs

(surface downwelling shortwave radiation) in ERA5 over the same period as HRRR (2018–2022) and found spatial patterns very similar to the 2013–2024 mean (Fig. R2). This suggests that HRRR solar CFs averaged over 2018–2022 would likely be comparable to those averaged over the EIA record period, implying that the observed differences primarily reflect meteorological model differences. Clouds, which dominate surface shortwave radiation, remain one of the largest sources of uncertainty in atmospheric modeling as they depend on multiple interacting physical processes.

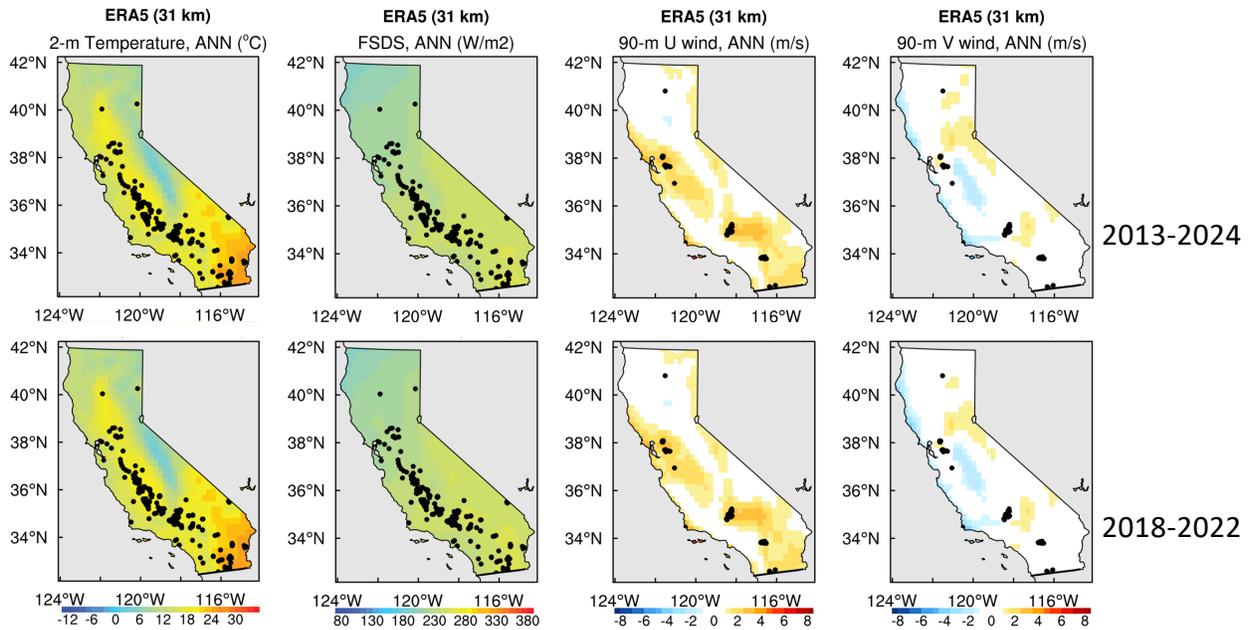


Fig. R2. Direct meteorological drivers of annual-mean solar and wind CFs. From left to right: total downwelling solar flux at surface (FSDS), 2 m temperature, 90 m zonal winds (U90M) and 90 m meridional winds (V90M).

We also noted that this section previously conflated axis-tracking sensitivity with meteorological model differences. Axis-tracking has now been moved to a separate subsection, and the above discussion has been added here.

12. Section 3.4.3 is titled “large-scale circulation,” yet large-scale circulation is not discussed in the current version of this section. In addition, the conclusions discuss generation modeling assumptions for both wind and solar energy, whereas the introduction and the corresponding research question frame this issue almost exclusively in terms of wind (“How much do generation modeling assumptions impact the wind energy?”). This inconsistency should be addressed.

Thank you for pointing out these issues! We have deleted large-scale circulation in the section title. And the introduction now clearly distinguishes generation modeling questions for wind and solar.

Minor comments

1. Introduce full names before abbreviations for SHOC, P3, RTE+RRTMGP, SPA, and MPAS.

Except for MPAS, we previously placed the full names of these terms in parentheses (followed by the citation, separated by “;”), because the abbreviations are often used more widely than their full names in practice. In this revision, we have reversed the order and added the full name of MPAS.

2. Section title 2.2.2: “Souce” or “Source”?

Corrected.

3. Line 176: replace “simulations” with “this type of experiment simulates”.

Corrected.

4. Line 185: provide the full name for MPAS.

Added: the Model for Prediction Across Scales (MPAS).

5. Line 221: “eia-processor” → “EIA-processor”.

Modified.

6. Lines 225–230: rewrite in a more formal, non-conversational style.

Thanks for this suggestion! We revised this paragraph to use a more formal and descriptive style to describe the filtering criteria:

“The EIA data first underwent the following plant-level quality control steps:

\begin{enumerate}

\item Only plants with a nameplate capacity of at least 2.5 MW were included.

\item Only plants located in the state of California were included.

\item For solar generation, only plants classified as “Solar Photovoltaic” were retained (accounting for 97.3% of all solar facilities), while those categorized as “Solar Thermal without Energy Storage” were excluded. In this study, the term *solar* refers specifically to *PV (photovoltaic)* power.

\item For solar PV plants, only those with a reported array configuration in the EIA dataset were retained, as indicated by at least one of the fields “Single-Axis Tracking?”, “Dual-Axis Tracking?”, or “Fixed Tilt?”.

\item Plants were excluded if CF was below 5% for more than nine months or if more than 12 months of data were missing.

\end{enumerate}”

7. Line 245: clarify what “independent-plant ratio > 15%” means and how it is computed.

We have added one sentence here:

“The independent-plant ratio is defined as the fraction of plants in a given year whose monthly CF time series are not perfectly correlated with those of other plants (identified using a Pearson correlation distance threshold of 10^{-6}).”

8. Line 262: clarify what “models” refers to.

Added:

"Air density correction for all models (ERA5, MERRA2, HRRR) was based on MERRA2 air density, interpolated to hourly resolution at hub height."

9. Line 274: "preprocessing" is duplicated.

The second "preprocessing" has been deleted.

10. Line 428: "suggests" is repeated.

The second "suggests" has been deleted.

11. Figure 7 caption needs revision, as large-scale circulation is one of the drivers influencing wind CF, rather than a separate category from the direct drivers.

We have deleted "large-scale circulation" from the summary sentence in Figure 7 caption.