



Subpixel variability and quality assessment of satellite sea surface temperature data using a novel High Resolution Multistage Spectral Interpolation (HRMSI) technique

Sandra L. Castro^{a,*}, Lucas A. Monzon^b, Gary A. Wick^c, Ryan D. Lewis^{b,1}, Greg Beylkin^b

^a Colorado Center for Astrodynamics Research, University of Colorado, Boulder, CO, USA

^b The Numericus Group, LLC, Boulder, CO, USA

^c NOAA Earth System Research Laboratory, Physical Sciences Division, Boulder, CO, USA

ARTICLE INFO

Keywords:

SST subpixel variability

Spatial variability

Interpolation techniques

ABSTRACT

A novel interpolation technique is applied to assessment of the quality of sea surface temperature (SST) observations and quantitative analysis of the subpixel variability within satellite footprints of different size. Using retrieved satellite data as input, the new, global, multistage interpolation technique generates a trigonometric polynomial, providing a representation of the underlying physical SST field in functional form. The resulting interpolating function can be efficiently and accurately evaluated anywhere within the domain over which it was derived and its moments calculated to estimate the mean and variance of the field over desired sub-regions. Application of the technique is demonstrated for SST retrievals from the Moderate Resolution Imaging Spectroradiometer (MODIS), Spinning Enhanced Visible and Infrared Imager (SEVIRI), and Advanced Microwave Scanning Radiometer - Earth Observing System (AMSR-E) sensors. Comparison of the functional form with the data from which it was derived demonstrates how the technique can potentially help to identify small observational artifacts such as MODIS scan striping and residual cloud contamination. Integrals of the interpolating functions over successively larger spatial scales successfully emulate the retrieved SST at the different effective spatial resolutions and the second moments are consistent with the direct sample variances, and hence representative of the spatial SST variability of the available finer-resolution observations over the coarser scales. Using the approach, the variability of 1-km-resolution SST observations on open ocean grids of both 5- and 25-km resolution is found to be ~ 0.07 K. In regions of sharper gradients such as associated with strong localized diurnal warming, the variability within 25-km-resolution grids increases to as much as 0.4 K for sampling at 1-km resolution. The variability of 1-km observations on a 25-km-resolution grid is about 2.4 times greater than that on a 5-km-resolution grid. Broader application of the technique globally could help better quantify regional variations in the spatial variability, which would subsequently improve uncertainty estimates for existing satellite-based SST products.

1. Introduction

Interpolation methodologies are widely employed in the analysis and application of sea surface temperature (SST) products and imagery. The most common usage is to fill gaps in the spatial coverage of the available products. The capabilities of certain new approaches, however, enable promising additional unique and valuable applications. In this paper, we show that a novel High Resolution Multistage Spectral Interpolation technique (HRMSI) can be applied to assessment of the quality of the underlying SST observations and analysis of the subpixel

variability within satellite footprints of different size.

Multiple techniques have been developed and utilized to fill gaps in satellite SST imagery, particularly for providing spatially complete, daily SST analyses (level 4 products, e.g. [Martin et al., 2012](#)). Gaps in the spatial coverage of the retrieved satellite products result from many interfering factors such as clouds, precipitation, and aerosols. Atmospheric constituents can confound the radiative signal from the earth's surface, either degrading the accuracy of the physical retrievals or obscuring them completely. The source and size of the gaps is specific to the retrieved variable and the portion of the electromagnetic spectrum

* Corresponding author at: Colorado Center for Astrodynamics Research, University of Colorado, 431 UCB, Boulder, CO 80309, USA.

E-mail address: sandrac@colorado.edu (S.L. Castro).

¹ Currently with Google Inc., Boulder, CO, USA.

being utilized. For infrared retrievals of SST, clouds are the primary challenge, completely attenuating the infrared radiation emitted by the surface. Within the microwave portion of the spectrum, the radiation can pass through non-precipitating clouds, but regions of precipitation still obscure the surface. Gaps also clearly result from limited sensor swaths, particularly in the equatorial regions for polar-orbiting satellites.

Interpolation methods utilized in the generation of spatially complete SST products have been based largely on variations of the so-called optimal interpolation technique (e.g. Reynolds and Smith, 1994; Reynolds et al., 2007). The review by Martin et al. (2012) provides a good overview of many current SST analyses and the different approaches used in their generation. An additional approach for the reconstruction of missing data using empirical orthogonal function (EOF) decomposition applied to several oceanographic variables was described by Alvera-Azcárate et al. (2007) and a more recent technique for multiscale interpolation involving the use of wavelets has been applied to SST fields by Chin et al. (2017). While largely consistent and extensively utilized, the quality of the analyzed products remains variable, particularly for regions of high spatial variability such as current systems, temperature fronts, and coastal and polar oceans (Reynolds and Chelton, 2010; Martin et al., 2012; Dash et al., 2012; and Castro et al., 2016), and further improvement in the quality of the products is desirable.

Many other interpolation/regression techniques address practical problems associated with fitting scattered and noisy data with potentially large gaps. One comparison of several different techniques applied to the gridding of multiple climate variables was presented by Hofstra et al. (2008). The techniques differ in complexity and their utilization of the available frequency content of the data, which, in turn, is determined by the sampling density.

Of particular interest here is the ability of an interpolation technique to provide an explicit functional representation of the data as opposed to discrete interpolated values at selected points of interest. Such a functional form enables not only the evaluation of the interpolant at any arbitrary point within the domain of interest, but also the computation of a variety of quantities such as derivatives and integrals in analytic form. While fully applicable to the traditional problem of gap filling in satellite imagery and construction of level 4 SST products, the functional form and the above-mentioned properties offer the potential to address other significant problems impacting the assessment of SST retrieval quality and better quantification of the full uncertainty budget for the SST products.

One such data quality issue is identification of instrumental or processing “artifacts” within the retrieved satellite data related to the way in which the sensors operate or the product is constructed. While highly useful, satellite products are known to have important issues that can be extremely difficult to identify. Scan striping is a particularly important example. Sensors employing arrays of detectors like the Moderate Resolution Imaging Spectroradiometer (MODIS, Esaias et al., 1998) and the Visible Infrared Imaging Radiometer Suite (VIIRS, Murphy et al., 2006) are subject to issues with consistency in the calibration of the individual array detectors that can lead to the appearance of “stripes” or distortions in the derived products along instrument scans (Bouali and Ignatov, 2014). Accurate identification and removal of these artifacts is complicated, and despite significant efforts by the data providers, current products in use can still exhibit some features of the artifacts.

Another prominent problem is accurate determination of the spatial SST variability within the footprints of different resolution products and its impact on validation and merging of the retrievals. Validation of satellite SST products relies primarily in situ measurements from drifting and moored buoys as well as radiometric measurements collected from research vessels and ships-of-opportunity. These in situ sensors provide “point” measurements, which are then compared with larger areal averages from the satellite-borne radiometers. This

mismatch in spatial resolution directly affects the perceived accuracy of the satellite retrievals when validated in this manner. The contribution of subpixel variability to the overall uncertainty budget of satellite SST retrievals has long been acknowledged (e.g., Minnett, 1991; Cornillon et al., 2010; Castro et al., 2017), but has yet to be fully quantified due to measurement limitations. Accurate estimation of the spatial variability is dependent on complete, high-resolution sampling of the region of interest. Satellite data, however, is often incomplete due to the gaps described above and is limited by the sensor resolution. Techniques such as variogram analysis (Castro et al., 2010; Kent et al., 1999; Cressie, 1993) enable estimates of variability on larger scales to be downscaled but the results are still highly dependent on the quality and density of the available observations. The availability of an accurate fit to the data in functional form would provide a new way of estimating the fundamental SST variability since the function could also be integrated over arbitrary satellite footprints. Moreover, since the variance, σ^2 , of the SST distribution within a satellite's footprint is a measure of the subpixel variability, the second moment of the interpolating function could be explicitly evaluated to yield the predicted value of σ^2 at different spatial scales.

In this paper, we address these problems through application of a novel global interpolation/regression methodology employing trigonometric polynomials (i.e. sinusoidal functions, hence the name trigonometric interpolation) that yields an explicit functional representation of the retrieved satellite data. While the utility of intelligently filling coverage gaps with a physically realistic and continuous functional form is clear, this paper does not focus on this aspect or a comparison of this methodology with other techniques. Rather, the primary emphasis here is on retrieval quality assessment and spatial variability of satellite SSTs. Following a brief description of the technique (Section 2) and the satellite data sets employed (Section 3), the basic capabilities and limitations of the technique are shown in Section 4. Broader application of the technique is illustrated in Section 5 along with a demonstration of how it can help identify problems associated with scan striping, instrument artifacts, and residual cloud contamination, and be employed as a tool for objective quality data assessment and flagging. In Section 6, the moments of the interpolants for the different satellite SST sensors and grid resolutions are evaluated and employed to explicitly quantify the spatial variability across broad regions and different spatial scales.

2. New trigonometric interpolation methodology

The new global interpolation methodology applied in this paper employs trigonometric polynomials obtained using a multistage approach. It is well understood that global interpolation/regression approaches are able to capture a significantly greater frequency range (and thus achieve higher resolution) than local techniques. For a simple example, assuming periodicity and uniform sampling, trigonometric polynomial interpolation requires only two points per wavelength, whereas any standard polynomial-based or local interpolation techniques would require a significant oversampling factor to attain the same spatial resolution. The problem with attempting a direct global interpolation on scattered data, however, is that variations in data density will force a lower global resolution, losing the advantage over local interpolation. By introducing a multistage approach, we can accommodate regions of very different data density without resorting to the oversampling factor present in purely local interpolation techniques. Starting from a coarse scale, we gradually increase the interpolant resolution in stages as we partially fill the gaps in the data. The resulting image has adequate resolution in the gaps and the best possible high resolution in the regions with sufficient data density. In so doing, high frequencies present in the input data associated with noise or other nonphysical effects can also be removed. The method shares many similarities with the approach of Chin et al. (2017) but the main difference is that our technique is global rather than local in nature. This section provides a brief mathematical introduction to the technique.

While its application to SST fields is the subject of this paper, the technique can be applied to scattered measurements of physical fields of any type.

A standard trigonometric interpolation assumes a functional form of the target image in terms of a real-valued trigonometric polynomial (that is, a finite (truncated) Fourier series),

$$I(x, y) = \sum_{|k| \leq K, |l| \leq L} c_{kl} e^{2\pi i k a x + 2\pi i l b y}, \quad (1)$$

for some coefficients c_{kl} , positive integers K and L , and positive (scaling) constants a and b . Note that we can also write Eq. (1) as

$$I(x, y) = \sum_{0 \leq k \leq K, 0 \leq l \leq L} a_{kl} \cos(2\pi k a x + 2\pi l b y) + b_{kl} \sin(2\pi k a x + 2\pi l b y),$$

using real-valued coefficients a_{kl} and b_{kl} . To simplify the description, we assume that $K = L = L_{final}$ and say that L_{final} is the *degree* of the trigonometric polynomial $I(x, y)$. Thus, L_{final} determines the highest resolution that can be achieved by this trigonometric polynomial representation. If the input data were available on a (sufficiently dense) uniform grid, then the coefficients c_{kl} could be easily estimated using the Fast Fourier Transform (FFT). The satellite retrieved values $\{t_n\}_n$, however, are available at irregular (scattered) points $\{(x_n, y_n)\}_n$ corresponding to coordinates of longitude, x_n , and latitude, y_n . For this reason, to determine the coefficients c_{kl} , we have to solve a least squares problem of the form $\sum_{|k|, |l| \leq L_{final}} c_{kl} e^{2\pi i k a x_n + 2\pi i l b y_n} = t_n$. Unfortunately, for a typical irregular grid, this formulation leads to an ill conditioned problem. As a result, we would be forced to use a lower order trigonometric polynomial (i.e., to accept a low degree L_{final}). Our multistage approach to build the function I avoids this problem as follows.

Our first step is to approximate the data using a low degree L_0 , i.e., we solve the weighted linear squares problem (with weights based on the local grid density)

$$\sum_{|k|, |l| \leq L_0} c_{kl}^0 e^{2\pi i k a x_n + 2\pi i l b y_n} = t_n, \quad (2)$$

to find coefficients c_{kl}^0 for a small L_0 . With these coefficients, we build the trigonometric polynomial

$$I_0(x, y) = \sum_{|k|, |l| \leq L_0} c_{kl}^0 e^{2\pi i k a x + 2\pi i l b y}.$$

This step, by itself, can also be interpreted as denoising the data since high frequency noise is suppressed in the representation. Once I_0 is estimated, we use it to partially fill existing gaps in the data, which allows us to generate an augmented data set consisting of the initial set and the newly filled locations and their corresponding estimated temperatures. At the next stage, we use the newly built, larger input data set to solve a system of the form given by Eq. (2), but where we can now increase the value of the degree of the trigonometric polynomial from L_0 to (a larger) L_1 without causing ill-conditioning. The improved condition number (a factor controlling the sensitivity of the solution to small variations in the input data) for the larger system is due to the information added/gained through data augmentation. Solving this augmented system, we find coefficients c_{kl}^1 , which we use to build a new estimate I_1 . This I_1 is used to further reduce the gaps in the data and build a new system of the form given by Eq. (2), now with $L_2 > L_1$. We continue in this manner until we reach the final degree, L_{final} , yielding our final interpolant $I(x, y)$. Although this final degree should never exceed the degree implied by the region of highest density of measurements, it is also chosen to avoid overfitting by accounting for the level of measurement noise.

As a step in solving the least squares problem, we evaluate trigonometric sums on unequally spaced grids. For this purpose, we use the Unequally-Spaced Fast Fourier Transform (USFFT) (Dutt and Rokhlin, 1993; Beylkin, 1995) to assure both speed and accuracy.

In order to efficiently and accurately evaluate the resulting interpolating function (the trigonometric polynomial in Eq. (1)), on any

uniform or non-uniform grid, we once again use the USFFT. The USFFT has the same complexity as the FFT and its actual cost is only a small factor greater than that of the FFT. Note that for strictly uniform grids, the FFT may be more efficient, but for cases where the input requires significant padding, the USFFT is also a better choice. Unlike other more common interpolation techniques, our method can be understood as a regression technique since the values of the interpolating function at the original sample locations may not coincide with the values of the input data. In other words, evaluation of the interpolant $I(x, y)$ at the original observation points yields predicted values, not the original retrievals. In this manner, analysis of the residuals (interpolation error) can help identify potential measurement errors and processing artifacts in the satellite retrievals.

Furthermore, the interpolating function can be accurately integrated over any domain. In particular, it permits the explicit computation of the statistical moments of the SST distribution. Of relevance for the application considered in this paper is the evaluation of the second moment (variance) as it yields a measure of the SST variability at selected points and spatial scales.

If several interpolating functions are available for overlapping regions, it is easy to combine them in a consistent way since they can be evaluated on a sufficiently dense common grid. In fact, any other convenient functional representation could be used for the result of such combination. Thus, global SST maps can be obtained for all ocean basins and then combined with a minimal additional computational cost. One of the remaining problems is to estimate the temperature near continental boundaries. With our technique small islands can be treated as gaps and, after filling the gaps, the resulting interpolation can be restricted to open ocean areas. In the case of a large land mass, additional techniques need to be developed to take into account relevant coastal features and irregular boundaries.

In summary, the new interpolation methodology does not rely on standard local interpolation techniques or local statistical analysis, but rather obtains a global trigonometric interpolant, within a Fourier Analysis framework, through gradual interpolation of the target region (a multistage approach), making it more robust to overfitting. The theoretical basis for this methodology was recently developed by The Numericus Group, LLC, employing state-of-the-art mathematical algorithms. The use of the code should be coordinated by contacting The Numericus Group at info@thenumericusgroup.com.

3. Satellite SST data sets

The capabilities of the technique are illustrated through application to retrievals of SST from several current, commonly employed satellite sensors that span the range of available spectral type and spatial resolution. The products include infrared (IR) retrievals from the polar-orbiting MODIS from the NASA Aqua satellite and the geosynchronous Spinning Enhanced Visible and Infrared Imager (SEVIRI) on the Meteosat Second Generation (MSG-2) satellite, and microwave (MW) retrievals from the Advanced Microwave Scanning Radiometer - Earth Observing System (AMSR-E) flown aboard the Aqua satellite. The testing described in this paper uses data from 8 February 2009 over the South Atlantic Ocean between 4N–40S and 34W–8E. Scenes from this day, as shown in Fig. 1, were chosen because of the unusually high thermal IR coverage (few clouds obscuring MODIS and SEVIRI) and the presence of interesting large-scale thermal features such as a strong diurnal warming filament, clearly visible from all three sensors. An initial demonstration of the capabilities and limitations of the new technique additionally employs a composite of several days of SEVIRI data surrounding 8 February.

The MODIS retrievals (Fig. 1a) are at 1-km resolution as obtained from the level 2 (original satellite scan line/spot geometry) product processed by the NASA Ocean Biology Processing Group (OBPG) and downloaded from the NASA Ocean Biology Distributed Active Archive Center (OB.DAAC, <https://oceancolor.gsfc.nasa.gov/>). The study area

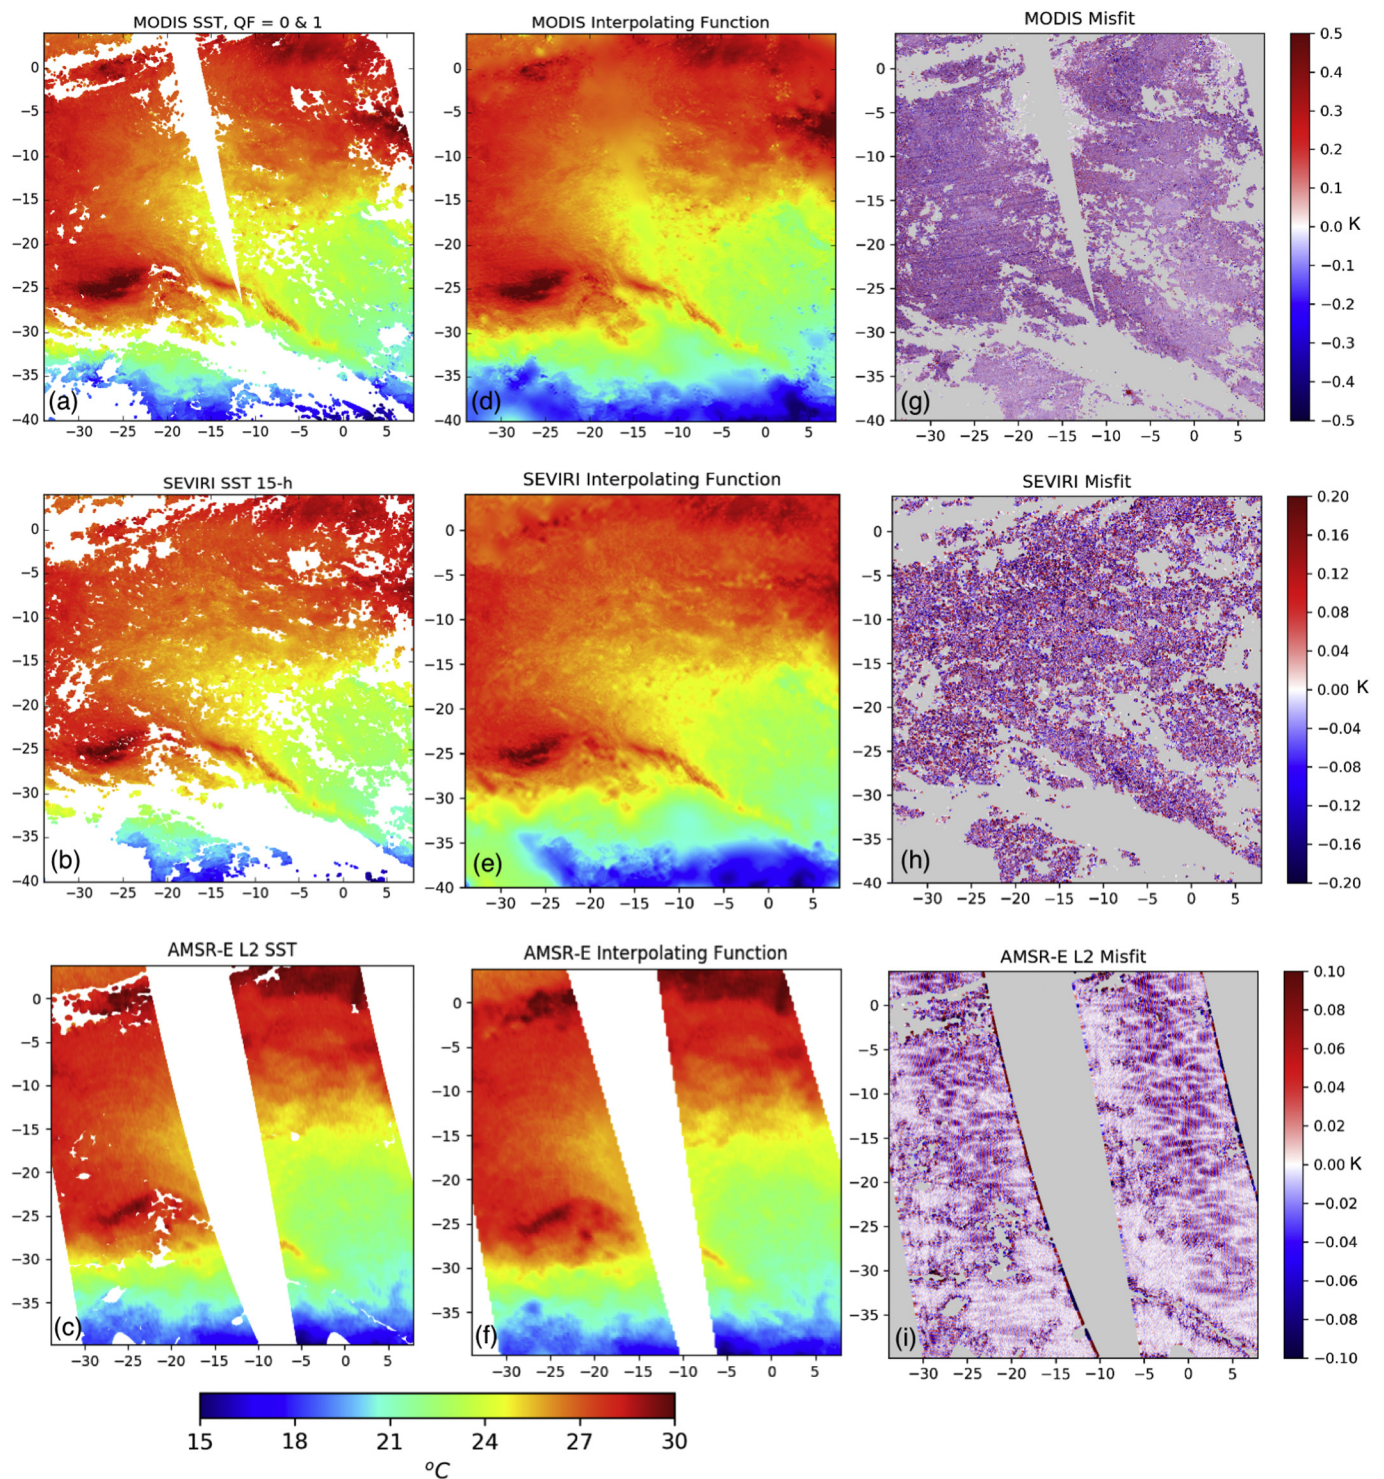


Fig. 1. Illustration of (a–c) the input data to the interpolating function, (d–f) the resulting interpolant evaluated everywhere in the domain, and (g–i) the corresponding misfit for (top) MODIS, (middle) SEVIRI, and (bottom) AMSR-E. The input data shown has outliers already removed as described in Section 5. Panels a–f all share a common color bar shown at the bottom of the figure. Panels g–i have distinct color bars shown just to their right.

was bisected by two ascending MODIS ground swaths, each composed of four 5-minute granules, over the period from 13:45 UTC in the bottom-right corner to 15:40 UTC in the upper-left corner. Initially, data with quality indices, QI, 0–1, i.e., cloud-free according to the OBPQ quality indexing conventions for their MODIS SST product (0: Good, 1: Questionable/suspect, 2: potentially cloud/sunglint contaminated, and 3: bad/cloud contaminated/failure), were considered as an effort to minimize the size of the gaps in MODIS coverage and test

the ability of the interpolant to recognize pixels with residual cloud contamination and potential inconsistencies in the quality indexing of the retrievals. It is important to emphasize that there is no uniform approach for satellite-derived SST quality flagging at present. For IR SST sensors, however, the quality indexing is generally tied to proximity to cloud.

The 5-km (0.05°) resolution SEVIRI SST product (Fig. 1b) is produced at the Meteo-France/Centre de Météorologie Spatiale (CMS,

Lannion, France) within the framework of the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) Ocean and Sea Ice Satellite Application Facility (OSI-SAF) and was obtained directly from IFREMER (The French Research Institute for Exploitation of the Sea). The product is a gridded (level 3 collated), hourly product obtained by averaging data obtained at full resolution every 15 min (EUMETSAT, 2011). From its geosynchronous orbit centered over 0° longitude, SEVIRI provides continuous spatial sampling over a grid from 60S–60N and 60W–60E. The selected image in Fig. 1 corresponds to 1500 UTC. Data corresponding to quality flags 2–5 were included. Additional quality indices beyond the “best” pixels (OSI-SAF QI convention: 0: no data, 1: bad, 2: very low quality, 3: low quality, 4: acceptable, 5: best) were again included to increase coverage and to test the ability of the interpolating function to evaluate the data quality assessment performed by the data producer.

The AMSR-E microwave SST data (Fig. 1c) were obtained in level 2 format from Remote Sensing Systems (RSS, <http://www.remss.com/>, Wentz and Meissner, 2004). The retrievals have a native resolution of approximately 38 km, but there is a high degree of oversampling. Initial analysis of the data considered quality indices 2–4 (RSS QI convention: 1: bad, 2: suspect, 3: should-be-good, 4: good) to evaluate the need for additional screening as described below. Note that AMSR-E retrievals are not affected by non-precipitating clouds and, thus, the quality indexing is not based on proximity to cloud, but rather on proximity to precipitation, ice, or land, and other factors affecting the quality of the MW retrieval such as sun glint, high winds, and electromagnetic interference. Since the MODIS and AMSR-E sensors fly on the same satellite, their spatial and temporal coverage is nearly coincident (see Fig. 1a and c).

4. Demonstration of the interpolation methodology

An initial test demonstrates the new interpolation methodology's ability to both reproduce existing data and fill in data gaps. For this experiment, a spatially complete gridded data (no gaps) was required along with some realistic cloud mask. The basis for the data grid was a maximum value composite (MVC) generated from all hourly SEVIRI SST scenes over the study area over four consecutive days starting on 8 February 2009. To fill in a small number (< 5%) of remaining gaps in the MVC, SEVIRI retrievals from the days prior to the starting day were inspected for additional observations at the empty grid cells. The first available observation, going backward in time, was used to fill the respective gap. The resulting SST scene is shown in Fig. 2a.

To test the interpolation methodology, SST values corresponding to representative gaps due to cloud cover were removed prior to generating the interpolating functions. The source of this cloud mask was taken as the screened clouds from the 1500 UTC SEVIRI hourly product from 8 February 2009 as shown by the gray areas in Fig. 2b. An interpolating function was generated using the remaining observations (clear pixels in Fig. 2b), and then evaluated at every 0.05° grid point corresponding to the original, spatially complete, SEVIRI composite (all obs in Fig. 2a). The resulting image is shown in Fig. 2c. It is important to emphasize that at the locations where input data were available, the values shown in 2c are from the interpolating function and not the original retrievals. Qualitatively, the interpolating function is observed to do a good job in reproducing the original SST data in all but the regions with the largest assumed cloud cover.

To quantify the accuracy of the interpolating function, the difference between the functional value of the interpolant and the original SEVIRI SST composite was evaluated at each grid point. Termed the “misfit,” these differences are plotted for the clear observations from which the function was derived (Fig. 2d) and for all the points including the assumed cloud gaps in the input data (Fig. 2e). The misfit shows a speckle-like noise pattern with very small amplitudes (within ± 0.2 K) where the input data were available, demonstrating that the clear observations are well reproduced by the interpolating function. As

expected, however, the misfit gets larger as the size of the gaps increases and, additionally, the speckle noise adds coherently and self-organizes in large constructive and destructive patterns (the saturated red and blue patches in Fig. 2e corresponding to the largest gaps within the cloud mask). This shows that the ability of the interpolating function to replicate the data within the gaps is constrained by the size of the gap.

The misfit values from Fig. 2e were then binned as a function of the Euclidean distance to the nearest available observation to show how closely the interpolator replicates input observations and to provide an indication as to how large of a gap can be effectively filled with this technique. The results for gaps of increasing size up to 25 pixels (125 km) are shown as the black trace in Fig. 3. The misfit has a mean value of 0.10 K for the points at 0-distance (i.e., where observations were available) and then increases monotonically as the gap size (distance to nearest observation) increases. If an acceptable accuracy for the interpolation error were assumed to be 0.2 K, these results demonstrate that the technique is able to effectively fill gaps of up to about 10 pixels (50 km) in size using the information contained in the available data.

While the focus of this paper is not on the skill of the interpolation at filling gaps, an additional experiment was performed to demonstrate how the method could be employed to fill larger coverage gaps as is done in existing SST analyses. In the regions where clouds were artificially introduced (Fig. 2b), the original 5-km resolution SEVIRI composite data was replaced with a 25-km average value (i.e., 5×5 -box averages of neighboring grid cells centered at each cloudy pixel) to simulate the availability of lower resolution AMSR SST retrievals under clouds. A new “blended” interpolating function (not shown) was generated based on this combined data. The method was only allowed to utilize the additional lower resolution input once the distance from the original higher resolution observations reached a user-specified distance of 10 pixels. The misfit resulting from evaluating the blended interpolator at the initial 5-km SEVIRI composite SSTs is displayed in Fig. 2f and the corresponding binning as a function of the Euclidean distance to the nearest available observation is shown as the red trace in Fig. 3. The binned misfit reaches its maximum value at a distance-from-observation of about 10 pixels corresponding to the transition to inclusion of additional data. Beyond this threshold, the blended interpolator can be evaluated to fill larger and larger gaps with remarkable accuracy and precision, as the misfit plateaus at 0.1 K for larger distances. For this idealized case, where 25-km resolution data are assumed to always be available to guide the interpolation when there is no high resolution data at hand and the coarser resolution data reasonably reflects the SST patterns at finer resolution, the simple addition of the complementary coarser data significantly improves the accuracy of the interpolator's spatial predictions at the unsampled locations. It is for this very same reason that existing SST analyses blend data from multiple satellite sensors.

Although we do not seek to directly compare the performance of different interpolation methods, an additional test places these results in the context of an independent, alternative method. A simple, distance-weighted, sphere-of-influence based interpolation technique was also tested based on the available “cloud-free” observations in Fig. 2b. For each grid point, the technique searched for all available clear ocean observations within a radius of 100 km and then computed the weighted average of those data with the weights given by one over the square of the radius (inverse distance). Observations directly at the grid point (0 distance) were assigned a weight of one. If there are no observations within the 100-km radius, the technique fails to produce an estimate, and thus it cannot be used to produce complete fields. The misfit between the weighted interpolation estimates and the original composite SEVIRI values was computed and further binned with distance from the observations and shown by the green trace in Fig. 3. Within smaller gaps, our high-resolution multistage interpolation technique outperforms the distance-weighted interpolation, but as the

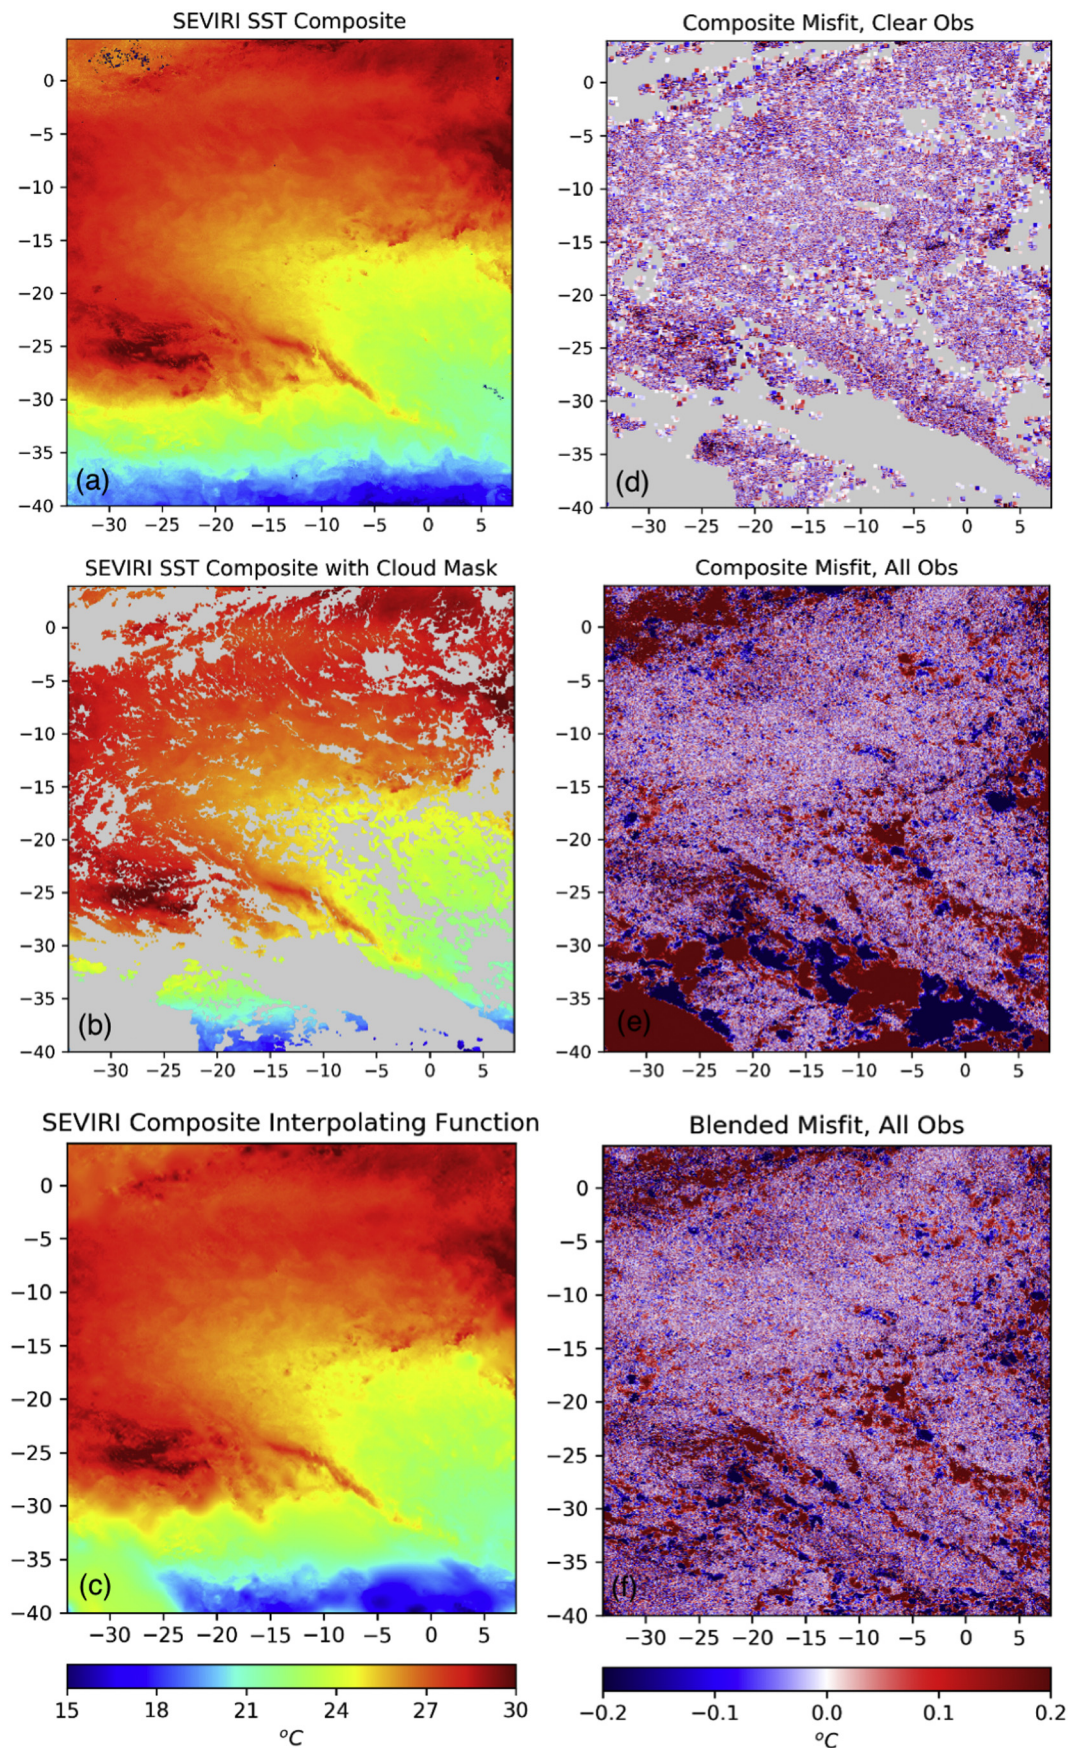


Fig. 2. Initial demonstration of the interpolation technique applied to a SST composite derived from SEVIRI. Panel (a) shows the original composite, (b) shows the composite with the cloud mask applied, and (c) shows the interpolating function evaluated everywhere on the composite grid. The resulting misfit is shown in panels d–e for both the clear values and all values respectively. The corresponding “blended” misfit in panel e was derived in an additional experiment where additional inputs to the interpolation in cloud screened regions were supplied based on simulated lower resolution data. See text for details.

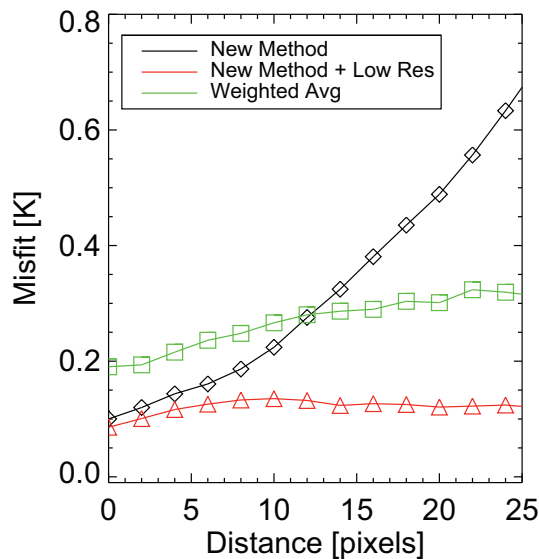


Fig. 3. Mean misfit evaluated as a function of Euclidean distance to the nearest available observation for the identified interpolation techniques. The black trace reflects the results obtained from our interpolating function derived from inputs in clear areas only. The red trace was obtained from an alternative interpolating function obtained with additional simulated lower resolution SST input in the cloudy regions. The green trace applies to an independent simple distance-weighted interpolation approach. Solutions for this weighted technique exist only up to a distance of about 30 pixels. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

gap increases beyond about 12 pixels, the localized weighted average does better than the interpolator derived from the cloud-free observations. Although counterintuitive at first glance, a global interpolation method such as ours uses global information content to fill in the small scale variability that in small gaps is smoothed by the average, whereas in large gaps, the new method is conservative opting for background values that are even smoother than a localized average. More detailed comparisons of different techniques and application to filling gaps in multi-sensor data are deferred to a separate paper.

5. Application of the Interpolant to scene quality assessment

Beyond the basic illustration of the capabilities of the technique shown in the previous section, the broader utility of the interpolation methodology can be demonstrated through application to the individual sensor SST scenes described in Section 3. One such unique application is helping to assess the quality of the data used to build the interpolating function, including the identification of any potential artifacts in the input data.

Interpolating functions were generated for each of the sensor products in a two-step approach: first using inputs with minimal quality control and a second step using refined inputs obtained after removing outliers identified based on the magnitude of the interpolation errors from the first step. In the preliminary screening, data from all the quality levels listed in Section 3 were used as inputs. This included data of more questionable quality that are typically excluded in other applications, such as the construction of SST analyses (level 4 products). The reason behind this approach is twofold: 1) to assess the ability of the interpolating function to help with additional screening of poor quality data, and, 2) to retain as much data as possible to build an interpolant able to accurately fill in gaps in coverage. So-called “questionable” data in IR retrievals can be extremely challenging to screen using the conventional cloud masking methodologies employed by most data producers. Traditional satellite cloud screening methods

Table 1

Percentage of original observations of specified quality index rejected in first stage screening.

Sensor	Best quality	Questionable quality
MODIS	0.01% (QI 0)	1.05% (QI 1)
AMSR-E	1.0% (QI 4)	4.5% (QI 2)

rely on a series of thresholds to identify large retrieval errors. These techniques are not perfect, however, as in many cases, such as with low stratus and polar maritime cloud, the clouds exhibit very little contrast from the underlying sea surface. Thus, flagged observations often include a mix of both contaminated and uncontaminated retrievals and valid data points may be discarded when an entire quality level is excluded from an analysis.

The preliminary screening was successful in identifying contaminated observations within the questionable levels while retaining other retrievals highly consistent with their surrounding values. Interpolation errors were computed as the difference between the actual satellite-retrieved SSTs used to derive the functions and the values of the interpolating function at the locations (nodes) of the SST input. Note that unlike exact interpolators, the interpolation error at the nodes is non-zero with the proposed method, accounting for uncertainty in the observations. Distributions for the interpolation errors were obtained from normalized histograms. Differences exceeding a specified threshold suggested unreasonable satellite retrievals inconsistent with a smoothly varying SST field. Residual outliers identified in this fashion were then rejected prior to generating the interpolating function. Sharp SST frontal features or narrow patches of diurnal warming pose challenges where valid data could also be rejected and represent a key tradeoff in selecting the rejection threshold. A high rejection threshold, nominally based on the 99% quantile of the distribution of the differences, was utilized in an effort to reject residual cloud contamination only.

The results reveal how the retrievals rejected in this manner compare to the quality indices assigned by the data producers. The percentage of “best” and “suspect” quality observations rejected is shown in Table 1 for MODIS and AMSR-E. No SEVIRI input observations were rejected during the screening process. For MODIS, the observations were screened using the SEVIRI-derived interpolating function rather than using direct comparison against values of the MODIS interpolating function, as the SEVIRI interpolator was more effective at identifying outliers given the higher levels of uncertainty in MODIS. For both MODIS and AMSR-E, a larger percentage of the observations indexed as “suspect” by the data producer were rejected, but some “best” quality observations were also removed.

In the second step, the input data retained after the exclusion of outliers in the screening step were used to compute the final, refined interpolating function for each sensor. The quality-controlled SST inputs are shown in Fig. 1a–c. The final degree, L_{final} , used for the interpolating trigonometric polynomials was 1024 for MODIS, 512 for SEVIRI, and 256 for AMSR-E. The degree was algorithmically derived as part of the multistage approach to build the interpolating function. The choice of the degree is the result of, at each stage, checking the size of the intermediate residuals and monitoring the condition number of the system to solve for the coefficients of the intermediate interpolating function. The polynomial degree reflects the ability of the interpolating functions to reproduce the observed variability without fitting residual noise in the data. The degree also reflects some correlation with the sensor product resolution. For efficiency in the current technique, degrees based on powers of 2 were used, but this will be generalized to any degree in the future. Because of the large gap between the AMSR-E swaths, separate interpolating functions were generated independently for the western and eastern swaths. The manual separation of input data for the two swaths was performed for AMSR-E only.

Table 2

Statistics comparing the SST values retrieved directly or obtained from the interpolating function against collocated buoy SST measurements. The bias is computed as retrieved value minus the buoy measurement.

Sensor	Satellite retrievals			Interpolator at input			Interpolator in gap		
	Counts	Bias (K)	STD (K)	Points	Bias (K)	STD (K)	Counts	Bias (K)	STD (K)
MODIS	64	−0.16	0.60	64	−0.11	0.53	40	−0.14	0.60
SEVIRI	26	0.06	0.28	26	0.01	0.21	50	−0.07	0.66

To illustrate the performance of the code, we recorded the execution time needed to generate the interpolating function with the highest degree on a workstation with an Intel i7-7700K CPU operating at 4.2 GHz. The MODIS input (Fig. 1a) contains 5,172,397 pixels and the interpolant constructed has degree 1024. Interpolant construction required 957 s (median of 3 timings, excluding I/O), that is, about 16 min. Next, we evaluated the interpolant on the whole domain displayed in Fig. 1d ($42 \times 44^\circ$) which contains 19,350,000 pixels. This evaluation required 37.4 s (median of 3 timings, excluding I/O). Our current implementation has not been heavily optimized, and we anticipate that significant acceleration is possible, especially in the construction of the interpolant. To expand the implementation of the technique to global applications would require partitioning the ocean basins into similar sub-domains.

The resulting values of the interpolating functions evaluated on full grids, corresponding to the resolution of the input data from each sensor, are shown in Fig. 1d–f. For AMSR-E, due to the high degree of oversampling, a regular 0.25° resolution grid was used, consistent with the corresponding level 3 products provided by Remote Sensing Systems. As before, at the nodes where original data were available, the values shown are from the interpolating function and not the original retrievals. Qualitatively, at the large image scale, the interpolating function for each sensor captures the spatial SST variability extremely well, further illustrating the strong potential of the technique. At nodes where input data were available, the interpolating function accurately reproduces all major features, including the region of diurnal warming. Within the original sampling gaps, the interpolated functional values generally appear physically realistic, while providing a continuous SST field. Not surprisingly, as emphasized in the initial demonstration in Section 4, the ability of the interpolating functions, derived based on single sensor data as done in this demonstration, is more limited in regions of extensive data gaps such as between satellite swaths. In these regions (e.g., Fig. 1d), the functional values are notably smoother, as the methodology takes a conservative approach in poorly sampled regions so as not to introduce spurious fine resolution features. At the edge of the image domains where there is no input data to bound the function, some larger errors can be observed, as the function is not designed to extrapolate outside the spatial domain for which it was built.

The general skill of the interpolating functions in reproducing the underlying SST structure and variability for the MODIS and SEVIRI satellite products can also be assessed through comparisons with independent drifting buoy observations. Buoys have long been a standard against which satellite SST retrievals have been validated. For this evaluation, quality controlled drifting buoy data obtained through the NOAA iQUAM in-situ SST monitoring system (<https://www.star.nesdis.noaa.gov/sod/sst/iquam/>; Xu and Ignatov, 2014) were compared separately against the direct satellite retrievals and against values of the interpolating functions for the scenes considered above both in regions where retrievals were available and in the gaps where no training data were available. For the direct satellite retrievals, buoy observations were matched with satellite retrievals within 1 h in time and either the closest grid cell (SEVIRI) or the nearest retrieval (MODIS) within a 20 km radius. With the interpolating function, the satellite SST is matched with the buoy by directly evaluating the function at the exact node corresponding to the buoy location. Since the MODIS image uses

granules from 13:40 to 15:45 UTC, buoys present in the domain from 1300 to 1600 UTC were extracted for the MODIS matchups. In the case of the 1500 UTC SEVIRI, buoys within 1400 and 1600 UTC were considered.

Statistics for the buoy comparisons are summarized in Table 2 and further support the skill of the derived interpolating functions. One notable result is that the rms difference with respect to the buoys within clear ocean pixels is smaller when using the interpolating function estimate than for the retrievals themselves, for both MODIS and SEVIRI. The biases are also generally similar. These findings reflect positively on both the overall accuracy of the interpolating function and its ability to be evaluated at precise point locations. While the agreement between the buoy measurements and estimates derived from the interpolating function is somewhat degraded in cloudy regions for SEVIRI, the values are quite similar between the clear and cloudy regions for MODIS suggesting some value of the interpolated MODIS results within the cloud-covered gaps.

A more quantitative pointwise evaluation of the resulting interpolation error at nodes where the input data were available better illustrates the ability of the methodology to further assess the quality of the underlying data. Residual differences (again termed the misfit) were computed between the functional value of the interpolant at the location of the available satellite SST retrievals and the actual retrieved SST values as in Section 4. Corresponding misfits for MODIS, SEVIRI and AMSR-E are plotted in Fig. 1g–i, respectively. The typically small values of the misfit again demonstrate how the function is able to accurately capture the observed SST variability throughout the scene. Only for MODIS are residual differences comparable to, or larger than, the expected satellite retrieval accuracies (about 0.5 K). Moreover, closer examination of patterns in the misfit reveals how powerful the technique can be in identifying potential remaining artifacts related to cloud contamination, instrumental effects, and sampling errors.

Locations of the largest values of the misfit (i.e., the top 1 or 5%) are further highlighted in Fig. 4 as the points plotted in black. It is notable that the patterns suggested by the largest misfit values are quite different for the different sensor types. This suggests that the differences cannot be simply ascribed to properties of the functional fit, but rather may be associated with the underlying data. Given that the exact cause of the misfit is not always known with certainty, the differences, however, must be carefully examined before concluding that they are indeed data artifacts. Statistics comparing these largest misfit values with all other points are shown in Table 3. The increase in the mean misfit for the largest differences is clear, particularly for MODIS.

The MODIS misfit maps (Figs. 4a and 1g) show stripes consistent with individual scan lines from the sensor, particularly in the western swath. Scan striping has long been a problem for MODIS SST retrievals resulting from the use of an array of independent detectors with imperfect relative calibrations and a multi-sided (cross-track) scan mirror (Gumley, 2002). Substantial effort has gone into reducing the impact of scan striping, and the results have improved to the point where striping cannot be visually detected in the input MODIS scene in Fig. 1a. Comparison of the smooth functional form and the retrieved data, however, suggests that striping is still present in the retrievals. Identification of striping at this low magnitude is challenging, and few other existing techniques have exhibited the ability to show the striping as clearly as the interpolating function does here.




Fig. 4. Locations of the largest misfit values for (a) MODIS, (b) SEVIRI, and (c) AMSR-E. The locations exceeding the percentile indicated in the title of each panel are plotted in black while all other observation locations are plotted in gray.

The other most coherent regions of misfit for MODIS (Fig. 4a) are found in regions just left of the scan center in the northern portion of the scene. The cause for this misfit pattern is not fully understood, but the region in the eastern swath does show correlation with a discrete portion of the scan geometry suggesting that the effect could again be instrumental in nature. Other isolated larger differences could have contributions both from sampling effects and limitations in the functional fit. Instrumentally, the differences could indicate random noise or potential problems with residual cloud contamination or subpixel cloudiness in the corresponding SST retrievals. With respect to the functional fit, very fine scale variability at the MODIS resolution might not be fully resolved given the tradeoff between distinguishing natural high frequencies from noise.

The magnitude of the misfit can again be compared to the original quality flags assigned by the data producer. Using MODIS as an example, the mean misfit expressed as a function of the quality index is presented in Table 4. Results are shown both for the absolute value of the misfit and for the true mean where positive and negative values can partially cancel out. As noted before, the mean misfit is very small in magnitude, reflecting the overall goodness of the functional fit. A slight increase in magnitude of the misfit is observed for the “questionable” data quality that was retained following the initial quality control where only the largest outliers were removed, but the fact that the change is so small further justifies the inclusion of the additional data in the analysis.

The SEVIRI misfit (Figs. 4b and 1h) exhibits a largely random pattern with the exception of a concentration of larger misfit values along the axes where diurnal warming is observed to be greatest. For these areas near strong diurnal warming, the source of the difference is likely algorithmic rather than physical as the interpolating function may be overly smoothing across the sharp SST gradients. The diurnal warming introduces stronger gradients than elsewhere within the scene and forcing the interpolating function to resolve these gradients would have resulted in the introduction of high-frequency noise. This is another clear instance of the balance between resolving high-resolution features (i.e., preserving the interpolator from over-smoothing) while keeping high-frequency noise to a minimum. Elsewhere, the large differences do not appear to suggest any correlation with potential residual cloud contamination in the retrievals with more suspect quality levels, suggesting that the majority of the cloud contaminated data have indeed been excluded by the data producer. This again suggests that use of lower level quality data does not necessarily have a negative impact; in fact for this application, the inclusion of additional data might have been highly beneficial since other results (Castro et al., 2014) have suggested that extreme diurnal warming events tend to be flagged as poor quality data by the automated quality control processes implemented with operational data.

The AMSR-E misfit values (Fig. 1i) are very small overall (about 0.1 K), but do show some tendency for larger values (Fig. 4c) to concentrate along the edges of screened precipitation regions (the main source of the gaps inside the swaths of the AMSR-E input), swath edges, and near the swath center. The concentration of elevated misfit values forming “halos” around the pre-screened features suggest that additional unflagged retrievals in the vicinity of the features might still be modestly impacted by the contamination source. While the functional fit could also potentially over-smooth across edges in data coverage, the fact that such similar halos are not observed for the other sensors lends support for unflagged rainfall contaminated pixels. The interpolating function is also able to highlight known anomalies in the retrieved values at the outermost spots on the right side of the scan. The increased misfit corresponds to known contamination caused by parts of the Aqua spacecraft entering the field of view of the first (right-most) pixels of each scan. These pixels are excluded in creation of the Remote Sensing Systems level 3 products but are included with the level 2 data (Wentz and Meissner, 2004). The misfit values just to the right of the scan center could possibly suggest very small calibration differences

Table 3

Statistics characterizing the difference in misfit for those points above and below the percentiles selected for the greatest differences between the observed and estimated SSTs.

Sensor	Counts	Mean (K)	STD (K)	Counts	Mean (K)	STD (K)
	< 99%			> 99%		
MODIS	5,120,673	0.12	0.01	51,724	0.58	0.11
SEVIRI	282,333	0.06	0.05	2852	0.28	0.04

Sensor	Counts	Mean (K)	STD (K)	Counts	Mean (K)	STD (K)
AMSR-E & L Swath	85,790	0.02	0.02	4515	0.19	0.13
AMSR-E & R Swath	95,639	0.02	0.01	5034	0.13	0.10

with scan spot, but these are much smaller than the expected retrieval accuracy. Interestingly, while there is some correlation between the larger errors and scan spot, there is no suggestion of correlation with scan line. The larger pixels at the left edge of the scan are due to numerical issues with the interpolator at the boundaries of the interpolated swath areas, as it is not equipped to extrapolate outside areas for which it was built. Recall that in the AMSR-E case, two separate interpolants were constructed due to the wide gap between swaths. Careful examination of the AMSR-E image in Fig. 1c reveals the presence of periodic conical scan patterns in the SSTs along the satellite tracks, especially on the left swath, which is mimicked in the interpolated data (see Fig. 1f). Trigonometric polynomial interpolants are particularly well suited for fitting periodic, smooth features. The fact that these patterns have smooth curvature with an apparent periodicity spanning over long distances along the satellite ground track means that it can probably be characterized by a sinusoid function, which explains the failure of the trigonometric interpolator to identify it as an artifact in the AMSR-E data unlike with the linear striping in MODIS.

6. Application to sub-pixel SST variability

Another novel application of the interpolating function is to explore subpixel variability within satellite products of differing spatial resolution. In geophysical phenomena, the measurement of a physical variable associated with a point, x , in a two-dimensional space represents an average value over some continuum area. From a practical standpoint, however, we tend to associate a measurement with its point value, $Z = z(x)$. For instance, it is common to implicitly treat the satellite-retrieved SST for a pixel as a point value, particularly when validating it against in situ point observations, even though the satellite measurement integrates the radiation coming from the entire area within the satellite's footprint. There is an estimation error stemming from this representation because there always is spatial variability in nature. This is what in satellite remote sensing has been broadly termed the point-to-pixel discrepancy. The unresolved spatial variability has an effect on the perceived uncertainty of satellite-derived products when the satellite retrieval is validated using observations with significantly finer resolution than the satellite. Vinogradova and Ponte (2013) presented a good recent discussion of this impact on sea surface salinity observations from the Aquarius satellite.

Table 4

Mean magnitude of the misfit in the MODIS interpolating function expressed as a function of quality index.

	Best quality	Questionable quality
Mean misfit (K)	0.00	−0.01
Mean absolute misfit (K)	0.12	0.16

Accurately quantifying the true variability is very challenging in practice and the results are inevitably tied to the size of the region/footprint and the resolution and density of the available observations. The concept behind an observationally-based estimation of spatial variability is to densely sample regions the size of the satellite's footprint using instrumentation of sufficiently high spatial resolution to capture all the scales of variability within the footprint. The variance of the high-resolution observations then provides a measure of the spatial variability within the coarser pixel. This challenge is currently met through advances in technology with small, high resolution instruments mounted on fast moving platforms, with the capacity to obtain measurements that fully resolve the finer scales of variability. While viable, these field campaigns are expensive and cannot be conducted globally. Another option, employed by Vinogradova and Ponte (2013) in the study of salinity, is to utilize high-resolution model output fields in the place of direct observations to represent the variability within a given footprint.

To estimate the subpixel variability within coarser resolution products, it is also possible to use measurements from higher resolution satellite observations. The limiting factors of this approach are the spatial resolution of the available satellite sensors, and obtaining sampling with adequate density to accurately estimate the true variability of the field. While the first problem is constrained by the available data, use of the interpolating function can potentially help with the latter. Estimates of variability obtained exclusively from satellite retrievals are limited by gaps and the underlying noise of the data. Collocation and direct aggregation of all observations within successively larger footprints is also computationally expensive. In contrast, interpolating functions derived from the high-resolution observations are continuous (gap free) and can be up-scaled or integrated easily over larger areas corresponding to the coarser resolution products. Thus, interpolating functions have the advantages of providing physically consistent values everywhere within the desired coarser footprints (aside from the largest gaps), ease of analysis, and smoothing out other sensor-specific noise that could affect direct averages and variability estimates.

Since the interpolant $I(x,y)$ described in Eq. (1) is an explicit trigonometric polynomial of two real variables, we can efficiently compute its moments over any desired region by reducing the problem to an application of the FFT or the USFFT. We are particularly interested in the second central moment of the function to estimate the variance of the field. The aim is to compare this estimated variance with the direct sample variance of the high-resolution pixels as they are aggregated over successively coarser resolution grids. This estimated variance can, in turn, inform estimation of the subpixel variability and its contribution to an uncertainty budget of satellite SST retrievals.

The differing resolutions of MODIS (1 km), SEVIRI (5 km), and AMSR-E (25 km for the level 3 grid) provide a means to investigate the subpixel SST variability on retrievals of scales from approximately 5–25 km. The MODIS interpolating function can be used to estimate

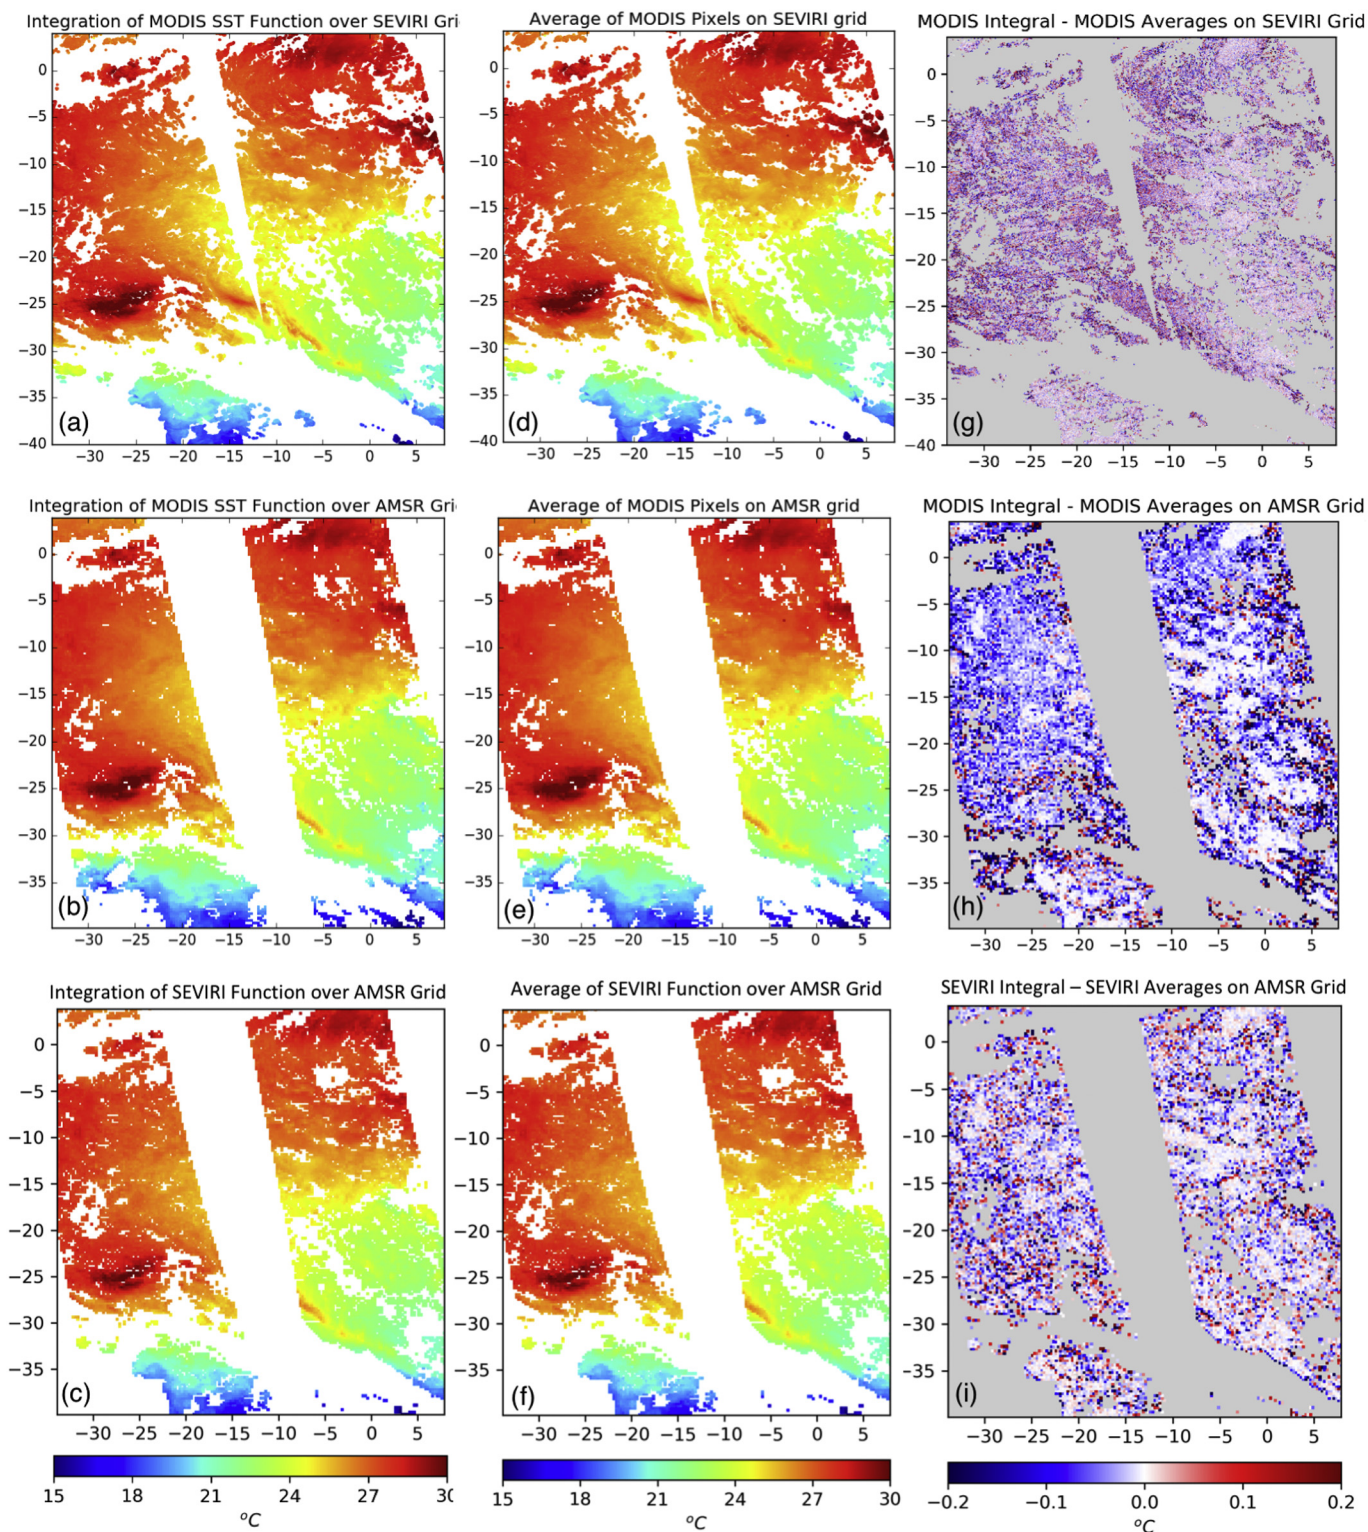


Fig. 5. Evaluation of whether the integral of the interpolating functions is a direct analog to the average of the underlying observations on the selected grids of differing resolution. The top panels correspond to the evaluation of MODIS resolution sampling on the SEVIRI grid, the middle panels to MODIS resolution sampling on the AMSR-E grid, and the bottom panels to SEVIRI resolution sampling on the AMSR-E grid. Panels a–c show the integrals of the interpolating functions over the grid cell, panels d–f show the average of the available observations within the grid cell, and panels g–i show the corresponding differences between the integrals and averages. Common color bars are shown at the bottom of each column.

variability within the SEVIRI and AMSR pixels and the SEVIRI interpolating function can be used as an independent estimate of variability within the AMSR-E footprint.

Prior to using the interpolating functions to infer the spatial

variability, it is useful to further verify their ability to accurately estimate the mean SST field over the selected coarser resolution footprints (e.g., the ability of the MODIS interpolating function to estimate an observation over the coarser AMSR-E footprint). If the interpolating




Fig. 6. Distribution of the differences between the integrals of the interpolating functions and averages of available observations for the results shown in Fig. 5 g–i. Results for the different sensor and grid combinations are plotted by color as shown in the legend.

functional accurately represents the underlying physical field, the integral of the interpolant over the footprint should approximate the satellite-retrieved value for that footprint. To test the functional approach, the integral of the interpolants over coarser resolution footprints was compared against direct arithmetic averages of all the available higher resolution observations within the same domain. There is a caveat, however, in that the comparison involves the mean of a continuous function over the integration domain and the average of a discrete sample with gaps in coverage. Missing observations within the gaps are expected to contribute to differences with the interpolant integral, but the comparisons can also provide insight into how consistently the gaps are filled in by the interpolant. Note that the integral could also be compared with the original coarser resolution observation on that footprint (say a 5-km integral of the MODIS interpolant with the corresponding native 5-km SEVIRI retrieval), but that would introduce an intercomparison between sensors with different characteristics that we want to avoid. Here, we wish to evaluate the ability of the interpolants to supplant the direct, higher resolution observations from which they were derived.

Results are shown in Fig. 5 for comparison of the MODIS integrated interpolating function and corresponding averaged MODIS observations on the SEVIRI and AMSR-E grids, and the SEVIRI function and observations on the AMSR-E grid. The differences do not, in any way, reflect measurement errors between different sensor types. To first order, the functional integrals agree very well with the averages of the data from which they were derived. Mean differences are near zero (< 0.04 K) and coherent differences are relatively small in magnitude. Overall, the largest differences, not surprisingly, occur on the edge of large data gaps. Increased differences do also occur in stronger gradient regions and near the diurnal warming feature (25S, 25W) where SST variability is greatest and accurate averages are dependent on all possible measurements (which may not be available). Those differences, however, are still generally within expected retrieval accuracy.

Distributions of the difference values are also shown in Fig. 6. While the mean difference between the interpolant integral and the direct MODIS average on the AMSR-E grid is very small (-0.04 K), there is a skewness of the difference distribution to negative values indicating that the value of the integral is typically smaller than the average. This could potentially result from the presence of cloud “halos” or slight residual cloud contamination on the edge of screened regions causing the interpolating function to project slightly cooler SST values into the enclosed regions with missing observations.

To better appreciate the fine scale differences, the corresponding results for a smaller region, centered on the large diurnal warming event, are shown in Fig. 7. Visually, the integrals of the interpolating functions, again, appear very similar to the direct averages of the measurements and the magnitude of the differences are reasonably small (mean differences all again < 0.04 K). For the MODIS difference on the SEVIRI grid (Fig. 7g), the effect of scan striping is again visible but no other coherent patterns emerge. For MODIS on the AMSR grid (Fig. 7h) the differences are elevated near the large gradients associated with the region of diurnal warming and also near the edges of coverage gaps. The differences in the SEVIRI results evaluated on the AMSR grid (Fig. 7i) are quite random except for some enhancement on the edge of data gaps. The mean difference between the integral and the sample average is < 0.01 K. Overall, the results suggest that the interpolating functions can accurately capture the variability within the coarser sensor footprints and support further use of the functions to quantify this subpixel variability.

Based on this success, the ability of the second moment of the interpolating function to characterize the variability within the coarser resolution pixels was explored next. The second moment of the interpolating functions (or variance functions) was computed and integrated over regions corresponding to the SEVIRI and AMSR-E footprints as for the first moment (the mean of the function) above. These explicit variances were then compared with direct computations of the sample variance for the corresponding higher resolution MODIS and SEVIRI measurements within the same domains. Graphical results are shown in Fig. 8 for the MODIS retrievals and interpolating function evaluated on both the SEVIRI and AMSR-E grids, and corresponding SEVIRI products on the AMSR-E grid. Distributions of the variance for the entire domain, computed in both ways, are further plotted in Fig. 9. A more detailed spatial comparison for the zoomed-in region with enhanced diurnal warming is also shown in Fig. 10.

The dominant feature of the functional and observational variance differences is their difference in magnitude. The variances derived from the interpolating functions are smaller in magnitude than the sample variances of the observations over the same domains, but they reflect similar regional patterns. No strong coherent spatial patterns are exhibited in the difference images in either Fig. 8g–i or Fig. 10g–i.

Given the differences in magnitude between the variance estimates, it is important to consider how representative the values are of the actual physical variability we wish to quantify. An observational measure of spatial variability is dependent on the density of the underlying sampling. Gaps in the observations can miss sources of variability. When observations are more abundant, it is possible to get an increasingly better representation of the variability. Use of the interpolating function allows estimation in the limit of complete sampling at the given resolution, since they are continuous everywhere in the spatial domain for which they apply. This assumes, of course, that the density of observations was at least sufficient to derive an accurate interpolating function. Observational variability estimates are also critically dependent on the noise level of the measurements. The presence of artifacts in the satellite retrievals noted in Section 5 implies that the data might not be completely reliable for direct estimation of the spatial variability. The interpolating function, in contrast, can largely mitigate the impact of sensor noise and other artifacts. Because of the filtering of noise and some inherent smoothing of the observations in the fitting process, the variability derived from the interpolating function is expected to be lower in magnitude. Under conditions where the density of observations is reduced (but not too low) and/or where the observations have elevated noise, the interpolating function could provide an improved estimate of the true subpixel variability. At minimum, the interpolating function should provide a lower physical bound on the spatial variability at the scale of the underlying observations.

Based on this reasoning and the similar patterns in the derived variability, the variance derived from integration of the second moment of the interpolating functions appears to provide a powerful mechanism

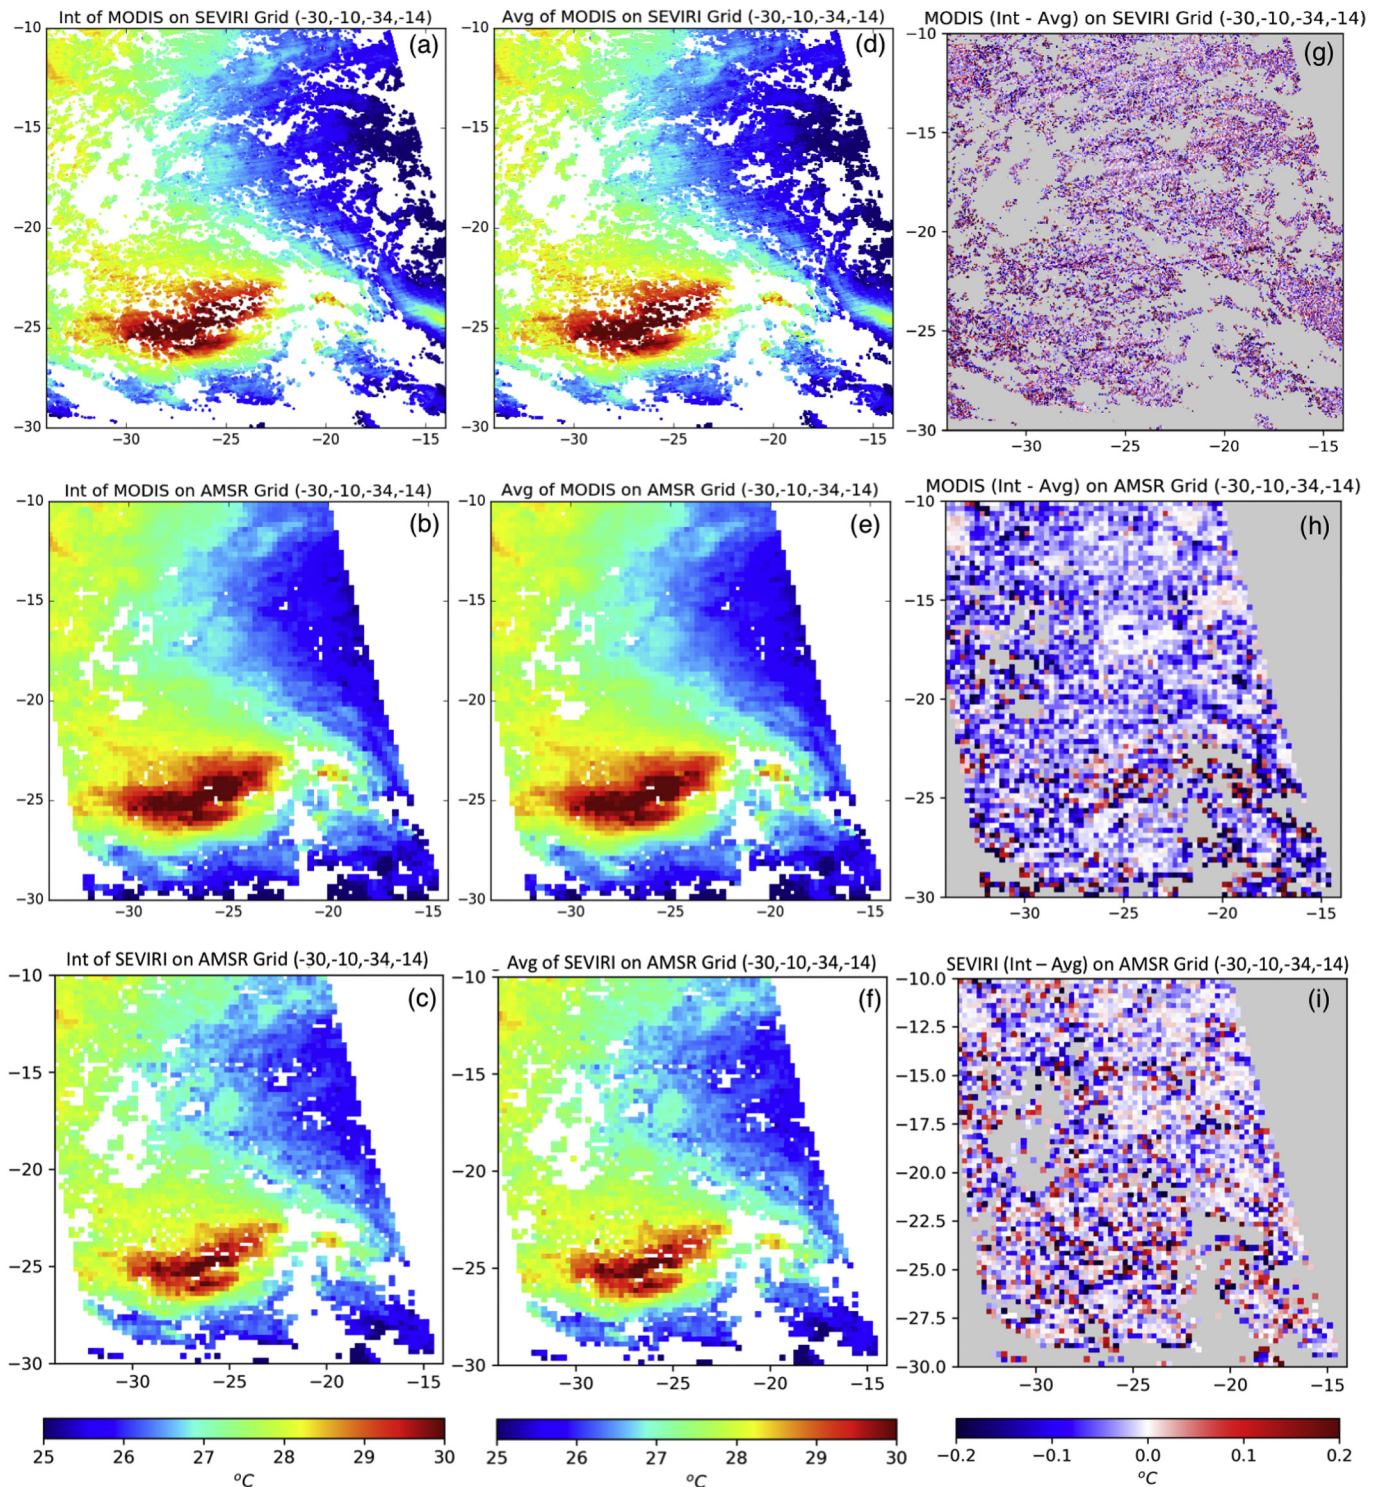


Fig. 7. As in Fig. 5, but for a localized region including the peak diurnal warming.

to help estimate the magnitude of, and regional differences in, the variability within the selected satellite footprints. Comparison of the quantitative variance estimates with grid resolution first provides insight into how the magnitude of the underlying physical variability changes with spatial scale. Comparing variance estimates from the MODIS interpolating function on the AMSR and SEVIRI grids shows increased variance over the larger 25-km AMSR cells (mean of 0.03 K^2 vs 0.005 K^2 for the full region). This is to be expected as larger regions generally encompass increased spatial variability. The variance computed from the SEVIRI interpolating function on the same AMSR grid

(mean of 0.01 K^2) is less than that from the MODIS interpolating function. This again makes sense as the higher resolution MODIS observations resolve more small-scale variability than does SEVIRI, though perhaps with increased noise. Lower limits on the derived variance, however, approach similar values in open ocean regions. The mode of the variance of the MODIS interpolating function evaluated on 25-km regions is 0.005 K^2 (Fig. 9) in agreement with the mean variance value evaluated over 5-km regions. This value corresponds to a standard deviation of 0.07 K , which is in excellent agreement with the spatial variability contribution to satellite SST retrieval accuracy found

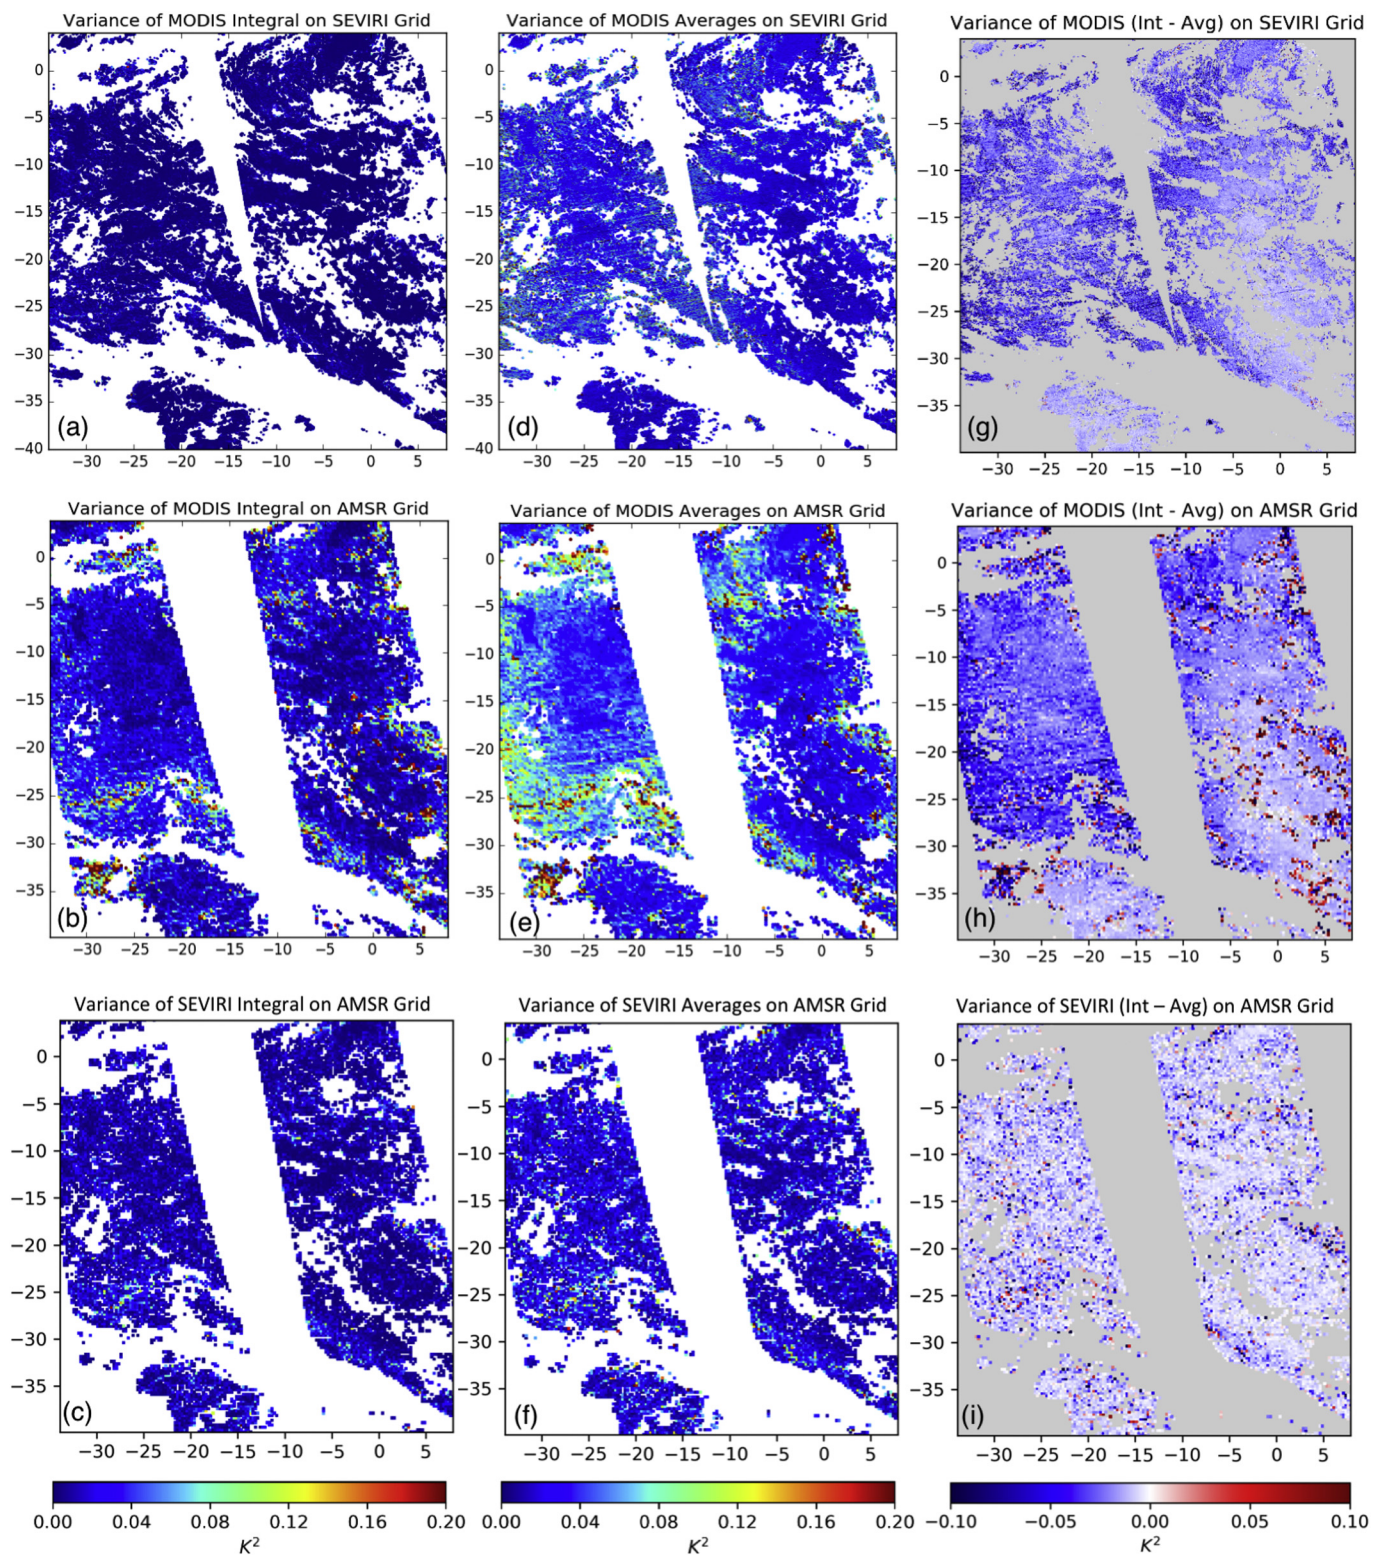


Fig. 8. Evaluation of the ability of the second moments of the interpolating functions to replicate the variance of the underlying observations on the selected grids of differing resolution. The top panels correspond to the MODIS resolution sampling on the SEVIRI grid, the middle panels to MODIS resolution sampling on the AMSR-E grid, and the bottom panels to SEVIRI resolution sampling on the AMSR-E grid. Panels a–c show the variance computed from the second moment of the interpolating functions over the grid cell, panels d–f show the variance of the available observations within the grid cell, and panels g–i show the corresponding differences between the functional and direct data approaches. Common color bars are shown at the bottom of each column.

by Castro et al. (2010). The smallest variance estimates from the direct retrievals also approach similar limits but show more variability across the satellite scans.

The variance estimates further permit closer examination of

regional variations in the spatial variability. For MODIS-scale 1-km observations aggregated on a 5-km resolution grid, the explicit variance is very homogenous over the entire domain, with a barely noticeable increase in the immediate vicinity of the peak diurnal warming




Fig. 9. Distribution of the variance estimates computed from the (solid lines) second moments of the interpolating functions and the (dotted lines) sample variance of the available observations for the results plotted in Figs. 8a–c and 8d–f, respectively. Results for the different sensor and grid combinations are plotted by color as shown in the legend.

(Fig. 10a). The spatial variability throughout is well captured by the mean variance value of 0.005 K^2 . On scales of 25 km, however, localized areas with increased variance up to 0.2 K^2 are visible, particularly near the diurnal warming. The peak values directly near the diurnal warming maximum obtained from the direct retrievals are quite similar to those from the interpolating function, further supporting this upper value. These results suggest that spatial representativeness errors on scales of 5 km can be reasonably assumed to be around 0.07 K over large expanses of the ocean, but that on scales of 25 km, localized variations should be considered and the errors can increase to near 0.4 K in regions with enhanced spatial variability. Overall, these results obtained using the new interpolating functions are broadly consistent with other localized estimates of spatial variability, but further enable a practical way to help quantify spatial variability and representativeness errors over much larger regions and at different scales.

7. Conclusions and future directions

A new powerful global interpolation technique yielding an interpolant in functional form has been developed with multiple potential applications to the generation and analysis of satellite-derived SST products. While directly relevant to the traditional application of gap filling for construction of blended gap-filled analyses, the technique has additional novel applications to the quality assessment of satellite SST retrievals and quantitative estimation of spatial variability on different scales. Use in addressing these problems could provide important guidance not readily available through other currently available methods.

The technique is based on the construction of a trigonometric interpolating function that fits the input satellite retrievals within a limited spatial wavenumber domain, providing adequate resolution in regions of larger data gaps and the highest possible resolution in regions with sufficient data density. This resulting interpolant provides a functional representation of the underlying physical SST field that can be evaluated anywhere within the domain over which it was derived. Importantly, moments of the interpolating function can be calculated accurately and efficiently, enabling estimation of the mean and variance of the field over desired sub-regions. In this application, our choice of functional form was trigonometric polynomials to assure speed of computation by using the FFT or the USFFT. Use of the technique was demonstrated by application to level 2 satellite SST retrievals from the MODIS and AMSR-E sensors and a level 3 gridded SEVIRI

product as well as a spatially complete composite SST product sampled with a realistic cloud mask.

The results illustrated how comparison of the functional form of the interpolant with the original retrievals upon which it was based could be employed as part of an operational processing scheme to help flag suspect retrievals for additional quality assessment. Large values of misfit highlight specific retrievals potentially inconsistent with smooth variations of the surrounding values. The technique was able to illustrate very small anomalies/artifacts such as MODIS sensor striping largely undetectable with other approaches.

The choice of the functional form strongly affects what patterns can be identified. The trigonometric polynomials employed in this application identified potential artifacts associated with linear scan striping but fit and reproduced small conical stripes associated with the scanning of AMSR-E. Alternate functional forms could be used to target other specific anomalies like conical striping. In general, there are many known ways in which additional constraints can be imposed to help filter out artifacts with specific characteristics and we plan to develop appropriate interpolants for such purpose. The overall framework is extremely versatile and powerful.

Direct comparisons between moments of the generated interpolating functions and the observations used in their derivation showed that the technique can be used to accurately represent spatial averages and quantify spatial variability in the underlying physical SST field. Integrals of the interpolating functions for MODIS, SEVIRI, and AMSR-E agreed closely with direct averages of the available retrievals over the same domains. This illustrates how the interpolating functions can be used to emulate retrieval of the SST field at different effective spatial resolutions. Moreover, the second moment of the interpolating functions was consistent with the variability of the available observations within grids of coarser resolution demonstrating that the functions can help provide quantitative estimates of, or bounds on, the spatial variability on different desired spatial scales. The spatial variance estimated from the interpolating functions was generally smaller than that of the direct observations due to observational noise and the smoothing nature of the functions, but the values represent a lower bound on the physical variability and could, in at least some cases, be more consistent with the actual variability on the scale of the sampling given the observational limitations. Comparison of the interpolated results with independent buoy SST measurements from the day assessed further supported the ability of the technique to accurately reproduce the underlying SST field. Broader tests over larger regions and additional days are, of course, desirable.

Quantitative estimation of the spatial variability of the SST on different scales and the associated representation error of point and finer scale measurements is particularly challenging with other traditional methods and much remains unknown about appropriate physical values. Application of the technique provided new insight into the spatial SST variability, at least within the limited region of this initial test. Within open ocean regions away from any frontal features, variability of 1-km-resolution observations on grids of both 5- and 25-km resolution was found to be $\sim 0.07 \text{ K}$ (as expressed by a standard deviation). In regions of sharper gradients such as associated with strong localized diurnal warming, the variability within 25-km-resolution grids increased to as much as 0.4 K for sampling at 1-km resolution. The variability of 1-km observations on a 25-km-resolution grid was about 2.4 times greater than that on a 5-km-resolution grid. Broader application of the technique globally could help better quantify true regional variations in the spatial variability, which would subsequently improve uncertainty estimates for existing satellite-based SST products.

The relative performance of the interpolation technique in the gap-filling problem as compared with other traditional methodologies was not explicitly examined in this paper. Detailed comparisons are required to accomplish this in a meaningful way. Additional activities related to this topic are planned and the utility of the technique for gap-filling in construction of level 4 SST analyses will be explored.

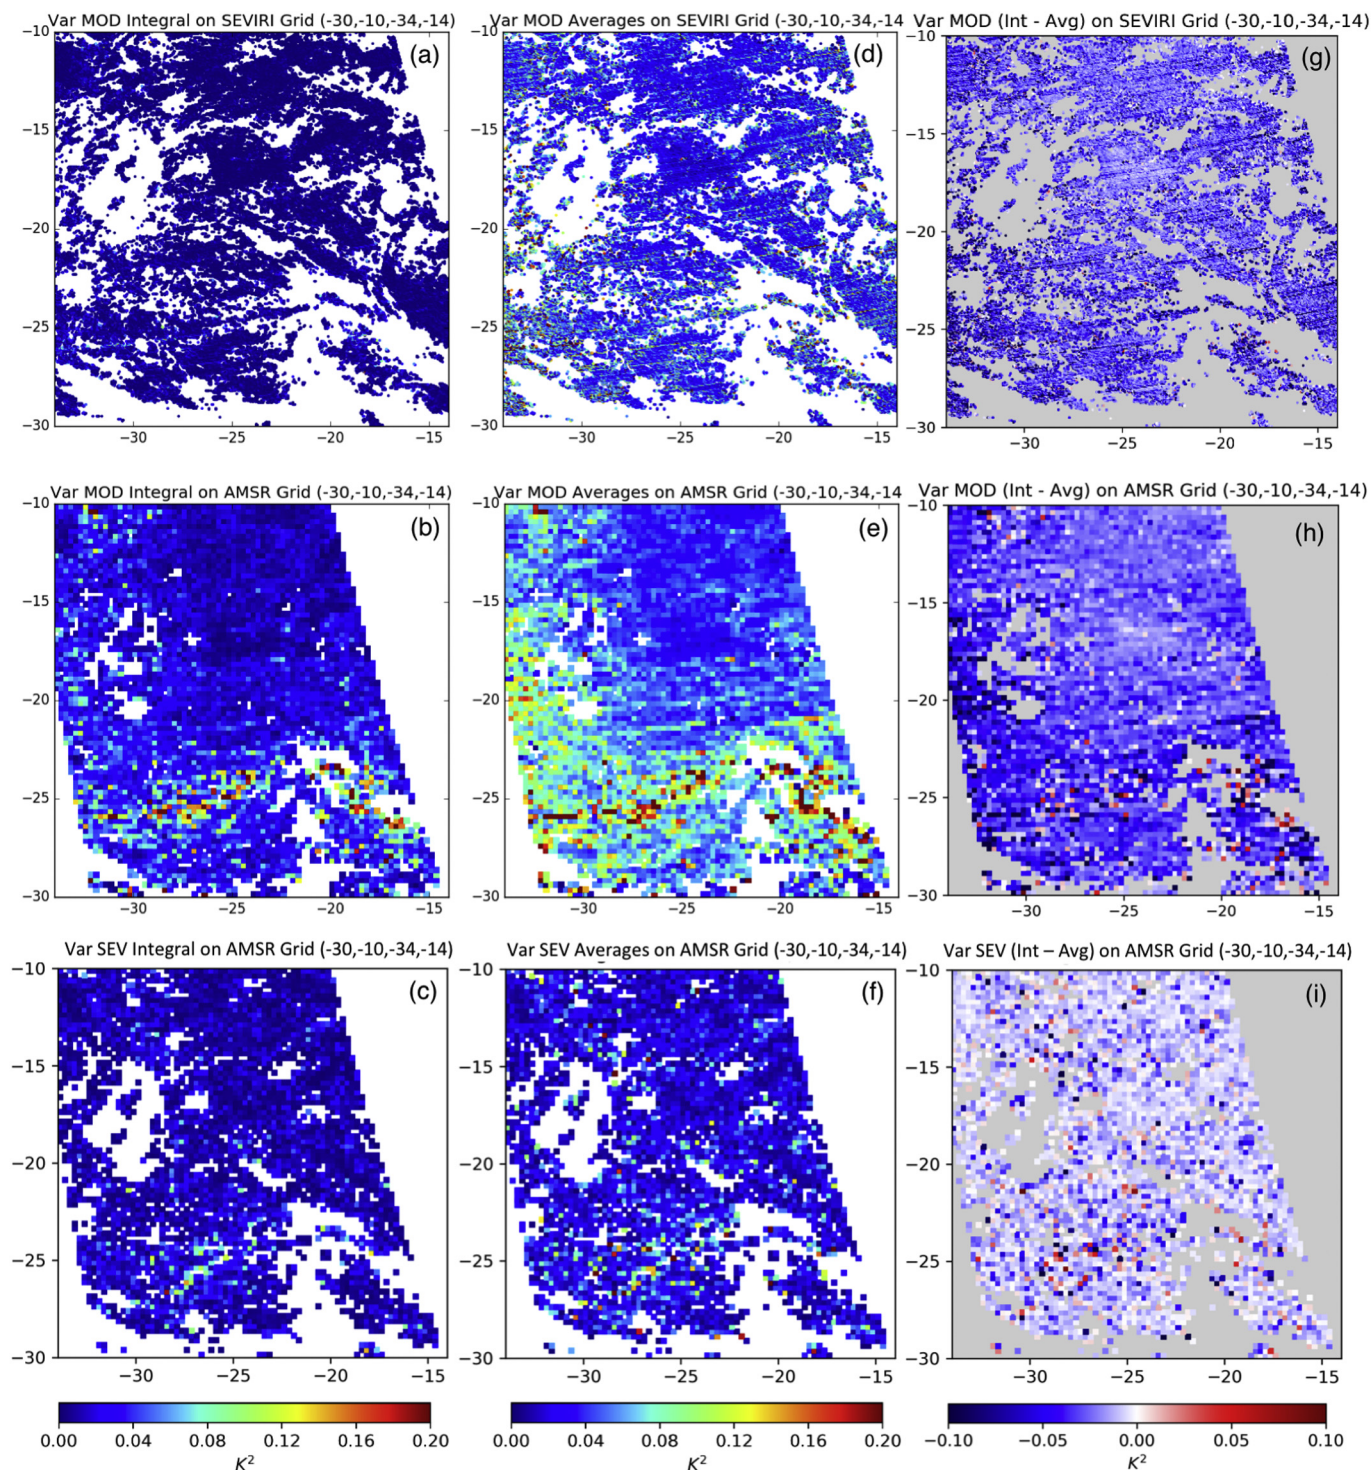


Fig. 10. As in Fig. 8, but for the localized region including the peak diurnal warming.

While applied here primarily to single-sensor data, the technique can also be applied with inputs from multiple sensors and observations collected across different measurement times to handle larger data gaps as previewed in Section 4. Successful application of the interpolation methodology is clearly dependent upon having sufficient observational data density and quality to accurately derive and constrain the interpolating functions. As expected, the quality of the interpolating functions is only as good as the data from which they are derived. The positive outcome of this investigation was, in part, enabled by the low amount of cloud cover in the scenes analyzed. Application of the

technique, based solely on single-sensor data, will be challenging in regions of extensive and persistent cloudiness where the infrared coverage can be limited. Additionally, the finest effective spatial resolution of the interpolating functions is naturally limited by the resolution of the data used in their derivation. Preliminary tests indicate that the technique can be easily expanded to blend multi-sensor data to help ameliorate any limitations associated with data density, which will further enhance the powerful new capability.

Acknowledgments

Participation of SLC and a portion of the time of LM was supported by a grant from the National Ocean Partnership Project for the Multi-sensor Improved Sea Surface Temperature (MISST) for the Integrated Ocean Observing System (IOOS) project (via NOAA award NA11NOS0120167). GAW was supported by NOAA in collaboration with the MISST for IOOS project. The satellite data sets utilized were all obtained from the providers as directly cited in the text. The efforts to ensure the open access to these data are greatly appreciated. The constructive comments of three anonymous reviewers are also acknowledged.

Declarations of interest

None.

References

- Alvera-Azcárate, A., Barth, A., Beckers, J.-M., Weisberg, R.H., 2007. Multivariate reconstruction of missing data in sea surface temperature, chlorophyll, and wind satellite fields. *J. Geophys. Res.* 112, C03008. <https://doi.org/10.1029/2006JC003660>.
- Beylkin, G., 1995. On the fast Fourier transform of functions with singularities. *Appl. Comput. Harmon. Anal.* 2, 363–381.
- Bouali, M., Ignatov, A., 2014. Adaptive reduction of striping for improved sea surface temperature imagery from Suomi National Polar-Orbiting Partnership (S-NPP) Visible Infrared Imaging Radiometer Suite (VIIRS). *J. Atmos. Ocean. Technol.* 31, 150–163.
- Castro, S.L., Wick, G.A., Minnett, P.J., Jessup, A.T., Emery, W.J., 2010. The impact of measurement uncertainty and spatial variability on the accuracy of skin and subsurface regression-based sea surface temperature algorithms. *Remote Sens. Environ.* 114, 2666–2678. <https://doi.org/10.1016/j.rse.2010.06.003>.
- Castro, S.L., Wick, G.A., Buck, J.J.H., 2014. Comparison of diurnal warming estimates from unpumped Argo data and SEVIRI satellite observations. *Remote Sens. Environ.* 140, 789–799.
- Castro, S.L., Wick, G.A., Steele, M., 2016. Validation of satellite sea surface temperature analyses in the Beaufort Sea using UpTempO buoys. *Remote Sens. Environ.* 187, 458–475.
- Castro, S.L., Emery, W.J., Wick, G.A., Tandy Jr., W., 2017. Submesoscale sea surface temperature variability from UAV and satellite measurements. *Remote Sens.* 9, 1089. <https://doi.org/10.3390/rs9111089>.
- Chin, T.M., Vazquez-Cuervo, J., Armstrong, E.M., 2017. A multi-scale high-resolution analysis of global sea surface temperature. *Remote Sens. Environ.* 200, 154–169. <https://doi.org/10.1016/j.rse.2017.07.029>.
- Cornillon, P., Castro, S., Gentemann, C., Jessup, A., et al., 2010. SST Error Budget – White Paper. Available online: <http://works.bepress.com/peter-cornillon/1/>.
- Cressie, N.A.C., 1993. *Statistics for Spatial Data*. John Wiley & Sons, Inc., New York.
- Dash, P., Ignatov, A., Martin, M., Donlon, C., Brasnett, B., Reynolds, R., et al., 2012. Group for high resolution sea surface temperature (GHRSSST) analysis fields inter-comparisons. Part 2: near-real time web-based level 4 SST quality monitor (L4-SQUAM). *Deep-Sea Res.* II 77–80, 31–43. <https://doi.org/10.1016/j.dsr2.2012.04.002>.
- Dutt, A., Rokhlin, V., 1993. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Comput.* 14, 1368–1393.
- Esaias, W.E., Abbott, M.R., Barton, I., Brown, O.B., Campbell, J.W., Carder, K.L., Clark, D.K., Evans, R.H., Hoge, F.E., Gordon, H.R., Balch, W.M., Letelier, R., Minnett, P.J., 1998. An overview of MODIS capabilities for ocean science observations. *IEEE Trans. Geosci. Remote Sens.* 36, 1250–1265.
- EUMETSAT, 2011. *Geostationary Sea Surface Temperature Product User Manual*. SAF/OSI/CDOP/M-F/TEC/MA/181. Available online: <http://www.osi-saf.org>.
- Gumley, L., 2002. Proc. MODIS workshop. In: URL: Western Australian Satellite Technology and Applications Consortium.
- Hofstra, N., Haylock, M., New, M., Jones, P., Frei, C., 2008. Comparison of six methods for the interpolation of daily European climate data. *J. Geophys. Res.* 113, D21110. <https://doi.org/10.1029/2008JD010100>.
- Kent, E.C., Challenor, P.G., Taylor, P.K., 1999. A statistical determination of the random observational errors present in voluntary observing ships meteorological reports. *J. Atmos. Ocean. Technol.* 16, 905–914.
- Martin, M., Dash, P., Ignatov, A., Banzon, V., Beggs, H., Brasnett, B., et al., 2012. Group for high resolution sea surface temperature (GHRSSST) analysis fields inter-comparisons. Part 1: a GHRSSST multi-product ensemble (GMPE). *Deep-Sea Res.* II 77–80, 21–30. <https://doi.org/10.1016/j.dsr2.2012.04.013>.
- Minnett, P.J., 1991. Consequences of sea surface temperature variability on the validation and applications of satellite measurements. *J. Geophys. Res.* 96, 18,475–18,489.
- Murphy, R.E., Ardanuy, P., Deluccia, F.J., Clement, J.E., Schueler, C.F., 2006. The visible infrared imaging radiometer suite. In: Qu, J.J., Gao, W., Kafatos, M., Murphy, R.E., Salomonson, V.V. (Eds.), *Earth Science Satellite Remote Sensing: Vol. 1: Science and Instruments*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 199–223.
- Reynolds, R.W., Chelton, D.B., 2010. Comparisons of daily sea surface temperature analyses for 2007–08. *J. Clim.* 23, 3545–3563. <https://doi.org/10.1175/2010JCLI3294.1>.
- Reynolds, R.W., Smith, T.M., 1994. Improved global sea surface temperature analyses using optimum interpolation. *J. Clim.* 7, 929–948.
- Reynolds, R.W., Smith, T.M., Liu, C., Chelton, D.B., Casey, K.C., Schlax, M.G., 2007. Daily high-resolution blended analyses for sea surface temperature. *J. Clim.* 23, 5473–5496.
- Vinogradova, N.T., Ponte, R.M., 2013. Small-scale variability in sea surface salinity and implications for satellite-derived measurements. *J. Atmos. Ocean. Technol.* 30, 2689–2914.
- Wentz, F.J., Meissner, T., 2004. *AMSR-E/Aqua L2B Global Swath Ocean Products Derived From Wentz Algorithm*. Version 2.
- Xu, F., Ignatov, A., 2014. In situ SST quality monitor (iQuam). *J. Atmos. Ocean. Technol.* 31, 164–180. <https://doi.org/10.1175/JTECH-D-13-00121.1>.