

BUILDING AN ACCESSIBLE, INTEGRATED EARTH OBSERVING AND INFORMATION SYSTEM

The International Satellite Cloud Climatology Project as a Pathfinder

WILLIAM B. ROSSOW AND JOHN J. BATES

Lessons learned in building the ISCCP observation–processing–archival system and making it operational could serve as a pathfinder for building a global observing and information system.

The recent release of a National Academies of Science, Engineering, and Medicine (NASEM) decadal survey of Earth science data needs (NASEM 2018) is only the latest of many calls for a global system to produce sustained, reliable observational data and information about Earth. Like previous such calls over the last 40+ years (Perry 1975; Zillman 1980; Global Climate Observing System 1995; Christian 2005), this survey mentions research-to-operations (R2O) actions that are needed to develop such a system but focuses mainly on making and collecting measurements with little

discussion of the processing system and the integrated team of talented scientists needed to turn raw observations into usable information or the archival system needed to make reliable information readily and widely accessible. Satellite observations by their nature provide more extensive coverage of Earth with finer space–time resolution than can be obtained from a collection of surface-based measurements, so a system to provide global information about Earth will necessarily be based on the constellation of satellites, especially the weather satellites operated by many countries, supplemented by more detailed local measurements. However, the current heterogeneity of the existing global collection of measuring assets, satellite and surface based, is a major obstacle to creating a truly integrated, globally uniform information system. Hence, most of the collected raw data volume is not used. The issues related to preservation of the observational data are discussed in Knapp et al. (2007). Here we discuss an example of addressing the problems in producing globally uniform information from such observations: the creation in 1982–83 of the data collection, processing, and archival system for the International Satellite Cloud Climatology Project (ISCCP; Schiffer and Rossow 1983, 1985; Rossow

AFFILIATIONS: ROSSOW—Franklin, New York; BATES—John Bates Consulting, Arden, North Carolina

CORRESPONDING AUTHOR: William B. Rossow, wbrossow@gmail.com

The abstract for this article can be found in this issue, following the table of contents.

DOI:10.1175/BAMS-D-19-0060.1

In final form 10 July 2019

©2019 American Meteorological Society

For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](#).

and Schiffer 1991, 1999). Extending the ISCCP data record beyond that produced by the original research project motivated transitioning the processing to a fully operational environment.

Observations of clouds, their global and regional characteristics, and variability on time scales from minutes to decadal, have been revolutionized by space observations over the past four decades. Transforming these observations into a uniform, global set of physical information about clouds that is readily accessible was and is challenging for several reasons. First, there was and continues to be wide international dispersal and variety of measurement capabilities and of information development responsibilities embedded in different kinds of institutions with different characteristics that affect the nature of the observations and the derivation of information. Second, there are differences between research data analyses and operational data processing and approaches to merging datasets from different providers: contrasting characteristics are experimental (short-term temporary funding) versus routine (long-term sustained funding), direct retrieval (more physically based) versus inferential (more statistically based) representations, and univariate versus multivariate analyses. Third, the end users of research and operational data products have different data and information needs with different quality requirements affecting the data product designs. In particular, using information to address policy questions or for regulatory applications raises questions about whether the data processing system and its products need independent certification and documentation of the whole chain of procedures from measurement through processing to analysis. All of these needs are currently met in part by the separate and disparate activities of an international set of institutions, producing different data products with different contents and formats that are derived by different procedures albeit from very similar measurements. Finally, there is the challenge of providing ready access to information at different, but physically self-consistent, levels of analysis, from raw measurements through various derived products to forecasts. Currently the raw measurements and analysis products are stored in a large number of locations with very limited access provided in some cases. Improved documentation and access to either very detailed regional or global climate information (long time records) is needed both to conduct scientific investigations and, increasingly, to fulfill legal requirements. In this short commentary, we reflect on the experiences and lessons learned in building the ISCCP observation–processing–archival

system and in carrying out the ISCCP R2O process to serve as a pathfinder for building a global observing and information system.

Although satellite imagery (usually in photographic format) had been used to identify weather systems and estimate cloud cover and cloud-top heights from the 1960s through the 1970s, it was not until the early 1980s that satellite agencies provided routine access to digital formats of the imaging data. ISCCP was initiated in 1982, originally motivated by research to study the roles of clouds in climate. This project required the collection of satellite imaging measurements to determine the variations of the physical properties of clouds covering spatial scales from local (about 10–30 km) to global and time scales from 3 h to multiyear. This task necessitated collecting the common “window” infrared and visible image data in multiple formats from multiple weather satellites (up to seven) operated by multiple agencies and using multiple ancillary datasets (five additional inputs from other satellite instruments) to isolate clouds from the other effects on the measurements and to support the needed diagnostic studies using the products (e.g., determining the radiative effects of clouds). For ISCCP these objectives were accomplished by adding the activities of research-oriented data processing centers for cross calibrating and analyzing the satellite radiances to the operational data collection activities of the separate satellite agencies (Schiffer and Rossow 1983, 1985). ISCCP data products (Schiffer and Rossow 1985; Rossow and Schiffer 1991, 1999) have been used to provide information about and monitor global/regional environmental conditions and changes on a range of time scales for climate research, monitoring weather and surface conditions—agriculture, air pollution, water quality—decision-making for infrastructure, transportation, and other business. ISCCP is widely cited in more than 12,000 papers.

The data analysis part of ISCCP was funded by NASA in support of research investigations of Earth’s radiation budget and climate, but initial quality checking, limited calibration, and subsampling of the imaging data and provision of them to NASA, along with the ancillary inputs, were performed by an international group of operational weather satellite agencies. Carrying out continuing and routine measurements and analysis is not a NASA mission so continuing the processing to extend the data record motivated the transition of the ISCCP processing to a more formal cooperative activity of the operational satellite agencies. Commitments to such an international observing system include both the production of the global

data products and specialized regional processing to produce products tailored to local needs. Thus, plans were made to transition the global processing to NOAA and to affirm the international commitments for continued production of regional products. Planning of the ISCCP transition began prior to the initiation of the NOAA Climate Data Records (CDR) program (Bates et al. 2015) in 2007, but the activity was subsequently merged into that program (see sidebar “Revisiting elements of a successful satellite CDR generation program”). Preparatory work started a few years later. Within the CDR program framework, there are two major transition status milestones: initial operational processing, into which over 40 products have been transitioned, and full operational processing, for which ISCCP is serving as a pathfinder. The ISCCP and other CDR teams involved in the transitions wrestled with many issues, both scientific and information-technology-related, which led to specific metrics for success (see sidebar “Elements required for a successful transition to full operational capability”). Since building a global observing system requires procedures for moving mature scientific analyses to operations (R2O), future systems for processing multiple input data streams throughout the geophysical sciences may benefit from the lessons learned in the ISCCP development and transition.

The design of the software processing system must carefully consider from the outset the practices that lead to production of a quality product including 1) a data flow that is as simple and logical as possible with a flexible modular structure; 2) clean, robust, and efficient coding following recognized coding standards; 3) quality monitoring at several stages of the processing; and 4) thorough documentation. Optimizing the flexibility of these aspects of the software, even in the initial design, can later lead to easier adaptation onto different computer systems and a significant improvement in processing speed, which is crucial for a sustained processing system. Here are some of the features of the ISCCP system that proved very useful.

- 1) To provide maximum flexibility in a processing system with multiple input data sources (two-wavelength images from up to seven satellites and

REVISITING ELEMENTS OF A SUCCESSFUL SATELLITE CDR GENERATION PROGRAM

The 2004 National Research Council report *Climate Data Records from Environmental Satellites* (National Research Council 2004) identified elements of a successful CDR program in the categories of organization, generation and stewardship, and sustaining the program. The organization elements recommended a separate advisory council for CDRs. NOAA chose to implement the program, which began in 2007, using existing NOAA climate program review boards already in place. Teams were competitively selected to provide fundamental CDRs (FCDRs; i.e., intercalibrated radiance data) and thematic CDRs (TCDRs; i.e., retrieved geophysical parameters of the essential climate variables). ISCCP provides both FCDRs and TCDCRs and the National Research Council (NRC) recommendations for generation and stewardship were followed. As part of the development of a new version of the data products and in preparation for the transition, the ISCCP processing code was reengineered to modernize and make it more flexible, including an option for near-real-time processing. Processing a standard dataset on different hardware to get similar results within machine error was sometimes challenging—the incompatibilities between software and computer operating systems can be very subtle. The program uses the term operational CDR to refer to a dataset that is systematically and routinely produced (typically updated quarterly) and that meets or exceeds maturity matrix level 5 (Bates and Privette 2012; see sidebar “Elements required for successful transition to full operational capability”). The transition process for ISCCP took about five years. The NRC report specifically identified sustaining the CDR program as a key for success. Indeed, this has been the most difficult challenge, in part because U.S. climate funding exists across multiple agencies and is not fungible. NOAA’s plans to establish a climate service line office were never approved by Congress making it difficult to continue the CDR program expansion. The NOAA CDR program was established as a companion (with a separate budget line) to the NOAA climate sensor program, established after the restructuring of the NPOESS program, for acquisition of climate sensors for Earth radiation budget and solar irradiance measurements. Following Superstorm Sandy and facing continued delays in the NOAA Joint Polar Satellite Program, Congress and the Administration in 2013 gave the lead in climate sensor continuity back to NASA (NASEM 2015). Although the CDR program was separately funded, this move further eroded the motivation for full operational transition of CDRs to NOAA. Most CDRs in the NOAA program have only achieved initial operations (for details see www.ncdc.noaa.gov/cdr). The transition of ISCCP to full operational capability was already well underway, so the decision was made to continue it to completion. The completion of the ISCCP transition is important as it means an essential climate variable used not only for monitoring, but also for applications as varied as hurricane intensity climatologies and solar power assessment, will continue to be available. Unfortunately, the vagaries of funding of the U.S. climate research and operations have limited the ability for more essential climate variables to be consistently produced.

TABLE I. Main recommendations.

- 1) Funding during transition period for both research and operations teams
- 2) Computer engineering and scientific expertise in both teams
- 3) Operations team should conduct research with data products for improved QC
- 4) Modular processing code (separate processing of each input source, restart, reprocessing)
- 5) Maximize information on processing logic/outcomes in level 2 output (input in output, labels)
- 6) Multistage QC (input, intermediate stages, output for each module)
- 7) Data QC performed by both teams during transition
- 8) Cowritten operations guide, scientific basis document, and user's guide

ELEMENTS REQUIRED FOR SUCCESSFUL TRANSITION TO FULL OPERATIONAL CAPABILITY

- 1) **Reproducibility:** The CDR's source code will have passed any required security review and will have been successfully made operational with any platform dependencies identified. The data generated will be stable and within machine rounding errors of predetermined test datasets for the CDR.
- 2) **Verification:** Source code and documentation are publicly available. Any discrepancies from original results will have been quantified and explained. Operational monitoring of CDR data will be in place. CDR data will be periodically reviewed against predetermined test datasets and may be reprocessed when improved algorithms or newly recalibrated data become available.
- 3) **Extensibility:** The CDR spans multiple decades and is capable of being extended in time to generate a stable climate time series. CDRs will be routinely extended as new data become available from existing or new sensors, providing long-term product consistency. Forward processing of data will be performed as near to real time as possible, allowing time for assessment and correction of any problems with the data. As new algorithms become available, new thematic climate data records may be created with different theoretical bases or the number of CDRs may be increased to encompass additional essential climate variables.
- 4) **Preservation:** The CDR package will be preserved in accordance with NOAA standards and includes data (output data, ancillary data, and any input data not currently archived), metadata, production software source code, software to read the data, documentation, calibration/validation information and data, and quality assurance information. The climate algorithm theoretical basis document and operational algorithm description of the CDRs are preserved to provide more complete provenance traceability of the scientific basis of the CDR. The CDR will meet NOAA archive standards for both data and metadata format.
- 5) **Accessibility:** Access to and distribution of CDRs to the user community will be assured and facilitated. CDRs will be available in as timely a manner as possible with user requests typically fulfilled within 24 h. Access will be full and open with provision for users to provide feedback. Data will be available in widely utilized formats, and proprietary environments for data usage will be avoided. Appropriate data manipulation and visualization tools may be provided.

Documentation of all these elements can be found online at www.ncdc.noaa.gov/cdr/atmospheric/cloud-properties-iscpp and <https://doi.org/10.7289/V5QZ281S>. These elements must meet or exceed CDR maturity matrix level 5 for full operational status.

five ancillary datasets), the processing structure was composed of separate modules for each input with merger into global products delayed as far down stream as possible. One way to ensure clear logic of the data flow is to embed the structure of a flowchart in the software as comments: the headings and subroutine blocks should be the same. Each of these modules can be run alone. The output from each is saved until the production has completed with quality checks performed at each stage of production. This approach allows for easier identification of the location of any problems and a restart of the processing without repeating all the other parts. The structure naturally produces a hierarchical data product with components that present the results at different levels of detail for different uses. In the case of the new version of ISCCP, in addition to archiving all of the input datasets, six products are saved. The most detailed output stage (level 2, image pixel level) is useful for development of improved retrieval methods

or, in combination with other measurements, the investigation of physical processes. Other kinds of research that employ statistics more heavily can use reduced-volume, mapped products that describe variations at larger space–time scales while providing summary information about the smaller-scale variations beyond simple averages. Two particularly important types of summaries are distributions, both univariate and multivariate, and conditional (situation dependent) information. Although the typical scientific proposal cycle is three years, the processing system design should still allow adaptation to new computer technology, different input data sources, and different product needs through a modular, multistage design. For instance, the ISCCP code, which was run on mainframe-type computers initially, was already multithreaded (each satellite month processed separately), making it easier to move to and manage on multiple workstations and then into a fully multithreaded environment with a large number of processors. This approach does not avoid incompatibilities between software and computer operating systems but does make it easier to diagnose and solve them. This design also made it easier to change input sources.

- 2) A particularly important feature of robust processing code is that it have a definite outcome for any possible input or result generated internally—no input or output is to be discarded but each output should include a quality label. This feature is aided by constraining the inputs allowed (labeled, not discarded)—a function of monitoring the quality at the input stage—and by constraining the outputs allowed (labeled, not discarded) by monitoring the quality at the output stage. Regarding the latter, the precision of outputs should reflect the precision of the inputs. Labeling of the output (see sidebar “Some details on output contents”) is efficient because it also allows for easy monitoring of the character of the processing: any change in the frequency distribution of the labels can indicate a problem or an unanticipated change that needs to be investigated. Another feature of level 2 products that is very useful for development of improved retrievals and evaluation of the accuracy of the retrievals is to include the input measurements in the output from the retrievals. This allows for easier post facto examination of the relation of

SOME DETAILS ON OUTPUT CONTENTS

Two especially useful features of level 2 output from an analysis that can be used for developing a product, monitoring its quality or for research investigations are to report the input pixel-level radiances in the pixel-level retrieval output and to provide labels that characterize each pixel. Reporting the input in the output allows for much easier post facto investigations of the relationship between the radiance measurements and the physical quantities retrieved from them. The labels are numerical codes that identify the pixel-level scene type (e.g., land or water, snow–ice covered, or snow–ice free) that are useful for sorting the results to investigate the scene dependence of the results. These labels can also indicate the pathway followed by each pixel through the processing code logic or the ending conditions of each subroutine in the processing chain. All of these types of labels are especially useful for monitoring large-volume processing because the distribution of label values should be stable across large spatial domains and over large time periods (e.g., one month)—any significant deviations may indicate unexpected changes in the input data.

input measurements to retrieved quantities, especially if additional labels are passed to the output that track the pathway of the input through the logic of the processing code, indicate the outcomes of various subroutines or tests, and also characterize the situation of each image pixel (land or water, high or low topographic height, snow–ice cover or snow–ice free, etc.). These labels can be part of the metadata for the data product and serve as quality indicators for each image pixel. The statistical distribution of these labels should also be kept for standard test cases to provide a baseline for monitoring the quality of the whole record.

- 3) The R2O transition and subsequent operational processing need procedures for continuous and diligent quality control (QC) of all aspects of processing from end to end, not just in the final products. Since the processing system is, at the transition to operations, very mature, problems noticed in the output are almost always caused by problems in the input, but it is still possible that a problem can arise because of situations that are not expected in the processing code (especially those produced by software and computer system incompatibilities). In fact, code testing should deliberately try to break the system by submitting extreme versions of the input data to see what happens. The ISCCP quality checking is done for each module at several stages of the processing, including the input data and the final

output products, but also at key stages within the processing stream to allow for diagnosis of any problems that are noticed. These quality-check results are also passed to the final data products as labels (metadata). Since the data volume being processed is very large (on the order of terabytes), outputs of the quality indicators are necessarily summaries but they can be designed to exhibit variations only when values or statistics change—well-behaved quantities and statistics are displayed as “boring” plots (e.g., straight lines) that exhibit deviations or structure only when unexpected values occur. This feature allows both for automated flagging of results and for operators to quickly visually scan a lot of detailed quality summaries because any problems are indicated by obvious deviations from the straight-line records. Such QC procedures and statistical plots of the results are also a good training tool to familiarize the operations team with the behavior of the processing system.

- 4) Research processing code often starts out as “dirty and quick” but over time can become “legacy” code that is increasingly difficult to maintain and debug, especially as the documentation is often incomplete. One useful aid to documentation is establishing configuration control by archiving all versions of the code that produce released versions of the products and maintaining a standard set of test cases with known outcomes and label statistics that can be run for each new version of the code. In the R2O situation, documentation has to provide more information than just the contents of the data products and the theory behind how they are produced, now usually contained in so-called algorithm theoretical basis documents (ATBDs). The transition process is facilitated by the writing of an operations guide but upon release of the data products, a user’s guide provides a useful complement to the ATBD. Even though the first versions of all these documents should be written by the research team and delivered with the code to the operations team, experience during the transition leads to revisions for clarity and accuracy by the operations team. The final versions of these documents thus result from the joint effort of the research and operations teams.

The modular processing stream design with quality checking and restart option at multiple stages provides the capability for reprocessing, which is a must for CDRs: this allows for redoing the processing

when problems have been identified. If a significant change in the input data or analysis procedure is introduced, then reprocessing the whole record, as was done for ISCCP twice, becomes necessary to maintain record uniformity. This capability requires constrained data volume especially because the previous versions of the products also have to be archived to document changes. Such a capability is very important for CDRs because some types of spurious variations in the data record, for example, caused by errors in radiance calibration or changes in the character of the input data, are not always recognized until the whole (or a large part of the) climatological record is examined. This is particularly true for identifying the effects of small instrument or orbit drifts. In other words, detecting flaws in the record is a matter of perspective. However, care must also be exercised to separate real, natural variability from instrument-processing-induced variability—the perspective provided by looking at longer records helps with this.

The R2O process in the ISCCP case was greatly aided by the fact that the research team (original developers, about 2 full time equivalents for R2O) and the operations team (about 2.5 full time equivalents) were both funded during the transition period. In addition to the scientific knowledge guiding the processing development, the research team also has to have sufficient computer expertise and programming and systems-level knowledge to prepare the code for transfer to other computer systems. In the case of ISCCP, such expertise was added to the research team for the transition. In addition to the requisite skill set to operate such a processing system, the operations team has to have both computer engineering expertise and scientific background to interpret the multistage QC results and to investigate and solve problems beyond simply running the production code. Programming experience is needed to trace back any feature in the output to its source and a scientific background is needed to evaluate whether any changes in the results are realistic (on Earth) or unrealistic (in the satellite or processing code). Monitoring of product quality is greatly aided if the processing team actually conducts research using the data, an activity begun in collaboration with the research team during the transition. Given the continually evolving characteristics of the satellite observing system and the changing state of information (computer) technology, changes to the input data and processing system (and reprocessing) will be required to extend the data record while maintaining its uniformity. This requires that

the operations team also has expertise in satellite operations and instrumentation. The main key to a successful R2O is good interpersonal interactions between the development and operations teams. This means extensive interaction through in-person and remote meetings where exchanges lead to deeper understanding of the processing system design and attributes. Rather than simply handing over the processing code, the goal is to transfer the experiences and knowledge of the original development team to the operating team, most particularly including the rationale behind the processing system and its design, the statistical and scientific interpretation of results, and the meaning of the quality control metrics. All this knowledge needs to be communicated and understood: it also has to be captured as completely as possible in the operations guide, for internal use by the operations team, and a user's guide, for external product users. Such documentation is a necessary feature for a long-term activity because the personnel on the operating team will also change with time.

Deciding when a research processing system has reached a level of maturity that conducting an R2O transition is appropriate is related to the characteristics that the system must have for a successful transition (see sidebar "Elements required for successful transition to full operational capability"). Discussions of what attributes the processing system and its products should exhibit to be mature, including considerations of the ISCCP experiences, have led to definitions of a maturity matrix (Bates and Privette 2012) that can be used to evaluate the readiness of a research processing system for being moved to operations. A maturity ranking of at least 3, in categories such as software readiness, documentation, validation, access, etc., is a good foundation for transition to operations, but completion of the process should raise the rank to at least 5. Of course, the decision for transition to operations also has to have a scientific and practical aspect to judge the value of the data products, whether to continue producing them (lengthening the record) and the cost of producing them. Regarding the latter issue, maintaining the homogeneous character of the data products usually implies over time that the cost of processing the same volume of data decreases: when ISCCP began processing in the 1980s, a small mainframe computer was used but now routine processing can be done on a small workstation (personnel required is 1–2 full time equivalents).

The success of ISCCP and its R2O is due to its design and the people involved at many satellite

agencies internationally over the past nearly 40 years. Its future success will depend on concerted governmental actions and cooperation at an international level to continue the collection and analysis of the data as part of a truly global observing system. The next generation of operational satellites is online with very large increases in data volume. The design of an operational processing system, along the lines of ISCCP, can prevent being overwhelmed by "big data" and can keep the additional analysis-to-information costs small relative to the cost of making and collecting the measurements. The key features of a practical Earth observing and information system are that it is "good enough"—it uses enough data rather than the maximum amount—and that it has the ability to grow by adding new capabilities developed from research so that emerging interdisciplinary geophysical questions can be addressed by a wide variety of scientists while requiring only modest computational resources.

Despite the rapidly growing data volumes, data storage technology now makes possible the archival of the complete level 1 data. This capability then makes it possible for later developments or upgrades of the processing and analysis procedures to be applied retrospectively to the whole record. The ISCCP products are now in their third version based on the archival of the sampled level 1 datasets: at each stage the whole record up to current time was processed to produce a homogeneous version.

Although there are now other research and operational satellite cloud datasets from several sources, only the ISCCP products provide globally uniform information with adequate diurnal time resolution from these same sources. Each of the satellite operators (experimental and operational) has their own objectives and requirements that can lead to similar information products, but the collection of these individual products is heterogeneous in properties and information content. A feasible and affordable approach to achieve global homogeneity of each item of information is to add one or more processing centers to the current satellite operating system, like ISCCP, to produce, in addition to the agency-specific products, a globally uniform version of the information that is based on the collective experience in producing separate products. Such a system would be less specialized than the individual agency products but would serve the needs for global and long-term uniform information. Such a background product would also provide a reference against which the other more specialized products can be compared. Such added centers would develop

new products as research projects and then, when mature enough (as judged by a maturity matrix for example), transition their processing to an operational mode for sustained information production. Sustained production of scientific information must also increasingly deal with legal requirements, such as the Data Quality Act (Section 515 of the Consolidated Appropriations Act, 2001—Pub. L. 106–554) and the Open, Public, Electronic, and Necessary (OPEN) Government Data Act (Foundations for Evidence-Based Policymaking Act—H.R. 4174), and maturity metrics assure compliance with these laws. The approach described above would serve to document the whole processing chain from observations to information.

ISCCP is now fully operational, producing a new version of the data products (at www.ncdc.noaa.gov/cdr/atmospheric/cloud-properties-isscp, <https://doi.org/10.7289/V5QZ281S>; Rossow et al. 2016), covering July 1983 through June 2017, and extending the record (Young et al. 2018). To build up the global observing system to provide uniform information, the research and operational satellite agencies could follow the ISCCP path for other information products by specifically funding development of global processing centers in research mode that use existing measurement systems to produce new information products and then funding the development and operational teams to transition a mature processing system as an addition to the operational data processing systems. These activities are needed, in addition to the other kind of R2O where a new type of measurement is moved from experimental to operational use, to produce useful information from the raw observations.

ACKNOWLEDGMENTS. The planning of an R2O process for ISCCP began when the authors were employed by NASA and NOAA, respectively. William B. Rossow retired from NASA Goddard Institute for Space Studies in 2006 and the City College of New York in 2017. John J. Bates retired from NOAA National Centers for Environmental Information—Asheville in 2016. The ISCCP transition team members on the research side were William B. Rossow, Alison W. Walker, and Cindy B. Pearl at the City College of New York and Violeta Golea and Joe Ferrier at NASA Goddard Institute for Space Studies and on the operations side were and are Alisa Young, Ken Knapp, Anand Inamdar, and Bill Hankins at NOAA National Centers for Environmental Information. We thank Ken Knapp and Otis Brown for very useful comments and suggestions. We also thank two anonymous reviewers, whose suggestion helped us improve the clarity of the text.

REFERENCES

- Bates, J. J., and J. L. Privette, 2012: A maturity model for assessing the completeness of climate data records. *Eos, Trans. Amer. Geophys. Union*, **93**, 441, <https://doi.org/10.1029/2012EO440006>.
- , —, E. J. Kearns, W. Glance, and X. Zhao, 2015: Sustained production of multidecadal climate records—Lessons from the NOAA Climate Data Record Program. *Bull. Amer. Meteor. Soc.*, **97**, 1573–1581, <https://doi.org/10.1175/BAMS-D-15-00015.1>.
- Christian, E., 2005: Planning for the Global Earth Observation System of Systems (GEOSS). *Space Policy*, **21**, 105–109, <https://doi.org/10.1016/j.spacepol.2005.03.002>.
- Global Climate Observing System, 1995: Plan for the Global Climate Observing System (GCOS). WMO Rep. GCOS-14, 49 pp.
- Knapp, K. R., J. J. Bates, and B. Barkstrom, 2007: Scientific data stewardship: Lessons learned from a satellite-data rescue effort. *Bull. Amer. Meteor. Soc.*, **88**, 1359–1361, <https://doi.org/10.1175/BAMS-88-9-1359>.
- NASEM, 2015: *Continuity of NASA Earth Observations from Space: A Value Framework*. National Academies Press, 118 pp., <https://doi.org/10.17226/21789>.
- , 2018: *Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space*. National Academies Press, 716 pp., <https://doi.org/10.17226/24938>.
- National Research Council, 2004: *Climate Data Records from Environmental Satellites: Interim Report*. National Academies Press, 150 pp., <https://doi.org/10.17226/10944>.
- Perry, J. S., 1975: The Global Atmospheric Research Program. *Rev. Geophys.*, **13**, 661–667, <https://doi.org/10.1029/RG013i003p00661>.
- Rossow, W. B., and R. A. Schiffer, 1991: ISCCP cloud data products. *Bull. Amer. Meteor. Soc.*, **72**, 2–20, [https://doi.org/10.1175/1520-0477\(1991\)072<0002:ICDP>2.0.CO;2](https://doi.org/10.1175/1520-0477(1991)072<0002:ICDP>2.0.CO;2).
- , and —, 1999: Advances in understanding clouds from ISCCP. *Bull. Amer. Meteor. Soc.*, **80**, 2261–2288, [https://doi.org/10.1175/1520-0477\(1999\)080<2261:AIUCFI>2.0.CO;2](https://doi.org/10.1175/1520-0477(1999)080<2261:AIUCFI>2.0.CO;2).
- , and Coauthors, 2016: International Satellite Cloud Climatology Project climate data record, H-series. NOAA National Centers for Environmental Information, <https://doi.org/10.7289/V5QZ281S>.
- Schiffer, R. A., and W. B. Rossow, 1983: The International Satellite Cloud Climatology Project (ISCCP): The first project of the World Climate Research

- Programme. *Bull. Amer. Meteor. Soc.*, **64**, 779–784, <https://doi.org/10.1175/1520-0477-64.7.779>.
- , and —, 1985: ISCCP global radiance data set: A new resource for climate research. *Bull. Amer. Meteor. Soc.*, **66**, 1498–1505, [https://doi.org/10.1175/1520-0477\(1985\)066<1498:IGRDSA>2.0.CO;2](https://doi.org/10.1175/1520-0477(1985)066<1498:IGRDSA>2.0.CO;2).
- White, R. M., 1979: The World Climate Conference: Report by the conference chairman. *WMO Bull.*, **28**, 177–178.
- Young, A. H., K. R. Knapp, A. Inamdar, W. Hankins, and W. B. Rossow, 2018: The International Satellite Cloud Climatology Project H-series climate data record product. *Earth Syst. Sci. Data*, **10**, 583–593, <https://doi.org/10.5194/essd-10-583-2018>.
- Zillman, J. W., 1980: The World Climate Programme. *Search*, **11**, 108–111.

NEW: SECOND EDITION


The Thinking Person's Guide to Climate Change

Robert Henson

EXPANDED AND UPDATED from the first edition and combining years of data with recent research, this accessible guide breaks down climate change into straightforward categories: the basics, the symptoms, the science, and solutions and actions. With a scientist's understanding and a journalist's voice, author Bob Henson has expanded this well-known primer to include global record highs, new research across the spectrum, and the Paris Agreement to cut greenhouse gasses.

Robert Henson is a meteorologist and writer at The Weather Company. His other books include *Weather on the Air: A History of Broadcast Meteorology*, also published by the American Meteorological Society.

© 2019, paperback, 570 pages
ISBN: 978-1-944970-39-0
List price: \$30 AMS Member price: \$20



AMS BOOKS
SCIENCE HISTORY SOCIETY
bookstore.ametsoc.org



AMS Washington Forum

*Connecting Science and Technology
with Business and Public Policy*

Save the Date
21–23 April, 2020 • Washington, DC

The 2020 AMS Washington Forum
Mark you calendars now for the 2020 AMS Washington Forum. This much-anticipated annual event provides an important platform to examine public policy issues across the weather, water and climate sciences.

This meeting is open for anyone interested in learning more about how weather, water, and climate information can help support risk management across the public and private sectors. This includes agencies and companies whose operations and planning are reliant on environmental factors.



AMS
American Meteorological Society