

High-resolution mapping of global surface water and its long-term changes

Jean-François Pekel¹, Andrew Cottam¹, Noel Gorelick² & Alan S. Belward¹

The location and persistence of surface water (inland and coastal) is both affected by climate and human activity¹ and affects climate^{2,3}, biological diversity⁴ and human wellbeing^{5,6}. Global data sets documenting surface water location and seasonality have been produced from inventories and national descriptions⁷, statistical extrapolation of regional data⁸ and satellite imagery^{9–12}, but measuring long-term changes at high resolution remains a challenge. Here, using three million Landsat satellite images¹³, we quantify changes in global surface water over the past 32 years at 30-metre resolution. We record the months and years when water was present, where occurrence changed and what form changes took in terms of seasonality and persistence. Between 1984 and 2015 permanent surface water has disappeared from an area of almost 90,000 square kilometres, roughly equivalent to that of Lake Superior, though new permanent bodies of surface water covering 184,000 square kilometres have formed elsewhere. All continental regions show a net increase in permanent water, except Oceania, which has a fractional (one per cent) net loss. Much of the increase is

from reservoir filling, although climate change¹⁴ is also implicated. Loss is more geographically concentrated than gain. Over 70 per cent of global net permanent water loss occurred in the Middle East and Central Asia, linked to drought and human actions including river diversion or damming and unregulated withdrawal^{15,16}. Losses in Australia¹⁷ and the USA¹⁸ linked to long-term droughts are also evident. This globally consistent, validated data set shows that impacts of climate change and climate oscillations on surface water occurrence can be measured and that evidence can be gathered to show how surface water is altered by human activities. We anticipate that this freely available data will improve the modelling of surface forcing, provide evidence of state and change in wetland ecotones (the transition areas between biomes), and inform water-management decision-making.

Between any two points in time, part of the Earth's surface is constantly underwater and part is never underwater, with the remainder fluctuating between these extremes. Coastlines and lake and river boundaries advance and retreat, rivers meander, new permanent lakes form and




Figure 1 | Different facets of surface water dynamics. **a**, Map of the USA showing Sacramento Valley location (red square). **b**, Surface water occurrence 1984–2015. **c**, Surface water occurrence change intensity 1984–2015. **d**, Surface water recurrence 1984–2015. **e**, Surface water seasonality 2014–2015. **f**, Transitions in surface water class 1984–2015. The Sacramento Valley is one of the major rice-growing

regions in the USA, extracted from the global data set. Seasonal water areas in the left and lower right of each panel correspond to flood irrigation, mainly rice paddies. The more permanent water features (centre and top right of each panel) are reservoirs. See Supplementary Information for a description of the water classes.

¹European Commission, Joint Research Centre, Directorate for Sustainable Resources, 20127 Ispra, Lombardy, Italy. ²Google Switzerland GmbH, Brandschenkestrasse 110, 8002 Zürich, Switzerland.




Figure 2 | Global surface water distribution and changes. Global maps, with 1° latitude/longitude summaries of surface water area shown on the right and underneath. **a**, Maximum water extent, permanent and seasonal surface water occurrence October 2014 to October 2015. **b**, Gains and loss

in permanent and seasonal surface water area between 1984 and 2015. All measurements made from inland and coastal waters are defined only by the GADM reference layer (see Methods).

others empty, while seasonal inundations and flood–irrigation cycles periodically create temporary bodies of water. When and where you find water on the planet’s surface is hugely important. The presence (or absence) of water influences the climate system, as accounted for in general circulation models², as well as CO_2 evasion and methane emissions³. Access to water influences movement, viable range and migrations for multitudes of species⁴; it is indispensable for sustainable development⁵ and can threaten the security of people, institutions and economies⁶.

Global surface water dynamics have been recorded from coarse-spatial-resolution satellite observations¹², higher-resolution seasonality maps have been produced using Landsat satellite imagery at 5- to 10-year intervals¹¹, and all Landsat images over multiple decades have been used to map seasonality and changes at continental¹⁹ and sub-continental²⁰ scales. The data set presented here (freely available from <https://global-surface-water.appspot.com/>) extends previous work by using the entire multi-temporal orthorectified Landsat 5, 7 and 8 archive spanning the past 32 years to map the spatial and temporal variability of global surface water and its long-term changes.

Each pixel in 1,823 terabytes of Landsat data (Extended Data Fig. 1a–c) was classified as open water, as land or as a non-valid observation using an expert system (see Methods, Extended Data Figs 2 and 3, and Supplementary Table 2). Open water is any stretch of water larger than 30 m by 30 m open to the sky, including fresh and

saltwater. Classification performance, measured using over 40,000 reference points (Methods and Extended Data Figs 4 and 5) confirmed that the classifier produces less than 1% of false water detections, and misses less than 5% of water (Extended Data Table 1).

The long-term water history was used to produce thematic products that document different facets of surface water dynamics. Figure 1 shows extracts from the global products for part of the USA’s Sacramento Valley (see Fig. 1a). Figure 1b shows occurrence (variations in persistence and location) between March 1984 and October 2015. The intensity with which occurrence increased or decreased over the 32 years documents gain, loss and constancy in persistence (Fig. 1c). The frequency with which water reappears from year to year across the time-series is mapped as recurrence (Fig. 1d), and water surfaces present throughout an entire year’s observations are separately mapped from those that are seasonal (Fig. 1e). Transitions between permanent water, seasonal water and land classes can be determined between any two years of observation; transitions between the first and last year of observation are shown in Fig. 1f. Temporal profiles document water history per pixel, per month and year, and change measurements at the global, continental and country scales are produced by combining these complementary information layers (Supplementary Table 1).

Three per cent (4.46 million km^2) of the Earth’s landmass was under water at some time between March 1984 and October 2015. Figure 2a




Figure 3 | Trends in annual permanent water surface area. a, Finland. b, New South Wales, Australia. c, Western states of the USA (Arizona, California, Idaho, Nevada, Oregon, Utah). d, Aral Sea (Kazakhstan, Uzbekistan). e, Iraq, Iran and Afghanistan. f, Tibetan plateau. Uncertainty

is estimated from the unobserved component of the maximum permanent water extent. The true surface water area is within this range. Trend lines are provided from years where the unobserved component is less than 5%.

shows that over half (52%) of this is found above 44° N, a pattern that corresponds overall with previous work^{9–11}. In 2015 permanent bodies of water covered 2.78 million km^2 , and 86% of this (2.4 million km^2) was geographically and temporally invariant, being consistently present across the entire observation record; the world's ancient lakes such as Baikal and Tanganyika, North America's Great Lakes and the Nordic region's 'land of a thousand lakes' are part of these truly permanent waters. But Fig. 2b also reveals striking patterns of surface water occurrence change.

Over the past three decades more than 162,000 km^2 of water bodies previously thought of as permanent have proved not to be so; almost 90,000 km^2 have vanished altogether and over 72,000 km^2 have transitioned from a permanent state to a seasonal state. During the same period almost 213,000 km^2 of new permanent water bodies came into existence; 29,000 km^2 of these used to be seasonally flooded but are now under water all year round, while 184,000 km^2 of permanent water formed in areas previously devoid of surface water. In the contemporary timeframe (October 2014 to October 2015) seasonal water covered 0.81 million km^2 .

Surface water is only a part of the water resource, but it is the most accessible to human populations²¹, and provides wide-ranging ecosystem services. Almost 52% of the planet's truly permanent surface water occurs in North America, home to less than 5% of the population in 2015²², and the continent also holds 18% of the contemporary seasonal water. Between 1984 and 2015 North America's permanent water area increased by 17,000 km^2 . In contrast, Asia, with 60% of the human population, accounts for only 9% of the truly permanent and 35% of the contemporary seasonal water. Asia has gained 71,000 km^2 of permanent water, which is a 23% increase for the continent. Africa and Latin America have almost the same share of the world's permanent water at around 9%, though their populations are very different, with Africa (16% of the total) supporting nearly twice as many people as Latin America (8.6%). Europe, including Russia, with 10% of the global population, has 22% of the permanent water and 18% of the

contemporary seasonal water. Oceania is the only continental region with a net loss of permanent water, albeit a tiny area at 229 km^2 .

The continental summaries obscure strong regional variation. Australia's millennium drought (between 2001 and 2009) substantially affected hydrology¹⁶, and measurable impacts of this can be seen in the permanent surface water area trends for the drought years (Fig. 3b). Regional variation is also apparent in the USA. Although the country's permanent surface water area overall has increased by 0.5% since 1984, a combination of drought and sustained demands for water¹⁷ have seen six western states lose 33%, more than 6,000 km^2 (Fig. 3c). These losses are modifying social behaviour, driving local water-management policy changes and causing shifts in agricultural production¹⁷.

Human behaviour may be changing in response to water distribution, but human action is itself changing surface water patterns. Over 70% of global net permanent water loss is concentrated in five countries. The marked negative anomaly in permanent water cover change centred at 45° N, 60° E (Fig. 2b) corresponds to this hotspot of change. Kazakhstan and Uzbekistan have lost much of the eastern lobe of the southern Aral Sea (Figs 3d and 4). The rate of loss was greatest between 1994 and 2009, though lately this has slowed and even partially reversed (Fig. 3d). Diversion of, and withdrawal from, the Amu and Syr rivers that once fed the lake are the main causes of loss, but changes in water management offer hope of stabilization and partial restoration¹⁵. Iran, Afghanistan and Iraq have also undergone major losses, having respectively 56%, 54% and 34% less permanent surface water in 2015 than in the first year of observation (Fig. 3e). These losses, which raise serious questions concerning water security and transboundary water management²³, are caused by factors including unregulated withdrawal, dams that change the flow rate and direction of rivers, and droughts¹⁶.

24 countries spread across all continents have each gained at least 1,000 km^2 of new permanent bodies of water (Supplementary Table 1). Much comes from reservoir construction (Extended Data Fig. 6),




Figure 4 | Surface water changes for the Aral Sea and Tibetan plateau. Transitions between the first year in which representative observations were acquired and the last year of observation. **a**, The Aral Sea. **b**, Tibetan

plateau. The regional context for each location is shown in the insets. The 32-year trends in water surface area for these mapped regions are shown in Fig. 3d and f. The key to the seasonal classes applies to both panels.

with most of the countries featuring prominently in the International Commission on Large Dams register of dam builders²⁴. Turkmenistan is an exception; there over 90% of the additional water surface—14,411 km²—comes from the recovery of the Garabogazköl Aylagy lagoon following the breach in 1992 of a dam between the lagoon and the Caspian Sea²⁵. Permanent surface water has also greatly increased without dam building on the Tibetan plateau (Fig. 4). Virtually all of the region's endorheic lakes are expanding and new lakes are forming, leading to a 20% increase, an extra 8,300 km² (Fig. 3f). Lake expansion on the plateau has been linked to increased run-off from accelerated snow-and-glacier melt caused by higher temperatures and annual precipitation¹⁴. Climate change adaptation challenges include grazing land reduction caused by inundation, grassland degradation linked to salination following the expansion of brackish waters, and threats to transport infrastructure²⁶. The climate-related gains in surface water in this region contrast markedly with the drought-related losses in Australia, the western states of the USA and in central Asia/the Middle East described above.

Seasonal water surfaces can show strong variability, moving between wet and dry years, even shifting geographically. Capturing such variability, especially for short-duration events, is challenging because cloud-free satellite observation must be concurrent with the water occurrence. The accuracy of seasonal water mapping is correspondingly lower (Extended Data Table 1). Changes in location, duration and timing depend on prevailing weather conditions (including changes driven by major perturbations such as the El Niño–Southern Oscillation), though erosion, sediment transport and deposition (particularly around coastlines and along river courses) and land-use choices also have an impact. Changes in any of these conditions can even drive transitions between permanent and seasonal classes (as defined in Methods). For example, many seasonally flooded paddy fields around the Sundarbans mangrove forest in Bangladesh and India have transitioned into permanently inundated fishponds (Extended Data Fig. 7). This may be through choice and market forces but may also be from necessity, as paddy field water and soils become increasingly saline with rising sea

levels and subsiding delta lands²⁷. Seasonally flooded lands can also be drowned in the backwaters of new dams, while downstream fluvial systems become increasingly fragmented and desiccated; drowning, fragmentation and desiccation are all evident for many dammed rivers such as the Paraná²⁸ (Extended Data Fig. 8), Colorado and Mekong rivers. These processes have broad impacts on humans and biodiversity²⁹, though accurate mapping will lead to improved management of water resources and a deeper understanding of connectivity, temporal characteristics and the consequences of land-management decisions³⁰.

The findings reported here reinforce the need for water-resource management strategies that integrate climate and socio-economic dimensions, as has already been proposed¹. This analysis applies a consistent algorithm to all 32 years of the Landsat observations to produce a validated data set that documents global surface water dynamics with new levels of spatial detail and accuracy. Linking this information to complementary data sets, such as satellite altimetry measurements, would produce fresh estimates of surface water volumes, river discharge and even sea-level rise¹². General circulation models that currently treat surface water in a simplistic fashion² should immediately benefit from the accurate location of truly permanent water surfaces. Mapping long-term changes in global surface water occurrence, documenting multi-decadal trends and identifying the timing (to within a given month or year) of events such as lake expansion and retreat or river-channel migration provides insights into the impacts of climate change and climate oscillation on surface water distribution, and concurrently captures the impacts humans have on surface water resource distribution.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 12 April; accepted 19 October 2016.

Published online 7 December 2016.

- Vörösmarty, C. J., Green, P., Salisbury, J. & Lammers, R. B. Global water resources: vulnerability from climate change and population growth. *Science* **289**, 284–288 (2000).
- Subin, Z. M., Riley, W. J. & Mironov, D. An improved lake model for climate simulations: model structure, evaluation, and sensitivity analyses in CESM1. *J. Adv. Model. Earth Syst.* **4**, M02001 (2012).
- Holgerson, M. A. & Raymond, P. A. Large contribution to inland water CO₂ and CH₄ emissions from very small ponds. *Nat. Geosci.* **9**, 222–226 (2016).
- Gardner, R. C. et al. *State of the World's Wetlands and Their Services to People: A Compilation of Recent Analyses*. Ramsar Briefing Note No. 7, <http://dx.doi.org/10.2139/ssrn.2589447> (Ramsar Convention Secretariat, SSRN, 2015).
- Vörösmarty, C. J. et al. in *Millennium Ecosystem Assessment Vol. 1 Ecosystems and Human Well-being: Current State and Trends* Ch. 7, 165–207, <http://www.unep.org/maweb/documents/document.276.aspx.pdf> (Island Press, 2005).
- World Economic Forum. *The Global Risks Report 2016* 11th edn, <http://www3.weforum.org/docs/Media/TheGlobalRisksReport2016.pdf> (World Economic Forum, 2016).
- Lehner, B. & Döll, P. Development and validation of a global database of lakes, reservoirs and wetlands. *J. Hydrol.* **296**, 1–22 (2004).
- Downing, J. A. et al. The global abundance and size distribution of lakes, ponds, and impoundments. *Limnol. Oceanogr.* **51**, 2388–2397 (2006).
- Verpoorter, C., Kutser, T., Seekell, D. A. & Tranvik, L. J. A global inventory of lakes based on high-resolution satellite imagery. *Geophys. Res. Lett.* **41**, 6396–6402 (2014).
- Feng, M., Sexton, J. O., Channan, S. & Townshend, J. R. A global, high-resolution (30-m) inland water body dataset for 2000: first results of a topographic-spectral classification algorithm. *Int. J. Digit. Earth* **9**, 113–133 (2015).
- Yamazaki, D., Trigg, M. A. & Ikeshima, D. Development of a global ~90m water body map using multi-temporal Landsat images. *Remote Sens. Environ.* **171**, 337–351 (2015).
- Prigent, C. et al. Changes in land surface water dynamics since the 1990s and relation to population pressure. *Geophys. Res. Lett.* **39**, L08403 (2012).
- Wulder, M. A. et al. The global Landsat archive: status, consolidation, and direction. *Remote Sens. Environ.* **185**, 271–283 (2016).
- Lutz, A. F., Immerzeel, W. W., Shrestha, A. B. & Bierkens, M. F. P. Consistent increase in High Asia's runoff due to increasing glacier melt and precipitation. *Nat. Clim. Chang.* **4**, 587–592 (2014).
- Micklin, P. The future Aral Sea: hope and despair. *Environ. Earth Sci.* **75**, 844 (2016).
- Zafarnejad, F. The contribution of dams to Iran's desertification. *Int. J. Environ. Stud.* **66**, 327–341 (2009).
- van Dijk, A. I. et al. The Millennium Drought in southeast Australia (2001–2009): natural and human causes and implications for water resources, ecosystems, economy, and society. *Wat. Resour. Res.* **49**, 1040–1057 (2013).
- MacDonald, G. M. Water, climate change, and sustainability in the southwest. *Proc. Natl Acad. Sci. USA* **107**, 21256–21262 (2010).
- Mueller, N. et al. Water observations from space: mapping surface water from 25 years of Landsat imagery across Australia. *Remote Sens. Environ.* **174**, 341–352 (2016).
- Tulbure, M. G., Broich, M., Stehman, S. V. & Kommareddy, A. Surface water extent dynamics from three decades of seasonally continuous Landsat time series at subcontinental scale in a semi-arid region. *Remote Sens. Environ.* **178**, 142–157 (2016).
- Postel, S. L., Daily, G. C. & Ehrlich, P. R. Human appropriation of renewable fresh water. *Science* **271**, 785–788 (1996).
- United Nations Department of Economic and Social Affairs, Population Division. *World Population Prospects: The 2015 Revision, Key Findings and Advance Tables*. Working Paper No. ESA/P/WP.241, https://esa.un.org/unpd/wpp/publications/files/key_findings_wpp_2015.pdf (United Nations, 2015).
- Najafi, A. & Vatanfada, J. Environmental challenges in trans-boundary waters, case study: Hamoon Hirmand Wetland (Iran and Afghanistan). *Int. J. Wat. Resour. Arid Environ.* **1**, 16–24 (2011).
- International Commission on Large Dams World Register* http://www.icold-cigb.org/GB/World_register/general_synthesis.asp?IDA=206 (IGIB/ICOLD, 2016).
- Kosarev, A. N., Kostianoy, A. G. & Zonn, I. S. Kara-Bogaz-Gol Bay: physical and chemical evolution. *Aquat. Geochem.* **15**, 223–236 (2009).
- Liu, B. et al. Outburst flooding of the moraine-dammed Zhuonai Lake on Tibetan plateau: causes and impacts. *IEEE Geosci. Remote Sens. Lett.* **13**, 570–574 (2016).
- Pethick, J. & Orford, J. D. Rapid rise in effective sea-level in southwest Bangladesh: its causes and contemporary rates. *Glob. Planet. Change* **111**, 237–245 (2013).
- Bonetto, A. A., Wais, J. R. & Castello, H. P. The increasing damming of the Paraná basin and its effects on the lower reaches. *Regul. Rivers Res. Manage.* **4**, 333–346 (1989).
- Vörösmarty, C. J. et al. Global threats to human water security and river biodiversity. *Nature* **467**, 555–561 (2010).
- Acuña, V. et al. Why should we care about temporary waterways? *Science* **343**, 1080–1081 (2014).

Supplementary Information is available in the online version of the paper.

Acknowledgements The USGS and NASA provided the Landsat imagery. R. Moore and her team provided the Google Earth Engine. R. Sargent and P. Dille from Carnegie Mellon University built the web interface to the global surface water occurrence maps, and M. Clerici and J. van 't Klooster built the web processing interface.

Author Contributions Each author contributed extensively and indispensably to the work presented in this paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to J.-F.P. (jean-francois.pekel@jrc.ec.europa.eu).

Reviewer Information Nature thanks I. Klein and D. Yamazaki for their contribution to the peer review of this work.

METHODS

Data. The entire archive of the Landsat 5 Thematic Mapper (TM), the Landsat 7 Enhanced Thematic Mapper-plus (ETM+) and the Landsat 8 Operational Land Imager (OLI) orthorectified, top-of-atmosphere reflectance and brightness temperature images (L1T)³¹ acquired between 16 March 1984 and 10 October 2015 was used^{32,33}.

Landsats 5, 7 and 8 are in near-polar orbit with repeat cover every 16 days; two satellites operate concurrently with an 8-day cycle. The ground area imaged by adjacent orbits overlaps by 7.3% at the Equator, increasing to 68.3% overlap at 70° latitude³⁴. Adjacent orbits to the west of a previous one are collected a week apart, thus pixels in the side-lap areas from both orbits are unique views. Sequential image frames along an orbit track also overlap, but these are identical data, not unique views; we exclude the end-lap but retain side-lap pixels for analysis.

Landsat 5, launched 1 March 1984, collected TM imagery until November 2011. Landsat 7 was launched 15 April 1999 and acquired imagery normally until 31 May 2003 when the scan line corrector (SLC) failed³⁵, and Landsat 8 began operational imaging in April 2013. The SLC failure causes around 22% of each scene to be lost³⁶. This loss increases from the centre, giving SLC-off images a slatted appearance at the edges.

Extended Data Fig. 1a–c shows variations in geographic coverage, first year of imaging and the number of images acquired from year to year. Landsat 5 had no on-board data recorders and its links with data relay satellites failed over time, so cover was often limited to the line of sight of receiving stations, and these did not provide complete global coverage³⁷. Commercial management of the programme from 1985 to the early 1990s³⁸ meant that acquisitions tended to be acquired only when pre-ordered³⁵. Geographic and temporal unevenness are particularly apparent up to 1999; the Americas, Western Europe and Africa were first imaged in 1984, Australia and South East Asia in 1986, but Kolyma was not imaged until 1995 and the northeastern part of the Siberian plateau not until 1999. These gaps are a feature of all global data sets based on the Landsat 30-m archives³⁹. Two Landsats operated from April 1999 onwards, which improved coverage, especially as Landsat 7 had on-board data recording capabilities. The programme's Long-Term Acquisition Plan also began in 1999, which further improved global coverage⁴⁰. Acquisitions fell between 2011 and 2012 as Landsat 5 operations were curtailed, though they increased in 2013 when Landsat 8 joined Landsat 7. More images were collected in June, July and August than December, January and February in all years, except for June 2003, following the SLC failure. However, by August 2003 data were once again routinely acquired (albeit with SLC-off gaps). From April 2013 the programme collected almost all images over land and in 2010 the USGS also began their Landsat Global Archive Consolidation, which ingests historic archive holdings from receiving stations around the world and produces orthorectified, top-of-atmosphere (L1T) data from these¹³. At the time of this study 3,066,080 L1T images (1,823 terabytes of data) covered 99.95% of the landmass. These images provided a local maximum number of unique views (including the side-lap areas) of 1881 and minimum of 11; the global mean was 537 and median 482. Pixels were excluded where over- or under-saturation occurred (manifested as random speckle or line-dropout), where one or more spectral bands were missing (typically occurring at image edges), and where scene geo-location was compromised.

Water detection. Although the goal of this study may seem simple—to separate water from other surfaces—consistent discrimination on the global scale, over multiple decades, is a non-trivial challenge. Water is a highly variable target and its spectral properties at the wavelengths measured by the TM, ETM+ and OLI sensors vary according to chlorophyll concentration, total suspended solids and coloured dissolved organic matter load, depths and water-body bottom material for shallow waters⁴¹, as well as variations in observation conditions (sun-target-sensor geometry, and optical thickness). On the global scale over 32 years all these conditions will be encountered somewhere at some time.

To address the challenge outlined above, techniques for big data exploration and information extraction less commonly used by the remote sensing community were exploited, namely expert systems^{42–44}, visual analytics⁴⁵ and evidential reasoning⁴⁶. Expert systems provide flexibility (to deal with the range of conditions encountered), visual analytics combine human cognitive and perceptual abilities with the storage and processing capacities of cloud computing platforms (the overall premise being that experts will have a better understanding of the problem to be solved if they interact with the data and view it through different representations and are thus able to design, test and fine-tune scenarios to extract the desired information), and evidential reasoning deals with problems related to both uncertainties and quality issues in the data set.

Expert systems are non-parametric classifiers that can account for uncertainty in data, incorporate image interpretation expertise into the classification process, and can be used with multiple data sources. The expert system outlined in Extended Data Fig. 2 was developed to assign each pixel to one of three target classes, either

water, land or non-valid observations (snow, ice, cloud or sensor-related issues). The inference engine of our system was a procedural sequential decision tree, which used both the multispectral and multitemporal attributes of the Landsat archive as well as ancillary data layers. Within the inference engine, expert knowledge was represented in the form of rules having the form: IF condition THEN inference. The condition contains equations describing the cluster hulls in a defined multispectral feature-space and can also be a combination of logical statements in which several components are linked through logical operators. The chaining of IF–THEN rules forms the problem-solving model that organizes and controls the steps and data used in the classification. Tracing the line of reasoning used by the inference engine during the development phase meant that the reason for the class choice associated with each pixel could be retrieved, which in turn allowed the reasons behind particular classification challenges to be identified. Solutions to identified failings could then be developed using evidential reasoning and addressed in subsequent iterations so that the overall performance of the classifier was progressively improved. When the performance of the expert system could no longer be noticeably improved it was applied to the entire Landsat L1T data set in a single run, and the output of this classification was then validated.

The equations describing the cluster hulls used in the expert system were established through visual analytics. The first step was to build a spectral library capturing the spectral behaviour of the three target classes across as wide a range of conditions as possible. 64,254 samples obtained through visual interpretation of 9,149 Landsat scenes recorded spectral variability of the target classes. Records held in the library comprised spectral values from all bands. These records were enriched by deriving the Normalized Difference Vegetation Index and Hue–Saturation–Value (HSV) colour-space transformations for the following band combinations: shortwave infrared (SWIR2); near-infrared (NIR); red; and NIR/green/blue using a standard transformation⁴⁷. The HSV colour model is well adapted for image analysis because the chromaticity (H and S) and the overall brightness (V) components are decoupled. This is highly desirable because changes in observation conditions first affect the V component and then the S component, while H remains relatively stable (except when the fundamental nature of the target changes, such as when land becomes water). Consequently, this property promotes temporal stability in the measurements and HSV-based classifications have been successfully used for near-real-time surface water detection at continental scales⁴⁸.

The information held in the spectral library was then analysed through visual analytics with the goal of extracting equations describing class cluster hulls in multidimensional feature-space for the expert system. An exact representation of clusters in feature-space helps to visualize class spectral overlap and thus to provide an understanding of the sources of thematic class uncertainty. Two-dimensional representations of the feature-space occupied by the samples held in the spectral library are shown in Extended Data Fig. 3. An exploratory data analysis tool was designed to support the interactive analysis. This provided multiple coordinated views⁴⁹ of the scatter plots and the source image from which any particular point within the plots was obtained. This two-way link between the representations of the samples within the multidimensional feature-space and source imagery meant that the geographic location (and context) of all samples could be considered along with their behaviour in feature-space. This was particularly valuable where the spectral properties of water coincided with those of other target classes. In these cases all dimensions of the feature-space could be examined to determine whether spectral separation was indeed possible.

The expert supervising the inference engine's development could then interactively draw the vertices of the hulls in this feature-space, which were then converted into equations through Delaunay triangulation⁵⁰ (Supplementary Table 2). Drawing the hulls directly into feature-space allowed the expert to account for irregular (even concave) shapes of clusters associated with the three target classes, and to deal adequately with the skewed data distributions so often encountered in remote sensing data sets. This approach also allowed cluster hulls to be defined that captured very infrequent—but nevertheless thematically important spectral behaviour. For example, the heavily sediment-laden seasonally occurring waters in parts of West Africa do not often appear in the time series, and are therefore poorly represented in the scatter plots, but these rare occurrences are important. The resulting equations are perforce complex because they constitute a precise representation of the complex shape of the three clusters within the multidimensional feature-space.

However, not all pixels could be unambiguously assigned to one or other of the target classes because overlap between the clusters was such that it could not be resolved anywhere in multidimensional feature-space. In these instances evidential reasoning—as part of the expert system—was used to guide class assignment. This took into consideration geographic location and the temporal trajectory in multispectral feature-space. Geographic location is important because spectral confusion between water and other surfaces can only occur in locations where

those other surfaces are found, and time is important because this spectral overlap may occur at specific times of the year and not at others.

The temporal trajectory in feature-space was used to gain information concerning the likelihood of a pixel being water. If a pixel sits unequivocally within a water hull for some of the time, then there is a high likelihood it will actually be water even if it occasionally occupies a hull where overlap occurs with other cover types, whereas if a pixel always sits in the overlapping domain or moves to the hull delimiting land this likelihood is removed. The frequency with which a pixel occupies the unequivocal portion of the water hull is used to estimate the likelihood of it actually being water.

Geographic locations linked to specific sources of spectral overlap were delineated using ancillary data layers. These constrain a portion of the data set within which specific decision rules could be identified and applied.

To detect water over glaciers spectral overlap from supraglacial moraines and shadow had to be resolved. The glacier areas were first delineated using the Randolph Glacier Inventory 5.0⁵¹. This defined the geographic regions where specific decision rules were required to confirm the water class assignment. Actual water over glaciers (melt ponds and surface streams) sometimes occupied the unequivocal portion of the water cluster in feature-space and the frequency with which this occurred determined the likelihood of water presence at that location.

Lava flow is also often wrongly assigned to the water class. Part of the spectral overlap between lava and water occurs in all dimensions of the multispectral feature-space, although again, over time, a pixel may move into a part of the feature-space where there is no overlap. A global-scale lava mask was established from both spectral characterization and visual interpretation of Landsat images, and within these boundaries the frequency with which pixels occupied 'lava overlap-free' portions of the water cluster were computed over the full span of the archive. This frequency was again used as an indicator of the likelihood of water presence for those locations.

Shadow, from whatever source, can cause false water detections because the underlying spectral characteristics of the surface are not truly represented, and may overlap with water. Three sources of shadow were addressed; buildings, terrain and clouds.

In the case of buildings, location and spectral behaviour over time were again brought into play. The Global Human Settlement Data Layer (GHSL)^{52,53} targeted areas where spectral confusion between water and the shadows cast by buildings needed to be resolved. Building shadows are seasonally dependent (related to changes in solar zenith angle), and thus a pixel may move in and out of the 'shadow-water overlap' cluster over time. This movement may be to a land cluster, or water. Over time, if a pixel exclusively moves to the water cluster, that location is very likely to be a permanent water feature within the urban area. The higher the frequency with which a pixel occupies the unequivocal water hull the greater the probability of water's presence at that location. This allows the expert system to map permanent water features in urban areas, and also to account for loss of water due to urban expansion and land reclamation. However, seasonal water detections within urban areas are more problematic, because these pixels will move between land and water in multispectral feature-space over time even in the absence of shadow.

Terrain shadows may be resolved using a threshold applied on slopes derived from a Digital Elevation Model (DEM). Unfortunately, this is only valid if the derived slopes represent the conditions on the date a given Landsat scene is acquired. Reservoirs are often built in steep terrain, and any dam built after the release date of the DEM would be masked, because the slopes recorded in the DEM are those prevailing before the reservoir filled. To detect new reservoirs, a mask covering areas where terrain shadow is expected was first derived from one of four DEMs^{54–57} (four were used to provide the best available DEM resolution at any given location). Within the masked area, pixels detected as water at multiple dates across the full year are likely to be water rather than seasonally cast shadow. Locations where water is detected in months when the sun is close to nadir are especially likely to be water. And again the movement into and out of the unequivocal water cluster in feature-space over time was used to identify actual water.

Cloud shadow is an even greater challenge than that from terrain or buildings, because it can occur anywhere and anytime. Thus the cloud shadow mask needed to be produced dynamically. Established cloud detection routines, such as FMASK⁵⁸, provide scene-by-scene identification of the location and extent of cloud shadow as this is linked to the clouds they detect. This supposes that the cloud producing the shadow is present in the image. But clouds outside a scene may still cast shadows in adjacent scenes. However, images from adjacent orbits cannot be used to resolve this because these occur a week apart. Consequently, cloud detection cannot be used as a robust indicator of cloud shadow. However, cloud shadows move over time, while water surfaces show more consistent temporal behaviour (even seasonal water can be expected in some months, and not at all in others). A temporal sliding

window was used across the water history record, post-classification. Within this sliding window, if the preceding and subsequent values were identified as water, then the observation under consideration was also likely to be water. Water detections bracketed by land detections are much more likely to be shadows. At these locations the frequency with which pixels occupied the unequivocal portion of the water cluster was again considered across the time series, and the lower this frequency the greater the likelihood of shadow.

Finally, visual inspection of the water maps identified scattered residual false detections, which were manually removed. Less than 0.002% (72 km²) of the maximum water extent was cleaned in this way. These errors were linked to industrial sites, photovoltaic farms and urban infrastructure not represented in the GHSL^{52,53}, such as airport runways. Some such errors may still remain, but their impact is accounted for in the validation.

While the drawing of the cluster hulls and the evidential reasoning are subjective, the visual analytics guide and inform the expert's decisions, which are based on objective use of the spectral library and the associated contextual information held in the image archive. Thus, this subjectivity is compensated for by an increased degree of confidence in the analysis⁵⁹ and the absence of biases associated with *a priori* assumptions concerning normal distributions associated with supervised classifiers such as the maximum likelihood or Mahalanobis distance.

The expert system was run in Google's Earth Engine, a computational infrastructure optimized for parallel processing of geospatial data plus a dedicated application programming interface and online access to the USGS L1T Landsat archive. Running the expert system on a single computer central processing unit (CPU) would have taken 1,212 years, but using 10,000 computers in Earth Engine the processing was completed in around 45 days, although building, testing and validating the expert system took almost two years.

Code availability. A web interface powered by Google Earth Engine allows the expert system to be run on any Landsat 5, 7 and 8 images. Access can be provided on request.

Validation. The expert system's performance was judged in term of errors of omission and commission at the pixel scale. The validation design took into account the small spatial extent of inland water surfaces (3% of the land surface) and its intrinsic spatio-temporal variation. The validation was performed using a total of 40,124 control points distributed both geographically (globally), temporally (across the 32 years), and across sensors (TM, ETM+ and OLI). Extended Data Fig. 4 summarizes the validation protocol.

Two reference data sets were produced: a sample of 27,268 pixels was dedicated to the estimation of the error of omission and 12,856 pixels were used to characterize errors of commission. Extended Data Fig. 5 shows the geographic distribution of all sample points and the associated errors.

To generate the omission error data set, a systematic sample frame (a grid 1° latitude by 1° longitude) was used. Under this frame the globe was stratified into areas with a high probability of water occurrence based on existing published global surface water maps, one for all latitudes up to 60° N (ref. 60) and one from 60° N to 78° N (ref. 10). A point was randomly selected within the water reference stratum in each latitude/longitude grid cell, and an image with less than 10% cloud cover corresponding to this location was then randomly selected for each sensor across the time span of the archive. In some locations the target of one sample per sensor at each location was not reached because of the absence of water at certain times, the total absence of observations over part of the archive or frequent cloud coverage.

Both of the published global water maps represent specific time periods and cannot guarantee the presence of water across all seasons and at all dates. The presence of water for each single validation point at each date was confirmed by visually checking all points using the validation tool described below. Only points confirmed as water were used in the estimation of the omission error.

To determine commission errors a grid 1° latitude by 1° longitude was again used as a systematic sample frame. For each latitude/longitude grid cell, images were randomly selected from the archives for all three sensors and the water detection expert system was run on these. One point per sensor from the areas classified as water was then randomly selected. The actual presence of water for each single validation point at each date was visually checked, again using the validation tool. If no water was in fact present, then this corresponded to an error of commission.

The validation tool was designed to facilitate the photointerpretation task. For each pixel in the database the full-resolution Landsat image selected for the validation (that is, randomly allocated in time to those pixel coordinates) was displayed and the specific pixel for interpretation highlighted. Images from the same location acquired at dates before and after that of the sample image were also presented, along with high-resolution satellite imagery and/or aerial photography from Google Earth and the Environmental Systems Research Institute (ESRI).

Using all these inputs an expert interpreter confirmed or refuted the presence of water at that location and time. The database was automatically populated with the results.

Errors of omission overall were less than 5% and commission less than 1%. Breaking down the validation sample by sensor showed that all three performed comparably well; commission errors ranged by only 0.2% and omission by 1.2%. Extended Data Table 1 provides details.

Breaking down the validation sample further by water seasonality class (using the seasonality class for the relevant year in which the validation point was randomly acquired) shows that although all three sensors performed similarly, omission for seasonal water classes was, as expected, higher than that for permanent water classes. Accuracy for TM, ETM+ and OLI when judged against errors of commission were respectively 99.6%, 99.5% and 99.7%, for permanent water, while accuracies considering errors of omission for permanent water were 98.8%, 97.8% and 99.1%. Errors of commission for seasonal water were very slightly higher, with accuracies of 98.8% (TM), 98.4% (ETM+) and 98.5% (OLI) and the greater errors of omission for seasonal water were reflected in accuracies of 74.9% (TM), 73.8% (ETM+) and 77.4% (OLI).

Errors of omission for seasonal water are higher than for permanent because there are fewer opportunities to observe each water body. These errors will result in an underestimation of its occurrence. But omission errors do not occur in the same place over time and over the 32 years multiple opportunities for observation arise. Thus sites where seasonal water can occur may be missed at one date, but may be correctly mapped at another. Overall, less than 1% of the points (214 samples out of 27,268) in the validation database where water was actually present remained entirely unmapped over time (Extended Data Table 1).

Thematic mapping. The maximum water extent (all locations ever mapped as water) during the 32 years was mapped. The frequency with which water was present on the surface from March 1984 to October 2015 was captured in a single product called surface water occurrence (SWO). To compute SWO, the water detections (WD) and valid observations (VO) from the same months are summed, that is, water detections and valid observations from March 1984 are added to water detections and valid observations from March 1985 and so on, such that $SWO_{month} = \sum WD_{month} / \sum VO_{month}$. Averaging the results of all monthly SWO_{month} calculations gives the long-term overall surface water occurrence. The month-by-month time step normalizes occurrence against seasonal variation in the number of valid observations across the year. Typically, more cloud-free observations (and thus valid observations) are available during dry seasons than wet. Without monthly weighting, the overall water occurrence (that is, computed over the full period) would be biased by temporal distribution of the valid observations (that is, giving more weight to the dry season than to the wet season). Extended Data Fig. 1c documents the number of scenes per month and year across the archive.

Change in water occurrence intensity between two epochs (16 March 1984 to 31 December 1999, and 1 January 2000 to 10 October 2015) was also produced (Extended Data Fig. 6a). This is derived from homologous pairs of months (that is, the same months contain valid observations in both epochs). The occurrence difference between epochs was computed for each pair and differences between all homologous pairs of months were then averaged to create the surface water occurrence change intensity map. Areas where there are no pairs of homologous months could not be mapped. The averaging of the monthly processing mitigates variations in data distribution over time (that is, both seasonal variation in the distribution of valid observations, temporal depth and frequency of observations through the archive) and provides a consistent estimation of the water occurrence change. This map shows where surface water occurrence increased, decreased or remained invariant between the two epochs.

The occurrence and occurrence change intensity maps provide a summary of the location and persistence of water on the surface, but they do not describe inter- and intra-annual variability. We propose water recurrence as a measurement of the degree of inter-annual variability in the presence of water. This describes how frequently water returned from one year to another (expressed as a percentage). Recurrence refers specifically to the temporal behaviour of water surfaces; unlike occurrence, recurrence is not systematically computed over the full span of the archive, because water may not have been present from the beginning to the end of the archive. Thus, we first have to define a 'water period'—that is, that part of the archive where water was present at least from time to time; the recurrence in fact quantifies this 'time to time'. The water period is established individually for each pixel. The water period runs from the first month in the first year in which water is observed to the last month of the last year in which water is observed of the entire 32-year period. In addition to defining the water period we also need to define a 'water season' (not equivalent to a 'wet season'). The water season is identified from the monthly water recurrence and is defined as those months of the year that from

time to time have water. A 'water year' is a year with at least one water observation, while an 'observation year' is a year with at least one valid observation within the water season. Water recurrence is then calculated as the ratio of the number of water years to observation years. The count of the number of years starts with the year in which water was first observed and ends with the most recent year in which water was observed. Years that contain only observations outside the water season are not counted; we have no way of knowing whether water might have occurred in the water season because we have no observations.

We also describe the intra-annual distribution of water, which discriminates between 'permanent' and 'seasonal' water surfaces. A permanent water surface is underwater throughout the year, while a seasonal water surface is underwater for less than 12 months of the year. In some places we do not have observations for all 12 months of the year (for example, because of the polar night in winter) and in these cases water is considered to be seasonal if the number of months where water is present is less than the number of months where valid observations were acquired. A second consideration is that some lakes freeze for part of the year. However, during the frozen period water is still present under the ice layer, both for lakes and the sea. The expert system treats ice as a non-valid observation, so the observation period corresponds only to the unfrozen months. If water is present throughout the observation period (that is, the unfrozen period), the lake is considered to be a permanent water surface. If the area of the lake contracts during the unfrozen period, then the pixels along the borders of the lake no longer represent water, and those pixels will be considered to represent seasonal water surface.

Seasonality is computed for every year. A single data set for the contemporary period (October 2014 to October 2015) is made available via the website and Extended Data Fig. 7a provides an example. The individual years' computations are used to produce the trends analyses provided in Fig. 3. Plotting the measured permanent surface water area for each year would provide such trends. But the gaps in the observation record are a source of uncertainty, especially in the early years of the archive. Consequently, part of the permanent water surface is potentially not taken into account (not observed), and constitutes a source of underestimation of the reported area. To account for this we combine the measured values of permanent surface water area with an estimate of the area of unobserved but potentially permanent surface water. This is computed using the maximum permanent water extent, that is, any pixel that has ever been identified as permanent in any year of the record across the full 32 years, minus the observed land values, minus any observed seasonal water, minus the observed permanent water. The true permanent surface water area will lie somewhere within this unobserved range, but the actual limits cannot be established. There is no uncertainty in those instances where the observation record is complete; conversely, in those years where no observations were made uncertainty is absolute.

Trends are calculated for defined geographic regions. The derived trend parameters (water loss per year, correlation, *P*-value) were computed using years where this unobserved area was less than 5% of the observed area. For countries and state-level reporting the Global Administrative Areas dataset (GADM)⁶¹ was used, for the Tibetan Plateau the boundary established by the Chinese Institute of Geographic Sciences and Natural Resources Research⁶² was used, and the trends over the Aral Sea were delimited by top-left coordinates 47° N, 58° E, bottom-right 43.8° N, 61.5° E.

Temporal profiles. Three histograms are generated. First, monthly recurrence shows the intra-annual distribution of the water, and characterizes water seasonality. It also provides information on the water recurrence for each single month. Second, a water history chart shows the class (land, seasonal water and permanent water) for each year in which valid observations were acquired. Third, month-by-month presence of water and observations within any single year can be extracted. Examples can be seen in Extended Data Fig. 6b.

Measuring change. The thematic maps and temporal profiles were used to identify a set of water classes that characterize transitions between the first year in which representative observations were acquired and the last year of observation. Representative years are identified by comparing each year in turn with the annual pattern of monthly recurrence from the temporal profiles. These profiles identify months in which water was observed, and indicate the percentage of valid observations classified as water in any given month. A year is flagged as representative if it contains sufficient valid observations from any combination of months to bring confidence to the determination of the presence or absence of water. The overall level of confidence is determined by the annual sum of the monthly long-term recurrences of observed months (per year). The rationale is that the likelihood of a real absence of water for a year is higher if the water is absent for months showing a high long-term water recurrence than from one showing small rates of recurrence. In the latter case the absence of water may be explained by a seasonal shift, and does not confer enough confidence to conclude that water was not present later.

Therefore, we considered that if the sum of the recurrence of the observed months is greater than 100, the absence of water observation brings enough confidence to consider that water was actually not present. Conversely, a single water presence is enough to demonstrate water presence. The water class in that representative year is then fixed as the 'first' year. The last year's water class is always the class assigned to the last year of observation (October 2014 to October 2015) because we have enough observation available within a year during this period.

The following transitions were mapped: unchanging permanent water surfaces; new permanent water surfaces (conversion of land into permanent water); lost permanent water surfaces (conversion of permanent water into land); unchanging seasonal water surfaces; new seasonal water surfaces (conversion of land into seasonal water); lost seasonal water surfaces (conversion of a seasonal water into land); conversion of permanent water into seasonal water; and the conversion of seasonal water into permanent water (Extended Data Fig. 8 provides an example).

These conversions refer to changes in state from the beginning and end of the time series; they do not describe what happened in the intervening years, so an unchanging water surface means that the seasonality at that particular point was the same in the first and last year it was observed, and not necessarily that it was stable throughout. Stability must be checked at the pixel scale by using the long-term water history described by the temporal profiles plus the recurrence and occurrence maps. There are instances where water is not present at the beginning or the end of the observation record but is present in some of the intervening years. By tracking the inter-annual patterns of such 'ephemeral' events and their intra-annual characteristics, each such pixel can be classified as either ephemeral permanent water (land replaced by permanent water that subsequently disappears) or ephemeral seasonal water (land replaced by seasonal water that subsequently disappears), depending on the majority of the observed seasonality during the period of water presence.

The GADM⁶¹ was also used to extract water area statistics at national, continental and global levels from the occurrence, recurrence and transition maps plus the long-term water history. Area measurements (in km²) for the following classes are reported in Supplementary Table 1: maximum surface water occurrence over 32 years; permanent water in the first year of observation; permanent water in the last year of observation; permanent water with 100% recurrence; transition from land to permanent water; transition from seasonal to permanent water; transition from permanent water to land; transition of permanent water to seasonal; seasonal water in the first year of observation; seasonal water in the last year of observation; and seasonal water with 100% recurrence.

Known issues and planned improvements. Bodies of water smaller than 30 m by 30 m, those obscured by floating, overhanging and standing vegetation or hidden by infrastructure such as tunnels and bridges were not included. Irrigated fields that stand in open water for some weeks were mapped but not when crop cover is well established. When observations coincide with paddy flooding, yet predate crop emergence and cover, then paddy fields are mapped, but inevitably some will be missed. Paddy fields in steep terrain present particular challenges because of their small size. Paddy fields are an example of short-duration events, but in fact short-duration seasonal water more generally is likely to be underestimated because of geographic and temporal discontinuities in the archive and gaps caused by persistent cloud cover.

The precision of the metrics at any location improves as the number of valid observations increases. The meta-information in the web interface documents the number of valid observations at each pixel location, which provides users with a proxy measure of confidence.

Long-term changes cannot be determined uniformly for the entire globe because the observation record varies; the first year of observation is 1984 for much of the world, but not for parts of the Siberian plateau and Kolyma (1999 and 1995 respectively). For these reasons we recommend caution in interpreting the changes in these particular regions, though given the geographic completeness of current L1T coverage we have confidence in the contemporary figures reported. The water detections are accurate (as determined by the validation), but the time range over which the transitions occurred in these regions is perforce shorter. Northern Hemisphere high latitudes are also problematic because the observation season is short, the solar zenith angles are low, and the archive is poorly populated above 78° N. One consequence is that at high latitudes, the seasonality of the ocean is occasionally wrongly reported as seasonal.

On the surface of glaciers, melt ponds and streams are identified here as permanent water surfaces, but this permanence is questionable, because unlike deeper lakes where only the surface freezes, here the full water volume may actually freeze.

The recurrence value may be overestimated for water surfaces that came into existence in the last few years of the time series. The temporal profiles identify

the years in which water occurred, and using these, decisions can be taken as to whether recurrence is considered relevant at any given pixel location.

Some isolated commission errors still occur in urban areas as the GHSL^{52,53} is itself being improved; roofs, coal and waste heaps and runways are the most common sources of confusion. These will be resolved as the urban information layers improve.


The SLC-off condition also introduces artefacts because the slatted appearance of the original images is occasionally carried into the water occurrence maps. These are not misclassified pixels but occur because the sampling between the gaps is greater than within them, allowing different water conditions to be captured. Techniques to fill the SLC-off gaps exist³⁶, but these create artificial values, and considering the strong temporal and spatial variability of water these techniques require careful use, as they may create false water detections. Water mapping using multitemporal time series on continental scales avoided the use of such techniques on these grounds¹⁹, though did note the same issue of SLC-off artefacts appearing in the final maps.

Analysis of coarse-spatial-resolution satellite data sets from systems offering daily revisit (nearly one observation a day over each location) first captured the inter- and intra-annual variability of surface water occurrence^{12,63,64}, and although the Landsat missions offer 8-day or 16-day rather than daily observations, they do provide high-resolution land surface observations spanning more than three decades. The spatial and temporal information reported in this data set complements that acquired in the past. Nevertheless, the biggest limitation to global surface water occurrence mapping from these data are undoubtedly the geographic and temporal discontinuities of the archive itself. The Landsat archive is continuously enriched through new acquisitions and recovery of old data from international receiving stations¹³. Imagery from other satellites in this resolution class could also be used to improve the temporal sampling. At least 24 other satellites have gathered multispectral imagery at resolutions of 20 m to 30 m from near-polar orbits concurrent with the Landsat programme⁶⁵. These are managed by at least 12 different sovereign states, and although data access is not always at the exemplary full, free and open level of Landsat, some systems do provide this, for example the European Union's Sentinel 2a satellite launched in 2015; the next version of the expert system will also ingest these data streams. Landsat 4, which was launched 16 July 1982 and in operation until 14 December 1993 also carried a 30-m-resolution TM, though the satellite contributes only a fraction to the total TM holdings in the archive, and much of this was restricted to the conterminous USA in the first two years of operation³⁶. Nevertheless, future reanalysis will include the Landsat 4 data and could possibly be extended back to 1972 through the inclusion of data from the Landsat Multi Spectral Scanner (MSS), though this would be challenging because of the more limited spectral, spatial and temporal dimensions of these data sets⁶⁶. Combining all available satellite observations with petabyte processing power would put real-time monitoring of change to Earth's inland and coastal waters within reach.


Data availability. The Landsat imagery used in this study is available from the USGS <http://earthexplorer.usgs.gov> and in Google Earth Engine <https://earthengine.google.com>. The data sets generated during the current study are available from <https://global-surface-water.appspot.com>. The data used to generate Fig. 2 and Extended Data Fig. 1c are provided as Source Data.

1. *Landsat 8 Data Users Handbook* <http://landsat.gsfc.nasa.gov/?p=10659>, USGS Publication LSDS-1574 (US Geological Survey, 2016).
2. Woodcock, C. E. *et al.* Free access to Landsat imagery. *Science* **320**, 1011 (2008).
3. Wulder, M. A. *et al.* Opening the archive: how free data has enabled the science and monitoring promise of Landsat. *Remote Sens. Environ.* **122**, 2–10 (2012).
4. *Landsat 7 Science Data Users Handbook* http://landsathandbook.gsfc.nasa.gov/orbit_coverage/prog Sect5_2.html (NASA, accessed 16 November 2016).
5. Markham, B. L., Storey, J. C., Williams, D. L. & Irons, J. R. Landsat sensor performance: history and current status. *IEEE Trans. Geosci. Remote Sens.* **42**, 2691–2694 (2004).
6. Chen, J. *et al.* A simple and effective method for filling gaps in Landsat ETM+ SLC-off images. *Remote Sens. Environ.* **115**, 1053–1064 (2011).
7. Goward, S. *et al.* Historical record of Landsat global coverage. *Photogramm. Eng. Remote Sensing* **72**, 1155–1169 (2006).
8. Loveland, T. R. & Dwyer, J. L. Landsat: building a strong future. *Remote Sens. Environ.* **122**, 22–29 (2012).
9. Gutman, G. *et al.* Assessment of the NASA-USGS global land survey (GLS) datasets. *Remote Sens. Environ.* **134**, 249–265 (2013).
10. Arvidson, T., Gasch, J. & Goward, S. N. Landsat 7's long-term acquisition plan—an innovative approach to building a global imagery archive. *Remote Sens. Environ.* **78**, 13–26 (2001).
11. Arst, H. *Optical Properties and Remote Sensing of Multicomponental Water Bodies* Vol. XII of *Marine Science and Coastal Management* Ch. 1 (Springer Science Praxis, 2003).
12. Lu, D. & Weng, Q. A survey of image classification methods and techniques for improving classification performance. *Int. J. Remote Sens.* **28**, 823–870 (2007).


43. Kartikeyan, B., Majumder, K. L. & Dasgupta, A. R. An expert system for land cover classification. *IEEE Trans. Geosci. Remote Sens.* **33**, 58–66 (1995).
44. Shoshany, M. Knowledge based expert systems in remote sensing task: quantifying gains from intelligent inference. *Int. Soc. Photogramm. Remote Sens. Arch.* XXXVII (B7) 1085–1088, http://www.isprs.org/proceedings/XXXVII/congress/7_pdf/6_WG-VII-6/06.pdf (XXIst ISPRS Congress, Technical Commission VII, 2008).
45. Keim, D. A. et al. in *Visual Data Mining* 76–90, http://kops.uni-konstanz.de/bitstream/handle/123456789/5631/Visual_Analytics_Scope_and_Challenges.pdf?sequence=1&isAllowed=y (Springer, 2008).
46. Yang, J.-B. & Xu, D. L. On the evidential reasoning algorithm for multiple attribute decision analysis under uncertainty. *IEEE Trans. Syst. Man Cybern. A* **32**, 289–304 (2002).
47. Smith, A. R. Color gamut transform pairs. *Comput. Graph.* **12**, 12–19 (1978).
48. Pekel, J.-F. et al. A near real-time water surface detection method based on HSV transformation of MODIS multi-spectral time series data. *Remote Sens. Environ.* **140**, 704–716 (2014).
49. Roberts, J. C. in *Coordinated and Multiple Views in Exploratory Visualization* (CMV'07 Fifth Int. Conf.) 61–71, <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4269947&isnumber=4269933> (IEEE, 2007).
50. Delaunay, B. Sur la sphere vide. *Bull. Acad. Sci. USSR* **7**, 793–800, <http://www.mathnet.ru/links/bf140e013bb2829a727614ee4e41051a/im4937.pdf> (1934).
51. Arendt, A. et al. *Randolph Glacier Inventory—A Dataset of Global Glacier Outlines: Version 5.0* <http://www.glims.org/RGI/> (Global Land Ice Measurements from Space, Digital Media, 2015).
52. Pesaresi, M. et al. A global human settlement layer from optical HR/VHR RS data: concept and first results. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **6**, 2102–2131 (2013).
53. Pesaresi, M. et al. *Operating Procedure for the Production of the Global Human Settlement Layer from Landsat Data of the Epochs 1975, 1990, 2000, and 2014* <http://publications.jrc.ec.europa.eu/repository/handle/JRC97705> (Publications Office of the European Union, 2016).
54. *Global 30-Arc Second Elevation Data Set (GTOPO30)* <https://lta.cr.usgs.gov/GTOPO30> (Department of the Interior, USGS, 1996).
55. Danielson, J. J. & Gesch, D. B. *Global Multi-Resolution Terrain Elevation Data 2010 (GMTED2010)*. USGS Report 2011–1073, <https://pubs.er.usgs.gov/publication/ofr20111073> (USGS Publications Warehouse, 2011).
56. Jarvis, A., Reuter, H. I., Nelson, A. & Guevara, E. *Hole-filled SRTM for the Globe Version 4* <http://srtm.csi.cgiar.org> (CGIAR-CSI SRTM 90m Database, 2008).
57. *Shuttle Radar Topography Mission (SRTM) 1 Arc-Second Global* <https://lta.cr.usgs.gov/SRTM1Arc> (Land Processes Distributed Active Archive Center (LP DAAC), USGS/EROS, accessed November 2016).
58. Zhu, Z. & Woodcock, C. E. Object-based cloud and cloud shadow detection in Landsat imagery. *Remote Sens. Environ.* **118**, 83–94 (2012).
59. Seipp, K., Ochoa, X., Gutiérrez, F. & Verbert, K. A research agenda for managing uncertainty in visual analytics. *Gesellsch. Inform.* 1–10 (Human Factors in Information Visualization and Decision Support Systems (HFIDSS), Mensch und Computer Workshopband, 2016).
60. *Shuttle Radar Topography Mission Water Body Data* https://lta.cr.usgs.gov/srtm_water_body_dataset (SRTM Water Body Data (SWBD), 2003).
61. *Global Administrative Areas (GADM) version 2.6*, <https://waterloo.ca/library/geospatial/collections/us-and-world-geospatial-data-resources/global-administrative-areas-gadm> (Univ. Berkeley, Museum of Vertebrate Zoology and the International Rice Research Institute, 2012).
62. Zhang, Y., Li, B. & Zheng, D. Datasets of the boundary and area of the Tibetan Plateau. *Glob. Change Res. Data Publ. Repository* <http://www.geodoi.ac.cn/weben/doi.aspx?id=135> (2014).
63. Papa, F. et al. Interannual variability of surface water extent at the global scale, 1993–2004. *J. Geophys. Res.* **115**, D12 (2010).
64. Klein, I. et al. Results of the Global WaterPack: a novel product to assess inland water body dynamics on a daily basis. *Remote Sens. Lett.* **6**, 78–87 (2015).
65. Belward, A. S. & Skøien, J. O. Who launched what, when and why; trends in Global Land-Cover Observation capacity from civilian Earth Observation satellites. *ISPRS J. Photogramm. Remote Sens.* **103**, 115–128 (2015).
66. Cohen, W. B. & Goward, S. N. Landsat's role in ecological applications of remote sensing. *Bioscience* **54**, 535–545 (2004).




Extended Data Figure 1 | Geographic and temporal coverage of the Landsat 5, 7 and 8 L1T archive between 16 March 1984 and 10 October 2015.
a, Total number of unique views. **b**, First year of imaging. **c**, Number of scenes per month and year.




Extended Data Figure 2 | Diagram of the expert system classifier.




Extended Data Figure 3 | Multispectral feature-space occupied by water and other surfaces. **a**, Hue (from SWIR2, NIR, red) versus NDVI. **b**, Hue versus Value (both from SWIR2, NIR, red). **c**, Hue versus Saturation (both from SWIR2, NIR, red). **d**, Hue versus Value (both from NIR, green, blue).



Extended Data Figure 4 | Diagram of the validation protocol. a, Omission error protocol. b, Commission error protocol.




Extended Data Figure 5 | Global geographic distribution of validation sample points and error.




Extended Data Figure 6 | Mapping the history of surface water occurrence. **a**, Examples of increasing surface water occurrence in Myanmar (see inset for regional context). **b**, Pixel-based temporal profiles showing recurrence by month over 32 years (top), water history by seasonality class and by year over 32 years (middle) and monthly water presence for each year in the water seasonality record, in this case 2011 (bottom). Collectively, the graphs show that at this location (latitude

18.3928°, longitude 96.2633°) there are no valid observations available for the period 1984–1986, in 1993, 1997 or 1998 (the gaps in the middle graph), that before 2011 this was dry land, that the dam formed in 2011 and this point was flooded sometime between April and September (bottom), but since then it has been permanent water (centre), and that in the 32 years of observation water has not been detected in June (no observations have been made in June since the dam filled (top)).



Extended Data Figure 7 | Mapping changes in intra-annual persistence (seasonality). **a**, Surface water seasonality between October 2014 and October 2015 in the Sundarbans in Bangladesh (see inset). **b**, Changes in inter-annual persistence between 1984 and 2015. The increase in

permanent surface water at the expense of seasonal is indicative of changes in land use from seasonally flooded paddy fields to permanently flooded fishponds.



Extended Data Figure 8 | Water transitions map the Parana river. The regional context is shown in the insets. River channel migration, changes to seasonal water across the floodplain and transitions from permanent to seasonal water ('New seasonal') and seasonal to permanent water ('New permanent'), visible in the figure, are symptomatic of habitat fragmentation and changing ecosystem service delivery.

Extended Data Table 1 | Validation results

a)

	Landsat 5			Landsat 7			Landsat 8		
	Overall	Seasonal	Permanent	Overall	Seasonal	Permanent	Overall	Seasonal	Permanent
Misclassified as land #	18	6	12	20	7	13	30	10	20
Correctly classified as water #	3226	493	2733	3037	424	2613	6525	671	5854
Commission accuracy	99.45%	98.80%	99.56%	99.35%	98.38%	99.50%	99.54%	98.53%	99.66%

b)

	Landsat 5			Landsat 7			Landsat 8		
	Overall	Seasonal	Permanent	Overall	Seasonal	Permanent	Overall	Seasonal	Permanent
Misclassified as land #	233	146	87	343	172	171	425	336	89
Correctly classified as water #	7561	436	7125	7808	485	7323	10898	1151	9747
Omission accuracy	97.01%	74.91%	98.79%	95.79%	73.82%	97.72%	96.25%	77.40%	99.10%

Accuracy results judged against: **a**, commission error by sensor and by seasonality class; **b**, omission error by sensor and by seasonality class.