☑ **Fact-Checked Science**

PAPER CHECK · UPPER ATMOSPHERE

Noctilucent clouds may sit lower than last century, but “99.5%” assumes the methods agree

Six researchers compared about 20,000 camera measurements of noctilucent clouds made since 2020 with triangulations going back to 1890, and found the clouds 0.69 km lower, a trend they report at “better than the 99.5% confidence level” [1 §5]. We rebuilt their fit from the preprint’s own figures: every number we could recompute reproduces, and we found no error. The drop holds at about the 95% level under most choices, but the 99.5% assumes that old photographs and today’s cameras give the same height for the same clouds, which the preprint doesn’t test and we found no test of. S0 News linked the preprint on November 6, 2026 [2].

By Fact-Checked Science · Published November 6, 2026

OUR ANALYSIS marks our own reasoning and calculations, including the values we read from the preprint’s figures and every number from our re-fit, and the verdicts are our assessment. Everything else comes from the numbered sources. For the preprint [1], § gives a section, Fig. a figure, Table its Table 1 and Eq. an equation. Heights are averages over a night of measurements unless we say otherwise.

The numbers we could recompute reproduce; the confidence rests on old and new methods agreeing

OUR ASSESSMENT

0.69 km

the drop, ±0.24 km, between the nights measured before 2000 and those since 2020, as the preprint finds it [1 §5]; our rebuild of its figures gives 0.68 ± 0.24. Our analysis

22 nights

of historical triangulation, in 14 summers from 1890 to 1992, against 48 camera nights since 2020, counted from the paper’s Fig. 4 [1 Fig. 4]. Our analysis

0.2 km

a difference between old and new methods, modern heights reading low, that would take the drop below the 95% level. Our analysis, a check added after our first run

THE PREPRINT’S CLAIMS

OVERSTATED

“reveals a descending trend in NLC altitude with a 99.5% confidence level”

The preprint’s abstract [1 Abstract]

Rebuilt from the paper’s own figures, every number we could recompute reproduces. Under our pre-set rules, grouping nights by summer or by era gives p from 0.0006 to 0.14, and counting each summer once gives p ≈ 0.02; the data show no sign that a summer’s nights are alike. But the ±0.18 [1 §5], our refit shows, covers only how much the nightly averages scatter. It leaves out any difference between 1890–1992 photographic triangulation and the 2020s camera network, which the paper doesn’t test: an offset of about 0.2 km would erase even the 95% level, by a check added after our first run. Our analysis. See sections 3 and 5.

UNCONFIRMED

“The current descent rate is −1.2 km per century, or half of the model prediction based on the current increase of CO_2 in the atmosphere.”

The preprint’s abstract [1 Abstract]

The rate is the model’s slope since 1960, 0.233 km a decade in Lübken and colleagues’ simulation [5 Fig. 3], times the fitted share K of 0.52 [1 Table 1]. The data fix one number, how far the 2020s sit below the older nights; the timing comes from the model, and seven modern summers can’t measure a current rate. A straight line through all the years gives 0.8 ± 0.3 km a century. A descent of this size is plausible: satellites see the summer mesosphere shrinking by 100–200 m a decade [17 Abstract]. Our analysis. See sections 3 and 7.

TRUE

“the minimum of mean altitude is reached near July 10”

The preprint, section 5 [1 §5]

Our refit puts the lowest clouds 17 days after the solstice, about 8 July. Lidar in Norway finds the clouds’ height varying with the season by up to 1–1.5 km [8 Abstract] and falling as the season goes on [14 Abstract], and a satellite instrument, SOFIE, finds them furthest below the mesopause in mid-season [23 Abstract]. Our analysis. See section 4.

PARTLY TRUE

“The change of the lower altitude is about twice the change of the mean altitude, being close to the descending rate predicted by Lübken et al. (2018).”

The preprint, section 6 [1 §6]

The first half holds in the plotted nights: allowing for the season, the lowest tenth fell by 1.5 km against about 0.7 km for the mean. The reading that the clouds’ bottom sank at the model’s rate doesn’t follow. Most of the pattern comes from the modern nights scattering more widely, a standard deviation of 1.04 km against 0.72 after allowing for the season, for reasons the paper doesn’t examine (checks added after our first run). And the model moves the whole layer down together, its lower edge staying 1.6 km below its peak [5 §3.2]. Our analysis. See section 4.

1 · WHAT THE PAPER DID

Twenty-two historical nights against forty-eight modern ones, fitted to a model’s curve

Noctilucent clouds are clouds of ice more than 80 km up, lit by the Sun in summer twilight after the ground has gone dark [1 §1][13 §1]. The preprint compares their heights as triangulated from 1890 to 1992 with about 20,000 heights from a network of cameras at 54–61°N, 29–50°E since 2020, and fits the gap with the shape of a model’s decline driven by carbon dioxide [1 Abstract][1 §4.1].

HISTORICAL TRIANGULATION

OUR ANALYSIS

22 nights

in 14 summers from 1890 to 1992, near Berlin, Oslo and Stockholm, in Estonia and Latvia, and in Scotland, counted from the paper’s Fig. 4 [1 §2][1 Fig. 4]

CAMERA NETWORK

OUR ANALYSIS

48 nights

averages over a night for one group of cameras, 2020 to 2026, 33 of them from the last two summers, and 21 more from the colour-based method [1 §3][1 Fig. 4]. In 2025 an exceptionally strong planetary wave shaped the season [4 Abstract]

THE FITTED TREND

K = 0.52 ± 0.18

the share of the model’s decline that fits the data: 1 would match the model, 0 would mean no trend [1 §4.1][1 Table 1]

How the old heights were chosen

Only “long-base” triangulations count: “the baseline must be S > H/2 = 40 km, and exact time synchronization is necessary”, since short baselines and clouds low over the horizon give large errors [1 §2]. The sources are Jesse’s measurements around Berlin in 1890–91, Størmer’s near Oslo in 1932–34, Swedish and Soviet ones in 1958–67, and Scottish ones late in the century [1 §2]. Each night’s measurements are weighted to add up to one, and where an old paper gives no sky position, the cloud is placed 2° north of the station [1 §2].

How the new heights are measured

About 20 cameras, mostly of the automatic Starvisor network, sit in groups 50–200 km across at 54–61°N, 29–50°E [1 §3]. Each cloud field is projected onto a surface at an assumed height and cut into fragments, and each fragment’s shift between cameras gives its height [1 §3]. A second method reads height from the clouds’ change of colour as they pass into the Earth’s shadow; the paper treats it as a check, and fits the trend on triangulation alone [1 §3].

The model they fit

The shape of the trend comes from the LIMA/MIMAS models of Lübken and colleagues. In their 2018 simulation, the clouds’ height, taken where they scatter most light back to a lidar, fell by about 1.3 km in July at 55–61°N between 1871 and 2008: 0.036 km a decade before 1960 and 0.233 km a decade after [5 §3.2][5 Fig. 3]. The preprint writes this as Eq. (3), a slow decline until 1960 and then one 6.5 times faster [1 §1][1 §4.1]. The fit scales that curve by K and adds terms for the season, local time, the lunar tide, latitude and sunspots. Only the trend and the season turn out significant [1 §5].

In the authors’ words NLC: noctilucent clouds.

> *“Comparison with published long-base triangulation measurements starting from 1890 reveals a descending trend in NLC altitude with a 99.5% confidence level.”*
>
> Ugolnikov et al. (2026), abstract [1 Abstract]

> *“The value of the long-term trend of noctilucent cloud altitude is 2.9 times higher than its error, which means that it is detected at better than the 99.5% confidence level.”*
>
> Section 5 [1 §5]

> *“However, the accuracy of the trend value will remain limited by the low number and precision of historical measurements.”*
>
> Section 6 [1 §6]

2 · OUR REBUILD

Rebuilt from the paper’s own figures, every number we could recompute comes out the same

The preprint publishes no table of its data. But its figures are drawn as vector paths, so each plotted night can be read from the file exactly and the fit run again [1 Fig. 4][1 Fig. 5].

STEP 1

Read the figures

Each marker in Figs. 4 and 5 is a separate shape in the PDF: a white circle for a historical night, a black one for a camera night, a ring for a colour-based one. Calibrated on the axes’ own tick marks, the two figures give every night’s height to within 0.003 km of each other, and Fig. 5 gives its place in the season [1 Fig. 4][1 Fig. 5].

STEP 2

Refit the model

With the paper’s model curve, Eq. (3), we fitted the 70 triangulation nights with the trend and the two seasonal terms, then again with the sunspot term, from SILSO’s daily numbers [1 §4.1][25]. The local-time, lunar and latitude terms need values for each night that the paper doesn’t publish; it finds all three insignificant [1 Table 1].

STEP 3

Compare

Every fitted value and both means land on the paper’s, to the printed digits or nearly, as the table shows. The rules for these tests were written down before any fit was run (section 8).

The paper beside our rebuild

QUANTITY	THE PAPER	OUR REBUILD OUR ANALYSIS
Height in 1880, H_0	82.8 ± 0.3 km	82.82 ± 0.34 km
Share of the model’s trend, K	**0.52 ± 0.18**	**0.520 ± 0.183**
Seasonal terms, A and B	−3.5 ± 0.9, 5.9 ± 1.2 km	−3.49 ± 0.89, 6.07 ± 1.28 km
Sunspot term, C	0.0013 ± 0.0019 km	0.0013 ± 0.0018 km
Mean before 2000	82.52 ± 0.16 km	82.51 ± 0.16 km, 22 nights
Mean since 2020	81.83 ± 0.18 km	81.83 ± 0.18 km, 48 nights
Difference	**0.69 ± 0.24 km**	**0.68 ± 0.24 km**
K with the colour-based heights	0.58 ± 0.18	0.56 ± 0.18 with the sunspot term (P10, added after our first run); 0.54 ± 0.18 without
Lowest point of the season	near 10 July	17 days after the solstice, about 8 July
Colour-based minus triangulated	−0.18 km	−0.16 km unweighted; −0.18 km weighted by the plotted error bars, as the paper finds (added after our first run)

The paper’s values are from its Table 1 [1 Table 1] and its sections 3 and 5 [1 §3][1 §5]. Ours come from the hosted re-fit, with the sunspot term for the first four rows; without it K is 0.503 ± 0.181. The last row is from its Fig. 3, the same nights measured both ways [1 Fig. 3].

OUR ANALYSIS

What the reproduction shows

The fit is computed correctly, and we found no error in the paper’s numbers. Its errors are those of an ordinary least-squares fit to the 70 nightly averages: their scatter about the fitted curve, 0.97 km, sets the ±0.18. So the question isn’t the arithmetic. It is what that error leaves out, which the next three sections take in turn.

What we don’t reproduce here

The preprint is under arXiv’s non-exclusive licence, so we redraw none of its figures and don’t republish the values read from them [1]. The hosted script reads them from the PDF again (section 8). The paper’s 20,000 single measurements aren’t public, so our fit uses its nightly averages, as its Figure 4 plots them [1 §5].

3 · HOW SURE?

The drop holds at about 95% under most choices; 99.5% needs the old and new methods to agree

The ±0.18 behind “99.5%” [1 §5] is, our refit shows, what the scatter of the 70 nightly averages around the fitted curve gives. We tested how the result moves when nights are grouped differently, parts are left out or the model changes: R-numbered tests follow rules written down before the first run, and P-numbered checks were added after it.

K, the share of the model’s decline, under each choice

OUR ANALYSIS

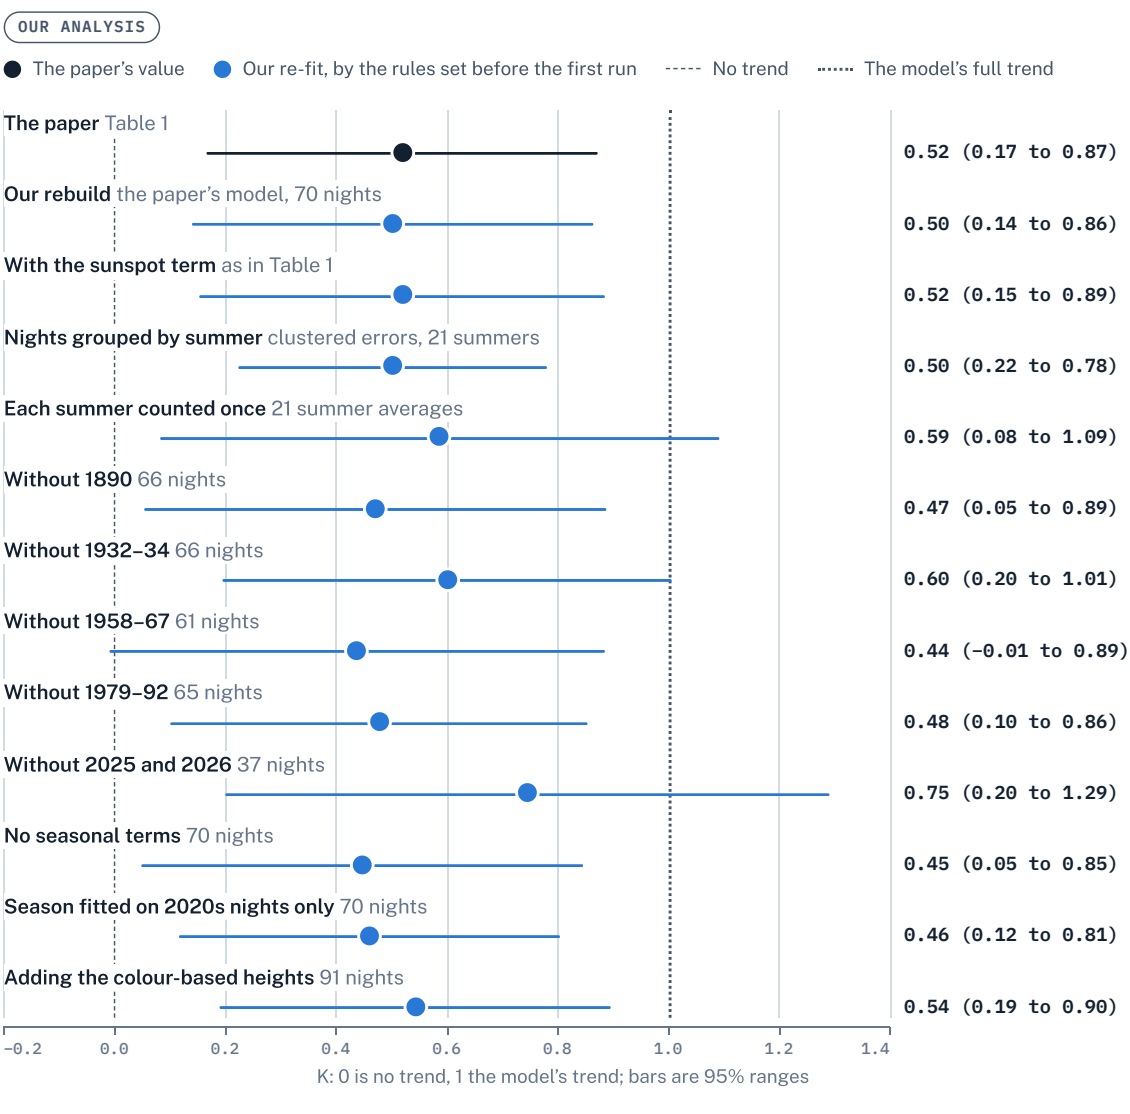

Our re-fit of the nightly averages read from Figs. 4 and 5 of [1], under the rules in section 8; the paper’s own value is from its Table 1 [1 Table 1]. The 95% ranges use Student’s t with each fit’s degrees of freedom, and 1.96 for the paper’s. “Without” rows leave out the nights of one era, or of 2025 and 2026, which supply 33 of the 48 modern nights.

The drop between the eras, in km, and what offsets between eras would do

OUR ANALYSIS

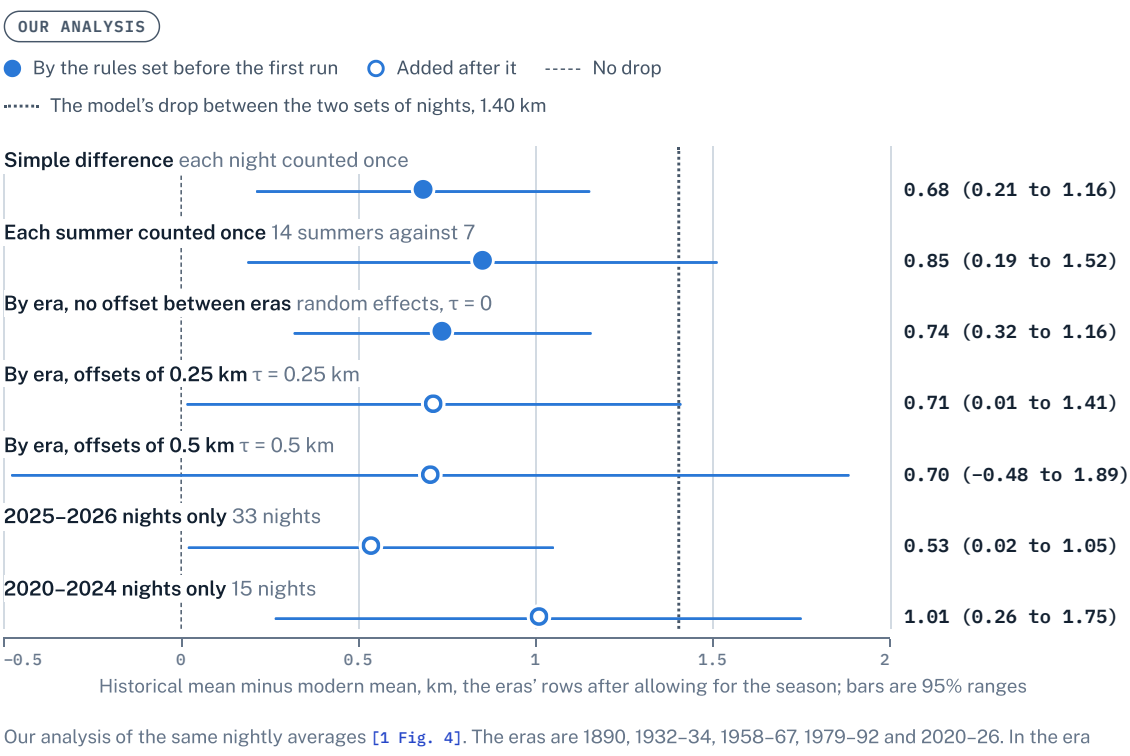

Our analysis of the same nightly averages [1 Fig. 4]. The eras are 1890, 1932–34, 1958–67, 1979–92 and 2020–26. In the era rows, each era, the modern network included, is given its own random offset of the size shown (τ), as a meta-analysis would treat separate studies; the four historical eras alone can’t pin that size down. The model’s 1.40 km is its curve, Eq. (3), averaged over each set of nights [1 §4.1]. The open rows, P3 and P6, were added after our first run.

OUR ANALYSIS

What “99.5%” means

At 2.9 standard errors, the chance of a gap this large if nothing had changed, and if the error told the whole story, is 0.4% on a normal curve, so “better than the 99.5% confidence level” holds [1 §5]; with Student’s t for 70 nights it is 0.52%, just short. It isn’t a 99.5% chance that the clouds descended. That depends on what the error leaves out, above all whether the old and new methods agree.

OUR ANALYSIS

Counting summers, not nights

Nights in one summer share that summer’s conditions aloft, so we tested whether they behave alike. Clustering the errors by summer, as our rules set, doesn’t widen them: K = 0.50 ± 0.13, p = 0.001. The nights show no sign that a summer’s nights are alike (P7: intraclass correlation −0.09, p = 0.75), though 21 summers can’t rule it out. If they shared offsets of 0.2–0.4 km, p would be 0.010–0.021 (P8). Counting each summer once, which gives a one-night summer the weight of a twenty-night one, gives p = 0.016–0.025. A straight line through all the years gives 0.80 ± 0.31 km a century (p = 0.012).

OUR ANALYSIS

Counting eras

The historical nights come from four eras, each with its own observers and instruments. Allowing for the season, their means, 82.90, 82.17, 82.74 and 82.50 km, all sit above the 2020s’ 81.87 km, and they agree with each other within their scatter (p = 0.49). Treated as separate studies, the eras give 0.74 ± 0.21 km, p = 0.0006. Counting each era as a single value gives p = 0.14, but that test would need a gap of 1.1 km to reach 95% (P9). Four eras also can’t show how far a whole era might be off: anything from 0 to 1.1 km fits them, and offsets of 0.25 or 0.5 km would give p = 0.045 or 0.24 (P3).

OUR ANALYSIS

How big a difference would undo it

If the camera network read 0.21 km lower than the old triangulations for the same clouds, the simple difference would fall to p = 0.05, and so would the fit of K (P1 and P11). On a normal curve the simple difference sits just inside the 0.5% that “99.5%” claims, and an offset of 6 m would take it past; with the Welch test our rules registered, it is already outside, p = 0.0062. Leaving out the 1958–67 nights gives K = 0.44 ± 0.22 (p = 0.055). The paper offers no test of whether the methods agree that closely; section 5 sets out why they might not.

The authors say as much about methods

“The basic difficulties are the natural scatter of NLC altitudes, exceeding the value of their mean depression over almost 150 years of observation, and the systematic shift of altitudes measured by different techniques, for example, lidars and photo imaging” [1 §6]. As an example they give the 1-km gap between the model’s height for the preindustrial clouds and the height their fit finds [1 §6]. The modellers had noted a similar gap against Jesse’s 1889–1891 mean, 82.08 km against their 83.8 km. Among their reasons: triangulating structures within a cloud gives heights well below the peak a lidar would see, and tides bias an observer who watches at fixed hours [5 §4]. The preprint’s answer is to use triangulation throughout; the next sections ask whether that is enough.

4 · SEASON AND SPREAD

The clouds are lowest in early July, and the modern nights scatter more widely than the old

The season matters as much as the decades: the clouds are higher early in the summer and sink to a low a few weeks after the solstice, as the paper and earlier work find [1 §4.2].

LOWEST CLOUDS

OUR ANALYSIS

≈ 8 July

17 days after the solstice in our refit; the paper says near 10 July [1 §5]

SPREAD OF THE NIGHTS

OUR ANALYSIS

1.26 km

standard deviation of the modern nightly averages, against 0.75 km for the historical ones; 1.04 against 0.72 km after allowing for the season (P2, added after our first run)

LOWEST TENTH OUR ANALYSIS

−1.5 km

the change in the lowest tenth of nightly averages, allowing for the season, against −0.4 km for the median

What others see through the season

Lidar at ALOMAR, in northern Norway, finds the clouds’ height varying through the season by up to 1–1.5 km [8 Abstract], and their altitude and thickness falling as the season goes on [14 Abstract]. SOFIE, on NASA’s AIM satellite, finds the clouds’ peak furthest below the mesopause, the coldest level, in mid-season [23 Abstract]. In the authors’ 2020–2024 data, every cloud seen before or at the solstice was above 82 km [3 §5].

OUR ANALYSIS

The season and the trend

The modern nights start earlier in the summer: four fall on or before 10 June, when the clouds are high, and no historical night is earlier than 21 June. The seasonal terms, fitted mostly on modern nights, deal with that, and the result doesn’t hinge on them. Without them K is 0.45 ± 0.20 (p = 0.03); with the season fitted on the 2020s alone, 0.46 ± 0.17. The authors’ 2020–24 pattern of high early clouds no longer holds strictly: three camera nights of 2025–26 before the solstice average below 82 km.

The top and bottom of the range, allowing for the season

NIGHTLY AVERAGE	1890–1992, 22 NIGHTS	2020–2026, 48 NIGHTS	CHANGE
Highest	83.81 km	83.88 km	+0.07 km
Highest tenth	83.62 km	83.22 km	−0.40 km
Median	82.43 km	82.03 km	−0.39 km
Lowest tenth	**81.87 km**	**80.34 km**	**−1.53 km**
Lowest	81.17 km	79.96 km	−1.21 km

Our analysis. Heights are moved to the solstice with the seasonal terms of our refit, from the nightly averages in Figs. 4 and 5 of [1]. The tenths are the 10th and 90th percentiles.

OUR ANALYSIS

Why the tails can’t carry the conclusion

Most of the pattern follows from the modern nights scattering more widely. On a normal curve, the wider spread and the 0.74 km shift of the mean predict a fall of 1.15 km in the lowest tenth and 0.33 km in the highest, against the 1.53 and 0.40 km seen (P12). Sample size matters only for the single lowest and highest nights, not for the tenths. The wider spread is unlikely to be chance before allowing for the season (p = 0.01) and less clearly so after (p = 0.07, P2). It may be real variety in the clouds, or the way they are measured. On ten dates, two or three modern averages exist for the same date, from different camera groups or twilights, and they differ by a median of 1.1 km, once by 2.9 km (P4). The colour-based and triangulated averages for the same nights differ with a spread of 0.83 km [1 Fig. 3]. These checks were added after our first run. The paper reads the lower tail as clouds sinking at the model’s rate [1 §6]; it doesn’t test the alternatives. And the model it cites moves the whole layer down together, its lower edge staying about 1.6 km below its peak [5 §3.2]. It says nothing about the spread of nightly averages.

5 · OLD AND NEW METHODS

The paper doesn’t measure the same clouds both ways, and an offset of about 0.2 km between them would undo the result

The paper rests the comparison on using one technique throughout: “To perform the correct trend estimation, we must analyze results from similar measurement techniques” [1 §1]. It applies the same baseline rule to old and new data [1 §2], but it doesn’t show that the two give the same height for the same clouds. The 0.2 km is our tipping-point check in section 3, added after our first run.

Then and now

	1890–1992	2020–2026	WHAT IT COULD DO TO THE COMPARISON OUR ANALYSIS
Where	Near Berlin, Oslo and Stockholm, in Estonia and Latvia, and in Scotland [1 §2]	54–61°N, 29–50°E [1 §3]	Little. Lidars find heights rising with latitude by only 43 ± 65 m a degree [8 Abstract], and the paper’s fit gives 0.01 ± 0.06 km a degree [1 Table 1].
When in the summer	21 June to 10 August	1 June to 12 August	Handled by the seasonal terms; the result holds without them (section 4).
Time of night	Not given in the paper	Every 2–3 minutes while the clouds are visible [1 §3]	The clouds’ heights swing with the tides through the day [13 §4], which can bias an observer who watches at fixed hours [5 §4]. The paper’s local-time term is 0.05 ± 0.08 km an hour [1 Table 1].
How	Photographs from stations 40–200 km apart, mostly of clouds close to the horizon [1 §2]	Automated cameras; fragments matched by correlation from a surface at an assumed height, 81.33 km in the 2025 paper [1 §3][3 §3]	Unknown. The paper doesn’t measure the same clouds both ways, and a search of the literature index found no study that does.
Which clouds	The brightest visible fragments [1 §1]	Bright events in 2020–2024 [3 Abstract]; fragments centred on brightness peaks [3 §3]	Unknown in size and direction.
How many	22 nights that pass the baseline rule, out of a larger record [1 §2]	48 camera-group nights	The excluded measurements aren’t listed.

Dates from our reading of Fig. 5 of [1], and the last column, are our analysis. The paper’s 1986 measurements had the longest baseline of the century, about 200 km [1 §2].

OUR ANALYSIS

Choices that move a mean by as much as the trend

The authors’ own work shows how much the way heights are averaged matters. In their 2020–2024 data, weighting fragments by area gave a mean of 81.4 km, weighting by brightness 80.9 km, and the brightest clouds sat at 80.1 km [3 §5]. Their colour-based heights read about 0.5 km above triangulation until a correction in this preprint [1 §3]. Lidars see the layer 1–1.5 km higher than their triangulation does [3 §4]. Next to shifts like these, 0.69 km is not large, and the old observers’ choices of which fragments to measure can’t be recovered.

OUR ANALYSIS

Which way would a bias go?

Either way. If the old heights were of the brightest, and so lowest, parts of the clouds, while the new ones average whole fields, the old values would read low and hide part of a real descent. If the new pipeline picks out lower parts of the clouds than the old observers did, the new values would read low and make a descent that isn’t there. The paper tests neither. Modern images measured the old way, or old plates measured the new way, would. The authors saw the risk: their 2025 paper called the “change of measuring techniques” a “Possible problem in long-term trend study of NLC” [3 §5], and the preprint’s conclusion speaks of “a confidence level better than 99%” [1 §6], more cautiously than its abstract.

OUR ANALYSIS

The record behind the 22 nights

The old record is bigger than the paper’s sample. The authors’ 2025 paper gives Jesse’s own mean height as 82.1 ± 0.1 km, a median “close to 83 km” for the 20th-century triangulations reviewed by Gadsden and Schröder, and 695 measurements from 1887 to 1964 summed up in a Russian book by Bronshten and Grishin [3 §1]. The preprint keeps the nights that pass its baseline rule and doesn’t list the rest [1 §2]; its four nights from 1890 average 82.5 km by our reading of its Fig. 4. Different selections from the same old record differ by 0.4–0.5 km, more than half the drop.

6 · EARLIER WORK

Earlier reviews found no change, but they couldn’t have seen one this small, and lidar and satellites probably can’t yet

For decades the clouds’ heights were cited as unchanged since the 1880s. The preprint sets out why it reaches a different answer [1 §1].

No change, said earlier reviews

A review of triangulated heights from 1889 to 1979, and a 1994 update, found no significant trend, as the preprint reports them [1 §1][12]. Von Zahn’s 2003 review put the heights between 82.2 and 83.6 km and ascribed the differences to latitude [1 §1][10]. Thomas and colleagues described the heights as ones “which have apparently remained constant” [11 Abstract], and Fiedler and colleagues wrote in 2011 that they seemed unchanged for more than 100 years [13 §1]. Lübken and colleagues applied a model link between temperature and cloud height to the historical heights and found “negligible temperature trends at PMC altitudes (~0.01–0.02 K/yr)” [7 Abstract].

OUR ANALYSIS

Why that doesn’t contradict the new result

The reviews ended around 1990. On the model’s curve, a trend half the model’s would have lowered the clouds by about 0.35 km from 1960 to 1990, too little to see in a few campaigns [1 §4.1]. And von Zahn’s comparison, as the preprint describes it, set old triangulations beside modern lidars, which find the clouds’ centre at 82.75 km at 54°N, about 1 km above the preprint’s triangulated 81.8 km [8 Abstract][1 §5]. The new result rests on the 2020s, after six decades of the model’s faster decline [5 §3.2]. At half the model’s rate, 25 years of lidar would see the clouds sink by about 0.3 km, less than their swing through one season [8 Abstract].

What the models expect

The model the paper fits simulates 138 years of clouds at 55–61°N, with winds and waves taken from one representative year [5 §2]. In it the clouds sank about 1.3 km, most of it after about 1960, and only in the runs where carbon dioxide rises: as the air below cools and shrinks, the level where ice can form moves down [5 §3.2]. At a fixed 83 km the air cooled by about 7 K, but on the clouds’ pressure level hardly at all [5 §3.1]. Methane’s extra water vapour barely moved the clouds but made them much brighter [5 §4]. A 2021 study with the same models gives about 260 m a decade after 1960, at every latitude [6 Abstract], close to the 2018 study’s 0.233 km [5 Fig. 3]. A study of 1961–2013 ties the height trend largely to cooling caused by stratospheric ozone [9 Abstract]; the 2018 runs include ozone’s changes after 1961 along with carbon dioxide’s, so the curve the preprint fits carries both [5 §2].

What lidar and satellites can say

Lidar at ALOMAR has watched the clouds since 1997, and its seasonal mean heights move with solar activity and temperature [13 Abstract]; the study of its long-term record is closed to us and has no public abstract [27]. The Solar Mesosphere Explorer satellite found no change in height from year to year in 1981–85 [22 Abstract], and SOFIE measured heights from 2007 [23 Abstract]. The longest satellite record, from SBUV instruments since 1979, gives the clouds’ brightness and ice, not their height [24 Abstract]. Ground-based records from the Moscow region found trends in NLC parameters over 1968–2018 small and insignificant, which their authors say contradict “large modeled NLC trends” [21 Abstract]; the abstract reports no heights.

OUR ANALYSIS

Does the paper engage with them?

With the triangulation reviews, yes: it is built as an answer to them [1 §1]. It doesn’t discuss the later model studies of the height trend beyond its 2018 source, or the satellite measurements of the shrinking atmosphere in section 7, which its authors’ 2025 paper did cite [3 §1]. Nor does it use any lidar or satellite record of the clouds’ heights over the years; its conclusion expects progress from such records as they lengthen [1 §6].

7 · THROUGH THE CATASTROPHISM LENS

Is the upper atmosphere shrinking? Yes, slowly and as predicted; it isn’t falling

A shrinking or falling upper atmosphere is a catastrophist image. This preprint bears on it directly: as the air below the clouds cools, it contracts, and the preprint expects that to show as falling cloud heights [1 §1].

SABER, 2002–2019

1,333 m

contraction of the layer between about 48 and 105 km, 342 m of it put down to carbon dioxide [18 Abstract]

POLAR SUMMER MESOPHERE

100–200 m

of shrinking a decade, from three satellite instruments combined over two to three decades [17 Abstract]

THE PREPRINT’S RATE

OUR ANALYSIS

≈ 1 m a month

12 m a year, or 1.2 km a century, if its K and the model’s timing are right [1 Abstract]

What has been measured

Cooling of the upper atmosphere by rising carbon dioxide and methane was predicted in 1989 [20 Abstract]. Satellites now see it: the polar summer mesosphere cools by about 1–2 K a decade and shrinks by 100–200 m a decade [17 Abstract]. The 2018 model study found its clouds’ descent consistent with radio-wave reflection heights observed since 1959 [5 §4], as shown in Berger and Lübken’s 2011 paper [26]. A 2023 review concludes that “cooling through mesospheric shrinking is responsible for a slight decrease in NLC heights” [19 Abstract].

OUR ANALYSIS

What it means

So the answer is yes, but slowly. At the preprint’s rate the clouds sink about a metre a month, a pace in line with the satellite measurements of the shrinking layer. Over a decade or two the Sun moves these levels more than carbon dioxide does: most of SABER’s 1,333 m came with the fading of the solar cycle [18 Abstract], and model clouds sit higher at solar maximum [16 Abstract]. The cooling aloft is a long-predicted greenhouse fingerprint, about ten times the surface’s warming and opposite in sign [15 §3.1]. Nothing in the preprint or the literature points to a sudden fall. The clouds’ brightening is a separate change, mostly from extra water vapour made from methane [5 §4].

8 · REPEAT OUR NUMBERS

One script fetches the preprint and the sunspot numbers, reads the figures and reruns every test

The rules came first, with their hash; what came after them is labelled in the log. Run from an empty folder, the script gave our outputs byte for byte.

- PY **nlc_pipeline.py** Fetches the preprint’s PDF from arXiv and SILSO’s daily file by URL, checks the PDF against our copy’s hash, runs every step and compares each output with our log; standard library only · 5 KB
- PY **pdf_objects.py, dump_page.py, vector_shapes.py and extract_points.py** Read the PDF as data, without running anything in it, and find the markers of Figs. 4 and 5 · 5 KB, under 1 KB, 4 KB and 4 KB
- PY **refit.py** The re-fit and the tests R1–R10 · 22 KB
- PY **posthoc.py, posthoc2.py and fig3_shift.py** The checks added after the first run, P1–P6, those added after the fresh review, P7–P12, and Fig. 3’s shift · 7 KB, 6 KB and 3 KB
- PY **chart_rows.py** The two charts in section 3 · 6 KB
- TXT **RULES_refit.txt** The rules, written before any fit was run · 5 KB
- TXT **refit-log.txt** The SHA-256 of the rules, the PDF, the values read from it and every output, in the order they were made, ending with these hosted copies · 9 KB
- JSON **refit-results.json, posthoc-results.json and posthoc2-results.json** The results of R1–R10 and P1–P12: derived numbers, without the nightly values read from the figures, and without the heights of the nine summers that had a single night · 12 KB, 3 KB and 2 KB
- JSON **fig3-shift.json and chart-rows.json** Fig. 3’s shift, and the charts’ numbers · under 1 KB and 5 KB

Run it

`python3 -I nlc_pipeline.py <folder>` , with the other scripts beside it, downloads the preprint (246 KB) and SILSO’s file (2.8 MB) into that folder, checks the PDF, and writes every output to its `out` folder, then compares each with our log. Once the files are there it takes a few seconds. SILSO adds a day to its file every day and sometimes revises recent days, which can change only the sunspot run, R1b. The full outputs include the values read from the preprint’s figures, for checking, not for republishing.

The data and their terms

The preprint: arXiv’s non-exclusive licence, so we host neither its PDF nor the values read from its figures; the script fetches the one and reads the other again [1]. **Sunspot numbers:** WDC-SILSO, Royal Observatory of Belgium, Brussels, under CC BY-NC 4.0; the script fetches them [25].

9 · LIMITS

What we couldn’t check

The old heights themselves

We read them from the preprint’s figures, which match its own means. We didn’t check them against Jesse, Størmer or the other original publications, and haven’t seen which measurements its baseline rule left out [1 §2].

Terms we couldn’t refit

The local-time, lunar-tide and latitude terms need values for each night that the paper doesn’t publish. It finds all three insignificant [1 Table 1].

The model’s curve

The rates in the preprint’s Eq. (3) are the 2018 model’s fitted slopes before and after 1960, which we read in its Fig. 3a [5 Fig. 3]. Joining them into one broken line from 1880 is the preprint’s own simplification [1 §4.1]. The model’s height is that of the clouds’ peak backscatter, with winds and waves held at one year’s pattern [5 §2][5 §3.1].

Records over time

The longest lidar record, ALOMAR’s, has a long-term study that is closed and has no public abstract, so we couldn’t read it [27]. The other records we found are too short to test a descent of about 0.1 km a decade. We read several of the studies above only as abstracts, as the source list marks.

Unpublished data

The 20,000 single measurements, and how the modern heights depend on the assumed height the projection starts from, 81.33 km in the 2025 paper [3 §3], aren’t public.

What would change these verdicts

01 **A cross-check of methods.** Modern images measured the way the old observers measured theirs, or old plates measured by the new pipeline, agreeing within about 0.2 km, the offset our post-hoc check in section 3 finds would undo the result. That would do more for the 99.5% than more modern nights.

02 **The historical table.** A published list of the old heights, with the measurements the baseline rule excluded, so that others can test the selection.

03 **An independent record.** Two more decades of camera, lidar or satellite heights showing the descent, or a peer review that confirms the analysis.

The authors write that “Significant progress is expected when long-term arrays of lidar, satellite, and imaging data during the 21st century become available” [1 §6].

Related checks

MISLEADING Mesospheric echoes and CO2 The Disaster Cycle, Fact-Checked

Sources

Bracketed numbers refer to this list. For the preprint [1], read in full as arXiv version 1, § gives a section, Fig. a figure, Table its Table 1 and Eq. an equation, and Abstract its abstract. For the authors’ earlier paper [3], read in its arXiv version, § gives a section and Abstract its abstract. For [5], [13], [15] and [19], read in full, § gives a section and Fig. a figure. Abstract marks the part of a source we rely on; a source marked “Abstract only” or “Record only” was read no further. NLC are noctilucent clouds; PMC, polar mesospheric clouds, are the same clouds seen from space. K is the share of the model’s decline that the preprint fits. Links checked November 6, 2026.

1 Ugolnikov, O. S., Yankovsky, I. S., Klimenko, M. V., Ugolnikova, E. O., Shtanov, A. M., and Zheltoukhov, S. G. (2026). Long-Term Trend of Noctilucent Cloud Altitude Detected: A Comparison of Positional Data with Historical Measurements. arXiv:2610.41750v1 [physics.ao-ph], submitted October 30, 2026, 14 pages, 5 figures. Not peer reviewed; on November 6, 2026, arXiv shows no journal reference. Under arXiv’s non-exclusive distribution licence, so no figure is redrawn here and the values we read from its figures are not republished. Full text.

2 SpaceWeatherNews (S0) (November 6, 2026). Sky Fall Continues, Solar Storms Linger for Days | S0 News Nov.6.2026. YouTube video, https://www.youtube.com/watch?v=VRbM1M44nik, published at 12:11 UT. No transcript was available.

3 Ugolnikov, O. S., Pertsev, N. N., Perminov, V. I., Yankovsky, I. S., et al. (2025). Five-years altitude statistics of noctilucent clouds based on multi-site wide-field camera survey. *Journal of Atmospheric and Solar-Terrestrial Physics* 269, 106491. Peer reviewed. Read in full as its arXiv version, arXiv:2412.04951v1 (December 2024), under arXiv’s non-exclusive licence.

4 Ugolnikov, O. S., Yankovsky, I. S., Pertsev, N. N., et al. (2026). Noctilucent clouds modulated by strong 5-day planetary wave in 2025: amplitudes, phases and altitudes based on ground-based observations and satellite temperature data. *Advances in Space Research* 77(12), 12466–12475. Abstract only, read in its arXiv version, arXiv:2511.20089v2.

5 Lübken, F.-J., Berger, U., and Baumgarten, G. (2018). On the anthropogenic impact on long-term evolution of noctilucent clouds. *Geophysical Research Letters* 45(13), 6681–6689. The source of the model curve the preprint fits. Full text, supplied by the site’s owner after the publisher’s site answered our own request with a bot check, which we do not attempt.

6 Lübken, F.-J., Baumgarten, G., and Berger, U. (2021). Long term trends of mesopheric ice layers: A model study. *Journal of Atmospheric and Solar-Terrestrial Physics* 214, 105378. CC BY-NC-ND 4.0. Abstract only, read through an index (OpenAlex), so paraphrased.

7 Lübken, F.-J., Berger, U., and Baumgarten, G. (2009). Stratospheric and solar cycle effects on long-term variability of mesospheric ice clouds. *Journal of Geophysical Research* 114. Abstract only.

8 Lübken, F.-J., Baumgarten, G., Fiedler, J., Gerding, M., Höffner, J., and Berger, U. (2008). Seasonal and latitudinal variation of noctilucent cloud altitudes. *Geophysical Research Letters* 35(6). Abstract only.

9 Berger, U., and Lübken, F.-J. (2015). Trends in mesospheric ice layers in the Northern Hemisphere during 1961–2013. *Journal of Geophysical Research: Atmospheres* 120(21). Abstract only.

10 von Zahn, U. (2003). Are noctilucent clouds a “Miner’s Canary” for global change? *Eos, Transactions American Geophysical Union* 84(28), 261–264. Abstract only; its review of heights is described as the preprint describes it.

11 Thomas, G. E., Olivero, J. J., DeLand, M., and Shettle, E. P. (2003). Comment on “Are noctilucent clouds truly a ‘Miner’s Canary’ for global change?” *Eos, Transactions American Geophysical Union* 84(36), 352–353. Abstract only.

12 Gadsden, M., and Taylor, M. J. (1994). Measurements of noctilucent cloud heights: a bench mark for changes in the mesosphere. *Journal of Atmospheric and Terrestrial Physics* 56(4), 461–466. Record only; described as the preprint describes it.

13 Fiedler, J., Baumgarten, G., Berger, U., Hoffmann, P., Kaifler, N., and Lübken, F.-J. (2011). NLC and the background atmosphere above ALOMAR. *Atmospheric Chemistry and Physics* 11, 5701–5717. CC BY 3.0. Full text.

14 Fiedler, J., Baumgarten, G., and von Cossart, G. (2003). Noctilucent clouds above ALOMAR between 1997 and 2001: Occurrence and properties. *Journal of Geophysical Research* 108(D8). Abstract only.

15 Vellalassery, A. (2024). Long-term trends and solar cycle response of noctilucent clouds. PhD thesis, University of Rostock. CC BY 4.0, as its repository record states. Full text.

16 Vellalassery, A., Baumgarten, G., Grygalashvyly, M., and Lübken, F.-J. (2024). Long-Term Evolution in Noctilucent Clouds’ Response to the Solar Cycle: A Model-Based Study. *Atmosphere* 15(1), 88. CC BY 4.0. Abstract only: the publisher refused our request for the PDF.

17 Bailey, S. M., Thurairajah, B., Hervig, M. E., Siskind, D. E., Russell, J. M., and Gordley, L. L. (2021). Trends in the polar summer mesosphere temperature and pressure altitude from satellite observations. *Journal of Atmospheric and Solar-Terrestrial Physics* 220, 105650. CC BY 4.0. Abstract only, read through an index (OpenAlex), so paraphrased.

18 Mlynczak, M. G., Hunt, L. A., Garcia, R. R., Harvey, V. L., Marshall, B. T., Yue, J., Mertens, C. J., and Russell, J. M. (2022). Cooling and contraction of the mesosphere and lower thermosphere from 2002 to 2021. *Journal of Geophysical Research: Atmospheres* 127(22), e2022JD036767. CC BY-NC-ND 4.0. Abstract only.

19 Laštovička, J. (2023). Progress in investigating long-term trends in the mesosphere, thermosphere, and ionosphere. *Atmospheric Chemistry and Physics* 23, 5783–5800. CC BY 4.0. Full text.

20 Roble, R. G., and Dickinson, R. E. (1989). How will changes in carbon dioxide and methane modify the mean structure of the mesosphere and thermosphere? *Geophysical Research Letters* 16(12), 1441–1444. Abstract only.

21 Dalin, P., Perminov, V., Pertsev, N., and Romejko, V. (2020). Updated long-term trends in mesopause temperature, airglow emissions, and noctilucent clouds. *Journal of Geophysical Research: Atmospheres* 125(5), e2019JD030814. Abstract only.

22 Thomas, G. E., and Olivero, J. J. (1986). The heights of polar mesospheric clouds. *Geophysical Research Letters* 13(13), 1403–1406. Abstract only.

23 Li, Y., Gao, H., Sun, S., and Li, X. (2024). Correlation between Peak Height of Polar Mesospheric Clouds and Mesopause Temperature. *Atmosphere* 15(10), 1149. CC BY 4.0. Abstract only.

24 DeLand, M. T., and Thomas, G. E. (2015). Updated PMC trends derived from SBUV data. *Journal of Geophysical Research: Atmospheres* 120(5), 2140–2166. Abstract only, read through an index (OpenAlex), so paraphrased.

25 WDC-SILSO, Royal Observatory of Belgium, Brussels. Sunspot Number Version 2.0: the daily total sunspot number, file SN_d_tot_V2.0.csv, retrieved November 6, 2026. CC BY-NC 4.0. Read in full as data.

26 Berger, U., and Lübken, F.-J. (2011). Mesospheric temperature trends at mid-latitudes in summer. *Geophysical Research Letters* 38(22). Record only; its comparison with radio reflection heights is cited as [5] describes it.

27 Fiedler, J., Baumgarten, G., Berger, U., and Lübken, F.-J. (2017). Long-term variations of noctilucent clouds at ALOMAR. *Journal of Atmospheric and Solar-Terrestrial Physics* 162, 79–89. Record only: closed, with no abstract in Crossref or OpenAlex.

METHOD

How this was checked

On November 6, 2026, we read the preprint in full, with its authors’ earlier paper on the same camera network [3], and read the values plotted in its Figures 3, 4 and 5 from the vector drawing in its PDF: 91 nightly averages, 22 historical, 48 by camera triangulation and 21 by colour. Four other sources were read in full, those marked so in the list, [5], [13], [15] and [19], and 19 more as abstracts or records. The full text of [5] was supplied by the site’s owner after Wiley’s site refused our script; MDPI refused our request for [16], and we didn’t retry. Every quotation was checked word for word against its saved source.

The verdicts and the items tagged “Our analysis” are our own assessment, built on the cited sources. They cover how the paper’s trend and its 99.5% hold up when nights are grouped by summer or by era, when parts are left out or the model changes, how large a difference between old and new methods would undo the result, how widely the nightly averages scatter, and what the literature says about these clouds’ heights and the shrinking upper atmosphere. No transcript of the S0 News episode that linked the preprint was available [2], so this page reviews the preprint on its own and rates nothing from the episode.

We wrote the rules of our re-fit after reading the plotted values and before running any fit, and logged the SHA-256 of the rules, the data and each output; checks added after the first run, and after a fresh review of this page, are labelled as such where the page uses them. That review found a bug in our reading of Figure 3’s error bars, which we fixed before publication; with it fixed, our weighted shift matches the paper’s −0.18 km. The re-fit reproduces the preprint’s Table 1, its means before 2000 and in the 2020s, and their difference. One hosted script fetches the preprint and the sunspot numbers by URL and recomputes everything; run from an empty folder, it gave our outputs byte for byte. We could not check the historical heights against their original publications, which we didn’t read, nor the paper’s local-time, latitude and lunar terms or its 20,000 single measurements, which aren’t published.

This page was researched and written with Claude, an AI assistant made by Anthropic, and then re-checked against its sources in a separate verification pass.

The verdicts on this page

OVERSTATED	Has some basis but claims much more than the evidence supports.
UNCONFIRMED	Rests on a preliminary result that hasn’t been confirmed.
PARTLY TRUE	Part of the claim holds up and part doesn’t.
TRUE	The evidence supports the claim.

All ten verdicts, and how we check

CORRECTIONS

Corrections and changes

If we get something wrong, we fix it and record the change below with the date. If a correction changes a verdict, the log says what the old verdict was.

Change log

Nov 6, 2026 Published on factchecked.science after a full check against the sources. Built from the preprint, the values plotted in its own figures, our pre-registered re-fit and the literature; the two charts are our own analysis, and no figure of the preprint is redrawn.

Fact-Checked Science is independent. It isn’t affiliated with the channels, presenters or authors whose claims it checks.

Last updated Nov 6, 2026