National Institute for Public Health and the Environment
Ministry of Health, Welfare and Sport

RIVM

A. van Leeuwenhoeklaan 9
3721 MA Bilthoven
PO Box 1
3720 BA Bilthoven
www.rivm.nl/en

T 088 689 89 89

Authors:
M.R. Lechner
J. Wesseling
Centre:
OLG
Contact:
michelle.lechner@rivm.nl
Reference:
KN-2026-0041
DOI:
10.21945/RIVM-KN-2026-0041
Date:
5 juni2026

Knowledge brief

Improving air quality forecasting visualisations

Summary

To further enhance the uptake of CAMS[1] products in the Netherlands, RIVM[2], TNO[3], KNMI[4] and NLSA[5] are collaborating in the **CAMS National Collaboration Programme** (**NCP**) to stimulate the uptake and adaptation of CAMS global and European-scale products, better informing official national environmental information systems and services. Part of the project (WP80) is a so-called "demonstrator" for the health user communities, which aims to link the CAMS information on air quality with improved publicly accessible information on health in the form of air quality forecasts. This knowledge brief describes the current progress and status of the demonstrator. Following a stakeholder workshop in June 2025, where user input was collected, two mock-up visualisations have been developed: a map-based forecast tool with a time slider (visualisation A) and a location-based timeline overview (visualisation B). Both are currently online and being tested with users. Feedback collected so far has already led to several improvements, and further refinement of the visualisations is ongoing.

Introduction

Within the CAMS National Collaboration Programme (NCP), RIVM is working on improving the communication of air quality forecasts for the health sector (WP80), in collaboration with TNO and KNMI. The project runs over three years and consists of three phases: user requirements analysis and design, development of the demonstrator, and exploitation and analysis.

The core purpose of the project is to develop better visualisations of air quality forecasts. To ensure these visualisations closely match the needs of end users, they are actively involved in the design process. This includes determining what information is most relevant to communicate (such as the selection of pollutants, the time horizon, and the preferred spatial resolution) and how this can be presented in a clear and actionable way. The intended audience (stakeholders) includes non-expert users such as individuals with respiratory conditions like asthma, as well as organisations like Longfonds and DCMR[6]. The underlying forecast data is based on model data available in CAMS, that has a resolution of several kilometres squared. A later version of the visualisations will use a version of the LOTOS-EUROS model developed by TNO, that can be downscaled to a higher spatial resolution of 1x1 km2.

To achieve better visualisations of air quality forecasts, stakeholders are actively involved throughout the process. At the start, a workshop was organised to gather input from citizens and representatives from relevant

[1] Copernicus Atmosphere Monitoring Service
[2] Rijksinstituut voor Volksgezondheid en Milieu
[3] Nederlandse Organisatie voor Toegepast-Natuurwetenschappelijk Onderzoek
[4] Koninklijk Nederlands Meteorologisch Instituut
[5] Netherlands Space Agency
[6] Dienst Centraal Milieubeheer Rijnmond

organisations on how they prefer to receive air quality forecasts. Based on these outcomes, two mock-up websites have been developed and are currently online. User feedback on these mock-ups is continuously being collected to improve the visualisations in an iterative way.

Workshop at the start

On the 18th of June 2025, RIVM hosted a stakeholder workshop as part of the first phase of the WP80 Health demonstrator. The workshop brought together participants from a range of backgrounds, including citizens, representatives from community air quality networks such as Hollandse Luchten and the SamenMeten community, Longfonds, and the regional environmental service DCMR. The goal was to gather input on how air quality forecast information can best be presented to non-expert users.

During the workshop, participants were first introduced to the current state of air quality forecasting in the Netherlands and the types of data available, both publicly through luchtmeetnet.nl and information that is available internally within RIVM. They were then asked to reflect on a series of questions: what kind of information is most useful, how far ahead should forecasts look, whether to present individual pollutants or a single air quality index, and how to handle uncertainty in a way that is informative rather than confusing.

A central exercise asked participants to sketch or describe their preferred way of presenting an air quality forecast. Working in small groups, they produced a range of ideas. Despite the variety, a clear common thread emerged: participants strongly preferred a simple, intuitive dashboard. The starting page should be immediately clear, show information for the user's own location, and include direct guidance on what the forecast means for their daily activities. More detailed information (such as graphs showing forecast uncertainty or breakdowns by individual pollutant) was seen as useful, but only as a second layer, accessible on request rather than shown upfront.

The workshop outputs pointed towards two distinct approaches to presenting air quality forecasts, which can broadly be described as a map-based view (‘buienradar’ style) and a timeline view.

Visualisation A (‘buienradar’)

The first concept (visualisation A) resembles the way weather radar applications like Buienradar present information: a spatial map of the Netherlands showing air quality across the country, with a time slider that allows users to scroll forward through the forecast. The starting view shows the entire country with a colour overlay indicating air quality at a given moment. Users can select the air quality index (LKI) or a single pollutant (NO2, O3, PM10, or PM2.5), enter a postcode to zoom to a specific location and use a time slider or play button to animate the forecast per 3-hour time step. Clicking on the map or entering a postcode shows the forecast value for that specific location. The colour scale runs from blue (good) through yellow and orange to red and purple (very poor), consistent with the existing Dutch air quality index categories that are also used at luchtmeetnet.nl.

Figure 1: Color scale for the LKI and individual concentrations, see https://www.luchtmeetnet.nl/informatie/luchtkwaliteit/luchtkwaliteitsindex-(lki)

LUCHTKWALITEIT	LKI [index]	NO_2 [µg/m³]	O_3 [µg/m³]	PM10 [µg/m³]	PM2.5 [µg/m³]
GOED	**0-1**	0-10	0-15	0-10	0-10
	1-2	10-20	15-30	10-20	10-15
	2-3	20-30	30-40	20-30	15-20
MATIG	**3-4**	30-45	40-60	30-45	20-30
	4-5	45-60	60-80	45-60	30-40
	5-6	60-75	80-100	60-75	40-50
ONVOLDOENDE	**6-7**	75-100	100-140	75-100	50-70
	7-8	100-125	140-180	100-125	70-90
SLECHT	**8-9**	125-150	180-200	125-150	90-100
	9-10	150-200	200-240	150-200	100-140
ZEER SLECHT	**>10**	>200	>240	>200	>140

Technical implementation

The visualisation is built as a single HTML file using several open-source libraries. The map itself is rendered using Leaflet, with background tiles provided by PDOK's grey topographic map (BRT Achtergrondkaart). The coordinate system used is the Dutch RD projection (EPSG:28992).

The underlying forecast data presently comes from the Finnish SILAM model that is available through CAMS. The model produces air quality predictions at 3 hourly intervals for multiple pollutants: LKI, NO2, O3, PM10, and PM2.5. For the Netherlands, these predictions are available on a grid and are stored in ASC format. The user can switch between pollutants using a dropdown menu; each selection loads a new set of raster files.

When the application loads, it reads all available forecast files for the selected LKI or pollutant and builds a complete time series from this data For each time step, the grid values are translated into colours based on fixed concentration thresholds. This coloured grid is then drawn onto the map as a transparent image layer on top of the background map of the Netherlands. By moving the time slider or pressing play, the image layer is replaced with the next one, creating the impression of an animation.

When a user selects a location (either by clicking on the map or entering a postcode, which is looked up via the PDOK location API) the application identifies which grid cell corresponds to that location and reads the predicted value directly from the data. In addition, a line chart appears showing the predicted values for all forecast hours at that location, making it easy to see how air quality is expected to develop over time. The highlighted point in the chart corresponds to the timestamp currently shown on the map and moves along as the user steps through the time slider.

Figure 2 shows a screenshot of visualisation A.

Figure 2: screenshot of visualisation A

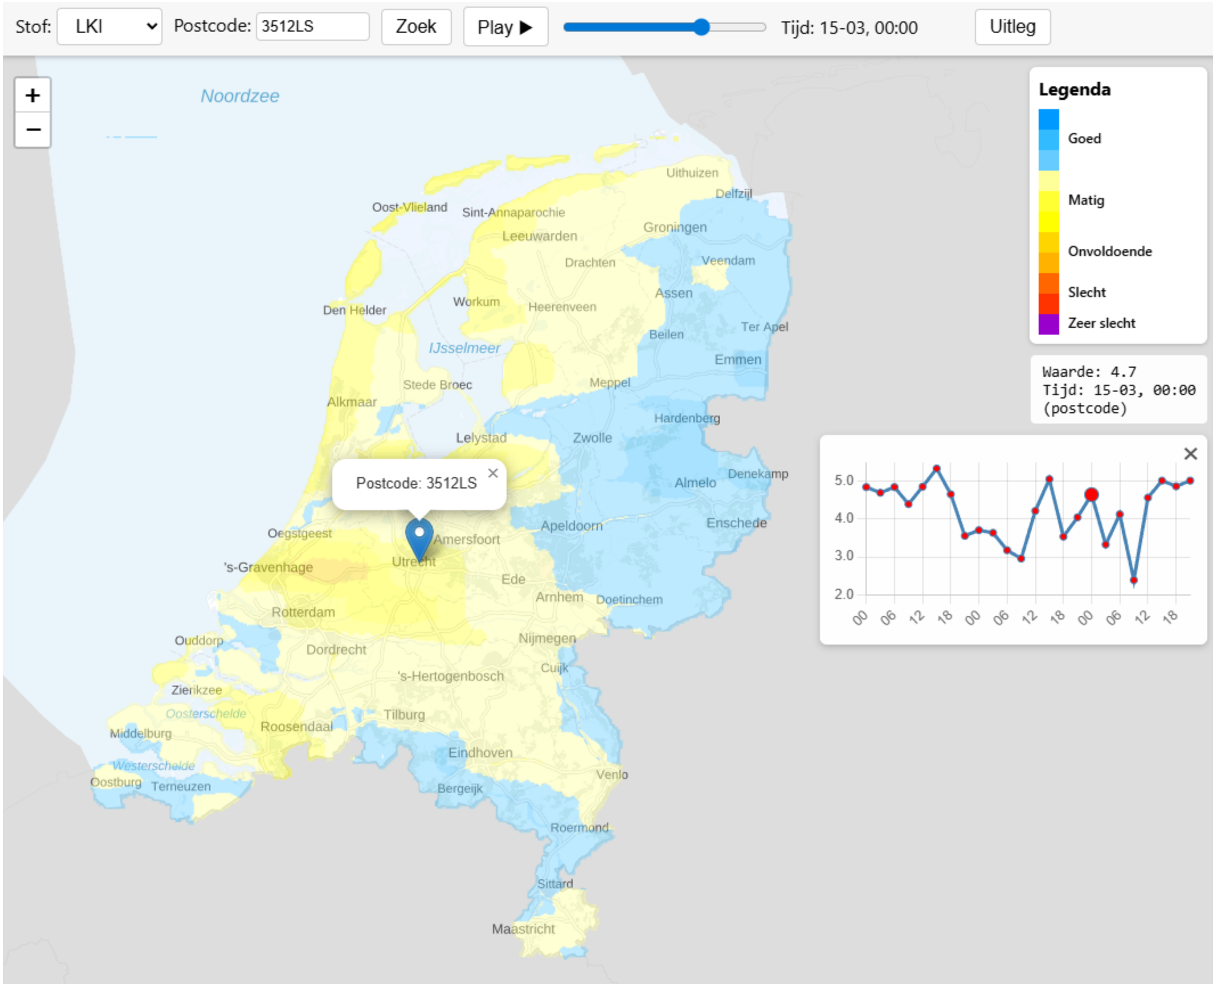

Visualisation B (Timeline)

The second concept is more timeline-oriented and starts from a default location (Utrecht). The map serves primarily to select or confirm a location, either by entering a postcode or clicking directly on the map. The core of the visualisation is a colour-coded table that displays the air quality forecast, organised by day and time of day. Each row represents a time of day (night, morning, afternoon, evening) while each column represents a day (today, tomorrow, the day after tomorrow). The colour of each cell follows the same LKI scale as used in visualisation A. Clicking on a cell shows a short advice, such as whether it is advisable to exercise outside during that period.

Technical implementation

The underlying forecast data is the same as in visualisation A: 3-hourly CAMS-based raster files in ASC format, where each file contains a grid of forecasted concentration values for the Netherlands at a given time. The key difference is in how the data is used. Rather than rendering the full spatial grid as a map overlay, this visualisation extracts only the values for a single selected location. For each forecast, it identifies the grid cell corresponding to that location and reads the predicted LKI value directly from the data.

The LKI-values are then grouped into four time-of-day periods (night (00:00–06:00), morning (06:00–12:00), afternoon (12:00–18:00), and evening (18:00–00:00)) and averaged within each period to produce a single representative value per period. The resulting summary is displayed as a heatmap using the Plotly library. When a user clicks on a cell, the application looks up the corresponding average LKI value and generates a short advice text based on fixed thresholds, displayed directly below the grid.

Figure 3 shows a screenshot of visualisation B.

Figure 3: screenshot of visualisation B

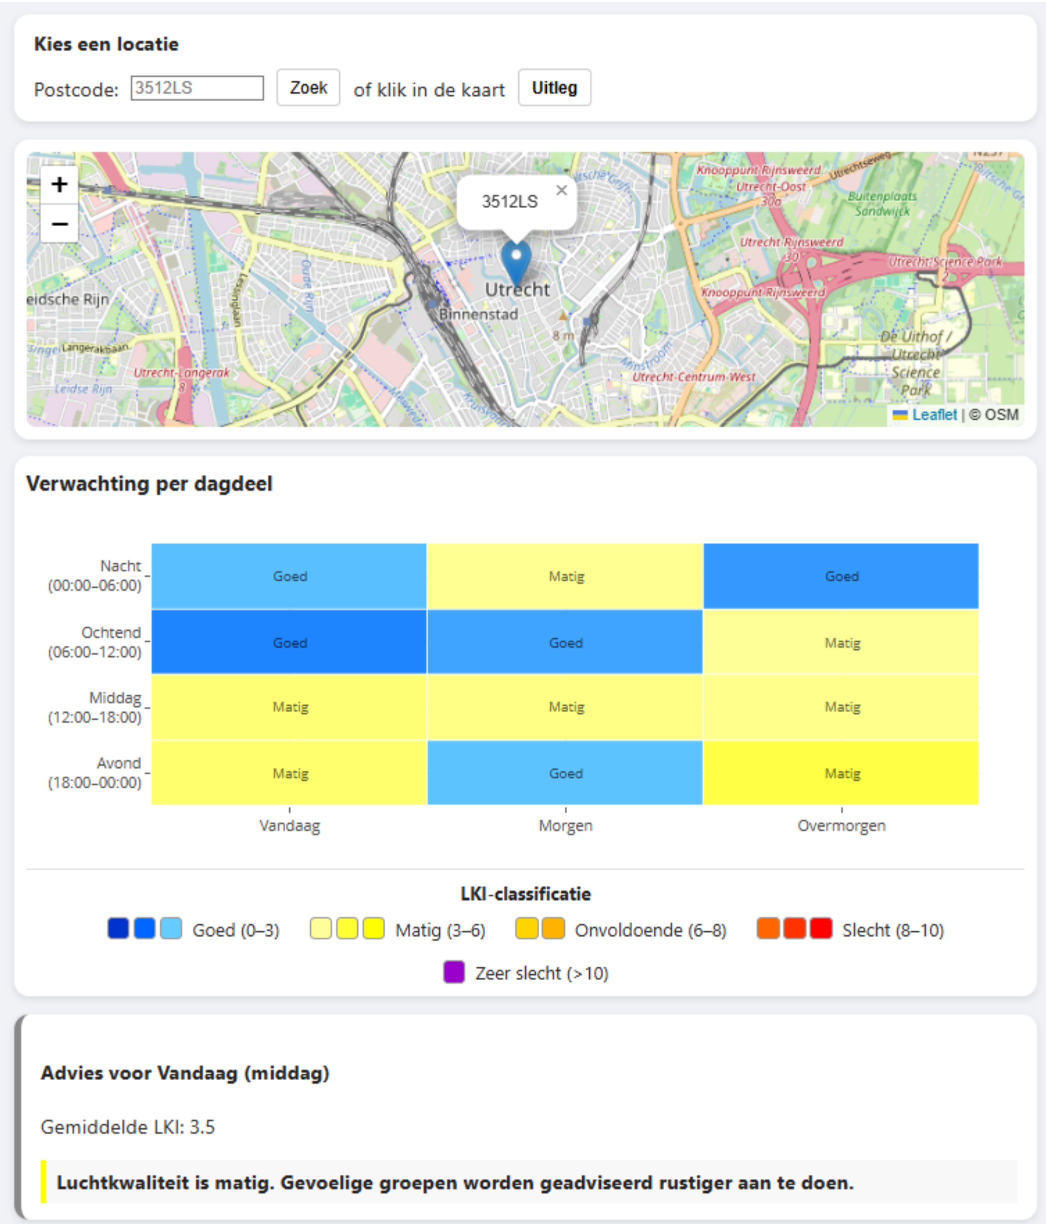

Feedback on the visualisations

After the mock-up websites went online, feedback was collected from a range of stakeholders, including representatives from Longfonds, DCMR, and the SamenMeten community, as well as colleagues from RIVM and KNMI. The feedback covered both visualisations and touched on readability, usability, and content.

Several points have already been addressed. Minor readability issues, such as label formatting and button labelling, have been corrected in both visualisations. In visualisation A, the LKI colour scale is now applied consistently across all pollutants, the application now starts automatically in play mode , and a time-series chart has been added to show predicted values for a selected location across all forecast hours- a feature suggested independently by multiple reviewers. In visualisation B, the time blocks have been changed to 6-hour intervals aligned with familiar tools such as the Stookwijzer.

A feedback session was organised on the 24th of March 2026. Participants provided feedback on both visualisations as well as general remarks. In terms of general feedback, suggestions included:

- G1. zooming in on the user's current location at startup,
- G2. adding a help button with details on all available options,
- G3. reconsidering the LKI colour scale as it was perceived as overly yellow, and
- G4. improving the display on mobile devices.

For Visualisation A, participants suggested:

- A1. enabling the display of multiple locations simultaneously,
- A2. removing the click pop-up,
- A3. smoothing the map using a weighted average,
- A4. adding numerical ranges to the legend categories,
- A5. displaying wind direction on the map, and
- A6. capping map loading times at a maximum of five seconds.

For Visualisation B, feedback focused on:

- B1. showing the dominant substance causing the LKI value (ideally all substances),
- B2. adding today/tomorrow/the day after tomorrow labels above the columns, and
- B3. reconsidering the layout as both the map and the blocks were considered too wide, potentially displaying them side by side.

During the meeting it was decided to start work on suggestions G1, G2 and G4. For G4, display on mobile devices, the location of the map and legends will be optimised.

For visualisation A, there was consensus that a smoother, more natural visualisation is preferable to emphasising local details in the maps. By interpolating hourly maps between the 3-hour time steps and simultaneously reducing the spatial resolution from 0.5 km to 1.0 km, the responsiveness of the visualisation is expected to improve slightly .

For visualisation B, the layout of the map, that is too wide will be looked at. The location of labels can be chanced quite easy.

The goal is to continue refining the visualisations iteratively , ensuring that the final tool aligns as closely as possible with the needs of its intended users.

Conclusion

A stakeholder workshop has been organised, two visualisation concepts have been developed based on the outcomes, and two mock-up websites are online. Two rounds of feedback have been collected and are being used to iteratively refine the visualisations. The active involvement of end users and relevant organisations throughout this process is considered essential to ensure that the final demonstrator meets the needs of its intended audience.