CLIMATE ANALYSIS IN AFRICAN CITIES (CAIAC)

D3.1 Core Scientific Activities Progress Report

Study accomplished under the authority of the European Space Agency

Reference: 2025/EI/R/3470

March 2026

Vision on technology for a better world

vito.be

D3.1 Core Scientific Activities Progress Report

VITO

Boeretang 200
2400 MOL
Belgium
VAT No: BE0244.195.916
vito@vito.be – www.vito.be
IBAN BE34 3751 1173 5490 BBRUBEBB

Nele Veldeman
Project Manager
nele.veldeman@vito.be

Koen De Ridder
Science Lead
koen.deridder@vito.be

Vision on technology
for a better world

vito.be

AUTHORS

VITO

Koen De Ridder

Tomas Crols

Benjamin Lanssens

Niels Souverijns

Inge Uljee

Nele Veldeman

University of Ottawa

Mohammad Alobaidi

Hossein Bonakdari

Ousmane Seidou

DISTRIBUTION LIST

Clement Albergel, ESA

Amy Campbell, ESA

SUMMARY

3.1 deliverable reports progress in the core scientific activities of the CAIAC project as of March 2026 (KO+9). It builds on the methodological framework defined in Deliverable D2.1 and on the first tailored data package documented in Deliverable D2.2.

Progress during the reporting period shows that the heat-modelling activities have reached a relatively advanced stage. Present-day UrbClim simulations have been completed for all target cities, together with post-processing, wet-bulb globe temperature (WBGT) calculation, and the derivation of key heat-stress indicators. These indicators have been downscaled using Earth Observation data, primarily ESA WorldCover, while a subset of cities benefited from the use of ESA CCI High Resolution Land Cover (HRLC). In addition, comparative work between WorldCover-based and HRLC-based downscaling has been carried out, and a Monte Carlo-based uncertainty assessment has been initiated for the HRLC-based approach.

At the same time, future-oriented scientific activities are progressing. Future land-use maps generated by the urban growth model have already been converted into terrain inputs for the UrbClim model for several cities, future climate signals based on SSP scenarios have been prepared, and future UrbClim simulations have been launched. Work on the Land Surface Temperature (LST) component remains ongoing, including method development related to thermal infrared EO data and the investigation of river-temperature initialisation issues identified for cities such as Niamey. On the flood side, substantial progress has also been achieved, including SWAT+ hydrological simulations, return-period analyses, hydrodynamic modelling refinement, and the development of a flood-threshold / flood-risk workflow.

Validation against station data and EO reference products is still being consolidated, uncertainty characterisation is partly mature, and the flood-modelling, heat-modelling and LST parts still require additional input before a final scientific synthesis can be presented. Therefore, D3.1 should be read as a scientific progress report, not as the final scientific assessment of CAIAC.

TABLE OF CONTENTS

LIST OF ACRONYMS

Acronym	Definition
ACMAD	African Centre of Meteorological Applications for Development
CCI	Climate Change Initiative
CMIP6	Coupled Model Intercomparison Project Phase 6
DART	Discrete Anisotropic Radiative Transfer
EO	Earth Observation
ERA5	ECMWF Reanalysis version 5
GCM	Global Climate Model
GeoDynamiX	Urban growth modelling framework
GHSL	Global Human Settlement Layer
GloFAS	Global Flood Awareness System
GWL	Global Warming Level
HEC-RAS	Hydrologic Engineering Center River Analysis System
HRLC	High Resolution Land Cover
LST	Land Surface Temperature
QDM	Quantile Delta Mapping
RIL	Rapid Inundation Labelling
SAHEL	Satellite-based African Hydrological Ensemble Learner
SSP	Shared Socioeconomic Pathway
SWAT+	Soil and Water Assessment Tool Plus
UHI	Urban Heat Island
WBGT	Wet-Bulb Globe Temperature
WP	Work Package

1 INTRODUCTION

The Science Requirements Document (D1.2) identified a clear need for high-resolution, forward-looking analyses of urban heat and flood risks in African cities. In particular, the project is expected to address how urban fabric amplifies extreme heat and flooding, how urban growth and climate change jointly shape future risk, how green infrastructure may mitigate impacts, and how climate information can eventually support socio-economic risk assessment. The proposal and technical background further position WP3 as the work package in which these scientific objectives are operationalised through the generation of urban climate simulations, flood risk maps, uncertainty estimates, and scientific analyses.

D3.1 reports the current state of these core scientific activities. It follows the methodological definition provided in D2.1 and the harmonised model-ready input preparation described in D2.2. The report is therefore situated at the point where scientific methodology and tailored input datasets are translated into simulations, derived indicators, preliminary scientific findings, and the first validation and uncertainty results.

The core objective of WP3 is to conduct the scientific analysis itself. This includes running urban climate simulations for present and future periods, generating climate indicators and flood risk maps, quantifying uncertainties through and addressing research questions on extreme heat and flooding.

2 MODELLING SETUP

This section provides only a brief summary of the modelling setup used in WP3; the full methodological description is provided in D2.1, Section 3.

2.1 GeoDynamiX (urban growth)

A distinctive feature of CAIAC is the explicit inclusion of urban growth in the scientific analysis. Future urban growth scenarios are simulated with GeoDynamiX, which combines land-use interactions, physical suitability, accessibility and zoning constraints to generate annual land-use and population maps. This step is essential because future heat and flood risk in rapidly growing African cities cannot be interpreted credibly without accounting for evolving urban extent and density.

The model is calibrated over the historical period using observed land-use and population changes and is then used to produce scenario-dependent projections to the end of the century. Pilot-city applications have been used to refine the calibration workflow and to support the broader roll-out across the target-city sample.

The outputs of GeoDynamiX serve a dual purpose in WP3. They are scientific outputs in their own right, allowing the project to characterise plausible urbanisation pathways, and they also act as dynamic inputs to the climate and flood chains by updating terrain and land-use conditions under future scenarios.

A more detailed description of the GeoDynamiX model configuration, integration of EO and land-use change data, simulation strategy and model outputs is provided in D2.1, Section 3.1.

2.2 UrbClim (heat)

UrbClim is the urban climate model used in CAIAC to simulate urban heat conditions at high spatial resolution. The model produces hourly near-surface meteorological fields, including air temperature, humidity, wind and LST, which are subsequently converted into climatological and impact-oriented indicators such as urban heat island intensity, heatwave metrics and Wet-Bulb Globe Temperature (WBGT).

In the CAIAC implementation, UrbClim relies on a harmonised EO-based land representation at city scale. The terrain inputs combine land-cover and built-up information from sources such as ESA CCI Land Cover, Copernicus WorldCover, Global Human Settlement Layer (GHSL) and Sentinel-2 NDVI, which are transformed into model-ready parameters describing vegetation, roughness, imperviousness and related urban-surface characteristics.

Present-day atmospheric forcing is based on ERA5. Future forcing is derived from CMIP6 data using a Quantile Delta Mapping workflow and delta matrices that preserve information on multi-model mean and spread. Future UrbClim simulations are coupled with scenario-dependent land-use updates from GeoDynamiX.

A more detailed description of the UrbClim model configuration, EO and climate-data integration, simulation strategy and output indicators is provided in D2.1, Section 3.2.

2.3 SAHEL (flooding)

The flooding component is based on the SAHEL framework and associated hydrological and hydrodynamic tools. In the CAIAC workflow, SWAT+ provides physically consistent runoff and discharge information, hydrodynamic modelling with HEC-RAS provides inundation-depth information where required, and a flood-threshold / flood-risk workflow translates these inputs into products that can support the scientific analysis of present-day and future flood hazard.

This approach enables flood analyses that are consistent with the broader scenario framework of the project. It couples hydroclimatic forcing, land-surface and urban-growth information, and event-frequency analysis into a coherent chain for assessing flood hazard in selected African cities.

A more detailed description of the SAHEL model configuration, EO and climate-data integration, simulation strategy and flood-map outputs is provided in D2.1, Section 3.3.

2.4 Climate scenarios

Climate projections in CAIAC are based on CMIP6 Shared Socioeconomic Pathways (SSP). The project methodology foresees direct use of CMIP6 GCM output for urban-climate and flood applications. For the flood model, additional calibrated products such as NASA NEX-GDDP-CMIP6 are also used. The baseline period for present-day simulations is 2001–2020. Future horizons are defined as 2021–2040, 2041–2060 and 2081–2100, and the scientific framing also considers Global Warming Levels of 1.5°C, 2.0°C and 3.0°C.

The methodological focus of CAIAC is primarily on SSP2-4.5, as a plausible central pathway, with SSP3-7.0 additionally used where required to represent higher warming levels such as 3°C. In the urban-growth model, scenario generation is framed around SSP2 and SSP3. In the flood model, multiple SSP pathways are considered, including SSP126, SSP245, SSP370 and SSP585.

At the current reporting stage, the scenario work is most advanced in the preparation of future forcing and land-use inputs and in the launch of future simulations. The scientific emphasis of D3.1 will therefore remain on baseline results and first scenario implementation status, while only the scenario-dependent outputs that have already reached a stable level of maturity are considered for inclusion as preliminary results.

3 CURRENT MODEL STATUS OVERVIEW

In this section, we present an overview of the current operational coverage of the three model across the target-city sample and the degree to which their outputs are already usable in the scientific analysis.

The urban-heat and urban-growth model covers the full target-city sample and provides the most advanced scientific basis at this stage. Flood models are already feeding the WP3 analysis, but their maturity remains more heterogeneous across scenarios, cities and modelling components. LST component still needs some work to be fully operational.

Table 1 summarises the present status of the principal modelling components and provides the operational context for the remainder of this report.

Component	Current status in D3.1	Outputs currently available	Remaining work
GeoDynamiX	Historical calibration and future scenarios generation are sufficiently advanced for integration into WP3.	Future land-use maps for all cities available for every scenario (SSP2 and SSP3).	–
UrbClim - baseline	Completed for the full target-city sample and fully usable in the present progress report.	Baseline simulations, post-processing, WBGT and key heat-stress indicators; EO-based downscaling products.	Future simulation based on land-use growth. Validation of present-day result. Uncertainty analysis.
UrbClim - future	Prepared and running for the first scenario combinations using future land-use and climate signals.	Scenario-ready forcing and terrain inputs; first future simulations in execution.	Completion and quality control of future runs; synthesis of the first scenario-based findings.
LST component	Method development and diagnostic work are ongoing; the component is not yet fully mature for consolidated results.	First river-temperature diagnostics and intercomparison work for large-river cities.	Apply the methodology (considering directional effect) to every city.
Flood chain	Hydrological and hydrodynamic components are operational, with initial outputs already available for selected flood cities.	Runoff/discharge analyses, return-period inputs, hydrodynamic simulations and initial flood-risk workflow outputs.	city-by-city maturity statement, representative figures and concise validation summary.

4 SCIENTIFIC ANALYSIS

4.1 Scope of the core scientific activities

The core scientific activities reported in D3.1 are aligned with the four key scientific questions established in D1.2: the amplification of heat and flooding by urban fabric; the relative contributions of urban growth and climate change; the mitigation potential of green infrastructure; and the eventual translation of climate hazards into socio-economic risk. WP3 is also framed around the generation of simulations, indicators, flood maps and uncertainty estimates in support of scientific analysis and publication.

At the current stage of the project, the maturity of evidence differs across these questions. The first question, concerning the spatial amplification of heat and flooding by urban form, is already being addressed through baseline UrbClim outputs, indicator maps, and the first flood-modelling products. The second question, concerning the relative roles of urban growth and climate change, is entering the analysis phase through the preparation of future climate signals and future land-use scenarios. The third and fourth questions, concerning green infrastructure and socio-economic impacts, remain scientifically relevant but are less mature in terms of project outputs currently available for inclusion in D3.1. This report therefore focuses on reporting the status and first findings that are already defensible at this stage.

4.2 Urban climate modelling progress

The urban-climate workflow has progressed from model preparation to operational scientific production. Baseline UrbClim simulations, post-processing and derivation of heat-stress indicators are now available for the full target-city sample. The resulting indicators have been downscaled from 300 m to 10 m with EO data, primarily using ESA WorldCover, while the higher-resolution ESA CCI HRLC product has been applied where the input data are available.

These outputs already support the first cross-city analyses. A first synthesis of nocturnal urban heat island behaviour across Köppen climate zones indicates that the more arid climate zones tend to display stronger median nocturnal UHI intensity than the more humid classes.

At the same time, the future heat chain has moved into active execution. Future land-use maps generated by the urban-growth model have been transformed into terrain inputs for UrbClim for the first set of cities, future climate signals have been prepared, and the corresponding future simulations have been launched. These runs are the key next step for disentangling the relative influence of climate change and urban growth on future urban heat stress.

4.2.1 Multi-city patterns of present-day urban heat and heat stress

A revised cross-city analysis was conducted using the CAIAC urban climate indicator average over the base period for the 54 target cities. The objective of this analysis was to identify a small set of robust and interpretable patterns that can already be supported by the baseline simulations. The analysis focuses on four complementary dimensions of heat risk: nighttime urban heat island intensity, hot days, warm nights, and humidity-sensitive heat stress expressed through WBGT exceedance hours (Figure 1).

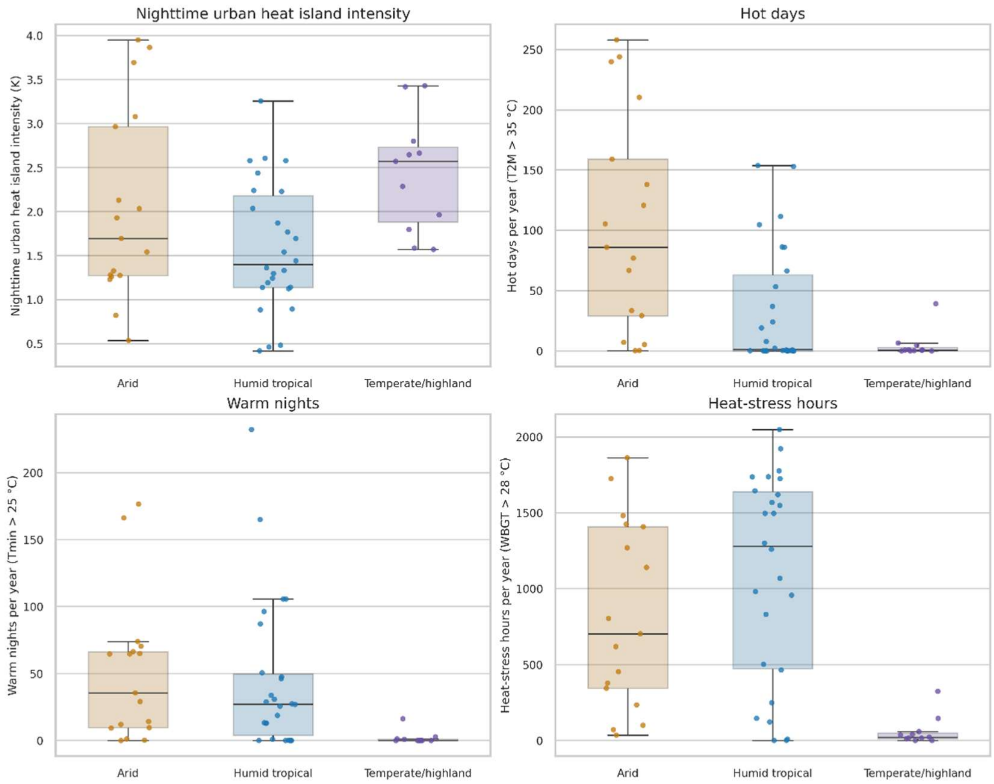

Figure 1 Distribution of four complementary indicators across the CAIAC city sample: nighttime urban heat island intensity, hot days, warm nights, and WBGT exceedance hours. The figure shows that climate family is a strong organizing factor, but that the dominant pattern depends on the indicator. Arid cities exhibit the strongest dry-bulb heat and strong nighttime urban warming, whereas humid tropical cities show the highest humidity-sensitive heat-stress burden.

The results show that present-day urban heat in African cities cannot be described by a single indicator. The strongest signal for nighttime urban heat island intensity is found in the arid and temperate/highland city groups, whereas the strongest signal for humidity-sensitive heat stress is found in the humid tropical group. In other words, the cities with the strongest nocturnal urban warming are not necessarily the cities with the highest annual heat-stress burden. This distinction is scientifically important because it shows that dry-bulb temperature metrics and heat-stress metrics capture different aspects of urban heat risk (see Figure 2).

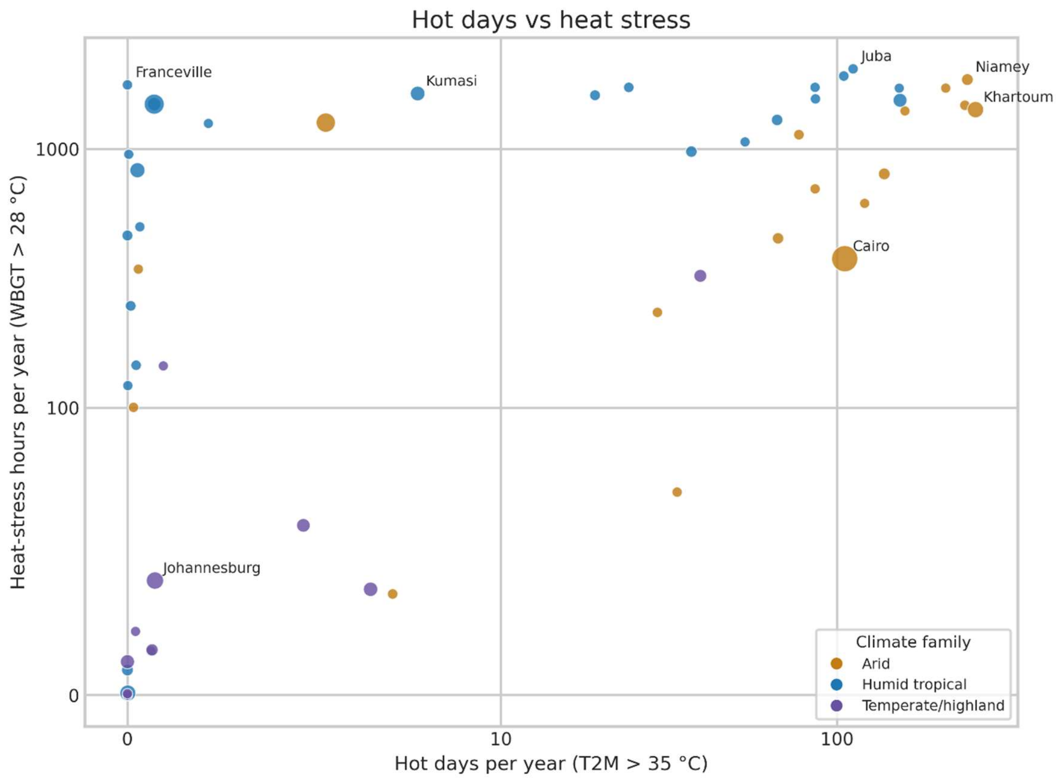

Figure 2 Cross-city comparison between a dry-bulb heat metric and a humidity-sensitive heat metric for the 54-city CAIAC sample. Dot sizes are relative to the city population. The plot highlights that cities with many hot days are not necessarily the same cities with the highest annual WBGT exceedance.

4.2.2 EO-based downscaling

In addition to the city-scale UrbClim simulations, a statistical downscaling workflow was applied to derive 10 m indicator maps from the coarser model output. In the baseline implementation, this downscaling relied primarily on ESA WorldCover, while the higher-resolution ESA CCI High Resolution Land Cover (HRLC) product was used for the subset of nine cities for which this dataset was available: Khartoum, Yaounde, Abeche, Mbale City, Juba, Asmara, Naqamte, Carnot and Ali Sabieh.

WorldCover-based and HRLC-based downscaling were explicitly compared for the city of Khartoum (see Figure 3). Note that WorldCover data are from 2021 while HRLC data are from 2019, the difference between both datasets could be largely explained by the difference in the date acquisition. The comparison shows that the use of HRLC does not fundamentally alter the core urban footprint, but it does modify the representation of several transitional and non-urban dryland classes. After remapping to UrbClim classes, HRLC increases the fraction of grassland by 7.85 % and the urban class by 0.60 %, while reducing cropland by 5.02 % and bare soil by 2.37 %. The largest discrepancies therefore concern mixed or transitional land-cover classes rather than the dominant urban core. This result is reflected on the indicator difference map that shows strong agreement for the main urban core, bare-soil background, and inland water bodies, while the largest differences occur in peri-urban and heterogeneous classes such as grassland, cropland, and shrubland.

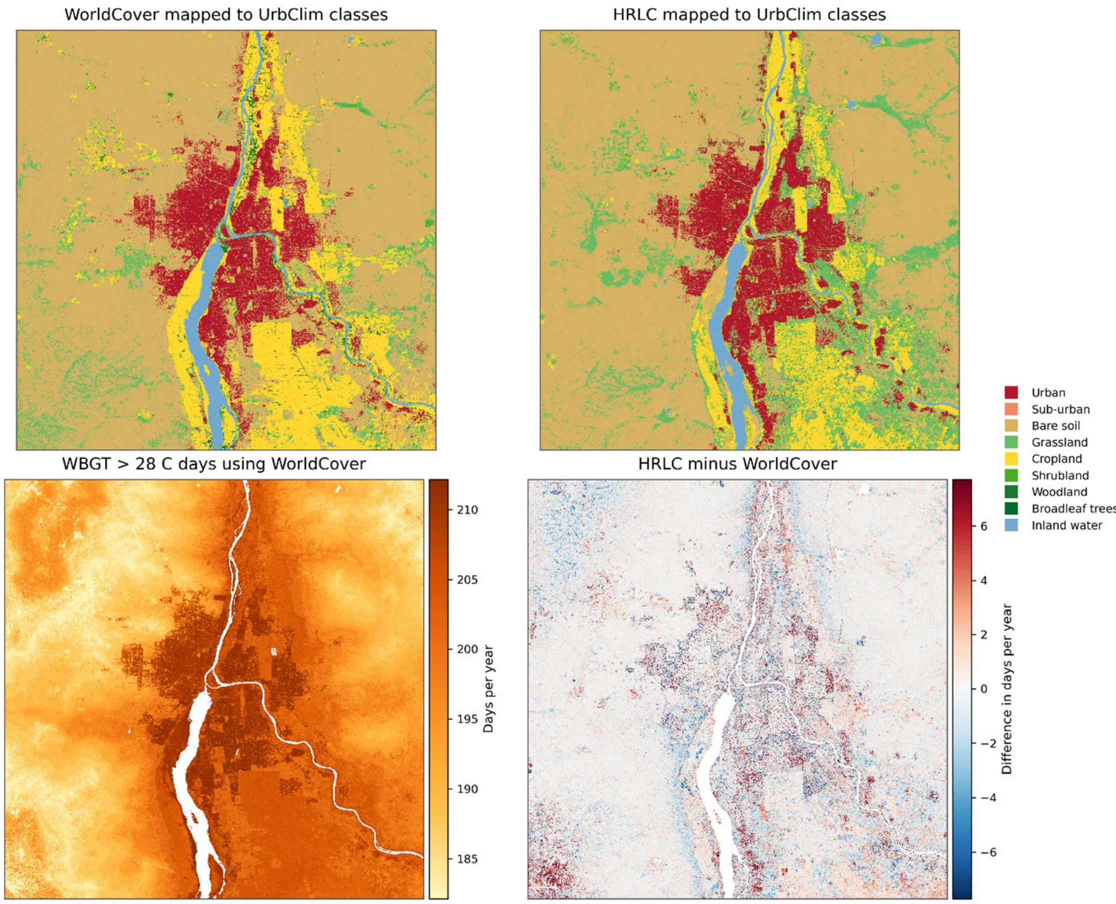

Figure 3 Top row: WorldCover- and ESA CCI HRLC-based land-cover maps after harmonisation to the UrbClim class system. Bottom row: example downscaled heat stress indicator obtained with WorldCover and the corresponding HRLC–WorldCover difference.

4.2.3 Uncertainty assessment of HRLC-based downscaling (Khartoum)

A first uncertainty-propagation experiment was carried out for Khartoum to assess how uncertainty in the ESA CCI HRLC product propagates into the EO-based downscaling of urban heat indicators. In practical terms, the experiment used the HRLC uncertainty information available at pixel level, including two class-flip probability layers expressing the probability that the most and second likely class occurs, the probability margin between the two leading classes and an additional quality index layer. These fields were used to generate 200 Monte Carlo land-cover realisations, which were then propagated through the Khartoum downscaling workflow. The resulting ensemble was analysed for two representative indicators: nighttime UHI and the annual number of days with WBGT above 28 °C.

The results show that land-cover uncertainty leads to measurable but still moderate output variability. For nighttime UHI, the mean uncertainty over urban pixels is 0.1 °C, compared with 0.08 °C over non-urban pixels. For number of WBGT days above 28 °C, the mean uncertainty is 0.9 days over urban pixels and 0.8 days over non-urban pixels. These values indicate that uncertainty is somewhat stronger within urban areas, but not dramatically so. At the same time, the downscaled urban signal itself remains much larger than these uncertainty levels: in Khartoum, the mean downscaled nighttime UHI over urban pixels is 3.7 °C (spatial standard deviation: 1.3 °C), while the mean value of number of days with WBGT above 28 °C over urban pixels is 204.4 days per year (spatial standard deviation: 7.6 days). Taken together, these first results suggest that HRLC uncertainty does affect local indicator values and should therefore not be ignored, but it does not appear to invalidate the broader spatial interpretation of the downscaled patterns (see Figure 4).

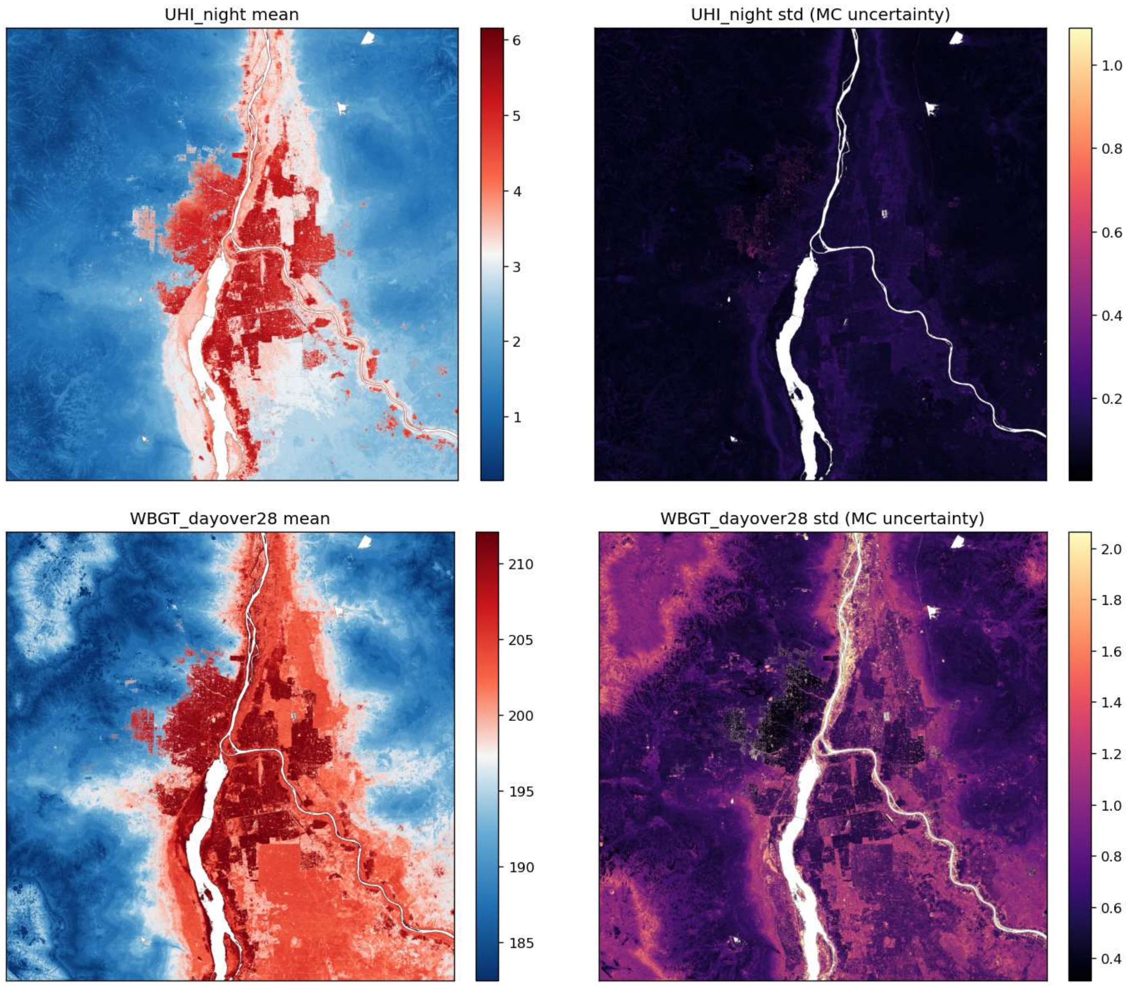

Figure 4 Ensemble mean and standard deviation of downscaled nighttime UHI (top) and annual WBGT > 28 °C days (bottom) for Khartoum, derived from 200 Monte Carlo land-cover realisations based on the ESA CCI HRLC uncertainty layer. In each row, the left panel show

4.2.4 Indicator maps

To complement the cross-city statistical analysis, a small number of city-scale indicator maps were retained to illustrate the spatial organisation of present-day urban heat and heat stress within individual domains. These maps are not intended to provide an exhaustive atlas of all indicators, but rather to highlight a few robust and visually clear examples of how the simulated signals are distributed across contrasting urban and climatic settings. Where indicated, the maps use the ESA CCI high-resolution land-cover downscaling version instead of the WorldCover based downscaling.

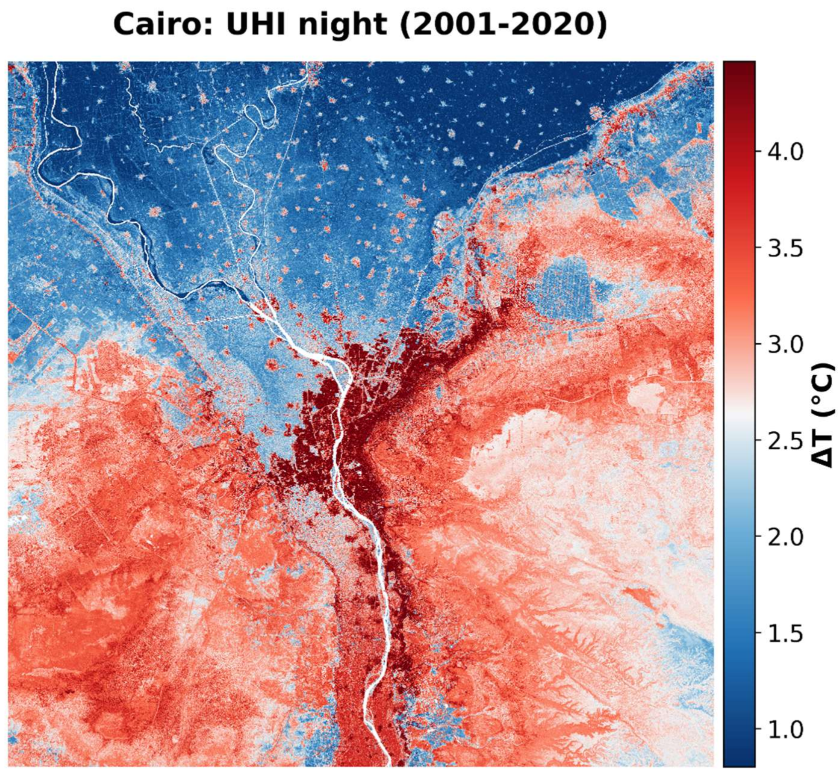

Figure 5 Cairo UHI night map (average over 2001-2020)

The Cairo nighttime urban heat island map (Figure 5) shows a very strong and spatially coherent nocturnal warming signal over the dense metropolitan core and along the main urban corridor following the Nile valley. Deserts on the east and west also show a homogeneous UHI landscape. The pattern is sharply contrasted with the cooler agricultural and deltaic areas to the north. This is a classic example of the nighttime urban heat island structure: dense built-up areas retain heat efficiently after sunset, while vegetated and irrigated surfaces cool more effectively. Cairo's broader physical setting helps explain this contrast, as the city occupies a hot, dry Nile corridor at the gateway to the delta, with the urban fabric embedded between desert and irrigated landscapes.

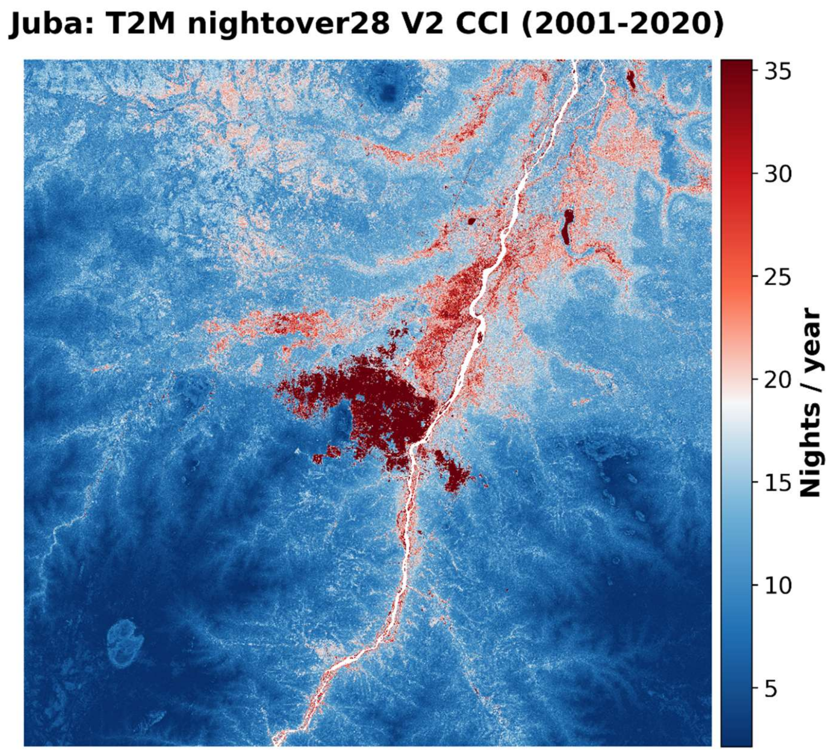

Figure 6 Juba number of tropical night (Tmin over 28 °C) per year (average over 2001-2020). This map was downscaled using ESA CCI HRLC data.

Juba map highlights the frequency of very warm nights, expressed here as the annual number of nights with minimum air temperature above 28 °C (Figure 6). The highest values are concentrated over the urban core, while the surrounding landscape shows distinctly lower frequencies. The geography of Juba is relevant for interpretation: the city is located on the Baḥr al-Jabal / White Nile corridor, in a tropical wet–dry environment where warm background conditions already favour limited nocturnal cooling during part of the year.

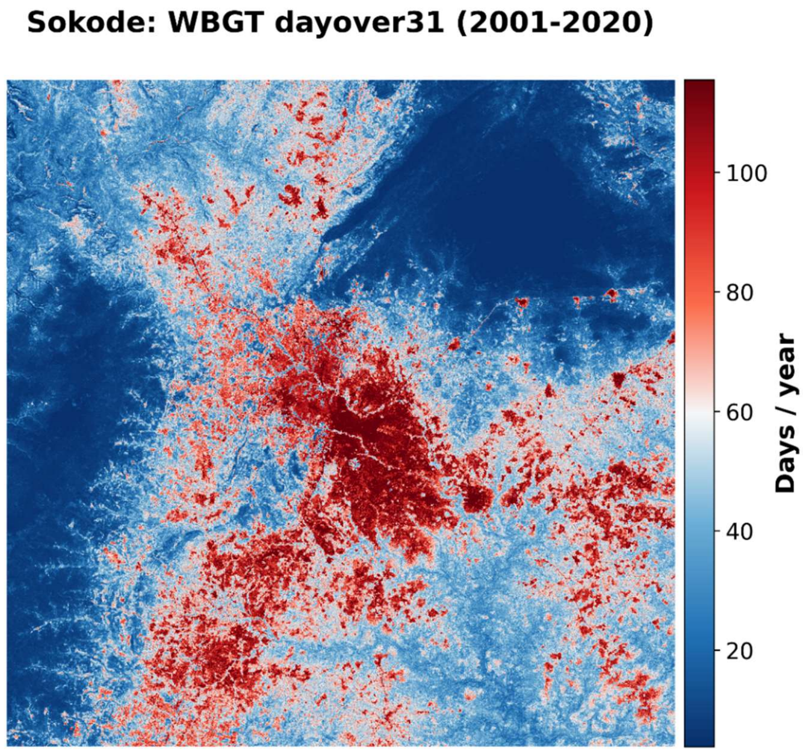

Figure 7 Sokode number of WBGT days over 31 °C (average over 2001-2020)

Sokodé map shows the annual number of days on which daytime WBGT exceeds 31 °C (Figure 7) which is usually considered as. a high heat-stress threshold. In contrast to the more compact signal seen for nighttime UHI, the heat-stress burden here extends beyond the city centre into a broader settled and cultivated landscape, although the urban core still stands out as one of the most affected zones. The map shows that daytime heat stress is not purely an inner-city phenomenon: it is shaped by both urban density and the wider climatic and land-surface context. The cooler areas around Sokodé appear to coincide with less urbanised and more topographically varied surroundings. This is consistent with Sokodé's location in central Togo near a gap in the Togo Mountains, where terrain and vegetation heterogeneity contribute to the pattern seen on the map.

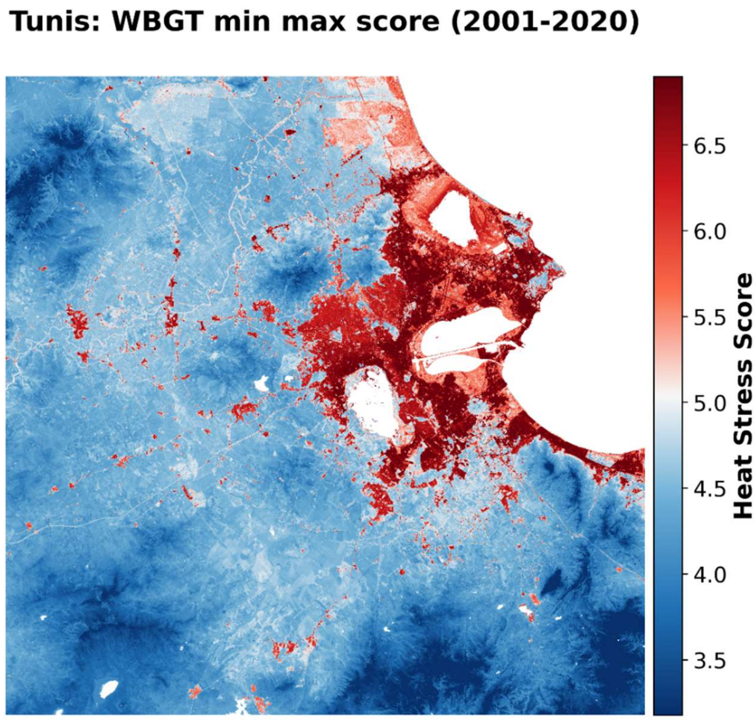

Figure 8 Tunis heat stress score (average over 2001-2020)

Tunis heat stress score map (Figure 8) highlights the parts of the city where daytime and nighttime heat stress combine most strongly, meaning that heat exposure remains elevated during the day while nighttime relief is also limited. In this case, the highest values are concentrated over the dense urbanised coastal sectors of Greater Tunis, especially around the low-lying built-up plain between the hills and the lagoon system. By contrast, lower values appear over the surrounding area and less urbanised sectors, where local topography and lower built density likely favour better nocturnal relief.

4.2.5 Future simulation progress

Future urban-climate simulations are currently being prepared and executed for the CAIAC cities. The EO-based terrain preparation for UrbClim has been completed, while future land-use simulations are being generated with the GeoDynamiX urban growth model under SSP2 and SSP3 scenarios. For the first available cities, these future land-use maps have already been converted into UrbClim-compatible terrain inputs, after which the future climate signal was applied and the corresponding UrbClim simulations were launched.

Within the CAIAC workflow, future simulations combine two complementary components. First, the future meteorological forcing is derived from bias-corrected CMIP6 climate projections using the percentile-based delta method described in D2.2 and D2.1. Second, the future urban surface state is derived from the future land-use maps produced by GeoDynamiX. In this way, future UrbClim simulations account jointly for large-scale climate change and city-scale urban expansion.

The future land-use maps are first projected to the fixed CAIAC UrbClim grid and converted to the UrbClim land-use taxonomy using the same class harmonisation rules as for present-day simulations. The resulting future land-use raster is then used as the updated land-use field in the UrbClim surface file.

The derivation of the remaining future terrain parameters follows a mixed static–dynamic approach. Parameters that are not expected to evolve directly with land use at this stage of the project, such as terrain elevation and soil texture or hydraulic parameters, are retained from the present-day setup. By contrast, surface properties that are strongly linked to land use are updated.

The focus is on four surface fields: soil sealing, vegetation cover fraction, roughness length and anthropogenic heat flux. From those fields, we derive present-day class statistics (class-wise mean, standard deviation, minimum and maximum values over the present-day domain).

For each changed pixel, the new future value is then sampled stochastically from a normal distribution centred on the present-day mean of the target future class, with a spread equal to the present-day standard deviation of that class. The sampled value is subsequently clipped to the present-day minimum and maximum of the same class, which prevents implausible out-of-range values while preserving within-class variability.

This procedure is applied separately to each variable. For soil sealing, sampled values are constrained to the physical range, and non-urban target classes are forced back to zero where required by UrbClim. For vegetation cover fraction, the update is performed month by month, again by sampling from the present-day monthly class statistics of the target class, after which the peak-greenness month is recomputed. For roughness length, changed pixels are sampled from the present-day class-specific roughness distribution and clipped to the reference class range. For anthropogenic heat flux, the same approach is used, based on the distribution of positive present-day AHF values, after which non-urban classes are set to zero and negative values are removed.

This method provides a pragmatic approach resulting in future surface that preserve the standard UrbClim structure while updating only the land-use-dependent parameters in locations where land use changes (see Figure 9).

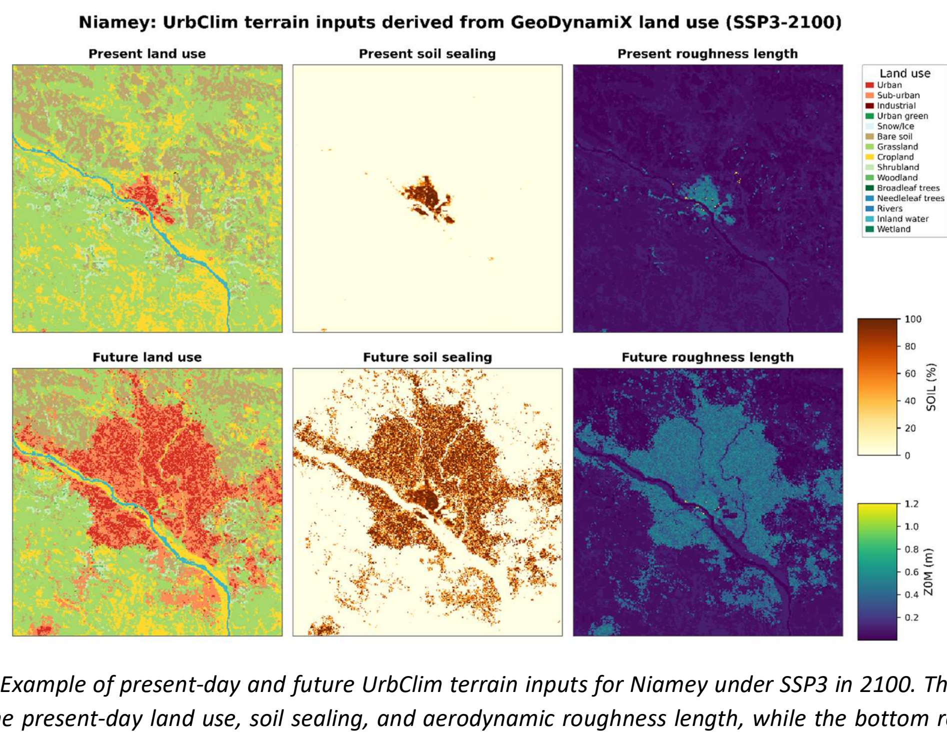

Figure 9 Example of present-day and future UrbClim terrain inputs for Niamey under SSP3 in 2100. The top row shows the present-day land use, soil sealing, and aerodynamic roughness length, while the bottom row shows the corresponding future fields. Note the formation of a new satellite city at the south-west domain border.

4.3 Land Surface Temperature (LST) modelling activities

4.3.1 Scope of the LST work within CAIAC

The LST component supports the scientific interpretation of urban heat from a thermal infrared perspective and offers an important route for model evaluation in a context where in-situ observations are scarce.

A key issue identified during the present reporting period concerns cities with very large rivers, such as Niamey, where unrealistically high nocturnal WBGT values were found over the river. The current interpretation is that this behaviour is linked to the bad treatment of river-water temperature in the forcing chain.

Therefore, this component is scientifically important because it can improve the physical realism of the urban-heat chain and demonstrate a clear added value of EO in a data-scarce environment.

Within the CAIAC methodology, LST data are not used as input to the UrbClim simulations, but exclusively as an independent validation and evaluation dataset: downscaled and satellite-derived LST imagery will be compared with UrbClim-simulated LST to assess model performance and support interpretation of the urban heat results.

4.3.2 Current model and method development status

The ESA CCI LST dataset provides a harmonised, consistently processed record of surface thermal conditions derived from multiple satellite platforms, including both geostationary (GEO) and low Earth orbit (LEO) sensors. The LEO dataset is delivered at a nominal spatial resolution of approximately 1 km, ensuring temporal consistency and inter-sensor comparability across a multi-decadal time series.

However, a spatial resolution of 1 km is fundamentally insufficient for urban climate applications, where the relevant surface heterogeneity, building materials, green spaces, impervious surfaces and water bodies, operates at spatial scales of 100–300 m. This scale mismatch is particularly critical for two high-priority applications: (i) characterisation of the urban heat island effect and its intra-urban spatial structure, and (ii) quantitative climate adaptation planning at the neighbourhood scale. Therefore, it becomes evident the need for a spatial sharpening framework for the ESA CCI LST datasets.

The fundamental physical principle governing thermal infrared (TIR) emission from a real surface is the Stefan–Boltzmann law for a graybody, which states that the total integrated emitted radiance L (W m^{-2}) is given by:

$$L = \varepsilon \, \sigma \, T^4 \qquad (1)$$

where $\varepsilon \in [0,1]$ is the surface broadband emissivity (dimensionless), $\sigma = 5.670 \times 10^{-8}$ W m^{-2} K^{-4} is the Stefan–Boltzmann constant, and T (K) is the true thermodynamic surface temperature. A satellite TIR sensor does not directly measure the true surface temperature T; rather, it registers a quantity known as the brightness temperature T_B, defined as the temperature of a hypothetical perfect blackbody ($\varepsilon = 1$) that would produce the same observed radiance L. Formally:

$$\sigma \, T_B^4 = L \qquad (2)$$

Combining Equations (1) and (2), the relationship between brightness temperature and true surface temperature is:

$$T_B = \varepsilon^{1/4} \, T \qquad (3)$$

or, equivalently, solving for the true surface temperature:

$$T = \varepsilon^{-1/4}\, T_B \tag{4}$$

This relationship is the core physical link exploited by the emissivity modulation approach (Nichol et al., 2009) and it shows that, for a surface with $\varepsilon < 1$, the brightness temperature systematically underestimates the true LST (since $\varepsilon^{1/4} < 1$ implies $T_B < T$).

Thus, the emissivity modulation method, following the framework of (Nichol et al., 2009), exploits Equation (3) to disaggregate coarse-resolution LST to a finer spatial scale. The key assumption underpinning the method is that, within a coarse 1 km pixel, the brightness temperature T_B is spatially uniform — that is, sub-pixel variations in observed LST are attributed entirely to spatial variations in emissivity rather than to independent temperature gradients. This is a physically motivated assumption because emissivity is primarily a function of surface material type and cover fraction, both of which are more slowly varying than temperature.

The algorithm implementing the emissivity modulation approach to sharpen ESA CCI LST datasets from 1km to 300m pixel size can be broken down into the following steps:

Step 1 — Derive brightness temperature at coarse resolution: Applying Equation (3) at the 1 km scale:

$$T_B = \varepsilon_{1\mathrm{km}}^{1/4} \cdot LST_{1\mathrm{km}} \tag{5}$$

Here, $LST_{1\mathrm{km}}$ is the surface temperature from the ESA CCI LST 1 km product (K), and $\varepsilon_{1\mathrm{km}}$ is the corresponding broadband emissivity (Wang & Liang, 2009) assembled from MODIS MYD21A1D climatological data at the same spatial resolution.

Step 2 — Transfer brightness temperature to fine scale: Under the spatial uniformity assumption, T_B computed in Step 1 is assigned to all sub-pixels within the coarse 1 km footprint.

Step 3 — Retrieve fine-scale LST: Applying Equation (4) at 300 m:

$$LST_{300\mathrm{m},i} = \varepsilon_{300\mathrm{m},i}^{-1/4} \cdot T_B \tag{6}$$

where the subscript i denotes the i-th sub-pixel within the coarse footprint, and $\varepsilon_{300\mathrm{m},i}$ is the fine-scale broadband emissivity derived from the ASTER AG100 climatology product at 100 m, spatially aggregated to 300 m. Substituting Equation (5) into Equation (6), the emissivity modulation downscaling relation can be written as:

$$LST_{300\mathrm{m},i} = LST_{1\mathrm{km}} \cdot \left(\frac{\varepsilon_{1\mathrm{km}}}{\varepsilon_{300\mathrm{m},i}}\right)^{1/4} \tag{7}$$

Equation (7) is the operational form of the emissivity modulation sharpening: the fine-scale LST is obtained by modulating the coarse-scale LST by the fourth root of the ratio of coarse-to-fine emissivity. For pixels where fine emissivity exceeds coarse emissivity (e.g., surfaces with higher emissivity at sub-pixel scale), the downscaled LST will be lower than the coarse value, and vice versa.

4.3.3 Results

Thus, from the above section, the main input datasets for the sharpening methodology are:

LST_{1km} is taken form the ESA CCI LST dataset esacci.LST.day.L3S.LST.multi-sensor.multi-platform.IRCDR.3-00.DAY [1]

ϵ_{1km} broad-band yearly seasonal climatology assembled from the MODIS MYD21AD[2]

$\epsilon_{300m,i}$ broad-band emissivity is assembled from the ASTER ϵ_{100m} AG100 climatology product.

We have implemented the sharpening methodology described in the afore mentioned subsection to estimate $LST_{300m,i}$, and Figure 10 provides a schematic overview of the end-to-end sharpening workflow.

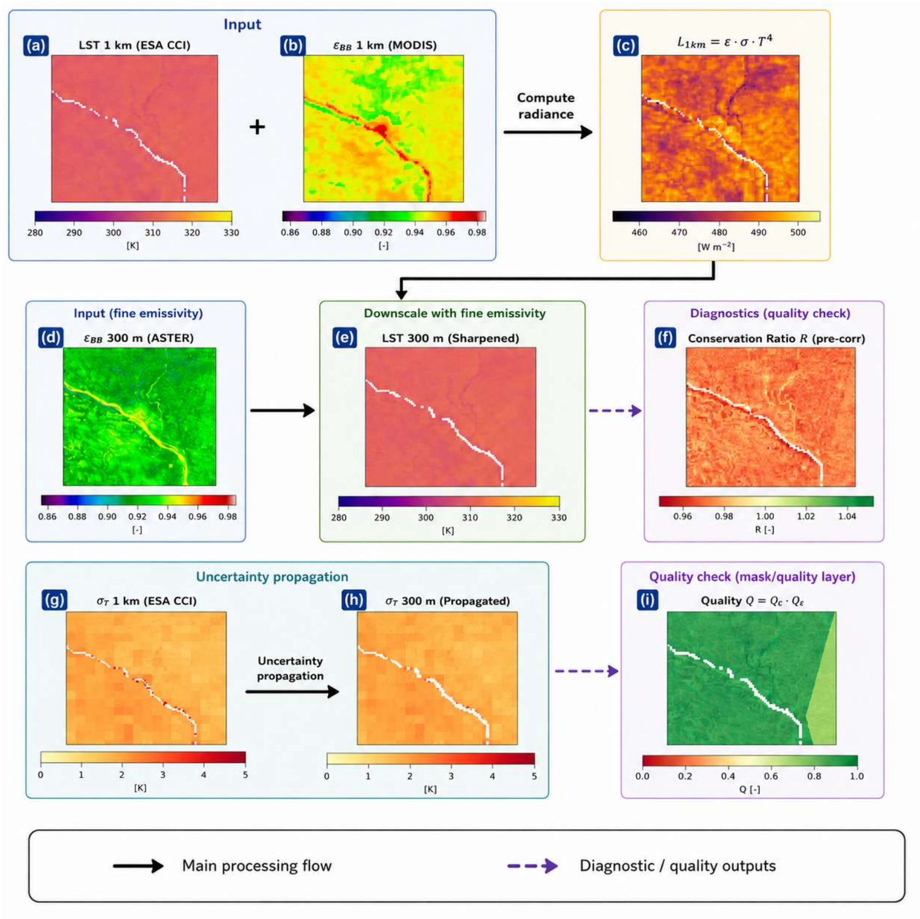

Figure 10 ESA CCI LST at 1 km (a) is combined with MODIS broadband emissivity (b) to derive coarse-resolution radiance (c), which is then downscaled using ASTER broadband emissivity at 300 m (d) to generate sharpened LST at 300 m (e). In parallel, LST uncertainty is propagated to 300 m (h) and diagnostic layers, including the conservation ratio (f) and final quality mask (i), are produced. Solid arrows show the main workflow and dashed arrows show diagnostic outputs

[1] https://catalogue.ceda.ac.uk/uuid/e67b02b814bf4a78b9ec90b4259f0741/?q=&sort_by=title_asc&results_per_page=20&objects_related_to_uuid=e67b02b814bf4a78b9ec90b4259f0741&jump=related-anchor

[2] https://www.earthdata.nasa.gov/data/catalog/lpcloud-myd21a1d-061

As a representative example, Figure 11 shows the time series of sharpened $LST_{300\mathrm{m}}$ maps for the city of Niamey, Niger, for the year 2003. Only acquisition dates for which at least two thirds of the pixels within the urban extent were cloud-free were retained.

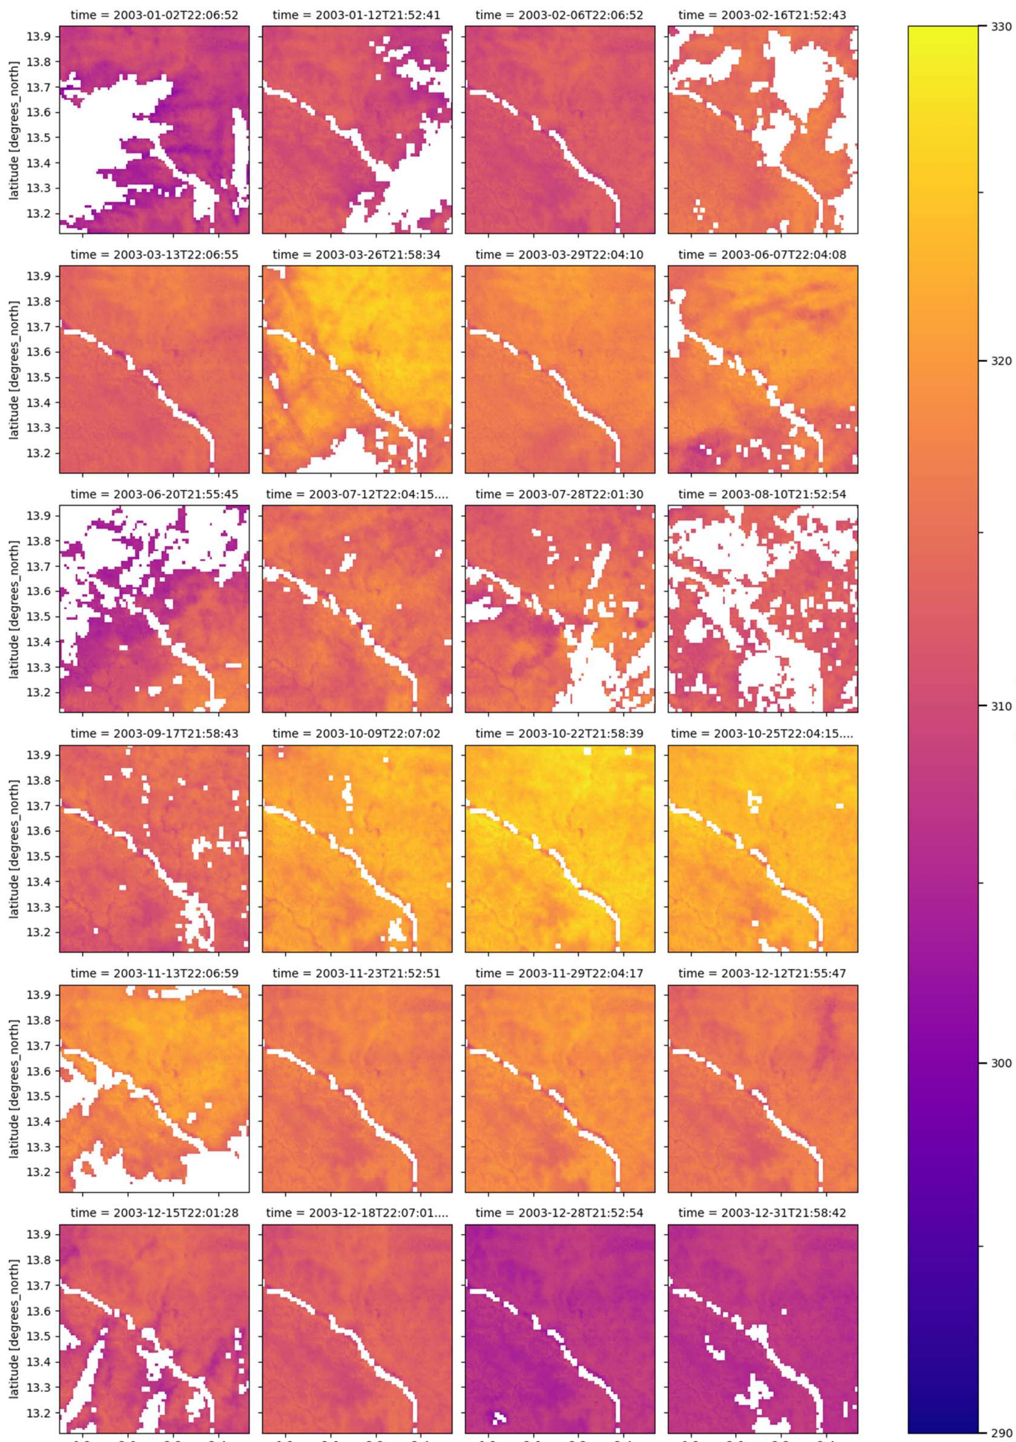

Figure 11 Time series of sharpened ESA CCI LST maps at 300 m resolution for the city of Niamey, Niger (2003). Each panel corresponds to a single cloud-screened acquisition date.

4.3.4 Next steps for LST integration in the scientific analysis

4.3.4.1 Energy conservation

Although the emissivity modulation approach has been used in the literature (Nichol et al., 2009), it suffers from a physically important shortcoming: it does not guarantee conservation of radiative energy. To explore this, we test whether the spatially averaged fine-scale radiance recovers the coarse-scale radiance.

Starting from Equation (1), the radiance at coarse scale is:

$$L_{1\mathrm{km}} = \varepsilon_{1\mathrm{km}}\, \sigma\, LST_{1\mathrm{km}}^4 \tag{8}$$

The spatially averaged fine-scale radiance (i.e., the radiance that should be recovered if energy is conserved) is:

$$\bar{L}_{300\mathrm{m}} = \frac{1}{N}\sum_{i=1}^{N} \varepsilon_{300\mathrm{m},i}\ \sigma\, LST_{300\mathrm{m},i}^4 \tag{9}$$

Substituting Equation (7) into Equation (9):

$$\bar{L}_{300\mathrm{m}} = \frac{\sigma\, LST_{1\mathrm{km}}^4\, \varepsilon_{1\mathrm{km}}}{N}\sum_{i=1}^{N}\frac{\varepsilon_{300\mathrm{m},i}}{\varepsilon_{300\mathrm{m},i}} = \frac{\sigma\, LST_{1\mathrm{km}}^4\, \varepsilon_{1\mathrm{km}}}{N}\cdot N = L_{1\mathrm{km}} \tag{10}$$

At first glance this appears to show energy conservation. However, this is only exact because we used the simplified approximation $T^4 \propto L/\varepsilon$. In practice, the emissivity modulation relation in Equation (7) is derived by transferring brightness temperature — a quantity linear in T — rather than radiance — a quantity proportional to T^4. The non-linearity of the T^4 relationship means that the spatial averaging of temperature is not equivalent to averaging of radiance, leading to an energy imbalance. Specifically, for a heterogeneous pixel with sub-pixel temperature variability δT_i, Jensen's inequality gives:

$$\frac{1}{N}\sum_{i=1}^{N} T_i^4 \geq \left(\frac{1}{N}\sum_{i=1}^{N} T_i\right)^4 \tag{11}$$

with equality only when all sub-pixel temperatures are identical. Equation (11) demonstrates that averaging temperature before raising to the fourth power consistently underestimates the mean radiance. For urban surfaces, where sub-pixel temperature contrasts can exceed 10 K, this bias is non-negligible.

Furthermore, Equation (7) assumes that T_B is identical for all sub-pixels within a 1 km footprint, which implicitly treats the coarse pixel as radiatively homogeneous — an assumption that is could be violated in heterogeneous urban scenes. To address the energy imbalance identified above, the sharpening procedure must be reformulated in radiance space, where the additive property of energy is satisfied by construction. The corrected methodology proceeds as follows.

Step 1 — Compute coarse-scale radiance. From Equation (1):

$$L_{1\mathrm{km}} = \varepsilon_{1\mathrm{km}}\, \sigma\, T_{1\mathrm{km}}^4 \tag{12}$$

Step 2 — Compute a first-guess fine-scale radiance by emissivity scaling. The spatial distribution of radiance at sub-pixel scale is initialised by scaling the coarse radiance by the ratio of fine-to-coarse emissivity:

$$L_{300\mathrm{m},i}^{(0)} = L_{1\mathrm{km}} \cdot \frac{\varepsilon_{300\mathrm{m},i}}{\varepsilon_{1\mathrm{km}}} \tag{13}$$

This first-guess allocates more radiance to sub-pixels with higher emissivity and less to sub-pixels with lower emissivity, consistent with the physical expectation that surfaces with higher emissivities radiate more strongly.

Step 3 — Enforce the energy conservation constraint. A physically consistent downscaling requires that the total radiance summed over all N sub-pixels within the coarse footprint equals N times the coarse-pixel radiance:

$$\sum_{i=1}^{N} L_{300\mathrm{m},i} = N \cdot L_{1\mathrm{km}} \tag{14}$$

In general, the first-guess $L_{300\mathrm{m},i}^{(0)}$ from Equation (13) does not satisfy Equation (14) because $\varepsilon_{1\mathrm{km}}$ is not equal to the arithmetic mean of $\varepsilon_{300\mathrm{m},i}$ when the two emissivity datasets are derived from different sensors and aggregation procedures. A multiplicative correction factor C is therefore defined as:

$$C = \frac{N \cdot L_{1\mathrm{km}}}{\sum_{i=1}^{N} L_{300\mathrm{m},i}^{(0)}} \tag{15}$$

Substituting Equation (13) into Equation (15):

$$C = \frac{N \cdot \varepsilon_{1\mathrm{km}}}{\sum_{i=1}^{N} \varepsilon_{300\mathrm{m},i}} \tag{16}$$

Equation (16) provides a clear physical interpretation: C is the ratio of the coarse emissivity (scaled by N) to the sum of fine emissivities. When the MODIS and ASTER emissivity products are perfectly consistent, $C = 1$ and no correction is needed.

Step 4 — Apply the correction to obtain energy-conserving fine-scale radiance. The corrected fine-scale radiance is:

$$L_{300\mathrm{m},i} = C \cdot L_{300\mathrm{m},i}^{(0)} \tag{17}$$

It is straightforward to verify that Equations (14) to (17) together guarantee energy conservation:

$$\sum_{i=1}^{N} L_{300\mathrm{m},i} = C \cdot \sum_{i=1}^{N} L_{300\mathrm{m},i}^{(0)} = \frac{N \cdot L_{1\mathrm{km}}}{\sum_i L_{300\mathrm{m},i}^{(0)}} \cdot \sum_i L_{300\mathrm{m},i}^{(0)} = N \cdot L_{1\mathrm{km}} \tag{18}$$

Step 5 — Retrieve fine-scale LST. The downscaled surface temperature $T_{300\mathrm{m},i}$ is obtained by inverting Equation (1) with the corrected fine-scale radiance:

$$T_{300\mathrm{m},i} = \left(\frac{L_{300\mathrm{m},i}}{\varepsilon_{300\mathrm{m},i}\,\sigma}\right)^{1/4} \tag{19}$$

4.3.4.2 Fine emissivity ϵ_{300m}

In the current implementation of the emissivity modulation approach, we rely on the climatological emissivity dataset derived from ASTER AG100, which represents land surface emissivity at 100 m resolution over the 2000–2008 period. As a result, for years beyond 2008, the sharpening process does not account for the temporal evolution of surface emissivity. This limitation may introduce biases in the sharpened LST maps, particularly in areas undergoing significant land surface changes. To address this, we plan to develop a yearly, seasonal emissivity dataset following a similar approach to that used for MYD21AD. Specifically, this will be based on the ASTER AST_05 product3, enabling a more temporally consistent representation of emissivity in the sharpening workflow.

4.3.4.3 Validation and quality layer

Direct validation of $LST_{300\mathrm{m}}$ against independent 300 m reference observations is currently infeasible for two reasons. First, there are no operational satellite missions providing daily LST observations at 300 m or finer spatial resolution — the highest spatial resolution daily LST product available remains at 1 km. Second, although sensors such as ASTER and ECOSTRESS can provide LST at 70–100 m resolution, their orbital revisit periods (of order 16 days for ASTER) make the probability of a temporally coincident matchup (time difference < 30 min) with any given $LST_{300\mathrm{m}}$ acquisition extremely low.

In the absence of direct validation, we propose an internally consistent quality indicator based on the energy conservation ratio R, defined as:

$$R = \frac{\sum_{i=1}^{N} L_{300\mathrm{m},i}}{N \cdot L_{1\mathrm{km}}} \tag{20}$$

By construction (Equation 18), $R = 1$ exactly after the correction step. However, R can be computed before applying the correction factor C (i.e., using $L_{300\mathrm{m},i}^{(0)}$ from Equation 13) to quantify the degree of emissivity inconsistency between the MODIS and ASTER datasets. Pixels for which R departs significantly from unity indicate strong inter-dataset emissivity discrepancies and should be flagged as lower-quality retrievals. The specific threshold on $|R - 1|$ will be calibrated using the available time series data and will be documented in the quality flag layer accompanying the $LST_{300\mathrm{m}}$ product.

Uncertainty propagation

The ESA CCI LST product includes per-pixel uncertainty estimates $u(LST_{1\mathrm{km}})$. These uncertainties propagate through the sharpening chain and must be consistently transferred to the output $LST_{300\mathrm{m},i}$ layer. Starting from Equation (19) and treating the input radiance $L_{300\mathrm{m},i}$ as the primary uncertain quantity, the uncertainty on the retrieved fine-scale temperature is obtained by first-order error propagation:

$$u(T_{300\mathrm{m},i}) = \left|\frac{\partial T_{300\mathrm{m},i}}{\partial L_{300\mathrm{m},i}}\right| \cdot u(L_{300\mathrm{m},i}) \tag{21}$$

Differentiating Equation (19):

$$\frac{\partial T_{300\mathrm{m},i}}{\partial L_{300\mathrm{m},i}} = \frac{1}{4}\left(\frac{L_{300\mathrm{m},i}}{\varepsilon_{300\mathrm{m},i}\,\sigma}\right)^{-3/4} \cdot \frac{1}{\varepsilon_{300\mathrm{m},i}\,\sigma} = \frac{T_{300\mathrm{m},i}}{4\,L_{300\mathrm{m},i}} \tag{22}$$

[3] https://www.earthdata.nasa.gov/data/catalog/lpcloud-ast-05-004

Substituting Equation (22) into Equation (21):

$$u(T_{300\mathrm{m},i}) = \frac{T_{300\mathrm{m},i}}{4\,L_{300\mathrm{m},i}} \cdot u(L_{300\mathrm{m},i}) \tag{23}$$

The uncertainty on the fine-scale radiance $u(L_{300\mathrm{m},i})$ is in turn derived by propagating $u(LST_{1\mathrm{km}})$ through Equations(12)–(17), accounting for both the emissivity uncertainty from the ASTER/MODIS products and the coarse-scale LST uncertainty from the CCI dataset. The full uncertainty chain will be implemented and documented at a later stage.

4.3.4.4 Directional effects assessment

The ESA CCI LST dataset is a composite assembled from multiple wide-swath sensors (MODIS, AATSR/SLSTR, SEVIRI, among others), which inherently means that the same surface location is observed under varying view zenith angles across different acquisitions. This is particularly relevant for urban areas, where the three-dimensional structure of the built environment creates directional anisotropy in the emitted thermal radiation field. In other words, the LST retrieved at off-nadir view angles differs systematically from the true nadir-equivalent LST, even when surface conditions are identical, because different proportions of sunlit facades, shadowed walls, and ground surfaces enter the sensor field of view.

This directional effect must be assessed and, if significant, corrected before meaningful cross-sensor comparisons can be made or before the sharpened $LST_{300\mathrm{m}}$ is compared with model outputs such as those from UrbClim.

To quantify this effect, we will perform TIR radiative transfer simulations for a representative subset of urban geometries and meteorological conditions using a 3D urban radiative transfer model. The simulations will systematically vary the view zenith angle while keeping surface temperatures, emissivities, and atmospheric conditions constant, thereby isolating the purely directional contribution to the observed LST signal.

4.4 Urban growth modelling progress

4.4.1 Current status of GeoDynamiX modelling

The GeoDynamiX workflow has now progressed from calibration set-up to full scenario production. Pilot-city applications were first used to test the automated calibration workflow and to verify the consistency of the scenario logic before scaling to the full set of target cities. At the time of writing, all SSP2 and SSP3 simulations have been completed, and the corresponding outputs have been organised for the simulated cities.

A first subset of the GeoDynamiX outputs has already been converted into terrain inputs for the future UrbClim chain, which means that urban-growth modelling is already contributing directly to the scientific analysis of future heat conditions.

For this deliverable purpose, stronger scientific conclusions on projected urban form, densification patterns, or divergence between SSP2 and SSP3 remains limited until the final SSP3 runs are fully analysed.

4.4.2 Limitations, challenges and mitigation related to modelling-grid extent

A practical challenge in applying GeoDynamiX across the full CAIAC city sample was ensuring that the fixed modelling grid remained large enough to contain not only the present-day urban area, but also the full projected urban expansion up to 2075. The initial modelling grids were defined using the present-day urban extent, population information and expected future growth as a first estimate. However, because some cities show strong or spatially uneven expansion, the first urban-growth simulations occasionally indicated that future built-up areas approached or exceeded the initial grid boundary.

This issue was addressed through an iterative grid-definition and quality-control process. After the first GeoDynamiX simulations, the final-year urban expansion maps were inspected. Where projected growth reached the edge of the modelling domain, the city grid was enlarged or adjusted, after which the urban-growth simulation and subsequent terrain-processing steps were repeated. This back-and-forth procedure ensured that the final grids include sufficient space for the projected urban expansion while maintaining consistency with the UrbClim grid constraints.

As a result, all final modelling grids now fully encompass the simulated urban expansion for the SSP2 and SSP3 scenarios considered in CAIAC. Remaining limitations are mainly linked to the trade-off between domain size, spatial resolution and computational cost. For very large urban areas, this was handled by adopting a coarser grid resolution where required (from 300 m to 600 m), so that both the current urban footprint and future expansion could be retained within the maximum model-domain size.

4.4.3 Preliminary observations from currently available scenario outputs

The vast majority of the simulated cities will grow strongly, in line with the projected demographic changes. Table 2 shows a summary of SSP2 growth numbers within the modelling grid. Mind that this is not the same as any administrative level.

Even in SSP2, the middle-of-the-road scenario, the expected average population of the 54 urban areas in 2100 is 3.6 times larger than the current population in 2020. The increase of the population is a combination of absolute growth in the country (birth ratio and international migration) and the expected urbanisation ratio (internal migration from rural areas). Translating population growth into projections for the urban area sizes (see D2.1 for more details), the urban areas would become on average 3.7 times larger in SSP2.

There are strong differences between the cities. Population changes between 2020 and 2100 ranges from an overall decline (Port Louis) to a growth ratio higher than 11 (Niamey). Urban areas usually have somewhat similar growth ratios under SSP2, but these can still differ for several reasons, as indicated in the ratio of urban growth to population growth. Some cities who currently have a strong densification process (e.g. Luanda and Kinshasa) see a significantly smaller increase of the built-up area than of the population. Values above 1 can indicate unnecessary urban sprawl (e.g. Abuja or Ouahigouya) but can also be a sign of a population decrease at some point in time as this does not lead to the demolition of buildings in the model. This is projected to happen soon in Port Louis, and by the end of the century in e.g. Victoria, Tunis, Johannesburg and Casablanca. Juba has the largest urban growth to population growth ratio, but this could be an overestimation as the model calibration during the 2005-2020 period is influenced by the effects of the war in South Sudan: the city first grows strongly, but later the population declines fast during this period. Also, in Galkayo, Somalia, this effect is somewhat present.

Table 2. Population and urban area in 2020 and 2100 (SSP2), their respective growth ratios, and the ratio of urban growth compared to population growth in all cities (within the modelling grid).

City	Population			Urban area (ha)			
	2020	2100 SSP2	ratio	2020	2100 SSP2	ratio	Urban growth to pop. growth ratio
Abeche	392,244	2,731,878	6.96	2,412	17,019	7.06	1.01
Abuja	2,905,989	15,496,460	5.33	45,864	288,981	6.30	1.18
Ali Sabieh	89,588	121,544	1.36	198	297	1.50	1.11
Antananarivo	4,477,264	20,440,820	4.57	14,364	52,506	3.66	0.80
Asmara	426,703	938,925	2.20	3,501	8,703	2.49	1.13
Bamako	3,922,627	21,698,210	5.53	35,055	159,831	4.56	0.82
Bata	763,546	2,747,276	3.60	4,221	13,095	3.10	0.86
Bo	390,656	1,100,689	2.82	2,646	8,235	3.11	1.10
Bouake	766,794	2,002,742	2.61	8,028	21,132	2.63	1.01
Brazzaville	2,766,534	7,699,460	2.78	15,588	39,123	2.51	0.90
Cairo	49,758,150	137,590,800	2.77	230,643	673,182	2.92	1.06
Carnot	90,135	479,148	5.32	576	2,844	4.94	0.93
Casablanca	4,929,256	6,989,457	1.42	17,217	25,182	1.46	1.03
Darou Mousty	93,126	380,329	4.08	621	2,736	4.41	1.08
Dodoma	366,526	1,698,194	4.63	7,830	41,175	5.26	1.13
Franceville	263,411	624,682	2.37	1,197	2,475	2.07	0.87
Gaborone	406,249	677,615	1.67	13,950	23,247	1.67	1.00
Gabu	122,088	341,288	2.80	792	1,971	2.49	0.89
Galkayo	184,822	818,700	4.43	1,215	10,044	8.27	1.87
Gbarnga	125,199	384,241	3.07	1,404	4,869	3.47	1.13
Gitega	554,365	3,114,548	5.62	972	5,040	5.19	0.92
Harare	2,566,414	6,298,063	2.45	35,793	100,026	2.79	1.14
Johannesburg	17,295,570	21,683,470	1.25	290,277	386,496	1.33	1.06
Juba	398,692	960,763	2.41	6,687	36,936	5.52	2.29
Kankan	299,171	1,087,129	3.63	3,429	12,402	3.62	1.00
Khartoum	7,134,878	25,774,570	3.61	69,003	224,883	3.26	0.90
Kigali	2,313,338	13,393,710	5.79	13,968	83,277	5.96	1.03

City	Population			Urban area (ha)			Urban growth to pop. growth ratio
	2020	2100 SSP2	ratio	2020	2100 SSP2	ratio	
Kinshasa	10,543,560	64,974,990	6.16	36,423	169,038	4.64	0.75
Kumasi	5,003,659	16,948,780	3.39	50,274	157,923	3.14	0.93
Luanda	10,599,450	97,779,670	9.22	17,100	83,304	4.87	0.53
Lusaka	3,360,621	14,261,260	4.24	27,495	92,376	3.36	0.79
Maseru	388,154	784,890	2.02	9,747	20,916	2.15	1.06
Mbabane	154,237	303,657	1.97	1,296	3,132	2.42	1.23
Mbale City	966,940	5,711,120	5.91	1,512	8,604	5.69	0.96
Moroni	268,787	896,300	3.33	1,890	5,787	3.06	0.92
Nairobi	7,604,771	26,416,540	3.47	45,045	147,528	3.28	0.94
Naqamte	387,957	1,811,271	4.67	1,269	6,426	5.06	1.08
Niamey	1,966,381	22,756,930	11.57	14,058	176,841	12.58	1.09
Nouakchott	1,334,219	4,726,097	3.54	13,671	50,490	3.69	1.04
Ouahigouya	214,291	965,429	4.51	2,187	18,126	8.29	1.84
Parakou	412,193	1,966,257	4.77	5,940	25,713	4.33	0.91
Port Louis	1,255,435	908,330	0.72	13,545	16,173	1.19	1.65
Praia	251,429	330,002	1.31	1,332	1,836	1.38	1.05
Sao Tome	202,590	358,599	1.77	954	1,719	1.80	1.02
Serrekunda	1,590,230	4,558,507	2.87	13,815	35,154	2.54	0.89
Sokode	212,996	657,764	3.09	2,025	6,111	3.02	0.98
Tete	488,182	2,853,066	5.84	4,995	24,525	4.91	0.84
Tiaret	321,574	549,074	1.71	2,340	4,311	1.84	1.08
Tripoli	1,769,428	2,784,520	1.57	44,964	77,625	1.73	1.10
Tunis	3,103,649	3,930,132	1.27	30,843	40,167	1.30	1.03
Victoria	90,070	101,789	1.13	297	423	1.42	1.26
Windhoek	391,831	895,243	2.28	7,524	16,947	2.25	0.99
Yaounde	4,678,073	17,489,740	3.74	20,538	65,529	3.19	0.85
Zomba	442,841	2,298,535	5.19	1,170	5,931	5.07	0.98

4.5 Flood modelling progress

4.5.1 Status of hydrological modelling with SWAT+

The SWAT+ modelling framework has been fully implemented across all 20 study cities. Watershed delineation and climate forcing integration are complete for all domains, covering historical conditions and four climate (SSP) scenarios (1-2.6, 2-4.5, 3-7.0, and 5-8.5). Daily runoff time series, return-period quantiles (2–200 years), and bias-adjusted discharge series (ISIMIP3b) are available for all cities and scenarios. The complete hydrological chain provides the spatially explicit input required for subsequent HEC-RAS hydrodynamic modelling.

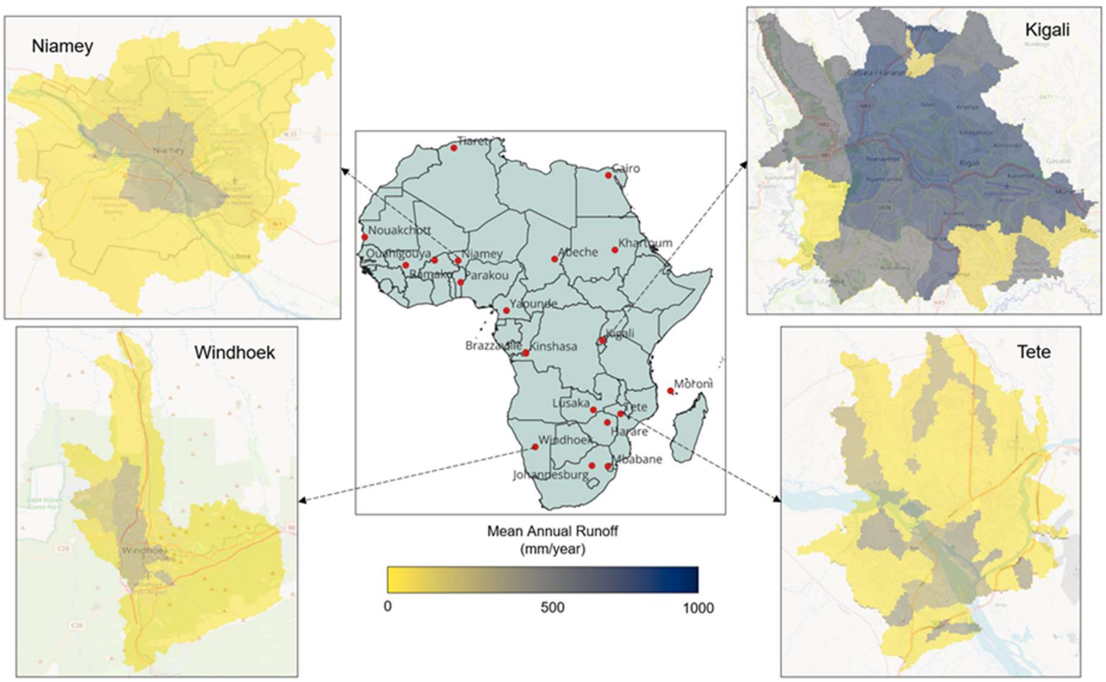

Figure 12 Historical mean annual runoff (mm/year; 1983–2014) simulated by SWAT+ for Niamey, Kigali, Windhoek, and Tete. The continental overview reveals a strong spatial gradient in mean annual runoff, reflecting the diversity of hydroclimatic regimes across the study domain

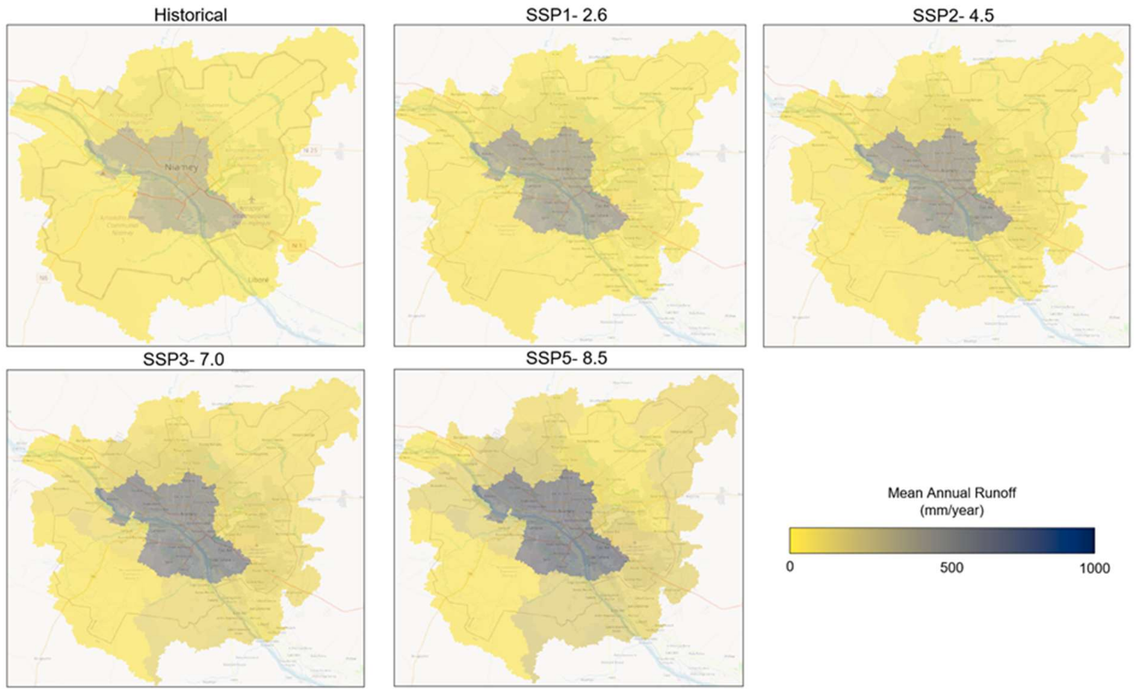

Figure 13 Mean annual runoff (mm/year) simulated by SWAT+ for Niamey across all five modelled periods

As depicted in Figure 13, the mean annual runoff for Niamey across all five modelled periods shows distinct changes. Under historical conditions, runoff exhibits a marked spatial gradient, with higher values concentrated in the urban core and lower values in the surrounding semi-arid periphery. This pattern reflects the combined influence of land use, soil infiltration capacity, and topographic routing within the delineated watersheds in SWAT+ (called HRUs). Across the future scenarios, a general intensification of runoff is projected, particularly under SSP3-7.0 and SSP5-8.5, where the spatial extent of high-runoff zones expands notably compared to the historical baseline. SSP1-2.6 and SSP2-4.5 show intermediate responses, with spatial distributions broadly similar to the historical period but with localized increases in peak runoff areas. These results are consistent across all 20 study cities, confirming the robustness of the modelling framework and its sensitivity to the full range of prescribed climate forcing.

4.5.2 Status of hydrodynamic modelling with HEC-RAS

Using return-period quantiles derived from the frequency analysis of surface runoff and river discharge, pluvial and fluvial flood scenarios were simulated for multiple return periods using HEC-RAS. The simulations were developed using 2D unsteady-flow analysis based on 30-meter DEMs and land cover data. For each study area, 2D flow areas were delineated and converted into computational meshes to represent the movement of water across the land surface.

Table 3. *Current status of the HEC-RAS models across all cities.*

City	Pluvial	Fluvial
Tiaret	Done	N.A.
Parakou	Done	N.A.

Ouahigouya	Done	N.A.
Yaounde	Done	N.A.
Abeche	Done	N.A.
Moroni	Done	N.A.
Brazzaville	Done	Done
Kinshasa	Done	Done
Cairo	Done	Under refinement
Mbabane	Done	N.A.
Nouakchott	Done	N.A.
Tete	Done	Under refinement
Windhoek	Done	N.A.
Niamey	Done	Done
Kigali	Done	N.A.
Johannesburg	Done	N.A.
Khartoum	Done	Under refinement
Lusaka	Done	N.A.
Harare	Done	N.A.
Bamako	Done	Under refinement

Fluvial flood scenarios are currently undergoing an additional refinement process to improve realism, as the 30-meter DEM introduces uncertainty in representing river bathymetry and channel geometry. To address this limitation, the riverbed idealisation is being adjusted by comparing the simulated 2-year flood extent with satellite imagery from flood events of similar magnitude, with the aim of ensuring that the modelled inundation extent more closely matches observed conditions. The following images show a refined map for Niamey for different return periods:

After the refinement of the fluvial maps, pluvials maps will also go under refinement to go to higher resolution (10 m).

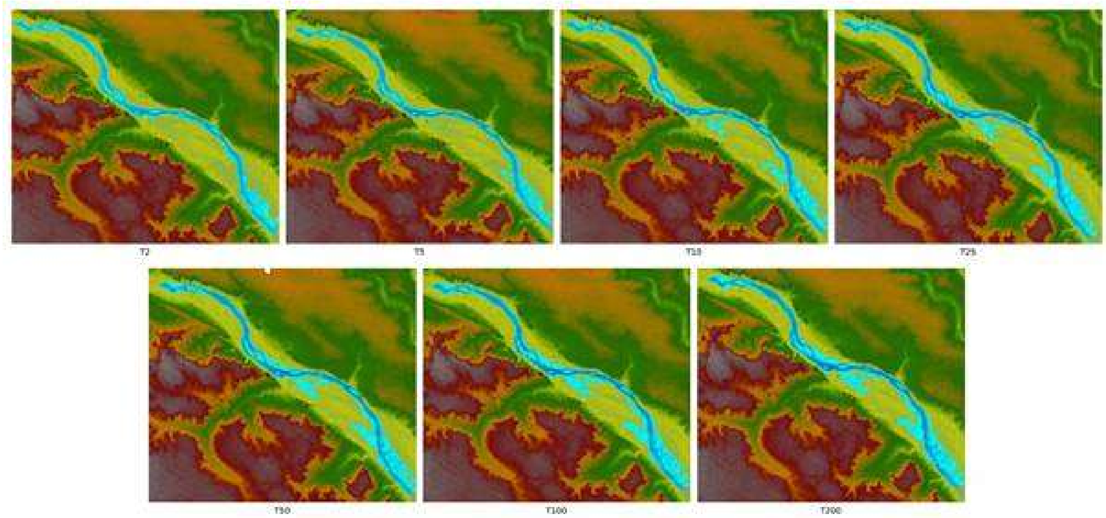

Figure 14 Refined maps for Niamey, each for different return period

4.5.3 Flood-threshold / flood-risk workflow status

As noted earlier, models of African cities were characterized by major limitations in observed data availability, and in many cases no database was available for surface runoff or river discharge. To address this gap, modeled and bias-corrected surface runoff data derived from the SWAT+ simulations were used for the watershed covering each city. A similar limitation was encountered for river discharge, which was also generated through SWAT+ modeling for cities crossed by rivers. For both variables, daily time series with acceptable accuracy were obtained, thereby resolving the major challenge of data scarcity.

On the other hand, implementation of the rapid inundation labeling (RIL) method, described in previous deliverable, required inundation depth data for different flood return periods. An initial attempt was made to use existing, freely available or commercial databases. Aqueduct was the first database examined. This global flood inundation dataset, freely available at a spatial resolution of 1000 m, was found to be of insufficient quality for the African context. After evaluation, it was observed that a large proportion of grid cells contained no data, which made the dataset unsuitable for application across most African cities. In addition, its coarse spatial resolution was considered inadequate for detailed urban-scale analysis and was expected to considerably reduce the accuracy of the work.

Further efforts were also made to obtain similar inundation datasets from other sources, but these were unsuccessful. Therefore, it was decided that the required inundation data should be generated independently. DEMs for the African cities, together with the available urban land-use data, were used as model inputs. In addition, river corridors and flow paths were defined as breaklines within the HEC-RAS environment. Based on this approach, flood inundation maps were simulated and extracted for multiple flood return periods at a much finer spatial resolution for the African cities.

The RIL method was based on the combined use of runoff data (or discharge data) and flood inundation depths associated with different return periods. In this project, flood events with return periods of 2, 5, 10, 25, 50, 100, and 200 years were considered. For runoff, the corresponding quantiles were extracted from the daily runoff time series for each cell within the city domain (similarly for discharge). The spatial resolution adopted for the extracted data and generated maps was 100 m. For each cell in the city, a curve was constructed. The plotted points represented the different flood return periods.

On the horizontal axis, the inundation depth of that cell for a given flood return period was placed, while on the vertical axis, the corresponding runoff quantile from the daily time series of the same cell was assigned for pluvial flooding. For fluvial flooding, a similar procedure was followed, except that discharge quantiles from the daily discharge time series of the corresponding cell were used instead of runoff quantiles.

Several mathematical equations and fitting approaches were tested for these point sets. However, a three-parameter Horton-type function was ultimately selected because it provided the most consistent and physically meaningful graphical behavior. These fitted curves were then extrapolated, and their intercept with the vertical axis was defined and extracted as the flood-generation threshold for each cell. In the next step, the threshold value of each cell was compared on a daily basis with the corresponding runoff value in the pluvial flood workflow, or with the discharge value in the fluvial flood workflow. If the threshold was equal to or lower than the runoff or discharge value on a given day, that cell was classified as wet. By repeating this procedure for all cells and all days, daily flood hazard maps could be generated.

4.5.4 Preliminary flood-related outputs available for WP3

During the initial phase of the work, data for all target cities were collected and organized. In parallel, the SWAT+ and GeoHECRAS simulations were completed for all cities. As a result, a large and comprehensive database was established for each city, representing a valuable outcome of the project in itself.

At present, the RIL method has been implemented for one city, namely Niamey, Niger. The results obtained were found to be acceptable and encouraging. The analysis was first conducted for the historical period. Thereafter, future projection data were incorporated, and daily flood hazard maps for the city of Niamey were generated up to the end of the year 2100 under four widely used CMIP6 SSP scenarios: SSP1-2.6, SSP2-4.5, SSP3-7.0, and SSP5-8.5. These are among the standard future climate scenarios commonly used in recent climate impact assessments. Two representative graphs of the fitted function under two different scenarios are shown below.

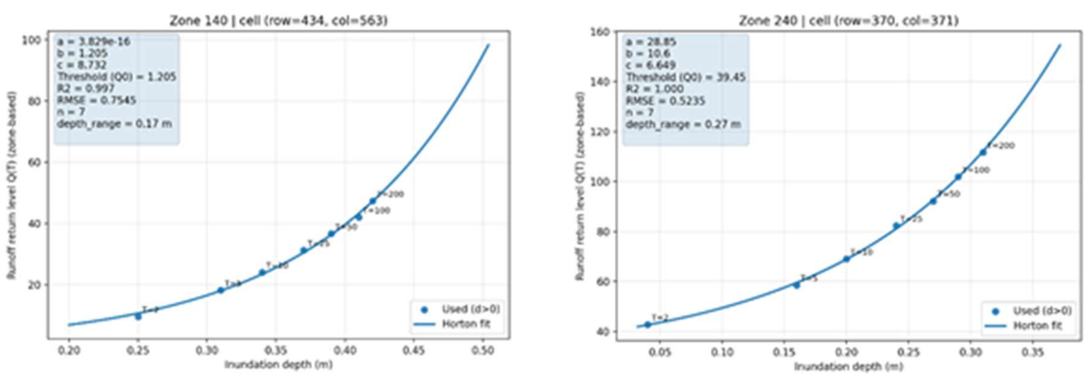

Figure 15. Representative graphs of the fitted Horton-type function, as part of the RIL framework, for two different cells in the city. The 3-parameter model performs very well and retains physics-based meaningful representation of the flood threshold values derived from the curve fitting procedure.

Overall, the results were considered promising. Across all scenarios, good fitting performance was achieved in more than 90% of cases, with coefficient of determination values generally exceeding %90. In addition, an examination of the annual number of wet cells under different scenarios indicated a clear increase in both the number of flood-prone cells and the number of flood-prone days as the scenarios became more severe. This behavior suggests that Niamey is expected to become increasingly flood-prone under more critical future climate conditions.

5 CONCLUSIONS AND NEXT STEPS

The CAIAC scientific workflow has now clearly progressed from methodological preparation to operational model production. For the urban-heat component, baseline UrbClim simulations have been completed for the full target-city sample, together with post-processing, WBGT computation and EO-based downscaling. These results already support a first cross-city assessment of present-day urban heat and heat stress in African cities. In parallel, the future heat-modelling chain is advancing well: future climate signals have been prepared, future land-use outputs have been translated into UrbClim terrain inputs, and the first future simulations are already running.

Urban-growth modelling has also reached an important level of maturity. GeoDynamiX scenario production for SSP2 and SSP3 has been completed for the target-city sample, and the resulting land-use projections are already being integrated into the future heat workflow. The first results confirm that many cities are expected to undergo very strong demographic and spatial expansion during the coming decades, underlining the importance of explicitly accounting for urban growth when assessing future climate risk. At the same time, a fuller scientific interpretation of differences between scenarios, densification patterns and city-specific development pathways will require additional synthesis work in the next reporting period.

For flooding, substantial progress has been achieved across the full modelling chain. SWAT+ hydrological simulations, return-period analyses and HEC-RAS hydrodynamic simulations are now available for the study cities, while the flood-threshold / flood-risk workflow has been successfully implemented for Niamey as a first demonstration case. These initial results indicate that the framework is capable of producing meaningful present-day and future flood-hazard information, even in data-scarce environments. However, some fluvial simulations still require refinement, and the RIL-based workflow must now be extended and consolidated for the wider city sample before broader scientific conclusions can be drawn.

The LST component remains the least mature part of the scientific workflow, but important methodological progress has been made. The sharpening framework for ESA CCI LST has been implemented, diagnostic work has advanced, and key next steps have been identified, including energy-conserving reformulation, emissivity improvements, uncertainty propagation and assessment of directional effects.

Overall, D3.1 shows that the CAIAC modelling chain is operational, scientifically productive and already able to generate meaningful intermediate findings on urban heat, urban growth and flood hazard. The next phase of work will focus on completing future simulations, finalising validation and uncertainty analyses, extending the flood-risk workflow to additional cities, and integrating the different model components into a more complete scientific assessment.

REFERENCES

Nichol, J. E., Fung, W. Y., Lam, K. se, & Wong, M. S. (2009). Urban heat island diagnosis using ASTER satellite images and 'in situ' air temperature. *Atmospheric Research*, *94*(2), 276–284. https://doi.org/10.1016/J.ATMOSRES.2009.06.011

Wang, K., & Liang, S. (2009). Evaluation of ASTER and MODIS land surface temperature and emissivity products using long-term surface longwave radiation observations at SURFRAD sites. *Remote Sensing of Environment*, *113*(7), 1556–1565. https://doi.org/10.1016/J.RSE.2009.03.009

vision on technology
for a better world

vito