Mapping global cooling and overheating risk evidence inequalities: understanding and building resilience in a heating world

Anna Conzatti,[1,*] Daniel Fosas[2]

[1] University of Exeter, European Centre for Environment and Human Health, Penryn, United Kingdom
[2] University of Edinburgh, School of Engineering, Edinburgh, United Kingdom

[*]*Corresponding email: a.conzatti@exeter.ac.uk*

SUMMARY
Rising temperatures and demographic pressures are intensifying the need for accessible and effective cooling. Despite scientific evidence growing exponentially, it is unclear whether research on residential cooling/overheating is commensurate with heat exposure risk and population growth under a heating climate. Hence, this study maps the geographical location of research on residential cooling/overheating and, to deliver a scalable method, we tested two approaches to extract study locations: rule-based Natural Language Processing (NLP) and Large Language Models (LLMs). LLM-based extraction matched or exceeded human accuracy (100% vs 98% for n=48 and n=54 papers), while rule-based NLP dropped to 39.8% due to the complexity of unstructured text. The resulting global maps reveal pronounced mismatches between where research is concentrated and where cooling needs are greatest. By integrating LLM-assisted synthesis with spatial analysis, this work provides a practical framework for identifying evidence gaps and supporting equitable, context-specific strategies for overheating resilience.

KEYWORDS
Cooling inequalities, overheating, AI-based text mining, spatial analysis, building resilience

1 INTRODUCTION
Climate change is increasingly affecting human health, with rising rates of health-related illness (Vicedo-Cabrera et al. 2021; Alrasheed and Mourshed 2023; Kenny et al. 2024). Among these, heat is emerging as a major threat to human health as extreme temperature events intensify and populations grow. Today, 56% of the global populations lives in urban areas – a figure expected to reach 68% by 2050 (UNDESA 2025). This expansion, often concentrated in regions already exposed to high thermal stress, amplifies vulnerability and exacerbates inequalities in climate resilience (Tian et al. 2025). In this context, access to cooling—particularly air conditioning (AC)—has become a critical public health need and a central component of climate adaptation (UNEP 2023; Falchetta et al. 2024).

Cooling demand is projected to increase sharply. According to the International Energy Agency (IEA) (2024), global space cooling demand will grow by 3.7% annually until 2035, with developing economies driving 90% of this rise due to increasing affordability and population growth. Yet access to active cooling remains highly uneven. As shown by Falchetta et al. (2024), although AC adoption may exceed 40% of households globally by 2050, billions will still lack access. Moreover, their findings highlight that cooling demand projections often underestimate wider social, infrastructural and building-level barriers to thermal resilience.

Despite the urgency of the challenge, the distribution of scientific knowledge on residential cooling and overheating may not necessarily be proportionally distributed according to heat exposure or risk. Owing to structural differences in research funding and diversity in the way research activity is framed, systematic differences in the body of evidence may emerge. If that were the case, many of the areas facing the greatest exposure, rapid urbanisation, and population growth may remain under-represented and poorly understood (UNDESA 2025). This could create a systematic gap between where cooling-related risks are most acute and where empirical research is produced. Moreover, previous work has demonstrated the feasibility of machine-learning–based geographical mapping for broad climate–health literatures; however, such approaches have not yet been evaluated and validated for extracting the geographical focus of residential cooling and overheating studies using full-text evidence (Berrang-Ford et al. 2021).

As a result, it remains difficult to assess whether research is aligned with present and future heat vulnerability, or to identify priority regions where evidence is most urgently needed. Despite a growing body of work on this topic, key connections between cooling demand, health, and social vulnerability remain underexplored.

This study addresses this gap by providing a global assessment of where scientific evidence on residential cooling is produced and how it aligns with patterns of heat exposure and population growth. The objectives are:

1. to develop and test a scalable method for identifying study locations in selected literature;
2. to approximate the geographical distribution of research on residential cooling;
3. to evaluate inequalities in research evidence between Global North and Global South.

2 MATERIALS/METHODS

This study examines where scientific evidence on residential cooling is produced and how it compares with heat exposure risk and demographic pressure. The unit of analysis is the study location, defined as the geographical setting (country) of the work reported in a publication. When a study covered more than one location, all reported locations were extracted.

Overall design

Given a collection of papers, multiple extraction approaches can be applied, each producing a potentially different classification. Manual extraction provides the ground truth but does not scale; hence we devised a two-step workflow. First, a feasibility study evaluated the performance of different automated extraction approaches against existing dataset from PRISMA-based literature reviews (Moher et al. 2010). Second, the most reliable scalable approach was applied to a new set of papers defined by the research question of this study. As classification performance is not expected to be perfect, the analysis focuses on relative spatial patterns rather than exact counts, particularly the proportion of studies conducted in the Global North versus the Global South. This provides a robust indicator even when some misclassification.

Search query

A PRISMA-informed bibliometric search was conducted in Scopus in September 2024 using the following query: TITLE-ABS-KEY (cooling AND adaptation AND (residential OR housing OR domestic)). The search returned 326 documents published across 150 journals. Only English-language, peer-reviewed papers presenting empirical evidence (field measurements, simulations, surveys, or case studies) were included. Reviews, editorials, and unrelated studies were excluded. For all included papers, study locations were extracted using

all tested methods. These decisions are congruent with the motivations of this study, but by their very nature introduce bias as it establishes a narrow notion of what constitute evidence. Thus, results are limited to what has been considered in-scope: peer reviewed papers in English that are reported in the Scopus database.

Extraction models

To capture the study locations at scale, three extraction approaches were tested: manual researcher classification, a rule-based Natural Language Processing (NLP) workflow, and a newly introduced Large Language Models (LLM)-based extraction model. The LLM model was chosen because it can interpret unstructured text and infer the geographical focus of a study even when the place is not stated in a clear, standard, or predictable pattern. For each paper, the model analysed the title, abstract, and main text. The results were validated through a cross-checking procedure. Two researchers independently coded a subset of papers and compared their selections against the AI output. Disagreement was resolved through discussion and reference to the full text.

NLP

The rule-based NLP approach was implemented using spaCy v3.0 (Honnibal et al. 2020) for Named Entity Recognition (NER; doc.ents, filtering GPE and optionally LOC labels), alongside dictionary and regex-based country matching. Extracted locations were normalised to country level. This approach performs well when locations are explicitly stated but is sensitive to implicit references, non-standard phrasing, and contextual descriptions.

LLMs

To extract location(s), we explored an LLM-assisted pipeline because they have the potential to consider nuances in the text to extract the information that is in scope here. At the core of the task is identifying the country or countries actually studied with in a publication rather than contextual information typically provided in introductions or state of the art sections.

The collection of papers was established in PDFs as it was the common format across all publications in the dataset. We extracted text as markdown to have a text file that can be easily consumed programmatically, while retaining semantic structure (e.g., headers, figures). This was automated using the Python libraries ‘docling’ and ‘markitdown’, as they support a wide range of strategies to automate tasks like text recognition of papers where the originals were scanned images. We postprocessed the exports to disregard the last section of the document titled “Reference” to shorten documents by removing bibliographic references as usage and billing of LLMs increases according to document length. Based on document length, a limit of 32,000 tokens was estimated to be adequate for this task. (A token is the unit of text processed by the LLM. Tokens are not the same as words, and a single word can map one or more tokens. OpenAI reports 1 token can be roughly equal to 0.75 words as a rule of thumb.)

For the LLMs, we explored two variants: an enterprise, online solution and a local, offline one. The first was The University of Edinburgh’s ELM service, which provides secure access to OpenAI’s ChatGPT-family models with no data retention. For this proof of concept, we selected ChatGPT-4.1-nano to evaluate whether a lower-cost model can reliably perform this constrained extraction task ($0.10 and $0.40 per 1M input and output tokens, respectively). In parallel, we tested a local offline alternative using the Ollama application, selecting GPT-OSS:20b (20B parameters) to provide an indicative comparison between a managed paid approach and a free self-hosted one.

Both models were used programmatically from Jupyter notebooks via Python APIs (OpenAI's for the online and Ollama's for the local). The prompt was: "Persona: You are a research assistant processing scientific journal papers. Question: What is the name of the country or countries the scientific paper is about? In doing so, consider only what the paper develops in its methods and results sections, not the contextual literature for introduction and background. Ensure the reply contains just a list of country names, nothing else.". The only difference was in the way papers were consumed. For the local model, papers were provided as context preceding the prompt to ensure the latter was included when using the 32k-token limit. For the online one, it was provided via the dedicated entry point in the API.

Outputs were stored as a list of country names matched to the paper, alongside statistics about token consumption: input text (publication and prompt) and output (country name, if any).

Materials for the feasibility study

The feasibility study used two independent datasets from prior PRISMA-based reviews (Moher et al. 2010), which differ in scope and geographical coverage. Manual classification in these datasets was treated as ground truth.

Key performance indicators for the method

Performance was evaluated as accuracy of the automated extraction method compared to manual classification (considered here ground truth). When there was a perfect match, a value of 1 (100%) was assigned, and we rewarded partial matches (e.g., if one out of two countries were identified that would be a score of 0.5 (50%), same as if there is only one country but the method identified two, one correct and one incorrect).

3 RESULTS

Results of the performance of location-extraction methods is introduced first with respect to testing datasets, followed by the global mapping of residential cooling/overheating research.

Feasibility study

Figure 1 reports the results of the feasibility study used to evaluate the performance of automated-location extraction approaches against manual classification, treated as ground truth. Across both validation datasets (Huebner and Conzatti 2025; Fosas et al. 2025), the LLM-based approaches outperformed the rule-based NLP approach. For the Huebner & Conzatti (2025). dataset (Figure 1a), both LLM variants achieve exact-match accuracy (100.0%), whereas the NLP approach reaches 55.7%. For the Fosas et al. (2025) dataset (Figure 1b), LLM performance remains consistently high (98.4%), while NLP accuracy drops further to 39.8%. Reasons for the mismatch in LLM-based ones for Fosas et al. were due to:

1. "Shamsi et al. (2021)", accuracy of 0.5: LLMs extracted 'USA' and 'Ireland'. The paper deals with the USA alone through the US DOE reference building stock models, whereas Ireland corresponds to the affiliation of authors and funder of the research.
2. "Bragadin et al. (2023)", accuracy of 0.5: LLMs extracted 'Greece' and 'Italy', whereas Fosas et al. (2025) reported only 'Italy'. Revaluating the study, there is a human error and the LLMs' response is in fact accurate. Thus, Figure 1 should really read 99.2%, not 98.4%.

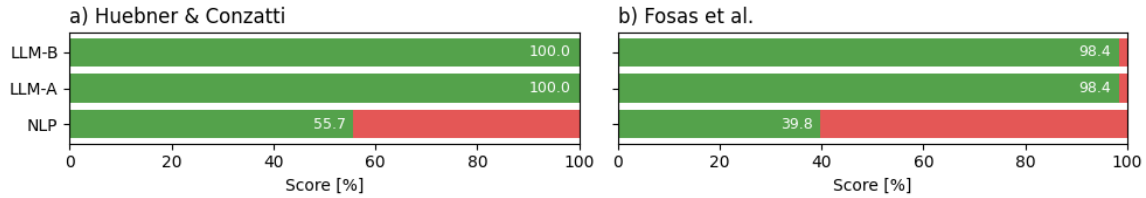


Figure 1 Performance of automated location-extraction approaches compared with manual classification across two validation datasets.

Geographical distribution of cooling research

The global distribution of residential cooling research is highly uneven (Figure 2). The top 10 countries account for approximately 55–60% of all identified study locations, while the top 20 countries contribute around 75–80%. Fewer than 40 countries—less than a quarter of all countries globally—account for over 80% of the total evidence.

Publications are heavily clustered in North America, Europe, East Asia, and Australia, which represent approximately 60–65% of all extracted study locations. A small group of countries dominates the literature: China, the United States, and the UK alone account for roughly 25–30% of all study locations, followed by Spain, Australia, Germany, and Italy. In contrast, large parts of the Global South—particularly Sub-Saharan Africa and South Asia—remain sparsely represented, together contributing less than 10% of all study locations. To further contextualise these spatial patterns, we examined author affiliation for Global South studies (Table 1).

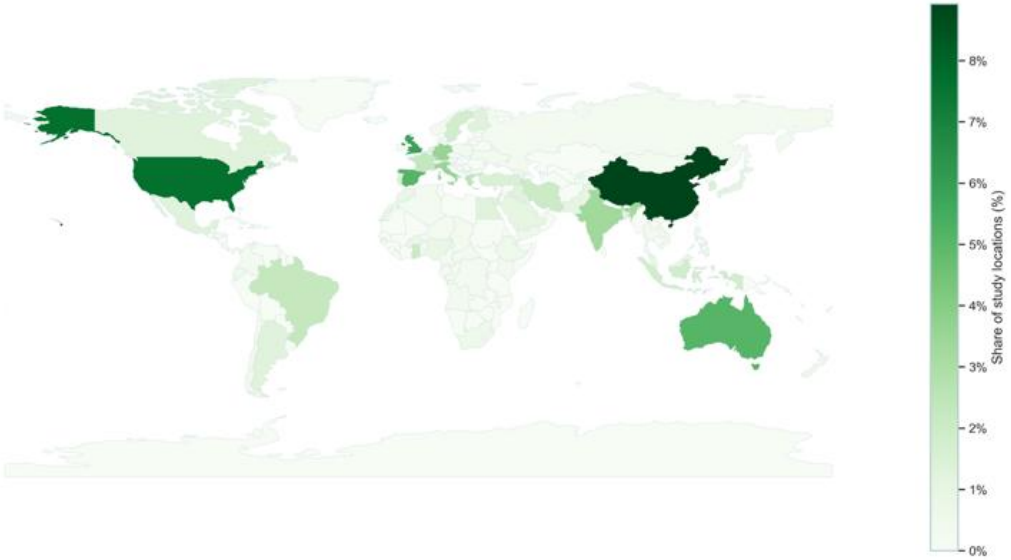


Figure 2 Global distribution of residential cooling research (n=326)

Table 1 Author affiliation regions for cooling studies conducted in the Global South

Study category: Studies with…	Number of studies	Share (%)
…1 or more Global North affiliation	56	58
…1 or more Global South affiliation	96	100
…both Global North and South affiliations	38	40
…only Global North affiliations	18	19
...only Global South affiliations	58	60

4 DISCUSSION

This study provides a global assessment of where empirical evidence on residential cooling is produced and highlights a mismatch between research activity and regions facing the greatest present and future heat exposure. Using a scalable AI-assisted extraction workflow, the results show that the global evidence base on residential cooling/overheating is dominated by a limited set of national context, while many countries undergoing rapid urbanisation and experiencing high projected heat stress remain weakly represented or entirely excluded.

Methodological contribution

A central contribution of this work lies in showcasing how local offline and free models enable rapid and scalable identification of evidence gaps in the literature. Manual classification may remain a desirable reference standard, but it does not scale to large corpora or frequent updates—and scientific outputs continue to grow exponentially and beyond what a single or modest research team is able to process timely. The limited mismatches observed indicate that automated approaches already match manual classification accuracy while being orders of magnitude faster and at a negligible cost, thus enabling researchers to appraise global research distributions. These findings align with broader evidence from the literature on automation and LLM-supported evidence synthesis, which highlights the potential of machine learning tools to accelerate mapping studies and support timely decision-making, provided that human oversight is retained (Ghali et al. 2025; Chen et al. 2025). Recent work on interactive LLM-based mapping studies further emphasises the value of human-in-the-loop approaches that combine automation with researcher control, transparency, and traceability (Petersen and Gerken 2025).

Implication of the Global North and Global South evidence gap

The observed Global North-Global South gap in residential cooling research has direct implications for wellbeing. Cooling strategies and building design guidance are highly context specific. When empirical evidence is disproportionately derived from a small number of countries, design and adaptation strategies risk being tailored for contexts that do not reflect conditions in regions where heat exposure and vulnerability are growing most rapidly.

This gap is particularly concerning considering recent projection showing that future cooling inequalities will be most severe in regions like Sub-Saharan African countries. Falchetta et al. (2024) demonstrated that up to 4 billion people may remain without access to air conditioning by 2050, with deprivation increasingly concentrated in Sub-Saharan Africa and South Asia despite rising heat exposure. Our findings suggest that these same regions are also underrepresented in the literature on residential cooling/overheating, limiting the availability of locally grounded evidence to inform effective and equitable overheating adaptation. If unaddressed, this risks to reinforce a cycle where regions facing the highest health burdens from heat also lack the evidence base needed to guide policy, building practice, and public health interventions (Mazzone et al. 2023; De Cian et al. 2025).

Knowledge production and research authorship

The analysis of author affiliations further indicates that, even when studies are conducted in the Global South contexts, a great share of research is led or co led by institutions from the Global North. While international collaboration plays an important role in advancing technical capacity and data availability, this pattern raises questions about agenda-settings, and the extent to which

local knowledge informs study design and interpretation. Recent work has shown that Global South researchers often act as intermediaries between local contexts and Global North academic systems, while contributing to essential contextual expertise (Djaafara et al. 2025). In the context of residential cooling/overheating, this suggests that addressing global cooling inequalities requires not only expanding the geographical coverage of empirical studies, but also strengthening locally led research and sustained South–North and South–South collaborations that support context-sensitive and policy-relevant evidence generation.

Future developments

If the current imbalance in research production persists, overheating adaptation strategies risk being guided by evidence that is poorly aligned with future climate risk and population growth. Bridging this gap requires targeted investment in empirical research in under-represented regions, support for local research capacity, and systematic mechanisms to monitor where evidence is and is not being generated. Contemporary feature extraction techniques like the ones discussed appear as a practical, affordable, and scalable solution. When it comes to trust and validation, implementation within so called "Retrieval-Augmented-Generation" workflows already enable local, offline queries of large corpora through LLMs while providing citations to specific locations within the retrieved documents, providing means for rapid validation.

5 CONCLUSIONS

This study shows that empirical research on residential cooling and overheating is unevenly distributed, with a strong concentration in a limited number of countries and substantial underrepresentation of regions facing the greatest current and future heat exposure. Therefore, this misalignment highlights a critical gap. The study also demonstrates that AI-assisted location extraction provides a reliable and scalable method for mapping geographical evidence gaps in the literature, matching manual classification accuracy while enabling rapid and reproducible global analysis.

Together, these findings suggest global cooling/overheating research could benefit from improved alignment between research activity and heat risk through scalable synthesis tools that help target efforts to guide equitable, context-sensitive strategies in a warming world.

ACKNOWLEDGEMENT

The authors would like to thank Louise Lebeault as her thesis project on climate change impacts on residential buildings sparked conversations that ended up motivating this work.

6 REFERENCES

Alrasheed, Mousa, and Monjur Mourshed. 2023. 'Domestic Overheating Risks and Mitigation Strategies: The State-of-the-Art and Directions for Future Research'. *Indoor and Built Environment* 32 (6): 1057–77. https://doi.org/10.1177/1420326X231153856.

Berrang-Ford, Lea et al. 2021. 'Systematic Mapping of Global Research on Climate and Health: A Machine Learning Review'. *The Lancet Planetary Health* 5 (8): e514–25. https://doi.org/10.1016/S2542-5196(21)00179-0.

Chen, Haichao et al. 2025. 'Can Large Language Models Fully Automate or Partially Assist Paper Selection in Systematic Reviews?' *British Journal of Ophthalmology* 109 (8): 962–66. https://doi.org/10.1136/bjo-2024-326254.

De Cian, Enrica et al. 2025. ‘The Impact of Air Conditioning on Residential Electricity Consumption across World Countries’. *Journal of Environmental Economics and Management* 131 (May): 103–22. https://doi.org/10.1016/j.jeem.2025.103122.

Djaafara, Bimandra A. et al. 2025. ‘The Bridge between Two Worlds: Global South Researchers’ Journeys through Global North Academic Training and Beyond’. *Epidemics* 53 (December): 100864. https://doi.org/10.1016/j.epidem.2025.100864.

Falchetta, Giacomo et al. 2024. ‘Inequalities in Global Residential Cooling Energy Use to 2050’. *Nature Communications* 15 (1): 7874. https://doi.org/10.1038/s41467-024-52028-8.

Fosas, Daniel et al. 2025. ‘Decarbonising Non-Domestic Buildings at Scale − A Systematic Review’. *Energy and Buildings* 336 (June): 115587. https://doi.org/10.1016/j.enbuild.2025.115587.

Ghali, Mohammed-Khalil et al. 2025. ‘Enhancing Knowledge Retrieval with In-Context Learning and Semantic Search through Generative AI’. *Knowledge-Based Systems* 311 (February): 113047. https://doi.org/10.1016/j.knosys.2025.113047.

Honnibal, Matthew et al. 2020. *spaCy: Industrial-Strength Natural Language Processing in Python*. https://doi.org/10.5281/zenodo.121230.

Huebner, Gesche, and Anna Conzatti. 2025. *D4.3 Wider Housing Inequalities Outcomes*. IEECP/HouseInc Project Deliverable. IEECP. https://ieecp.org/wp-content/uploads/2025/11/D4.3_Wider-housing-inequalities-outcomes_fv_updated.pdf.

International Energy Agency. 2024. *World Energy Outlook 2024*. International Energy Agency. https://iea.blob.core.windows.net/assets/140a0470-5b90-4922-a0e9-838b3ac6918c/WorldEnergyOutlook2024.pdf.

Kenny, Glen P. et al. 2024. ‘Indoor Overheating: A Review of Vulnerabilities, Causes, and Strategies to Prevent Adverse Human Health Outcomes during Extreme Heat Events’. *Temperature* 11 (3): 203–46. https://doi.org/10.1080/23328940.2024.2361223.

Mazzone, Antonella et al. 2023. ‘Understanding Systemic Cooling Poverty’. *Nature Sustainability* 6 (12): 1533–41. https://doi.org/10.1038/s41893-023-01221-6.

Moher, David et al. 2010. ‘Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement’. *International Journal of Surgery* 8 (5): 336–41. https://doi.org/10.1016/j.ijsu.2010.02.007.

Petersen, Kai, and Jan M. Gerken. 2025. ‘On the Road to Interactive LLM-Based Systematic Mapping Studies’. *Information and Software Technology* 178 (February): 107611. https://doi.org/10.1016/j.infsof.2024.107611.

Tian, Peng et al. 2025. ‘Spatial Inequalities in Global Population Exposure to Extreme Heats and Heatwaves’. *Applied Geography* 174 (January): 103474. https://doi.org/10.1016/j.apgeog.2024.103474.

UNDESA. 2025. *World Urbanization Prospects 2025: Summary of Results.* No. 12. UN DESA/POP/2025/TR. UNDESA.

UNEP. 2023. *Keeping Cool in an Increasingly Hotter World: UNEP Spotlight*. United Nations Environment Programme. https://doi.org/10.59117/20.500.11822/44243.

Vicedo-Cabrera, A. M. et al. 2021. ‘The Burden of Heat-Related Mortality Attributable to Recent Human-Induced Climate Change’. *Nature Climate Change* 11 (6): 492–500. https://doi.org/10.1038/s41558-021-01058-x.