Preprint Cambridge Centre for Computational Chemical Engineering ISSN 1473 – 4273

Semantic agent-workflow for cross-city heat planning and sector coupling

Jingfeng Zhou[1,2], Yi-Kai Tsai[2], Ralf Müller[3], Kushagar Rustagi[3], Jiying Chen[1], Xiaochi Zhou[1], Feroz Farazi[1,2], Sebastian Mosbach[1,2], Jethro Akroyd[1,2], Markus Kraft[1,2,3]

released: April 2, 2026

[1] Department of Chemical Engineering and Biotechnology
University of Cambridge
Philippa Fawcett Drive
Cambridge, CB3 0AS
United Kingdom

[2] CARES
Cambridge Centre for Advanced Research and Education in Singapore
1 Create Way
CREATE Tower, #05-05
Singapore, 138602

[3] CMPG
GRIPS – Gründerinnenzentrum Pirmasens
Delaware Avenue 1–3
66953 Pirmasens
Germany

Preprint No. 344

UNIVERSITY OF CAMBRIDGE

Keywords: Knowledge graph, Urban building energy modelling, City-scale energy demands, Regional planning, Solar potential

Edited by
Computational Modelling Group
Department of Chemical Engineering and Biotechnology
University of Cambridge
Philippa Fawcett Drive
Cambridge, CB3 0AS
United Kingdom

E-Mail: mk306@cam.ac.uk
World Wide Web: https://como.ceb.cam.ac.uk/

Abstract

Heating is a major component of urban energy use. The German Federal Heat Planning Act requires municipalities to prepare transparent, spatially consistent heat plans. Current practices rely on disparate datasets and separate models, making it difficult to compare cities or coordinate action beyond municipal boundaries. To address this gap, this study demonstrates a unified, semantic and agent-based data architecture based on dynamic knowledge graphs that automates data retrieval, harmonises building information and supports reproducible simulations. The framework is applied to two neighbouring German cities to quantify heating demand, rooftop solar-heating potential and the techno-economic impact of adding thermal storage. Results show that 69% of buildings in Pirmasens and 56% in Kaiserslautern achieve a positive net present value for solar heating without storage, while coordinated operation increases usable surplus power and can supply up to 93 tonnes of green hydrogen per year, avoiding nearly 1 kilotonne of CO_2. By linking municipal knowledge graphs through a shared semantic layer, the framework enables interoperable and consistent cross-city assessment of heating demand, solar potential, and surplus-electricity coordination. The work demonstrates how semantic infrastructure can operationalise national heat-planning mandates and provide a transferable digital backbone for integrated low-carbon urban and industrial development through sector coupling.

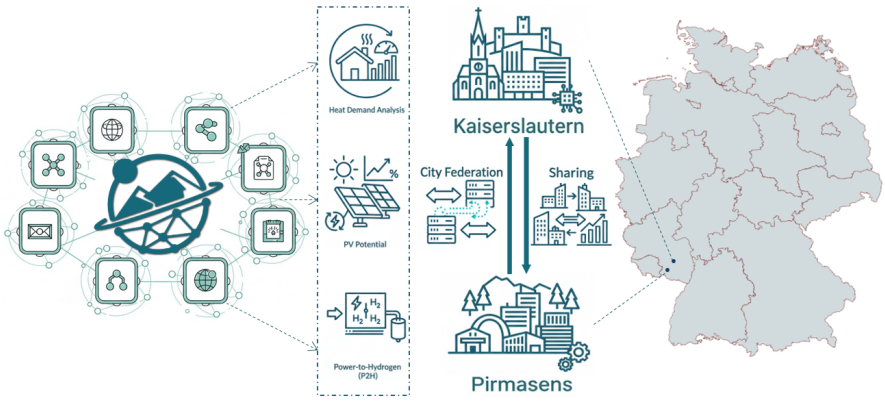

Highlights

- Unified workflow for building heat-demand simulation is demonstrated in two cities.
- Quantify rooftop solar heating and evaluate storage effects on project economics.
- Semantic data architecture supports consistent cross-city analysis.
- Conduct regional coordination to share surplus power and increase hydrogen output.
- Support municipal planning in line with the German Federal Heat Planning Act.

Contents

1 Introduction

In recent years, the decarbonisation of urban heat systems has emerged as a critical component of national and municipal climate strategies. This is largely due to the scale of energy used for heating, which continues to dominate end-use consumption globally. Based on statistical estimates, heating remains the world's largest single end-use of energy, with space and water heating in buildings accounting for approximately 46% of global energy consumption [33]. As local authorities increasingly take responsibility for meeting climate targets, cities have become key actors in designing and implementing low-carbon heat transition pathways.

To support this shift, several countries have introduced policy frameworks to support locally tailored approaches to urban heat decarbonisation. For instance, Denmark has long enabled municipalities to coordinate district heating expansion using renewable and waste heat [39]. In the United Kingdom, local governments are now expected to prepare spatially explicit Local Area Energy Plans that integrate emissions targets, building stock data and infrastructure constraints [70]. The Netherlands requires neighbourhood-level "heat transition visions" to guide the phasing out of natural gas [63]. China has promoted city-level action through the Clean Winter Heating Plan for Northern China, combining targets for "clean heating" uptake with measures spanning fuel switching, district heating upgrades and building efficiency improvements [34]. In Latin America, Chile has linked residential heating interventions to local air-quality governance, using Atmospheric Decontamination Plans and associated policies to reduce firewood-based space heating emissions via technology replacement and fuel quality improvements [32]. Germany has recently adopted a similar approach through the Federal Heat Planning Act, which mandates the development of municipal heat strategies based on geospatial analysis and renewable integration [6]. While these national approaches differ in scope and governance, Germany's legislation is particularly notable for its attempt to standardise municipal heat planning across the country under a unified policy framework. Despite these advances, significant technical challenges remain in urban energy modelling and heat planning. Data fragmentation and the lack of common modelling standards often lead to isolated analyses that are difficult to compare or replicate across cities [47]. In addition, many planning models are embedded within highly localised analytical contexts and are not designed to be adapted to other urban settings, which limits their generalisability and the broader application of insights. A related gap is the lack of modelling frameworks that can support spatial coordination or assess inter-city complementarities and shared infrastructure, despite the increasing recognition of regional heat planning needs [25].

Urban building energy modelling (UBEM) has emerged as a powerful approach for representing heating demand across entire urban areas at high spatial resolution [64], offering detailed insights into spatial energy patterns and providing valuable support for heat planning. Yet, UBEM is often implemented as a stand-alone approach and rarely integrated with data infrastructures that can evolve with changing inputs, support flexible scenario analysis, enable cross-city comparison, or facilitate coordinated planning. This separation limits the ability of municipal planners to produce consistent, transparent, and transferable outputs. These limitations could be addressed by a data framework that integrates

heterogeneous sources while providing semantic consistency, interoperable interfaces, and mechanisms for continuous updates and reconfiguration as datasets, schemas, and planning questions change. Knowledge-based semantic approaches, particularly when combined with semantic technologies, offer opportunities to meet these requirements.*The World Avatar* (TWA) provides a dynamic semantic approach that can harmonise heterogeneous datasets through explicit, machine-interpretable structures and can reorganise information into planning-relevant representations as needed, while supporting automated retrieval and integration [28]. When coupled with UBEM, this approach enables building-by-building energy simulations at the city scale that remain linked to up-to-date information and interoperable across modelling components. Embedding UBEM within such a dynamic semantic infrastructure provides a viable pathway for flexible, transparent, and scalable municipal heat-planning tools across different urban contexts.

The **purpose of this paper** is to investigate options for supporting the low-carbon transition of municipal heating systems and related regional energy uses, and to demonstrate how the knowledge-based approach espoused by *The World Avatar* (TWA) can provide a robust and transferable data backbone for such analyses. Using TWA's semantic data structures, harmonised city-wide datasets can be retrieved and used for transparent, reproducible building-by-building simulations. Building on earlier work for Pirmasens, a city in the German federal state of Rhineland-Palatinate [69], the refined framework is applied to the neighbouring city named Kaiserslautern to evaluate adaptability under different local conditions, and the analysis is extended to beyond the single-city scale to cross-city coordination by linking municipal datasets through a shared query layer to assess surplus-power sharing and power-to-hydrogen potential. Overall, the study indicates that a scalable semantic infrastructure can enable consistent heat-planning assessments across municipalities and unlock additional decarbonisation value through regional coordination.

The structure of this paper is as follows. Section 2 introduces the conceptual and policy background. Section 3 describes the modelling framework and workflows for heating-demand simulation, rooftop solar utilisation, surplus-power coordination, and power-to-hydrogen conversion. Section 4 presents and discusses the simulation validation, city-level heat-planning outcomes, and the effects of regional coupling. Section 5 summarises the key insights and implications for municipal planning and sector-coupled decarbonisation.

2 Background

This section outlines the principles and applications of UBEM, which provides a characterisation of building-level energy demand at city scale and serves as a core analytical basis for low-carbon heating strategies. Subsequently, TWA is introduced as an approach to integrate heterogeneous datasets and promote semantic interoperability across modelling domains. Within TWA, the City Energy Analyst (CEA) agent connects semantic data representations with building- and district-scale energy simulations, enabling automated, reproducible, and scalable workflows, which is also elaborated herein. Finally, approaches to regional energy synergy analysis are considered, extending city-scale capabilities to multi-city contexts to support coordinated infrastructure planning, inter-city

energy exchange, and cross-boundary optimisation. Together, these elements establish the conceptual and technical foundation for the integrated modelling framework employed in this study.

2.1 Urban building energy modelling

UBEM extends the principles of conventional building energy modelling from individual buildings to the city scale, enabling the estimation of energy demand across large urban areas with high spatial granularity [74]. By integrating information on building geometry, typology, construction characteristics, occupancy behaviour, and climatic conditions, UBEM provides a bottom-up representation of energy use across the building stock. This fine-grained characterisation captures spatial and temporal variations in energy consumption and thus forms a critical analytical basis for assessing urban decarbonisation pathways [64].

Data for UBEM are typically sourced from Geographic Information Systems (GIS), cadastral databases, remote sensing, and statistical building archetypes [73, 78], while simulation outputs include spatially resolved heating and cooling loads, end-use energy distributions, and the potential impacts of retrofit or decarbonisation scenarios [10, 77]. As such, UBEM serves as an essential analytical foundation for urban heat planning, allowing local authorities to quantify baseline energy needs, assess intervention pathways, and evaluate alternative supply options at multiple spatial and temporal scales [36]. Hence, UBEM has become an indispensable foundation for municipal and regional heat planning.

Over the past decade, a number of UBEM platforms have been developed, reflecting distinct modelling philosophies, data requirements, and computational strategies. Among these, City Energy Analyst (CEA) [16, 17] is an open-source platform that supports city-scale energy analysis through a set of modules. For example, its demand module computes hourly thermal and electrical loads at the building level, and its resource potential module estimates renewable energy potentials, including photovoltaics, solar thermal collectors, and hybrid photovoltaic-thermal systems. CEA has been used in a range of case studies [23, 24, 48, 52, 53, 67], illustrating its application in academic and planning contexts.

To support use across different data environments, UBEM workflows commonly rely on assumed parameters generated from design guides that allow simulations to be conducted even in the absence of detailed local data [80]. While this approach ensures accessibility and comparability, it also introduces limitations in terms of accuracy and contextual relevance. Model outputs that rely heavily on generic assumptions that do not necessarily reflect local construction standards, occupancy schedules, or energy system configurations, thereby constraining the representativeness and reproducibility of results [82]. Consequently, the capacity of UBEM to support decision-making depends strongly on the quality, consistency, interoperability of underlying data inputs, and knowledge of local building codes.

These challenges could be addressed by a more dynamic and semantically consistent data infrastructure that can harmonise heterogeneous datasets, support automated data exchange, and ensure reproducibility across modelling environments. Integrating UBEM within such a framework would allow for continuously updated, transparent, and transfer-

able energy simulations at city scale. The following section therefore introduces TWA, a dynamic knowledge-based approach designed to provide precisely these capabilities and to enhance the interoperability and scalability of UBEM-based heat planning.

2.2 The World Avatar approach

The World Avatar (TWA) is a community-driven open-source project that aims to explore how explicit knowledge models can be used to achieve a general knowledge representation of the world at large, acted upon by a closed system of semantic computational agents that act autonomously to maintain and enrich the knowledge representation [1, 7, 68]. TWA is based on the idea that data, models, and analytical tools form a dynamic semantic system in which information from heterogeneous sources can be represented in a machine-interpretable and interoperable form [45]. This representation supports consistent querying and reasoning across domains, and it can evolve as new data and modelling requirements arise, which is particularly relevant for urban energy system analyses.

The current implementation of TWA is built on Semantic Web technologies, using the Resource Description Framework (RDF) and domain ontologies to represent statements as subject-predicate-object triples [40, 43], allowing entities and relations to be expressed in a structured form. Ontologies formally define domain concepts and their relationships [20], providing semantic consistency across datasets and tools by enforcing shared meanings and constraints. Together, RDF and ontologies enable heterogeneous information to be integrated, queried (*e.g.*, via SPARQL), and extended without losing semantic coherence.

The "dynamic" aspect of TWA is realised through autonomous software agents that interact with the knowledge graph [83]. Agents can retrieve, validate, transform, and write data, and can orchestrate model execution by linking external sources and simulation outputs back into the same semantic structure. This mechanism confers dynamic behaviour, allowing the knowledge graph to be continuously updated and, when needed, reorganised into analysis-ready representations while preserving semantic consistency through shared ontologies.

TWA can also operate under varying data fidelities. Where high-resolution local inputs are unavailable, agents can populate required fields using proxy datasets, inferred values, or simulation-derived quantities to support continuity of analysis. As more detailed observations or authoritative records become available, they can be incorporated using the same semantic structure, enabling progressive refinement while maintaining traceability of data provenance and assumptions.

By combining semantic technologies, ontology-based consistency, and an agent-enabled update mechanism, TWA provides an interoperable and extensible infrastructure for integrating heterogeneous urban energy data and models. In the context of UBEM, these capabilities directly address data fragmentation and improve reproducibility and transferability of modelling workflows. The next section introduces how CEA is integrated within TWA as an agent to connect semantic data representations with building- and district-scale energy simulations.

2.3 City Energy Analyst agent within The World Avatar

The CEA agent acts as an interface between TWA and the CEA simulation engine [17]. It provides an autonomous component within TWA that enables building- and district-scale energy simulations and supports the automated assessment of energy demand and renewable energy potential, thereby facilitating applications such as municipal and regional heat planning [69]. In this work, the agent is understood not merely as a script or workflow wrapper, but as a modular semantic service that retrieves semantically defined inputs from TWA, executes predefined simulation tasks in CEA, and returns structured outputs to the knowledge graph. By using TWA's dynamic knowledge graph, the agent can automatically retrieve, integrate, and update the input datasets required for simulations, thereby reducing manual configuration and improving model reproducibility.

A central feature of the CEA agent is its ability to replace default assumptions with context-specific, semantically linked data obtained from TWA. Instead of relying on generic archetypes, the agent accesses building-level information, such as geometry, type, and occupancy, as well as site-specific weather and terrain data. These inputs are automatically retrieved and mapped to the formats required by CEA. As demonstrated in previous implementation [69], this capability substantially improves simulation accuracy and local representativeness, a finding corroborated by other studies on data-driven urban energy modelling [15]. By dynamically adapting its inputs to reflect real-world conditions, the agent bridges the gap between model abstraction and empirical data.

Through the semantic integration mechanism of TWA, the CEA agent continuously utilises the most recent and contextually relevant information. When new or updated datasets are introduced, including revised building stock information, updated meteorological records, or infrastructure extensions, the agent can automatically incorporate them into subsequent simulations without requiring reconfiguration. Conversely, the simulation results generated by the agent are written back into the knowledge graph as structured semantic data, ensuring that outputs are interoperable, traceable, and reusable by other agents or analytical workflows. This bidirectional data flow closes the loop between modelling and knowledge management, creating a self-updating ecosystem.

The CEA agent contributes to a highly interoperable and extensible modelling environment. Multiple agents can interact through shared ontologies, exchanging data and coordinating computations to support multi-domain and multi-scale analyses. While the agent provides consistent, comparable, and data-rich outputs at the city scale, the broader question of how regional energy synergies should be conceptualised, assessed, and planned extends beyond any specific tool or platform. The next section therefore presents a general perspective on regional energy synergy analysis, including the clarification of key definitions and typical coordination mechanisms.

2.4 Municipal and regional energy system analysis

Municipal and regional energy system analysis focuses on planning interactions between heating and power systems within and beyond municipal boundaries to improve efficiency, resilience and decarbonisation outcomes. At the municipal scale, a city-scale characteri-

sation of energy demand and renewable generation potential provides the analytical foundation for heat planning [30]. Building-level modelling yields spatially and temporally resolved heating and electricity demand that reflects variations in building type, usage patterns and microclimatic conditions, while parallel assessments of photovoltaic potential estimate the renewable supply that can be integrated within local boundaries [79]. Taken together, these single-city analyses identify local self-sufficiency, imbalances and technology options within municipal borders and support the design of targeted renovation strategies, network expansions and decentralised renewable deployment.

The relevance of such municipal scale assessments has increased following the introduction of Germany's Federal Heat Planning Act (*Bundeswärmeplanungsgesetz*), which came into force in 2024 [6]. The Act establishes a uniform national framework for municipal heat planning by requiring all municipalities above 100,000 inhabitants to complete a full heat plan by 2026, and smaller municipalities by 2028. Each plan must include a geospatially explicit representation of current and future heating demand, the technical potential of renewable and waste heat, the suitability of areas for district heating, and scenarios for achieving climate neutral heat supply by 2045. The legislation further mandates the use of transparent, reproducible and data based analytical methods, with datasets structured so that they can be aggregated consistently at state or regional level. Consistency of input data and modelling assumptions is essential for meaningful comparison [56, 81], because only comparable datasets and aligned assumptions allow differences in results to be attributed to real structural and contextual factors rather than artefacts of mismatched temporal resolution, system boundaries, or data provenance. Approaches that employ federated data retrieval and harmonised semantic structures have therefore become increasingly relevant [4], as they reduce manual data wrangling and version drift, preserve traceable links to heterogeneous sources, and allow models and analyses to be executed and updated consistently across administrative boundaries, thereby supporting reproducible assessments in addition to cross-city comparison. Modelling approaches that derive building-level heating demand and rooftop photovoltaic potential on this basis directly support the legal requirements, as they provide the type of structured and comparable information needed for statutory municipal heat plans.

Building on this basis, regional energy synergy analysis focuses on coordination across multiple cities or urban districts in order to achieve collective efficiency, resilience and decarbonisation goals. While much of the existing research in urban energy modelling has focused on individual cities as self-contained systems, recent studies increasingly highlight the potential of inter-city cooperation to balance renewable generation and demand at a broader spatial scale [25]. This regional perspective aligns with emerging policy frameworks that emphasise integrated planning across administrative boundaries and the joint utilisation of distributed energy resources [5], exemplified by Germany's Federal Heat Planning Act, which standardises municipal heat planning and encourages regional coordination. It recognises that cities are not isolated units but interconnected nodes within wider energy, industrial, and environmental networks.

Most existing work still models cities as isolated entities [26, 37], focusing on internal load-generation balancing without explicitly representing inter-city interconnections or surplus sharing across municipal boundaries. Yet, considering cities as nodes within an interconnected regional network is both useful and justified. Urban areas often display

complementary load and generation patterns [75], primarily due to variations in building stock composition, land-use, industrial activity and renewable resource availability. Facilitating electricity sharing between cities allows excess renewable power generated in one location to offset deficits in another, reducing curtailment and enhancing regional self-sufficiency. Although the practical implementation of inter-city exchanges depends on local grid conditions, this conceptual framework could still illustrate the potential benefits of coordination at the regional level [76], and can support the regional scenario development envisaged in the Federal Heat Planning Act.

When both cities simultaneously produce excess electricity, the remaining surplus can be directed towards flexible industrial applications within the wider region. More broadly, a growing body of research examines how surplus renewable electricity can be utilised through electrochemical conversion pathways, with hydrogen production via electrolysis being a particularly prominent option [54, 55]. This is especially relevant when electrolysers are co-located with chemical or biorefinery facilities, where they can operate as controllable loads that absorb intermittent renewable generation, thereby mitigating grid imbalances while supporting industrial decarbonisation. Hydrogen produced in this way can be used on-site for heat and feedstock substitution, or stored and transported for later use, effectively linking the power and industrial sectors [42]. Beyond hydrogen, other electrochemical or power-to-X pathway, such as the production of synthetic fuels or chemicals, offer similar opportunities for flexible utilisation of surplus energy [9].

This sector-coupled perspective reflects a broader trend in regional energy research that moves beyond single-domain optimisation towards integrated system design. By connecting urban and industrial processes through shared energy and material flows, regions can increase resource efficiency, reduce emissions and improve economic resilience [18]. Sector coupling also enables the creation of local value chains for renewable electricity, strengthening regional economies and facilitating the transition to a circular, low-carbon energy system. Taken together, municipal and regional energy system analysis integrates detailed single city characterisation, spatial coordination between cities and cross sectoral coupling within a consistent analytical framework. By quantifying demand, generation and exchange potentials on comparable terms, it identifies where and when cooperation can deliver mutual benefits and how shared infrastructures may facilitate a more efficient, low-carbon regional transition [21]. The conceptual and empirical insights gained from such analyses provide a foundation for policy-makers and planners seeking integrated multi-city and multi-sector strategies that maximise renewable utilisation and strengthen regional self-sufficiency [57], directly aligning with the data transparency, cross municipal coordination and climate neutrality pathways mandated by the Federal Heat Planning Act.

3 Method

This study investigates options to support the low-carbon transition of municipal heating systems and energy use in two independent cities (kreisfreie Städte) in the German federal state of Rhineland-Palatinate, namely Pirmasens and Kaiserslautern. Both are located in the south-western part of the state at the northern edge of the Palatinate Forest

(Pfälzerwald) low-mountain range and share a temperate oceanic climate.

Pirmasens is a medium-sized city close to the French border with a municipal area of about 61 km^2 and an estimated 39,700 inhabitants in 2023. The settlement structure is characterised by a compact historic core surrounded by post-war residential quarters and small-scale industrial and commercial areas. Kaiserslautern, located north-east of Pirmasens, is the larger regional centre with an area of approximately 140 km^2 and about 100,000 inhabitants. It hosts a university campus, research institutes and major commercial and military facilities, resulting in a more heterogeneous land-use pattern and a broader mix of building typologies. The straight-line distance between the two urban cores is about 30 km, and typical road travel distances are about 35 km. The location and key characteristics of the two cities are shown in Figure 1. An example of a 3D visualisation of Pirmasens and the regional view of these two cities could be found in the Appendix.

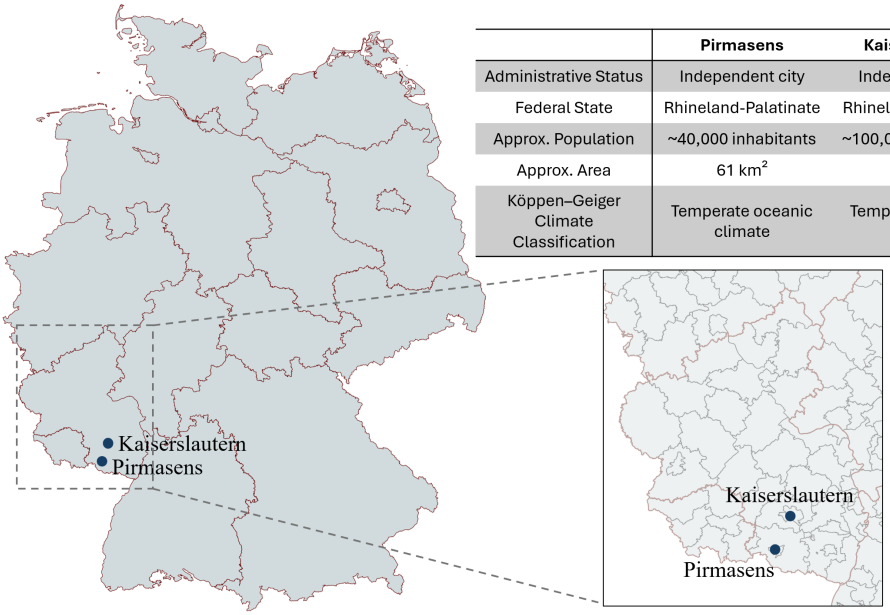

	Pirmasens	Kaiserslautern
Administrative Status	Independent city	Independent city
Federal State	Rhineland-Palatinate	Rhineland-Palatinate
Approx. Population	~40,000 inhabitants	~100,000 inhabitants
Approx. Area	61 km^2	140 km^2
Köppen–Geiger Climate Classification	Temperate oceanic climate	Temperate oceanic climate

Figure 1: *Geographic location and key characteristics of Pirmasens and Kaiserslautern.*

Despite the differences in size and functional role, the two cities share similar climatic conditions and broadly comparable Central European building traditions, while exhibiting distinct urban forms and demand patterns. In particular, the short distance between them implies that coordinated use of electricity and hydrogen infrastructure is, in principle, technically conceivable. This combination makes them a suitable pair of case studies for testing the transferability of a knowledge-graph-based municipal heat planning workflow and for exploring potential regional synergies between neighbouring urban areas within a unified modelling framework.

3.1 Knowledge-based data representation, agent-based simulation, and cross-city retrieval

In this work, all data required for the CEA agent simulations are managed through TWA so that both Pirmasens and Kaiserslautern share a common, machine-readable representation of inputs and outputs. Building data, including geometry, footprints, heights, and ground elevations, are integrated into TWA from municipal 3D city models and cadastral datasets. Building usage and property types are likewise integrated from local land-use and address registries, and terrain elevation is incorporated from a digital elevation model. Hourly weather time series are instantiated by a dedicated weather agent that queries the Open-Meteo API to generate typical meteorological conditions for each location.

These heterogeneous datasets are harmonised using a small set of ontologies as shown in Fig. 2. In this study, all ontological concepts are expressed in typewriter font. Buildings and their geometry are represented using OntoBuiltEnv and GeoSPARQL ontologies [3, 27], with quantitative attributes and their corresponding units represented using the OM units of measure ontology [65]. Weather inputs are represented with OntoEMS ontology and stored as hourly time series using OntoTimeSeries ontology [27, 29]. Energy demand and rooftop PV concepts required by the UBEM workflow are represented using OntoUBEMMP ontology, which links buildings to `EnergyConsumption` and `ElectricitySupply` profiles [29].

On top of this semantic data layer, the CEA agent is deployed as a containerised web service within TWA. A simulation is triggered by an HTTP request to the agent specifying a set of `Building` Internationalised Resource Identifiers (IRIs), e.g. all buildings within the municipal boundary of Pirmasens or Kaiserslautern. The agent then issues SPARQL queries to retrieve target and surrounding building geometries, building usage and archetype information, the relevant weather station and associated weather time series, and local terrain elevation. The retrieved RDF data are translated into CEA input files, including shapefiles for target and context buildings, usage and construction databases, EPW-formatted weather files and raster elevation models, thereby replacing CEA's default assumptions with location-specific information wherever available. If optional inputs, such as detailed usage schedules, cannot be found in TWA, the agent falls back to CEA's internal defaults while still requiring geometry from the knowledge graph.

After preparing the inputs, the agent runs CEA's building energy demand and solar potential modules for a full meteorological year. Upon completion, the update endpoint parses the CEA output and instantiates the results back into TWA using the same ontology patterns. For each simulated building, hourly profiles for heating, cooling, electricity and grid consumption are written back as semantically annotated time series with explicit units. Rooftop PV results are likewise stored as `ElectricitySupply` profiles linked to the corresponding building surfaces. The agent also records metadata such as simulation date, scenario identifier and CEA version, allowing runs for different cities and scenarios to coexist and be queried consistently.

To enable cross-city retrieval while keeping municipal deployments independent, the TWA stacks for Pirmasens and Kaiserslautern are connected through a shared query layer based on RDF4J [11]. Each city exposes its own SPARQL endpoint, and a federation

layer provides a single logical endpoint that dispatches sub-queries to the respective city repositories and merges results. Reusable query templates retrieve building inventories, simulated demand profiles, and PV generation time series across both cities, with OM units ensuring consistent interpretation of numerical values. In this way, TWA provides key inputs to CEA, CEA outputs are written back as queryable time series, and the federated endpoint supports subsequent cross-city analysis and regional synergy assessment. The schematic diagram of the data structure is shown in Figure 2. All referenced namespaces are declared in the Appendix.

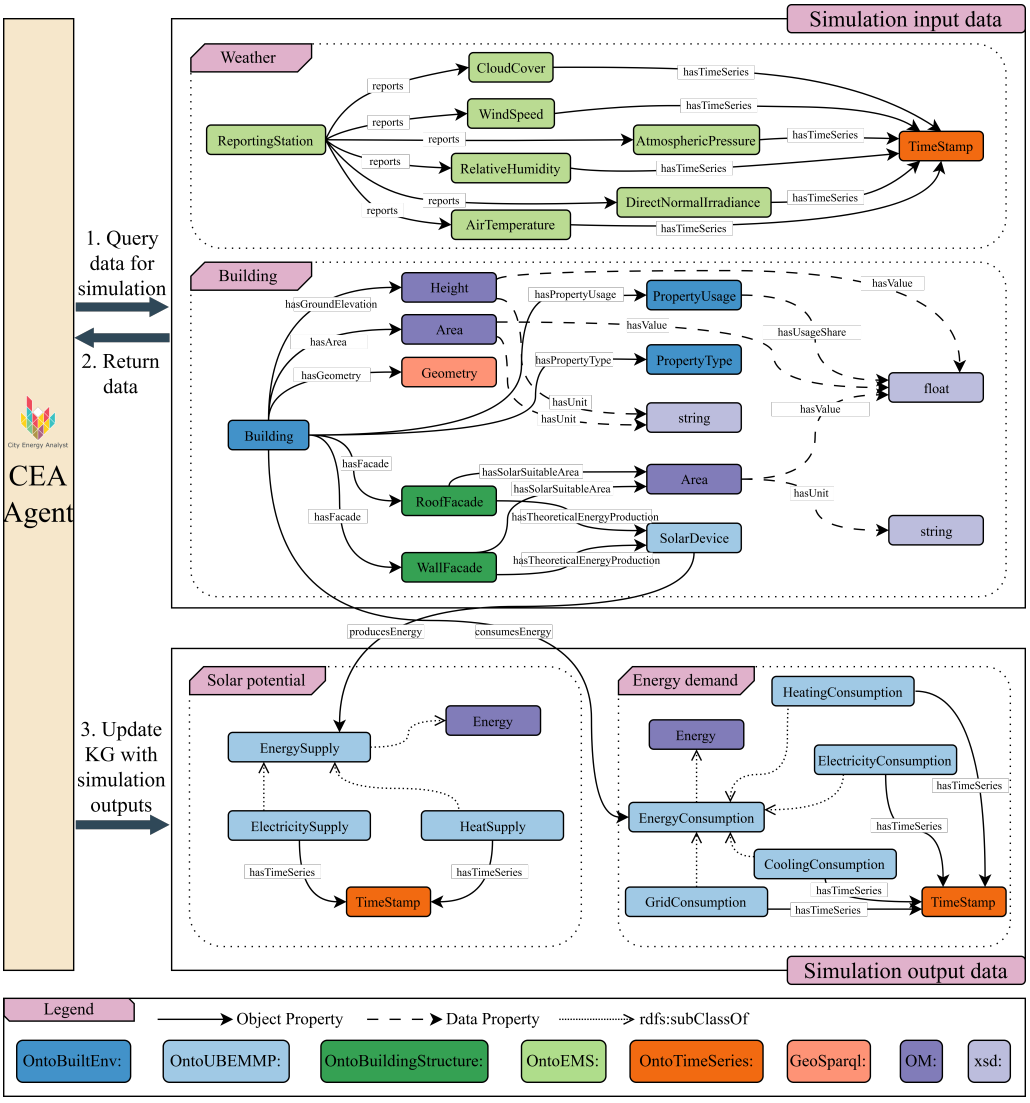

Figure 2: *Ontology-based data structure linking CEA agent simulation inputs and outputs within the knowledge graph.*

Specific indicators are used to characterise heating demand and rooftop PV generation at building and city scales. Let b index buildings in the analysed set $\mathscr{B}$, and t index time steps in $\mathscr{T}$, which span the simulation year. The simulated heating power demand of building b at time step t is $P_{\text{heat}}(b,t)$, and the electrical power generated by rooftop PV on

building b at time step t is $P_{\mathrm{PV}}(b,t)$. Let Δt be the time-step duration (here $\Delta t = 1$ h), and let A_b be the gross floor area of building b.

The total annual energy use (EU) and the energy-use intensity (EUI) for heating and PV are defined as

$$\mathrm{EU}_{\mathrm{city,heat}} = \sum_{b\in\mathscr{B}} \sum_{t\in\mathscr{T}} P_{\mathrm{heat}}(b,t)\,\Delta t, \tag{1}$$

$$\mathrm{EU}_{\mathrm{city,PV}} = \sum_{b\in\mathscr{B}} \sum_{t\in\mathscr{T}} P_{\mathrm{PV}}(b,t)\,\Delta t, \tag{2}$$

and

$$\mathrm{EUI}_{b,\mathrm{heat}} = \frac{1}{A_b} \sum_{t\in\mathscr{T}} P_{\mathrm{heat}}(b,t)\,\Delta t, \tag{3}$$

$$\mathrm{EUI}_{b,\mathrm{PV}} = \frac{1}{A_b} \sum_{t\in\mathscr{T}} P_{\mathrm{PV}}(b,t)\,\Delta t, \tag{4}$$

where $\mathscr{B}$ denotes the set of buildings included in the analysis and $\mathscr{T}$ denotes the set of time steps over the simulation year. The indicators in Eqs. (1)-(4) are used to compare Pirmasens and Kaiserslautern and to support subsequent assessment of rooftop solar utilisation, regional coordination, and power-to-hydrogen potential.

3.2 Heating consumption and rooftop solar thermal potential analysis

The assessment of heating consumption and rooftop photovoltaic potential is conducted at the individual building scale using the hourly simulation outputs produced by the CEA agent. These outputs serve as the foundation for all subsequent analyses and provide the consistent and comparable dataset required for later cross city assessments. For each `Building` instance represented in TWA, the simulation returns hourly values of `HeatingConsumption` and `ElectricityConsumption` together with theoretical solar potential expressed through `SolarDevice` and `ElectricitySupply` instances linked to `ts:TimeSeries` and `ts:TimeStamp`. All numerical values are annotated with OM units to ensure internally consistent processing.

The analysis workflow follows the logic shown in Figure 3. An analysis request specifies a target building and a scenario. The workflow begins by validating the completeness of simulation inputs. Buildings with missing geometry, missing time series or inconsistent units are excluded from further analysis. For each building, the feasible rooftop area is identified using the `RoofFacade` elements stored in TWA together with their `hasSolarSuitableArea` attribute. The hourly `HeatingConsumption` time series and solar potential time series associated with the building are retrieved from TWA and used to conduct a building level balancing analysis. This analysis is carried out under two scenarios. The first scenario represents systems without thermal energy storage (TES). In this case, the hourly solar production is used directly to offset heating demand and any remaining unmet demand is supplied by the baseline gas system. The second scenario represents

systems that include a water based TES vessel. Under this configuration, hours with solar surplus result in storage charging while hours with solar deficit trigger storage discharge, subject to the vessel's capacity and standing loss. Any residual deficit after storage discharge is supplied by gas. For each scenario, the algorithm computes the residual heating demand that must be met by gas and converts this demand into an annual fuel cost using a consumption dependent gas tariff. Annual operational savings are determined by comparing the baseline gas cost to the post installation residual gas cost. Capital and operational expenditures for the solar collector and storage system depend on the installed collector area and storage size. These costs together with the annual savings yield the annual cash flow. The net present value (NPV) over the system lifetime is computed as

$$\mathrm{NPV} = -C_0 + \sum_{y=1}^{Y} \frac{\mathrm{CF}_y}{(1+r)^y}, \tag{5}$$

where C_0 is the initial investment cost, CF_y is the net cash flow in year y, r is the discount rate, and Y is the project lifetime. Multiple collector sizes are evaluated for each building, and the configuration that yields the highest NPV is selected as the optimal design.

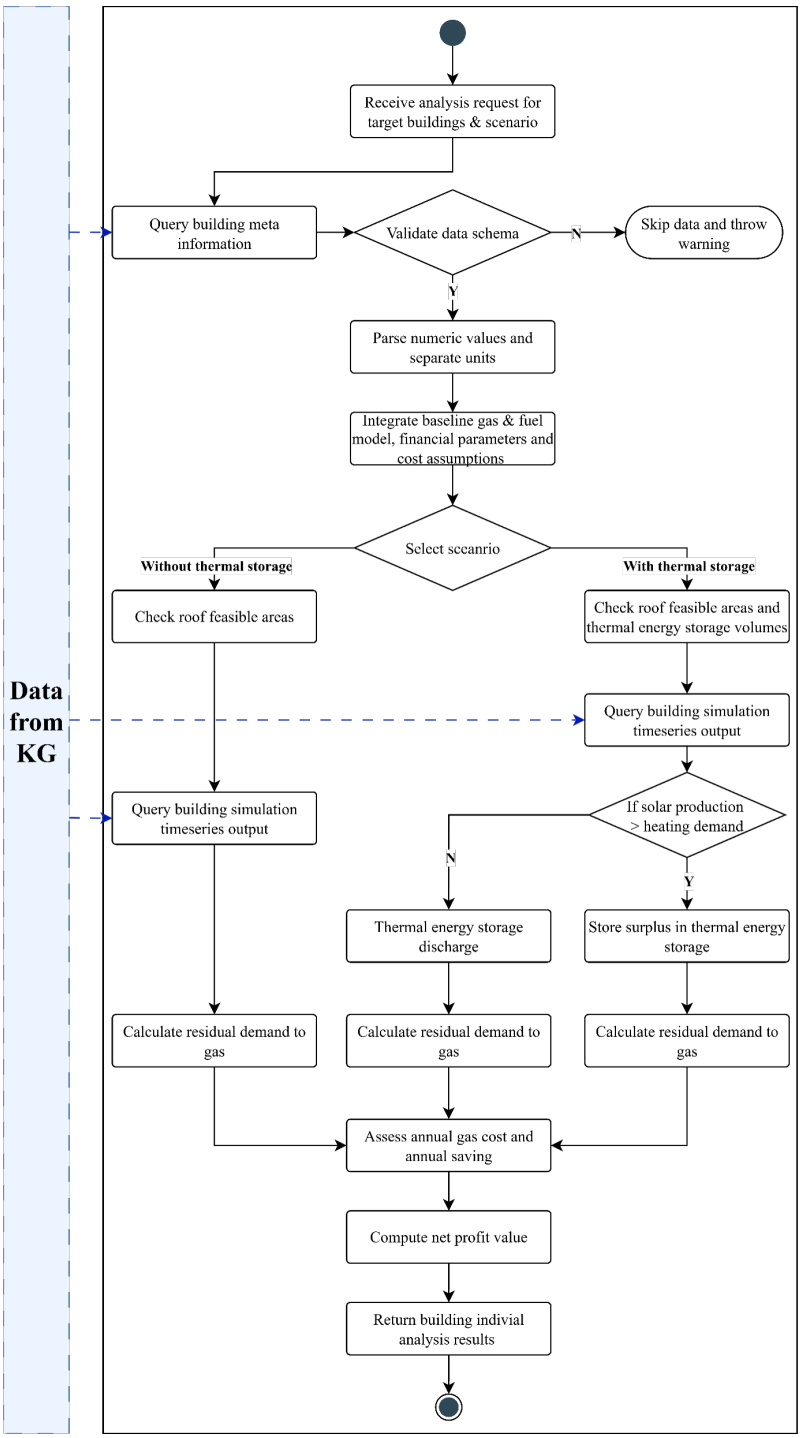

Figure 3: *Workflow of the building-level assessment of heating demand, rooftop solar utilisation and economic performance under scenarios with and without thermal energy storage.*

The parameters required for heating and rooftop solar analysis are applied consistently to all buildings within a city, which ensures comparability and forms a coherent basis for subsequent regional synergy assessment. Tables 1 summarise the technical, financial and

operational parameters adopted in the analysis, sourced from [8, 12–14, 31, 46, 69]. Baseline fuel is assumed as natural gas. The installed collector area A_{inst} determines the usable solar surface and directly influences thermal output, system cost and storage requirements. The discount rate r and system lifetime T define the financial evaluation horizon and the present value of annual cash flows. The gas boiler efficiency governs the conversion of fuel into useful heat and therefore affects both baseline heating cost and residual gas demand after solar integration. In the storage scenario, the TES volume, specific capacity and standing loss determine the amount of solar heat that can be buffered and later utilised. Equipment, installation and operation costs follow area dependent functions that ensure a consistent representation of system expenditures. All values are applied uniformly across the building stock to maintain comparability and support subsequent analyses. Detailed calculations could be found in the supplementary file.

Table 1: *Key technical, financial and operational parameters*

Parameter	**Value**
Gas boiler efficiency	0.80
Annual gas price increase	4.49%
Discount rate r	3.37%
System lifetime T	25 years
Collector module area	2.5 m^2
Minimum collector area	10 m^2
TES volume	$38.9A_{\text{inst}} + 216.7$ L
Minimum TES volume	500 L
Maximum TES volume (domestic)	1500 L
Maximum TES volume (non-domestic)	3000 L
TES specific capacity	70 kWh m^{-3}
TES standing loss	1% per hour
Equipment cost (no TES)	$300A_{\text{inst}} + 800$ EUR
Equipment cost (with TES)	$400A_{\text{inst}} + 2500$ EUR
Installation cost	25% of equipment cost
O&M cost (no TES)	$2.5A_{\text{inst}} + 100$ EUR yr^{-1}
O&M cost (with TES)	$3.5A_{\text{inst}} + 147.5$ EUR yr^{-1}

3.3 Regional synergy analysis

Regional energy synergies are evaluated using a post-processing workflow that operates on the hourly simulation outputs stored in TWA. Building-level data from the different municipalities is accessed through queries that combine the individual TWA instances into a single logical dataset. The workflow is implemented as a dedicated analysis component that receives a request specifying the set of cities and the selected scenario for synergy analysis. The logic is summarised in Figure 4. After receiving a request, the component queries building level meta information and checks the completeness of the data schema. Entries with missing geometry, missing time series or inconsistent OM units are discarded

and recorded as warnings.

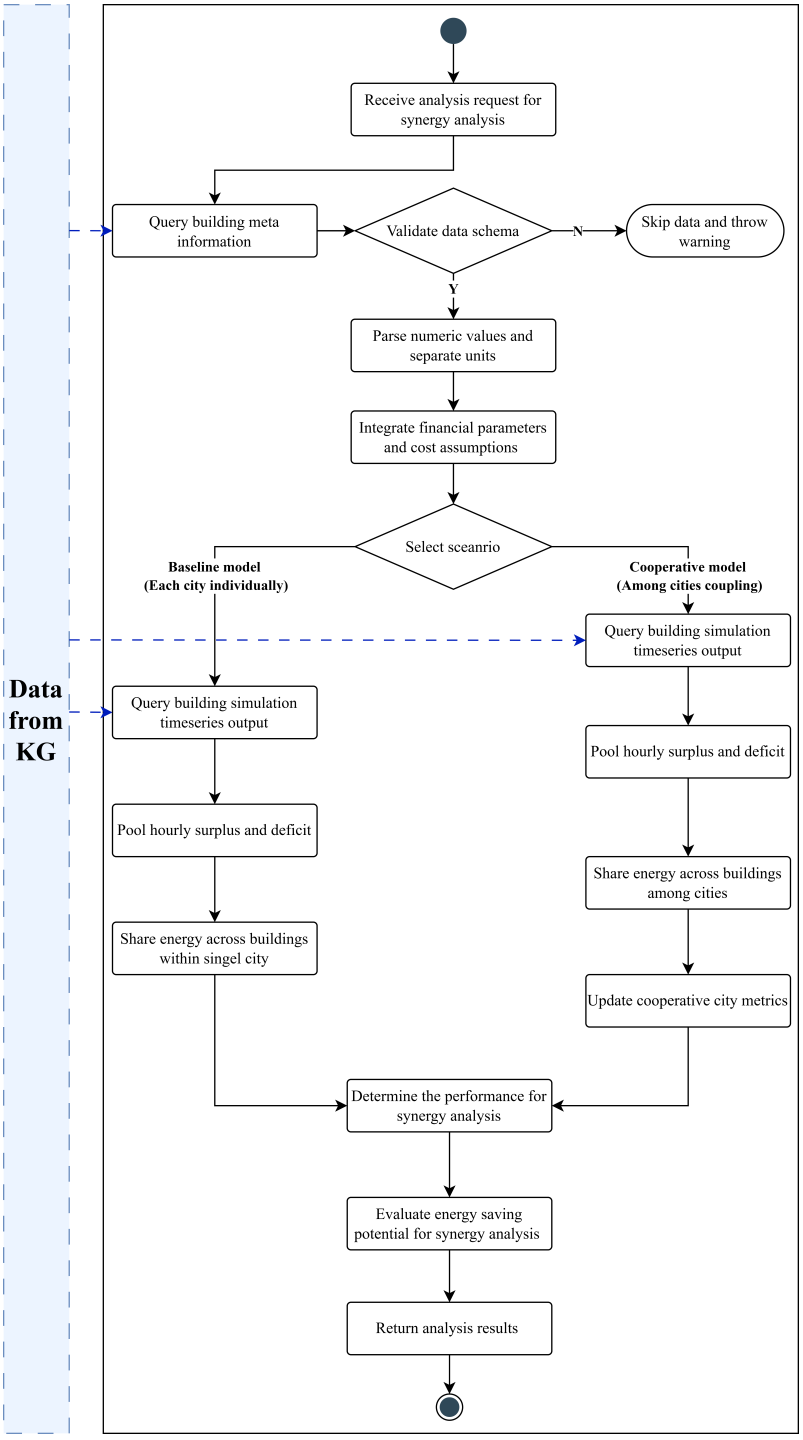

Figure 4: *Workflow of the regional synergy assessment comparing independent city operation with cooperative surplus sharing across Pirmasens and Kaiserslautern.*

For each city, the workflow retrieves hourly electricity demand and rooftop photovoltaic generation for every `Building` instance through the linked `ElectricityConsumption`

and `ElectricitySupply` time series. At each time step t, the net electrical balance of building b is computed as

$$P_{\text{net}}(b,t) = P_{\text{PV}}(b,t) - P_{\text{load}}(b,t), \tag{6}$$

where $P_{\text{PV}}(b,t)$ denotes rooftop photovoltaic output and $P_{\text{load}}(b,t)$ denotes electrical demand. Positive values of P_{net} correspond to surplus electricity that can potentially be shared, and negative values correspond to deficits that require external supply. Building level surpluses and deficits are aggregated to the city scale to obtain hourly net profiles for each municipality.

Two stylised scenarios are considered. In the baseline scenario each city is treated individually. For every time step, building-level surpluses within a city are pooled and used to cover building-level deficits in the same city. Any residual surplus is assumed to be curtailed and any residual deficit is assumed to be supplied by the external grid. This procedure yields city-specific indicators such as annual curtailed energy, annual grid imports and city-level self-sufficiency. In the cooperative scenario, building-level surpluses and deficits are first aggregated to the city level and then combined across cities. Hourly residuals after intracity balancing are exchanged between cities so that surplus in one municipality can reduce the deficit in the other. Remaining surplus and deficit after this intercity balancing correspond to regional curtailment and regional grid import. City-level metrics are updated to reflect the cooperative operation.

In the final step, performance metrics for synergy analysis are derived for both scenarios. For each city and for the combined region, annual energy indicators such as total curtailed photovoltaic generation, total grid imports and the ratio of locally supplied demand are computed. An economic evaluation applies uniform and time-independent electricity prices to all cities in order to estimate annual costs associated with grid imports and the value of locally used photovoltaic electricity. The difference between baseline and cooperative cases represents the potential benefit of regional coordination.

The regional synergy analysis is designed to assess technical feasibility rather than to reproduce a detailed market or network model. Several simplifying assumptions are therefore imposed. Electricity exchange between cities is limited only by the hourly availability of surplus and deficit, while transmission capacity constraints, network losses and reinforcement costs are neglected. A single average electricity price and a single valuation of photovoltaic electricity are applied across the entire region, and regulatory and transaction costs are not represented explicitly. Under these conditions, the cooperative scenario can be interpreted as an upper bound on the energy and cost savings that would be achievable if surplus renewable electricity could be freely shared across the region. The results therefore quantify the technical potential for future regional coordination and provide a basis for subsequent studies that include more detailed market design and grid constraints.

3.4 Power-to-hydrogen and industrial integration analysis

The power-to-hydrogen analysis uses the surplus electricity profiles derived in the previous sections to estimate the corresponding hydrogen output under the assumed electrolyser parameters, providing an indicative upper-bound assessment of the potential contribution to industrial hydrogen supply. The workflow is implemented as a post-processing

step operating on hourly time series stored in TWA, as summarised in Figure 5. An analysis request specifies the cities, the underlying surplus scenario and the use of surplus electricity for hydrogen production. The component first validates the data schema and discards records with missing geometry, incomplete time series or inconsistent OM units. For the remaining records, numerical values are combined with electrolyser parameters aligned with those of the water electrolysis plant planned at the Energiepark in Pirmasens, with a rated electrical input of 1.8 MW_{el} [62]. The electrolyser is planned in connection with the on-site biorefinery, which converts biomass-based feedstocks into fuels and chemical products and can use hydrogen as a flexible industrial energy carrier and process input.

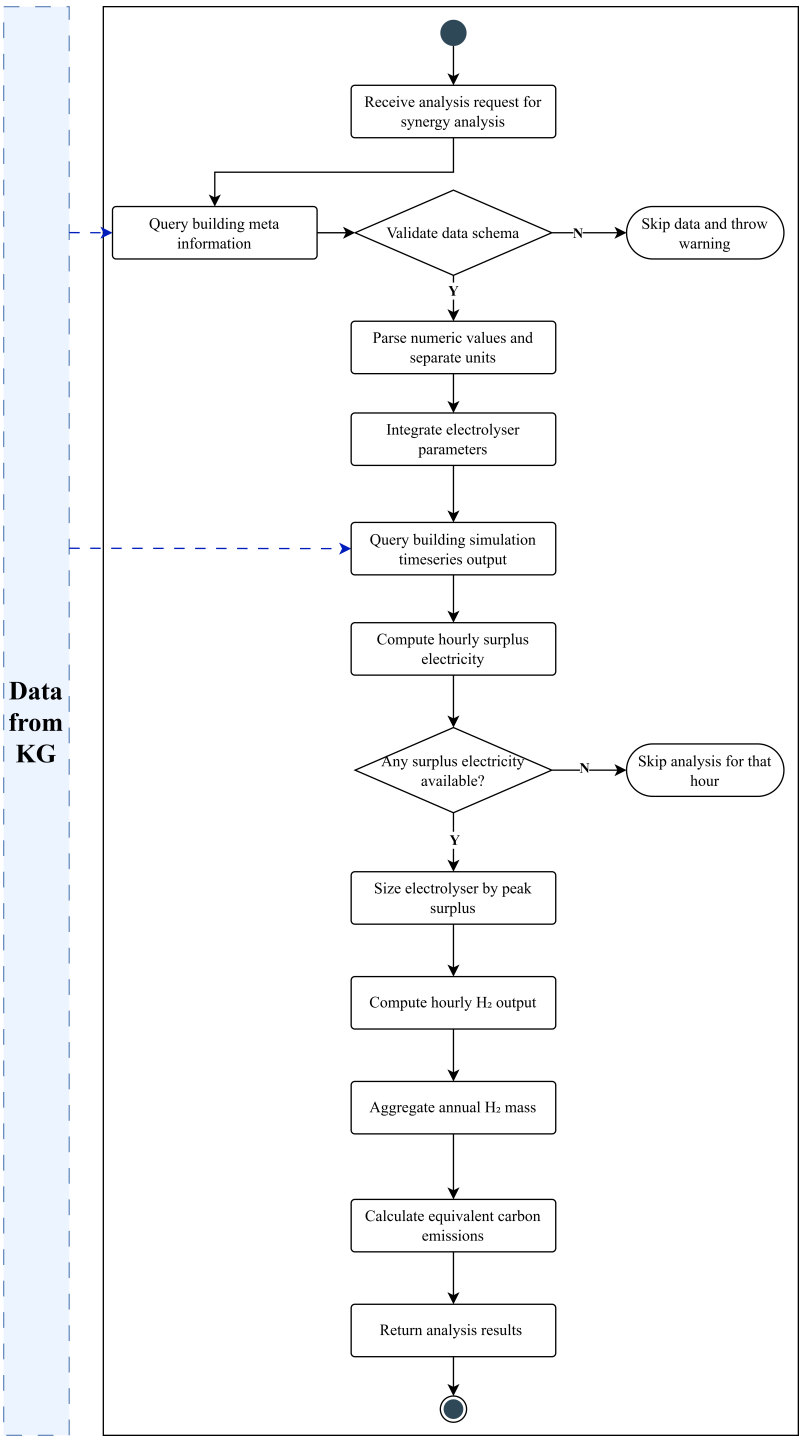

Figure 5: *Workflow of the surplus-based hydrogen production assessment using aggregated hourly electricity surpluses from Pirmasens and Kaiserslautern.*

The analysis is conducted in two steps. In the first step, only the surplus electricity from Pirmasens is used. Hourly surplus power for Pirmasens is obtained from the regional synergy analysis after intra-city balancing of photovoltaic generation and electrical demand. This yields a local surplus time series $P_{\mathrm{sur}}(t)$ that represents the electricity that could be

diverted to hydrogen production without affecting local demand coverage. In the second step, surplus electricity from both Pirmasens and Kaiserslautern is combined. Hourly surplus profiles from the cooperative regional scenario are aggregated to obtain a regional surplus time series $P_{\mathrm{sur}}(t)$ for the combined system. In both steps, the same electrolyser representation and operating rules are applied.

The electrolyser is modelled as a single industrial-scale unit with a fixed rated power, consistent with the scale of the planned biorefinery project at the Energiepark. For a given surplus time series $P_{\mathrm{sur}}(t)$, the dispatched electrical power $P_{\mathrm{el}}(t)$ is defined as

$$P_{\mathrm{el}}(t) = \min\left(P_{\mathrm{sur}}(t), P_{\mathrm{el,max}}\right), \tag{7}$$

where $P_{\mathrm{el,max}}$ is the rated power of the electrolyser. Only surplus electricity is used and the nominal capacity is never exceeded. Hours without surplus correspond to $P_{\mathrm{el}}(t) = 0$. Start-up behaviour, minimum-load constraints and ramping limits are not represented explicitly. Instead, a simplified dispatch assumption is adopted in which surplus power can be converted up to the rated capacity on an hourly basis. This provides an indicative upper-bound estimate of utilisation and hydrogen output at the temporal resolution of the regional synergy analysis, and it avoids introducing additional operational assumptions that are not supported by the available data.

Hydrogen production is calculated using a constant electrolyser efficiency η_{el} and the lower heating value of hydrogen $LHV_{\mathrm{H_2}}$. The lower heating value here is consistent with standard practice for electrolysis-based energy balances and reporting of conversion efficiencies [49], where the recoverable energy at typical end-use conditions is referenced without assuming condensation of product water. In line with reported performance for present-day electrolysers, the efficiency is set such that approximately 65 % of the electrical input can be stored in hydrogen [72]. For each time step t, the hydrogen mass flow $\dot{m}_{\mathrm{H_2}}(t)$ is obtained from

$$\dot{m}_{\mathrm{H_2}}(t) = \frac{\eta_{\mathrm{el}} P_{\mathrm{el}}(t)}{LHV_{\mathrm{H_2}}}. \tag{8}$$

Annual hydrogen production $M_{\mathrm{H_2}}$ is calculated by summing $\dot{m}_{\mathrm{H_2}}(t)$ over all hours of the year. Electrolyser utilisation is quantified through full load hours (FLH), defined as

$$\mathrm{FLH} = \frac{\sum_t P_{\mathrm{el}}(t)\,\Delta t}{P_{\mathrm{el,max}}\,H_{\mathrm{year}}}, \tag{9}$$

where Δt is the time step (one hour) and $H_{\mathrm{year}} = 8760$ is the number of hours in a year. These indicators are computed for the Pirmasens only case and for the combined regional case, which allows the impact of regional coordination on hydrogen production and electrolyser utilisation to be assessed without changing the electrolyser size or efficiency.

Industrial integration is represented at an aggregated level. The electrolyser is interpreted as part of the planned biorefinery at the Energiepark Pirmasens, where hydrogen is used internally for process energy and feedstock substitution. All hydrogen produced in the analysis is therefore assumed to be fully utilised within this industrial facility. Detailed modelling of hydrogen storage, compression and on-site distribution is not included. Within this framing, the results describe the technical potential of using local and regional

surplus electricity to supply hydrogen to an industrial biorefinery of fixed, site-specific scale.

In addition to hydrogen volumes and full load hours, the results are expressed as order-of-magnitude indicators that link the power-to-hydrogen potential to industrial decarbonisation. Annual hydrogen production is converted into an equivalent share of a reference industrial hydrogen demand. Indicative annual CO_2 reductions are estimated by comparing the simulated green hydrogen production with a counterfactual supply of hydrogen from natural-gas-based steam methane reforming (SMR). For this purpose, a representative emission factor for steam methane reforming, denoted as EF_{SMR}, is introduced. This parameter expresses the amount of CO_2 emitted per kilogram of hydrogen produced via conventional SMR, and therefore provides the benchmark against which the avoided emissions from green hydrogen can be assessed. A commonly used value of $EF_{\mathrm{SMR}} = 10.5\ \mathrm{kgCO_2/kgH_2}$ is applied, consistent with life-cycle assessments of natural-gas-based hydrogen production [35]. The resulting CO_2 savings are interpreted as approximate ranges. In this way, the power-to-hydrogen analysis not only quantifies the technical potential to convert surplus electricity into hydrogen but also indicates the order of magnitude of industrial decarbonisation that such a conversion could support.

4 Results and discussion

This section presents the results of the analyses. First, the reliability of the building energy simulations is assessed by comparing energy-use intensities across building types. Second, municipal heat planning outcomes are analysed for each individual city, focusing on heating demand, rooftop solar-thermal utilisation and building-level economic performance. Third, the feasibility of regional coordination is examined by comparing the isolated and cooperative operation of the two cities with respect to surplus sharing. Finally, the potential for power-to-hydrogen conversion of surplus electricity and its industrial relevance is evaluated.

4.1 Building energy simulation results and reliability assessment

Figure 6 compares the simulated annual EUI across building-use categories for the cities of Pirmasens and Kaiserslautern.

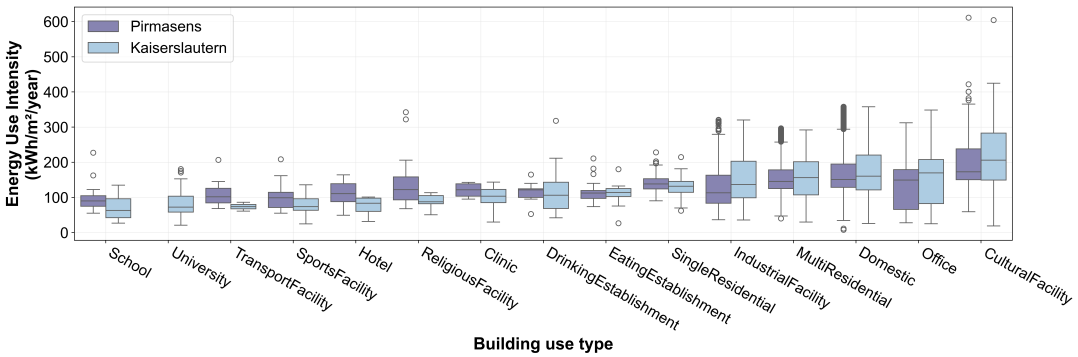

Figure 6: *Simulated annual EUI distributions across building-use categories for Pirmasens and Kaiserslautern. Boxplots summarise the median and interquartile range, and colours distinguish the two cities.*

The EUIs calculated by the CEA agent mostly fall within the range commonly reported for European building stocks, including offices, hospitals, industrial and commercial buildings, domestic buildings, and educational facilities, as shown in previous empirical studies [2, 22, 38, 50, 51, 66, 71]. This overall alignment indicates that the underlying thermal archetypes, climatic boundary conditions and simulation setup lead to demand estimates that are consistent with external benchmarks. For most building-use types such as schools, sports facilities, offices and domestic services, the distributions for Pirmasens and Kaiserslautern show substantial overlap. Median values differ only moderately and the interquartile ranges display similar spread, suggesting that the modelling approach produces comparable outcomes across the two independent TWA stacks. High-intensity categories such as industrial facilities and domestic-service buildings show wide variability in both cities. This is expected for these heterogeneous categories and agrees with the wide ranges reported in benchmark studies. A small number of outliers above 300-600 kWh m^{-2} a^{-1} appear in both cities, but they correspond to specialised or high-load buildings and remain within plausible empirical bounds. A notable distinction is observed in the university category. Kaiserslautern hosts a major university campus whereas Pirmasens does not have an equivalent facility. This structural difference naturally leads to a university EUI distribution that appears only for Kaiserslautern and is absent for Pirmasens.

Taken together, the comparison shows that the simulated EUIs are broadly consistent with published empirical ranges, that the observed differences between the two cities are explainable by genuine structural characteristics. Since the integration only combines query results from the underlying repositories without altering data values, it does not add an additional source of data error. These findings provide confidence that the building-level simulations form a reliable basis for the subsequent analyses of rooftop solar utilisation, regional coordination and hydrogen-production potential.

4.2 Municipal heat planning at individual city scale

The municipal heat planning analysis at individual city scale focuses on how heating demand and rooftop solar-thermal potential translate into feasible decarbonisation options in

each municipality. The comparison in Figure 7 shows that Kaiserslautern exhibits higher total heating demand and a larger rooftop solar potential than Pirmasens. This difference is primarily driven by the larger building stock and the presence of industrial and university buildings in Kaiserslautern, whereas Pirmasens is characterised by a higher share of small domestic buildings. Despite these structural differences, both cities display a broadly favourable ratio of rooftop potential to heating demand. This pattern is consistent with previous findings that overall solar resource or aggregate roof area is often not the binding constraint. Instead, realisable contributions are typically limited by the building-level distribution of suitable roof surfaces and by the temporal mismatch between solar availability and heating load [19].

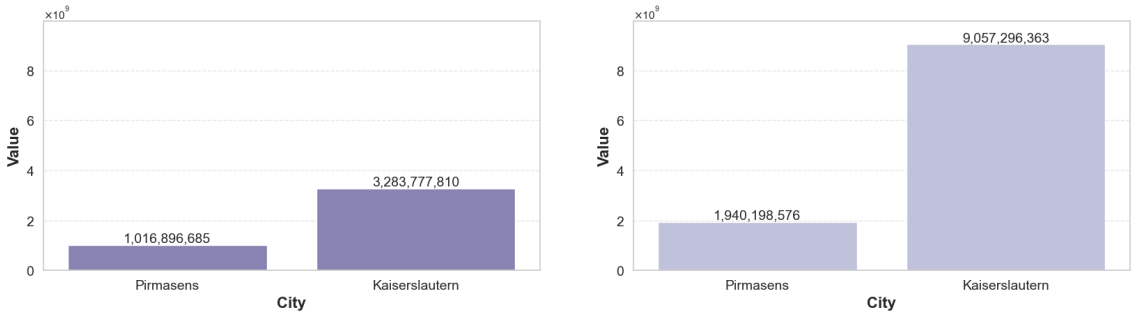

Figure 7: *Comparison of annual building heating demand and rooftop solar-thermal potential for Pirmasens and Kaiserslautern.*

The net present value analysis, which is shown in Table 2, indicates that the economic viability of rooftop solar thermal deployment is strongly conditioned by these structural factors. Pirmasens exhibits a larger share of buildings with positive NPV than Kaiserslautern, which is consistent with its higher heating intensities in domestic buildings and the comparatively favourable roof geometries of low-rise stock. Kaiserslautern, while benefiting from a higher absolute number of profitable projects because of its size, shows a lower proportion of economically attractive systems within the stock. Introducing thermal energy storage increases capital expenditure and therefore slightly reduces the number of buildings that remain profitable. However, for those buildings where the investment remains attractive, storage substantially raises the net present value, indicating that improved temporal matching between solar gains and heat demand can more than compensate for the additional cost. The city-wide increase in total NPV in both municipalities demonstrates that storage acts as a targeted enhancement that benefits the best-suited buildings rather than a uniform solution across the entire stock.

For Pirmasens the overall pattern closely resembles the outcomes reported in an earlier single-city assessment based on the same CEA agent, both in terms of the prevalence of profitable systems and the relative impact of storage [69]. The fact that the extended analysis reproduces these characteristics while also capturing the distinct profile of Kaiserslautern supports the credibility of the modelling workflow and its transferability between municipal contexts. The use of a shared semantic TWA set-up is central in this respect, since data for both cities are accessed through the same semantic structure and query logic rather than through separately maintained model implementations. This reduces the risk that differences arise from inconsistent data handling and ensures that variations observed

Table 2: *Summary of building-level NPV results for rooftop solar-thermal systems with and without TES.*

Metric	Pirmasens		Kaiserslautern	
	Without TES	With TES	Without TES	With TES
Total number of buildings with positive NPV after 25 years	8 371	8 141	17 042	16 799
Percentage of buildings with positive NPV after 25 years	69.10%	67.20%	56.05%	55.25%
Mean NPV of buildings with positive NPV after 25 years (EUR)	28 582.21	36 543.57	28 761.87	36 517.57
Total NPV of buildings with positive NPV after 25 years (million EUR)	239.26	297.50	500.51	613.46
Mean capital expense of installed solar systems per building (EUR)	20 402.79	28 181.31	20 516.97	28 368.24

between the two cities emerge from genuine differences in building-stock composition and use rather than from methodological artefacts.

From the perspective of the German Heat Planning Act, these results illustrate how a knowledge-based urban energy model can support the inventory and potential phases of municipal heat planning. The combination of aggregated indicators and building-level NPV results enables municipalities to identify which parts of the stock are most suitable for rooftop solar thermal deployment, where storage can deliver additional value and where alternative measures are likely to be required. Federated access to the underlying data further allows these analyses to be extended to additional municipalities without redesigning the data interfaces, thereby supporting a more standardised implementation of heat planning across administrative boundaries. In this way, the individual city analyses for Pirmasens and Kaiserslautern provide not only a quantitative estimate of local decarbonisation potential but also a set of spatially explicit priorities that can be directly fed into municipal heat plans prepared under the new shared knowledge-based framework.

4.3 Regional coordination and cross-city surplus sharing

The regional synergy assessment examines how surplus electricity and heating deficits evolve when Pirmasens and Kaiserslautern are treated as a coupled system rather than as independent municipal units. Using the shared knowledge-based structure, hourly time series for both cities are retrieved in a harmonised form, so that all indicators are constructed from consistent simulation outputs and isolated and cooperative configurations can be compared on a like-for-like basis.

Daily-average heating loads for each city reveal that Kaiserslautern carries the majority of winter demand, while Pirmasens exhibits lower absolute loads but a similar seasonal pattern. Here, the heating deficit denotes the residual heat demand that remains after the available local solar heat contribution has been applied and must therefore be supplied by the auxiliary source. When considered separately, both cities experience pronounced

heating deficits in winter and shoulder seasons. These deficits are most severe in January and December and remain non-negligible through early spring, as shown in the individual deficit curves in Figure 8.

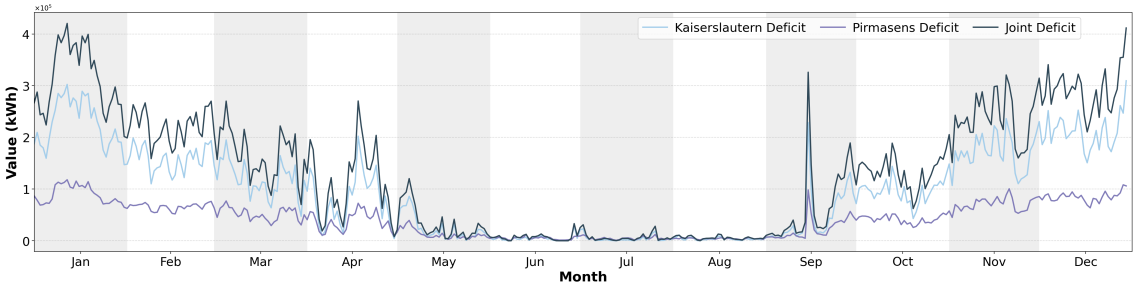

Figure 8: *Annual daily average heating deficit metrics for Pirmasens, Kaiserslautern and the joint regional system.*

A key result emerges when the two cities are assessed jointly. The combined deficit curve is consistently lower than the arithmetic sum of the two isolated-city deficits during large parts of the year. This behaviour is visible in the joint-series trace in Figure 8, where the "Joint Deficit" line lies below the individual city curves across most transitional months. The reduction indicates that temporal complementarity between the two cities plays a meaningful role. At many hours, surplus electricity from one municipality partially offsets the instantaneous deficit of the other, thereby lowering the joint unmet load.

Surplus electricity profiles provide even clearer evidence of cross-city synergy. Peaks of photovoltaic surplus in the two cities do not perfectly coincide due to differences in roof geometry, shading conditions and local weather patterns, but when the profiles are combined the joint series displays both higher peaks and fewer zero-surplus hours. This smoothing effect is shown in Figure 9, where the "Joint Surplus" curve exceeds the individual profiles during many hours of the year. The cooperative surplus series therefore expands the number of hours in which surplus electricity is available for secondary uses such as hydrogen production or thermal storage charging.

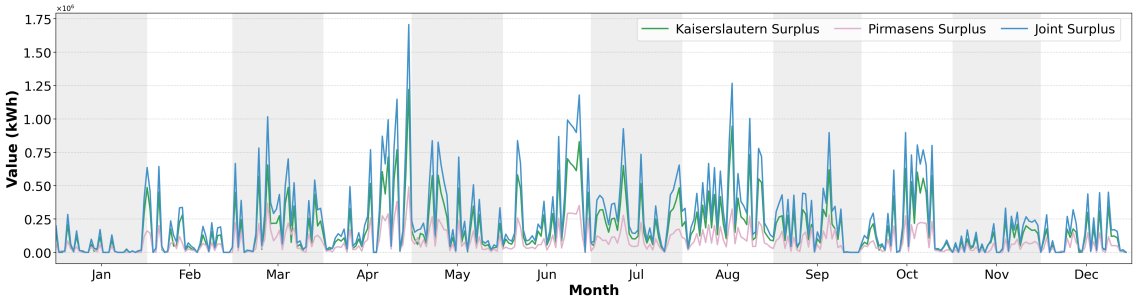

Figure 9: *Annual daily average electricity surplus metrics for Pirmasens, Kaiserslautern and the joint regional system.*

Taken together, these results show that meaningful coordination benefits can arise even between cities located only 30 km apart. This suggests that similar complementarities

are likely to exist across many neighbouring municipalities, particularly where demand profiles, building stocks and renewable potentials differ. Realising such opportunities at scale requires more than ad hoc, manually curated multi-city studies. It requires a uniform knowledge model that allows consistent indicators to be computed across jurisdictions and a consistent cross-city data access approach that makes it possible to discover and evaluate synergies systematically by querying distributed municipal knowledge graphs through a single interface. Within the context of the German Heat Planning Act, these findings therefore highlight the value of coordinated planning across municipalities that form part of the same functional region while preserving local autonomy in data management.

4.4 Power-to-hydrogen potential analysis and industrial relevance

Figure 10 compares the hydrogen production of the electrolyser when supplied only by surplus electricity from Pirmasens with the case where it is operated on the combined surplus from Pirmasens and Kaiserslautern. In both configurations hydrogen output exhibits a pronounced seasonal pattern, with limited production during the winter months and sustained operation from late spring to early autumn. The regional case consistently yields higher daily production levels, reflecting the fact that the aggregated surplus from the two cities increases both the number of hours with available surplus and the magnitude of the surplus in those hours.

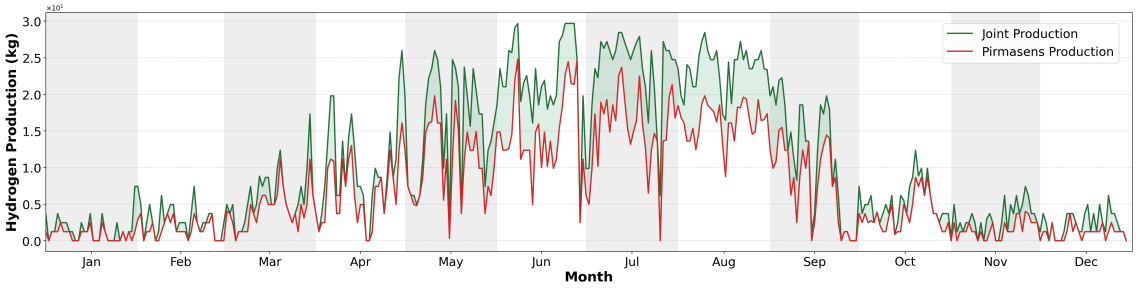

Figure 10: *Daily average hydrogen production for an electrolyser supplied by surplus electricity from Pirmasens only and from the combined regional surplus of both cities. The patched area shows the differences.*

The FLH of the electrolyser highlight the impact of regional coordination on utilisation. In the Pirmasens-only configuration the 1.8 MW_{el} unit achieves an equivalent of 1 883 full-load hours per year. When surplus electricity from Kaiserslautern is added, the FLH increase to 3 142 h, corresponding to an increase in utilisation by roughly two thirds. This improvement arises because periods without surplus in one city often coincide with at least some surplus in the other, as shown by the joint surplus series in Figure 9, which smooths the supply to the electrolyser and reduces idle time.

The higher utilisation directly translates into increased hydrogen output. Annual production rises from about 55.9 t H_2 in the Pirmasens-only case to 93.3 t H_2 when both cities contribute surplus electricity, which is shown in Table 3. Although these volumes are modest in absolute terms when compared with national hydrogen strategies, they are non-

negligible at the scale of a single industrial site and illustrate the order of magnitude of hydrogen that could be produced from surplus rooftop generation in a medium-sized urban region. In more detailed studies including grid constraints and market effects, these values would represent an upper bound.

Table 3: *Electrolyser utilisation, annual hydrogen production and indicative avoided CO_2 emissions for the Pirmasens-only and regional surplus configurations.*

Indicator	Pirmasens only	Regional joint case
Equivalent full-load hours ($h\,yr^{-1}$)	1 882.6	3 142.3
Annual hydrogen production ($kg\,yr^{-1}$)	55 905	93 312
Avoided CO_2 if displacing SMR ($t\,yr^{-1}$)	587.0	979.8

To indicate the associated decarbonisation potential, the simulated green hydrogen production is compared with a counterfactual supply from natural-gas-based steam methane reforming. Using a representative emission factor for conventional SMR, consistent with the methodological assumptions, the Pirmasens-only configuration would avoid approximately 0.59 $kt\,CO_2\,yr^{-1}$ if the entire output displaced conventional hydrogen, while the regional configuration would avoid around 0.98 $kt\,CO_2\,yr^{-1}$. These avoided emissions are small compared with regional CO_2 inventories but significant at plant scale, corresponding to a meaningful reduction in the fossil footprint of a single biorefinery.

From an industrial and governance perspective, the results suggest that power-to-hydrogen can act as a regional complement to otherwise predominantly building-focused heat planning. Cross-city coordination turns spatially scattered rooftop surpluses into a single, more continuous resource stream that can be aligned with the load profile of an energy-intensive site, thereby strengthening the local industrial base while making use of electricity that would otherwise remain unused. Within the framework of the German Heat Planning Act, which explicitly encourages municipalities to consider inter-municipal solutions where they are beneficial, this points to a role for hydrogen-oriented sector coupling as one of the regional options that can be examined alongside network expansion and efficiency measures. Rather than treating industrial decarbonisation and municipal heat planning as separate processes, the integrated knowledge-based approach used here allows both to be analysed within a common data and modelling environment, highlighting where regional surplus sharing can simultaneously support low-carbon heat supply and the substitution of fossil feedstocks in local industry.

5 Conclusions and future work

This paper has developed and applied a knowledge-graph-based framework for municipal and regional heat planning using The World Avatar (TWA) and its City Energy Analyst (CEA) agent. Two independent cities, Pirmasens and Kaiserslautern, are represented using common ontologies and analysed through a shared query framework to run building-level energy simulations. On this basis, the study carries out three layers of analysis, namely municipal assessments of heating demand, rooftop solar-thermal potential and

building-scale techno-economic performance with and without thermal energy storage, a regional coordination analysis that compares isolated and cooperative operation of the two cities, and a power-to-hydrogen assessment in which surplus electricity is converted to hydrogen for an industrial biorefinery.

At the individual city scale, the first step is to verify that the transferable workflow produces credible baseline patterns before interpreting the subsequent techno-economic and regional coordination results. The simulations reproduce plausible energy-use intensities and known structural differences between the two building stocks, indicating that the workflow can be transferred between municipalities without reconfiguration of the core model. Both cities possess rooftop solar-thermal potential, but this potential is unevenly distributed. To evaluate building-scale techno-economic performance, multiple collector sizes are assessed for each building and the configuration that maximises NPV is selected, with cases compared for operation with and without thermal energy storage. Pirmasens has a higher proportion of economically attractive systems, with around 69 % of buildings achieving a positive net present value (NPV) without storage, compared with about 56 % in Kaiserslautern. This reflects more heat-intensive domestic buildings and favourable low-rise roof geometries in Pirmasens, while the larger and more heterogeneous stock in Kaiserslautern offers a greater absolute number of viable projects. Thermal energy storage proves most beneficial when it is targeted at well-suited buildings. For those buildings with positive NPV, storage raises the mean project value from roughly 0.029 million EUR to 0.037 million EUR in both cities by improving the temporal match between solar gains and heating demand, while only slightly reducing the share of profitable projects.

The regional analyses show that treating the two municipal systems as a coupled region yields additional benefits that are not accessible under purely local planning. Complementary load and generation patterns reduce joint deficits and increase the number of hours with usable surplus, which in turn improves the utilisation of an industrial electrolyser connected to the planned biorefinery in Pirmasens. In the Pirmasens-only configuration, the 1.8 MW_{el} unit operates for the equivalent of about 1,900 full-load hours per year and produces roughly 56 $t\,H_2\,yr^{-1}$. When surplus electricity from both cities is pooled, full-load hours increase to more than 3,100 $h\,yr^{-1}$ and annual hydrogen production rises to about 93 $t\,H_2\,yr^{-1}$. If this hydrogen replaces conventionally produced supply, the associated plant-scale CO_2 savings are close to 1 $kt\,yr^{-1}$, showing that even the surplus from two medium-sized cities can already support kiloton-scale emission reductions at an individual industrial site and that larger regional portfolios could scale this effect further.

Beyond these quantitative findings, the study highlights the broader value of a knowledge-based semantic framework for the decarbonisation of urban and industrial energy systems. By embedding UBEM, municipal heat planning and power-to-hydrogen in a single semantic and computational environment, the framework shows how building-scale, city-scale and regional-scale decisions can be analysed in a consistent way and directly related to industrial decarbonisation options. In the context of the German Heat Planning Act, this offers a concrete pathway towards heat plans that are transparent, reproducible and capable of incorporating regional coordination and sector coupling alongside more traditional measures such as refurbishment and network expansion. The same approach is readily extensible to additional cities, technologies and sectors, suggesting that dynamic knowledge graphs can serve as an enabling infrastructure for evidence-based heat planning and

integrated energy transition strategies in other regions as well.

The constructed knowledge-based digital twin provides a scalable foundation for integrated energy planning and power-to-hydrogen analysis. As energy infrastructures increase in complexity and evolve towards cyber-physical systems, the future work is expected to extend beyond routing planning to support AI-enabled asset operational safety [59], risk early warning [60], and resilience management [61]. Mature AI surrogates, such as deep representation learning [60], tree ensemble methods [58, 59], that have proven track records in energy generation forecasting, fault diagnosis, and condition monitoring can be encapsulated as computational agents within a dynamic knowledge graph. This will enable the digital infrastructure to support a broader range of application scenarios. Beyond computation scenarios, the integration of Generative AI [41] can introduce a cognitive layer by positioning the joint architecture as a tool endpoint of the Model Context Protocol [44], which allows large language models to access and interpret distributed physical knowledge and computational agents. This strengthens semantic understanding and task orchestration, so that complex natural language instructions can automatically trigger the appropriate physics-based safety models for cross-domain reasoning and support human-centred, collaborative decision support.

Acknowledgements

This research was supported by the European Union's Horizon Europe research and innovation programme under grants 101058732 (JIDEP), 101074004 (C2IMPRESS), and 101188248 (CLIMATE-ADAPT4EOSC).

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT-5.2 in order to enhance the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Data and code availability

The codes developed for this work are available under an open-source licence on GitHub in The World Avatar repository https://github.com/TheWorldAvatar and https://github.com/cambridge-cares/TheWorldAvatar.

A Appendix

A.1 Namespaces

be: <https://www.theworldavatar.com/kg/ontobuiltenv/>
bs: <https://www.theworldavatar.com/kg/ontobuildingstructure/>
ems: <https://www.theworldavatar.com/kg/ontoems/>
ts: <https://www.theworldavatar.com/kg/ontotimeseries/>
ub: <https://www.theworldavatar.com/kg/ontoubemmp/>
geo: <http://www.opengis.net/ont/geosparql#>
om: <http://www.ontology-of-units-of-measure.org/resource/om-2/>
owl: <http://www.w3.org/2002/07/owl#>
rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
rdfs:<http://www.w3.org/2000/01/rdf-schema#>
time:<http://www.w3.org/2006/time#>
xsd: <http://www.w3.org/2001/XMLSchema#>

A.2 Visual interfaces for municipal and regional heat-planning analyses

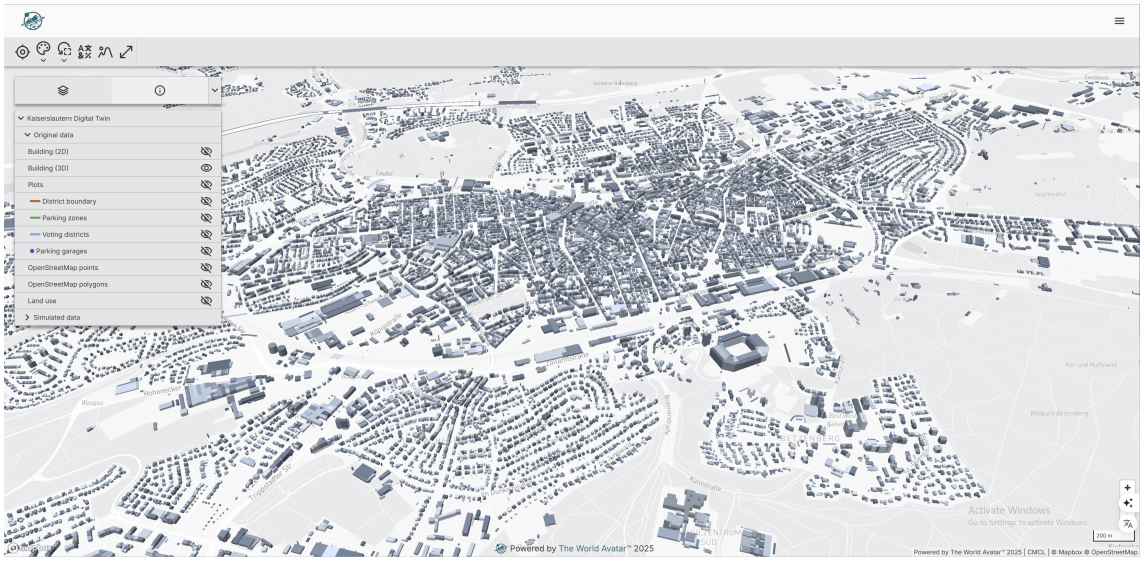

Figure A.1: *Interactive 3D visualisation of the Kaiserslautern digital twin, showing building geometries and associated urban layers.*

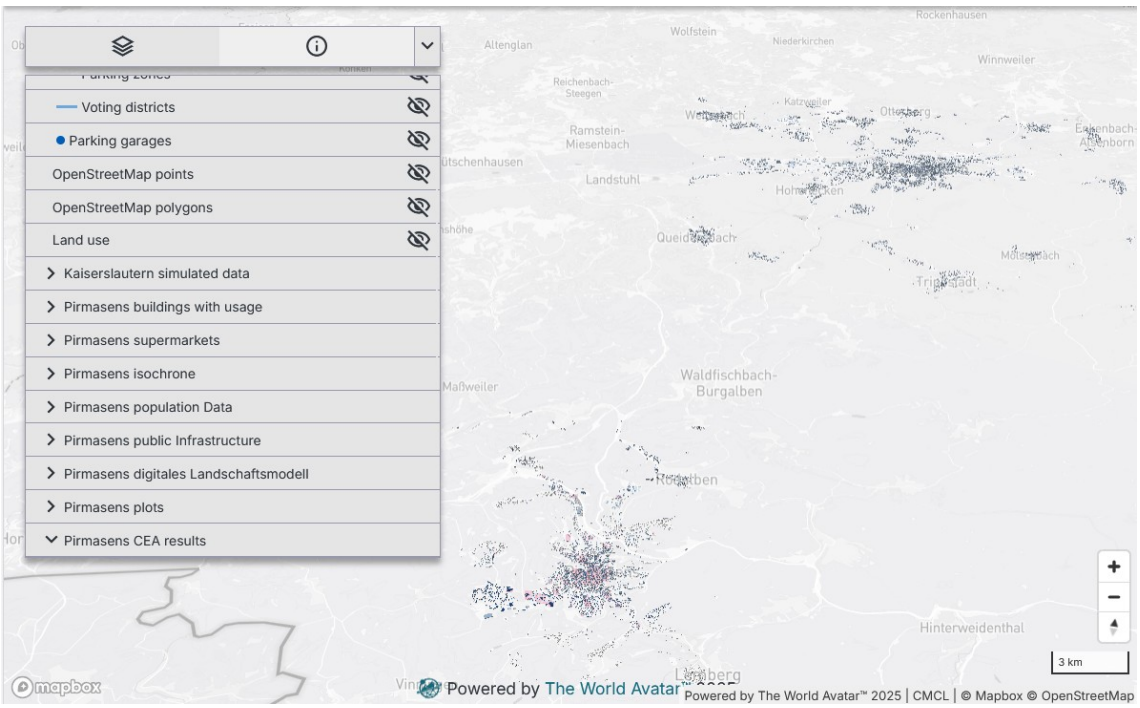

Figure A.2: *Regional map view linking the digital twins of Pirmasens and Kaiserslautern within the federated knowledge-graph environment.*

A.3 Implementation details and equations

Data streams and notation This section describes the data streams, variables, and notation used throughout the analyses, and how simulation outputs from the City Energy Analyst (CEA) agent are translated into inputs for the subsequent analyses. Unless stated otherwise, all quantities are evaluated as hourly time series over a full simulation year, with a fixed time-step, $\Delta t = 1$ h. All time series were generated by the CEA agent and retrieved via The World Avatar (TWA).

Two distinct but related data streams are used. The *heat stream* underpins the building-level assessment of solar-thermal systems and thermal energy storage (TES). For a given building, H_t denotes the heating energy demand in time step t. The quantity G_t represents the contribution at time step t to the reference solar-thermal energy generation profile obtained from CEA for a standard collector configuration. To represent alternative system sizes, the reference profile is linearly scaled such that

$$\tilde{G}_t = s\, G_t, \tag{A.1}$$

where the scaling factor s is defined by the ratio between the assumed installed collector area A_{inst} and the collector area associated with the reference profile G_t. In cases where G_t is reported per unit collector area, the reference area is 1 m^2 and $s = A_{\text{inst}}$. The scaled profile $\tilde{G}_t$ is used in the hourly heat-balance calculations described in the following sections.

The *electricity stream* is used for the analysis of photovoltaic (PV) surplus sharing and for the power-to-hydrogen assessment. For each building b and time step t, $P_{\text{PV}}(b,t)$ denotes

the rooftop PV electricity generation and $P_{\text{load}}(b,t)$ denotes the corresponding electricity demand. These quantities form the basis for computing hourly surplus and deficit power, which are subsequently aggregated at city and regional levels.

The thermal energy storage (TES) operation is described by the storage state of charge B_t and the usable thermal energy storage capacity E_{TES}, both expressed in units of energy. The electrolyser operation in the power-to-hydrogen analysis is characterised by the dispatched electrical power $P_{\text{el}}(t)$, which is constrained by the rated electrolyser capacity $P_{\text{el,max}}$.

All techno-economic parameters used in the calculations are treated as fixed, exogenous assumptions and are applied uniformly across the entire building stock. These include, but are not limited to, the gas boiler efficiency η_{boiler}, the discount rate r, the economic lifetime T, the TES sizing rule, cost coefficients for equipment, installation and operation, and the TES standing-loss factor q_l. For transparency and reproducibility, the numerical values of these parameters are summarised in Table 1 in the main text.

A.3.1 Heating consumption and rooftop solar-thermal potential analysis: implementation details

This section describes the building-level calculations used to assess heating demand coverage by rooftop solar-thermal systems, the operation of optional thermal energy storage (TES), and the resulting techno-economic performance. It provides the full implementation details corresponding to Section 3.2 in the main text.

Purpose and scope For each building, the algorithm evaluates alternative solar-thermal system sizes, optionally coupled with TES, using hourly heating demand and generation profiles. For a given system configuration, the procedure computes (i) hourly heat balances, (ii) residual fuel demand supplied by an auxiliary gas boiler, (iii) annual operating costs, and (iv) the net present value (NPV) over the system lifetime. The installed collector area is treated as a discrete design variable, and the configuration that maximises NPV is selected for each building.

Solar-thermal generation scaling The CEA agent provides a reference hourly solar-thermal generation profile G_t for each building. Alternative system sizes are represented by linearly scaling the reference profile to calculate a scaled profile $\tilde{G}_t$ as per Eq. (A.1).

Hourly heat balance without storage At each time step, scaled solar-thermal generation and heating demand are balanced to determine direct self-consumption, surplus heat, and residual unmet demand,

$$\begin{aligned} Q_{\text{self},t} &= \min\left(\tilde{G}_t, H_t\right), \\ Q_{\text{over},t} &= \max\left(\tilde{G}_t - H_t, 0\right), \\ Q_{\text{need},t} &= H_t - Q_{\text{self},t}. \end{aligned} \tag{A.2}$$

where H_t denotes the hourly useful heating energy demand. The quantity $Q_{\text{self},t}$ represents solar heat used instantaneously on site, $Q_{\text{over},t}$ is surplus heat energy available for storage, and $Q_{\text{need},t}$ is the remaining demand that must be met by the auxiliary heating system.

TES sizing and capacity When a storage-equipped configuration is considered, the usable energy storage capacity E_{TES} is calculated (in kWh),

$$E_{\text{TES}} = \frac{70 \cdot V}{1000}, \tag{A.3}$$

where the factor of 70 $\text{kWh}\,\text{m}^{-3}$ corresponds to the assumed usable energy content per unit volume of a water-based TES over a representative temperature swing and is applied uniformly across all buildings, and the tank volume V (in litres) is determined using an empirical linear sizing rule based on the installed collector area A_{inst} (in m^2), subject to minimum and maximum bounds,

$$V = \max\Big(V_{\min},\ \min\left(V_{\max},\ 38.9 \cdot A_{\text{inst}} + 216.7\right)\Big). \tag{A.4}$$

The lower bound $V_{\min}$ enforces a minimum viable storage size, while $V_{\max}$ limits the maximum volume and depends on building type (domestic or non-domestic). All numerical values are listed in Table 1 in the main text. The collector area A_{inst} is treated as a design variable and evaluated in discrete increments of one module (2.5 m^2), starting from a minimum of 10 m^2 and capped by the available solar-suitable roof area of the building.

TES operation Storage operation is simulated sequentially over each time step. First, the state of charge decays due to standing losses. Second, stored heat is discharged to cover unmet demand. Finally, surplus solar heat is used to charge the storage subject to the capacity constraint,

$$\begin{aligned} B_t^- &= q_l\, B_{t-1}, \\ w_t &= \min\left(Q_{\text{need},t},\, B_t^-\right), \\ B_t &= \min\left(B_t^- - w_t + Q_{\text{over},t},\, E_{\text{TES}}\right), \end{aligned} \tag{A.5}$$

where B_{t-1} and B_t denote the storage state of charge at the end of time step $t-1$ and t, respectively, B_t^- is an intermediate state after losses, q_l is the hourly retention factor and w_t is the energy discharged to meet heating demand.

Auxiliary heating and fuel demand Residual heating demand after solar utilisation and TES discharge is assumed to be supplied by existing on-site natural-gas boilers. The required fuel input is computed,

$$Q_{\text{gas,input}} = \frac{Q_{\text{heat}}}{\eta_{\text{boiler}}}, \tag{A.6}$$

where Q_{heat} is the useful heat delivered to the load and $\eta_{\text{boiler}} = 0.8$ is a constant assumed thermal efficiency.

Gas cost model The annual gas cost is modelled as a tiered, piecewise-linear function,

$$C_{\text{gas}} = \begin{cases} a_1 + 0.01h \cdot b_1, & h \leq 2{,}000, \\ a_2 + 0.01h \cdot b_2, & 2{,}000 < h \leq 10{,}000, \\ a_3 + 0.01h \cdot b_3, & 10{,}000 < h \leq 30{,}000, \\ a_4 + 0.01h \cdot b_4, & 30{,}000 < h \leq 100{,}000, \\ a_5 + 0.01h \cdot b_5, & 100{,}000 < h \leq 300{,}000, \\ a_6 + 0.01h \cdot b_6, & h > 300{,}000, \end{cases} \tag{A.7}$$

where a_k denotes the tier-specific fixed component of the annual gas tariff (EUR yr^{-1}) and b_k denotes the corresponding marginal gas price (EUR per 100 kWh of fuel input) for tier k. The tier boundaries are defined by the annual fuel use h, and all tariff parameters are adopted from Tsai et al. [69].

Investment and operating costs Capital and operating expenditures are computed as area-dependent functions,

$$\begin{aligned} I &= (400 \cdot A_{\text{inst}} + 2{,}500)\,(1 + 0.25), \\ O &= 3.5 \cdot A_{\text{inst}} + 147.5. \end{aligned} \tag{A.8}$$

where I is the upfront investment cost in EUR including a 25% installation multiplier, and O is the annual operating and maintenance cost in EUR yr^{-1}.

Economic evaluation For each candidate system size, annual cash flows are derived from the reduction in gas expenditure relative to the no-solar baseline, minus operating costs. The net present value is then computed over the system lifetime using Eq. (3) in the main text. The configuration that maximises NPV is selected as the optimal design for the building. Buildings with negative optimal NPV are classified as economically unattractive under the assumed parameters.

A.3.2 Regional synergy analysis: PV surplus sharing and coordination metrics

This section details the implementation of the regional photovoltaic (PV) surplus-sharing analysis described in Section 3.3 of the main text. The objective is to quantify how locally generated rooftop PV electricity can offset electrical demand within and across buildings, municipalities, and the combined region.

Purpose and scope The analysis operates on the electricity demand and generation profiles retrieved from TWA. It evaluates two configurations: isolated operation, where sharing is limited to buildings within a single city, and cooperative operation, where surplus electricity can be exchanged across cities. The resulting metrics characterise self-consumption, self-sufficiency, curtailment, and residual grid dependence.

Building-level electricity balance Rooftop PV generation and electrical load are partitioned into self-consumed energy, surplus, and deficit,

$$\begin{aligned} P_{\text{self}}(b,t) &= \min\Big(P_{\text{PV}}(b,t),\, P_{\text{load}}(b,t)\Big), \\ P_{\text{sur}}(b,t) &= \max\Big(P_{\text{PV}}(b,t) - P_{\text{load}}(b,t),\, 0\Big), \\ P_{\text{def}}(b,t) &= \max\Big(P_{\text{load}}(b,t) - P_{\text{PV}}(b,t),\, 0\Big). \end{aligned} \tag{A.9}$$

where $P_{\text{PV}}(b,t)$ is the rooftop PV generation and $P_{\text{load}}(b,t)$ the electrical demand of building b and time step t.

Aggregation and surplus sharing The surpluses and deficits are aggregated across all buildings under consideration,

$$\begin{aligned} P_{\text{sur}}(t) &= \sum_b P_{\text{sur}}(b,t), \\ P_{\text{def}}(t) &= \sum_b P_{\text{def}}(b,t). \end{aligned} \tag{A.10}$$

When surplus sharing is allowed, the amount of demand that can be covered by local or regional surplus is

$$P_{\text{sh}}(t) = \min\Big(P_{\text{sur}}(t),\, P_{\text{def}}(t)\Big), \tag{A.11}$$

and the remaining deficit supplied by the external grid is

$$P_{\text{def}}^{\text{rem}}(t) = P_{\text{def}}(t) - P_{\text{sh}}(t). \tag{A.12}$$

Annual performance indicators Based on the hourly balances, three aggregate indicators are defined. The self-consumption ratio,

$$\alpha = \frac{\sum_t \sum_b P_{\text{self}}(b,t)}{\sum_t \sum_b P_{\text{PV}}(b,t)}, \tag{A.13}$$

quantifies the fraction of PV generation used locally. The self-sufficiency ratio,

$$\beta = \frac{\sum_t \sum_b P_{\text{self}}(b,t)}{\sum_t \sum_b P_{\text{load}}(b,t)}, \tag{A.14}$$

measures the share of electrical demand covered by rooftop PV. The curtailment ratio,

$$\gamma = \frac{\sum_t P_{\text{sur}}(t) - P_{\text{sh}}(t)}{\sum_t \sum_b P_{\text{PV}}(b,t)}, \tag{A.15}$$

represents the fraction of PV generation that cannot be utilised and is therefore curtailed. These indicators are evaluated consistently for isolated and cooperative configurations to quantify the benefits of regional coordination.

A.3.3 Power-to-hydrogen and industrial integration analysis: implementation details

This section describes the implementation of the power-to-hydrogen (P2H) analysis used to assess the potential for using surplus photovoltaic (PV) electricity for hydrogen production. It provides the computational details corresponding to Section 3.4 in the main text.

Purpose and scope The P2H analysis operates on the surplus electricity profiles obtained from the regional PV surplus-sharing calculations. These profiles represent electricity that cannot be absorbed by local or regional electrical demand at each time step and are assumed to be available for hydrogen production via electrolysis. The electrolyser is modelled as a single industrial-scale unit with fixed rated power and constant efficiency, and the analysis yields annual hydrogen production and indicative avoided CO_2 emissions under these assumptions.

Electrolyser dispatch and hydrogen production At each time step t, the electrolyser is supplied with a dispatched electrical power $P_{\mathrm{el}}(t)$, which is determined by the available surplus electricity and constrained by the rated electrolyser capacity, as defined in the main text. Hydrogen production is calculated assuming a constant electrolyser efficiency η_{el} and using the lower heating value of hydrogen, $\mathrm{LHV}_{\mathrm{H_2}}$. The resulting hydrogen mass flow at time step t is

$$\dot{m}_{\mathrm{H_2}}(t) = \frac{\eta_{\mathrm{el}} P_{\mathrm{el}}(t)}{\mathrm{LHV}_{\mathrm{H_2}}}, \tag{A.16}$$

where $P_{\mathrm{el}}(t)$ is the electrical input to the electrolyser in kW, $\mathrm{LHV}_{\mathrm{H_2}}$ is expressed in $\mathrm{kWh\,kg^{-1}}$, and $\dot{m}_{\mathrm{H_2}}(t)$ is the hydrogen mass flow in $\mathrm{kg\,h^{-1}}$. The formulation assumes steady-state operation within each time step and neglects start-up, ramping and part-load efficiency effects.

Annual hydrogen production For the fixed time step $\Delta t = 1$ h, the total annual hydrogen production is obtained by summing the hourly mass flows over the simulation year,

$$M_{\mathrm{H_2}} = \sum_t \dot{m}_{\mathrm{H_2}}(t), \tag{A.17}$$

where $M_{\mathrm{H_2}}$ denotes the total hydrogen mass produced in one year, expressed in kg. This quantity is used directly in the performance indicators reported in Section 4 of the main text, such as annual hydrogen output and equivalent full-load hours of the electrolyser.

Indicative avoided CO_2 emissions To provide an order-of-magnitude indication of the potential decarbonisation impact, the annual green hydrogen production from Eq. A.17 is compared with a corresponding counterfactual supply from conventional natural-gas-based steam methane reforming (SMR). Introducing an emission factor $\mathrm{EF}_{\mathrm{SMR}}$ in units of

$\mathrm{kg\,CO_2/kg\,H_2}$, the corresponding annual avoided $\mathrm{CO_2}$ emissions are calculated as

$$E_{\mathrm{CO_2,avoided}} = \mathrm{EF_{SMR}} \cdot M_{\mathrm{H_2}}, \tag{A.18}$$

where $E_{\mathrm{CO_2,avoided}}$ is expressed in $\mathrm{kg\,CO_2\,yr^{-1}}$. The resulting emissions figures should be interpreted as indicative upper-bound estimates, as upstream emissions, system losses and alternative hydrogen supply pathways are not represented explicitly.

A.4 Nomenclature

Table A.1: *Nomenclature.*

Symbol	Description	Unit
Indices and sets		
b	Building index	–
t	Time step index	–
y	Year index	–
$\mathscr{B}$	Set of buildings	–
$\mathscr{T}$	Set of time steps	–
Temporal parameters		
H_{year}	Hours per year	$\mathrm{h\,yr^{-1}}$
Δt	Time step length	h
Building and PV parameters		
A_b	Building or roof area	$\mathrm{m^2}$
A_{inst}	Installed PV area	$\mathrm{m^2}$
CF_y	Capacity factor in year y	–
$\mathrm{EU}_{\mathrm{city,PV}}$	City-level PV electricity generation	$\mathrm{kWh\,yr^{-1}}$
$\mathrm{EUI}_{b,\mathrm{PV}}$	Building PV electricity generation intensity	$\mathrm{kWh\,m^{-2}\,yr^{-1}}$
G_t	Reference solar-thermal energy generation profile at time step t	kWh
$\tilde{G}_t$	Scaled solar-thermal generation profile at time step t	–
Energy flow variables		
H_t	Heating energy demand at time step t	kWh
$P_{\mathrm{PV}}(b,t)$	PV power generation of building b	kW
$P_{\mathrm{load}}(b,t)$	Electrical load of building b at time step t	kW
$P_{\mathrm{net}}(b,t)$	Net building electricity balance	kW
$P_{\mathrm{self}}(b,t)$	Self-consumed PV electricity	kW
$P_{\mathrm{sur}}(b,t)$	Surplus PV power from building b	kW
$P_{\mathrm{sur}}(t)$	Total system surplus power	kW
$P_{\mathrm{def}}(b,t)$	Power deficit for building b at time step t	kW
$P_{\mathrm{def}}(t)$	Total system power deficit at time step t	kW
$P_{\mathrm{def}}^{\mathrm{rem}}(t)$	Remaining unmet power deficit	kW
$P_{\mathrm{heat}}(b,t)$	Heat supplied to building b at time step t	kW
Q_{heat}	Useful heat output	kWh
$Q_{\mathrm{need},t}$	Heat demand at time step t	kWh
$Q_{\mathrm{over},t}$	Excess heat at time step t	kWh
$Q_{\mathrm{self},t}$	Self-consumed heat	kWh
$Q_{\mathrm{gas,input}}$	Gas energy input to boiler	kWh
$\mathrm{EU}_{\mathrm{city,heat}}$	City-level annual heat energy use	$\mathrm{kWh\,yr^{-1}}$
$\mathrm{EUI}_{b,\mathrm{heat}}$	Building heat energy use intensity	$\mathrm{kWh\,m^{-2}\,yr^{-1}}$

Continued on next page

Table A.1 – continued from previous page

Symbol	Description	Unit
Storage variables		
η_{boiler}	Boiler efficiency	–
B_t	Storage state of charge at time step t	kWh
B_{t-1}	Storage state of charge at time step $t-1$	kWh
B_t^-	Storage discharge at time step t	kWh
E_{TES}	Usable thermal energy storage capacity	kWh
V	Storage volume or state variable	m^3 or kWh
V_{max}	Maximum storage capacity	m^3 or kWh
V_{min}	Minimum storage level	m^3 or kWh
q_l	Thermal loss rate	–
w_t	Energy discharged to meet heating demand at time step t	–
Hydrogen and conversion systems		
η_{el}	Electrolyser efficiency	–
EF_{SMR}	Emission factor for steam methane reforming hydrogen	kg CO_2 kg_{H2}^{-1}
$E_{\text{CO}_2,\text{avoided}}$	Avoided CO_2 emissions	kg CO_2
FLH	Electrolyser full load hours	h yr^{-1}
LHV_{H_2}	Lower heating value of hydrogen	kWh kg^{-1}
$\dot{m}_{\text{H}_2}(t)$	Hydrogen mass production rate at time step t	kg h^{-1}
M_{H_2}	Hydrogen mass produced	kg
$P_{\text{el,max}}$	Maximum electrolyser power capacity	kW
$P_{\text{el}}(t)$	Electrolyser power consumption at time step t	kW
$P_{\text{sh}}(t)$	Power supplied to hydrogen system	kW
Economic variables and model parameters		
α	Cost or scaling coefficient (model parameter)	–
β	Cost or scaling coefficient (model parameter)	–
γ	Discount or degradation factor	–
NPV	Net present value	€
I	Investment cost	€
O	Operating cost	€ yr^{-1}
C_0	Initial capital cost	€
C_{gas}	Gas price	€ kWh^{-1}
r	Discount rate	–
a_k	Cost coefficient for gas consumption tier k	€
b_k	Cost coefficient for gas consumption tier k	€

References

[1] J. Akroyd, S. Mosbach, A. Bhave, and M. Kraft. Universal digital twin - a dynamic knowledge graph. *Data-Centric Engineering*, 2, 2021. doi:10.1017/dce.2021.10.

[2] F. Bandeiras, M. Carvalho, B. Bahmankhah, C. Silva, J. Gomes, D. Aelenei, et al. Towards net zero energy in industrial and commercial buildings in portugal. *Renewable and Sustainable Energy Reviews*, 119:109580, 2020. doi:10.1016/j.rser.2019.109580.

[3] R. Battle and D. Kolas. Enabling the geospatial semantic web with parliament and geosparql. *Semantic Web*, 3(4):355–370, Nov. 2012. doi:10.3233/SW-2012-0065.

[4] M. Booshehri, L. Emele, S. Flügel, H. Förster, J. Frey, U. Frey, M. Glauer, J. Hastings, C. Hofmann, C. Hoyer-Klick, L. Hülk, A. Kleinau, K. Knosala, L. Kotzur, P. Kuckertz, T. Mossakowski, C. Muschner, F. Neuhaus, M. Pehl, M. Robinius, V. Sehn, and M. Stappel. Introducing the open energy ontology: Enhancing data interpretation and interfacing in energy systems analysis. *Energy and AI*, 5:100074, Sept. 2021. doi:10.1016/j.egyai.2021.100074.

[5] C. Brandoni and F. Polonara. The role of municipal energy planning in the regional energy-planning process. *Energy*, 48(1):323–338, Dec. 2012. doi:10.1016/j.energy.2012.06.061.

[6] Bundesministerium für Wohnen, Stadtentwicklung und Bauwesen. Kommunale wärmeplanung. für eine deutschlandweit zukunftsfeste und bezahlbare wärmeversorgung, 2023. URL https://www.bmwsb.bund.de/Webs/BMWSB/DE/themen/stadt-wohnen/WPG/WPG-node.html. Accessed: 25 June 2024.

[7] J. Chen, J. Bai, J. Xu, F. Farazi, S. Mosbach, J. Akroyd, and M. Kraft. Transforming building retrofits: Linking energy, equity, and health insights from the world avatar. *Advances in Applied Energy*, page 100230, 2025. doi:10.1016/j.adapen.2025.100230.

[8] Co2online.de. Solarthermie: Preise, Kosten & Amortisation. https://www.co2online.de/modernisieren-und-bauen/solarthermie/solarthermie-preise-kosten-amortisation/, 2025. Accessed: 19 Nov 2025.

[9] A. R. Dahiru, A. Vuokila, and M. Huuhtanen. Recent development in power-to-x: Part i - a review on techno-economic analysis. *Journal of Energy Storage*, 56: 105861, Dec. 2022. doi:10.1016/j.est.2022.105861.

[10] J. Dong, Y. Schwartz, A. Mavrogianni, I. Korolija, and D. Mumovic. A review of approaches and applications in building stock energy and indoor environment modelling. *Building Services Engineering Research & Technology*, 44(3):333–354, May 2023. doi:10.1177/01436244231163084.

[11] Eclipse. Eclipse RDF4J, 2024. URL https://rdf4j.org/. Accessed: 25 June 2024.

[12] Effizienzhaus-online.de. Solarthermie – Tipps zur Anlagengröße. https://www.effizienzhaus-online.de/solarthermie-tipps-anlagengroesse/, 2025. Accessed: 19 Nov 2025.

[13] Energieheld.de. Solarthermie Größe: Die optimale Solarthermie Kollektor Größe — Solaranlage Größe. https://www.energieheld.de/solaranlage/solarthermie/groesse-dimensionierung, 2025. Accessed: 19 Nov 2025.

[14] Energieheld.de. Kosten für Solarthermie: Alles zu Solarthermie Kosten — Solarthermieanlage Kosten. https://www.energieheld.de/solaranlage/solarthermie/kosten, 2025. Accessed: 19 Nov 2025.

[15] H. Fang, H. Tan, R. Kosonen, X. Yuan, K. Jiang, and R. Ding. Study of the data augmentation approach for building energy prediction beyond historical scenarios. *Buildings*, 13(2):326, Jan. 2023. doi:10.3390/buildings13020326.

[16] J. Fonseca, D. Thomas, R. Mok, S. Hsieh, B. Krishna Sreepathi, G. Happle, L. Rogenhofer, M. Niffeler, Z. Shi, M. Mosteiro Romero, Jack-Hawthorne, F. Khayatian, E. Riegelbauer, lguilhermers, B. Lie Ong, orenkiwi, jarunan, M. MeshkinKiya, M. Sulzer, R. Molony, paulneitzel, A. Elesawy, A. Bosova, J. A. BELLO ACOSTA, T. H, luzpaz, AlexJew, A. Ceruti, Juveria-shah, and T. G. Badger. architecture-building-systems/cityenergyanalyst: Cityenergyanalyst v.3.35.4. *Zenodo*, Feb. 2024. doi:10.5281/zenodo.10697424.

[17] J. A. Fonseca, N. Thuy-An, A. Schlueter, and F. Marechal. City Energy Analyst (CEA): Integrated framework for analysis and optimization of building energy systems in neighborhoods and city districts. *Energy and Buildings*, 113:202–226, Feb. 2016. doi:10.1016/j.enbuild.2015.11.055.

[18] L. Fraccascia and I. Giannoccaro. What, where, and how measuring industrial symbiosis: A reasoned taxonomy of relevant indicators. *Resources, Conservation and Recycling*, 157:104799, June 2020. doi:10.1016/j.resconrec.2020.104799.

[19] D. Gao, B. Zhao, T. H. Kwan, Y. Hao, and G. Pei. The spatial and temporal mismatch phenomenon in solar space heating applications: status and solutions. *Applied Energy*, 321:119326, 2022. doi:10.1016/j.apenergy.2022.119326.

[20] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, June 1993. doi:10.1006/knac.1993.1008.

[21] B. Guler, E. Çelebi, and J. Nathwani. A 'regional energy hub' for achieving a low-carbon energy transition. *Energy Policy*, 113:376–385, Feb. 2018. doi:10.1016/j.enpol.2017.10.044.

[22] H. Guy, S. Vittoz, G. Caputo, and T. Thiery. Benchmarking the energy performance of European commercial buildings with a Bayesian modeling framework. *Energy and Buildings*, 299:113595, 2023. doi:10.1016/j.enbuild.2023.113595.

[23] G. Happle, J. A. Fonseca, and A. Schlueter. Effects of air infiltration modeling approaches in urban building energy demand forecasts. *Energy Procedia*, 122: 283–288, Sept. 2017. doi:10.1016/j.egypro.2017.07.323.

[24] G. Happle, Z. Shi, S. Hsieh, B. Ong, J. A. Fonseca, and A. Schlueter. Identifying carbon emission reduction potentials of BIPV in high-density cities in Southeast Asia. *Journal of Physics: Conference Series*, 1343(1):012077, 2019. doi:10.1088/1742-6596/1343/1/012077.

[25] Y. He, Y. Zhou, J. Liu, Z. Liu, and G. Zhang. An inter-city energy migration framework for regional energy balance through daily commuting fuel-cell vehicles. *Applied Energy*, 324:119714, Oct. 2022. doi:10.1016/j.apenergy.2022.119714.

[26] L. Hofbauer, W. McDowall, and S. Pye. Challenges and opportunities for energy system modelling to foster multi-level governance of energy transitions. *Renewable and Sustainable Energy Reviews*, 161:112330, June 2022. doi:10.1016/j.rser.2022.112330.

[27] M. Hofmeister, J. Bai, G. Brownbridge, S. Mosbach, K. F. Lee, F. Farazi, M. Hillman, M. Agarwal, S. Ganguly, J. Akroyd, and M. Kraft. Semantic agent framework for automated flood assessment using dynamic knowledge graphs. *Data-Centric Engineering*, 2024. doi:10.1017/dce.2024.11.

[28] M. Hofmeister, G. Brownbridge, M. Hillman, S. Mosbach, J. Akroyd, K. F. Lee, and M. Kraft. Cross-domain flood risk assessment for smart cities using dynamic knowledge graphs. *Sustainable Cities and Society*, 101:105113, Feb. 2024. doi:10.1016/j.scs.2023.105113.

[29] M. Hofmeister, K. Lee, Y.-K. Tsai, M. Müller, K. Nagarajan, S. Mosbach, J. Akroyd, and M. Kraft. Dynamic control of district heating networks with integrated emission modelling: A dynamic knowledge graph approach. *Energy and AI*, 17, 2024. doi:10.1016/j.egyai.2024.100376.

[30] K. Hori, T. Matsui, T. Hasuike, K.-i. Fukui, and T. Machimura. Development and application of the renewable energy regional optimization utility tool for environmental sustainability: Reroutes. *Renewable Energy*, 93:548–561, Aug. 2016. doi:10.1016/j.renene.2016.02.051.

[31] IEA SHC Task 54. Info sheets. https://task54.iea-shc.org/info-sheets, 2025. Accessed: 19 Nov 2025.

[32] International Energy Agency. Chile policy for use of wood and derivatives for heating. https://www.iea.org/policies/6553-chile-policy-for-use-of-wood-and-derivatives-for-heating, 2019. Accessed: 20 Jan 2026.

[33] International Energy Agency. Renewable heat – renewables 2021. https://www.iea.org/reports/renewables-2021/renewable-heat?mode=market&region=World&publication=2021&product=Total, 2021. Accessed: 23 July 2025.

[34] International Energy Agency. Clean winter heating plan in northern china (2017 - 2021). https://www.iea.org/policies/7906-clean-winter-heating-plan-in-northern-china-2017-2021, 2025. Accessed: 20 Jan 2026.

[35] M. Jahami, P. Singh, and B. Khandelwal. Life cycle assessment of smr and electrified-smr with renewable energy systems: Projecting emissions and optimizing hydrogen production for california's 2035 goals. *Renewable Energy*, 243:122501, Apr. 2025. doi:10.1016/j.renene.2025.122501.

[36] F. Johari, G. Peronato, P. Sadeghian, X. Zhao, and J. Widén. Urban building energy modeling: State of the art and future prospects. *Renewable and Sustainable Energy Reviews*, 128, 2020. doi:10.1016/j.rser.2020.109902.

[37] J. Keirstead, M. Jennings, and A. Sivakumar. A review of urban energy system models: Approaches, challenges and opportunities. *Renewable and Sustainable Energy Reviews*, 16(6):3847–3866, Aug. 2012. doi:10.1016/j.rser.2012.02.047.

[38] S. Kim, H. Shin, and J. Ahn. Energy performance analysis of airport terminal buildings by use of architectural, operational information and benchmark metrics. *Journal of Air Transport Management*, 83:101762, 2020. doi:10.1016/j.jairtraman.2020.101762.

[39] Klima- Energi- og Forsyningsministeriet. Bekendtgørelse af lov om varmeforsyning, 2020. URL https://www.retsinformation.dk/eli/lta/2020/1215. Accessed: 24 July 2025.

[40] G. Klyne. Resource description framework (RDF): Concepts and abstract syntax. *http://www. w3. org/TR/rdf-concepts/*, 2004.

[41] J. Kocoń, I. Cichecki, O. Kaszyca, M. Kochanek, D. Szydło, J. Baran, J. Bielaniewicz, M. Gruza, A. Janz, K. Kanclerz, A. Kocoń, B. Koptyra, W. Mieleszczenko-Kowszewicz, P. Miłkowski, M. Oleksy, M. Piasecki, L. Radliński, K. Wojtasik, S. Woźniak, and P. Kazienko. ChatGPT: Jack of all trades, master of none. *Information Fusion*, 99:101861, 2023. doi:10.1016/j.inffus.2023.101861.

[42] L. Kong, S. C. L. Koh, V. Sena, D. Robinson, and M. Wood. Towards a resilience evaluation framework for hydrogen supply chains: A systematic literature review and future research agenda. *International Journal of Hydrogen Energy*, 99:589–606, 2025. doi:10.1016/j.ijhydene.2024.12.164.

[43] M. Kraft and S. Mosbach. The future of computational modelling in reaction engineering. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1924):3633–3644, 2010. doi:10.1098/rsta.2010.0124.

[44] H. Li, Y. Xu, and T. Hong. Energyplus-mcp: A model-context-protocol server for ai-driven building energy modeling. *SoftwareX*, 32:102367, 2025. doi:10.1016/j.softx.2025.102367.

[45] M. Q. Lim, X. Wang, O. Inderwildi, and M. Kraft. *The World Avatar—A World Model for Facilitating Interoperability*, page 39–53. Springer International Publishing, 2022. ISBN 9783030862152. doi:10.1007/978-3-030-86215-2_4.

[46] A. Lyden, C. S. Brown, I. Kolo, G. Falcone, and D. Friedrich. Seasonal thermal energy storage in smart energy systems: District-level applications and modelling approaches. *Renewable and Sustainable Energy Reviews*, 167:112760, 2022. doi:10.1016/j.rser.2022.112760.

[47] P. R. Mazzarino, S. Finocchiaro, L. Barbierato, D. S. Schiera, L. Bottaccioli, and E. Patti. An automated tool for urban building energy modelling: From sparse datasets to CityJSON. In *2024 IEEE International Conference on Environment and Electrical Engineering and 2024 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, pages 1–7, 2024. doi:10.1109/EEEIC/ICPSEurope61470.2024.10751147.

[48] P. Mehta, D. Griego, A. Nunez-Jimenez, and A. Schlueter. The Impact of self-consumption regulation on individual and community solar PV adoption in Switzerland: an agent-based model. *Journal of Physics: Conference Series*, 1343:012143, 2019. doi:10.1088/1742-6596/1343/1/012143.

[49] C. Minke, M. Suermann, B. Bensmann, and R. Hanke-Rauschenbach. Is iridium demand a potential bottleneck in the realization of large-scale pem water electrolysis? *International Journal of Hydrogen Energy*, 46(46):23581–23590, 2021. doi:10.1016/j.ijhydene.2021.04.174.

[50] E. Moreci, G. Ciulla, and V. Lo Brano. Annual heating energy requirements of office buildings in a european climate. *Sustainable Cities and Society*, 20:81–95, 2016. doi:10.1016/j.scs.2015.10.005.

[51] P. Morgenstern, M. Li, R. Raslan, P. Ruyssevelt, and A. Wright. Benchmarking acute hospitals: Composite electricity targets based on departmental consumption intensities? *Energy and Buildings*, 118:277–290, 2016. doi:10.1016/j.enbuild.2016.02.052.

[52] M. Mosteiro-Romero, J. A. Fonseca, and A. Schlueter. Seasonal effects of input parameters in urban-scale building energy simulation. *Energy Procedia*, 122:433–438, Sept. 2017. doi:10.1016/j.egypro.2017.07.459.

[53] M. Mosteiro-Romero, D. Maiullari, F. Collins, A. Schlueter, and A. V. Timmeren. District-scale energy demand modeling and urban microclimate: A case study in The Netherlands. *Journal of Physics: Conference Series*, 1343(1):012003, 2019. doi:10.1088/1742-6596/1343/1/012003.

[54] G. Nascimento da Silva, P. R. R. Rochedo, and A. Szklo. Renewable hydrogen production to deal with wind power surpluses and mitigate carbon dioxide emissions from oil refineries. *Applied Energy*, 311:118631, Apr. 2022. doi:10.1016/j.apenergy.2022.118631.

[55] S. G. Nnabuife, A. K. Hamzat, J. Whidborne, B. Kuang, and K. W. Jenkins. Integration of renewable energy sources in tandem with electrolysis: A technology review for green hydrogen production. *International Journal of Hydrogen Energy*, 107:218–240, Mar. 2025. doi:10.1016/j.ijhydene.2024.06.342.

[56] A. Oraiopoulos and B. Howard. On the accuracy of urban building energy modelling. *Renewable and Sustainable Energy Reviews*, 158:111976, Apr. 2022. doi:10.1016/j.rser.2021.111976.

[57] J. C. Osorio-Aravena, A. Aghahosseini, D. Bogdanov, U. Caldera, N. Ghorbani, T. N. O. Mensah, J. Haas, E. Muñoz-Cerón, and C. Breyer. Synergies of electrical and sectoral integration: Analysing geographical multi-node scenarios with sector coupling variations for a transition towards a fully renewables-based energy system. *Energy*, 279:128038, Sept. 2023. doi:10.1016/j.energy.2023.128038.

[58] B. Patnaik, M. Mishra, R. C. Bansal, and R. K. Jena. Modwt-xgboost based smart energy solution for fault detection and classification in a smart microgrid. *Applied Energy*, 285:116457, 2021. doi:10.1016/j.apenergy.2021.116457.

[59] Y. Peng, G. Fu, B. Sun, J. Chen, W. Zhang, M. Ren, and H. Zhang. Data-driven collapse strength modelling for the screen pipes with internal corrosion defect based on finite element analysis and tree-based machine learning. *Ocean Engineering*, 279:114400, 2023. doi:10.1016/j.oceaneng.2023.114400.

[60] Y. Peng, G. Fu, J. Chen, B. Sun, and X. Sun. Bottom-hole pressure inversion method for nature gas wells based on blowout combustion flame shape parameters. *Energy*, 294:130673, 2024. doi:10.1016/j.energy.2024.130673.

[61] Y. Peng, G. Fu, B. Sun, X. Sun, J. Wang, and J. Chen. Well blowout flame's thermal radiation prediction under environmental wind based on multi-point heat sources and inverse analysis. *Geoenergy Science and Engineering*, 234:212609, 2024. doi:10.1016/j.geoen.2023.212609.

[62] PFI Germany. Bioraffinerie pirmasens: Nachhaltige energie aus biomasse, 2025. URL https://pfi-germany.de/en/biotechnologie/bioraffinerie-im-energiepark/. Accessed: 19 Nov 2025.

[63] Planbureau voor de Leefomgeving. Overzicht transitievisies warmte, 2022. URL https://www.pbl.nl/publicaties/overzicht-transitievisies-warmte. Accessed: 24 July 2025.

[64] C. F. Reinhart and C. Cerezo Davila. Urban building energy modeling – a review of a nascent field. *Building and Environment*, 97:196–202, Feb. 2016. doi:10.1016/j.buildenv.2015.12.001.

[65] H. Rijgersberg, M. van Assem, and J. Top. Ontology of units of measure and related concepts. *Semantic Web*, 4(1):3–13, Feb. 2013. doi:10.3233/SW-2012-0069.

[66] A. Schuler, C. Weber, and U. Fahl. Energy consumption for space heating of west-German households: Empirical evidence, scenario projections and policy implications. *Energy Policy*, 28(12):877–894, 2000. doi:10.1016/S0301-4215(00)00074-4.

[67] Z. Shi, H. Silvennoinen, A. Chadzynski, A. von Richthofen, M. Kraft, S. Cairns, and P. Herthogs. Defining archetypes of mixed-use developments using google maps api data. *Environment and Planning B: Urban Analytics and City Science*, 50(6): 1607–1623, Dec. 2022. doi:10.1177/23998083221141428.

[68] Y. R. Tan, M. Hofmeister, S. Z. Phua, G. Brownbridge, K. Rustagi, J. Akroyd, S. Mosbach, A. Bhave, and M. Kraft. Beyond connected digital twins – can digital twins really deliver sustainable cities? *Sustainable Cities and Society*, 131: 106596, 2025. doi:10.1016/j.scs.2025.106596.

[69] Y.-K. Tsai, M. Hofmeister, S. Ganguly, K. Rustagi, Y. R. Tan, S. Mosbach, J. Akroyd, and M. Kraft. Municipal heat planning within the world avatar. *Energy and AI*, 20:100479, May 2025. doi:10.1016/j.egyai.2025.100479.

[70] UK Parliament. Local area energy planning for net zero, 2021. URL https://post.parliament.uk/research-briefings/post-pn-0703/. Accessed: 24 July 2025.

[71] S. Vaisi, F. Pilla, and S. J. McCormack. Recommending a thermal energy benchmark based on CIBSE TM46 for typical college buildings and creating monthly energy models. *Energy and Buildings*, 176:296–309, 2018. doi:10.1016/j.enbuild.2018.07.041.

[72] T. J. Wallington, M. Woody, G. M. Lewis, G. A. Keoleian, E. J. Adler, J. R. R. A. Martins, and M. D. Collette. Hydrogen as a sustainable transportation fuel. *Renewable and Sustainable Energy Reviews*, 217:115725, July 2025. doi:10.1016/j.rser.2025.115725.

[73] K. Wang, M. Wang, and J. Zhou. Integrating heritage conservation into urban building energy modelling for retrofit decision-making in historic districts. *npj Heritage Science*, 14(1):1, 2026. doi:10.1038/s40494-025-02186-9.

[74] M. Wang, J. Zhou, Y. Liang, H. Yu, and R. Jing. Climate change impacts on city-scale building energy performance based on gis-informed urban building energy modelling. *Sustainable Cities and Society*, 125:106331, May 2025. doi:10.1016/j.scs.2025.106331.

[75] W. Wang, R. Jing, Y. Zhao, C. Zhang, and X. Wang. A load-complementarity combined flexible clustering approach for large-scale urban energy-water nexus optimization. *Applied Energy*, 270:115163, July 2020. doi:10.1016/j.apenergy.2020.115163.

[76] H. Yin, Z. Zhang, Y. Wan, Z. Gao, Y. Guo, and R. Xiao. Sustainable network analysis and coordinated development simulation of urban agglomerations from multiple perspectives. *Journal of Cleaner Production*, 413:137378, Aug. 2023. doi:10.1016/j.jclepro.2023.137378.

[77] J. Zhou, P. Fennell, I. Korolija, Z. Fang, R. Tang, and P. Ruyssevelt. Review of non-domestic building stock modelling studies under socio-technical system framework. *Journal of Building Engineering*, 97:110873, Nov. 2024. doi:10.1016/j.jobe.2024.110873.

[78] J. Zhou, J. Chen, K. Wang, B. Xu, X. Lin, J. Yan, and M. Wang. Projecting city-scale energy performance under future weather scenarios through urban building energy modelling incorporating cluster archetypes. *Sustainable Cities and Society*, 134: 106941, Nov. 2025. doi:10.1016/j.scs.2025.106941.

[79] J. Zhou, Z. Fang, G. Yin, J. Chen, X. Zhou, R. Jing, and M. Wang. Automated roof segmentation for assessing urban scale energy performance and indoor thermal environment of photovoltaic and green roofs. *Energy*, 338:138958, Nov. 2025. doi:10.1016/j.energy.2025.138958.

[80] J. Zhou, P. Fennell, I. Korolija, and P. Ruyssevelt. Percentage-based thermal zoning approach for enhanced stock-level building energy performance modelling. *Energy and Buildings*, 329:115231, 2025. doi:10.1016/j.enbuild.2024.115231.

[81] J. Zhou, I. Korolija, P. Fennell, and P. Ruyssevelt. Supplementing building envelope information for physics-based modelling with data-driven approaches based on public datasets. In *Multiphysics and Multiscale Building Physics*, page 316–321. Springer Nature, 2025. doi:10.1007/978-981-97-8313-7_43.

[82] J. Zhou, J. Li, J. Xie, X. Dong, K. Wang, R. Jing, R. Tang, and M. Wang. State-of-the-art review of urban building energy modelling on supporting sustainable development goals. *Applied Energy*, 402:126924, Dec. 2025. doi:10.1016/j.apenergy.2025.126924.

[83] X. Zhou, A. Eibeck, M. Q. Lim, N. B. Krdzavac, and M. Kraft. An agent composition framework for the j-park simulator - a knowledge graph for the process industry. *Computers & Chemical Engineering*, 130:106577, 2019. doi:10.1016/j.compchemeng.2019.106577.