Computers & Graphics 140 (2026) 104747

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on CGVC'25

# Steering the transition: A visual analytics approach to interactive energy scenario modeling ☆

Dany Laksono *, Radu Jianu, Aidan Slingsby

*City St. George's, University of London, Northampton Square, London, EC1V 0HB, United Kingdom*

ARTICLE INFO

*Keywords:*
Visual analytics
Multivariate visualization
Scenario modeling

ABSTRACT

Navigating the transition to net-zero requires energy planners to rapidly assess and contrast the impacts of competing policy decisions. While computational models excel at large-scale optimization, they often lack the context to navigate the nuanced trade-offs and conflicting priorities of public policy. To bridge this gap, we present a visual analytics system that integrates human-in-the-loop expertise into the planning process, ensuring data-driven insights are tempered by practical judgment. Our approach introduces a modular planning paradigm designed to reduce the cognitive load of balancing multi-objective goals. Building upon previous work in human-in-the-loop decarbonization platforms, we present a visualization framework using model-driven glyphs to provide a consistent representation of multivariate data across geographic scales, ensuring an equitable assessment of interventions. Finally, we provide a tightly integrated workflow that allows planners to move fluidly between composing interventions, simulating decarbonization impacts, and performing cross-scenario comparisons. Through a detailed application scenario, we illustrate how this system accelerates the sensemaking process, allowing planners to benchmark strategies and deliver evidence-based policies with confidence.

## 1. Introduction

The transition to net-zero emissions presents local authorities and urban planners with a multifaceted policy challenge. Effective energy scenario modeling requires a holistic approach that accounts for carbon reduction potential, installation and operational costs, and energy demand alongside crucial socio-demographic dimensions such as fuel poverty and housing inequality [1,2].

The effort to design inclusive decarbonization plans is currently hindered by two fundamental limitations in existing decision-support tools. **First**, regarding the planning process, many systems rely on "black-box" optimization models that generate opaque, monolithic strategies [3,4]. This lack of transparency and flexibility is ill-suited for the highly negotiable and iterative nature of public policy creation. Such rigidity fails to accommodate the nuanced trade-offs and stakeholder dialogue essential for developing publicly accountable and socially equitable policies [5]. **Second**, geographic maps are indispensable for energy decision-making. However, current platforms struggle to visualize the high-dimensional data required for effective scenario modeling. Most rely on a traditional "map layer and slider" metaphor that is poorly suited for the multivariate and temporal datasets inherent to energy modeling. This approach forces planners to mentally fuse information from disjointed, sequential layers while toggling through time-series sliders, which might obscure the complex interrelationships between social, economic, and technical variables [6,7].

To bridge these gaps, we propose a visual analytics framework that empowers planners to construct, test, and dynamically tune localized decarbonization initiatives. Operating under a "Filter-Prioritize-Upgrade" methodology, the system allows users to interactively segregate building stocks, rank specific cohorts for improvement, and evaluate modular interventions. Expanding on our initial proof-of-concept [8], this study details three major technical and design contributions: **(i) A Multi-scale Visual Framework:** We utilize a model-driven glyph design to uniformly represent multivariate and temporal data across geographic hierarchies, ensuring consistent analysis at every scale. **(ii) A Modular Scenario Builder:** We introduce a component-based paradigm that decomposes complex planning into manageable units, allowing for the comparison of diverse scenario interventions against competing policy goals. **(iii) A Human-in-the-loop Workflow:** Our system provides a tightly integrated environment where users can fluidly move between data exploration, simulation, and real-time strategy refinement.

☆ This article is part of a Special issue entitled: 'CAG_CGVC'25' published in Computers & Graphics.

* Corresponding author.
*E-mail address:* dany.laksono@city.ac.uk (D. Laksono).

https://doi.org/10.1016/j.cag.2026.104747
Received 28 February 2026; Received in revised form 26 August 2026; Accepted 26 August 2026
Available online 2 September 2026

Developed through interdisciplinary collaboration with a UK-based energy consultant, our platform extends previous human-in-the-loop iterations by enabling the parallel comparison of distinct decarbonization pathways. This evolution reflects real-world policy workflows and aims to support the creation of transparent, adaptable, and evidence-based decarbonization strategies.

## 2. Related work

Our research occupies the intersection of geovisualization and energy system modeling. We distinguish our approach by merging a composable "plan-building" workflow with a highly-interactive, multi-scale visual interface. Moving beyond black-box optimization, our framework integrates a component-based planning paradigm with multivariate glyph visualizations capable of presenting high-dimensional data within a single spatial context.

### 2.1. Decision support in climate planning

Energy System Analysis (ESA) and Modeling (ESM) tools are critical for charting decarbonization trajectories. These systems aim to assist decision-makers by applying computational models to determine prioritization strategies for building stocks [9–12]. However, the inner workings of these models often remain obscured, prompting demands for more transparent and interpretable modeling processes to determine transition pathways [3].

Devising an energy strategy is intrinsically a multi-objective optimization (MOO) dilemma, requiring planners to weigh financial constraints against environmental targets and social fairness. Recent strides in visual analytics address MOO by enabling the interactive exploration of the Pareto front [13–15]. While these systems empower users to identify desirable outcomes from a pre-computed solution space, our work shifts the focus toward the authorship process itself.

Rather than navigating a predetermined mathematical frontier, our component-based paradigm allows planners to manually assemble, layer, and simulate modular interventions. This "*Filter-Prioritize-Upgrade*" methodology transforms the user from a passive observer of model outputs into an active architect of the decarbonization plan. By tightly coupling data-driven cohort selection with immediate visual feedback, we offer a transparent and cognitively manageable alternative to opaque, deterministic solvers.

### 2.2. Geospatial interfaces for planning

Geographic Information Systems (GIS) are foundational to urban and energy planning, typically employing a map-layer metaphor where thematic data are superimposed onto a basemap. This approach provides the technical means for decision-makers to grasp connections among geographic, societal, and cultural elements [16–18]. However, the standard 'layer-stacking' paradigm becomes cumbersome when dealing with the high-dimensional data inherent to energy modeling [19]. As users add layers representing metrics such as energy use, EPC rating, and social factors, they face significant cognitive overhead in mentally fusing disjointed information to understand interrelationships [6]. Furthermore, representing temporal change often requires cumbersome time-sliders or animations that hinder direct comparison across multiple time points. Interactive visual analytics has nonetheless been applied successfully to spatial planning and decision-making tasks—for instance, selecting billboard locations from large-scale trajectory data [20] and analyzing urban performance [21]. Our work brings a comparable interactive, decision-support approach to local energy planning.

To address this 'layer-overload', researchers in geovisualization and visual analytics have explored integrated representations, most notably glyphs: small, data-driven graphical symbols that encode multiple variables into a single visual entity [22]. Glyph-maps have been effectively utilized to explore spatio-temporal patterns across diverse domains [23,24], including the visualization of multi-criteria priorities for decarbonization [25] and building stock attributes [8]. Our work builds on this tradition. While existing studies have used glyphs for energy data — such as rose bar glyphs for geothermal benefits [26] — we build on the mirrored streamgraph glyph introduced in our earlier work [8], which contrasts costs against carbon savings over a simulated timeline. This facilitates an at-a-glance understanding of complex, time-series decarbonization plans without the context-switching required by traditional sliders. Here, these glyphs serve as a shared visual vocabulary: applied consistently across geographic scales and across synchronized scenario comparisons, they allow competing pathways to be read side by side in the same visual terms.

Finally, we incorporate abstraction techniques drawn from research on value-by-area cartograms and squarified grids [23,24,27]. By abstracting geography into a regular grid, we transform the map into a functional canvas for comparison. This approach ensures that geographically small but densely populated urban zones are not visually overwhelmed by larger rural areas, giving each spatial unit equal prominence [28,29]. This mitigation of area-based bias is critical for ensuring a "fair and just" assessment of policy interventions across diverse geographic scales.

### 2.3. Visual analytics in energy domains

The broader visualization community has increasingly directed its focus toward the energy sector, transitioning from operational monitoring to long-term strategic support. Early research in this space primarily targeted real-time building management, such as integrating sensor data with occupant feedback to identify localized inefficiencies [30]. However, recent studies emphasize the growing need for sophisticated Scenario Modeling to navigate decarbonization pathways [10].

A significant challenge in this domain is navigating the 'solution space' created by Multi-Objective Optimization (MOO). Recent frameworks such as ParetoLens [31] and the visual analytics platform for solar and wind energy modeling [32] address this by allowing analysts to inspect the distribution of optimal solutions across decision and objective spaces. Despite these advancements, many existing systems focus on the exploration of pre-computed optimal frontiers, where users select from a set of results already generated by an algorithm [13,14]. This approach can be restrictive for policy-making, where the "best" solution often emerges from an iterative, creative process instead of a purely mathematical one [33].

Our work builds directly on Laksono et al. [8], which utilized visual analytics to integrate human-in-the-loop mechanisms into the interactive scenario-building process. Relative to this earlier proof-of-concept, the central advance of this paper is the move from authoring a single strategy to comparing several. This study extends that foundation by shifting the user's role from the generation of isolated strategies toward a more robust multi-scenario authorship and comparison framework. Scenarios are now first-class objects that can be branched, versioned, and examined side by side in synchronized map views at any geographic scale—and it is this comparison workflow that, in practice, drives iterative composition and the eventual selection of a defensible plan. By incorporating a dedicated provenance tracking mechanism, each user-defined scenario is preserved, allowing for the auditing of the decision-making history. Furthermore, we employ comparative workflows, moving beyond single-scenario analysis to accommodate side-by-side evaluations of distinct decarbonization pathways. The improved dashboard extends prior Filter-Prioritize-Transform paradigm into a highly iterative workflow for constructing, revising, comparing, and justifying multiple decarbonization pathways.

Several further additions support this iterative workflow: a Field Calculator through which planners define their own metrics, eligibility flags, and prioritization keys; more flexible cohort targeting through lasso and custom boundary selection; and a decision-episode provenance model that records how each pathway was constructed, not just its final state. By allowing planners to manually assemble, simulate, and compare modular interventions in real time, the system provides a transparent and flexible alternative to static scenario evaluation, in which scenarios are typically defined and assessed as fixed configurations.

## 3. System architecture and methodology

### 3.1. Design methodology and requirements

This work was developed as a problem-driven design study [34], grounded in a multi-year embedded partnership with a UK-based energy consultancy. Over a three-year period, the first author worked alongside the consultancy's analysts several days per week, observing how Local Area Energy Plans are produced and where existing tools create friction. This embedded co-design was complemented by semi-structured interviews with seven practitioners spanning energy consultancies and local government, together with observation of network-operator planning workshops; the interviews were examined thematically to characterize the domain's data, analytical tasks, and decision requirements.

Given that these requirements emerged through continuous dialogue rather than a single elicitation exercise, we make their derivation explicit. Interview material and the running observation record were analyzed thematically, yielding five recurring findings about how local energy plans are produced and where current tools obstruct that work (F1–F5, Table 1). We then examined each finding to identify the action it implied at a particular point in the planning cycle: constraining the candidate set, deciding priorities, applying an intervention, evaluating its consequences, or preserving and comparing alternatives. These actions were then consolidated into the five tasks described below:

- **T1: Targeted Cohort Identification**: Sifting through granular property records to isolate specific demographics — such as inefficient housing in low-income areas — using multivariate glyphs and interactive data tables.
- **T2: Modular Strategy Composition**: Assembling comprehensive policy frameworks from focused, discrete intervention "building blocks" that represent specific decarbonization actions.
- **T3: Spatio-Temporal Outcome Evaluation**: Reviewing projected financial and environmental impacts across diverse geographic boundaries and long-term planning horizons.
- **T4: Iterative Hypothesis Testing**: Continually refining intervention parameters to discover the most effective and equitable solutions through a rapid feedback loop.
- **T5: Multi-Scenario Comparison**: Juxtaposing distinct strategies side-by-side to evaluate the relative merits of different scenario compositions and policy trade-offs.

To support these high-level tasks, we developed an integrated visual analytics system that tightly couples a responsive front-end interface with a robust scenario-building engine. This architecture ensures that user-defined interventions are immediately processed by the underlying simulation model, providing the real-time visual feedback necessary for complex decision-making. By bridging the gap between raw data exploration and predictive modeling, the system transforms static building data into an actionable, "what-if" planning canvas.

### 3.2. The visual analytics dashboard

The interface (Fig. 1) comprises five tightly coupled, coordinated multiple views (CMV) components designed to facilitate a fluid, iterative planning workflow:

The **Scenarios** Panel (Fig. 1A) is the entry point of the dashboard. It serves as the user session and scenario management interface, allowing users to create new scenarios, or branch and version existing ones, and to load saved scenarios. The provenance feature allows planners to experiment with, and later compare, multiple scenarios (Fig. 7). Committed authoring actions are recorded in a Decision Trail — an inspectable ledger of what changed, when, and in which scenario version — so that a finalized plan carries a traceable history of the decisions that produced it (detailed in Section 3.4).

The **Interventions** panel (Fig. 1B) facilitates intervention composition within each scenario. Planners can allocate different budgets, durations, or technological interventions — building insulation, rooftop photovoltaics, air- and ground-source heat pumps, hydrogen boilers, district heating or a mixed retrofit, drawn from the catalogue of Table 2 — to focused subsets of buildings. Existing interventions can also be modified (changing budget, duration, or start year), enabling the planner to iteratively refine intervention compositions and instantly observe the impact.

The **Strategy Explorer** (Fig. 1C) serves multiple purposes as a spatial filtering tool (T1), a canvas for multi-scale outcome evaluation (T3), and a facilitator for comparison (T5). For spatial targeting, the map provides polygon, rectangle, and lasso selection, as well as upload of an external boundary (e.g., a ward or a candidate heat-network zone supplied as GeoJSON) to delineate a custom region; the enclosed buildings form a cohort that can be reused as a named profile. The map employs multivariate glyphs to visualize high-dimensional data:

- *Decarbonization Potential*: Each glyph acts as a compact information graphic where visual channels (e.g., spoke length and color) represent key performance indicators such as technology adoption rates and cost. A hover-projection feature allows users to overlay a selected glyph's profile onto all other geographic units, enabling direct, localized comparison (Fig. 2).
- *Decarbonization Timeline*: In this panel (Fig. 1E), the map visualizes temporal trajectories using mirrored streamgraphs that contrast projected costs against carbon savings over several years.
- *Multi-Scale Abstraction*: To ensure analytical clarity across geographic hierarchies — ranging from individual buildings and gridded units to LSOA (Lower Layer Super Output Area), MSOA (Middle Layer Super Output Area), and Ward levels — the system can render data as squarified grids (Fig. 3) while animating the transition. These cartograms prioritize visual equity and legibility over strict spatial fidelity.

The **Table View** (Fig. 1D) provides a spreadsheet-style interface with histograms for the granular sorting and selection of raw building data. Selections can be made via these histograms or through direct row selection. It is bidirectionally linked with the Strategy Explorer to support multi-stage filtering, enabling planners to transition seamlessly between statistical and spatial selections to isolate specific property cohorts (T1). A built-in Field Calculator (Fig. 4) lets planners derive their own metrics and eligibility flags from the raw attributes through a validated expression language with live preview—for example, a cost-effectiveness ratio relating a building's capital cost to its baseline annual energy spend, or a composite fuel-poverty targeting flag. Computed fields behave as ordinary columns, available for filtering, map encoding, and, as described below, intervention prioritization, so planners can encode domain knowledge without leaving the dashboard.

The **Scenario Composition View** (Fig. 1E) serves as the primary authoring canvas. In coordination with the Intervention Panel (Fig. 1B), planners can configure and specify discrete "plan components". By layering these modular blocks, users can construct complex, multi-objective strategies (T2). This view also functions as a direct interface to the simulation engine: adjusting the composition (e.g., reordering blocks or shifting timelines) triggers an immediate model recalculation, with results reflected instantly on the map. To assist in real-time evaluation, each intervention block embeds miniature bar charts that encode annual deployment outputs, while the overall scenario impact is summarized within the 'Workspace' tab in the panel.

**Table 1**
Derivation of the analytical tasks. Each recurring finding from the fieldwork is paired with the design implication drawn from it and the task it motivates. *Source* records the engagement stream in which the finding was established: I—semi-structured practitioner interviews, analyzed thematically; E—embedded observation of planning practice over the partnership.

| | Finding → design implication | Source | Task |
|---|---|---|---|
| F1 | Plans combine building fabric, network headroom, deprivation and fuel poverty, and political priority; the modeling stage is experienced as a black box. → *Selection and evaluation must account for multiple attributes within a single linked view, keeping the underlying criteria visible rather than reducing them to a single score.* | I | T1, T3 |
| F2 | Plans are commissioned as static deliverables, yet budgets, datasets and policies keep changing. → *A strategy must be assembled from discrete, independently parameterized components that can be revised and re-run.* | I | T2, T4 |
| F3 | Reasoning moves between building stock, statistical geographies, political boundaries and network catchments, which do not align. → *Outcomes must be readable at several nested scales under one consistent encoding.* | I, E | T1, T3 |
| F4 | Formal optimization is regarded with skepticism: slow to configure and run, opaque in its internal logic, and aggregated too coarsely for neighborhood-level feasibility. → *Replace monolithic optimization with transparent, steerable blocks fast enough to sustain an interactive feedback loop.* | I, E | T2, T4 |
| F5 | Plans emerge through multi-stage negotiation among several parties and must remain defensible afterwards. → *Alternatives must be preservable, re-loadable and directly comparable, with the reasoning behind a chosen plan inspectable after the fact.* | I, E | T5 |

**Table 2**
The intervention catalogue. Eligibility is evaluated per building against the named suitability attribute; capital cost is drawn from the corresponding cost field of the building stock. The carbon factor is applied to each building's baseline; the demand factor is a per-deployment constant, applied identically to every building. Both are indicative rather than calibrated engineering estimates (see Section 5.4). A positive demand factor denotes added electrical load, a negative one a reduction.

| Technology | Eligibility | Capital cost | Carbon | Demand |
|---|---|---|---|---|
| Building insulation | Any fabric measure suitable[a] | Sum of suitable measures | 0.20 | −0.6 |
| Rooftop PV | `rooftop_pv` | `rooftop_pv_total_cost` | 0.15 | −0.8 |
| Air-source heat pump | `ashp` | `ashp_total_cost` | 0.28 | +1.2 |
| Ground-source heat pump | `gshp` | `gshp_total_cost` | 0.32 | +1.0 |
| Hydrogen boiler | `hydrogen_boiler` | `hydrogen_boiler_total_cost`[b] | 0.12 | +0.5 |
| District heating[b] | `district_heating` | `district_heating_total_cost`[b] | 0.24 | −0.4 |
| Mixed retrofit | (ASHP ∧ PV) ∨ district | ASHP + PV cost[c] | 0.36 | +0.3 |

[a] Cavity-wall, external-wall, roof, or under-floor insulation.
[b] Placeholders in the present dataset: a flat £10,000 and £3,953 respectively, identical for every building.
[c] Falling back to the stock's aggregate cost field where both components are zero.

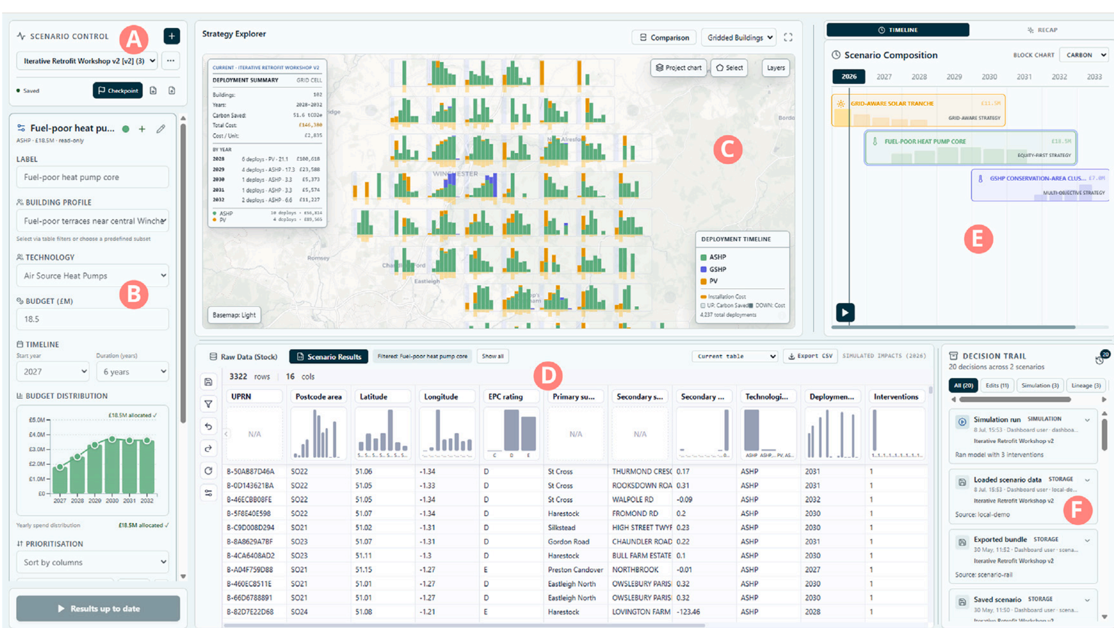


**Fig. 1.** The dashboard interface: Planners create or load saved scenario configurations via *Scenarios* panel (A); *Interventions* panel (B) allows composition of a new decarbonization intervention or modify existing one; The *Strategy Explorer* (C) serves as the main panel which visualize both building stock and deployment output; The *Interactive Table* (D) can be used to filter buildings which are targeted for interventions; Interventions within a scenario can be modified via *Scenario Composition* (E). The *Decision Trail* records how pathways were authored, constructed and modified (F).

The **Decision Trail** (Fig. 1F) serves as the plan's inspectable audit record. Working in concert with other dashboard panels, it captures each committed authoring decision — scenario creation, branching and versioning, intervention edits, computed-field definitions, simulation runs, and import/export — as a timestamped, lineage-aware entry instead of a low-level event log (T4). By design it logs discrete decision

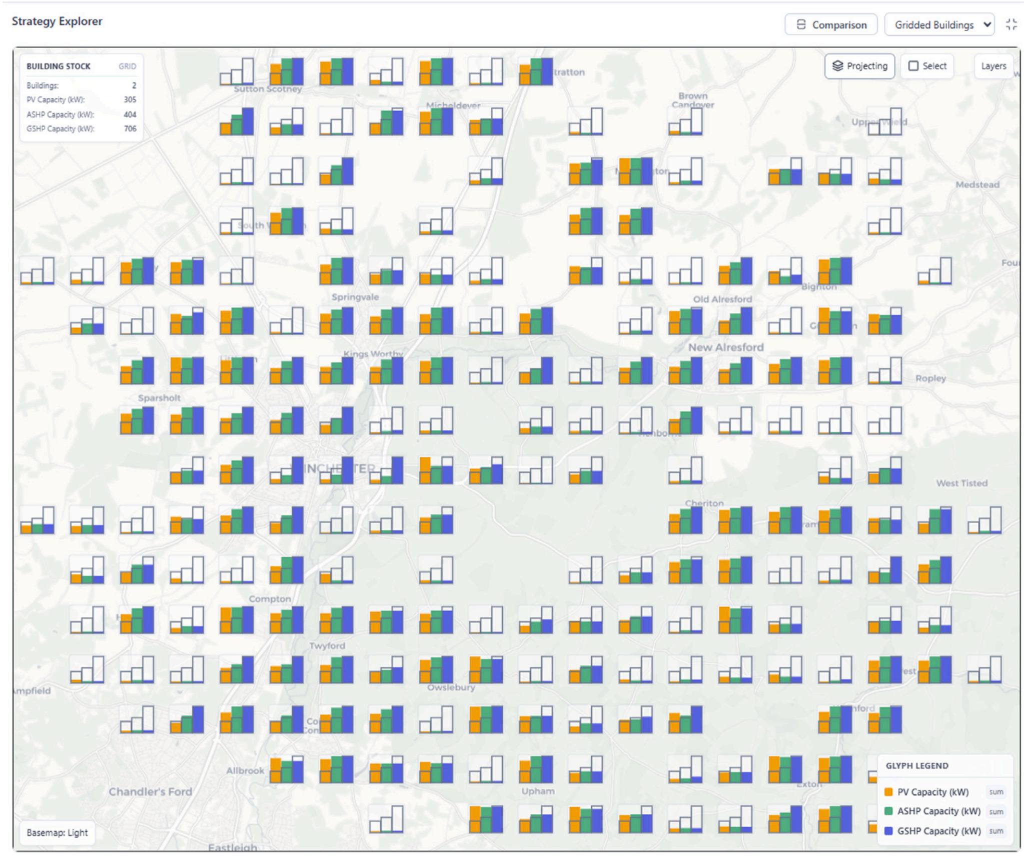


**Fig. 2.** Comparing renewable potentials across gridded geography by projecting glyphs on hovered mouse interaction.

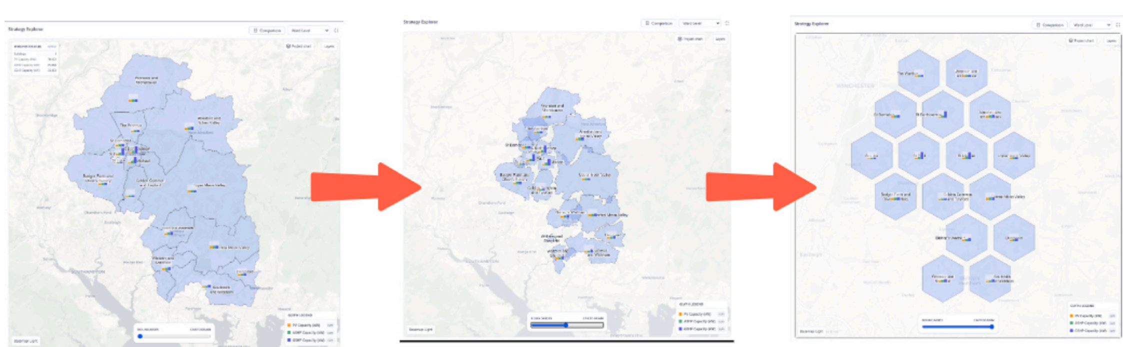

**Fig. 3.** Adjusting the morphing slider from the original geographic boundaries (left), through an intermediate morphing state (center), to a cartogram layout (right), reducing glyph crowding in small or spatially compressed areas and making local patterns easier to compare.

*episodes*, so the trail reads as a coherent narrative of how a scenario was constructed and refined over time. This view also functions as a diffing interface onto scenario state: expanding any entry reveals the field-level before/after evidence and an automatic diff summary for that decision, alongside its source context and workflow stage. To support recall and defensibility, the trail threads these entries along the scenario's branching lineage, allowing a planner to reconstruct not just the finalized plan but the sequence of trade-offs and rejected alternatives that produced it.

### *3.3. The scenario modeler simulation engine*

The simulation framework is architected around a flexible, data-agnostic *Intervention* primitive designed to model longitudinal decarbonization trajectories. Central to this is the Scenario Modeler, which simulates decarbonization pathways by iteratively applying structural and technological changes to a bespoke 'facet'. We use *facet* in a specific sense: a minimal row-access adapter over a tabular dataset, requiring only a row count, an indexed row accessor, and a list of column names. Any dataset satisfying this interface can be simulated, with domain-specific facets supplied programmatically at the adapter layer. This makes the domain-portability argument below a concrete property of the architecture rather than an aspirational claim. This developer-level extension is distinct from the in-session affordance available to planners, who add derived columns to the already-bound facet through the Field Calculator; those computed columns are materialized back onto it and thereafter behave as ordinary attributes. The deployed dashboard binds a single facet: the Winchester residential building stock. Its design is governed by three primary functional stages, supported by

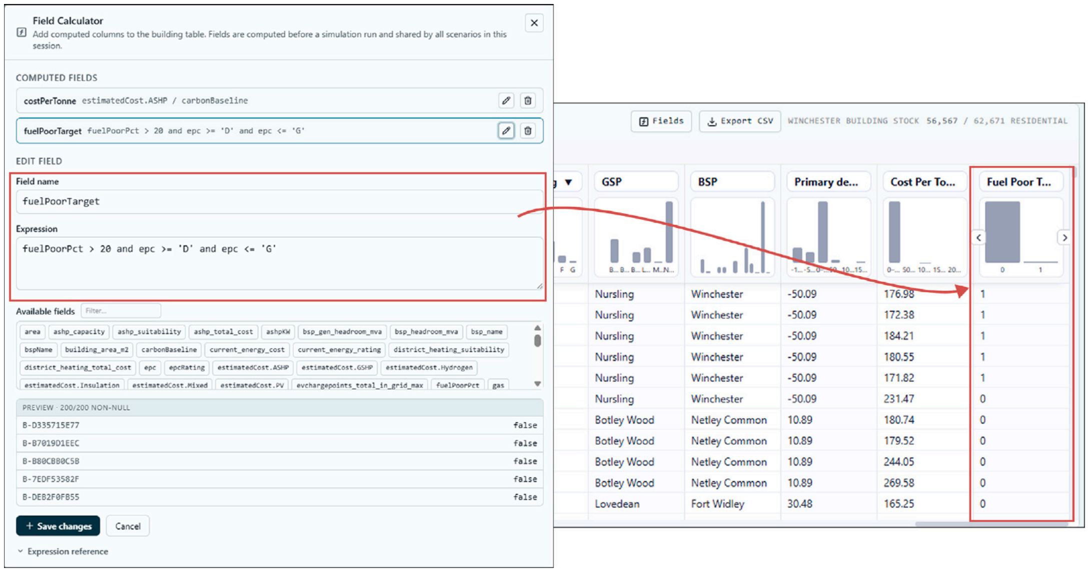


**Fig. 4.** The Field Calculator, in which planners author bespoke metrics and eligibility flags through a validated expression language with live preview. The `fuelPoorTarget` column is calculated as a Boolean metric of EPC rating and Fuel poor percentage fields. The computed `Fuel Poor Target` field become first-class column usable for filtering, map encoding, and intervention prioritization.

an extensible plugin architecture, defining every 'intervention' through three core functions executing iteratively per simulation year:

1. **Filter (Eligibility)**: Determines intervention eligibility. It answers the question: "Which buildings are candidates for this intervention right now?" This dynamic predicate isolates eligible structures based on multifaceted criteria. At the programmatic level, it evaluates a building's array of currently deployed technologies against its `suitableTech` constraints, verifies membership within specified Unique Property Reference Number (UPRN) profiles, and enforces programmatic thresholds such as `targetEPC` metrics (e.g., proactively targeting dwellings with EPC grades below 'C'). Eligibility for network-scoped technologies is treated differently because network reach is a planning decision rather than a building attribute. When a planner defines the intervention using a saved cohort or uploaded boundary, the resulting zone determines where the network can be deployed, replacing the usual per-building suitability attribute.
2. **Prioritize (Sequencing)**: Determines spatial and operational order answering: "In what order should we process the eligible buildings?" This logic is critical when temporal resources, operational capacities, or fiscal budgets are strictly constrained, and competing priorities needs to be determined. The framework employs a multi-variate sorting algorithm that calculates a comprehensive rank by evaluating three weighted parameters, exposed in the interface as `carbonSavingPotential` (Carbon), `fuelPovertyScore` (Equity), and `gridCapacityEfficiency` (Grid). These rank respectively on the building's baseline annual energy cost — used as a proxy for carbon-saving potential — a composite fuel-poverty and deprivation score, and the available headroom at its supplying substation. The parameters use different units, so they cannot be added together directly. Each is therefore first converted to a value between 0 and 1 across the eligible stock before weighting. Rescaling is bounded due to a small number of heavily constrained substations with large negative headroom present in the dataset for a placeholder of empty rows, and would otherwise compress almost the whole stock into a narrow band of the scale. Values are mapped between the 1st and 99th percentiles of the stock and clamped to the unit interval. The rank of building $b$ is then the weighted linear combination
$$S(b) = \sum_k w_k\, \hat{x}_k(b), \qquad \hat{x}_k(b) = \mathrm{clamp}\left(\frac{x_k(b) - q_k^{0.01}}{q_k^{0.99} - q_k^{0.01}},\, 0,\, 1\right), \quad (1)$$
where $x_k$ is the raw value of parameter $k$, $q_k^p$ its $p$-quantile over the stock, and $w_k$ the weight set by the corresponding slider. Buildings are deployed in descending order of $S$, so adjusting a weight immediately re-orders deployment.
A sort-by-columns mode is also possible, ordering eligible buildings by one or more chosen columns — including fields authored in the Field Calculator — with later columns breaking ties. This parameterized weighting empowers planners to seamlessly shift deployment priority from maximum carbon mitigation to maximum poverty alleviation in real-time.
3. **Upgrade (Impact)**: Determines actionable impact, answering the question: "What change is applied, and what are the quantifiable results?". The framework orchestrates the physical modification (e.g., Air Source Heat Pump installation) through a composable sequence of focused plugins — budget guarding, intervention logging, carbon impact, and demand impact — invoked in turn for each eligible building. A shared simulation context carries mutable run state, including the active intervention and remaining annual budget. The budget guard runs first: if a building's technology cost exceeds the remaining allocation the upgrade yields no metric delta; otherwise it deducts the cost and lets the subsequent plugins record deployment, carbon savings, and demand change. Because the plugins update this shared state in a fixed order, deterministic prioritization produces reproducible outcomes, and new effects (e.g., grid-capacity checks) can be added as further plugins without altering the core Filter–Prioritize–Upgrade loop. The resulting per-building deltas are aggregated into yearly and intervention summaries, simulated result facets, and each building's intervention history.

The loop uses a fixed catalogue of deployable technologies (Table 2), which is defined when the dashboard is deployed and cannot be edited by planners. Planners instead configure modules for these technologies by specifying their building profile, budget, timeline and prioritization rule, as described in Section 4. Each catalogue entry specifies the building attribute used to determine suitability, the cost used by the budget constraint, and the coefficients used to calculate changes in carbon and energy demand. Insulation combines four fabric treatments: a building is eligible if any treatment is suitable, and its cost is the sum of the applicable treatments. Mixed retrofit similarly combines measures, requiring either suitability for both a heat pump and rooftop PV or existing district-heating suitability. New technologies can therefore be added by adding a catalogue entry, without changing the engine or the Filter–Prioritize–Upgrade loop. The engine only requires the relevant suitability and cost fields.

The coefficients in Table 2 are illustrative: They were chosen only to reflect the relative relationships reported in the literature: ground-source heat pumps generally have higher seasonal performance than air-source heat pumps [35,36]; electrified heating can increase peak electricity demand, whereas fabric measures and on-site generation can reduce it [37,38]; and hydrogen has the weakest and most uncertain abatement potential among the options considered [39]. The coefficients therefore represent the intended ordering between technologies, not measured performance. Their absolute values have no independent meaning (Section 5.4).

Given the data-agnostic nature of the *Intervention* primitive and its modular, functional components — Filter, Prioritize, and Upgrade — the underlying simulation engine is not intrinsically bound to housing decarbonization. The separation of domain-specific logic into discrete plugins and extensible profiles ensures that the core architecture remains highly generic. Consequently, this abstract modeling paradigm can be readily adapted to simulate longitudinal interventions across a variety of spatiotemporal domains, such as urban infrastructure retrofitting, epidemiological modeling, or public transport network optimization, provided the underlying dataset can be mapped to an equivalent facet.

### 3.4. Technical implementation

The application relies on a client-side architecture using React and MapLibre GL JS. This structure supports the dashboard's user interface and interactive elements directly within the browser, avoiding the need for continuous server communication during interaction.

Beyond persisting scenario state, the system records *decision-episode provenance*: a timestamped, lineage-aware log of committed authoring decisions—scenario creation, branching, versioning, intervention edits, computed-field creation, simulation runs, and import/export. Events are annotated with source context and workflow stage, while edit events additionally carry field-level before/after evidence and automatic diff summaries. In the terms of Ragan et al. [40], this couples *interaction* and *rationale* provenance to serve recall, presentation, and meta-analysis, while its branching and versioning lineage is analogous to VisTrails-style workflow provenance [41]. By design the trail captures committed decision episodes, keeping it a meaningful audit record rather than a low-level event log. Work persists in browser storage and in a self-contained export bundle — sessions, scenarios, cohorts, computed fields, workflow state, and provenance events — that can be re-imported to share a plan between officers or submit it as evidence; an optional remote backend is provided to support future multi-user collaboration.

A core component of the data processing pipeline is the use of DuckDB-WASM, an in-browser columnar database engine. DuckDB-WASM loads data from compressed `.parquet` files and uses WebAssembly workers to run SQL queries locally. This setup allows operations such as spatial bounding box filtering (`getGridAggregates`) and boundary-level aggregations (`getBoundaryAggregates`) to execute locally, reducing latency during data filtering and processing events. At the authoring layer, the Field Calculator extends this client-side pipeline by allowing planners to create derived attributes from existing building fields; once saved, these computed columns are materialized back into the working building facet so they can drive histograms, map encodings, cohort filters, CSV exports, and intervention prioritization without requiring external tools.

The map interface is implemented with MapLibre GL JS. To handle the rendering of dense structural data points, it uses a custom gridded-glyphmap library, `screengrid`, which combines D3 and the HTML Canvas API to aggregate buildings into screen-space grid cells and draw multivariate glyphs over the map base layer. Additionally, the dashboard incorporates JavaScript libraries such as Turf.js to handle visual transitions between different aggregation boundaries, such as transitioning from Lower Layer Super Output Areas (LSOA) to its cartogram representation. This enables the smooth morphing of geospatial polygons (Fig. 3), allowing the planner to examine data at various spatial levels while maintaining a consistent visual context.

## 4. Application scenario

To illustrate the workflow end-to-end, we walk through a planner constructing and benchmarking two contrasting decarbonization pathways for Winchester City. The walkthrough is reproducible within the dashboard: the scenarios, saved cohorts, computed fields, and provenance record are provided as a preloaded demo workspace, allowing readers to replay each step using the same data. Fig. 5 provides the structure for this section, with each of the five steps presented as a stage showing its tasks, interface panels, and resulting artefact. Panel letters throughout refer to Fig. 1(A–F).

**Step 1 — Target a cohort (T1).** In the *Scenarios* panel the planner creates a new scenario, *"Cost-Effective Renewables",* whose goal is to maximize emissions reduction in the most fuel-deprived, inefficient stock. To isolate this cohort they cross-filter in the *Table View*, brushing the EPC and fuel-poverty histograms, and refine the selection spatially on the *Strategy Explorer* using lasso selection or an uploaded ward boundary. Where a custom targeting rule is needed, they write one in the Field Calculator (Fig. 4) — here a Boolean flag, `fuelPoorTarget`, combining a fuel-poverty share above 20% with EPC bands D–G — and save the resulting selection as a reusable named building profile (*"Fuel-poor EPC D–G cohort"*), shown as stage 1 of Fig. 5.

**Step 2 — Compose interventions (T2).** In the *Interventions* panel the planner assembles the pathway from discrete building blocks: a *high-yield rooftop PV* cohort (£12.0M from 2026) targeting the stock with the strongest generation potential, followed by a *mixed-retrofit follow-on* (£8.5M from 2029) directed at the highest-carbon EPC D homes. Both draw on the intervention catalogue of Table 2 — the *rooftop PV* and *mixed retrofit* rows respectively — and each carries its own budget, duration, start year, and prioritization weights. Each configured intervention appears as a block on the *Scenario Composition* timeline, which the planner layers to build the multi-year plan (Fig. 5, stage 2).

**Step 3 — Simulate and iterate (T3, T4).** Running the model recalculates deployment across the full building stock, and the map re-encodes to the deployment timeline, showing capital cost against carbon saved per cell as mirrored streamgraphs. The planner notes a deployment pattern scattered across the rural and suburban periphery of Winchester and, finding the total cost exceeds their constraints, enters a rapid feedback loop — trimming the follow-on retrofit budget (£9.5M to £8.5M) and shortening its window from 2029–2033 to 2029–2032, then reordering the priority sequence of intervention blocks — until the efficiency metrics stabilize. Each committed edit is recorded with before/after evidence in the *Decision Trail*. Reviewing the first PV-only pass the next morning, they checkpoint the refined plan as version v2, which concentrates the solar spend into a shorter, larger PV tranche (£12.0M to £14.5M over 2026–2030), retires the mixed-retrofit follow-on altogether, and pulls heat-pump delivery earlier with an *ASHP infill* block (£16.0M over 2028–2033; the *air-source heat pump* row of Table 2). Its deployment order is a sort-by-columns rule ranking on a `costEffectivenessRatio` computed field authored in the Field Calculator — capital cost per pound of each building's baseline annual energy spend, best ratio first — with fuel-poverty share as the tie-break (Fig. 5, stage 3).

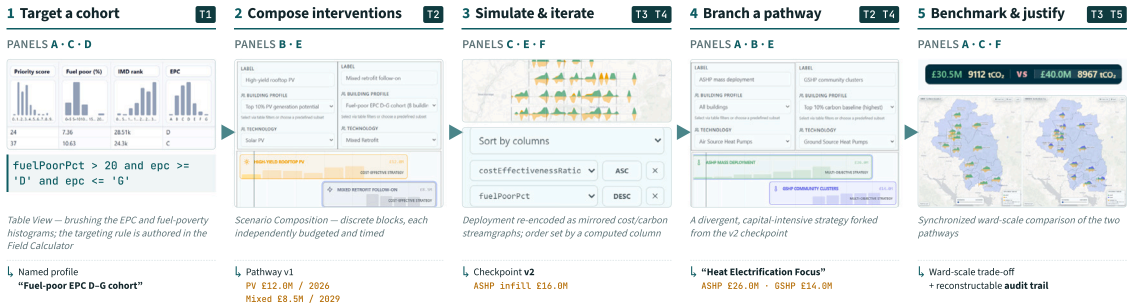


**Fig. 5.** The five-step planning workflow of Section 4, shown as a single spine. Each stage names its tasks (T1–T5), the dashboard panels it uses — keyed to the letters of Fig. 1(A–F) — a crop of the interface at that moment, and the artefact it hands to the next stage: a named cohort, an intervention block, a simulated checkpoint, a branched pathway, and a ward-scale benchmark. All values are taken from the scenarios that ship preloaded with the public dashboard: stage 1 from the shared building stock and its computed fields; stages 2 and 3 from *Cost-Effective Renewables* at versions v1 and v2; stage 4 from *Heat Electrification Focus*, branched from the v2 checkpoint; and stage 5 from the two compared side by side.

**Step 4 — Author a divergent pathway (T2, T4).** To explore an alternative policy extreme, the planner returns to the *Scenarios* panel and branches a second scenario from the v2 checkpoint, *"Heat Electrification Focus"*—a more capital-intensive strategy simulated with an expanded £40.0M budget that, via the *Interventions* panel, heavily prioritizes air-source heat pumps (a £26.0M mass deployment) complemented by £14.0M of ground-source community clusters. They again iterate on the deployment sequence in the *Scenario Composition* view — retiming the GSHP clusters from 2028–2032 to 2029–2033 and raising their budget from £12.0M after the first model pass — while reviewing on the *Strategy Explorer* map how the spatial footprint of heat-pump deployment evolves relative to the renewables-led pathway (Fig. 5, stage 4). This is the second fork of the week: a *Fuel Poverty Alternative* had already been branched from v1 the previous day, so the lineage now holds four related pathways.

**Step 5 — Benchmark and justify (T3, T5).** Finally, from the *Scenarios* panel the planner loads the *"Cost-Effective Renewables v2"* checkpoint alongside the active *"Heat Electrification Focus"* and activates side-by-side comparison mode (Fig. 7). Exploring the synchronized *Strategy Explorer* maps across hierarchical scales — switching to the aggregated Ward view via the morphing control (Fig. 3) — they isolate geographical friction points. At the headline level the two pathways are closely matched on abatement — 8967 $tCO_2$ for £40.0M under electrification against 9112 $tCO_2$ for £30.5M under the renewables-led plan — so the more informative contrast is spatial: the heat-led pathway concentrates costly deep interventions in fewer wards, while the renewables-led one spreads cheaper measures across more of the stock, and individual wards diverge sharply in the capital they demand and the offset they return (Fig. 5, stage 5). Because every step above is captured in the *Decision Trail*, the resulting geographically targeted recommendation is backed by a transparent, reconstructable audit trail. Expanding the trail into the *Decision History* view (Fig. 6) makes this audit explicitly temporal: decisions are grouped chronologically into a dated ledger, while a *pathways-over-time* overview plots each decision along a per-scenario lane—with branch connectors marking the moment each alternative diverged, including the fork of *Heat Electrification Focus* from the v2 checkpoint in Step 4. Selecting a lane filters the ledger to a single pathway, so the planner — or a colleague auditing the plan — can reconstruct not only *what* was decided, but *when*, and under which alternative.

## 5. Discussion

As illustrated in the application scenario, while optimization algorithms and deterministic energy models provide a mathematical foundation for decarbonization, they often lack the nuance required for public policy implementation. The visual analytics approach developed in this research reconceptualizes local authority planning by positioning the human planner — rather than an opaque algorithm — at the center of the decision-making process. By supporting tasks T1 through T5, the system highlights several critical shifts in how local authority energy strategies can be formulated and evaluated.

### 5.1. Bridging the gap between optimization and policy (T1, T2)

A recurring challenge in municipal planning is the cognitive friction between data science outputs and socio-political reality. "Optimal" systemic solutions are frequently difficult to justify to local communities or adapt to shifting political mandates. By transitioning toward a modular authoring environment (T1, T2), the framework reduces reliance on black-box optimization. Instead, it enables a 'human-in-the-loop' paradigm where planners act as active authors of policy. This empowers domain experts to inject localized, qualitative knowledge into the models — such as deliberately targeting specific social housing cohorts — ensuring that the resulting decarbonization strategies reflect political intent and, as our partners emphasized, are intended to be easier to defend to stakeholders.

### 5.2. Mitigating spatial biases in resource allocation (T3, T5)

Traditional Geographic Information Systems (GIS) used in energy planning commonly rely on visualization technique such as choropleth maps, which inadvertently introduce spatial biases: geographically large but sparsely populated rural areas dominate the visual field, while densely populated urban wards recede. For local planning, this visual distortion can subconsciously skew resource allocation and perceived equity.

By employing squarified grids and multivariate glyphs for benchmarking (T3, T5), the dashboard actively counteracts these cognitive biases. Abstracting boundaries into equal-area representations (Fig. 3) prioritizes demographic and structural data over mere land mass. Furthermore, the synchronized side-by-side comparative views (Fig. 7) allow planners to evaluate the systemic impacts of shifting resources (e.g., pivoting funding from rural heat pumps to urban solar networks) without the risk of "layer fatigue". This approach promotes more equitable spatial reasoning, directly aiding planners in balancing trade-offs between dense and sparse environments.

### 5.3. Institutional accountability and strategic provenance (T4)

In public sector governance, the process of arriving at a decision is often as heavily scrutinized as the final decision itself. Interactive analytics provides the agility necessary for iterative hypothesis testing (T4), but rapid exploration is only valuable if the reasoning is captured.

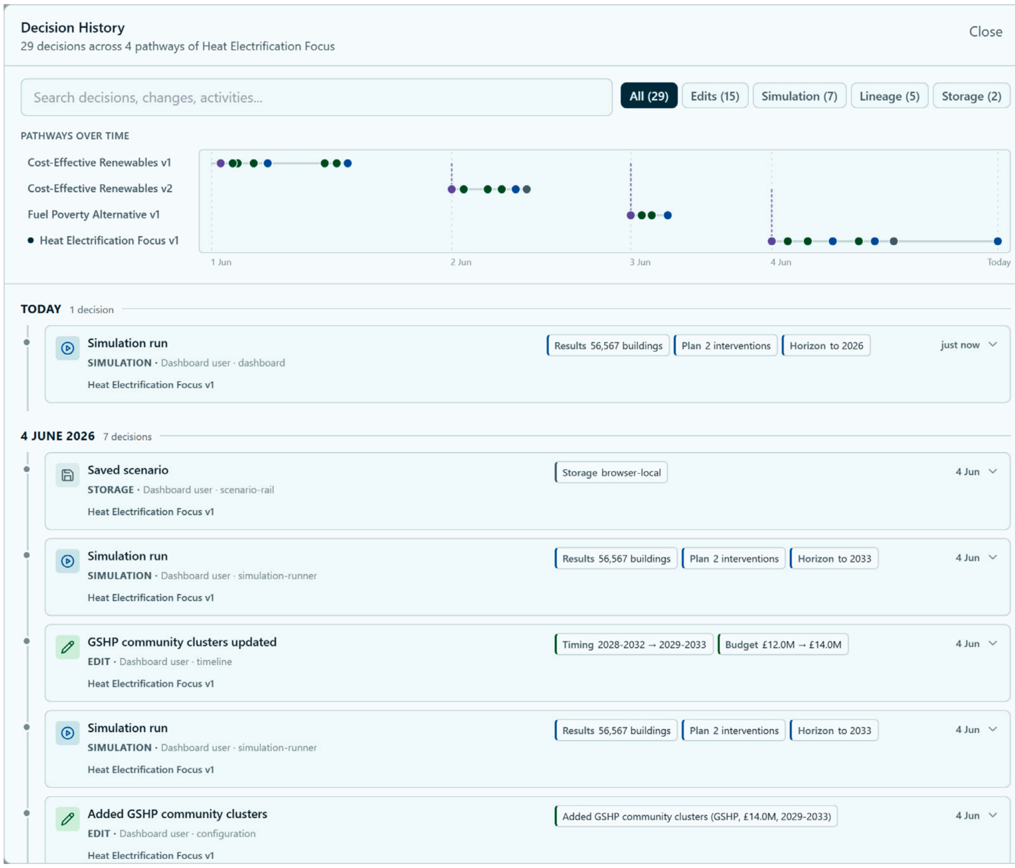


**Fig. 6.** The expanded Decision History view for the demo session's scenario lineage. The *pathways-over-time* overview (top) plots each recorded decision on a per-scenario lane — color-coded by type (edit, simulation, lineage, storage) — with dashed connectors marking where the v2 version, the *Fuel Poverty Alternative*, and the *Heat Electrification Focus* branch diverged from their parents. Idle gaps between working sessions are visually compressed so each day's burst of activity remains legible; dashed gridlines mark day boundaries. Below, the ledger groups decision episodes under dated headers, newest first, with each card carrying its relative timestamp, provenance, and field-level change evidence.

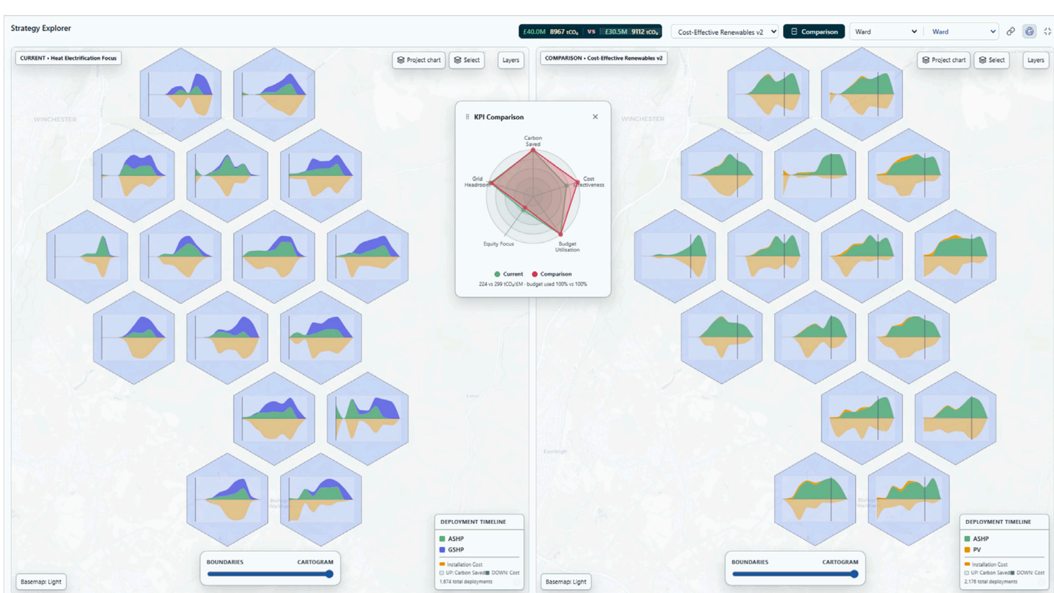


**Fig. 7.** Ward-level cartogram comparison of *Heat Electrification Focus* (left) and *Cost-Effective Renewables v2* (right). Mirrored streamgraphs show carbon savings (top) and installation costs (bottom) per ward over the scenario timeline, while the KPI radar compares the strategies on normalized axes. Both scenarios use their full budgets, but achieve different outcomes: *Heat Electrification Focus* spends £40.0M to reduce 8967 $tCO_2$ across 1674 deployments, whereas *Cost-Effective Renewables v2* reduces 9112 $tCO_2$ from £30.5M across 2176 deployments. The resulting cost-effectiveness is 224 and 299 $tCO_2$ per £M, respectively. Glyphs distinguish air- and ground-source heat pump deployments (green and purple, left) and heat pumps and PV (green and amber, right).

The integration of continuous state tracking and persistence elevates the dashboard from a purely exploratory tool into a strategic ledger. By maintaining computational provenance of how a scenario was authored, modified, and refined over time, the system establishes an audit trail. This audit trail is the decision-episode record described in Section 3.4: because each committed action is stored with its before/after evidence, a planner can reconstruct not just the final scenario but the sequence of decisions and trade-offs behind it. This is critical for institutional memory, enabling local authorities to explicitly demonstrate why alternative decarbonization pathways were rejected in favor of the finalized plan. Consequently, the tool supports both the creation of policy and the bureaucratic defense of it.

### 5.4. Validation and limitations

Consistent with the problem-driven design study through which it was developed [34], the framework was validated formatively rather than through a controlled summative experiment. Over the multi-year partnership, design decisions were shaped and corroborated by contextual expert feedback on authentic planning tasks—for example, practitioners' preference for radial glyph encodings, the utility of the mirrored streamgraph for reading the cost–carbon trade-off over time, and a consistent demand for transparent, adjustable prioritization in place of opaque optimization.

Both the initial prototype [8] and the version reported here were presented to the industrial partner in dedicated review sessions, each followed by extended discussion. Three outcomes are worth recording. **First**, the practitioners affirmed that a transparent, prioritization-led approach of this kind matches how deployment decisions are reasoned about in practice, and can support decision-making in local area energy planning. **Second**, the discussions singled out the transparency of the approach, and the ability to stack and explore interventions, as its most valuable qualities. **Third**, the partner expressed interest in adopting the dashboard and its underlying scenario-modeling library within their own platform. In design-study terms, such resonance and prospective adoption constitute meaningful evidence of utility [34,42], as practitioners' willingness to consider integration suggests that the system aligns with a recognizable need in their existing workflow.

Rigor in design studies need not rest on a single summative experiment; it can also be established through sustained, reflexive engagement with the domain [42]. Moreover, prematurely applying controlled evaluation to an evolving design can be misleading or counterproductive [43]. Accordingly, we treat the reduction of cognitive friction as a design goal motivated by our domain findings, rather than as an empirically measured outcome. A controlled study of cognitive load and decision quality with local-authority planners in an operational setting therefore remains an important direction for future work.

A further limitation concerns uncertainty. The current system presents deterministic projections of cost, carbon, and demand, and does not yet communicate the uncertainty inherent in installation costs, technology-adoption rates, or the underlying model assumptions. Because these outputs may be used to justify policy decisions, conveying such uncertainty — for example through confidence ranges on the streamgraph glyphs or sensitivity indicators on the deployment outputs — is an important direction for future work.

Relatedly, the impact model is deliberately coarse. As Table 2 shows, carbon savings are estimated by applying a fixed per-technology coefficient to each building's baseline energy use, which is derived from its current annual energy cost rather than a thermal simulation. Demand change is represented by a fixed per-technology value that is the same for every building receiving that technology. This simplification keeps the model transparent and fast enough to support the interactive feedback loop. It is sufficient for the comparative judgements the system is designed to support: which pathway deploys where, in what order, and at what relative cost. The outputs should not therefore be interpreted as calibrated estimates of absolute savings. Using them for delivery forecasting would require replacing the illustrative coefficients with measured or physics-based estimates.

A notable distinction between our framework and traditional optimization models lies in the direction of the analytical workflow. While our system empowers planners to author bespoke strategies, it does not currently support back-casting or goal-oriented optimization—such as automatically calculating the 'most optimistic' pathway for a specific target, like the maximum possible solar deployment by 2050 regardless of budget. In such 'unconstrained' scenarios, a deterministic solver might more rapidly identify the upper bounds of technical feasibility. However, our focus remains on the negotiable middle ground of policy-making, where the goal is not just to find a mathematical maximum, but to interactively explore how specific, budget-constrained interventions accumulate over time to reach localized targets

Furthermore, the decarbonization model implemented in this study is currently focused on the deployment of specific residential renewable technologies. In practice, local authority planning involves significantly more complex analyses, such as balancing grid-level electricity demand, optimizing the placement of electric vehicle charging points (EVCPs), or the strategic siting of utility-scale assets like wind turbines and ground-mounted solar arrays. While the underlying model is theoretically capable of supporting these diverse datasets, designing a unified graphical user interface that remains intuitive across all use cases presents a significant challenge. To address this, future iterations could explore the integration of Large Language Models (LLMs) to translate natural language intent into executable predicate functions, potentially bypassing the constraints of traditional menu-driven interfaces.

Consequently, future development will focus on conducting formal user studies to evaluate the framework's effectiveness in real-world policy environments. Technical expansions will include exploring LLM-based interfaces for scenario creation, incorporating dynamic grid constraints into the simulation engine, and introducing collaborative features to support multi-stakeholder scenario composition.

## 6. Conclusion

Navigating the energy transition requires more than just raw computation; it demands a decision-making environment that respects the nuanced, negotiable nature of public policy. Building upon our previous work in human-in-the-loop modeling [8], this paper presented an enhanced visual analytics framework that moves beyond the limitations of "black-box" optimization. By transitioning the user's role from a passive observer of algorithmic outputs to an active architect of modular strategies, our system restores the agency necessary to produce accountable and equitable decarbonization plans.

Our work introduced three primary technical and design contributions: (i) a multi-scale visual framework using model-driven glyphs to support comparison across multiple spatial indicators without relying on repeatedly toggled GIS layers; (ii) a modular planning paradigm for composing complex, multi-objective scenarios; and (iii) a synchronized comparative workflow supported by decision-episode provenance tracking. This study specifically extends the prior state-of-the-art by enabling parallel scenario benchmarking and persistent versioning, allowing planners to fluidly move from granular data filtering to real-time strategy comparison.

Ultimately, this research demonstrates that the most effective energy scenarios are those derived transparently. By providing a "what-if" planning canvas that balances data-driven precision with human judgment, our framework offers a scalable path toward achieving evidence-based, net-zero targets.

## CRediT authorship contribution statement

**Dany Laksono:** Writing – review & editing, Software, Methodology, Data curation. **Radu Jianu:** Writing – original draft, Supervision, Software, Methodology, Conceptualization. **Aidan Slingsby:** Visualization, Supervision, Software, Methodology.

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

During the preparation of this work, the author(s) used Gemini (Google) in order to assist with paraphrasing, general sentence structuring, and improving the overall clarity of the manuscript. After using this tool, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Dany Laksono reports financial support was provided by City St George's University of London. Dany Laksono reports a relationship with Advanced Infrastructure, Ltd. that includes: funding grants. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

We thank Advanced Infrastructure Technology Ltd. (AITL) for the data and for co-funding this work along with City St. George's, University of London.

## Data availability

The authors do not have permission to share data.

## References

[1] Collins A, Walker A. Local area energy planning: Achieving net zero locally. 2023, https://researchbriefings.files.parliament.uk/documents/POST-PN-0703/POST-PN-0703.pdf, (Accessed 09 June 2024).

[2] Gudde P, Oakes J, Cochrane P, Caldwell N, Bury N. The role of UK local government in delivering on net zero carbon commitments: You've declared a Climate Emergency, so what's the plan? Energy Policy 2021;154:112245. http://dx.doi.org/10.1016/j.enpol.2021.112245.

[3] Scheller F, Wiese F, Weinand JM, Dominković DF, McKenna R. An expert survey to assess the current status and future challenges of energy system analysis. Smart Energy 2021;4:100057. http://dx.doi.org/10.1016/j.segy.2021.100057.

[4] Sheikhzadehbaboli P, Pakka VH, Bustamante JS, Lee S. Electricity Substation Supply Area based data collation and visualization techniques for local area energy planning: Perspectives from UK distribution network datasets. In: 2024 IEEE international conference on communications, control, and computing technologies for smart grids. smartGridComm, IEEE; 2024, p. 47–52. http://dx.doi.org/10.1109/SmartGridComm60555.2024.10738027.

[5] Hendriks P, Vriens D. From geographical information systems to spatial group decision support systems: A complex itinerary. Geogr Environ Model 2000;4(1):83–104. http://dx.doi.org/10.1080/136159300111388.

[6] Smallman HS, Rieth CA. ADVICE: decision support for complex geospatial decision making tasks. In: International conference on virtual, augmented and mixed reality. Springer; 2017, p. 453–65. http://dx.doi.org/10.1007/978-3-319-57987-0_37.

[7] Illingworth DA, Feigh KM. Impact mapping for geospatial reasoning and decision making. Hum Factors 2022;64(8):1363–78. http://dx.doi.org/10.1177/0018720821999021.

[8] Laksono D, Jianu R, Slingsby A. Interactive visual analytics for local decarbonisation planning: Empowering policy-aligned scenario exploration. In: Computer graphics & visual computing (CGVC) 2025. The Eurographics Association; 2025, http://dx.doi.org/10.2312/cgvc.20251218, URL https://diglib.eg.org/handle/10.2312/cgvc20251218.

[9] Bouw K, Noorman KJ, Wiekens CJ, Faaij A. Local energy planning in the built environment: An analysis of model characteristics. Renew Sustain Energy Rev 2021;144:111030. http://dx.doi.org/10.1016/j.rser.2021.111030.

[10] Molnar S, Johnson G, Gruchalla K, Potter K. Opportunities and challenges in the visualization of energy scenarios for decision-making. In: 2024 IEEE workshop on energy data visualization (energyVis). IEEE; 2024, p. 7–11. http://dx.doi.org/10.1109/EnergyVis63885.2024.00006.

[11] Kotzur L, Nolting L, Hoffmann M, Groß T, Smolenko A, Priesmann J, Büsing H, Beer R, Kullmann F, Singh B, et al. A modeler's guide to handle complexity in energy systems optimization. Adv Appl Energy 2021;4:100063. http://dx.doi.org/10.1016/j.adapen.2021.100063.

[12] Ahmed W, Al-Ramadan B, Asif M, Adamu Z. A GIS-based top-down approach to support energy retrofitting for smart urban neighborhoods. Buildings 2024;14(3):809. http://dx.doi.org/10.3390/buildings14030809.

[13] Huang Y, Zhang Z, Jiao A, Ma Y, Cheng R. A comparative visual analytics framework for evaluating evolutionary processes in multi-objective optimization. IEEE Trans Vis Comput Graphics 2023;30(1):661–71. http://dx.doi.org/10.1109/TVCG.2023.3326921.

[14] Chen L, Wang H, Ouyang Y, Zhou Y, Wang N, Li Q. Fslens: A visual analytics approach to evaluating and optimizing the spatial layout of fire stations. IEEE Trans Vis Comput Graphics 2023;30(1):847–57. http://dx.doi.org/10.1109/TVCG.2023.3327077.

[15] Zhang Z, Yang F, Cheng R, Ma Y. ParetoTracker: Understanding population dynamics in multi-objective evolutionary algorithms through visual analytics. IEEE Trans Vis Comput Graphics 2024. http://dx.doi.org/10.1109/TVCG.2024.3456142.

[16] Sugumaran R, Degroote J. Spatial decision support systems: principles and practices. Crc Press; 2010.

[17] Bush RE, Bale CSE. Energy planning tools for low carbon transitions: An example of a multicriteria spatial planning tool for district heating. J Environ Plan Manag 2019;62(12):2186–209. http://dx.doi.org/10.1080/09640568.2018.1536605.

[18] Urrutia-Azcona K, Usobiaga-Ferrer E, De Agustín-Camacho P, Molina-Costa P, Benedito-Bordonau M, Flores-Abascal I. ENER-BI: Integrating energy and spatial data for cities' decarbonisation planning. Sustainability 2021;13(1):383. http://dx.doi.org/10.3390/su13010383.

[19] Kawakami Y, Yuniar S, Ma K-L. HexTiles and semantic icons for MAUP-aware multivariate geospatial visualizations. 2024, arXiv, abs/2407.16897, arXiv:2407.16897. URL https://arxiv.org/abs/2407.16897.

[20] Liu D, Weng D, Li Y, Bao J, Zheng Y, Qu H, Wu Y. SmartAdP: Visual analytics of large-scale taxi trajectories for selecting billboard locations. IEEE Trans Vis Comput Graphics 2016;23(1):1–10. http://dx.doi.org/10.1109/TVCG.2016.2598432.

[21] Chen J, Huang Q, Wang C, Li C. SenseMap: Urban performance visualization and analytics via semantic textual similarity. IEEE Trans Vis Comput Graphics 2023;30(9):6275–90. http://dx.doi.org/10.1109/TVCG.2023.3333356.

[22] Borgo R, Kehrer J, Chung DH, Maguire E, Laramee RS, Hauser H, Ward M, Chen M. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In: Eurographics (state of the art reports). 2013, p. 39–63. http://dx.doi.org/10.2312/conf/EG2013/stars/039-063, URL http://diglib.eg.org/handle/10.2312/conf.EG2013.stars.039-063.

[23] Slingsby A. Tilemaps for summarising multivariate geographical variation. In: Workshop on Visual Summarization and Report Generation at VIS 2018. 2018, URL https://openaccess.city.ac.uk/id/eprint/20884/.

[24] Beecham R, Dykes J, Hama L, Lomax N. On the use of 'glyphmaps' for analysing the scale and temporal spread of covid-19 reported cases. ISPRS Int J Geo-Information 2021;10(4):213. http://dx.doi.org/10.3390/ijgi10040213.

[25] Laksono D, Slingsby A, Jianu R. Gridded-glyphmaps for supporting geographic multicriteria decision analysis. In: EuroVis 2024 - Short Papers. The Eurographics Association; 2024, http://dx.doi.org/10.2312/evs.20241062, URL https://diglib.eg.org/handle/10.2312/evs20241062.

[26] Xu H, Liu X, Lian J, Li Y, Malhotra M, Omitaomu OA. A geo-visual analysis for exploring the socioeconomic benefits of the heating electrification using geothermal energy. In: Proceedings of the 5th ACM sIGSPATIAL international workshop on advances in resilient and intelligent cities. 2022, p. 11–5. http://dx.doi.org/10.1145/3557916.3567823.

[27] Nusrat S, Alam MJ, Kobourov S. Evaluating cartogram effectiveness. IEEE Trans Vis Comput Graphics 2016;24(2):1077–90. http://dx.doi.org/10.1109/TVCG.2016.2642109.

[28] Kobakian S, Cook D, Roberts J. Mapping cancer: the potential of cartograms and alternative map displays. Ann Cancer Epidemiology 2020;4. http://dx.doi.org/10.21037/ace-19-31.

[29] Cano RG, Buchin K, Castermans T, Pieterse A, Sonke W, Speckmann B. Mosaic drawings and cartograms. In: Computer graphics forum. Vol. 34, Wiley Online Library; 2015, p. 361–70. http://dx.doi.org/10.1111/cgf.12648.

[30] Prouzeau A, Dharshini M, Balasubramaniam M, Henry J, Hoang N, Dwyer T. Visual analytics for energy monitoring in the context of building management. In: 2018 international symposium on big data visual and immersive analytics. BDVA, IEEE; 2018, p. 1–9. http://dx.doi.org/10.1109/BDVA.2018.8534026.

[31] Ma Y, Zhang Z, Cheng R, Jin Y, Tan KC. Paretolens: A visual analytics framework for exploring solution sets of multi-objective evolutionary algorithms [application notes]. IEEE Comput Intell Mag 2025;20(1):78–94. http://dx.doi.org/10.1109/MCI.2024.3487134.

[32] Pammi R, Afzal S, Dasari HP, Yousaf M, Ghani S, Venkatraman MS, Hoteit I. A visual analytics framework for renewable energy profiling and resource planning. EuroVis Work Vis Anal (EuroVA) 2023;49–54. http://dx.doi.org/10.2312/eurova.20231096, URL https://diglib.eg.org/handle/10.2312/eurova20231096.

[33] Hakanen J, Miettinen K, Matković K. Task-based visual analytics for interactive multiobjective optimization. J Oper Res Soc 2021;72(9):2073–90. http://dx.doi.org/10.1080/01605682.2020.1768809.

[34] Sedlmair M, Meyer M, Munzner T. Design study methodology: Reflections from the trenches and the stacks. IEEE Trans Vis Comput Graphics 2012;18(12):2431–40. http://dx.doi.org/10.1109/TVCG.2012.213.

[35] Aprianti T, Tan E, Diu C, Sprivulis B, Ryan G, Srinivasan K, Chua HT. A comparison of ground and air source heat pump performance for domestic applications: A case study in Perth, Australia. Int J Energy Res 2021;45(15):20686–99.

[36] Energy Systems Catapult. Electrification of Heat Demonstration Project. Project reports, Energy Systems Catapult; 2025, URL https://es.catapult.org.uk/project/electrification-of-heat-demonstration-project/, Funded by the Department for Energy Security and Net Zero. Project completed January 2025.

[37] Love J, Smith AZP, Watson S, Oikonomou E, Summerfield A, Gleeson C, Biddulph P, Chiu LF, Wingfield J, Martin C, et al. The addition of heat pump electricity load profiles to GB electricity demand: Evidence from a heat pump field trial. Appl Energy 2017;204:332–42.

[38] Department for Energy Security and Net Zero. National Energy Efficiency Data-Framework (NEED) report: summary of analysis 2026. Accredited official statistics, Department for Energy Security and Net Zero; 2026, URL https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-report-summary-of-analysis-2026.

[39] Rosenow J. Is heating homes with hydrogen all but a pipe dream? An evidence review. Joule 2022;6(10):2225–8.

[40] Ragan ED, Endert A, Sanyal J, Chen J. Characterizing provenance in visualization and data analysis: An organizational framework of provenance types and purposes. IEEE Trans Vis Comput Graphics 2016;22(1):31–40. http://dx.doi.org/10.1109/TVCG.2015.2467551.

[41] Callahan SP, Freire J, Santos E, Scheidegger CE, Silva CT, Vo HT. VisTrails: Visualization meets data management. In: Proceedings of the 2006 ACM SIGMOD international conference on management of data. 2006, p. 745–7. http://dx.doi.org/10.1145/1142473.1142574.

[42] Meyer M, Dykes J. Criteria for rigor in visualization design study. IEEE Trans Vis Comput Graphics 2020;26(1):87–97. http://dx.doi.org/10.1109/TVCG.2019.2934539.

[43] Greenberg S, Buxton B. Usability evaluation considered harmful (some of the time). In: Proceedings of the SIGCHI conference on human factors in computing systems. 2008, p. 111–20. http://dx.doi.org/10.1145/1357054.1357074.