Geo-Topo Maps: Hybrid Visualization of Movement Data over Building Floor Plans and Maps

Quentin Ventura*   Michael J. McGuffin†

École de technologie supérieure

Abstract

We demonstrate how movements of multiple people or objects within a building can be displayed on a network representation of the building, where nodes are rooms and edges are doors. Our representation shows the direction of movements between rooms and the order in which rooms are visited, while avoiding occlusion or overplotting when there are repeated visits or multiple moving people or objects. We further propose the use of a hybrid visualization that mixes geospatial and topological (network-based) representations, enabling focus-in-context and multi-focal visualizations. An experimental comparison found that the topological representation was significantly faster than the purely geospatial representation for three out of four tasks.

Index Terms: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

1 Introduction

Various technologies now allow people, vehicles, and other objects to be tracked over space and time. GPS is commonly used to track outdoor movements, while active badges, RFID, and security cameras can be used to track movements inside a building. Such data can be challenging to visualize and understand, even for small numbers of moving objects over short time spans.

For example, Figure 1.a shows simulated movements of just 3 people within a building. As in typical office buildings, the doorways and hallways are narrow compared to the size of rooms. This creates chokepoints or bottlenecks that people must pass through, causing trajectories to overlap each other. It is also difficult to convey the progression of time within such a 2D map. Although we can see that the red trajectory passes through all 5 rooms labelled A through E, it is difficult to see in what order these rooms are traversed, or if any rooms are visited more than once by the red trajectory. Thus, we have two main problems: first, it is difficult or impossible to see the direction of individual motions and the order in which locations are visited; and second, occlusion makes it difficult to distinguish individual trajectories. A third problem, not immediately apparent but nonetheless important, is that the user may be interested in seeing the details of activities and movements within two rooms that are far from each other, and not need to see the details of movements between these two rooms.

Figure 1.b shows a topological view that provides a solution to some of these problems. Rather than showing the detailed shape of the raw movement data, it shows transitions between nodes, where each node represents a room (or corridor), and each edge is a doorway. Movements are drawn in such a way as to disambiguate them, eliminating occlusion. The diamond icons mark the beginning of each trajectory. Furthermore, movements follow the "right-hand

*e-mail: quentin.ventura.1@ens.etsmtl.ca

†e-mail: michael.mcguffin@etsmtl.ca

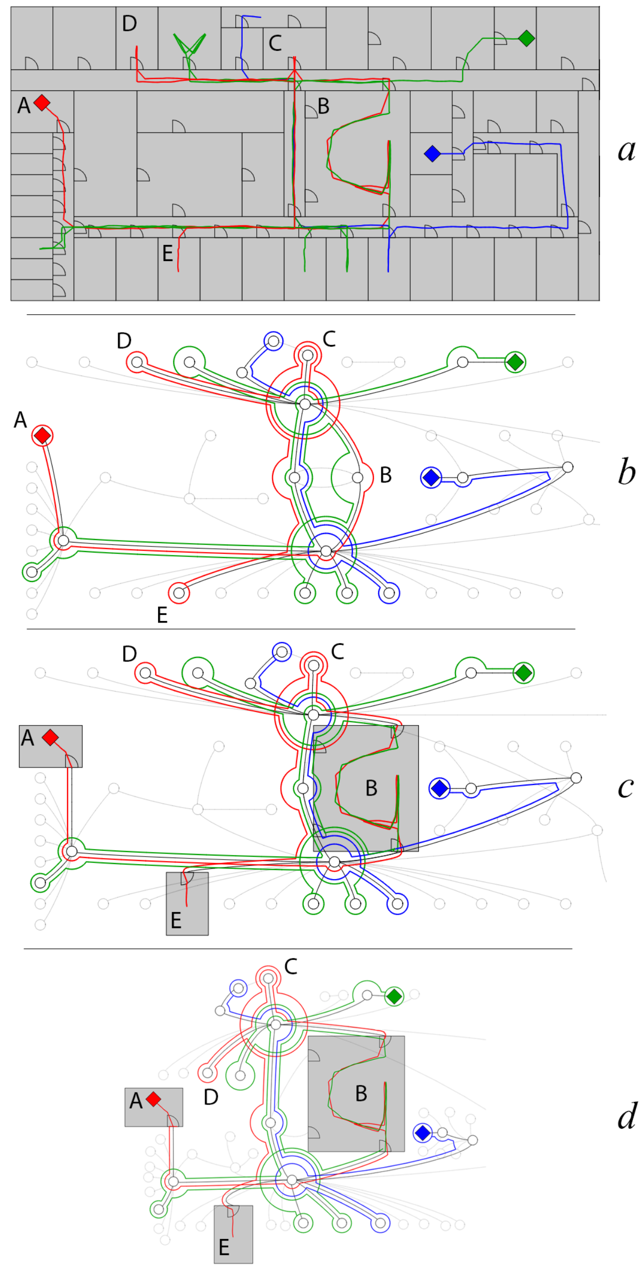

Figure 1: *a:* Movements of 3 people within a floor plan. *b:* Topological view. *c:* Hybrid view showing details and context. *d:* Rooms of interest moved closer together, reducing the visualization's total size, allowing the user to zoom in while preserving context.

traffic" rule, making the direction of every movement segment implicit. The red person travels from A, through B, C, D, back to C, and then ends at E. As long as the number of moving people or objects is not too large, and the user is more interested in the temporal ordering of transitions than detailed motion within each room, this view solves the problems of showing direction, temporal ordering, and occlusion.

Figure 1.c shows a hybrid mixture of the two previous views, yielding a focus-in-context visualization, where the user can see the details of movements in the rooms of interest, and also see the surrounding context in its simplified topological form. We call this hybrid a *Geo-Topo Map*. In Figure 1.d, the user has moved the rooms of interest closer together, reducing the total area taken up by the visualization. This would allow the user to subsequently zoom in without losing any information, and can be thought of as a kind of multi-focal technique where multiple foci can be placed closer together. In this way, all three of the aforementioned problems can be solved.

Such topological and hybrid views could be useful for supporting analytic tasks in understanding the movements of workers, equipment, or robots within a factory during typical days or during an emergency, or movements of fire fighters or police officers responding to an emergency, or care providers, patients, and equipment within a hospital, or the movements of workers or visitors within a museum or building with a mixture of public (open) and private (secure) rooms. Certain building layouts, such as airports or hospitals or campuses, may cover large geographic areas, hence a multi-focal technique could be useful in simultaneously viewing the detailed activity in two or more rooms separated by long hallways, intermediate rooms, or courtyards that are of less interest.

Our work investigates the use of these topological and hybrid visualizations of movement data. Other recent work has proposed several ways of visualizing spatio-temporal data [2], some of which use aggregation to scale up to large movement datasets [1, 17, 27, 25]. In contrast, our approach does not scale up to large numbers of moving objects, but has the advantage of showing each individual moving object or person. Our approach could thus provide an intermediate level of abstraction, in between aggregation of many objects at one extreme, and showing the detailed shape of the trajectories at the other extreme.

Our contributions are (1) the use of circular arcs around nodes to show temporal ordering of movements in topological views, (2) a hybrid mixture of geospatial and topological views, (3) a discussion of design issues, and (4) an experimental comparison of geospatial (i.e., floor plan) views and topological views that found topological views to be significantly faster for certain tasks.

2 Related Work

Andrienko and Andrienko [2] present an overview of approaches for visualizing spatio-temporal data. We briefly survey a few here.

One way to analyze previous work is in the way it maps data variables to the visualization. Movement data can be characterized as having a single temporal variable and at least two spatial variables (latitude and longitude, and sometimes altitude). 2-dimensional line graphs, with a horizontal time axis, can be used to show changes in movements over time, such as in [6], where spatial position is shown as a distance (with respect to some reference location) on the vertical axis, as a function of time. Other systems use a 3D visualization, showing time, latitude, and longitude simultaneously [15]. Certain systems allow the user to interactively change the mapping of data variables to axes [13]. TripVista [11] computes multiple attributes on movement segments, and uses parallel coordinates to visualize the resulting multivariate data. Our own work maintains the familiar 2D spatial layout in the visualization, making it easy for users to relate the visualization to a floor plan, however our approach also uses graphic techniques to show temporal information within the 2D layout.

Previous work has also provided the user with a lens for filtering or selection, or some kind of magnification or focal capability [24, 12, 16]. Our system does not implement a lens, however our hybrid visualization can be thought of as showing one or more foci within a surrounding topological view.

General techniques for aggregation are surveyed in [10]. Aggregation of many movements and/or many moving objects can provide the user with an overview of large movement data sets, and has been demonstrated in [1, 17, 27, 25]. Our current work is not focused on aggregation of multiple moving objects. However, Shneiderman's classic mantra [22] recommends that the user be able to drill down into the details of an overview, and our approach might be useful to invoke when the user wishes to drill into an overview of a massive movement dataset.

Our approach is more closely related to techniques that simplify the shape of spatial trajectories, and that show movements as discrete transitions between locations. Henry Beck's London Underground Tube map is famous for having made subway maps easier to read. Such maps can be seen to discretize locations (subway stations), and have been updated in recent work [24]. However, we are unaware of previous topological maps of subways that show the movements of individuals within the subway system. Visits [23] displays two views of data: a geographic map, and a linear sequence of local maps, with the two views linked with curves. Compared to our hybrid approach, Visits has the advantages of making the linear sequence of locations very apparent, and also shows the full geographic information of the surrounding context. Our hybrid approach, however, has the advantages of embedding geospatial details within the topological view (yielding a single, integrated view), and can also show the movements of multiple individuals.

Finally, our hybrid approach is comparable to other hybrids like TreeMatrix [20], where one type of visualization is *nested* [14] within another. It can also be thought of as a focus-in-context, multi-focal visualization technique [5], with geospatial (floor plan) views providing the details, and a topological view providing the surrounding context.

3 Topological Representation

Our topological views are based on *space-portal* graphs, a term used by [26]. The space of the original movement data is partitioned into regions (such as rooms in a floor plan, or countries in a geographic map). The space-portal graph is then defined with one node for each region, and one edge for each "portal" (doorway, border crossing, etc.) Note that in the case of a floor plan, the rooms *and hallways* (corridors) each correspond to a node.

Blaas et al. [3] introduced "smooth graphs" to display transitions across 2 or more nodes as curved links with animated textures to show direction. The curved links in their system often overlap. We instead chose an approach that would eliminate all overlap between trajectories in the topological view, at the expense of requiring a band of space along each edge whose thickness is linear in the number of traversals. This band of space, within which traversals are drawn, is comparable to the space reserved for edge bundles in [18]. Figure 2 illustrates. Around each node is a local polar coordinate system. Within this coordinate system, the radius corresponds to *time of arrival* of each person or object, increasing outward. Thus, trajectories with a smaller radius arrived earlier in the room (or hallway). The angle of the trajectory within the polar coordinate system is determined simply by the direction toward the previous (or next) node in the trajectory.

This use of circular arcs of increasing radius around each node is somewhat similar to how kelp diagrams [7] increment the "thickness" (radius) of subsets around nodes to visually disambiguate them, or to how circular arcs are laid out in AlertWheel [9].

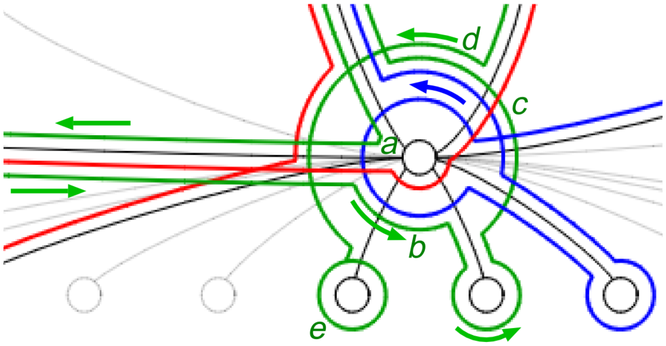

Figure 2: Enlarged view of Figure 1.b, centered on a node corresponding to a hallway. Movement along graph edges follows the "right-hand traffic" rule, and movement around nodes is always counterclockwise. Arrows were added for illustration only, and are implicit in the visualization. Movement segments farther out from graph edges or from node centers occur later in time. Green travels through points *a*, *b*, *c*, and *d* in chronological order, finally ending at *e*.

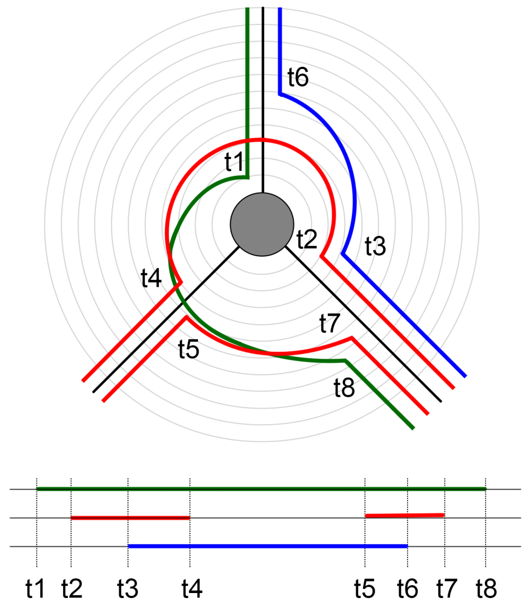

Figure 3: *Top:* a mock-up of an alternative design for the circular arcs around a node, showing progression of time within each node. *Bottom:* a gantt chart for the node.

To convey the direction of trajectories in our system, we employed a metaphor based on how cars drive on the right side of the roads in most countries. Once the user is accustomed to this, they can easily tell the direction of motion, without any arrow heads or animations being displayed. (In our user study, described later, most users appeared to understand this metaphor with little explanation.) Furthermore, trajectories always move counterclockwise around nodes. Thus, the circular arcs sometimes cover an angle greater than 180 degrees, such as the blue arc in Figure 2, to maintain this convention.

In our system, the radius of each circular arc around the nodes is held constant and determined solely by arrival time, thus the user cannot know in what order people or objects *leave* a node. An alternative design is sketched in Figure 3, where the radius of the arcs is gradually increased as time progresses. In theory, this conveys the arrival and departure time of each trajectory, and should allow the user to see if two trajectories are within the same node simultaneously. However, we found this design to be somewhat difficult to interpret, and so implemented the one in Figure 2.

3.1 Alternative Graph-Based Models of Topology

In this section, we briefly consider other kinds of graphs for modeling the topological relationships in a floor plan, that could plausibly result in simpler and easier-to-understand visualizations. The analysis that follows will help ensure that we have properly considered the potential alternatives before settling on a final choice.

As a brief reminder, given a graph of N nodes and E edges, if the graph is connected, then $E \geq N-1$. If the graph is planar, then the embedding of the graph in the plane divides the plane into R regions (including the exterior region), where $R = E - N + 2$ (by Euler's formula).

Now, consider a floor plan with r rooms and d doors. Let $r_i < r$ be the number of rooms with i doors each, e.g., r_1 is the number of rooms with 1 door each. Also let r_{i+} be the number of rooms with i or more doors each, e.g., r_3+ is the number of rooms with at least 3 doors each. In the example of Figure 4.a, we have $r = 11$, $d = 14$, $r_1 = 5$, $r_2 = 3$, $r_{3+} = 3$.

The space-portal graph for the floor plan is constructed with r nodes and d edges. If a pathway exists between every pair of rooms, then the space-portal graph is connected, and $d \geq r-1$. If furthermore the floor plan is planar (e.g., the floor plan is for a single floor of a building), then the number of regions in the space-portal graph is $d-r+2$. In the example of Figure 4.b, there are $14-11+2 = 5$ regions.

Other graphs can be constructed that contain the same information as in the space-portal graph, but that may be visually simpler. First, consider the *geometric dual* of the space-portal graph, which is constructed by replacing each region with a node, and adding an edge between two nodes if and only if the corresponding regions in the original space-portal graph are adjacent along an edge. The resulting geometric dual has $d-r+2$ nodes, d edges (including r_1 self-loops), and r regions, i.e., the same number of edges as the original space-portal graph, but with the numbers of nodes and regions swapped. Figure 4.c illustrates.

Consider also the *hypergraph dual* of the space-portal graph, which we call the portal-space graph. The hypergraph dual is constructed by interchanging nodes and edges, yielding a portal-space graph with d nodes (each of degree 2) and r hyperedges (each incident on 1 or more nodes). In Figure 4.d, nodes are drawn as unlabelled points, and hyperedges as closed curves.

Finally, Figure 4.e shows an alternative drawing of the hypergraph dual, where rooms having 2 doors are drawn more simply as open curve segments connecting their two doors, and rooms having only 1 door are drawn only as a label near the corresponding door.

To compare the visual complexity of these graphs, we assume that small, localized elements (points, nodes, and labels) incur a negligible cost, and we only count up the number of extended curves or line-like elements in the following table.

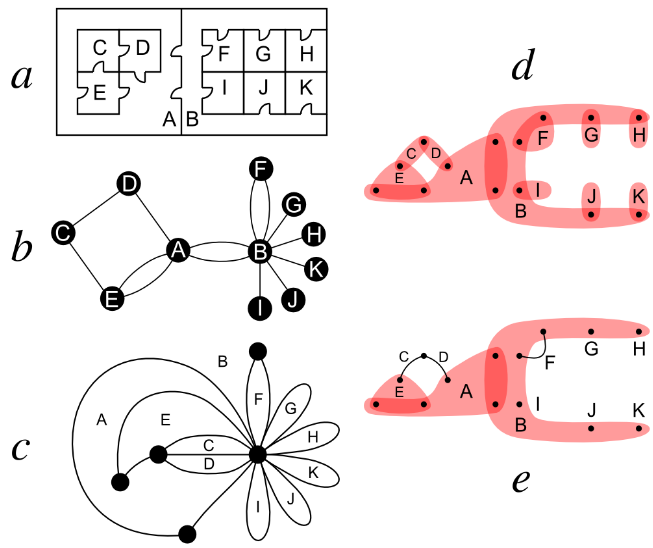

Figure 4: *a:* Original floor-plan. *b:* Space-portal graph, where each node is a room, and each edge is a door. *c:* The geometric dual of the space-portal graph, where each region is a room, and each edge is a door. *d:* The portal-space graph, where each node is a door, and each hyperedge is a room. *e:* A simpler way of drawing the portal-space graph.

	number of open curves or non-loop edges	number of closed curves (self-loops or subsets)	total
Space-Portal Graph (r nodes, d edges, $(d-r+2)$ regions)	d	0	d
Geometric Dual ($(d-r+2)$ nodes, d edges, r regions)	$d-r_1$	r_1	d
Portal-Space Graph (d nodes, r hyperedges)	0	r	r
Simplified Portal-Space Graph (d nodes, r hyperedges)	r_2	r_{3+}	$r-r_1$

Comparing the "total" column, and keeping in mind $d \geq r-1$, we see that the simplified portal-space graph generally requires fewer curved elements to be drawn, especially if there are many 1-door rooms. In our personal experience, 1-door rooms are common in office buildings. The simplified portal-space graph may be a promising avenue for future research into simplified topological representations of floor plan data, especially large building floors with many small offices.

Despite the possible advantages of the portal-space graph, we must keep in mind that we seek not only a representation of topological relationships between rooms, but also a way to show movements through rooms without occlusion. In the portal-space graph, because the doors in it are drawn as points, it is unclear how to draw multiple trajectories traveling through doors while avoiding occlusion or overplotting. For this reason, we finally decided to stick to using normal space-portal graphs in our visualizations, where nodes represent rooms.

4 Geo-Topo Maps

As already discussed, Figure 1 shows how the topological and geospatial (e.g., floor plan) views can be mixed into a hybrid. We call these hybrid visualizations *Geo-Topo Maps*. Figure 5 is a mock-up illustrating that our topological and hybrid techniques are not limited to building floor plans. They can be applied to any movement data over a space that has been partitioned into regions, whether these be rooms or countries.

4.1 Implementation

Our prototype implementation can display data in all the forms shown in Figure 1. The user can lasso select multiple nodes (or rooms) and toggle their representation between the topological style and the geospatial (floor plan) style.

In addition, a time slider allows the user to navigate through time, and see small colored diamond icons move along the trajectories to indicate the current positions of people. In the geospatial mode, these diamond icons move smoothly in response to the time slider. However, in the topological mode, the icons jump instantaneously across edges, from one node to another, because such jumps correspond to the instantaneous traversal of a doorway. In future work, we suspect it would be worthwhile to introduce piecewise "ramp" functions that cause the icons to gradually travel over edges in response to the time slider, making them easier to follow by the user, even though such smooth motion would not be a completely accurate reflection of the true movement data which moves from room to room instantaneously.

4.2 Example Applications

Hybrid Geo-Topo maps could be used to summarize the activities of a single person. For example, a typical factory worker might move between two machine rooms, a cafeteria, and a washroom, and these different rooms might be separated by many intermediate rooms where the worker does not stop. A hybrid Geo-Topo map could display detailed activities within each room of interest, and summarize the intermediate movements with the topological representation. A Geo-Topo map could also display the movements of 2 or 3 different kinds of employees, making clearer which rooms are used exclusively by one type of employee, and which ones are used by multiple types of employees, and in what order. Such visualizations might be useful for planning changes to building layout, comparing alternative layouts, or improving assembly lines. Geo-Topo maps could also be used to summarize the peculiar movements of an outlier (e.g., a thief or spy). Hospitals, campuses, and airports cover large areas and could particularly benefit from topological summarization of intermediate regions.

In certain scenarios, multiple people are involved in passing a physical item from person to person. For example, a bomb or suitcase might be passed from accomplice to accomplice throughout a large building, and being able to retrospectively visualize the meetings between one suspect and other people with detailed geospatial views could help an analyst reconstruct a chain of people and then visually summarize their analysis with a Geo-Topo map. Infections that are spread through physical contact might also be visualized with a Geo-Topo map, using geospatial views to show possible contacts between people, and topological views to summarize intermiate movements.

A purely topological view could also be used to show movements of buses within a road network, helping users to see which buses they can take to travel from one location to another.

5 Experimental Comparison

To investigate the performance advantages that the topological view might have over a geospatial view, we compared both in a controlled experiment. The hybrid mixture of visualizations was not evaluated in this experiment, since we wished to first compare each visualization technique in its "pure" (non-hybrid) form.

Four kinds of tasks were given to users. Each task required the user to answer a multiple-choice question about the dataset being

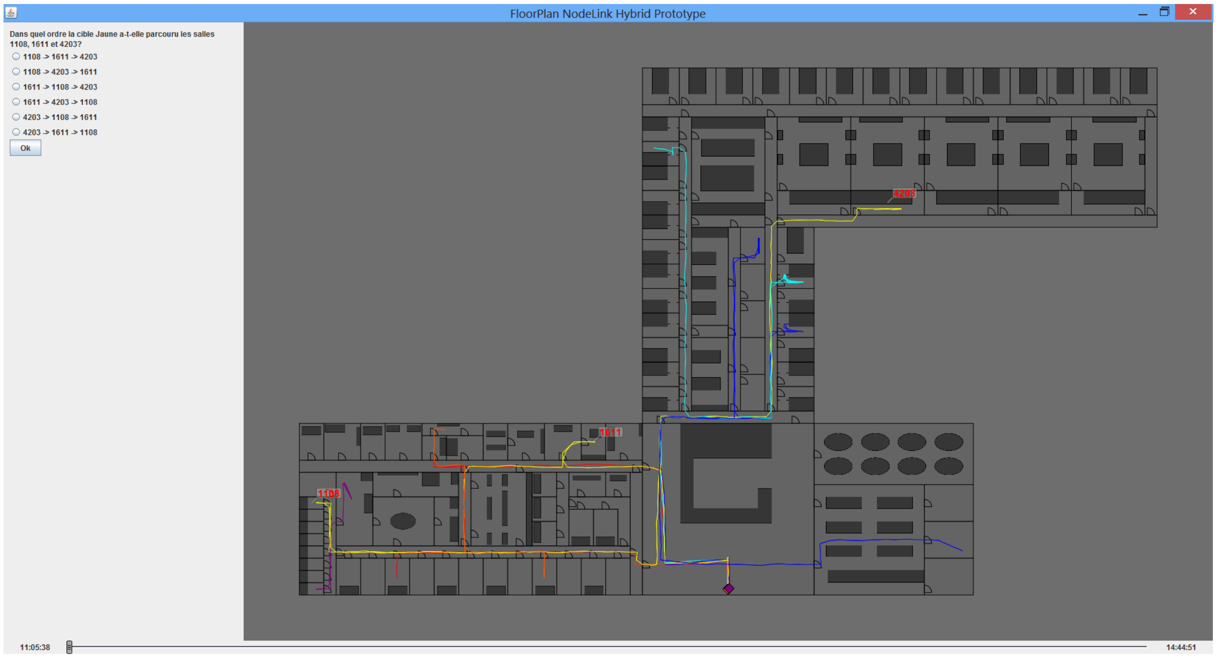

Figure 6: The user interface shown to the user during the experiment, in the geospatial (floor plan) condition. The multiple-choice question appears in the upper left corner, and the time slider widget is along the bottom. Note that the large room in the middle bottom contains virtual furniture (not displayed), requiring individuals to move around the furniture within the room.

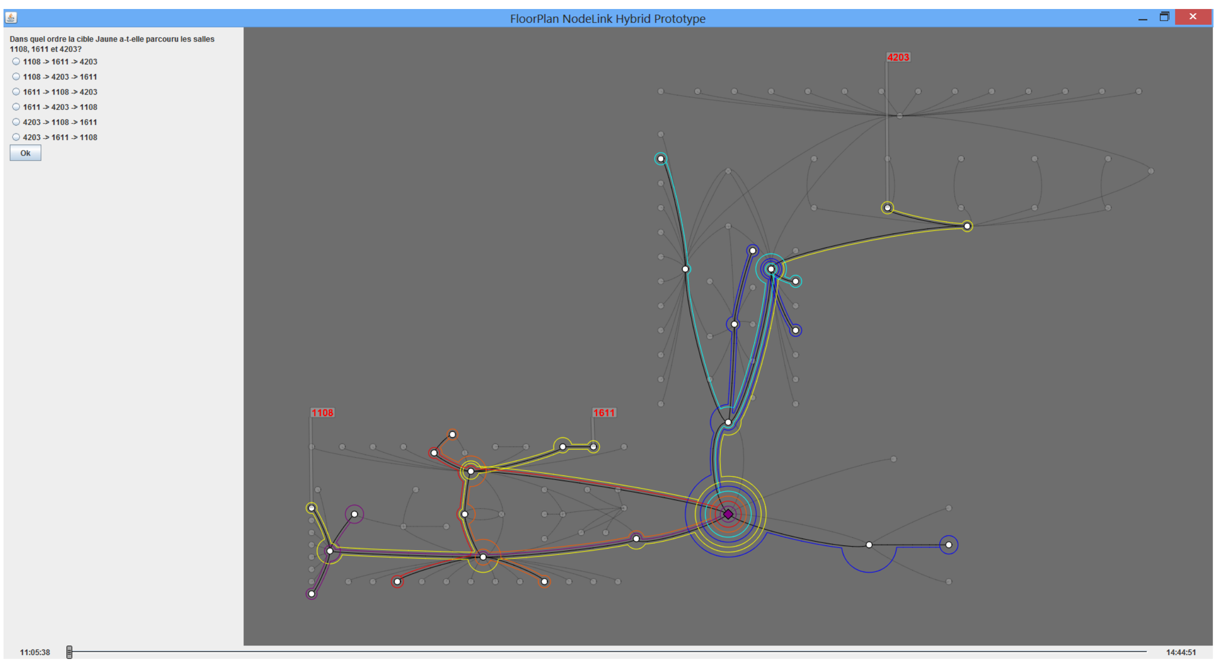

Figure 7: The user interface shown to the user during the experiment, in the topological condition.

viewed, either using the geospatial (floor plan) or topological visualization. The four tasks were based on the following template questions:

- T1: How many times did person X visit room Y? (6 possible answers, varying from “1 time” to “6 times”.)
- T2: In what order did person X visit rooms U, V and W? (6 possible answers, covering all permutations such as “U → V → W”, “U → W → V”, etc.)
- T3: Which person arrived first in room U? (6 possible answers, covering all 6 people in the dataset.)
- T4: How many persons visited both rooms U and V? (6 possible answers, varying from “1” to “6”.)

Note that the locations of rooms mentioned in the questions were indicated with red labels on the visualizations, in both conditions.

Synthetic datasets were generated for the experiment. Although this carries the risk of being less realistic as data, it has the advan-

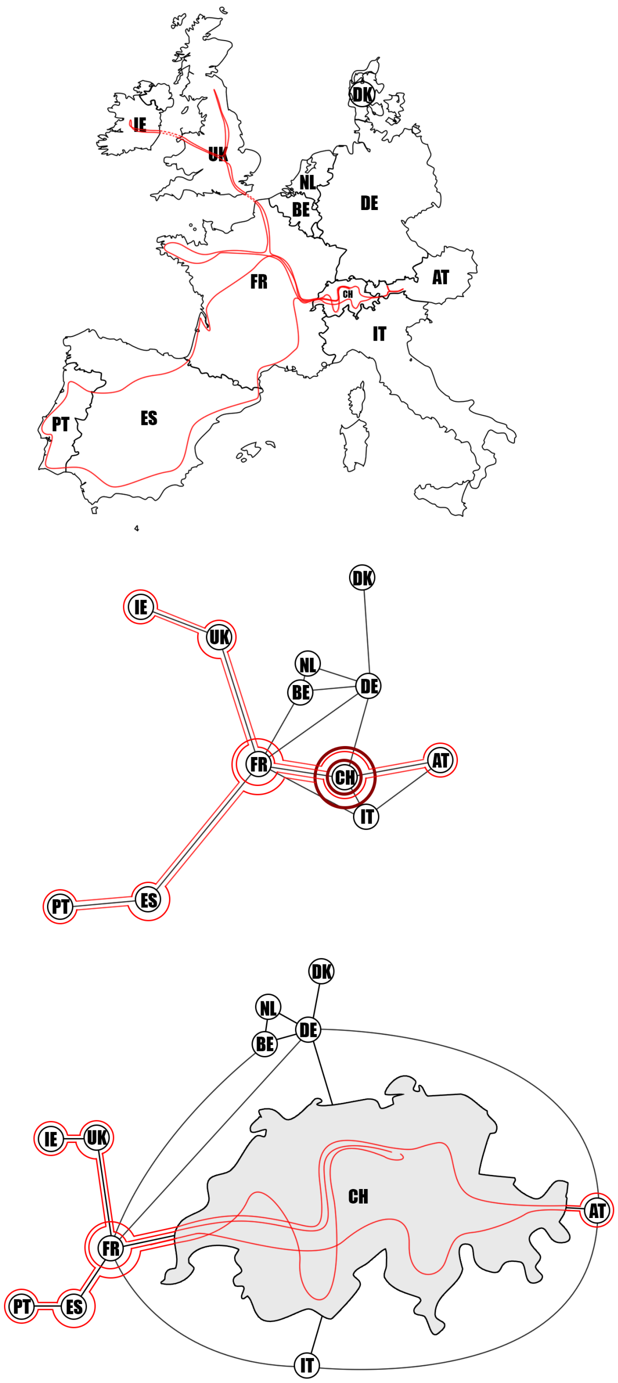

Figure 5: *Top:* A geographic map of a trip, obscuring small details and failing to show temporal ordering. *Middle:* A topological representation shows that the traveler began in Switzerland (shown by the small, dark circle around CH) and traveled through various countries to finally return to the same country (shown by the large, dark circle around CH). *Bottom:* A hybrid view allows the user to zoom in on details within Switzerland, while preserving the essential information about the surrounding contextual countries.

tage of being far more amenable to controlling the characteristics of the data, which is important for a controlled experiment. For each dataset, movements of 6 people were simulated within the building layout shown in Figures 6 and 7. Each simulated person was made to walk through between 3 and 6 target rooms, using Dijkstra's algorithm to find shortest paths between rooms to move along. In addition, each room was modeled with a subgraph (not shown in the topological visualization), representing the possibly pathways within the room around pieces of furniture. Within each room, between 1 and 11 nodes were randomly selected in the room's subgraph for the person to visit, again using Dijkstra's algorithm to move within the subgraph of each room. The person was made to stay between 5 to 15 minutes within each target room, and move with an approximately constant speed between nodes. Finally, a small amount of noise was added to movements, computed from a *wandering* behavior [19].

8 datasets (D1 through D8) were generated for the experiment, each containing generated movements of 6 people. 4 questions were prepared for each dataset, based on the tasks T1-T4 above. Each user performed tasks with half of the datasets in the geospatial condition, and the other half in the topological condition. The assignment of datasets to conditions, and the ordering of conditions, was fully counterbalanced. In other words, each quarter of users performed tasks according to one of the following orderings:

- Geo+(D1,D2,D3,D4), Topo+(D5,D6,D7,D8)
- Geo+(D5,D6,D7,D8), Topo+(D1,D2,D3,D4)
- Topo+(D1,D2,D3,D4), Geo+(D5,D6,D7,D8)
- Topo+(D5,D6,D7,D8), Geo+(D1,D2,D3,D4)

Each user thus performed a total of 8 datasets $\times$ 4 trials/dataset = 32 trials. Of these, the tasks performed with D1 and D5 were warmup tasks and were not counted in the final analysis, leaving 6 $\times$ 4 = 24 trials per user. 12 users participated, for a total of 288 trials collected.

At the start of each trial, the visualization window showed only the building layout (in the form of a floor plan or a space-portal graph), with no visible trajectories. The question for the trial was also visible, in the upper-left corner of the window. The participant was asked to read the question, and once they understood the question and were ready to begin, they clicked, causing the trajectories to be displayed, as well as the 6 possible answers to the question, and causing the timer for the trial to begin. The user could then interact with the visualization, by zooming, panning, or by dragging on the time slider. (The user could not reposition nodes or rooms; the layout was fixed.) Once the user had determined the answer to the question, they selected the radio button beside the desired answer and clicked the "Ok" button to confirm, at which point the timer was stopped. Visual feedback was given to tell the user if they succeeded or failed the question. Clicking then moved the user on to the next trial. Thus, the user had only one chance to get each question correct.

Users ranged in age from 20 to 36 years, and all but one had a background in engineering.

5.1 Results

Results were analyzed using analysis of variance (ANOVA). Visualization technique had a significant effect ($F_{1,11} = 17.17$, $p < 0.001$) on time, as did task ($F_{3,33} = 10.26$, $p < 0.00001$). Ordering had no significant effect ($p > 0.05$) on time.

Neither the visualization technique nor the task had a significant effect on error rate ($p > 0.05$). In the geospatial condition, the error rates for the four tasks T1-T4 were 5.6, 0.0, 11.1, and 11.1%, respectively. In the topological condition, they were 5.6, 16.7, 2.8, and 5.6%, respectively.

Users were asked to give subjective ratings of the interface, and found it "intuitive" (3.91/5), "easy to learn" (4.36/5), and were "able to accomplish what [they] wanted to do" (4.82/5).

Further analyzing time, the topological visualization was significantly faster for tasks T1, T3, and T4, but significantly slower for task T2 (Figure 8).

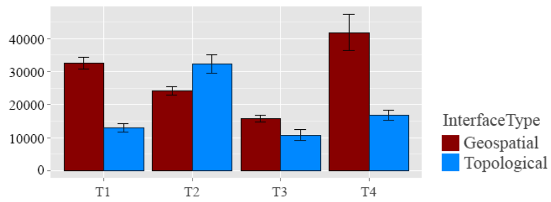

Figure 8: Median times in milliseconds, with 95% confidence intervals. Within each task type, the two visualization techniques were significantly different.

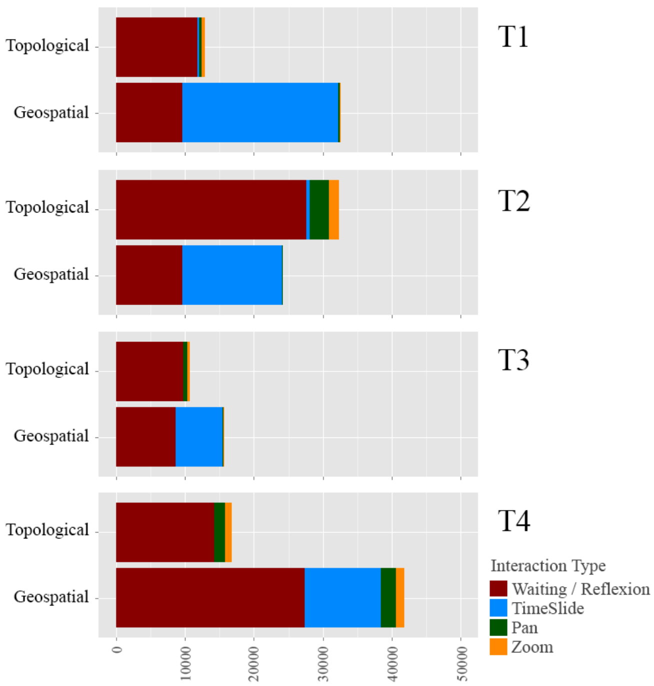

Figure 9: Average times in milliseconds, broken down by the fraction of time spent by the user not dragging on anything (red), interacting with the time slider (blue), performing pan operations (green), and zooming (yellow).

Figure 9 breaks down times by how much time the user spent performing camera or time slider operations. Notice that in the topological condition, the user spent very little time dragging on the time slider, in contrast with the geospatial (floor plan) condition. In the topological condition, tasks T1 and T3 could be performed by simply looking at the circular arcs surrounding the node mentioned in the question, and we suspect this is what users actually did, allowing them to complete the task quickly, whereas the same tasks in the geospatial condition required interacting with the time slider. Task T4 could be completed, in theory, simply by looking at the colors of trajectories present in the two rooms mentioned in the question, however in practice this was difficult to do in the geospatial condition due to occlusion hiding the trajectories.

Task T2 was the only one where the geospatial condition was faster. This task could be completed by the user in the topological condition by using one's eyes to follow the trajectory through the rooms, and from Figure 9, it seems that users made very little use of the time slider in this condition. However, using the time slider to complete this task in the geospatial condition turned out to be faster, as the user could watch the motion of the diamond icons as they dragged on the slider.

Returning to tasks T1 and T3, and noticing the large difference in time slider use between geospatial and topological, we conjecture that the topological condition would continue to be faster even with datasets involving longer durations of time. Such datasets would presumably require more time slider scrolling when viewed in geospatial form, while not requiring more time to interpret the circular arcs around nodes in the topological form.

5.2 Potential Improvements

The experimental results suggest that the topological view may be more appropriate for summarizing activity within a single room (tasks T1 and T3), whereas the time slider in the geospatial condition was best for showing the order in which rooms are visited (task T2). With other data sets, however, dragging on a time slider to cause the diamond icons to re-enact the movements of individuals could be problematic: long time spans might make the time slider's gain too high, and individuals who spend many hours in one room could be reflected in diamond icons that are motionless for most of the user's drag and then suddenly jump to another room. Two ways this could be improved, in both the topological and geospatial views, are: (1) allowing the user to drag directly on a trajectory (either in the geospatial or topological forms) to navigate through time, similar to dragging on moving objects within videos [8] and animations [21] to implicitly move from frame to frame; and (2) displaying rapid animations of the movements of individuals whenever the user rolls their cursor over a trajectory (e.g., repeatedly "plucking" a trajectory with the cursor could cause the animation to replay over and over, giving the user a sense of the order in which locations are visited).

In the topological view, at least three avenues exist to improve scalability: (1) hallways could be broken up into shorter, simpler segments connected by virtual doors, making the topological view easier to understand and less distorted with fewer individuals traversing each part of the hallway; (2) multiple similar trajectories might be aggregated to compute a "median trajectory" [4] that is displayed to summarize the entire group or cluster; (3) rooms visited many times might be enlarged whenever the cursor passes over the room, similar to a magnification lens or popup view, making it easier to examine the multiple circular arcs surrounding the room.

6 Conclusions and Future Directions

The topological visualization we have presented can be applied to movement data where space has been partitioned into regions, such as rooms or countries. It has the advantage of simplifying the shapes of trajectories to show the order in which regions (rooms, countries, etc.) are visited, while avoiding occlusion between multiple individuals. The direction of these trajectories can be clearly interpreted, thanks to the use of a metaphor based on the right-hand traffic rule, and the order in which each trajectory arrives at a node is also clearly shown.

Furthermore, when the topological view is mixed with the geospatial view, we obtain a hybrid visualization affording multiple foci-in-context, allowing the user to see detailed movement trajectories where desired.

Our experimental comparison established that the topological visualization is superior to the geospatial visualization for certain tasks, but not all. This further motivates the investigation of combinations of the two visualizations, to benefit from the advantages of both techniques.

As a next step, future work could experimentally evaluate different combinations of geospatial and topological visualizations. For example, a user interface with two coordinated views (one purely geospatial, one purely topological, with coordinated highlighting linking the two) could be compared against the hybrid Geo-Topo Map proposed in this paper. The coordinated view interface would theoretically have the advantage of allowing the user to quickly switch (with a fast eye movement) to whatever visual representation is best for the task (or subtask) at hand, whereas the hybrid would allow more screen space to be devoted to the one representation, and also avoid the user having to move their eyes back and forth between two views.

Other future directions include modifying the visual design of the circular arcs in the topological visualization to indicate when two people are in the same room at the same time (allowing the user to easily perceive meetings), and developing ways to scale the topological visualization up to a larger number of nodes (perhaps by allowing nodes to be collapsed into meta-nodes) and a larger number of moving people or objects.

Acknowledgements

Thanks to the participants in our study for their valuable time, and to the members of the HIFIV research group at ÉTS for their feedback. This research was supported by an NSERC Strategic Project Grant.

References

[1] N. Andrienko and G. Andrienko. Spatial generalization and aggregation of massive movement data. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 17(2):205–219, 2011.

[2] N. Andrienko and G. Andrienko. Visual analytics of movement: An overview of methods, tools and procedures. *Information Visualization*, 12(1):3–24, 2013.

[3] J. Blaas, C. P. Botha, E. Grundy, M. W. Jones, R. S. Laramee, and F. H. Post. Smooth graphs for visual exploration of higher-order state transitions. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 15(6):969–976, 2009.

[4] K. Buchin, M. Buchin, M. van Kreveld, M. Löffler, R. I. Silveira, C. Wenk, and L. Wiratma. Median trajectories. *Algorithmica*, 66(3):595–614, 2013.

[5] A. Cockburn, A. Karlson, and B. B. Bederson. A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys (CSUR)*, 41(1), 2008.

[6] T. Crnovrsanin, C. Muelder, C. Correa, and K.-L. Ma. Proximity-based visualization of movement trace data. In *Proceedings of IEEE Visual Analytics Science and Technology (VAST)*, 2009.

[7] K. Dinkla, M. J. van Kreveld, B. Speckmann, and M. A. Westenberg. Kelp diagrams: Point set membership visualization. *Computer Graphics Forum*, 31(3):875–884, 2012.

[8] P. Dragicevic, G. Ramos, J. Bibliowitcz, D. Nowrouzezahrai, R. Balakrishnan, and K. Singh. Video browsing by direct manipulation. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, pages 237–246, 2008.

[9] M. Dumas, J.-M. Robert, and M. J. McGuffin. AlertWheel: Radial bipartite graph visualization applied to intrusion detection system alerts. *IEEE Network, Special Issue on Network Visualization*, 26(6):12–18, 2012.

[10] N. Elmqvist and J.-D. Fekete. Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 16(3):439–454, 2010.

[11] H. Guo, Z. Wang, B. Yu, H. Zhao, and X. Yuan. TripVista: Triple perspective visual trajectory analytics and its application on microscopic traffic data at a road intersection. In *Proceedings of IEEE Pacific Visualization (PacificVis)*, 2011.

[12] J.-H. Haunert and L. Sering. Drawing road networks with focus regions. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 17(12):2555–2562, 2011.

[13] C. Hurter, B. Tissoires, and S. Conversy. FromDaDy: Spreading aircraft trajectories across views to support iterative queries. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 15(6):1017–1024, 2009.

[14] W. Javed and N. Elmqvist. Exploring the design space of composite visualization. In *IEEE Pacific Visualization Symposium (PacificVis)*, pages 1–8, 2012.

[15] T. Kapler and W. Wright. GeoTime information visualization. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, 2004.

[16] R. Krüger, D. Thom, M. Wörner, H. Bosch, and T. Ertl. TrajectoryLenses: A set-based filtering and exploration technique for long-term trajectory data. *Computer Graphics Forum*, 32(3pt4):451–460, 2013.

[17] H. Liu, Y. Gao, L. Lu, S. Liu, H. Qu, and L. M. Ni. Visual analysis of route diversity. In *Proceedings of IEEE Visual Analytics Science and Technology (VAST)*, pages 171–180, 2011.

[18] S. Pupyrev, L. Nachmanson, and M. Kaufmann. Improving layered graph layouts with edge bundling. In *International Symposium on Graph Drawing (GD)*, 2010.

[19] C. W. Reynolds. Steering behaviors for autonomous characters. In *Game Developers Conference*, pages 763–782, 1999.

[20] S. Rufiange, M. J. McGuffin, and C. P. Fuhrman. TreeMatrix: A hybrid visualization of compound graphs. *Computer Graphics Forum*, 31(1):89–101, 2012.

[21] S. Santosa, F. Chevalier, R. Balakrishnan, and K. Singh. Direct space-time trajectory control for visual media editing. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, pages 1149–1158, 2013.

[22] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of IEEE Symposium on Visual Languages (VL)*, pages 336–343, 1996.

[23] A. Thudt, D. Baur, and S. Carpendale. Visits: A spatiotemporal visualization of location histories. In *Proceedings of Eurographics Conference on Visualization (EuroVis), Short Papers*, pages 79–83, 2013.

[24] Y.-S. Wang and M.-T. Chi. Focus+context metro maps. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 17(12):2528–2535, 2011.

[25] Z. Wang, M. Lu, X. Yuan, J. Zhang, and H. van de Wetering. Visual traffic jam analysis based on trajectory data. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 19(12):2159–2168, 2013.

[26] E. Whiting, J. Battat, and S. Teller. Topology of urban environments. In *Computer-Aided Architectural Design Futures (CAADFutures) 2007*, pages 114–128. Springer, 2007.

[27] W. Zeng, C.-W. Fu, S. M. Arisona, and H. Qu. Visualizing interchange patterns in massive movement data. *Computer Graphics Forum*, 32:271–280, 2013.