J geovis spat anal
DOI 10.1007/s41651-017-0001-7

An Interactive Approach for Exploration of Flows Through Direction-Based Filtering

Katerina Vrotsou[1] · **Georg Fuchs**[2] · **Natalia Andrienko**[2,3] · **Gennady Andrienko**[2,3]

© The Author(s) 2017. This article is an open access publication

Abstract This paper is concerned with the representation and exploration of flows, defined as spatial interactions between geographic locations. Flows are challenging to display in a comprehensible manner due to the nature of the data, which are characterized by many crossings and overlaps leading to clutter. A number of different strategies have been suggested for addressing this problem, which commonly involve reducing the search space, aggregating the data or simplifying the representations often at the cost of information loss or distortion of spatial context. We propose an interactive approach for exploring large and highly connected networks of flows without distorting the geographical space and without losing the context overview in the process. The approach is based on a flow-specific interaction technique for filtering the data by direction, that enables an analyst to successively identify underlying spatial arrangement patterns. We illustrate our approach through exploring flows of tourists and locals in the Greater London area.

Keywords Flow analysis · Spatial arrangement pattern · Interactive query · Direction-based filtering

✉ Katerina Vrotsou
katerina.vrotsou@liu.se

Georg Fuchs
georg.fuchs@iais.fraunhofer.de

Natalia Andrienko
natalia.andrienko@iais.fraunhofer.de;
natalia.andrienko.1@city.ac.uk

Gennady Andrienko
gennady.andrienko@iais.fraunhofer.de;
gennady.andrienko.1@city.ac.uk

[1] Department of Science and Technology, Linköping University, SE-601 74 Norrköping, Sweden

[2] Fraunhofer Institute IAIS, Schloss Birlinghoven, 53757 Sankt-Augustin, Germany

[3] City University London, London EC1V OHB, UK

Introduction

A flow is a representation of movement as a spatial interaction between a pair of locations (Rodrigue et al. 2009). Hence, flows are a type of origin-destination data which are characterized by: (1) a set of locations, P, which can be points or regions in a two dimensional space, (2) a set of links, L, connecting these locations and composing origin/destination pairs, and (3) a function, M, that defines the interaction between locations by assigning *magnitudes* to the links (Boyandin 2013; Gibson et al. 2012). Each link is directed and is specified as an ordered pair of locations (p_i, p_j) connecting the location of origin, p_i, to the location of destination, p_j, where $(p_i, p_j) \neq (p_j, p_i)$. The magnitudes defining the interactions between these locations correspond to non-negative values, R^+, which are mapped onto the links through a function $M : L \rightarrow R^+$. Consequently, a flow is defined by a pair of linked locations and the corresponding magnitude of this link, and can be represented as a triple (p_i, p_j, m_{ij}). Examples of flow data include migration, commuting, transactions of goods, and mobile phone communication and each flow in these examples may represent number of individuals migrating, numbers of passengers per day, weight of goods, and number of phone calls respectively.

Published online: 09 May 2017

The overall goal in the analysis of flow data is to study the spatial distribution of the flows and build a (mental) representation of this distribution as a combination of spatial patterns. A spatial pattern, in this context, is an explicit or mental construct characterizing essential features of the spatial distribution of a subset of flows in a comprehensive yet frugal manner (Andrienko and Andrienko 2006). The general types of patterns sought when analysing data, according to Andrienko and Andrienko (2006), are association, differentiation, and arrangement patterns. Association patterns refer to similarity of characteristics and differentiation refer to dissimilarities.

Our focus, in this paper, lies in the exploration of arrangement patterns in flow data. Arrangement in general refers to ordering or other configuration of data items with respect to the organization of a reference set. In our case, the reference set is a two dimensional space which is organized so that its elements (i.e., the spatial locations) are connected by relations of distance and direction. Hence, arrangement of data items in space is a configuration with respect to spatial distances and directions.

Three types of spatial arrangement patterns that are commonly of interest during the exploration of flows are *spatial concentration*, *spatial trend*, and *spatial alignment*. We explicitly define these as follows:

1. **Spatial concentration**: a subset of flows with close origins and/or destinations. Such patterns emerge, for example, in areas attracting many migrants or commuters (concentration of incoming flows); in areas repelling people or areas exporting goods to many other areas (concentration of outgoing flows); in hub areas, or centres of activities (concentration of both incoming and outgoing flows).
2. **Spatial trend**: an increase or decrease of flow density and/or magnitudes along some direction. Spatial trends are displayed, for example, by an increase of commuter flow density and magnitudes in the direction from rural North to industrial South of a country; or by a decrease of tourist flows in the directions from a touristic city centre towards the outskirts.
3. **Spatial alignment**: a subset of flows forming a line. Alignment patterns are formed, for example, by flows following natural linear features (e.g., rivers, coastlines) or transportation corridors (highways, metro lines), or flows limited in their direction by natural constraints (e.g., ridges in mountain regions).

Our goal is to support the discovery and exploration of these three types of spatial arrangement patterns of flows.

Flows are commonly represented graphically in two dimensions (2D) on maps by straight or curved, directed or undirected lines or arrow symbols linking origins to destinations. The major challenge of representing flows is twofold and arises from (1) the, often, very large size of the data, and (2) from their character which can include links between thousands of locations forming large and highly connected networks that are prone to numerous crossings and overlaps. Due to this large number of intersections between the links, the underlying spatial patterns tend to be concealed by the emerging visual clutter.

To address this challenge and assist an analyst in finding spatial arrangement patterns in the data, we introduce a novel interaction interface specially designed for the exploration of flows which successively filters the data by direction. We implement the proposed interface within an interactive environment for the exploration of flows in geographical space, called *'Flowcube'*, which employs a three dimensional (3D) representation for displaying connections between locations. The main contributions of our work hence are:

- the design of a flow-specific interaction technique based on filtering flow data by direction, implemented within an interactive exploratory interface for flow data, and
- observations of the use of the interface by analysts leading to suggestions of strategies for the effective exploration of flows.

We demonstrate the functionality of our approach by exploring differences in flows of tourists and local residents in the Greater London area using data extracted from geolocated Flickr photographs.

Related Work

Flows are traditionally represented on 2D maps by band, line or arrow symbols connecting pairs of origin-destination locations (Fig. 1). An example of a directed arrow representation weighted by the magnitude of the flows can be seen in Fig. 2a where 13,986 flows are displayed. Early examples of computer generated flow maps are presented by Tobler (1987) and several representation approaches have followed (Phan et al. 2005; Rae 2009; Guo 2009; Boyandin et al. 2010; Wheeler 2015). A challenge posed by the representation of flows is that they can connect not only neighbouring locations but any locations at any distance from each other forming this way large densely connected graphs. This massive number of intersections, hence, results in representations which are very cluttered and thus hard to comprehend. Short links are often masked completely by this clutter, while long links are occluded making it nearly impossible to visually trace them and see which regions they are actually connecting. Additional mappings of relevant flow attributes to the links' visual variables can be used, such as the magnitude of the flow to width or colour, but even these are often not discernible reliably due the heavy clutter, as can be seen

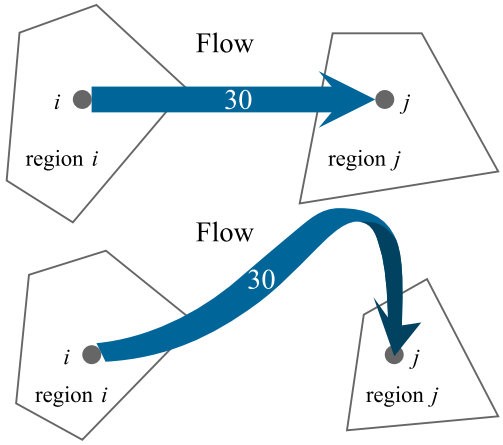

Fig. 1 Representations of flows as spatial interactons. A flow from region i and region j represented as a directed arrow (*top*) and as a directed arc (*bottom*) connecting the centroids of regions i and j and weighted by the flow magnitude which, in this case, is set to the number of transitions between regions (30)

in Fig. 2. In graph drawing, layout algorithms are commonly used to optimize the position of the nodes and thus improve the readability of the graph (Gibson et al. 2012). In flow representations, however, nodes correspond to geographical locations so re-positioning these is not a welcomed alternative. Instead, proposed approaches often involve reducing or simplifying the data or reference space to improve their representation.

A widely adopted strategy is the use of filtering and selection in order to show a reduced amount of data at a time. Single or few locations are chosen and flows to and from only these are shown, or the data are filtered by a set threshold quantity (Tobler 1987; Rae 2009; Guo 2009). An example of applying filtering in a 2D representation is shown in Fig. 2b where only flows matching a minimum magnitude threshold are displayed. Reducing the data set in this manner, however, can lead to loss of detail and of information that may be significant for the characterization of the movement. Filtering out, for example, connections that do not satisfy a certain threshold can lead to a situation where a place connected with many other places by weak links or a cluster of parallel weak links is ignored. This also leads to patterns such as spatial concentration and trend remaining concealed. Furthermore, such approaches assume that the analyst exploring the data has an initial idea of what to look for. Even though this is a reasonable assumption, it can lead to neglected locations. Aggregation of regions and flows in order to reduce the represented data size while still being able to uncover general patterns of flow has been proposed by Guo (2009). Another approach proposed by Guo and Zhu (2014) uses a kernel-based model for estimating flow density between locations and a selection and generalization approach for selecting representative flows in order to uncover generalized major flow patterns in the data. These approached, however, focus on identifying the most prominent flow patterns while potentially valuable detail is lost in the process.

We want to avoid these situations and instead facilitate exploration where an analyst can make sense of the data and detect interesting patterns without loss of information. Lens techniques are a way to overcome the drawbacks of global filtering or aggregation by providing a means to interactively select confined regions in view space in which the representation is adapted. By moving the lens the user can examine arbitrary portions of the data, while the transient nature of the tool makes it feasible to connect what is shown in it with the context information briefly covered by it. A first systematic treatment of lens techniques, or 'magic glasses', was presented by Bier et al. (1993). One application of lens tools to cluttered line data, with similar goal but different approach, is the Sampling Lens for Parallel Coordinate plots by Ellis and Dix (2006).

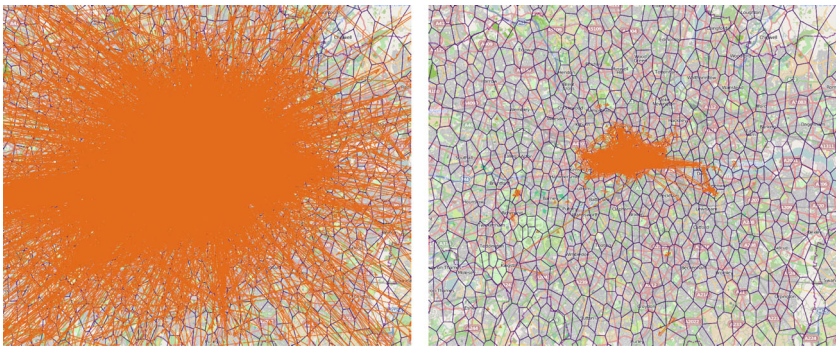

(a) Tourist flows. All flows. (b) Tourist flows with minimum magnitude of 10.

Fig. 2 Flows represented as directed arrows weighted by flow magnitude on a 2D map. The example includes 13,986 flows representing the movement of tourists between regions in London. (**a**) All flows are represented. (**b**) Flows are filtered with respect to a minimum magnitude threshold of 10. In both views the large number of cross-overs of links render cluttered displays. Short links are masked completely by this clutter, while longer links are occluded making it difficult to visually trace them and see which regions they are actually connecting

Another strategy that has been employed in the representation of flows is to merge or group edges that follow similar paths. Phan et al. (2005) propose an algorithm that uses hierarchical clustering in order to connect a single origin to its destination locations. Their method works on a reduced space since it focuses on single-source flow maps and, furthermore, the edge routing strategy used for producing the flow map layout distorts the positions of the nodes which can be confusing. Verbeek et al. (2011) propose an alternative algorithm for producing similar flow trees that does not alter the position of the nodes. This method, however, also concentrates on the creation of single-source flow map layouts.

Techniques for merging edges span much wider than flow mapping; they are a core topic in the graph drawing community, an extensive survey of which can be found in von Landesberger et al. (2011). Numerous algorithms for performing edge bundling (Holten 2006; Holten and Wijk 2009; Selassie et al. 2011; Ersoy et al. 2011; Luo et al. 2012), edge clustering (Cui et al. 2008), and edge routing (Lambert et al. 2010a, b) exist and are applied to large, dense graphs for reducing the visual clutter caused by edge crossings. These techniques are applicable to flow representations. However, displaying the entire set of flows at once still results in highly cluttered views despite the optimization achieved by merging the edges. Furthermore, such methods usually work well for representing flows from a few selected places or for flows exhibiting specific movement structures, for example a radial structure. However, since they tend to distort the spatial positioning of the represented data, it becomes harder to distinguish arrangement patterns.

Besides flow maps with linear symbols, many alternative abstract representations of flows have been suggested. Density maps have been used for revealing concentrations of flows at the cost of loss of detail concerning the links themselves (Nielsen and Hovgesen 2008; Rae 2009, 2011; Lampe and Hauser 2011). Combinations of choropleth maps, point and linear symbols, and custom made glyphs are used by Specht and Hanewinkel (2011) to analyse commuting patterns in Germany. Filtering or aggregation of the data are used, though, in order to create interpretable views. An abstract representation combining flow maps with a heatmap for representing temporal flows was proposed by Boyandin et al. (2011). In this representation, however, the spatial continuity of the flows is disrupted by the heatmap making spatial patterns hard to distinguish. Flows have also been displayed in the form of origin-destination matrices (OD matrices), where each cell represents a flow from an origin (a row) to a destination (a column) and is coloured according to the magnitude of this flow (Guo et al. 2006). Sorting and reordering of cells helps reveal flow structure and spatial interaction (Bertin 1983; Wilkinson and Friendly 2009). Overall, however, even though these displays are free from occlusions, they lack spatial context. Wood et al. (2010, 2011) address this problem by reordering the OD matrix in a space preserving manner into their Origin-Destination map (OD map). Yang et al. (2017) address the same problem through MapTrix that combines an OD matrix with a map and uses leader lines to link each row and column of the matrix with its geographic location on the map. Both these approaches reveal an overview of the flow structure but specific flow patterns are hard to follow.

In this work we use a three dimensional (3D) representation. By doing this we aim to reduce clutter which occurs by link crossings and at the same time use the third dimension for displaying additional information. 3D has been used in the past for the representation of network data (Cox et al. 1996; Munzner et al. 1996). In fact our work, discussed in the following sections, can be seen as an extension of 'Arc Maps' presented by Cox et al. (1996).

Flowcube

Flowcube is an interactive environment for the exploration of flow movement which combines interactive visualization techniques tailored for the representation and analysis of flows. Flowcube is implemented as part of a much larger geo-visual analytics framework (Andrienko et al. 2013).

Design Considerations

Within Flowcube we have chosen to represent flows as three dimensional (3D) arcs over a map, as can be seen in Fig. 3. Using a 3D representation addresses the critical issues arising from heavy clutter by (1) allowing to de-conflict link cross-overs along the vertical dimension, and (2) using this extra dimension for mapping attributes. For example, the length of link can be mapped onto the height of the flow's arc, either proportionally or inverse proportionally. This way long or short links, respectively, are accentuated while at the same time visually separating cross-overs of links with different lengths. Furthermore, the mapping of flow attributes onto visual variables of the links becomes more legible compared to the 2D version since links are often easier to distinguish in 3D.

A problem that arises by this 3D layout is that links tend to occlude each other. Using an interactive representation where the view point can be changed reduces this drawback to some degree since it allows the user to navigate around in the representation and look at the links from different directions. This does not yet, however, help to gain insight into dense "haystacks" of tightly grouped links of similar height, as can be observed in Fig. 3a,d. Hence, to better address this drawback we introduce, in this work, a novel interaction

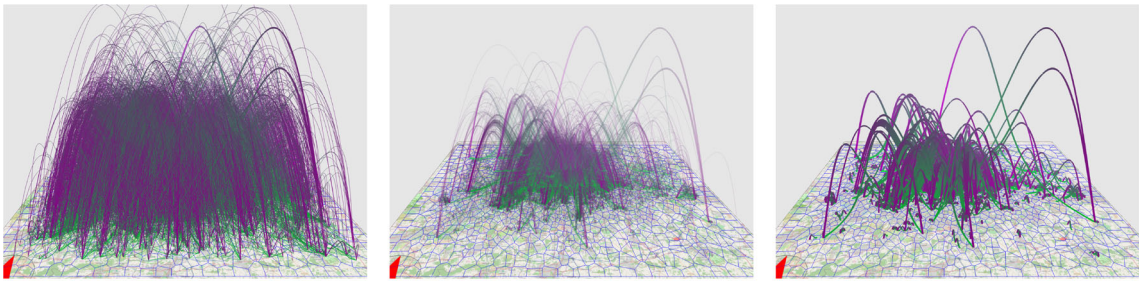

(a) Locals. All flows. (b) Flow magnitude mapped to link opacity. (c) Flows with magnitude 5 or more.

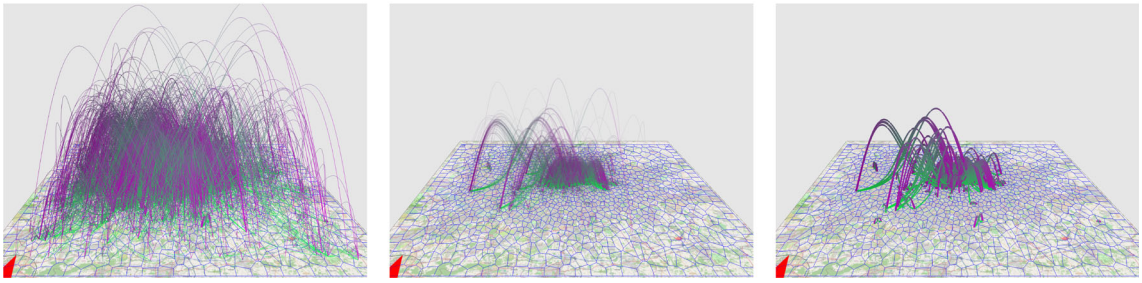

(d) Tourists. All flows. (e) Flow magnitude mapped to link opacity. (f) Flows with magnitude 5 or more.

Fig. 3 Three dimensional representations of flows as 3D arcs over a map in Flowcube. Movements of individuals in the Greater London area are displayed. Flows of locals are shown on the top row of the figure and of tourists on the bottom row. To the left the flow data are displayed without any manipulation. In the middle the magnitude of the flows is mapped to the opacity of the links revealing general concentration patterns. To the right the data are filtered with respect to a minimum flow magnitude threshold revealing the same general spatial concentration patterns

technique that has been developed to specifically match the structural characteristics of a flow data set. This technique is described in "Direction-Based Filtering". A closer description of the Flowcube environment shall be given first.

Input Data Characteristics

Flowcube is designed to operate on flow data given in origin-destination (OD) link format. The locations constituting the graph nodes of such a data set are typically actual geographic regions, either pre-defined administrative regions such as states, municipalities, and boroughs; or they can be computed regions such as cells of a spatial tessellation. Conceptually, however, regions can comprise the partition of any other, potentially abstract two-dimensional space.

Note that in principle, point coordinates of link start and end points suffice; but since flows are the result of (spatial) aggregation of interactions, points nonetheless imply a spatial partitioning.

Complementary information characterizing each region or link, such as areas' number of visits or unique visitors and flows' average movement speed respectively, can also be input into Flowcube. Such information can then be mapped onto different visual attributes in the representation in order to simultaneously visualize selected aspects of the data.

Representation

Flowcube includes a three dimensional display for representing flows and an interface for altering aspects of the representation. The main representation is composed of three elements: (1) a two dimensional (2D) base map, (2) a representation of the region boundaries, and (3) the representation of flows as links between these regions (Fig. 3).

Base Map This element is essential to communicate the spatial frame of reference. For geographic maps, we use OpenStreetMap (Haklay and Weber 2008) as a base map in Flowcube since it is free and easily accessible. Depending on the spatial extent of the currently explored territory, the appropriate map tiles, in terms of level of detail (LOD), are downloaded and included in the representation. This implies that if the flow data set covers a small area, such as a city, then a map displaying street level LOD will be downloaded, if instead the flows stretch across an entire country then a less detailed map will be used. However, the user can always override the automatic LOD selection of the displayed map.

Region Boundaries The base map is overlaid with a contour representation of the regions. The region boundaries are drawn according to the geometry defined in the input data. If geographical point locations are provided as input

instead of region geometries, we apply Delaunay triangulation (Delaunay 1934) to retrieve a contour representation composed of Voronoi cells (Voronoi 1908). Links could, of course, be drawn between points instead of regions. However, points are harder to interactively select, especially in a 3D representation as Flowcube, we therefore prefer the use of regions.

Flow Representation The core idea of Flowcube is to utilize the vertical dimension to reduce the inevitable crossover clutter between links. Thus, flows are represented using cubic splines (Bartels et al. 1987) linking connected regions (Fig. 1 bottom). The user can choose between line or band representations for these links. In order to better distinguish between their start and end points, similar to Fekete et al. (2003), the curves are asymmetric having a mild slope at their start and a steep one at their end (Fig. 1 bottom). By default, the height of each curve is proportional to the distance between the places it is connecting, so that longer links are drawn higher. This can, however, be adjusted by the user interactively to e.g. draw the links inversely proportional to their length instead, or to map another attribute onto the height of the representation. In addition to spline arcs, the user can also switch to straight lines. Flows with identical origin and destination, are not represented in Flowcube.

We use colour, width and opacity of the links to encode additional information (Fig. 3). The user can select attributes, such as magnitude and length of flow, to map onto the representation of the links. In the default representation, colour is used to signify the start (green) and end (purple) points of the curve, and the magnitude of each flow is mapped onto the width. Such mapping of attributes is an important feature of the representation since it gives flexibility to the analyst to uncover flow patterns.

Finally, as mentioned earlier, Flowcube is integrated with a larger visual analytics framework. All views in the framework are linked and therefore data depicted in Flowcube can also be seen in any of the available displays, including a traditional 2D flow representation of arrows on a map (Fig. 2). This can provide the user with (1) an additional, alternative view if needed, and (2) allows to synchronise Flowcube with further, e.g. data-driven, filters active on other displays of the system. Vice versa, it allows to interactively feed back findings made in Flowcube to other displays.

Standard Interaction

Flowcube offers flexibility of exploration to the user. A user can interact with the representation using standard techniques such as panning, zooming, and rotation. Rotation is performed around any point of interest selected with the mouse. Several selection alternatives are provided in order to allow querying of the data and retrieval of details on demand. A user can select regions by clicking on their corresponding cells on the map. Upon selection of a region a user can choose to view it as a *destination* by showing only links connecting into the selected region (in-links, Fig. 4a), as an *origin* by showing only links connecting out of the selected region (out-links, Fig. 4b), or as both by showing simultaneously in- and out-links (Fig. 4c). Single (Fig. 7a) or multiple (Fig. 7d) regions can be selected. When multiple regions are selected the in- and/or out-links connecting to all the selected regions are displayed. Furthermore, there is an option available that allows the user to 'follow the path' of the flows instead. This means that upon selecting multiple regions the order of selection is considered and only links connecting regions $i \rightarrow i+1$, corresponding to the selection order, are displayed. In this manner the user can concentrate and inspect the magnitude of specified flow paths. Moreover, upon selection it is possible to look at a close-up of the selected area in a separate 2D map view available in the framework. This functionality is included as an alternative to zooming into the 3D Flowcube representation and inspecting the base map, which is of course also possible. Having this additional option can provide smoother

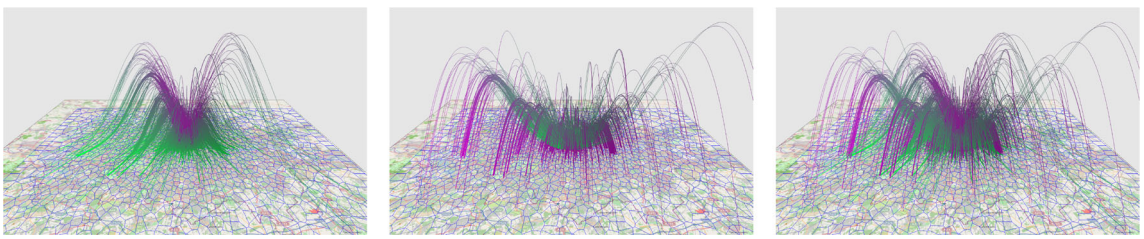

(a) In-links. (b) Out-links. (c) All links.

Fig. 4 Examples of selecting regions in Flowcube. Movement of tourists in the Greater London is displayed with central regions selected. (**a**) Links having a destination that coincides with the selected regions are displayed. (**b**) Links having an origin that coincides with the selected regions are displayed. (**c**) Links having an origin and/or destination that coincide with the selected regions are displayed

exploration since the orientation of the representation does not have to be altered in order to inspect the geographical details of a selected region.

Finally, utilizing the host framework, filtering the displayed data by any number of available attributes, including temporal references, is also possible. This allows users to reduce the amount of visible data if they so wish. Figure 3c, f show examples of filtering the data by the number of transitions between regions. At the same time, the current link and/or region selection can be propagated back to the framework as new filter criterion to serve as a starting point of further detail anaylsis using complementatry visualization methods.

Direction-Based Filtering

Apart from standard interaction available in Flowcube, we introduce in this work a flow-specific interaction technique based on directional filtering for systematically scanning through a flow data set. We will refer to this interaction technique as 'direction-based filtering'.

The inherent geometric structure of a flow representation is linear. As a consequence, and as discussed earlier, when analysing flows the spatial patterns and relationships that are of interest are provided by the positions of the regions and the directions of the pairwise links connecting them. Moreover, flow data sets are often large and can display links between most regions, which renders their depiction using connecting lines very cluttered, even when utilizing the third (vertical) dimension, as can be seen in Figs. 2a and 3. These facts have been the starting point of our design.

We take advantage of the geometric properties of the link representations and create a band-shaped lens, the 'Direction-Based Filtering Lens' (DBF lens) (Fig. 5), to match this geometry and enable direction-based filtering of the flows. The lens operates as a rectangular slice in view space, similar to clipping planes often used in scientific visualization, bringing into focus a subset of links at a time.

DBF Lens Design

The DBF lens is represented as a horizontal, rectangular band vertically positioned in the current Flowcube view (Fig. 5). This rectangle in view space is projected onto the 2D base map (using the inverse modelview-projection matrix), thereby defining a selection area oriented perpendicular to the view axis. The link set is then filtered to show only those links that are completely contained in the area spanned by the DBF lens, i.e., when both their start and end location fall within the DBF area. Therefore with an active DBF lens only links that are parallel to each other to a certain degree are displayed. Links outside of the area as well as longer perpendicular links which would cause occlusion and clutter are culled, making it easier to visually trace the selected links than in unfiltered flow representations.

The size and position of the DBF lens, i.e. its height and its vertical position in view space can be interactively adjusted by the user. This way the amount of data being displayed can be altered during the course of the exploration. A narrow DBF lens can be used to isolate few regions and explore only the nearly perfectly parallel links connecting them, while a broader lens can reveal links between a larger number of neighbouring regions with a higher variation in link directions. The principal filter orientation can be quickly and intuitively changed by rotating the map underneath the DBF lens in the 3D view.

Furthermore, it is possible to constrain the DBF lens to display only flows following the same principal direction; specifically, to only show flows moving from right to left in the lens. Depending then on the orientation of the map under the projected lens area and the projected area's width, only links more or less aligned with the principal direction are visible. An example can be seen in Fig. 8b which displays

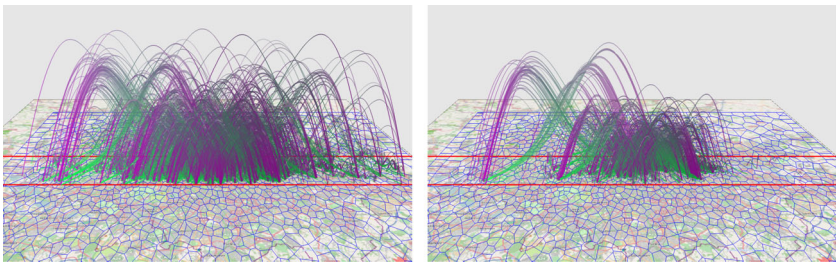

(a) Locals. (b) Tourists.

Fig. 5 Flows of locals vs. tourists are queried along an east/west direction across central London using the DBF lens. (**a**) The flows of locals are spread across most of the region defined by the DBF lens and display dense concentration and alignment patterns. (**b**) The flows of tourists display a more narrow concentration in the central regions of London and a clear trend to/from the west of the city towards Heathrow airport

flows going in the east-southeast direction within the DBF lens.

To facilitate smooth transitions between explored flow slices while interacting with the representation, the displayed links are drawn with varying opacity depending on their location within the DBF lens. A link's opacity is determined by the average opacity of its starting and ending point. Each point's opacity, in turn, depends on its distance from the DBF lens' horizontal axis of symmetry C (Fig. 6). Points lying directly on C evaluate to maximum opacity, with opacities being gradually modulated towards fully transparent for points sitting on either of the DBF lens' edges (i.e, falling on either the top or bottom boundaries of the projected lens rectangle). Hence, a link is most visible when centred in, and perfectly aligned with, the DBF lens. By contrast, links closer to the lens' borders or with deviating orientations are mostly transparent. Interpolation from C to the lens edges uses linear interpolation (Fig. 6) as we found it resulting in a sharper distinction of links centred in the DBF lens with closely matching orientation from less concisely matching links. Experiments were also made using a Hermite interpolation (Spitzbart 1960) function $f(t) = 3t^2 - 2t^3$, with $t = ||dist(x, C)|| \in [0, 1]$ resulting in a smoother but also less discriminating selection with respect to distance and orientation to C.

DBF Lens Interaction

Interaction using the DBF lens differs from that of a traditional lens. Traditionally a lens has a position centred under the mouse pointer and is applied to the data as the user mouses over the display. The DBF lens, instead, has a fixed position on the screen. The data are brought into focus when they intersect with the band area defining the lens as the user interacts with the representation of the flows, by rotating, translating or zooming the display. Instead of rotating and aligning the lens across different directions, the map with data is aligned under the lens as the user interacts with the display. We have chosen this design layout in order to maintain the 'natural' interaction of the user with the geospatial representation of the flows. We believe that this way the context of the represented movement is better preserved.

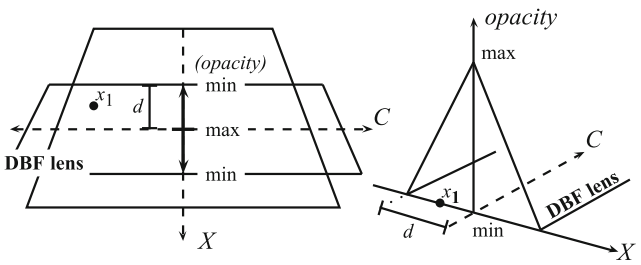

Fig. 6 Opacity value computation within DBF lens. A link's opacity is set equal to the average opacity of its start and end point. Each point's opacity depends on its distance from the DBF lens' horizontal axis of symmetry C

Apart from allowing the user to freely translate the representation through mouse interaction, we have made it possible to 'scroll' the map back and forth along the view direction using the 'up' and 'down' keys, making it easier for all regions of the territory to enter and be explored in the DBF lens in the same direction. It is also possible to initiate an animation which automatically scrolls the representation 'back and forth'; the user can decide the starting direction of the animation, and initiate or terminate it at any time. This allows the user to systematically slice and inspect the data with respect to a specific common flow direction.

The user can control the map orientation by interactively selecting an arbitrary pivot point on the map plane with the mouse and then using either the mouse or the 'left' and 'right' keys to rotate the representation in-plane around this chosen pivot point. Similarly to map translation and scrolling, rotation around the last selected pivot point can also be performed by an automated animation.

Implementation Notes

The DBF lens implementation utilizes the programmable OpenGL pipeline to perform updates of the opacity-interpolated link selection at interactive rates even for very large link sets. It utilizes the fact that link selection is determined by the DBF lens' (linear) projection onto the map plane, so instead of using clipping planes in 3D it performs fast half-plane inclusion queries on a spatial index structure. Upon rendering, the simple line segments used to approximate the spline arc are first processed by a geometry shader to generate volumetric lines as proposed by Hillaire (2012), considering values of the attribute mapped to link width. The link arcs' height mapping is evaluated during vertex processing, link colour and opacity mappings during the fragment processing stage on a per-fragment basis, thus utilizing built-in GPU hardware-accelerated interpolation functionality. Even on a mobile GeForce GTX 670M-equipped system with 1.5GB of VRAM, this allows to render the data set with 45,016 links, used in our example exploration, and active DBF lens with frame rates in excess of 35fps while rotating the map.

DBF Discussion

Direction-based filtering is a straightforward approach that allows the user to view the data a few at a time while follow-

ing the direction of the flows. It's implementation through the DBF lens is novel and effective in that the lens' geometry matches the spatial relationship of analytic interest in the analysis of flows, meaning the pairwise connections between regions.

The DBF lens makes it possible to uncover arrangement patterns of interest in the following manner:

- A *spatial concentration* pattern is exhibited as an area of increased density of flows within the lens. By rotating the view around the centre of the area, the extent of the spatial concentration in different directions can be investigated.
- A *spatial trend* pattern is exhibited as increasing/decreasing flow density in the direction to the left and/or to the right of some position within the lens. By rotating the view around this position, the user investigates in which spatial directions the trend takes place and finds the directions along which the trend is the most prominent.
- A *spatial alignment* pattern is exhibited as multiple links with close spatial directions covering a large part of the lens. The effect disappears as the view is turned in a different direction approximately around the position of the "mass centre" of the links.

Furthermore, defining a band-shaped lens and bringing subsets of flows into focus successively with smooth transitions preserves the context between views. Finally, the fact that interaction in the DBF lens involves the manipulation of the map offers a continuous filtering effect as flows appear and disappear naturally as they enter the lens. Hence, by using the DBF lens in combination with the additional attribute mapping functionality of Flowcube, a user can systematically explore flows in a manner that does not disrupt their normal interaction with the representation.

Exploration and Analysis Example

We demonstrate our direction-based filtering approach using a data set of individuals' flows in the Greater London area, extracted from the photo-sharing website Flickr.[1] We explore the flows of locals vs. tourists in the area and identify interesting patterns within and between the groups. The following section describes how we extracted flow movement data from the set of geo-referenced photographs and distinguished between movement of local residents and tourists. "Analysis" then discusses the analysis process we applied to this data set.

[1] http://www.flickr.com

Retrieval of Flow Data

Photographs shared on Flickr are often enriched with metadata including when (timestamp) and where (geographical coordinates) a particular photo was taken. This information is typically appended to the photo file's EXIF metadata section directly by the camera; but it can also be specified manually by the user upon uploading a photograph to Flickr. The public Flickr API allows a crawling operation through the service's database of photographs made publicly available by its users, as well as their public 'friend lists'. In our crawler implementation, we seed the crawl with manually selected users, obtain information such as user and photo id, date of capture, and geographical coordinates of their published photos, then recursively crawl users on their friend list in the same way. For our example we have extracted a data set consisting of 1,339,817 photographs taken by 32,292 individuals in the Greater London area between 2005 and 2010.

Using the timestamp information available, the photos from each user are brought into temporal sequence that in turn gives a sequence of photo locations, which we take as the user's movement trajectory. This trajectory is not continuous since the actual path followed by a person between taking photographs is unknown. This data can therefore not be used to infer detailed information on individual persons' movement, but can be used in aggregated form to analyse the general flow patterns of people across the cityscape (Jankowski et al. 2010).

We further assume that these patterns differ between local residents and visiting tourists. To be able to compare the two, we used the following simple strategy to distinguish each user as either a local or a tourist: if there was a gap of 365 days or more between two consecutive photographs, or the entire trajectory duration was less than 14 days inside the Greater London area the corresponding Flickr user is considered a visiting tourist; otherwise, a local resident. After having split the data, we preprocessed both sets further to avoid trajectories with unrealistically long durations or long gaps between positions; thus trajectories with a temporal gap of more than 7 days between consecutive positions were split up at those gaps. Preprocessing resulted in a data set of 30,008 tourist, and 102,863 local resident trajectories.

The final step of data preparation comprised extraction of flows from the trajectories using the method for spatio-temporal aggregation proposed by Andrienko and Andrienko (2011). Characteristic trajectory points are extracted and clustered first. Cluster medoids are then used as seeds for a Voronoi tessellation of geographic space, reflecting the spatial distribution of the photos. Clustering and tessellation were performed on all trajectories together to avoid biasing spatial aggregation towards either

of the two subsets. Afterward, trajectories of each set were separately aggregated into flows between pairs of regions (Voronoi cells) by counting the number of transitions between them. The resulting flow data set comprised 1,367 regions connected, respectively, by 13,986 flows in the tourist and 45,016 flows in the resident subsets.

Analysis

The two data sets containing flows of locals and tourists in the Greater London area are displayed in Flowcube (Fig. 3). The default view is cluttered with links, very much resembling a haystack (Fig. 3a,d). Mapping the number of transitions between regions to the opacity value of the links (Fig. 3b,e) or filtering the data by a transition count threshold (Fig. 3c,f) reveals more of the general structure of the movement. The views reveal a very wide *spatial concentration* in the central parts of London, with tourists displaying a narrower *concentration* pattern compared to the locals. Some smaller *spatial concentration* patterns can be identified, revealed by major flows between regions on the outskirts which, at this point, can only be seen since they are outside the blob of data crowding the central regions of the map. For example, the visible flows between two regions to the south in Fig. 3f represent movement in Southeast Croydon where there is a concentration of golf clubs, parks, and open spaces. Any detailed patterns, however, are hard to discern. The main general conclusion drawn from these first data views is that the movement of tourists appears to be concentrated in the central areas while locals' movement is more spread across the entire Greater London area.

In order to get a better overview of the data, we activate the DBF lens and start slicing our way through the haystack in search of flow patterns. Comparing the movement of the two groups along an east/west direction across the central areas of London (Fig. 5), we notice a distinctive *spatial alignment* pattern in both groups. The movement of locals almost fills the entire band area (Fig. 5a). The tourists, on the other hand, show a clear dense *concentration* pattern in the centre and also a distinctive *spatial trend* towards regions in the west (Fig. 5b). Selecting these regions and zooming to inspect the map reveals that Heathrow airport is situated there (Fig. 7a).

Continuing the exploration of the tourist data set, interesting patterns along different directions are revealed. Rotation around central London reveals an interesting *alignment* pattern of tourist flows from the centre towards Wembley stadium in the northwest, and Greenwich park in the southeast (Fig. 7b). Such flows, originating and stretching not far from the centre, are hard to explore using conventional flow representations since they are obscured by longer flows. Furthermore, these flows are composed of a number of parallel links starting from several central regions, which may not have been considered as significant if the data were filtered by a minimum threshold. An interesting *spatial trend* pattern of tourists is detected starting from the centre and following a southwest direction (Fig. 7c). The flow reveals an initially quite dense movement to Kingston and Hampton Court Palace and then continues further in the same southwest direction towards the nature reserve areas Wisley Common, Ockham and Chatley Heath with lower density (Fig. 7c). Selecting the areas of Kingston and Hampton

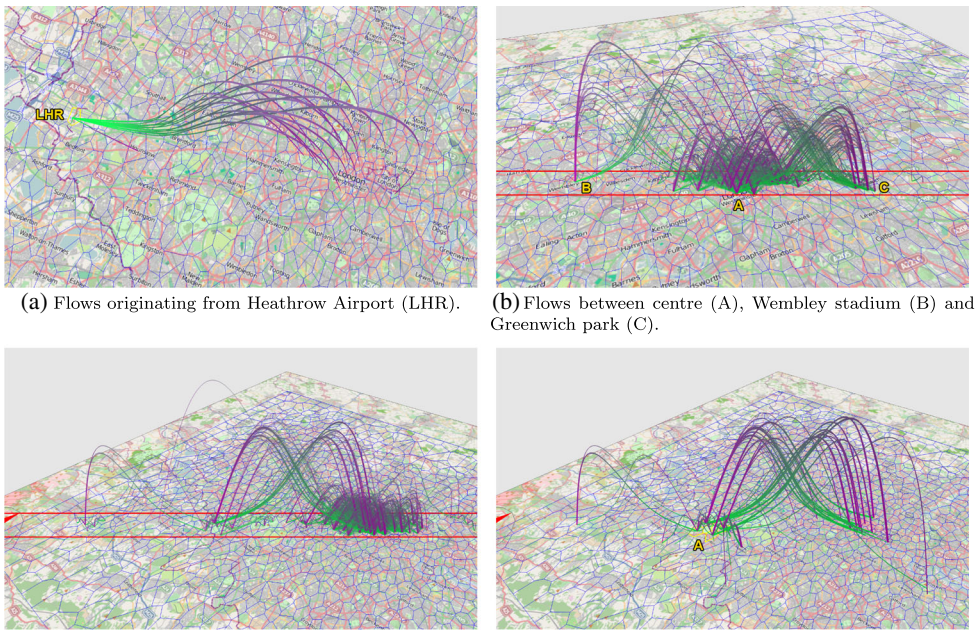

(a) Flows originating from Heathrow Airport (LHR).

(b) Flows between centre (A), Wembley stadium (B) and Greenwich park (C).

(c) Flows to/from city centre (to the right) along southwest direction (left).

(d) Flows to/from Kingston and Hampton Court (selected regions at (A)).

Fig. 7 Flowcube exploration examples of tourists' movements in Greater London. (**a**) An alignment flow pattern between Heathrow airport and neighbouring regions in the city centre. (**b**) Flows aligned along a northwest/southeast direction displaying a concentration in the centre a trends along the two directions. (**c**) A spatial trend pattern along a southwest direction starting with a dense concentration of flows in the city centre going through Kingston and Hampton Court to Wisley Common, Ockham and Chatley Heath. (**d**) Example of selecting a region of interest and displaying all flows to/from it. In this case Kingston and Hampton Court are selected which are part of the trend of Fig. 7(c)

Court Palace and turning off the DBF lens reveals all links to/from these locations and further confirms the identified *trend* (Fig. 7d).

Quite different patterns can be observed in the movement of locals. Since the places involved in these patterns are not widely known, we interpreted and verified the findings by searching for place-related information on the Web and by consulting a London resident. Moving the Flowcube representation back and forth and bringing the regions from north to south into focus in the DBF lens reveals distinct *spatial concentration* in the outskirts surrounding the city centre of London. An example of two such patterns in the north, which are also connected to each other by a minor flow, can be seen in Fig. 8a. These are the boroughs of Watford in northwest and Barnet and Enfield in the north, which are urban residential areas. Rotation of the map around the centre while looking at flows directed from right to left in the DBF lens reveals an *alignment* pattern of flows and a *trend* from several central areas towards a group of regions on the east (Fig. 8b). Selection of these regions on the map reveals that there are flows to/from many different parts of the city to this area (Fig. 8c). Closer inspection of the selected regions shows that most flows arrive to the large 'Bluewater' shopping centre, which attracts shopping activities from a wide area and is well known to the locals, both Londoners and commuters. Using the same exploration approach to identify *concentration* patterns, we have discovered another interesting destination for locals: Brands Hatch, a popular race circuit in the southeast outskirts of London, which often hosts Grand Prix and Formula 2 and 3 events. Not surprisingly, this destination is connected to various places across the Greater London area.

Even though the exploration resulted in numerous findings, we have included only a few examples that reflect the overall results, due to lack of space. The type of flows and pattern characteristics detected in the two groups have been distinctively different. In summary, tourists tend to constrain their movement to the centre and their flow patterns arrive mostly to popular tourist and leisure attractions. So, spatial patterns emerging from the tourist data set consist of dense *concentration* of strong flows in the central regions of London, and display strong *trend* patterns, composed of large magnitude flows, between a smaller number of popular destinations. Such patterns would, therefore, also be identifiable by filtering the data by a relative high threshold. Even though the tourists' flow patterns are not surprising, the fact that these are easily and successfully detected using Flowcube is a positive proof of concept. The exploration of locals, on the other hand, has revealed an overall movement pattern that spans across large parts of the Greater London area. *Spatial concentration* patterns in the areas surrounding London have been detected, indicating popular residential areas, and also several non-tourist destinations have been identified. The locals' flow patterns involve to great extent movement to and from neighbouring areas which result in a larger number of roughly parallel flows. These *alignment*

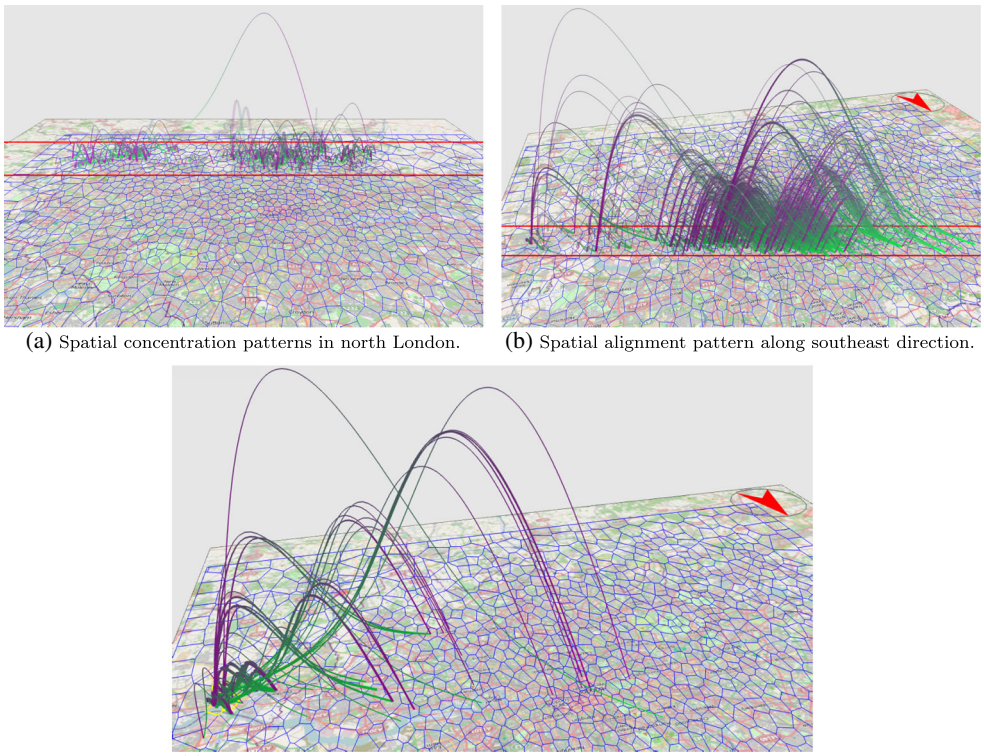

(a) Spatial concentration patterns in north London.

(b) Spatial alignment pattern along southeast direction.

(c) Flows to/from selected regions of spatial concentration in east London.

Fig. 8 Examples of locals' spatial patterns in Greater London identified through exploration in Flowcube. (**a**) Spatial concentration patterns detected in the northern outskirts and identified as urban residential areas. (**b**) A spatial alignment pattern along a southeast direction displayed a high concentration of flows in the centre and a clear trend pattern towards regions farther east. (**c**) The regions of interest identified in the trend pattern (in Fig. 8(b)) are selected and all links to/from them are displayed. A closer inspection of the selected regions of interest (in Fig. 8(c)) on the map reveals that the flows arrive to the Bluewater shopping centre

patterns would potentially be neglected by the use of filtering but are distinctively revealed using the DBF lens. Finally a general observation, made using the DBF lens, about the overall movement in London is that flows tend to be mostly aligned in an east/west direction and less in a north/south one. This could be an effect of the river Thames which flows across London and defines the city structure.

User Feedback and Exploration Strategies

In order to collect feedback and improve the functionality of our approach we conducted 'think aloud' exploration sessions with potential users and followed them up with semi-structured interviews. Our main objective was to investigate the usefulness of direction-based filtering within a 3D flow representation. A secondary objective was to assess whether the interaction mode suggested for filtering with the DBF lens, namely manipulating the map instead of the lens, was in fact a welcome alternative.

Exploration Session Setup

Four individuals participated in the 'think aloud' sessions, in agreement to the recommendation by Nielsen (1994) of 4 ± 1 subjects for such studies. The participants were two male and two female, all of whom were employees or students at a university. All participants were experienced in interactive visualization and visual analysis of data, but none of them had previous experience in using the Flowcube environment and DBF lens. The participants performed two subsequent exploration sessions each. The two data sets described in "Retrieval of Flow Data" were used for each of these. Two participants (one male and one female) used the data set of locals for the first session and the one of tourists for the second one, while the other two participants explored the data sets in reverse order.

The participants were first introduced to our flow exploration environment, which was done by explaining and demonstrating the representations, interaction and attribute mapping functionality which was available. Thereafter, the first exploration session commenced. During this part they were asked to freely explore the flow data, as they would in a normal exploratory analysis situation, while an observer was taking notes. The participants were encouraged to continuously describe what they were doing or were trying to do and to communicate their hypotheses and findings. After the exploration had proceeded for 10 min the participants were interrupted and were asked a set of pre-defined questions regarding aspects of the representation and interaction design and following this they were invited to give additional feedback they saw fit.

After the first exploration session completed the participants were introduced to the direction-based filtering approach for exploring the data and were orally instructed on how to use the DBF lens. They then performed a second exploration session of a different data set having this time access also to the DBF lens. Also during this session participants were allowed to freely explore the data for 10 min and were encouraged to continuously 'think aloud' while they were being observed. Following this, a second short semi-structured interview took place where questions concerning the new available filtering and the interaction mode were asked. Again participants were encouraged to leave any additional feedback they had to give. Altogether participation lasted between 45–60 min.

Feedback and Improvements

Invaluable feedback was gathered through this evaluative process both concerning the interaction approach itself, and the interface and exploration environment. The overall feedback was positive and several ideas and suggestions for improvements were established, some of which were addressed immediately and some are planned as future work. All participants found it difficult to detect useful knowledge in the default flow representation (Fig. 3a,d) due to the large number of links and crossings between them. Applying various sorts of data manipulation, such as mapping attributes to the size and opacity of the links and filtering out values not meeting certain thresholds made the display manageable. This however occurred at the cost of less prominent flows disappearing entirely from the screen. The use of the DBF lens made the display manageable without needing to use high filtering thresholds was among the comments from the participants. Other comments on the approach included that the DBF lens lets you detect structure even in very densely connected regions and that neighbouring regions connected by many weak links remained visible. There were positive comments about the fact that links within the DBF lens were faded in/out gradually while interacting with the view instead of instantly appearing/disappearing. Participants liked the fact that they could be 'prepared' for what was 'coming and going'.

Some informal feedback was already collected during prototype development. An explicit request was that map interaction should optionally be performed using key strokes in addition to the mouse, due to the former offering more precise rotation control constrained to a specific axis when the DBF lens is active, while the latter was preferred for general view navigation due to its ease and speed. Another addition is the integration of automatic animation around a user-selected pivot point (see Direction-Based Filtering) for systematic direction scanning around an area of interest, which was also directly asked for in the initial prototype.

Exploration Strategies

During the exploration sessions with the users and during our own experimentation with direction-based filtering within Flowcube we observed that some exploration strategies were more or less typically employed by everyone.

Before introducing the DBF lens, the general strategy followed by all participants was to first use filtering to reduce the amount data and identify areas connected by strong links and then move on to get details about these areas by selecting them and exploring their connections separately. There was very little interaction with the map, the view remained mostly static during the exploration.

After introducing the DBF lens the exploration strategy was somewhat altered, not so much with regard to the intention of the actions performed, but with respect to the execution and result. No initial thresholding was applied to reduce the amount of visible data. Instead, the strategy adopted for exploring flows was to switch on the DBF lens and translate the map along different directions in order to identify areas of high flow concentration. Once identified, rotation around the central points of these areas was performed in order to explore their connections. Following this, when interesting flows were established, the common strategy was again to select regions of interest and inspect their connections to the entire territory by de-activating the DBF lens. When satisfied with the inspection the DBF lens was re-activated and the exploration continued.

The difference between the two approaches is that without the DBF lens only regions connected by prominent links, i.e. flows with large magnitudes and thus high weights, were inspected. This strategy of inspecting only major links proved to be suboptimal as strong connections composed from multiple less prominent flows could be missed. Using the DBF lens instead resulted in more regions being explored, since neighbouring regions connected by many weaker links could also be identified as interesting. One such incident, for example, was the spatial concentration pattern around the boroughs of Watford and Barnet & Enfield depicted in Fig. 8a.

Discussion and Conclusions

We have presented an interactive environment for exploring flows, called Flowcube, in which a three dimensional arc representation of links is used. Using a 3D representation reduces clutter occurring from link crossings but this is partly replaced by clutter due to occlusion of the links. In order to overcome this occlusion and focus on the analysis objective, we have proposed a novel direction-based filtering approach as an effective interaction alternative and we implement it through the DBF lens.

We adopted a task-oriented approach in the design of our visualization and interaction techniques. We first identified the spatial patterns commonly sought in the analysis of flows as arrangement patterns and considered the possibilities and shortcomings of traditional approaches for detecting them. We then developed our proposed techniques so that these types of patterns were preserved and could be uncovered. As a result, the geometry of the proposed DBF lens closely follows the spatial patterns of interest in the analysis of flow data sets. By bringing into focus flows that follow the same direction parallel links are accentuated in the representation revealing this way spatial alignment and trend patterns. Filtering by direction also helps the analyst concentrate on their analysis task by removing clutter caused by crossing links.

The analysis can be further enhanced by such features as allowing the mapping of attributes to the appearance of the links. Particularly, mapping the magnitude of flows to the opacity results in parallel links becoming even more prominent. More flow patterns are accentuated this way: both alignment patterns created by a few strong flows, and also patterns composed of several weaker ones. Filtering out data that don't meet a certain threshold would neglect these latter patterns and possibly so would aggregation, depending on the level of detail at which regions are merged together. Aggregation would also be subject to the fact that the observed patterns may radically change depending on how smaller units are aggregated into larger ones is well-known in geographic analysis as the Modifiable Areal Unit Problem (MAUP) (Openshaw 1984).

A limitation of the proposed DBF lens is that it only accentuates arrangement patterns along the directions that are explored by the user. This implies that if the user does not rotate the map to inspect the flows from multiple directions then patterns can be missed. In order to aid the exploration with regard to this limitation, a feature that is included in the proposed approach is that arbitrary centres of rotation can be chosen in the representation which means that spatial concentrations of flow of any origin can be explored along all directions using the DBF lens. Another limitation of the DBF lens is that the identification of patterns is subject to the chosen size of the lens. Choosing a very narrow lens will only uncover few and strictly parallel flows while choosing the lens too broad may allow crossing links to clutter the view and obstruct the identification of arrangement patterns. The choice of lens size depends on the structure and appearance of the data to be explored and there are therefore no optimal settings. Consequently, the proposed approach, as probably any analytic task, relies largely on the expertise and observation skills of the analyst performing the exploration. This can be seen as an advantage with respect to the flexibility of analysis but is also a limitation since it does not guarantee completeness of results.

In future work we, therefore, intend to investigate the potential of aiding the analysis process by guiding the exploration and suggesting candidate regions for closer investigation. To achieve this, we plan to take advantage of the graph structure of our data and make use of the plethora of graph measures that can be computed for assessing the strength and importance of nodes (regions) and links (Vrotsou et al. 2011). Regions identified as interesting through such computations will then be highlighted in the representation as suggestions for potential exploration starting points. Finally, we also intend to consider temporal dynamics of movement by exploring interactive methods for representing and analysing flows over time within Flowcube.

Acknowledgements This work was supported in part by EU in projects VaVeL (grant agreement 688380), SoBigData (grant agreement 654024) and BigData4ATM (grant agreement 699260) and by DFG in Priority Research Programme "Volunteered Geographic Information: Interpretation, Visualisation and Social Computing" (SPP 1894).

Compliance with Ethical Standards The submitted research is in accordance with the ethical standards and the ethical responsibilities of the authors as stated by Springer.

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Ethical Approval The presented research has not involved human subjects, human material, or human data and therefore ethical approval has not been required.

Informed Consent Informed consent was obtained from all individual participants included in the study.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

Andrienko N, Andrienko G (2006) Exploratory analysis of spatial and temporal data: a systematic approach. Springer, Berlin

Andrienko N, Andrienko G (2011) Spatial Generalization and Aggregation of Massive Movement Data. IEEE Trans Vis Comput Graph 17(2):205–219. doi:10.1109/TVCG.2010.44

Andrienko G, Andrienko N, Bak P, Keim D, Wrobel S (2013) Visual analytics of movement. Springer, Berlin

Bartels RH, Beatty JC, Barsky BA (1987) An introduction to splines for use in computer graphics and geometric modeling. Morgan Kaufmann, San Mateo

Bertin J (1983) Semiology of graphics: diagrams, networks, maps. University of Wisconsin Press, Madison

Bier E, Stone M, Pier K, Buxton W, DeRose T (1993) Toolglass and magic lenses: the see-through interface. Comput Graph 27:73–80

Boyandin I (2013) Visualization of temporal origin-destination data. PhD thesis, Department of Informatics, University of Fribourg, Fribourg

Boyandin I, Bertini E, Lalanne D (2010) Using flow maps to explore migrations over time. In: Geospatial visual analytics workshop in 13th AGILE international conference on geographic information science, pp 1–4

Boyandin I, Bertini E, Bak P, Lalanne D (2011) Flowstrates: an approach for visual exploration of temporal origin-destination data. Comput Graph Forum 30(3):971–980. doi:10.1111/j.1467-8659.2011.01946.x

Cox KC, Eick SG, He T (1996) 3D geographic network displays. SIGMOD Record 24(4)

Cui W, Zhou H, Qu H, Wong PC, Li X (2008) Geometry-based edge clustering for graph visualization. IEEE Trans Vis Comput Graph 14(6):1277–1284. doi:10.1109/TVCG.2008.135

Delaunay BN (1934) Sur la sphère vide. Bull Acad Sci USSR (VII) Classe Sci Mat 6:793–800

Ellis G, Dix A (2006) Enabling automatic clutter reduction in parallel coordinate plots. IEEE Trans Vis Comput Graph 12:717–724

Ersoy O, Hurter C, Paulovich F, Cantareiro G, Telea A (2011) Skeleton-based edge bundling for graph visualization. IEEE Trans Vis Comput Graph 17(12):2364–2373. doi:10.1109/TVCG.2011.233

Fekete JD, Wang D, Dang N, Aris A, Plaisant C (2003) Overlaying graph links on Treemaps. In: IEEE symposium on information visualization poster compendium, pp 82–83

Gibson H, Faith J, Vickers P (2012) A survey of two-dimensional graph layout techniques for information visualisation. Inf Vis. doi:10.1177/1473871612455749

Guo D (2009) Flow mapping and multivariate visualization of large spatial interaction data. IEEE Trans Vis Comput Graph 15(6):1041–108. doi:10.1109/TVCG.2009.143

Guo D, Zhu X (2014) Origin-destination flow data smoothing and mapping. IEEE Trans Vis Comput Graph 20(12):2043–2052. doi:10.1109/TVCG.2014.2346271

Guo D, Chen J, MacEachren AM, Liao K (2006) A visualization system for space-time and multivariate patterns (VIS-STAMP). IEEE Trans Vis Comput Graph 12(6):1461–1474. doi:10.1109/TVCG.2006.84

Haklay M, Weber P (2008) Openstreetmap: user-generated street maps. IEEE Pervasive Comput 7(4):12–18

Hillaire S (2012) Antialiased volumetric lines using shader-based extrusion. In: Cozzi P, Riccio C (eds) OpenGL insights. CRC Press, Boca Raton, pp 163–172

Holten D (2006) Hierarchical edge bundles: visualization of adjacency relations in hierarchical data. IEEE Trans Vis Comput Graph 12(5):741–748. doi:10.1109/TVCG.2006.147

Holten D, Wijk JJV (2009) Force-directed edge bundling for graph visualization. Comput Graph Forum 28(3):983–990

Jankowski P, Andrienko N, Andrienko G, Kisilevich S (2010) Discovering landmark preferences and movement patterns from photo postings. Trans GIS 4(6):833–852

Lambert A, Bourqui R, Auber D (2010a) 3D edge bundling for geographical data visualization. In: 14th international conference information visualisation, pp 329–335. doi:10.1109/IV.2010.53

Lambert A, Bourqui R, Auber D (2010b) Winding roads: routing edges into bundles. Comput Graph Forum 29(3):853–862. doi:10.1111/j.1467-8659.2009.01700.x

Lampe OD, Hauser H (2011) Interactive visualization of streaming data with Kernel Density Estimation. In: 2011 IEEE pacific visualization symposium, pp 171–178. doi:10.1109/PACIFICVIS.2011.5742387

Luo SJ, Liu CL, Chen BY, Ma KL (2012) Ambiguity-free edge-bundling for interactive graph visualization. IEEE Trans Vis Comput Graph 18(5):810–821. doi:10.1109/TVCG.2011.104

Munzner T, Hoffman E, Claffy K, Fenner B (1996) Visualizing the global topology of the MBone. In: IEEE symposium on information visualization. IEEE Comput. Soc. Press, pp 85–92. doi:10.1109/INFVIS.1996.559226

Nielsen J (1994) Estimating the number of subjects needed for a thinking aloud test. Int J Hum Comput Stud 41:385–397

Nielsen TS, Hovgesen H (2008) Exploratory mapping of commuter flows in England and Wales. J Transp Geogr 16(2):90–99. doi:10.1016/j.jtrangeo.2007.04.005

Openshaw S (1984) The modifiable areal unit problem. In: Concepts and techniques in modern geography

Phan D, Xiao L, Yeh R, Hanrahan P, Winograd T (2005) Flow map layout. In: IEEE symposium on information visualization, vol 17

Rae A (2009) From spatial interaction data to spatial interaction information? Geovisualisation and spatial structures of migration from the 2001 UK census. Comput Environ Urban Syst 33(3):161–178. doi:10.1016/j.compenvurbsys.2009.01.007

Rae A (2011) Flow-data analysis with geographical information systems: a visual approach. Environ Plann B Plann Des 38(5):776–794. doi:10.1068/b36126

Rodrigue JP, Comtois C, Slack B (2009) The geography of transport systems. Routledge, Evanston

Selassie D, Heller B, Heer J (2011) Divided edge bundling for directional network data. IEEE Trans Vis Comput Graph 17(12):2354–2363. doi:10.1109/TVCG.2011.190

Specht S, Hanewinkel C (2011) Visualizing large sets of socio-economic flow data. In: 25th international cartographic conference. Paris

Spitzbart A (1960) A generalization of hermite's interpolation formula. Am Math Mon 67(1):42–46

Tobler WR (1987) Experiments in migration mapping by computer. Am Cartogr 14(2):155–163

Verbeek K, Buchin K, Speckmann B (2011) Flow map layout via spiral trees. IEEE Trans Vis Comput Graph (TVCG) 17(12):2536–2544. doi:10.1109/TVCG.2011.202

von Landesberger T, Kuijper A, Schreck T, Kohlhammer J, van Wijk J, Fekete JD, Fellner D (2011) Visual analysis of large graphs: state-of-the-art and future research challenges. Comput Graph Forum 30(6):1719–1749. doi:10.1111/j.1467-8659.2011.01898.x

Voronoi G (1908) Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire Recherches sur les parallélloèdres primitifs. J Reine Angew Math (Crelles Journal) 1908(134):198–287. doi:10.1515/crll.1908.134.198

Vrotsou K, Andrienko N, Andrienko G, Jankowski P (2011) Exploring city structure from georeferenced photos using graph centrality measures. In: Gunopulos D, Hofmann T, Malerba D, Vazirgiannis M (eds) Machine learning and knowledge discovery in databases, lecture notes in artificial intelligence, vol 6913. Springer, Berlin, pp 654–657. doi:10.1007/978-3-642-23808-6_50

Wheeler AP (2015) Visualization techniques for journey to crime flow data. Cartogr Geogr Inf Sci 42(2):149–161. doi:10.1080/15230406.2014.890545

Wilkinson L, Friendly M (2009) The history of the cluster heat map. Am Stat 63(2):179–184. doi:10.1198/tas.2009.0033

Wood J, Dykes J, Slingsby A (2010) Visualisation of origins, destinations and flows with OD maps. Cartogr J 47(2):117–129. doi:10.1179/000870410X12658023467367

Wood J, Slingsby A, Dykes J (2011) Visualizing the dynamics of London's bicycle hire scheme. Cartographica 46(4):239–251

Yang Y, Dwyer T, Goodwin S, Marriott K (2017) Many-to-many geographically-embedded flow visualisation: an evaluation. IEEE Trans Vis Comput Graph 23(1):1–1. doi:10.1109/TVCG.2016.2598885