



HAL
open science

Natural Scales in Geographical Patterns

Telmo Menezes, Camille Roth

► **To cite this version:**

Telmo Menezes, Camille Roth. Natural Scales in Geographical Patterns. Scientific Reports, 2017, 7 (45823), 10.1038/srep45823 . hal-01501818

HAL Id: hal-01501818

<https://hal.science/hal-01501818>

Submitted on 4 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Natural Scales in Geographical Patterns

Telmo Menezes^{1,*} and Camille Roth^{2,1,3,†}

¹*Centre Marc Bloch Berlin e.V., Friedrichstr. 191, 10117 Berlin, Germany*

²*Sciences Po, médialab, 84 rue de Grenelle, 75007 Paris, France*

³*Centre National de la Recherche Scientifique, France*

Human mobility is known to be distributed across several orders of magnitude of physical distances, which makes it generally difficult to endogenously find or define typical and meaningful scales. Relevant analyses, from movements to geographical partitions, seem to be relative to some ad-hoc scale, or no scale at all. Relying on geotagged data collected from photo-sharing social media, we apply community detection to movement networks constrained by increasing percentiles of the distance distribution. Using a simple parameter-free discontinuity detection algorithm, we discover clear phase transitions in the community partition space. The detection of these phases constitutes the first objective method of characterising endogenous, natural scales of human movement. Our study covers nine regions, ranging from cities to countries of various sizes and a transnational area. For all regions, the number of natural scales is remarkably low (2 or 3). Further, our results hint at scale-related behaviours rather than scale-related users. The partitions of the natural scales allow us to draw discrete multi-scale geographical boundaries, potentially capable of providing key insights in fields such as epidemiology or cultural contagion where the introduction of spatial boundaries is pivotal.

I. INTRODUCTION

Geographical scaling has been at the core of a wealth of studies of human mobility. On one hand, physical distances between connected individuals or between related places have repeatedly been shown to hardly obey any distinctive scale, let alone exhibit distinct phases. Distance frequencies observed in large geotagged datasets of human behaviour usually follow strongly heterogeneous distributions spanning several orders of magnitude, be it for links based on cell phone movements^{14,17} and calls^{9,18,31}, social media “check-ins”^{10,22} or postings⁴, commutes^{20,29} or taxi rides²¹, or circulation of artifacts⁸. On the other hand, this type of data has more recently been used to uncover geographically consistent areas based on clusters of places, movements or interactions^{3,7,11,13,19,22,27,31,34,35} where, in essence, the relevant literature generally proceeds on the assumption that all empirical measurements, irrespective of their diverse spatial scales, should be taken into account to form a single global picture. The choice of the appropriate description scale is left to the beholder: *ex ante*, when gathering data within a given bounding box, and often *ex post*, by focusing on a proper description scale. Here, behavioural traces spanning several orders of magnitudes are typically aggregated independently of the physical scale they correspond to; then, geographical areas or patterns are uncovered by community detection algorithms; a final level of description is finally chosen according to some criterion. In practice, these methods generally produce dendrograms defining an embedded series of geographical partitions, where lower-level partitions include higher-level ones in a continuum of increasingly coarse description scales. An appropriate level of the dendrogram is eventually selected because it either maximises some quantity (typically modularity in network-based methods³⁴), yields a clear-cut

dichotomy³¹, or best matches some a priori known description scale¹³. Results are therefore single-scale rather than scale-free: the aim generally consists in discovering or, rather, recovering a gold standard geographic partition of a given area — such as the partition of Belgium into two linguistic communities³¹, or the breakdown of administrative regions in Great Britain²⁷. While some studies showed that long- and short-distance connections play distinct roles in defining clusters³⁴, the quest for an ideal set of clusters which have to be discovered once and for all (and possibly aggregated into larger blocks¹³ until a binary dichotomy is reached³¹) remains pregnant.

We show here that the choice of observation scales is neither exogenous nor univocal. To this end, we demonstrate that it is possible to endogenously uncover a small number of meaningful description scale ranges from apparently scale-free raw data. In other words, geographical data on human behaviour encloses several coexisting and natural phases which we recover despite the absence of scale at the lower level of link distance distributions.

Empirically, we rely on human mobility data stemming from Instagram, an online photo-sharing service targeted at smartphone users. Distributions of link distances between successive user locations are unsurprisingly devoid of a typical scale (see Fig. S2). Admittedly, this heterogeneity corresponds to an entanglement of a variety of human behaviours, ranging for instance from local commuting to long-range travel. The ubiquitous scale-free distribution of distances certainly aggregates links of diverse nature (and, in turn, depends on diverse geographical patterns, such as borders). However, we draw markedly distinct conclusions than the current state of the art as regards its significance. Rather than using all the data and discussing the optimality of a high-level observation scale a posteriori, we work the other way around by a priori relying on link scale to blindly define a series of scale-dependent networks. These networks are based on

an increasing link distance threshold and thus configure an increasing movement radius. This then yields a series of geographical partitions from which we derive a small number of remarkably consistent high-level observation scales. This simultaneously defines a small set of relevant low-level scales, in terms of link distances or movement radii.

In other words, we disentangle the endogenous scale structure by exhibiting phase transitions based on the similarity of geographical patterns. Therefore, while we acknowledge that territories are potentially structured by partially overlapping partitions, we contend that the underlying behaviour corresponds to only a few sensible scales. We liken this implicit finding to the explicit, man-made hierarchies which can be found in more traditional top-down approaches relying on discrete ontologies featuring a small number of embedded spatial scales, such as administrative divisions (e.g. NUTS)¹³.

II. EMPIRICAL APPROACH

We obtained datasets of human movement in a variety of geographical regions from public data extracted from Instagram over a period of 16 months (See *Materials and Methods* section for a complete and formal description of the techniques and algorithms discussed here). Modern smartphones contain geolocation technology that can be employed to geotag photos, and many users accept this setting. Instagram associates photos with the identifier of the users who took them, and it also timestamps them.

The global adoption of Instagram in the last few years make it a considerable trove of geosocial data for a nascent literature. Regarding mobility in particular, several recent works demonstrated that reliable mobility data could be extracted from photo-posting and check-in platforms^{4,10,11,22}, while Instagram appears to be both. Focusing on individual movements,³⁰ shows that geographical patterns inferred from Instagram are similar to those found in data stemming from Foursquare, a so-called “check-in” platform where users typically broadcast their position to their friends (however, temporal patterns and posting behaviours differ between the two platforms, suggesting that they correspond to distinct Internet uses). A more recent study⁷ deals more precisely with sociospatial patterns and divisions at the level of cities. It describes how Instagram may be used to reveal clusters of users and places which are qualitatively meaningful and consistent with the manual analysis of the information made available by users on their account (describing for instance their profession, affiliations, interests).

For the purpose of this work we are not interested in the photos themselves, but only in the metadata. By tracking the places where a given user took photos we can infer the plausible relatedness between any two given locations in a region.

We focus on nine different areas, that were chosen to offer a diversity of cases according to several criteria:

Belgium, Portugal, Poland, Ukraine, Israel, the wider Benelux region and the cities of Berlin and Paris. We have thus five countries of various sizes, a transnational region in Western Europe and two cities.

A. Voronoi diagrams and human movement graphs

We work with two main formalisms: Voronoi diagrams¹ and weighted graphs of intensity of human movement between places. Given a pre-defined set of geographical positions in a region – let us call them seeds – the Voronoi diagram partitions the plane in cells, such that every cell corresponds to a seed. Voronoi diagrams have the property that any point in a cell is closer to the cell seed than to any other seed. This allows us to discretise the space, by assigning each photo in a region to a Voronoi cell. The pre-defined positions are also the vertices in the graph, while the weight of an edge is simply the number of users that the two vertices share — that is, the number of users that took at least one photo in each location.

From these weighted graphs we generate distributions of edge distances which we divide into 100 percentiles. This constitutes the full spectrum of scales that we work with. The use of percentiles of the distance distribution allows for the definition of comparable scales across regions, while the absolute distance values that correspond to each scale are determined by the aggregate of human movement in a given region. A graph connecting the positions in a region is generated for each one of these percentiles, by considering only connections with a distance up to the given percentile. In other words, for each region, the graph at scale s gathers $s\%$ of all observed movements ranked by increasing distance, i.e. movements up to the distance radius that corresponds to percentile s .

B. Geographical clusters and boundaries

Scale-dependent graphs are partitioned using a *community detection algorithm*. The detection of geographical clusters often relies on network-based community detection method, for instance in⁹ where county-level borders are reconstructed by maximising modularity of graph communities in US-based cell phone data, in²² which is based on a weighted network of check-in trajectories, or¹³ where Louvain is applied on a commuter network. This class of algorithms attempts to find the partition of a graph with the highest modularity, which is the fraction of edges inside the partitions normalised by the expected fraction on a random graph with the same degree distribution. Modularity thus measures the *strength* of a given partition²⁴.

At this stage we switch to Voronoi diagrams. Their cells can now be assigned to geographical clusters according to the results of the above step. To remove residual

noise, a smoothing process is applied over the partitioned diagrams. The actual geographical boundaries can now be computed. Two locations are considered neighbours if their respective Voronoi cells share a boundary. Edges shared by two Voronoi cells assigned to different communities are drawn as boundaries.

III. RESULTS

A. Phase transitions and natural scales

After community detection and smoothing for every percentile-scale, we are in a position to analyse the similarity between scales. More specifically, we are interested in seeing if there are well-defined ranges of scales that are sufficiently similar amongst each other and sufficiently distinct between ranges so that we can talk of natural scales, and reduce the 100 percentiles to a smaller number of scales.

Using a simple parameter-free breakpoint detection algorithm we are able to find phase transitions in scale space. We understand “phase transition” in a generic way, i.e. it corresponds to an abrupt change in the behaviour of partitions when slightly increasing the movement radius, going from a scale to the next ones. For the nine regions under study, our algorithm finds that the scale space is divided into no more than 2 or 3 well-differentiated intervals of scales characterised by very similar patterns. Aggregating increasingly long links while remaining below the upper bound of a given interval does not alter significantly the space partition typical of that interval. We call these intervals *natural scales*. Moreover, the breakpoints automatically found by our algorithm mostly match the visual intuition: in figure 1, we see that these phase transitions are also quite obvious simply by visual inspection.

B. Multi-level partitions and prototypical scales

Given intervals of similar scales or natural scales, it is now desirable to have a method to visualise the boundaries defined by the partitions in those intervals. We propose a simple solution: identify the percentile that best represents the entire interval. We call this prototypical percentile a *prototypical scale* of the region under study. The prototypical scale of a given interval is the percentile of the interval with the corresponding partition that is the most similar to all other partitions in the interval. Prototypical scales found for each region are also represented in figure 1, along with natural scales. By construction, partitions the various scales of a given natural scale should thus roughly resemble the partition of the corresponding prototypical scale. In the following maps, prototypical scales are thus used as visual representations of natural scales.

Figure 2 uses the results for Belgium to illustrate how natural scales correspond to partitions in the map, and how the several natural scales can be combined in a single multiscale map, which provides richer information about the geographical patterns of the region than what is possible with more traditional methods. By using the full graph (percentile 100) and forcing the community detection algorithm to find the best partition in two communities, we present a bipartite division of the territory. As can be observed, the resulting partition matches almost perfectly the border between the two largest linguistic communities in Belgium. This is a well known result^{5,31} and it shows two things. On one hand, when we simplify our method this way, thus making it equivalent to previously published approaches, we obtain similar results, which provides some evidence of correctness. On the other hand, adopting these simplifications is most likely not the best way to unravel the structure of movement patterns in Belgium: to the contrary, for all scale phases, a bipartition does not achieve the best modularity, which usually corresponds to a larger number of geographical areas (Fig. 2).

Following our method, Belgium may be more precisely decomposed as an overlay of three territorial partitions of increasing fineness. The largest natural scale features a partition based on a small number of broad areas whose boundaries correspond to inter-urban mobility, as it only emerges when links longer than 81.5 km are included. By contrast, the smallest natural scale is such that most boundaries surround and enclose a local capital; it is based on links smaller than $r \leq 38.6$ km. The middle natural scale appears when links between 38.6 and 81.5 km are considered. Interestingly, it diverges from the long-distance scale by only a few boundaries: for instance, while Hasselt is part of a broader Dutch-speaking cluster on the large-scale map, it belongs to the same cluster as Liege at the medium scale.

Finally, in figure 3, we further present the multi-scale maps for all nine regions. We also depict the absolute physical distances for all natural scale thresholds. Note that the absolute physical meaning of “large” or “small” scale is heavily region-dependent: for Paris, which is a quite dense metropolis extending over a comparably small area, the smallest natural scale typically covers a range of pedestrian movements ($r \leq 2.1$ km). For Berlin, the switch between the small and large natural scales occurs at a radius of $r \leq 10.0$ km which could correspond to “local” foot, bike or metro trips. For the largest regions such as Poland, Romania or Ukraine, they seem to correspond to a wider range of motorised inter-urban displacements, roughly around the order of magnitude of a hundred of kilometers.

A set of high-resolution maps for all natural scales as well as multi-scale representations of all regions may be found in the *Supp. Info*. Performing a thorough socio-geographic analysis of these maps is beyond the scope of this article, but we can identify some features that confirm folk knowledge about certain regions. In Portu-

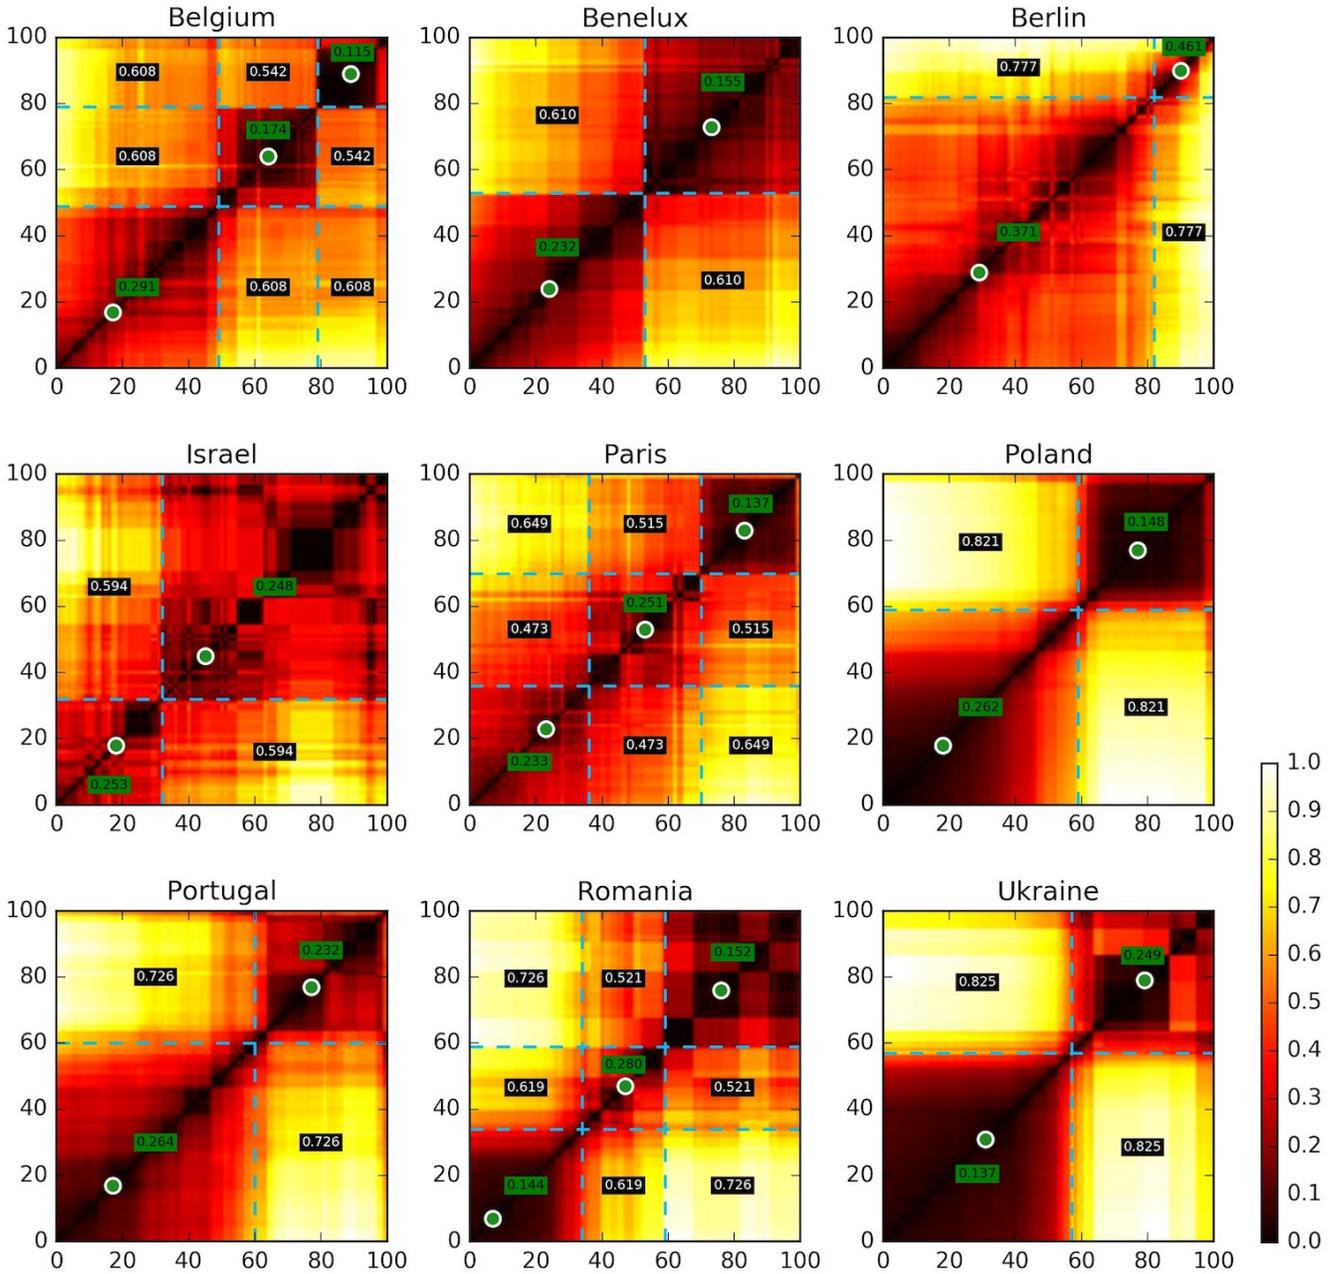


FIG. 1. *Scale dissimilarity heat maps*. Dissimilarity values are normalized per region to a $[0, 1]$ scale. Lighter colors represent higher dissimilarity. Pure black (0.0) corresponds to a perfect match, bright yellow (1.0) to the maximum dissimilarity found for the region. Dashed blue lines indicate the discontinuities identified by the breakpoint detection algorithm and, accordingly, natural scales; green dots represent the prototypical scale for each natural scale interval. The mean absolute dissimilarity value per pair of intervals is shown. A value in a green background corresponds to an internal mean dissimilarity (the interval is being compared to itself); a black background indicates a mean dissimilarity between different intervals.

gal, large scale boundaries delineate the highly touristic beaches of Algarve in the south and fuzzily divide the country into north and south regions, while the short scales provide sensible local partitions, for example the dense city of Oporto and the socio-economic divide between the capital city of Lisbon and the neighbouring but more affluent Cascais/Estoril coastal area. The Benelux

map enriches the previous insight on Belgium by providing a broader picture on potential cross-national interfaces – an achievement not possible with country-specific datasets traditionally used in the literature – here, the highest scale exhibits a mix of expected international borders (for instance between Belgium and the Netherlands) and fuzzy cross-national spaces (such as the wide

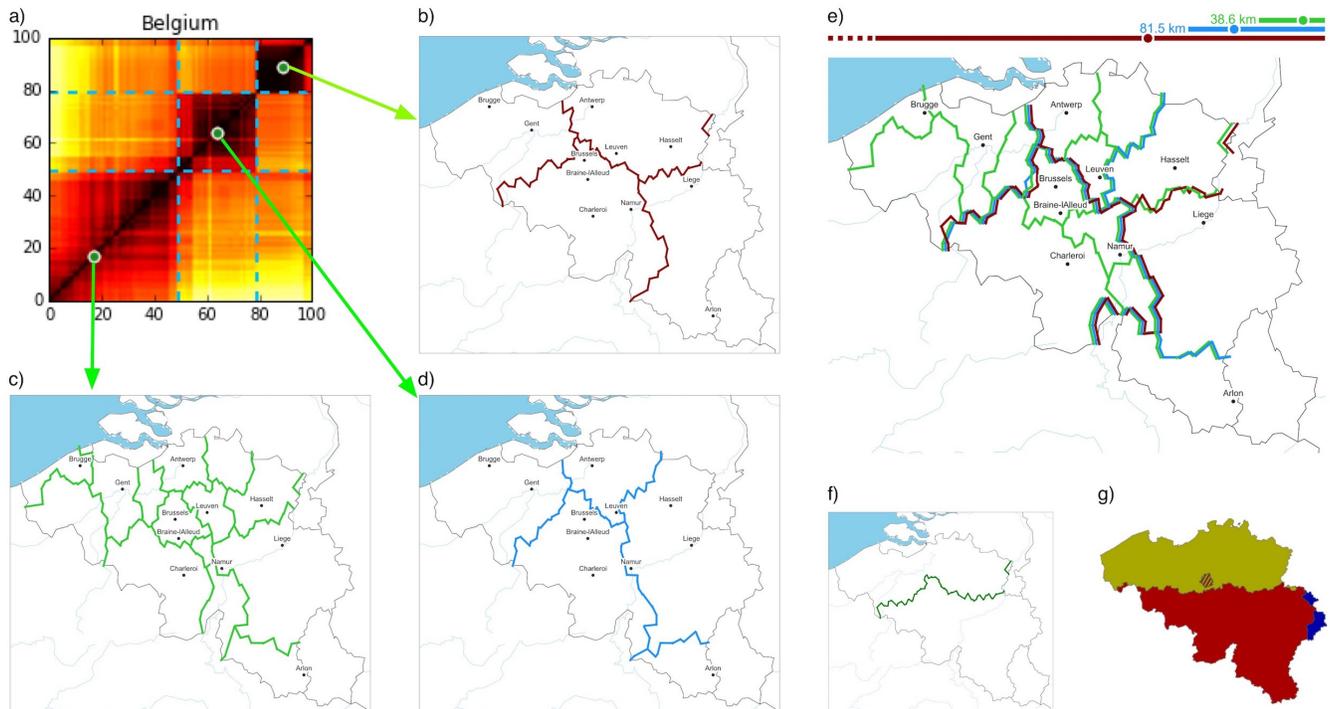


FIG. 2. *Belgium borders at different scales.* a) Heat map extracted from figure 1; b) Borders for the long distance scale; c) Borders for the short distance scale; d) Borders for the middle distance scale; e) Multiscale borders; f) Borders based on optimal two community partitioning of the full graph; g) Language communities of Belgium. All maps except g) were generated by the authors using the Basemap Matplotlib Toolkit ver. 1.0.8 (<http://matplotlib.org/basemap/>). Map in g) ©Wikimedia contributors Vascer and Knorck, licensed under CC BY-SA 3.0. The licence terms can be found on the following link: <https://creativecommons.org/licenses/by-sa/3.0>

commuting area surrounding Luxembourg, or the narrow strips adjacent to the French-Belgian border, e.g. around Lille), while leaving room for cross-border low-scale patterns. Paris features both the traditional east-west sociological partition of the city, while exhibiting more specific activity neighbourhoods at the lower level (Quartier Latin, Belleville, the governmental area).

C. Scale-dependent user behavior

Natural scales thus describe geographical areas and boundaries operating within a broad range of scale percentiles, though not beyond. In this respect, they correspond to a discrete spectrum of mobility behaviours which most likely unveil consistent yet distinct spatial practices of the underlying region. How are scales, boundaries and user behaviour related? For one, we observe on figure 3 that some regions such as Poland or Romania appear to exhibit a much higher proportion of smaller, lower-scale patterns than other regions such as Paris or Benelux. We find that these discrepancies have an interpretation in terms of user-level mobility behaviour: regions where movement distance distributions are broadest (i.e. where low and high percentiles corre-

spond to markedly distinct physical distances) also exhibit a much larger amount of small-radius geographical patterns at the shortest scale (see figures S3 and S4 of *Supp. Info.*). In other words, we show that the relative amount of patterns across the spectrum of natural scales corresponds to a relative spread of actual physical link distances across that same spectrum.

We further examine the relationship between natural scales and user-level behaviour by assigning to users the set of natural scales that they contributed to. We consider that users contribute to a natural scale if they perform at least one movement with a distance within the corresponding scale interval. Figure 4 shows a mixed picture. Overall, the proportion of users contributing exclusively to the highest scales is generally small, while the shortest scales are the most populated. At the same time, the most active users in terms of visited locations (as well as posted photos, see *Supp. Info.*) are those who span the most scales. From this we conclude that there exists a wide core of users active in all scales, which additionally always gathers a sizable proportion of all users (often the highest proportion). This hints at the fact that natural scales are based on scale-related behaviours rather than scale-related users.

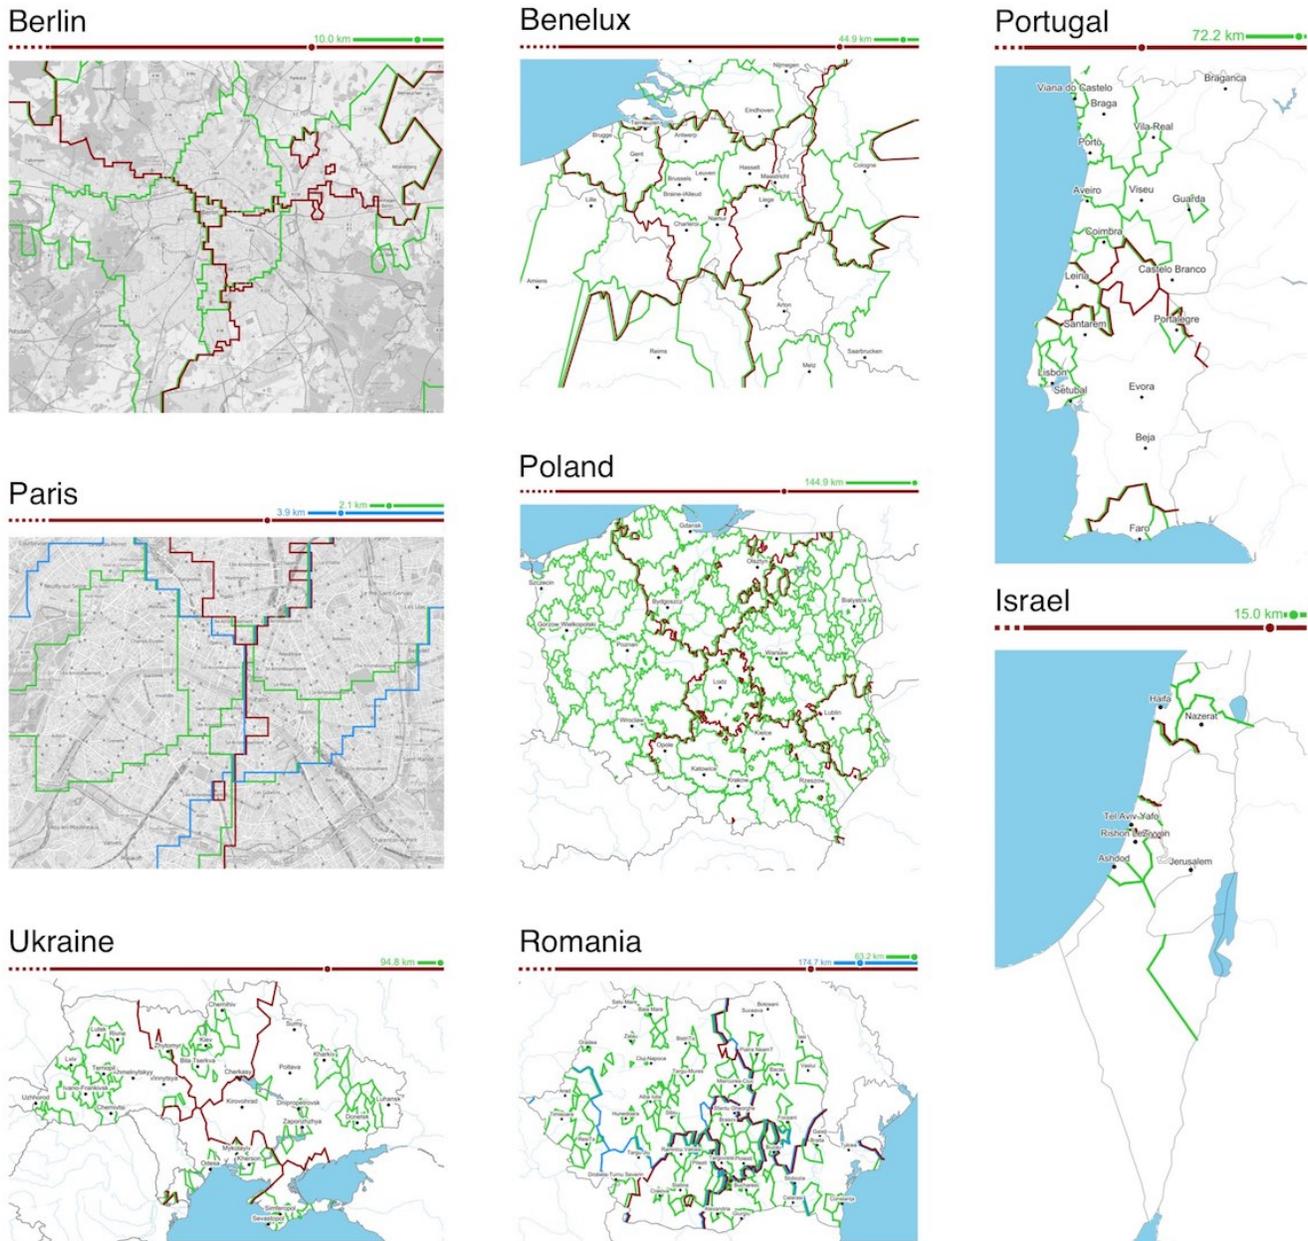


FIG. 3. *Several multi-scale maps.* Green corresponds to the smallest natural scale, blue to the middle (if it exists) and red to the largest. All maps were generated by the authors using the Basemap Matplotlib Toolkit ver. 1.0.8 (<http://matplotlib.org/basemap/>). Map tiles used in the background of the Berlin and Paris maps ©OpenStreetMap contributors, licensed under CC BY-SA (www.openstreetmap.org/copyright). The licence terms can be found on the following link: <http://creativecommons.org/licenses/by-sa/2.0/>.

IV. CONCLUDING REMARKS

By effectively distinguishing link scales and defining an increasing series of more and more global networks, we show that territories are automatically decomposable into a partially overlapping hierarchy of geographical partitions and, further, that this hierarchy exhibits a remarkably small number of natural scales. Besides, we

fulfilled in the case of spatial mobility networks the ambition of finding natural phases in community partitions based on some notion of resolution (see ref.³⁶ for non-geographical scale-free networks). In contrast with the classical expectation that aggregate mobility data is essentially scale-free, we were able to uncover a discrete number of distance thresholds and radii configuring consistent movement patterns.

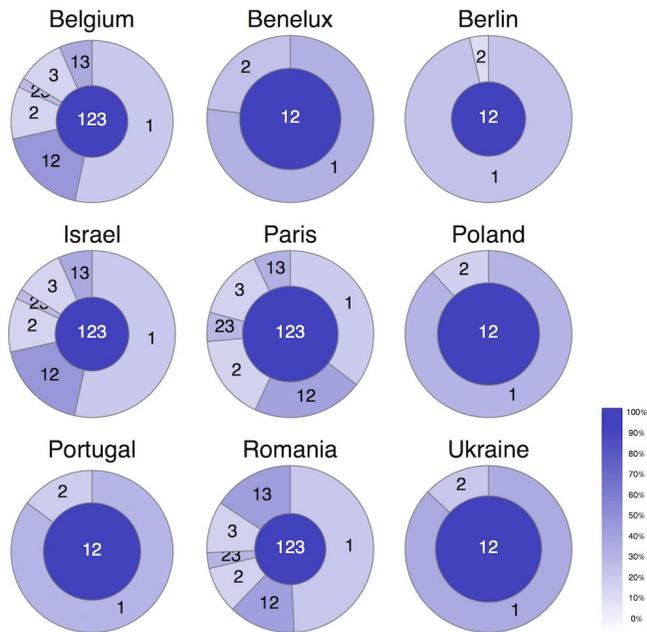


FIG. 4. *Fraction of users contributing to each natural scale.* The area of each slice/circle is proportional to the number of users active in the set of scales that it represents (for instance, “12” corresponds to users contributing exclusively to scales 1 and 2). User activity is represented by slice darkness, which is proportional to the number of visited locations relative to the maximal activity of a given region (100%): here, “123” users are always the darkest / most active slice, they consistently visit many more locations than other users.

More broadly, understanding and breaking up mobility patterns as an overlay of a small number of endogenous scale-specific behaviours bears important consequences in diverse fields such as epidemiology, cultural contagion and public policy^{2,15,25} where the low-level modeling of displacements^{4,10,28,32,33} is pivotal: here, the introduction of boundary conditions based on a scaffolding of a small number of natural scales emerging endogenously from the data could prove to be particularly fruitful.

V. METHODS

The datasets used in this paper can be obtained by request to the corresponding author. The source code of the programs that implement all the data retrieval and processing tasks discussed in this paper was released as open source and is available in a public repository²³.

A. Data collection and preprocessing

The Instagram API allows for the collection of all the meta-data and comments of photos on a maximal 5 km radius around a given geographical coordinate. To define points of interest, for countries we use a worldwide database of geographical locations with more than 500 inhabitants. For cities we simply define a dense enough grid that guarantees that the entire territory is covered, given the 5 km radius around each point. We then query the API for all photo meta-data within the maximal radius around each point. Given the possibility of overlap for locations close to each other, we perform further data processing to remove duplicates and associate each photo with the closest known location.

B. Networks, scales and boundaries

We use the previous data to generate a weighted graph connecting the set of locations. This graph is undirected and based on user movement, by considering the set of locations where a same user took photos. The weight of an edge represents the number of users who took photos in both the locations connected by the edge. Conventionally, we can represent this graph as $G = \{V, E\}$, where V is the set of locations and $E \subset V \times V \times \mathbb{N}$ the set of weighted edges.

We finally remove all vertices with a very low degree (which we define in *ad hoc* fashion as less than 5 for all regions). Intuitively, this means that we only consider locations where at least 5 different users took photos. These very low activity locations are highly susceptible to sampling distortion and introduce noise in the community detection process.

1. Percentile graphs

One graph is generated for each distance percentile, which thereby defines a scale. Let us consider m_s to be the maximum absolute distance for percentile or scale s , and $d(e)$ a function that gives the distance between the two vertices of edge e . The graph G_s for scale s is then defined as:

$$G_s = \{V, E_s\}, \text{ where } E_s = \{e \in E \mid d(e) \leq m_s\} \quad (1)$$

2. Network partitions

We employ the well-known Louvain method⁶ — a *de facto* gold standard in network community detection, widely used for the high quality of its results at a low computational cost — as implemented in the igraph software package^{12,37}. Optimal community detection, like many clustering problems, is probably NP-hard¹⁶. The

Louvain method is thus an approximation algorithm. It is also non-deterministic. To both achieve higher quality partitions and increase the stability of partitions across scales, we perform 100 runs of Louvain for each graph, and choose the result that attains the highest modularity. Another common approach is to consider all the outcomes of a large number of runs, and visualise the partitions in a way that assigns visual weights to boundaries in proportion to the number of times they appear³⁴. Given that we are working with the extra dimension of scale, we avoid this approach for the sake of simplicity.

Notice that community detection is performed on the graph of locations, with no information on the geographic proximity of the vertices. Fortunately, we do find that the communities detected are mostly contiguous, with some noise (see *Supp. Info.*).

The Louvain method can produce an arbitrary number of partitions. To validate our results, we are also interested in producing bi-partitions. To achieve this, we take the best partition found by Louvain and exhaustively try all possible merges of the given partitions into two. The merge with the highest modularity (although typically lower than the result produced by Louvain) is chosen.

C. Geographical boundary smoothing

We use this notion of Voronoi neighbourhood to define a smoothing process. From the partition process of the previous section, every location is assigned to a community. If the majority of the neighbours of a location (including the location itself) belong to a different community, then the cell is assigned to this majority community. The process is repeated iteratively, until the previous condition is not triggered.

The geographical boundaries are finally defined by the Voronoi cell boundaries for which the two neighbouring cells' locations do not belong to the same community.

D. Scale similarity, breakpoint detection and natural scales

1. Measuring partition similarity

Firstly we define a metric of similarity between two partitions of a same set of locations using a Rand index²⁶. Consider V the set of locations (as before) and P_s and $P_{s'}$ two partitions of V produced at scales s and s' by community detection followed by smoothing. A partition is defined as a set of subsets of locations, it is thus included in $\mathcal{P}(V)$.

Let us define a function $\mu^P(i, j)$ that takes the value 1 if both i and j belong to the same subset of a partition P , 0 otherwise:

$$\mu^P(i, j) = \begin{cases} 1, & \text{if } \exists X \in P \text{ such that } i, j \in X \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

We can then define the similarity of P_s and $P_{s'}$ as the ratio between the number of pairs of locations in V that have the same value of μ for both P_s and $P_{s'}$ (i.e. they are classified similarly at scales s and s'), and the total number of possible location pairs:

$$\delta(V, P_s, P_{s'}) = \frac{|\{(i, j) \in V^2, i \neq j, \mu^{P_s}(i, j) = \mu^{P_{s'}}(i, j)\}|}{\binom{|V|}{2}} \quad (3)$$

2. Intervals of similar scales

The above δ metric allows us to compare graph partitions for each percentile against every other percentile. An immediate application is visual inspection, by generating heatmaps as the ones show in figure 1. A central question to the research being presented in this article is whether partitions P_s change smoothly as s increase, or if there are clear discontinuities. The heatmaps indicate quite clearly that the discontinuities do exist.

To identify the breakpoints in partition similarity we introduce another metric, somewhat similar to the concept of modularity in graphs – albeit even simpler. This metric measures *interval separation*, given a set of breakpoints $B = \{b_0, \dots, b_n\}$. Let us also consider the set of intervals defined by these breakpoints: $\mathcal{I}(B) = \{[0, b_0], [b_0, b_1], \dots, [b_n, 100]\}$. The interval separation for a given B can thus be defined as:

$$\sigma_B = \frac{\sum_{I \in \mathcal{I}(B)} |I| \cdot \sum_{s, s' \in I} \delta(V, P_s, P_{s'})}{\max_{b \in B \setminus \{b_0\}} \delta(V, P_{b-1}, P_b)} \quad (4)$$

Intuitively, this is a ratio between the mean similarity within intervals (weighted by the interval size) and the maximum similarity between consecutive partitions in different intervals. The higher the σ , the greater the similarity between partitions in the same interval compared to the worst case similarity between partitions on both sides of a breakpoint between consecutive intervals.

Using this metric, we define a simple algorithm that iteratively adds breakpoints until σ can no longer be improved. We define a minimum interval size of 5 to avoid isolating noisy outliers. In practice, the minimum interval only has an effect on the two cities, for which the very final scales are indeed quite noisy.

3. Prototypical scales

Intervals thus define natural scales and for a given interval I , we define the prototypical scale s_I as the percentile of I with the partition that is the most similar to

all other partitions in I . To formalise:

$$s_I = \operatorname{argmax}_{s \in I} \sum_{s' \in I} \delta(V, P_s, P_{s'}) \quad (5)$$

Multi-scale smoothing

The smoothing method that was previously described can be extended to a set of partitions at s different scales. To each Voronoi tile we assign a tuple consisting of the community number (c_i) that the tile x belongs to at each scale after applying Louvain:

$$t_x = \langle c_0, \dots, c_s \rangle \quad (6)$$

Such tuples are treated as values, and the majority rule is applied as before. If a certain tuple is in the majority in the neighbourhood of t_x , then t_x takes the value of that entire tuple.

The advantage of this approach is that it leads to a greater overlap of borders from different scales. The cost

is that some precision is lost. As can be seen, for example in figure 2, there are some deviations from the borders at individual natural scales to the borders on the same scales of the multiscale map. We contend that this is a reasonable compromise for the purpose of apprehending the relationship between the different natural scales in a map.

ACKNOWLEDGEMENTS

This paper has been partially supported by grants “Phantomgrenzen” and “Algodiv” (ANR-15-CE38-0001), funded respectively by the BMBF (German Federal Ministry for Education and Research) and by the ANR (French National Agency of Research). We warmly thank Béatrice von Hirschhausen, Sabine von Löwis and Karin Casanova for useful discussions and remarks. We are also grateful for constructive anonymous review comments.

* menezes@cmb.hu-berlin.de

† roth@cmb.hu-berlin.de

¹ Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)*, 23(3):345–405, 1991.

² Duygu Balcan, Vittoria Colizza, Bruno Gonçalves, Hao Hu, José J. Ramasco, and Alessandro Vespignani. Multi-scale mobility networks and the spatial spreading of infectious diseases. *PNAS*, 106(51):21484–21489, 2009.

³ Aleix Bassolas, Maxime Lenormand, Antònia Tugores, Bruno Gonçalves, and José J. Ramasco. Touristic site attractiveness seen through twitter. *EPJ Data Science*, 5(1):1–9, 2016.

⁴ Mariana G. Beiró, André Panisson, Michele Tizzoni, and Ciro Cattuto. Predicting human mobility through the assimilation of social media traces into mobility models. *EPJ Data Science*, 5(30), 2016.

⁵ Vincent Blondel, Gautier Krings, Isabelle Thomas, et al. Regions and borders of mobile telephony in belgium and in the brussels metropolitan zone. *Brussels Studies*, 2010.

⁶ Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008.

⁷ John D. Boy and Justus Uitermark. How to study the city on instagram. *PLoS ONE*, 11(6):e0158161, 2016.

⁸ D. Brockmann, L. Hufnagel, and T. Geisel. The scaling laws of human travel. *Nature*, 439:462–465, 2006.

⁹ Francesco Calabrese, Dominik Dahlem, Alexandre Gerber, Deirdre Paul, Xiaoji Chen, James Rowland, Christopher Rath, and Carlo Ratti. The connected states of america: Quantifying social radii of influence. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 223–230. IEEE, 2011.

¹⁰ Eunjoon Cho, Seth A. Myers, and Jure Leskovec. Friendship and mobility: User movement in location-based social networks. In *Proc. 17th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, KDD ’11, pages 1082–1090, New York, NY, USA, 2011. ACM.

¹¹ Justin Cranshaw, Raz Schwartz, Jason I Hong, and Norman M Sadeh. The livelihoods project: Utilizing social media to understand the dynamics of a city. In *ICWSM*, 2012.

¹² Gabor Csardi and Tamas Nepusz. The igraph software package for complex network research. *InterJournal, Complex Systems*, 1695(5):1–9, 2006.

¹³ Andrea De Montis, Simone Caschili, and Alessandro Chessa. Commuter networks and community detection: a method for planning sub regional areas. *arXiv preprint arXiv:1103.2467*, 2011.

¹⁴ Pierre Deville, Chaoming Song, Nathan Eagle, Vincent D. Blondel, Albert-Laszlo Barabási, and Dashun Wang. Scaling identity connects human mobility and social interactions. *PNAS*, 113(26):7047–7052, 2016.

¹⁵ Neil M. Ferguson, Derek A. T. Cummings, Christophe Fraser, James C. Cajka, Philip C. Cooley, and Donald S. Burke. Strategies for mitigating an influenza pandemic. *Nature*, 442:448–452, 2006.

¹⁶ Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3):75–174, 2010.

¹⁷ M. C. González, C. A. Hidalgo, and A.-L. Barabási. Understanding individual human mobility patterns. *Nature*, 453:779–782, 2008.

¹⁸ G. Krings, F. Calabrese, C. Ratti, and V. Blondel. A gravity model for inter-city telephone communication networks. *Journal of Statistical Mechanics: Theory and Experiment*, L07003, 2009.

¹⁹ Balázs Lengyel, Attila Varga, Bence SÁgvári, Ákos Jakobi, and János Kertész. Geographies of an online social network. *PLoS ONE*, 10(9), 2015.

- ²⁰ M. Lenormand, M. Picornell, O.G. Cantú-Ros, A. Tugores, T. Louail, R. Herranz, M. Barthelemy, E. Frías-Martínez, and J.J. Ramasco. Cross-checking different sources of mobility information. *PLoS ONE*, 9(8):e105184, 2014.
- ²¹ X. Liang, J.C. Zhao, L. Dong, and K. Xu. Unraveling the origin of exponential law in intra-urban human mobility. *Scientific Reports*, 3(2983):1–7, 2013.
- ²² Yu Liu, Zhengwei Sui, Chaogui Kang, and Yong Gao. Uncovering patterns of inter-urban trip and spatial interaction from social media check-in data. *PLoS ONE*, 9(1):e86026, 2014.
- ²³ Telmo Menezes and Camille Roth. Ghostborders. <https://github.com/telmomenezes/ghostborders>, 2016.
- ²⁴ Mark EJ Newman. Modularity and community structure in networks. *PNAS*, 103(23):8577–8582, 2006.
- ²⁵ Alexander M. Petersen, Joel Tenenbaum, Shlomo Havlin, and H. E. Stanley. Statistical laws governing fluctuations in word use from word birth to word death. *Scientific Reports*, 2(313):1–9, 2012.
- ²⁶ William M Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850, 1971.
- ²⁷ Carlo Ratti, Stanislav Sobolevsky, Francesco Calabrese, Clio Andris, Jonathan Reades, Mauro Martino, Rob Claxton, and Steven H Strogatz. Redrawing the map of great britain from a network of human interactions. *PLoS ONE*, 5(12):e14248, 2010.
- ²⁸ I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chong. On the levy-walk nature of human mobility. In *INFOCOM 2008. The 27th Conference on Computer Communications*. *IEEE*, April 2008.
- ²⁹ Camille Roth, Soong Moon Kang, Michael Batty, and Marc Barthélémy. Structure of urban movements: polycentric activity and entangled hierarchical flows. *PLoS ONE*, 6(1):e15923, 2011.
- ³⁰ Thiago H. Silva, Pedro O. S. Vaz de Melo, Jussara M. Almeida, Juliana Salles, and Antonio A. F. Loureiro. A comparison of foursquare and instagram to the study of city dynamics and urban social behavior. In *Proc. UrbComp’13*, 2013.
- ³¹ Stanislav Sobolevsky, Michael Szell, Riccardo Campari, Thomas Couronné, Zbigniew Smoreda, and Carlo Ratti. Delineating geographical regions with networks of human interactions in an extensive set of countries. *PLoS ONE*, 8(12):e81707, 2013.
- ³² Chaoming Song, Tal Koren, Pu Wang, and Albert-László Barabási. Modelling the scaling properties of human mobility. *Nature Physics*, 6:818–823, 2010.
- ³³ Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.
- ³⁴ Christian Thiemann, Fabian Theis, Daniel Grady, Rafael Brune, and Dirk Brockmann. The structure of borders in a small world. *PLoS ONE*, 5(11):e15422, 2010.
- ³⁵ Jameson L. Toole, Michael Ulm, Marta C. González, and Dietmar Bauer. Inferring land use from mobile phone activity. In *Proc. ACM SIGKDD Intl. Workshop on Urban Computing, UrbComp ’12*, pages 1–8, New York, NY, USA, 2012. ACM.
- ³⁶ V. A. Traag, G. Krings, and P. Van Dooren. Significant scales in community structure. *Scientific Reports*, 3(2930):1–10, 2013.
- ³⁷ Vincent A. Traag. *louvain-igraph: v0.5.3*, December 2015.