Nokia MDC Atlas: An Exploration of Mobile Phone Users, Land Cover, Time, and Space

André Skupin
Department of Geography
San Diego State University
San Diego, CA 92182-4493
+1-619-594-6946
skupin@mail.sdsu.edu

Harvey J. Miller
Department of Geography
University of Utah
Salt Lake City, UT 84112-9155
+1-801-585-3972
harvey.miller@geog.utah.edu

ABSTRACT

A total of 21 maps were produced that together form a type of atlas of the Nokia Mobile Data Challenge. Like in a traditional geographic atlas, we present a limited number of base map configurations, onto which various thematic elements are then overlaid. Two of those base maps are derived from MDC data, the third is placed in geographic space. Thematic overlays serve several purposes, including elaborating different elements from which the base map geometry had been derived, as well as linking other data to the base map.

Keywords

Spatialization, Visualization, High-dimensional Space, Self-Organizing Maps, Mobility, Cartography.

1. INTRODUCTION

Mobile telephony is among a group of emerging *space-adjusting technologies*: these change the nature and experience of geographic space and consequently lead to a rearranging of human activities [1, 3]. The loosening of formerly tight bonds between people, place and activity allows other, sometime hidden, activity spaces to flourish. Rather than indicating the "death of distance" and spatial uniformity, these activity spaces relate to geographic space in interesting and complex ways [2, 4].

The purpose of this project is to explore the hidden spaces of human activity by producing proof-of-concept visualizations. These visualizations highlight a conceptualization of mobile phone users as simultaneously existing in different spaces. Our study elaborates on three of these:

(1) High-dimensional attribute space derived from mobile phone users' questionnaire responses

(2) High-dimensional attribute space derived from mobile phone users' time spent in different types of land cover

(3) Geographic space as a canvas for time spent in different activity modes

A total of 21 maps were produced that together form a type of atlas of the MDC project [5]. Like in a traditional geographic atlas, we present a limited number of base map configurations, onto which various thematic elements are then overlaid. Two of those base maps are derived from MDC data, the third is placed in geographic space. Thematic overlays serve several purposes, including elaborating different elements from which the base map geometry had been derived, as well as linking other data to the base map.

2. DATA AND METHODS

2.1 Data

This project utilized five main data sets:

- Questionnaire of MDC subjects
- GPS- and WLAN-based coordinates of MDC-enrolled mobile phones
- CORINE land cover data for Switzerland
- Bing maps and OpenStreetMaps for reference backdrop

2.1.1 Questionnaire of MDC Subjects

The questionnaire data set contained information for 29 participants. Since our approach is based on determining the similarity of subjects based on their responses, we tried to limit the number of incomplete responses entering the analysis. Questions 1, 3, 5, and 7 contained answers from nearly all the participants (e.g., only one of the 29 subjects did not answer the gender question), which is why we focused on those questions. Transformation of nominal data into a form suitable for similarity computation yielded 15 variables. In other words, we analyze the similarity of subjects in a 15-dimensional attribute space.

2.1.2 GPS-Based and WLAN-Based Coordinates of MDC-enrolled Mobile Phones

The provided gps.csv files of 38 participants contained a total of more than 1.52 million locational records. The WLAN-based `wlan_loc.csv` files contained a further 1.87 million records. The combined data set of almost 3.4 million records is the initial input to location-based analysis, allowing an intersection of participants with land cover data and analysis of aggregate activity patterns. All of those records included time stamps, which is the basis for delineating distinct journeys and for determining the time spent at particular locations.

2.1.3 CORINE Land Cover Data for Switzerland

One of the goals of the project was to demonstrate how one could represent similarity patterns among mobile phone users in terms of geographic space, but without compromising geographic privacy as such. To that end, we aimed to intersect user locations with the attributes of the geographic locations encountered. Due to ease of availability, we are demonstrating this here with a land cover data set. CORINE (Coordination of Information on the

Environment) is an initiative of the European community that has resulted in a comprehensive land cover data set spanning much of the continent, at a resolution of up to 100 meters. Switzerland is not officially part of that program, but there is obvious interest in having Swiss land cover data available in a form aligned with European efforts. At this point, a land cover layer derived from a 250 m resolution data set seems to be the best publically available approximation, in form of 13,542 polygons covering all of Switzerland at CORINE land cover level 2 (up to 15 land cover categories).

2.2 Software

The following are key software packages used in the course of this project:

• Various pre-processing steps: Microsoft Office Excel 2007, Processing development environment (http://processing.org/)

• Multidimensional scaling (MDS): IBM SPSS Statistics

• Self-organizing map (SOM) preprocessing/postprocessing: SOM Analyst (http://code.google.com/p/somanalyst/)

• SOM training: SOM_PAK (modified after http://www.cis.hut.fi/research/som_pak/)

• Visualization: ESRI ArcGIS (http://www.esri.com/software/arcgis/)

2.3 Analysis of User Questionnaires

Projection of MDC subjects from the 15-dimensional space derived from the questionnaire data into a two-dimensional display space is meant to allow visual examination of patterns among subjects. This is the basis for the point geometry in maps 1 – 7. In addition, elements of the 15-dimensional questionnaire data set are projected onto the land cover based layout of subjects in maps 14 – 16. In choosing among dimensionality reduction methods, multidimensional scaling (MDS) and self-organizing maps (SOM) are key candidates, both of which we experimented with. While MDS is very well suited to such a small data set, a large number of subjects are so similar to each other (e.g., males that work full-time for whom a private motor vehicle is the main means of transport) that they end up being extremely clustered in the display space. That is why we opted for SOM, which recognizes local variations among globally similar entities and thus spreads out the high-dimensional inputs a bit more.

2.4 Preprocessing of Location Records

The initial data set of GPS- and WLAN-based locations consists of almost 3.4 million records. As described in the MDC contest materials, some coordinates had been intentionally truncated, due to privacy concerns. This is problematic in terms of one of our approaches to delineating distinct journeys, namely to begin/end a journey when a user had not changed location in 30 or more minutes. With truncated coordinates, the user location remaining constant could either indicate that the user indeed did not move or that he/she moved within a radius of several hundred meters. Note that being able to distinguish among different journeys (along the time axis) is a crucial element of our ability to determine the time spent by a user at particular locations. Since we do not have a continuous record of users – due to GPS or WLAN functionality or the phone itself being switched off – we have to be careful about being able to determine either spatial stationarity (as described above) or temporal discontinuity (here defined as a gap of 10 minutes or more between sequential location records). Due to the importance of detecting spatial stationarity, we eliminated approximately 956,000 truncated records. With a view of intersecting locations with Swiss land cover data, a further 3,238 records outside of the spatial footprint of the CORINE data were eliminated, resulting in the final set of 2,442,048 records. Using spatial stationarity and temporal discontinuity as defined above, the sequential set of locations for each of the 38 participants was then chunked into a total of 48,959 journeys. *Within* each thus delineated journey, the time stamps of sequential records were then compared to determine an estimate of the time spent at each location.

2.5 Analysis of Land Cover

CORINE-equivalent land cover data for all of Switzerland were obtained as a polygon layer from the Swiss *Bundesamt für Statistik* and intersected with the full set of 2.4 million location records. As a result, each location record becomes associated with a land use type. Together with the “time spent” attribute computed for each record, this makes it possible to derive user-level statistics on the time spent per land cover type. Out of the fifteen CORINE land cover level 2 categories, thirteen are actually encountered in the data set (as a land-locked country, the categories “marine wetlands” and “marine waters” is absent in Switzerland). Each user can thus be represented in a 13-dimensional attribute space that expresses the proportion of time spent in each type of land cover. For exploratory visualization, users are then projected from that high-dimensional space into a two-dimensional display space, resulting in the second base map.

2.6 Geographic Mapping of Activity Patterns

The final set of visualizations relates an aggregate view of mobile phone users’ activities to geographic location. The driving idea here is that there may be distinct patterns in when and where certain activities are concentrated. A well-known critical application scenario for this type of approach is the need to distinguish between nighttime and daytime population concentrations for emergency management. Typical population census data tend to account only for the residential population by home address. That is a fine estimate for concentrations of nighttime population, but tells us very little about where people are concentrated during a typical workday. Meanwhile, other applications, such as in tourism planning, would benefit from more detailed information about geographic patterns at various times of the day, week, or year. Our proof-of-concept here attempts to highlight some examples with a focus on three distinct activity windows. For simplicity sake, we refer to these as *home*, *work*, and *leisure*:

1. ***home***: Mondays – Thursdays 3:00-4:00
2. ***work*** (which is meant to include university studies): Mondays – Thursdays 11:00-12:00
3. ***leisure***: Saturdays & Sundays 11:00-12:00

In the absence of actual activity-based attributes for the 2.4 million location records, such narrowly defined time slices would seem reasonably conservative approximations of likely activity patterns.

Through a combination of kernel density estimation and map algebra, the study then proceeds to work out explicit differences among the relative intensity with which a particular geographic location is used for the three activity types. Throughout, a fine-grained notion of density is employed that relates *time spent* to area size. Specifically, we measure intensity in seconds per square kilometer.

3. MAPS

The twenty-one maps presented here are not meant to be a comprehensive collection of visual artifacts that could be generated from the types of data that were made available for the MDC Open Challenge. Instead, we chose to implement certain sequences of conceptual, computational, and semiotic transformations that highlight the power of visualization, within the limits of available resources. We hope to demonstrate that raw data on mobile phone users and their location can be the source of novel intersections between notions of location-based computing, high-dimensional attribute space, and geographic space. Despite being designed as a series of proofs-of-concept, we deliberately chose computational techniques that will readily scale to realistically sized data sets. For example, the SOM neural network model can be applied to hundreds of thousands of mobile phone users (as compared to the examples involving 29 and 38 users, respectively). Meanwhile, the density-based analyses shown here is not only capable of effectively aggregating millions of locational records in space and time, but would also serve to effectively shield the privacy of users, especially – and to some degree only – if one was dealing with truly large data sets.

Below we give additional explanations for how each of the maps in this atlas was generated. Where appropriate, we also discuss possible interpretations of observed patterns. Given the limited amount of data available in the contest, such interpretation is meant to suggest ways of using these kinds of maps rather than suggesting generalizable insights about mobile phone users.

Map 1: Base Map of Study Participants Based On Questionnaire.

This is the first of two base map layouts of mobile phone users. The 15-dimensional vectors representing subject responses to questionnaire questions 1, 3, 5, and 7 were used to train a self-organizing neural network of 900 neurons. Following training, every input vector is placed at the location of their best-matching (i.e., most similar) neurons in the SOM. Subject IDs are used as labels. As a result of neural network training, topological structures existing in the high-dimensional input space tend to be preserved in the two-dimensional display space. Neighboring labels thus indicate that the respective subjects gave similar answers in the questionnaire.

Subjects 082 and 160 provided identical answers to questions 1, 3, 5, and 7 and thus are in the same location in the 15-dimensional input space and have the same best-matching neuron. A small random offset was used to separate their geometry.

Maps 2 through 7 project various thematic contents onto this base map. All overlays are derived from one or more input variables.

Map 2: Gender and Age of Study Participants.

The gender portion illustrates that two thirds of study participants were male. Subject 179 did not answer the gender question. Subjects 089 and 127 are the oldest participants answering the questionnaire. Matching age must be one of the reasons for them being pulled towards each other in the map, despite difference in gender.

Map 3: Status of Study Participants.

Derived from answers to questionnaire question 5, map 3 shows that most participants work fulltime. Notice how Subjects 063, 169, 179 are neighbors despite appearing in different categories. That points to them having similarities in other respects, which later maps demonstrate to be a mix of age and transport mode (see Map 6).

Map 4: Private Vehicle Ownership versus Public Transport Use. Distinction between Bus, Metro, and Train Use.

Different modes of transport are highlighted in the panels of Map 4. Note how use of a private motor vehicle versus use of public transport almost forms a binary distinction, with the only overlaps occurring in subjects 060, 083, and 185. Since different modes of transport were fed to this model as separate dimensions, it is perhaps surprising that public transport forms a contiguous region. Which factors bind public transport users in this manner? We can see that train users (red point symbol) and metro users (blue point symbol) form almost exclusive sets, with only subject 077 using both modes. What binds them is the use of buses (green point symbol) by 86% of public transport users.

Map 5: All Transport Modes Combined.

One can see here that half of the participants tend to use a single mode of transport and only two participants use four different modes. Among single-mode subjects, car users (cars here standing for any kind of motor vehicle, including motorcycles, scooters, etc.) are the only ones occurring with great frequency. Only 15% of bus riders do so exclusively, while the other bus riders also utilize additional modes of transport.

Map 6: Participants According to Questionnaire Similarity. Overlay of Participants' Age, Status, and Transport Modes.

In looking for explanation for the patterns observed so far one could flip back and forth between different maps. For example, one could ask which factors might relate to users' exclusive use of private motor vehicles and flip between Map 5 and the earlier maps in the atlas. Alternatively, one could layer a larger number of variables on top of the base map. In this manner, Map 6 combines age, status, and transport modes. We can now see that those using only private motor vehicles as modes of transport tend to fall into the higher age group and are working. The only exclusive car users not working are one younger student and the person that had not given her/his gender. Note also that students tend to be in the younger age categories, not surprisingly.

Map 7: Participants According to Questionnaire Similarity. Overlay of Age, Gender, and Transport Modes.

Observations made here include that only one of the females uses a private motor vehicle exclusively and she is part of the oldest age category. All except one of the train users are female, the only exception being a younger male that indicated four different modes of transport. Meanwhile, all except one of the bike riding MDC subjects are male.

One important note to add is that the kind of visual set operations seemingly performed when making these interpretations (1) may be unfeasible for realistically sized data sets of thousands of users and (2) could/should instead be implemented using computational means. The real power of visualization in dealing with these questionnaire-type data does not lie in *answering* questions of a numerical type, but in providing means to engage and catalyze analysts' creative potential for *generating* such questions.

Map 8: Participants According to Similarity in the Relative Time Spent Per Land Cover Type.

Again using the self-organizing map (SOM) method, 38 users are here laid out according to similarity in the relative time spent per land cover type. For most users, one land cover type is dominant, with *urban fabric* being the most common dominant type, followed by *industrial, commercial, and transport units*. Two pairs of users' are dominated by time spent near *pastures* and *arable land*, respectively, which explains those two pairs being represented on the map (068/185, 060/172).

A somewhat even mix of land cover types is quite uncommon, with subjects 082, 083, 169, and 179 as notable exceptions. We will revisit this issue later.

One major advantage of the SOM method over other dimensionality reduction methods is its efficient use of the low-dimensional space, such that finer distinction among input vectors can be worked out. Notice, for example, how one can make out finer nuances among the many users for whom *urban fabric* was the dominant land cover type. Meanwhile, the high-dimensional distance separating outliers from the main body of users in the 13-dimensional space can be dramatically compressed in the eventual map. That makes for an efficient use of display space, but one has to be careful about judging relative distances in that space. For example, subjects 109 and 117 seem to be extreme outliers. The name of the dominant land cover type for subject 109, *artificial non-agricultural vegetated areas*, is a bit opaque; urban parks are the prime example for this cover type. Subject 109 is pulled towards subject 068 because of a very similar proportion of *urban fabric*, while the complete lack of commonalities in other land cover categories avoids the kind of pushback that 109 would have received elsewhere on the map.

In terms of topological relationships in the high-dimensional space, subject 117 likely belongs with the other subjects in the upper-right corner of this map, due to significant, but not dominant, proportion of *urban fabric*, mixed with *some industrial, commercial, and transport units* and *forests*. Oddly though, this is the only subject for which our land cover overlay determined a significant portion of *inland waters*. According to the map, subject 117 spent two-thirds of time in the *inland waters* category. That seems to deserve further scrutiny. It turns out that the subject did indeed spend a fair amount of time on and near Lake Geneva. For example, the subject seems to have used the ferry system along eastern Lake Geneva, which would not have been explicitly captured by the questionnaire. More importantly, the subject seems to use public transport a fair bit (see also Map 3-7). As compared to the A9 motorway, which runs a bit inland from the Lake and would be an obvious choice for a private motor vehicle, public transport between Montreux and Lausanne follows the lakeshore very closely. Inspection of the subject's 118,676 location records shows that he (his gender being known) traveled extremely closely to the lakeshore, with the extremely close proximity between highway 9 (not the A9) and train tracks making it difficult to discern whether buses or trains were used. The critical issue is that this close proximity to the lake requires a higher-resolution land cover data set than we were able to obtan. At a land cover polygon granularity equivalent to 250 m resolution, many of the subject's near-shore locations end up intersecting with the *inland waters* category. Again, this is an issue that could be easily addressed with higher resolution land cover data, which should be available in the near future publically or are likely already existing in non-public formats. Another implication is that all of the land cover assignments would suffer from this resolution restriction, implying that we should be talking about users spending time near instead of in particular land cover types.

Map 9: Participants According to Similarity in the Relative Time Spent Per Land Cover Type. Symbols Scaled According to Total Time Spent by Participant While in Location Capture Mode.

Fine-grained mobile phone locations determined via GPS or WLAN tend to not be captured continuously since users can make the choice to switch such functionality on/off, depending on specific application needs (e.g., navigation apps versus others) and battery status. We would expect that different users have different use patterns when it comes to GPS and WLAN and that these will be reflected in the land cover patterns encountered during the generation of location records. In Map 9, the pie charts of Map 8 are scaled to reflect the total time spent by each user in location capture mode.

The most obvious lesson is that different users did in fact spend quite different amounts of time in a location-recording mode, with user 141 (bottom right) as an extreme example. The next apparent patterns is that users with a single dominant land cover tend to have recorded locations for longer periods of time. Meanwhile, the few users with highly varied and mixed proportions of different land cover types (users 082, 083, 169, and 179) tended to spend overall little time in location capture mode. Why is that? Why would users with longer location capture periods tend to encounter more monotone environments? One reasonable explanation may be that the latter users simply kept their phones in location capture mode for longer periods. Note in conjunction that users that predominantly captured locations within the *urban fabric* tended to spend longer times in location capture mode. In explaining that, we would speculate that *day-to-day* GPS navigation and use of WLAN access points tend to be more associated with activities within the *urban fabric*. Those activities should also occur with relatively even proportion across the different days of the week. Meanwhile, users only operating occasionally in location capture mode may switch on such functionality more deliberately, perhaps in association with activities that they are only engaged in on certain days of the week. This more nuanced view of the use of location capture modes over time is explored in Maps 12 and 13.

Map 10: Illustration of Land Cover Based Neural Network Model.

Maps 1 through 9 were based on projections of 29/38 users onto two-dimensional models of the 15/13-dimensional input spaces. Arguably, data sets of a few dozen entities are extremely small. It is important to point out that the computational technique employed for dimensionality reduction – the SOM method – is applicable to far larger number of input vectors, up to several million. The models themselves could (1) be trained with a far larger number of input vectors (up to several million) and (2) the trained model could be used to project data that were not part of the training data set (i.e., a more typical neural network application).

To illustrate what the *model* itself looks like, Map 10 includes a visualization of the relative weights associated with each of the 900 neurons of the land cover based user presence data set. This is the model onto which users are projected, as seen in Maps 8 and

9. Where finer nuances exist in the data, such as among users with *urban fabric* as the dominant land cover, the neural model preserves and elaborates on such nuances. Meanwhile, outliers (e.g., subjects 109 and 117) are separated from their neighbors by a rapidly changing series of neurons, which traditional multivariate methods would pick up as cluster boundaries.

Map 11: Land Cover Types According to Similarity in How Relative Time Spent Was Distributed Across 38 Participants.

Contemporary visualization is all about allowing analysts to gain multiple perspectives on data, through a series of computational and semiotic transformations. Map 11 takes this idea a bit further. It is based on transposing the very same data used to generate Maps 8-10. If users can be analyzed in terms of the relative amount of time spent in different land cover types, then one should be able to compare land cover types in terms of the relative time spent in them by the different users. In other words, we are transposing a data set consisting of 38 objects in a 13-dimensional space into one consisting of 13 objects in a 38-dimensional space. We then project from that space into a 2-D display space, leading to the geometric distribution of land cover types shown in Map 11.

This map is included in the MDC atlas for two reasons: (1) it shows the total time spent by users in each land cover type (the resolution-related issues notwithstanding) and (2) it illustrates how one might embed the analysis of individual movement behavior within much broader investigations of geographic phenomena. For example, one could address rather complex questions surrounding cyclical activity patterns (e.g., seasonally varying uses of city parks) or long-term changes in activity patterns (e.g., skateboarding or parkour as changing the roles played by different urban environments). Such studies would of course require much larger data sets.

Map 12: Time Spent in Location Capture Mode Relative to Total Time Passed Between First and Last Location Capture of Each Participant.

Much of the earlier discussion regarding map 9 would be moot if some users simply operated their devices for an overall much shorter period than other users. Smaller amounts of time spent in location recording mode would thus be simply a function of total mobile phone use. Map 12 rejects that idea. The total time spent in location recording mode is here seen in proportion to the length of time between the first and last time stamped location. Users with a dominant proportion of time spent in the urban fabric (see Map 9) did indeed tend to spend a large proportion of their time in a location recording mode (typically in the 20%-40% range). Meanwhile subjects with a more even mix of land cover types (users 082, 083, 169, and 179) tended to spend much smaller proportion of time in that mode (2%-7%).

Map 13: Relative Time Spent in Location Capture Mode on Different Days of the Week.

In the discussion of Map 9 we speculated that the differences in the relative dominance of certain land cover types might be reflected in differences in how location capture was distributed across days of the week. Map 13 confirms that. Users that spent most of their time in the *urban fabric* (see Map 9) seem to have locations recorded quite evenly across the range of weekdays. The percentage of time during which locations were recorded on a weekend day (Sat and Sun) is also roughly in accordance with the expected value (2/7 = 28.6%). Compare that to user 083, who recorded almost half of her locations on weekend days, or user 082, with disproportionately large location capture on Wednesdays and Thursdays.

Map 14: Participants' Transport Modes Indicated in Questionnaire Overlaid on Geometry Derived from Relative Time Spent Per Land Cover Type.

A crucial idea behind the base map notion – whether in its traditional geographic context or for high-dimensional information spaces – is that one can map onto the base geometry other data that were generated independently from the base map. Maps 14 through 16 demonstrate this with an overlay of user questionnaire data onto the land cover based user geometry. Patterns observed in this map would point to possible relationships between how subjects characterized themselves in the questionnaire and their on-the-ground activity patterns. Even more than for the exploration of base map variables (where the set of variables is fixed and known), this kind of complex overlay is best done within an interactive environment, where different overlays can be performed in a highly exploratory manner. Maps 14 through 16 are simply meant to demonstrate some of the potential of that approach, including the ability to generate meaningful narratives from data through visualization.

Map 15: Gender and Transport Mode (Bus And Metro) Overlaid on Land Cover Based Map of Participants.

The two panels in Map 15 exemplify the search for meaningful patterns one might perform, such as:

• “Do females and males separate into meaningful groups on the basis of land cover types they spent time in?”

> Probably not, given the lack of visible organization.

• “Does use of buses indicated in the questionnaire relate to land cover types?”

> We wouldn’t expect to, since regular use of buses was claimed by 86% of public transport users and those users should reflect a diversity of land cover types. Bus users are indeed quite widely distributed across land cover types, with the only noticeable pattern being formed by group of male bus riders along the bottom-left edge of the map.

• “Does use of the metro indicated in the questionnaire relate to land cover types?”

> The answer seems to be yes! Of the seven metro users, five male bus riders form a contiguous group of metro riders.

Map 16: Gender, Bike Use, and Relative Time Spent Per Land Cover Type Overlaid on Land Cover Based Map of Participants.

When thematic overlays onto a base map are performed, one will often want to relate questions regarding the thematic layers back to base attributes. For example, in Map 16 one observes that bike riders are placed within a band of males from the bottom left towards the mid-right edge. The absence of bike riders anywhere else in the map is likewise noticeable. Adding a land cover layer (containing the variables from which the base geometry was created) helps to see that all but one of the bike riders spent the majority of recorded time in the urban fabric. The only exception is user 109, a male who spent most of his recorded time in or near parks (per the *artificial non-agricultural vegetated areas* category).

Map 17: Relative Time Spent in Location Capture Mode During Typical Work Hours as Compared to Home.

The final set of maps (17-21) are derived from the density of location records within certain time windows, with density expressed as seconds spent per square kilometer. Subtraction of density surfaces from each other allows elaborating the specific geographic areas in which certain activities are concentrated. In Map 17, work time density is subtracted from the density at a time when most users were presumed to be at home. Overlaid on a standard geographic base, EPFL and central Lausanne clearly emerge as centers of workday activity, while Prilly serves as more of a residential area (for users in this study).

Map 18: Relative Time Spent in Location Capture Mode During Typical Work Hours as Compared to Leisure Time.

Direct comparison between work and leisure time density yields a very similar patterns to Map 17, with workday concentrations around EPFL and central Lausanne. However, leisure time concentrations extend further outward than workday concentrations or residential concentrations (see Map 17).

Map 19: Regional Overview of Relative Time Spent in Location Capture Mode During Typical Work Hours as Compared to Leisure Time.

The regional view presented here further explores the pattern suggested in Map 18 regarding leisure time extending geographically further than work time activities. Map 19 confirms that users covered much larger and varied areas during the weekend day time window than during the same time Mondays through Thursdays. Notice, for example, the bands of green tracing various roads as well as a wide sprinkling of green clusters in outlying areas.

Map 20: Relative Time Spent in Location Capture Mode During Typical Work Hours, Home, and Leisure Time. Focus on Lausanne Region.

The final two maps illustrate how one could combine the relative density within three activity windows (home, leisure, work) in a single visualization through a trivariate color composite. Yellowish tones indicate a dominance of work time concentration, while cyan points towards mostly home and magenta towards mostly leisure time activities. Regions near EPFL and central Lausanne expectedly are depicted in yellow tones. However, we notice a more nuanced depiction, with the region immediately north of EPFL leaning towards a greenish tint, indicating that this area has plays a secondary role as residential location. Central Lausanne, especially towards the southern half of its work-dominated region, has more of an orange tilt, which points to leisure activities being part of the mix. Meanwhile, the intense blue shading of Prilly indicates an even mix of home and leisure time activity and implies that none of the subjects spent much work time in that area.

Map 21: Relative Time Spent In Location Capture Mode During Typical Work Hours, Home, and Leisure Time. Focus On Geneva Region.

The concentration of MDC subjects in the Lausanne area made it possible to generate a nuanced aggregate picture of temporally sliced activities, as indicated by a range of color tones. In contrast, the Geneva region is depicted in pure tones of cyan, magenta, and yellow. A single cluster each of home and work time activity further points to the likely reason: that only one user was active in Geneva. That user works in central Geneva and resides in a suburb. However, when it comes to privacy concerns, it is important to realize that the generation of density landscapes eliminates any information that could personally identify a particular person. The map merely indicates that *some* MDC user likely works and lives in certain neighborhoods, but information about *who* that user specifically is becomes lost during density computation. Notice in this map also the leisure time cluster bordering Lake Geneva and the faint magenta glow along motorway A1, leading to a temporal cluster at the airport.

4. REFERENCES

[1] Abler, R. 1975. Effects of space-adjusting technologies on the human geography of the future. In *Human Geography in a Shrinking World.* R. Abler, D. Janelle and A., Philbrick, Eds. Duxbury Press.

[2] Couclelis, H. 1996. The death of distance? *Environment and Planning B: Planning and Design*, 24, 387.

[3] Janelle, D. G. 2004. Impact of information technologies. In *The Geography of Urban Transportation*, 2ed. S. Hanson and G. Giuliamo, Eds. Guilford Press.

[4] Skupin, A. 2007. Where do you want to go today [in attribute space]? In *Society and Cities in the Age of Instant Access*, H. J. Miller, Ed. Springer.

[5] Skupin, A. and Miller H. J. 2012. Nokia MDC Atlas (Addendum). [http://research.nokia.com/mdc]