On Inferring the Time-Varying Traffic Connectivity Structures of an Urban Environment

Sanjay Chawla
QCRI and University of Sydney
schawla@qf.org.qa

Somwrita Sarkar
University of Sydney
somwrita.sarkar@sydney.edu.au

Javier Borge-Holthoefer
Qatar Computing Research Institute (QCRI)
borge.holthoefer@gmail.com

Shameem Ahamed
UTC France and QCRI
shameem.ahamed@gmail.com

Hossam Hammady
Qatar Computing Research Institute (QCRI)
hhammady@qcri.org.qa

Fethi Filali
Qatar Mobility Innovation Centre (QMIC)
ffilali@qmic.com

Wassim Znaidi
Qatar Mobility Innovation Centre (QMIC)
wznaidi@qmic.com

ABSTRACT

Road networks shape traffic mobility in a city. These dynamics are often represented as traffic flows in and out of defined urban travel zones. The functional dynamics of traffic zones can be represented by time-dependant correlations between time series of traffic flows in and out of these zones. In this paper we address the question: given the dense time-varying functional correlations of traffic flow in a city, how can we derive a sparse representation that explains the time-varying structural connectivity of traffic zones in a city? We call this sparse representation the *time-varying effective traffic connectivity* of the city. We formulate an optimization problem to infer the sparse effective traffic network from dense functional correlations of traffic flow for arbitrary levels of temporal granularity, and demonstrate the results for the city of Doha, Qatar on data collected from several hundred bluetooth sensors deployed across the city to record vehicular activity through the city's traffic zones. Preliminary experiments suggest that our framework can be used by urban transportation experts and policy specialists to take a real time data-driven approach towards urban planning and real time traffic planning in the city, especially at the level of administrative zones of a city.

1. INTRODUCTION

For the first time in human history more people are living in urban than rural areas[1]. The economies of scale that an urban environment provides creates a natural incentive for urban populations to grow and for other people to migrate to urban centers for employment and a better quality of life [6]. It has thus become compelling and important to understand the dynamics of urban living and manage the growth of cities in order to provide its citizens both opportunity for work and a healthy lifestyle.

(a)

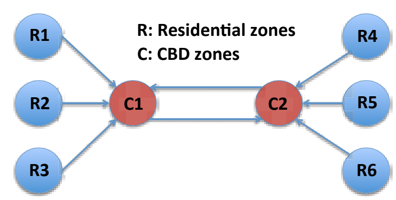

(b)

	C1	C2	R1	R2	R3	R4	R5	R6
C1	0	t_{12}	t_{13}	t_{14}	t_{15}	0	0	0
C2	t_{21}	0	0	0	0	t_{26}	t_{27}	t_{28}
R1	t_{31}	0	0	0	0	0	0	0
R2	t_{41}	0	0	0	0	0	0	0
R3	t_{51}	0	0	0	0	0	0	0
R4	0	t_{62}	0	0	0	0	0	0
R5	0	t_{72}	0	0	0	0	0	0
R6	0	t_{82}	0	0	0	0	0	0

(c)

	C1	C2	R1	R2	R3	R4	R5	R6
C1	1	H	L	L	L	0	0	0
C2	H	1	0	0	0	L	L	L
R1	L	0	1	H	H	H	H	H
R2	L	0	0	1	H	H	H	H
R3	L	0	0	H	1	H	H	H
R4	0	L	0	H	H	1	H	H
R5	0	L	0	H	H	H	1	H
R6	0	L	0	H	H	H	H	1

Figure 1: Concept: Predicting similarity structure by correlations: (a) Traffic flow between residential and CBD zones. (b) The traditional trip distribution matrix. (c) Dense zonal traffic flow (time-series) correlation based structure; H = High, L = Low. Static trip distributions do not capture high level similarity between traffic characteristics of zones. In this paper, the dynamic sparse optimal structure of the dense (c) is inferred in real time, for arbitrary temporal granularity.

The availability of highly granular sensor data which records urban activity, often at very fine spatial and temporal scales, provides a timely opportunity to understand the patterns of an urban environment in a deeper way than is possible through traditional census tracts (that are temporally long term) or detailed surveys (that provide only samples of populations). For example, the movement of vehicles is now

[1]http://www.un.org/en/development/desa/index.html

UrbComp'15, August 19th, 2015, Sydney, Australia

electronically recorded at very fine spatio-temporal granularity and sometimes made publicly available.

An important characteristic of any urban region is the emergence of correlations between activities and traffic flows in its subregions or zones during certain time periods. For example, there will be temporal (lagged) positive correlations between suburbs and the Central Business District (CBD). Similarly, in a polycentric city with multiple CBDs, each CBD may have very similar spatial patterns of correlated activity. In other words, the high level pattern we wish to understand is: which areas or zones of a city behave similarly, spatially or temporally? If these high level patterns are understood for arbitrary levels of fine temporal granularity, real time traffic planning will gain a valuable tool.

Travel demand in the city is often inferred using a four-step model, inferring: (a) Trip generation in N traffic zones, or the total number of trips originating and ending in zone $i = 1 \ldots N$, (b) Trip distribution, or the number of trips actually going from zone i to zone j, $i, j = 1 \ldots N$, (c) Modal split, or the proportion of trips split according to the different transportation modes, and (d) Trip assignment, the actual path through the spatial networks for trips originating in zone i and ending in zone j.

Policies and interventions in traffic and transportation planning are then planned using this demand. However, traditionally, OD matrices are developed over long time spans, and are typically useful for long term traffic planning. To compute demand, demographic data other than traffic flow counts is used, and assumptions (such as the well known gravity model) are used to compute trip distribution [Fig. 1](b). With new types of sensor data, real time OD matrices can be developed for short term traffic planning on a day to day basis. Further, for arbitrarily fine levels of temporal granularity, even when a full dynamic OD cannot be developed, it could be useful to understand the correlation structure of traffic flow through travel zones [Fig. 19(c)], and as shown in Fig. 1, could provide high level insights into the short term time-varying similarity of how travel zones behave dynamically, which a long term trip distribution matrix may not provide per se.

In this paper, we develop an approach where real time traffic data is used to compute dense functional correlations of traffic flow between administrative zones of a city. This dense functional dynamics signature is then used to infer a sparse optimal effective traffic connectivity structure, varying over time, providing a dynamic picture of which zones in the city are most correlated by traffic flow during different times of an average weekday.

We do not produce traditional OD matrices (though in the future our work can be extended to develop dynamic real time ODs), but focus on the related problem of inferring the sparsest structure that can explain the dense functional correlations of traffic flow in different zones in the city, at arbitrary levels of temporal granularity, and inferring which areas in the city are most similar in terms of their traffic flow characteristics.

Specifically, we formulate a sparse optimization problem to infer the effective structural connectivity of an urban traffic network (Doha, Qatar) from its functional data. Figure 2 captures the workflow of our approach: (i) The city of Doha is divided into structural/physical zones and data from bluetooth sensors in the zones will be used to form zone time series which record traffic activity in and out of the zone, (ii)

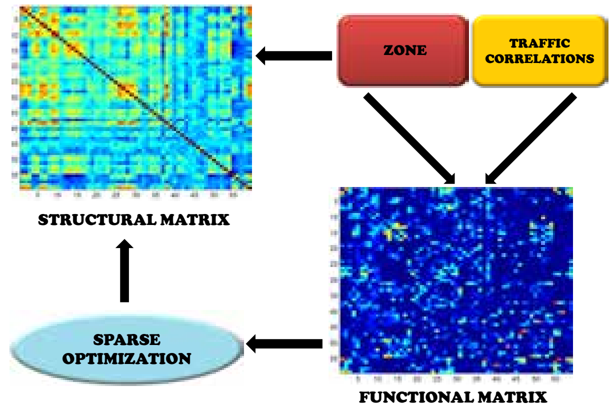

Figure 2: Inferring the Effective Structural Connectivity from Functional Data. Traffic sensor data from different zones is aggregated and correlated to describe the functional connectivity. A sparse optimization problem is then used to elicit the latent effective structural connectivity matrix.

the recorded time series (at different aggregation levels) will be the basis of a zone by zone functional correlation matrix which describes the global functional connectivity between any pairs of zones, (iii) we will formulate a sparse optimization problem where the unknown variable will be the effective structural matrix which is *minimal*, non-negative and best explains the observed functional matrix.

The inferred structural matrix is a highly sophisticated time-varying summary of a city's activity. It can be derived for different time granularities ranging from a few minutes to days and years. The matrix can be queried by decision makers and end users to explore how activity in nearby and distant zones tend to influence and impact each other. To the best of our knowledge this is the first work in the urban computing literature which takes a holistic, mathematically rigorous and global view of explaining the functional connectivity of a city.

The rest of the paper is structured as follows. In Section 2 we overview different forms of connectivity that are prevalent in the literature on network analysis. In Section 3 we formally define the optimization problem and discuss our design choices. After a brief overview of Alternating Direction Method of Multipliers (ADMM), in Section 4, we provide the complete derivation of the iterative algorithm to solve the optimization problem. In Section 5 we provide preliminary experimental results to validate our approach. In Section 6 we cover related work and we conclude in Section 7 with a summary and a road map for future work.

2. FORMS OF URBAN CONNECTIVITY

Different forms of connectivity can exist in an urban environment. A city is characterized by its movement behavior. For example, most monocentric cities have a central downtown area where major businesses and government offices are located. There is vehicular traffic coming into the downtown area in the morning and leaving in the afternoon or evening. Typically the morning rush hour peak into the downtown is more sharply concentrated than the afternoon/evening one, which is more spread out in time. Over time, a city can expand and other central areas can emerge, leading to a more polycentric city structure. For example, in Doha, besides the central downtown area (known as West Bay), several other areas including around the airport, the industrial area and Qatar Foundation have emerged. Often there is a complex causal relationship between the road network in a city and how new pockets of urban activity arise. Sometimes, a new areas will emerge on the sidelines of a road artery, other times roads are constructed to service active areas in the city. We will define three forms of connectivity: *structural, functional* and *effective* [8]. These connectivities are well defined and established in biological spatio-temporal systems modeling, such as the brain, but their introduction will be equally relevant to urban spatio-temporal modeling.

Structural The structural connectivity is defined by the physical road network and spatial locality. It could be defined on the basis of structural adjacency or contiguity of two spaces, but also equally in terms of ease of access of one place to another in the city. For example, two contiguous zones are structurally connected but so are two zones that are connected directly by a highway. The complete road network in conjunction with the spatial contiguity/adjacency relationships characterizes the structural connectivity of a region.

Functional The functional connectivity is defined by the actual transportation movement dynamics supported on top of the structural layer. Depending upon the time scale used for analysis, two distant zones can be either correlated or uncorrelated. For example, two areas that are far from each other and connected only through several intermediate locations or stops, but hosting complementary activities generating a lot of connected traffic flow would show high functional correlation, even if their structural contiguity score is low by both distance or ease of accessibility. The functional correlation can be both positive and negative. Functional connectivity is derived using time series data. For example, in our particular case, statistical correlation will be extracted using the data recorded by bluetooth sensors. Many zones in Doha have bluetooth sensor embedded at different road intersections. As vehicles (which contain bluetooth devices, like smartphones) go past sensors, a counter is incremented. As is well known, statistical correlation between two sensors in zones is not an indication of causality. However, an important observation is that the functional connectivity data tends to be dense but of low effective matrix rank. We will leverage this property in our problem formulation.

Effective The time-dependant effective or *effective structural* connectivity is the latent and generative network process which explains the functional connectivity of the system in terms of the structural connectivity. While the physical structural connectivity is static, the effective structural connectivity changes with time and has both a periodic and a non-stationary component. The easiest way to explain this is by observing that even though the road network represents a static physical capacity, effective structural connectivity for the same road network is different when the network is hosting low flow and when it is hosting a traffic jam or congestion situtation. In the first case, effective connectivity is lower than the structural, in the latter case, it is higher than what is permitted by the structural. For example, in an urban setting, during morning rush hour a certain part of the network activity is highly correlated. During off-peak hours, the traffic activity diffuses throughout the network and then concentrates again during the evening rush hour. However there are pockets of connectivity between other zones in the urban environment which have their own dynamic. For example, the airport maybe the center of another effective connectivity process which interacts with other generative processes. The aggregation of all *latent* and *generative* activity is termed as the effective structural connectivity. The main objective of this work indeed is to infer the effective structural connectivity that characterizes the city of Doha.

3. PROBLEM DEFINITION

To set up the problem we begin with a graph $G = (V, E)$ to represent zones in a city. Two nodes u and v have an edge between them if they are neighbors. Associated with each node v is a time series $S_v(t)$ to represent the time varying activity at the node, for example the number of vehicles crossing the node v at time t.

Let F_s^I be the $|V| \times |V|$ functional correlation matrix derived from the time series $S_v(r)$ at time aggregation level s in time interval I. Each element of $F_s^I(u, v)$ represents the correlation between u and v. Going forward we will just represent F_s^I as F.

Using F as an input, our objective is to infer a *minimal* effective structural matrix X which captures the correlations in F. The dimensionality of X is the same as that of F.

A key modeling decision that has to be made is to hypothesize the relationship between the F and the X matrix. Since X is supposed to capture the effective correlation in F, we can treat elements of X as the building blocks (or basis elements) which combine together to produce F. We thus hypothesise that

$$FX = F$$

In the computer vision community, Elhamifar [3] refers to this relationship as the self-expression property of high dimensional data, i.e., even though sensor data tends to be high dimensional, it usually lives in a union of low-dimensional affine spaces and each point in the affine space can be expressed as a linear combination of other points that share the space.

To make the discussion more concrete consider the following 2×2 matrix

$$\begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix}$$

We would like to infer the X matrix given by

$$\begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{pmatrix}$$

such that $FX = F$. Now FX is

$$\begin{pmatrix} f_{11}x_{11} + f_{12}x_{21} & f_{11}x_{12} + f_{21}x_{22} \\ f_{21}x_{11} + f_{22}x_{21} & f_{21}x_{12} + f_{22}x_{22} \end{pmatrix}$$

In order to prevent self-referencing, i.e., we would like to avoid an element being represented in terms of itself, we will force x_{11} and x_{22} to be zero. Now, if we minimize the sum of the square loss between elements of FX and X and solve for x_{12} and x_{21} we obtain

$$x_{12} = \frac{2f_{11}f_{12}}{f_{11}^2 + f_{12}^2} \quad (1)$$

$$x_{21} = \frac{2f_{21}f_{22}}{f_{22}^2 + f_{21}^2} \quad (2)$$

$$(3)$$

However, our aim is not just to choose an X which minimizes the difference between F and FX but we would like to obtain a sparse and non-negative X to make our solution interpretable. Sparsity can be enforced by minimizing the ℓ_1 norm of X. However, in order to achieve a solution that is both sparse and non-negative we need to design an iterative algorithm as a closed form solution is not achievable.

However, we carry out a further simplification by working with a low rank representation of F rather than the full $|V| \times |V|$ matrix. The rationale for using a low-dimensional version is as follows: Row i of F captures the connectivity of node i with all the other connections. Due to the nature of road traffic distribution many other nodes will have a similar form of connectivity as node i. Thus we seek a representation where a node can be represented as a point in $k-$dimensional space, with $k << N$, with the property that if two nodes have a similar connectivity then they should be close to each other in the $k-$dimensional space. This can be achieved in several ways, but here we use the simplest possible representation: the spectral representation of F, or the co-ordinates given by the first k eigenvectors of F; $F = VDV^T$, $F_k = V_k D_k V_k^T$, where V_k which is $k \times N$ now represents the positions of N nodes in $k-$dimensional space $\mathbb{R}^k$. When the transformation X is applied to F, the components of sparse X will be like weights on the columns of V_k, "picking out" the most relevant connections of node i to all other nodes. Thus we have

$$V_k X = V_k \quad (4)$$

.

This is also the simplest possible representation where the following constraints can be captured: (a) high functional connectivity must imply high probability of structural connectivity, but (b) the structural connectivity must be sparse, since only direct connections must be inferred, and (c) both the short range and the long range structural connectivity must be inferred using the functional connectivity.

Thus, the final optimization problem can be stated as

$$\begin{aligned} \min\ & ||X_p||_1 \\ \text{sub to } & V_k(X_p) = V_k \\ & diag(X_p) = 0 \\ & X_p \geq 0. \end{aligned} \quad (5)$$

Note that, similar to [3], we want to enforce the diagonals of the solution variables to be 0, so as to avoid the trivial solution of each node expressing itself as its own linear combination and none of the others.

4. OPTIMIZATION ALGORITHM

We solve the above optimization problem using *Alternating Direction Method of Multipliers* (ADMM) framework [1]. In the next section, we give a brief introduction to ADMM framework and work out the different intermediate steps in solving a problem in ADMM framework.

4.1 Alternating Direction Method of Multipliers

ADMM is an extension of lagrangian method for solving a class of constrained optimization problems. Like in the lagrangian method, in ADMM the constraints are moved into the objective and multiplied by a penalty term which captures the "price" of a solution not satisfying the constraint. Thus the original constrained problem is transformed into a series of unconstrained optimization problem. However, ADMM has two additional variations over the standard lagrangian method: (i) an additional strongly convex penalty term (often quadratic) is added to the objective to help improve the convergence of the iterative solution and (ii) a new variable is introduced to make the objective separable and a consensus constraint is added which forces the old and the new variable to agree at convergence. For example, consider the following constrained optimization problem:

$$\begin{aligned} \text{minimize} \quad & f(x) + g(y) \\ \text{subject to} \quad & Ax + By = C \end{aligned} \quad (6)$$

Augmented Lagrangian for the above equation is

$$L(x, y, \alpha) = f(x) + g(y) + \alpha^T(Ax + By - C) + \frac{\rho}{2}\|Ax + By - C\|_2^2$$

ADMM iterates through the following three steps until the dual variable converges with ϵ approximation

$$\begin{aligned} x^* &= \underset{x}{\text{argmin}}\ L(x, y, \alpha) \\ y^* &= \underset{y}{\text{argmin}}\ L(x^*, y, \alpha) \\ \alpha^* &= \alpha + \rho(Ax^* + By^* - C) \end{aligned}$$

4.2 Solving Our Problem Using ADMM

We present an ADMM based solution framework for solving the optimization problem given in Equation 5. The complete algorithm is shown in Algorithm 1. We introduce auxiliary variable A corresponding to the optimization variable X_p, and indicator function $I_+(X_p)$ for the non-negative constraint of X_p, we have:

$$\begin{aligned} \text{minimize } & L[A, X_p, \Delta_1] = \\ & \|X_p\|_1 + \frac{\lambda_t}{2}\|V_k - AV_k\|_2^2 + \\ & \text{tr}[\Delta_1^T(A - (X_p - diag(X_p)))] + \\ & \frac{\rho_1}{2}||A - (X_p - diag(X_p))||_2^2 + I_+(X_p), \end{aligned} \quad (7)$$

where $I_+(X_p) = 0$ when $X_p \geq 0$ and ∞ otherwise.

Now we solve this using a standard iterative ADMM process, in which, we minimize L by differentiating it with respect to one primary variable at a time by keeping others constant (partial gradient equating to zero, since the problem is convex) and followed by updating the Lagrange multiplier (dual variable) Δ_1. We represent with A^*, X_p^*, Δ_1^*, the updated variables corresponding to A, X_p, Δ_1 at each iteration.

4.2.1 Updating A

A can be updated by computing its gradient and setting it to zero. Since L is differentiable, considering terms containing A only, we get

$$\nabla_A L = \nabla_A \left[\frac{\lambda_t}{2} \|V_k - AV_k\|_2^2 + \frac{\rho_1}{2} \|A - X_p + diag(X_p)\|_2^2 + tr\left[\Delta_1^T (A - X_p + diag(X_p))\right] \right]$$

$$\begin{aligned}\nabla_A \|V_k - AV_k\|_2^2 = & \\ &= \nabla_A \mathrm{tr}\big((V_k - AV_k)^T (V_k - AV_k)\big) \\ &= \mathrm{tr}\Big(\nabla_A \big(- V_k^T A V_k - V_k^T A^T V_k + V_k^T A^T A V_k\big)\Big) \\ &= -2V_k V_k^T \ + \ 2AV_k V_k^T\end{aligned}$$

Similarly,

$$\begin{aligned}\nabla_A \|A - X_p + diag(X_p)\|_2^2 \ = & \\ &= \mathrm{tr}\Big(\nabla_A \Big(A^T A - A^T X_p + A^T diag(X_p) - X_p^T A + diag(X_p) A\Big)\Big) \\ &= 2A - 2X_p + 2\ diag(X_p)\end{aligned}$$

Lastly,

$$\nabla_A \Big(\mathrm{tr}\left[\Delta_1^T (A - X_p + diag(X_p))\right] \Big) = \Delta_1$$

Combining all terms and setting it to zero, we get

$$A^* = (\lambda_t V_k V_k^T + \rho_1 I)^{-1} [\lambda_t V_k V_k^T + \rho_1 X_p - \Delta_1]$$

4.2.2 Updating X_p

$$\nabla_{X_p} L = \nabla_{X_p} \left(\|X_p\|_1 + \mathrm{tr}\left[\Delta_1^T \Big(A - \big(X_p - diag(X_p)\big)\Big)\right] + \frac{\rho_1}{2} ||A - (X_p - diag(X_p))||_2^2 \ + I_+(X_p) \right)$$

$$\nabla_{X_p} \|X_p\|_1 = \mathrm{sign}(X_p). * \mathbb{I} \ \ (X_p \neq 0)$$

In the above .* indicates the entrywise product and sign function is applied entrywise.

$$\nabla_{X_p} \mathrm{tr}\left[\Delta_1^T \Big(A - \big(X_p - diag(X_p)\big)\Big)\right] = - \Delta_1 + diag(\Delta_1)$$

$$\nabla_{X_p} \|A - X_p + diag(X_p)\|_2^2 = - 2A + 2\ diag(A) + 2X_p$$

combining the individual terms and equating it to zero, we get

$$\mathrm{sign}(X_p) - \Delta_1 + diag(\Delta_1) + \rho_1(-A + diag(A) + X_p) = 0$$

Now, considering the positivity constraint of X_p, we get

$$X_p^* = \frac{1}{\rho_1} \Big[\rho_1 \big(A - diag(A)\big) + \Delta_1 - diag(\Delta_1) - \mathbf{1}\mathbf{1}^T \Big]_+$$

4.2.3 Updating Δ_1

Finally, we update the dual variables as

$$\nabla_{\Delta_1} L = A - X_p$$

We update dual variable as mandated by the ADMM procedure as,

$$\Delta_1^* = \Delta_1 + \rho_1 (A^* - X_p^*),$$

Algorithm 1 ADMM Algorithm

Input: F, k, λ_t, ρ_1
Output: X_p
Initialize X_p, A, Δ_1
1: $F \leftarrow VDV^T$
2: Use V_k from $V_k D_k V_k^T$ (first k eigenvectors)
3: **repeat**
4: $A \leftarrow (\lambda_t V_k^T V_k + \rho_1 I)^{-1} [\lambda_t V_k^T V_k + \rho_1 X_p - \Delta_1]$
5: $X_p \leftarrow \frac{1}{\rho_1} \Big(\rho_1 (A - diag(A)) + \Delta_1 - diag(\Delta_1) - \mathbf{1}\mathbf{1}^T \Big)_+$
6: $\Delta_1 \leftarrow \Delta_1 + \rho_1 (A^* - X_p^*)$
7: **until** convergence

5. EXPERIMENTS

We have carried out preliminary experiments on both synthetic and real world dataset to demonstrate the effectiveness of our proposed inference scheme. We used synthetic data to demonstrate the general applicability of our mathematical formulation based on ADMM to retrieve a sparse representation of data. For the real world data, we used bluetooth sensor data of Doha city provided by QMIC [2]. We briefly describe our dataset and provide empirical results at the end.

5.1 Synthetic Data

We create synthetic data based on a simple formulation based adjacency matrix. When we represent a graph using adjacency matrix, the k^{th} power of an adjacency matrix gives us the number of paths of length k between two nodes. Since the observed functional connectivity of any network is a result of the dynamics (traffic flow) occurring on both direct as well as indirect paths between nodes, a good approximation to the functional connectivity would be the sum of the powers of the adjacency matrix. Starting with a synthetically generated dense functional matrix, we show that our mathematical formulation retrieves an approximate sparse matrix which captures most of the original data in the adjacency matrix, and thus enables inferring vital information contained in the original matrix.

Given a sparse adjacency matrix S, we create a dense representation F as

$$F = S + S^1 + S^2 + \ldots + S^i$$

We multiply the value of the original matrix S by a constant term, $\frac{0.9}{\sigma}$, such that the sum converges for large values of i.

[2] http://www.qmic.com/

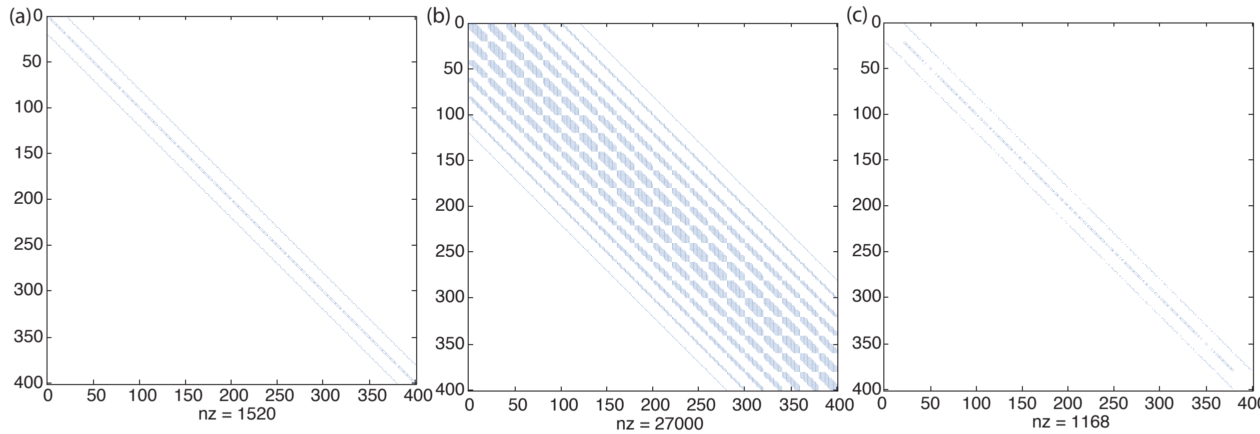

Figure 3: Synthetic Data Experiments. (a) Original sparse matrix. (b) Generated dense matrix. (c) Inferred sparse matrix that accurately reconstructs (a), the original sparse matrix.

Here σ denotes the largest singular value of the matrix S i.e. the positive square root of the largest eigen value of the matrix $S^T S$.

We retrieved an almost identical sparse representation from the dense matrix. Precision of our sparse representation was 100% and recall was 76.8%. The plots of the sparsity pattern of the original sparse matrix S, derived dense matrix F and derived sparse matrix S' is given in Figure 3.

5.2 Traffic Data

The traffic data consists of sensor readings from different locations of the city of Doha. The city is officially divided into 66 zones, and the zones are equipped with bluetooth sensors which logs the vehicle traffic in the corresponding zones. Whenever a vehicle equipped with a bluetooth device passes a zone sensor, the bluetooth id is recorded. This approach not only allows us estimate the aggregate count of vehicles in the zone during a time interval but we can also track if a particular vehicle appears in another zone after a certain time lag. The zone details of the Doha city is given in Figure 4. Except for few zones, all zones are equipped with at least one sensor. For our experiments, we excluded zones which either did not have a sensor or where the sensor was faulty. As a result we used traffic data from 59 zones. Sensor data was represented as a time series containing the bluetooth counts for every second. In practice, working with data of such small granularity (here per second data) can be problematic as the inherent random noise can overwhelm the signal. Also, traffic undergoes different phases during the day which suggests the relevant, informative time-scales are dynamic: one can not study a large city always at the same temporal resolution, because the traffic "accelerates" and "calms down" at different times. Sometimes having a tick in the time series every 2 seconds is enough to capture patterns, sometimes time scales are longer.

Here it will be important to differentiate between short range and long range spatio-temporal effective connectivities. Intuitively, we expect that shorter time scales will capture local spatial and temporal traffic dynamics while the longer time periods will capture long range spatial and temporal traffic dynamics. It is worth noting that even the definition of short and long time scales is hierarchically organized: whole day traffic data is long term with respect to a 2 second time window data, but short term with respect to a weekly, monthly, or yearly time window data.

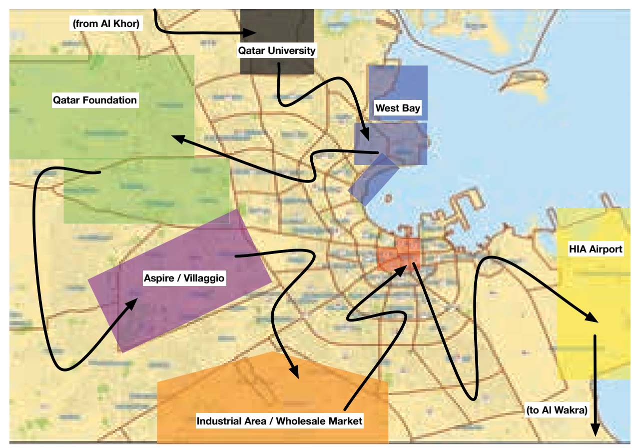

Figure 4: The zones of Doha City. The black arrows indicate the pattern that was used to linearly order the zones. Important areas of the city (e.g., West Bay and Airport) are highlighted.

We note here that our analysis can be done for both *(i)* short term effective structural connectivity and *(ii)* long term effective structural connectivity, where short term and long term can be relatively defined with respect to each other. The short term connectivity in hours between zones might explain the short-term spatial and temporal latent causal relationships between road traffic and other dependent factors. For example, a small temporary closure of a service road, maintenance work etc. may lead to small local diversions for a limited span of time. Long term effective structural connectivity over weeks, months, or years, gives a bigger picture of the evolution of the relationship between road traffic and other entities. It can clearly provide an indication of the rise of new pockets of urban activity in the city.

In this work, we focus on a two-hour window traffic flow data to form the functional connectivity matrix. This will provide us with an effective connectivity picture that is for the spatial range of the entire city, and for a temporal range of daily weekday variations of traffic flow. We averaged the data over a period of two hours with one hour overlap i.e. a sliding window of two hours with one hour overlap between consecutive frames. Thus for each two hour window with one hour overlap, for a day, we have total of 59 vectors each with size 23 containing the aggregated traffic data.

We denote the functional matrix representing the inter-traffic dependencies between zones using F_i $(i = 1 \ldots 23)$ measured by the correlation of hourly aggregated data. We aim to derive the effective structural connectivity matrix from the given functional matrix. The functional matrix explained above captures the traffic flow among different zones for the two hour window and may contain dependent components. We approximate the original function matrix F_i to low ranked approximation matrix preserving the 95% of the information contained in the original matrix using spectral decomposition. We empirically estimate the number of dimensions in the low rank approximation to be 44. We apply our algorithm onto the low rank approximation of F_i, and perform the inference of the resulting time-varying effective structural connectivity signatures over the overlapping two hour windows.

5.3 Results

Figure 5 visualizes the time-varying effective structure evolution of Doha traffic over a day. Since effective connectivity is a matrix, to visualize it spatially, we can define one zone as a "query" region (colored blue in Fig. 5), and use a color scale to define a chloropleth map where the color of each zone depicts the derived effective connectivity at a particular time point for the "query" region. The plots depict the derived sparse effective connectivity of the zones in Doha, showing the latent relationships among different zones at specific points of time with respect to a reference point. We considered a zone in West-Bay area, the business and financial center of Qatar, as our reference point, for demonstration.

At 5AM, early morning of the day, the effective connectivity between West-Bay and northern neighbourhood are very strong, indicating a strong correlation between West-Bay and public beach area including the area of Lusail and Pearl. This connectivity can be explained as a result of strong activity in the Lusail area where construction is in full swing for FIFA world cup 2022. The effective connectivity of the West-Bay and the public beach area carries over for the next hour as well. By 7AM and 8AM, the effective connectivity concentrates just around West-Bay area, and other connectivity evolves between West-Bay and other residential areas like "Al Sad". The effective connectivity between "Al Sad" and West-Bay can be explained as the daily working class commuting.

By 11 in the morning the effective connectivity subsides and concentrates again over the West-Bay area. The connectivity gets stronger by mid-noon between the neighbouring area and West-Bay area. The couple of hours just after mid-noon keeps the same type of effective connectivity. We can see patterns similar to the office starting hours (7AM-9AM) at the office end hours (3PM-5PM).

Here, we explained the connectivity with respect to a single fixed point. However, any zone can be picked up as a query zone, and the resultant effective connectivity visualized. Thus, we can do the same for multiple focal points like University, Industrial Area etc. In our experiments, we could infer university timing connectivity pattern for the zones corresponds to Qatar University and Qatar Foundation Education City.

In addition to the regular patterns, we could also infer some uncommon activities developed during specific time period at specific zones. Since our data does not include an activity log of the events, a one-to-one mapping between the patterns and activity is hard to establish. But our model captures the heterogeneous dynamics of the zone-traffic data.

A main point to note is that the derived effective structural connectivity is sparse, as opposed to the dense functional dynamics. Thus, this analysis performed for arbitrary short, medium or long time spans will provide different perspectives of time-varying effective connectivities for any query zone for any time window. This will be particular useful to understand how the local and global dynamics of the traffic of the whole city is behaving with respect to a single area in the city.

6. RELATED WORK

The literature on Urban Computing is rapidly growing and blends research methodologies from diverse areas including data mining, machine learning, transportation, visualization, data management,urban planning, network analysis. See the recent special issue in ACM TIST [12]. A comprehensive survey article gives a broad overview of the area [13].

The elicitation of effective structural connectivity as proposed in this paper, has been more actively pursued in computational neuroscience [8, 5, 4]. However, even here, most approaches are derived using the physiological basis in the domain in conjunction with techniques related to community detection and network analysis [8]. In the data mining and ML literature, the work on inferring functional connectivity from fMRI data is partially relevant[9, 2, 7]. Note however our work is inferring the structural from functional connectivity.

There has been some work on inferring the functional regions of city. For example, Yuan et. al. [11] have proposed the DRoF framework which uses techniques from topic modeling to predict the function of a zone in a city based on mobility patterns. Similarly the work by Toole et. al. [10] uses spatio-temporal change detection in mobile phone activity to infer land use of regions and zones in a city. Integrating inference of functional regions in our framework can be potential direction of future work.

7. CONCLUSION AND FUTURE WORK

The paper proposes a novel optimization based framework to infer the effective structural connectivity of Doha, Qatar using functional data acquired from bluetooth sensors. The optimization problem is based on the premise that a latent structural connectivity gives rise to the functional observations as measured by correlation between activity time series between different zones. The inferred effective connectivity is a sparse time-varying explanation of how traffic between different zones of Doha interact. For example, by fixing one zone as the query region we can clearly see the evolution of the effective structure over time with respect to the fixed zone. The effective structural matrix consists of both short

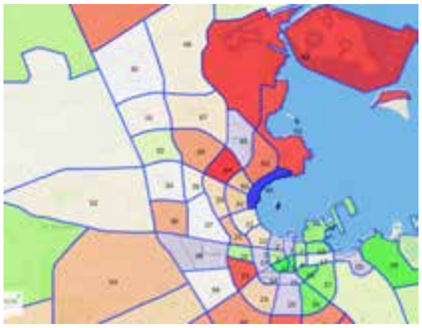

(a) Effective Connectivity 5AM

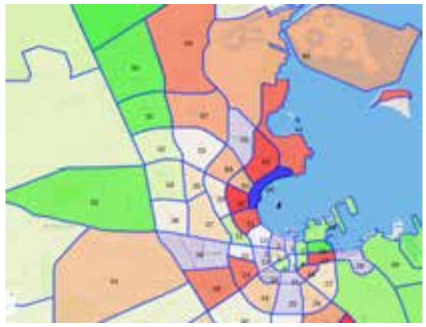

(b) Effective Connectivity 6AM

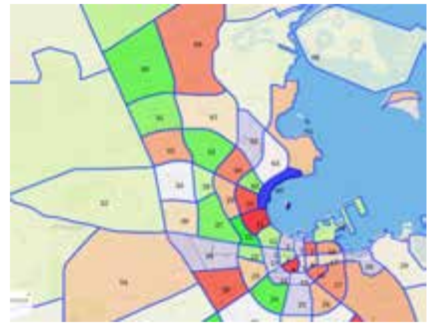

(c) Effective Connectivity 7AM

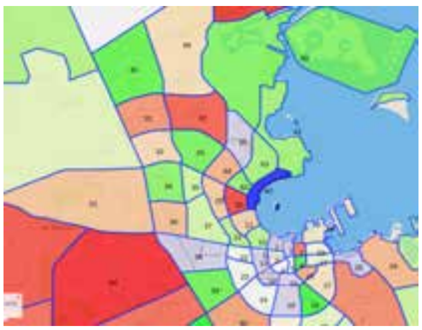

(d) Effective Connectivity 8AM

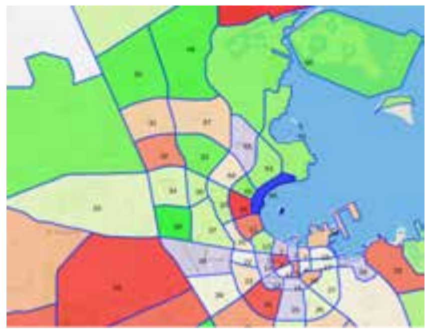

(e) Effective Connectivity 9AM

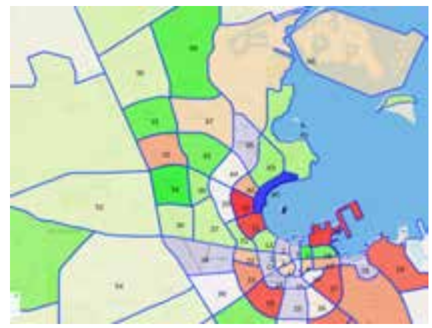

(f) Effective Connectivity 10AM

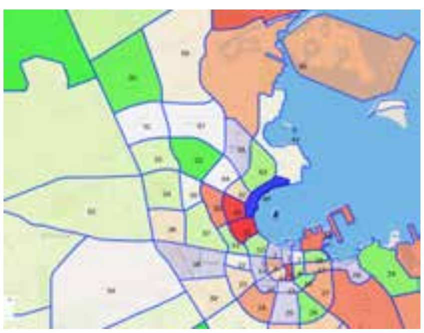

(g) Effective Connectivity 11AM

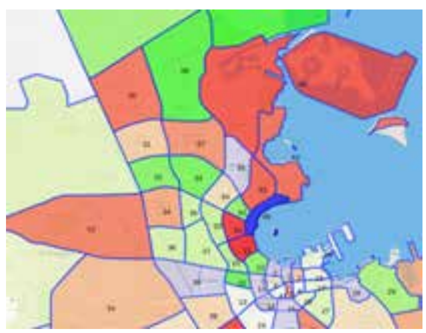

(h) Effective Connectivity 12PM

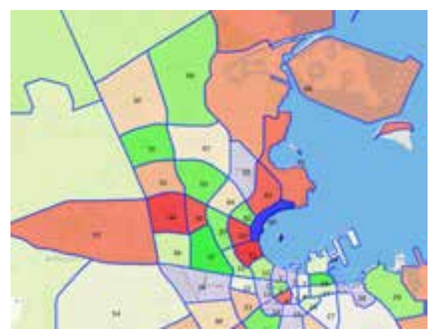

(i) Effective Connectivity 1PM

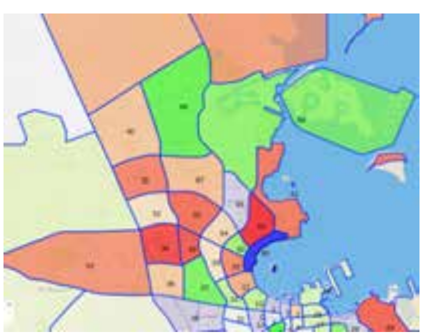

(j) Effective Connectivity 2PM

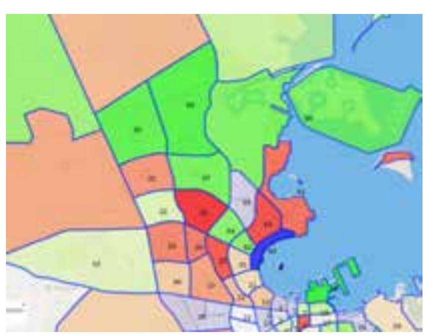

(k) Effective Connectivity 3PM

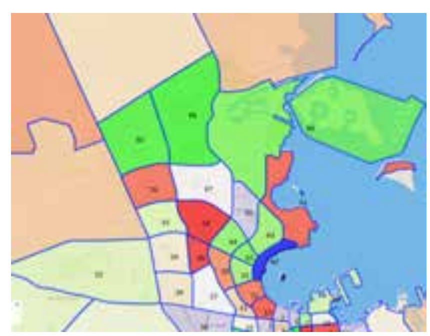

(l) Effective Connectivity 4PM

Figure 5: Doha zone traffic effective structure evolution over a day. The blue zone is the query zone. Red signifies high connectivity, green signifies low connectivity. It can be seen that in the morning rush hour, parts (b) to (e), strong positive structure is observed between the western suburbs and the CBD. Further, throughout the day, positive structure is observed in regions closer to and around the CBD, for example, Lusail and the airport zones to the north and south respectively.

range and long range dependencies as would be expected in an urban environment where different dynamics are simultaneously at play.

At the moment the experiments provide a preliminary and qualitative validation of our approach. For future work we will design a suite of experiments (both synthetic and real) to both calibrate the model (e.g., the choice time scale and aggregation level) and determine if the structural matrix inferred has properties which are physically realizable. One approach would be to relate it to a forward dynamical model rooted in the physics of transportation theory. Another possible application of our approach is to help model and simulate traffic staggering mechanisms as a solution to alleviate rush hour congestion.

8. REFERENCES

[1] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3:1–122, 2011.

[2] I. Davidson, S. Gilpin, O. Carmichael, and P. Walker. Network discovery via constrained tensor analysis of fmri data. In *ACM SIGKDD*, pages 194–202. ACM, 2013.

[3] E. Elhamifar and R. Vidal. Sparse subspace clustering: Algorithms, theory and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2765–2781, 2013.

[4] R. T. Gray, C. K. C. Fung, and P. A. Robinson. Stability of random brain networks with excitatory and inhibitory connections. *Neurocomp.*, 72:1565–1574, 2009.

[5] X. Kong and P. S. Yu. Brain network analysis: a data mining perspective. *SIGKDD Explorations*, 15(2):30–38, 2013.

[6] P. Krugmann. *Geography and Trade*. MIT, 1991.

[7] E. E. Papalexakis, A. Fyshe, N. D. Sidiropoulos, P. P. Talukdar, T. M. Mitchell, and C. Faloutsos. Good-enough brain model: Challenges, algorithms and discoveries in multi-subject experiments. In *ACM SIGKDD*. ACM, 2014.

[8] O. Sporns. *Networks of the brain*. MIT Press, 2011.

[9] L. Sun, R. Patel, J. Liu, K. Chen, T. Wu, J. Li, E. Reiman, and J. Ye. Mining brain region connectivity for alzheimer's disease study via sparse inverse covariance estimation. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, June 28 - July 1, 2009*, pages 1335–1344, 2009.

[10] J. L. Toole, M. Ulm, D. Bauer, and M. C. González. Inferring land use from mobile phone activity. *CoRR*, abs/1207.1115, 2012.

[11] J. Yuan, Y. Zheng, and X. Xie. Discovering regions of different functions in a city using human mobility and pois. In *The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012*, pages 186–194, 2012.

[12] Y. Zheng, L. Capra, O. Wolfson, and H. Yang. Introduction to the special section on urban computing. *ACM Trans. Intell. Syst. Technol.*, 5(3):37:1–37:2, July 2014.

[13] Y. Zheng, L. Capra, O. Wolfson, and H. Yang. Urban computing: Concepts, methodologies, and applications. *ACM Trans. Intell. Syst. Technol.*, 5(3):38:1–38:55, Sept. 2014.