Social and Spatial Ethnic Segregation:

A Framework for Analyzing Segregation With Large-Scale Spatial Network Data

Joshua Blumenstock
University of Washington
Seattle, Washington
joshblum@uw.edu

Lauren Fratamico
University of California, Berkeley
Berkeley, California
fratamico@berkeley.edu

ABSTRACT

While ethnic segregation plays an important role in determining the development trajectories of many countries, empirical measures of the dynamics of segregation remain rudimentary. In this paper, we develop a new computational framework to model and measure fine-grained patterns of segregation from novel sources of large-scale digital data. This framework improves upon prior work by providing a method for decomposing segregation into two types that previous work has been unable to separate: *social segregation*, as observed in interactions between people, and *spatial segregation*, as determined by the co-presence of individuals in physical locations. Our primary contribution is thus to develop a set of computational and quantitative methods that can be used to study segregation using generic spatial network data. A secondary contribution is to discuss in detail the strengths, weaknesses, and implications of this approach for studying segregation in developing countries, where ethnic divisions are common but data on segregation is often plagued by issues of bias and error. Finally, to demonstrate how this framework can be used in practice, and to illustrate the differences between social and spatial segregation, we run a series of diagnostic tests using data from a single city in a large developing country in South Asia. The case study we develop is based on anonymized data from a mobile phone network, but the framework can generalize easily to a broad class of spatial network data from sources such as Twitter, social media, and networked sensors.

Categories and Subject Descriptors

K.4.4 [**Computers and Society**]: General; J.4 [**Social and Behavioral Sciences**]: Economics; J.4 [**Social and Behavioral Sciences**]: Sociology

General Terms

Human Factors, Measurment, Economics, Algorithms

DEV-4 December 6-7, 2013 Cape Town, South Africa

Keywords

Ethnic segregation, integration, ethno-linguistic fractionalization, call detail records, development

1. INTRODUCTION

Ethnic diversity, segregation, and fractionalization have long been thought to play a critical role in the socioeconomic structure and overall stability of many developing countries. Studies based on cross-country regressions generally point to a detrimental effect of ethnic diversity on economic performance [17, 12]. These results are not limited to developing countries [2, 3], but the negative impact of fractionalization are particularly acute in countries with weak political and legal institutions [19]. As a result, a great deal of attention has focused on polices to promote integration and interaction in heterogeneous societies [9].

However, research on ethnic segregation has been hindered by a lack of reliable data. Posner [28], for instance, notes that "most measures of ethnic diversity... are inappropriate for testing [the effect] of ethnic diversity on economic growth" (p.849). A recent review article by Blattman [4] similarly concludes with a "plea for new and better data" (p.3). As these, and others, point out, a weakness of this empirical literature on segregation is a general over-reliance on standard government survey data. Census data, for instance, can capture coarse patterns of physical settlement, but rarely tracks patterns of social interaction which are necessary to develop a more nuanced understanding of the micro-dynamics of social interaction and cohesion.

In this paper, we develop and present a set of methods for measuring and analyzing ethnic segregation at a level of detail that has not been achieved in the prior literature. In addition to adding specificity and resolution to the geographic analysis of segregation, a key advantage of our approach is that it permits a comparative analysis of two distinct types of segregation that most prior work has conflated: *social* segregation, or the extent to which individuals of different ethnicities are observed to interact with one another, and *spatial* segregation, or the extent to which individuals of different ethnicities visit the same locations.

The fact that these two forms of segregation are rarely differentiated in the literature is striking, since the differences between social and spatial segregation can have important implications for public policy. To take an extreme example, resettlement programs and other policies designed to promote physical integration may be less effective in places where only a weak correlation exists between social and spatial segregation. By contrast, in contexts where spatial integration can be shown to have a strong causal influence on social integration, forced physical integration may be more effective. While this paper stops well short of making any such specific claims, it provides an empirical framework that we believe can be used to investigate these and other matters of real importance. To our knowledge, this is the first study to demonstrate how new forms of data can be used to separately analyze social and spatial segregation in a large population, and to relate these concepts to the broader development discourse.

1.1 Related Work

Our analysis relates most closely to two strands in the literature focused on the empirical analysis of spatial and social segregation, respectively. As noted above, spatial segregation has been the dominant approach to measuring segregation, largely due to the relative ease of measurement [33, 16]. However, as digital technologies have enabled more precise measurements of peoples' locations, so too has the empirical study of segregation become more sophisticated. For instance, in a series of early studies of community structure in call data, Blondel et al. [5, 24] find that Belgian mobile phone users communicate more frequently with individuals who speak the same language. In the work perhaps most closely related to our analysis, Toomet et al. [30] use landline communications data from Estonia to analyze patterns of geographic segregation at different times of the day. They find that large and urban areas in Estonia tend to be more segregated than smaller rural ones, and that places visited during leisure hours (i.e., after work but before sleep) tended to be the most integrated. While [30] thus represents an important contribution, as with much of the existing demographic research the focus is exclusively on spatial segregation.

More broadly, our work fits into a growing literature that relies on large-scale digital traces to understand social phenomena [25], and in particular to a set of recent work has analyzed political polarization and ideological segregation using data from internet blogs and Twitter [1, 13]. Our case study relies on the analysis of anonymized mobile calling records, but in principle could be extended to other spatial network datasets (such as those from popular social media sources) that contain geographic and linguistic markers in the dataset. This work is thus similar in vein to recent work that relies on digital traces to measure patterns of migration [34], the spread of epidemics [31], and infer socioeconomic status [7, 6, 20].To our knowledge, however, this paper represents the first effort to use such data to analyze ethnic segregation in a developing country.

A final, related strand of literature is concerned with developing generative models of network formation that can account for observed levels of segregation and integration. For instance, building on classic models of network attachment, Jackson et al. [23] and Cheng and Xie [10] show that a simple model that combines same-race preference with multidimensional preferences would imply that larger "contexts" (in their case, schools) lead to increased segregation. Similarly, in one of the only studies of which we are aware that explicitly accounts for the relationship between social and spatial segregation, Currarini et al. [14] develop a formal model of friendship formation where individuals have types, and then exhibit a preference for people of similar types and for people who are physically close to them. Using data gathered from a single U.S. high school, these authors find that observed patterns of segregation and integration are consistent with the predictions of their model.

2. FRAMEWORK AND METHODS

2.1 Measuring Segregation and Homophily with Spatial Network Data

For our empirical analysis, our primary concern is measuring levels of segregation within a population. In related literature, a few common empirical standards have been used to measure segregation, integration, and homophily[32, 33]. However, as noted by [29], arbitrary geographic segmentation, as is often imposed by administrative units, is a poor way to measure the geographic area surrounding a person within it. In particular, such segmentation does not permit a more fine-grained analysis of local (i.e., geographically disaggregated) or dynamic (i.e. time-variant) aspects of segregation.

The approach we propose enables a quantitative analysis of segregation that is both dynamic and disaggregated, and which, in principle, can be applied to any spatial network dataset. The required features of such a dataset are the following:

1. *Ethnic identifiers* that can be used to infer the ethnicity or other group-level characteristic of individuals
2. *Spatial markers* that allow individuals to be assigned to physical locations. Such markers can be in the form of GPS coordinates, mobile phone BTS locations, census tracts, etc.
3. *Network connections* that link individuals within the network. Example markers might be call or email traffic, friendship connections, retweets, etc.
4. *Temporal markers (optional)*, which are only necessary if one wishes to analyze the time-varying dynamics of segregation

In Section 3 we use data from a mobile phone network in a South Asian city as a case study, but other examples of spatial network data are becoming increasingly common as digital devices proliferate in developing countries. Examples of spatial network data that would be suitable for this framework include call detail records, twitter or Facebook network data, data collected from specific ODK deployments, etc.

2.2 Spatial and Social Segregation

Spatial network data allows for the measurement of two different, but related, measures of segregation: *social segregation*, the extent to which a person is more likely to interact with someone of the same ethnicity; and *spatial segregation*, the amount a person is more likely to be physically located near someone of the same ethnicity. We will further discuss the compound measure of *disproportionate segregation*, a concept related to inbreeding homophily, which reflects the extent to which individuals are more socially segregated than would be expected at a given level of spatial segregation.

Formally, we index individuals by the subscript i and locations by r. We assume a total population of N individuals of K different types (ethnicities) that exist in a physical space containing R different discretized locations. Each subpopulation (ethnicity) k contains N_k individuals, and in a given window of time each region contains N_r individuals such that $N = \sum_{k \in K} N_k = \sum_{r \in R} N_r = \sum_{r \in R} \sum_{k \in K} N_{rk}$.

2.2.1 *Spatial Segregation*

We define the *spatial segregation* of a given region as simply the fraction of the population in r that is of type t:

$$w_{tr} = \frac{N_{tr}}{N_r} \tag{1}$$

Values of w_{tr} close to 0.5 are indicative of low segregation, while values near the minimum of zero or maximum of one indicate high levels of segregation. Such a definition additionally allows for measurement of heterogeneous spatial segregation along a temporal dimension. For instance, following the example of Toomet et al. [30], it would be straightforward to segment w_{rt} by t into time of day/week buckets to separately quantify patterns of segregation on work days, weekends, and in the evening.

2.2.2 *Social Segregation*

Using observed interactions in a given time window, we define s_t as the number of contacts that a person of type t has of the same type and d_t as the number of contacts that people of type t have that are of a different type. We can now define *social segregation* as H_{tr}, sometimes referred to as the homophily index, i.e., the fraction of contacts the people of type t (in region r) form with the same type (in any region).

$$H_{tr} = \frac{s_t}{s_t + d_t} \tag{2}$$

2.2.3 *Disproportionate segregation:*

We say that a population is disproportionately segregated if $H_{tr} > w_{tr}$. This indicates that people of type t in region r have more contacts of their own type than would be expected by random pairing - i.e., if they made friends just based on the relative percentage of types in the population. To formalize this concept, we define *disproportionate segregation* DS_{tr} as

$$DS_{tr} = \frac{H_{tr} - w_{tr}}{1 - w_{tr}} \tag{3}$$

This index of disproportionate segregation DS, sometimes referred to as *inbreeding homophily* [11], measures the unexpected social segregation observed for a given level of spatial segregation. The index of disproportionate segregation will be positive when individuals interact most with individuals of the same type (despite a high presence of individuals of other types), 0 when individuals interact by type at the rate expected based on the relative fraction of their type in the population, and negative if individuals interact more with individuals of other types (despite a high presence of individuals of the same type).

2.3 A stylized example

To help make the above metrics more intuitive, we illustrate how each statistic is computed using data from two observed networks in South Asia. Figure 1 shows the schematic networks of two small, geographically-contained networks. Each node with an outbound arrow represents a unique mobile phone user that made a call from this area. Grey nodes represent users in ethnic group A and black nodes represent those in ethnic group B, as determined through their language preference. A directed edge from node i to node j shows that one or more calls were placed from user i to user j. Thus, users with outbound arrows can be said to be occupying the same geographic area (users with inbound arrow may be located elsewhere).

Table 1: Example computations of segregation indices for networks in Figure 1

	Figure 1 (left)	Figure 1 (right)
N_{Ar}	11	11
N_{Br}	3	18
w_{Ar}	$\frac{11}{11+3} = .786$	$\frac{11}{11+18} = .379$
w_{Br}	$\frac{3}{3+11} = .214$	$\frac{18}{18+11} = .621$
s_A	10	11
d_A	1	5
s_B	1	8
d_B	2	11
H_{Ar}	$\frac{10}{10+1} = .909$	$\frac{11}{11+5} = .688$
H_{Br}	$\frac{1}{1+2} = .333$	$\frac{8}{8+11} = .422$
DS_{Ar}	$\frac{.909-.786}{1-.786} = .575$	$\frac{.688-.379}{1-.379} = .192$
DS_{Br}	$\frac{.333-.214}{1-.214} = .094$	$\frac{.422-.621}{1-.621} = -.075$

Table 1 shows all the calculated metrics for the two small networks in Figure 1. For instance, in the network on the left, w_{Ar} is simply the number of grey nodes with outbound edges divided by the total number of nodes with outbound edges, with $w_{Ar} = \frac{11}{11+3}$ and $w_{Br} = \frac{3}{3+11}$. Using equation 2, we compute social segregation by looking at the network of contacts with whom nodes in the region interact. s_A is defined as the number of edges between two grey nodes, and d_A is computed as the number of edges originating at a grey node and directed toward a black node. In this example, $s_A = 10$, $d_A = 1$, and thus social segregation for ethnicity A is $H_{Ar} = \frac{10}{10+1}$. Disproportionate segregation can then be trivially computed from these already-computed metrics.

2.4 Segregation and Network Formation

The ability to separately measure these different types of segregation is not solely a descriptive exercise; it can also inform our theoretical understanding of the processes which determine the formation and structure of existing social networks. This is perhaps most evident by contrasting three canonical, stylized models of network formation.

Random attachment: Under a naive model of completely random attachment, where neither physical distance nor ethnic homophily plays a role in whether two people choose to interact, interactions would simply be proportional to the overall population share and we would expect levels of social segregation to be bimodal at $f = N_A/N_B$ for people of ethnicity A and $1/f$ for people of ethnicity B.

Geographic constraints: Similarly, if individuals are neither homophilic (preferring people of the same ethnicity) nor heterophilic (preferring people of a different ethnicity), but only interact with people in their immediate geographic surroundings, then we expect the indices of social segregation to be very closely correlated to the indices of spatial segregation computed above. This "geographically constrained" model is, in a sense, desirable, as it is indicative of a context where people are equally willing to interact with people of different types.

Social preferences: By contrast, if individuals instead do prefer to interact with people of their own ethnicity, as formalized by [14], then we expect levels of social segrega-

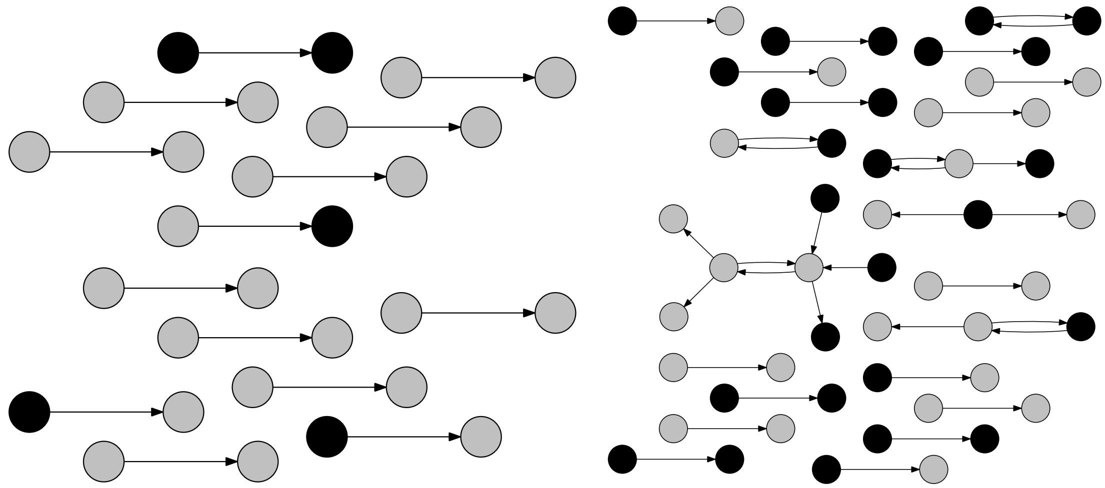

Figure 1: Social graphs from two representative networks. Each node represents a unique mobile phone subscriber, with grey nodes indicating people of ethnicity A and black nodes indicating people of ethnicity B. Arrows (directed edges) between node i and node j indicate that at least one call was made from i to j.

tion in excess of spatial segregation. In the few cases where researchers have attempted to fit data to such models, this "social preference" model is typically a better approximation of the true population dynamics.

In the following section, we use data to compute the different types of segregation described above, and to illustrate how such statistics can generate broader implications for understanding the underlying processes of network formation.

3. CASE STUDY: SEGREGATION IN A SOUTH ASIAN CITY

The previous section provides a formal framework for measuring different types of segregation when it is possible to observe three features: individual type or ethnicity; individual location; and connections between individuals. While such data can be generated from sources such as Twitter [22], Facebook [18], and email providers [34], in the context of developing countries, the most utilized technology is the mobile phone. As a result, the past few years have seen a fast rise in the amount of ICTD-related research based on the analysis of data generated through mobile phone use [20].

The case study we use to illustrate the above methods employs data similar in vein to that used by these researchers, but taken from a large mobile phone network in South Asia. In addition to carefully adhering to ethical and IRB guidelines for anonymizing and ensuring the confidentiality of the data used for this case study, we are contractually prohibited from disclosing, among other things, the identity of the mobile phone operator and the name of the country in which they operate. As a result, the contextual details we can provide are limited. However, our goal in presenting this case study is not to make inferences about the specific population being analyzed, but rather to illustrate the manner in which the framework can be implemented, and the types of insights that might result from its application.

3.1 Context and Data

For the purposes of the analysis that follows, we focus on a population of individuals in a developing country in South Asia. In the region we study, approximately 80% of individuals own mobile phones, and we utilize data obtained from the nation's largest mobile phone network. While the complete dataset contains the communication records of millions of individuals, here we focus on on a subset of the population in one of the largest urban centers in the country, and the subset of those subscribers whose language preference is known. In total, we focus on a population of 65,225 individuals, who are observed to make 695,706 calls from several thousand unique mobile phone towers.

These data allow us to observe the three necessary features as follows (note that a full discussion of the limitations of these data is deferred to section 5.1):

Ethnic Identifiers: When individuals register their phone or device SIM card on the operator's network, they have to select one of two languages to be their language of preference. This setting is used by the operator to communicate with the subscriber, for instance for marketing communications and for technical support. In this country, there is a strong correlation between language spoken and underlying ethnicity, so we have reason to believe that the selection an individual makes on his SIM card is indicative of the ethnicity with which he identifies. As qualitative evidence, we conducted interviews with 12 different individuals and found a correlation ($\rho = 0.845$) between an individual's "primary ethnicity" and the language selected on the SIM card. We did not, however, rigorously test this correlation, and thus we want to be careful not to over-emphasize the validity of the analytic results from the case study. As noted above, the primary intent of case study is to illustrate the operationalization of the framework developed in Section 2, rather than to draw specific conclusions about the the particular South Asian country for which we have data.

Spatial markers: To measure the approximate locations of individuals in space, we use geographic markers contained in the call records. Since every communication event on the network must pass through one of the networks mobile phone towers, it is possible to use the communication logs to roughly assign individuals to places at the moment

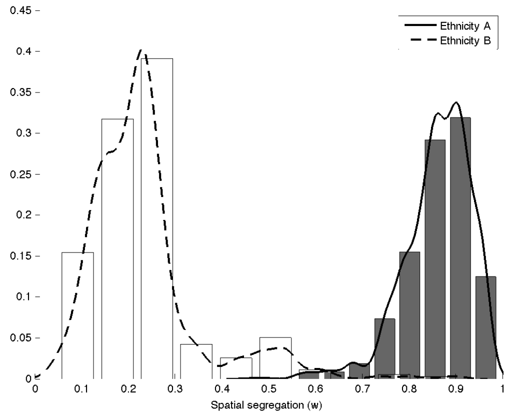

Figure 2: Distribution of spatial segregation w in a major urban area in a developing country in South Asia. Values close to 0.5 imply integration, values close to zero and one imply segregation.

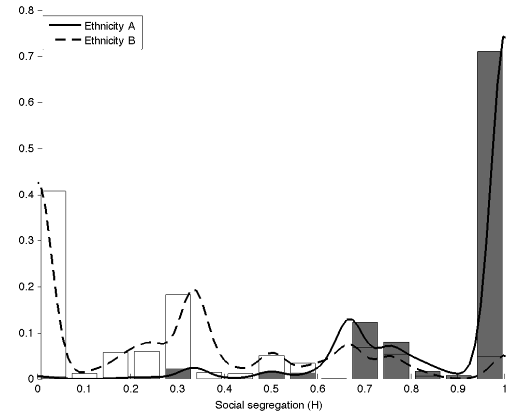

Figure 3: Distribution of social segregation H in the same urban area. Large masses at both ends of the distribution indicate a general tendency for both ethnicities to interact with ethnicity A.

when they interact with one another. In this city, there are approximately one thousand different mobile phone towers, which are distributed in space as depicted in Figure 6. In this figure, plus (+) signs represent the location of mobile phone towers, and the white polygons indicate the approximate coverage area of each tower.

Network connections: Through the mobile operator dataset, we are able to observe all phone call interactions between individuals on the network. We interpret this as one of multiple forms of social interaction, providing a noisy signal of the structure of the underlying social network. Most importantly, we use this interaction data to identify the set of contacts of each individual i, which we use to calculate s_i and d_i, the share of i's contacts of same and different ethnicity, respectively.

3.2 Implementation details

The original dataset obtained from the operator consists of two files: one which contains records of each transaction, and another which contains metadata on each subscriber, including the language preference on the subscriber's SIM card. The transactional records contain observations of the form $< i, j, r, t >$ where i and j are the masked identities of the calling and recipient parties, r is the tower nearest to i, and t is a timestamp for the call. To compute the metrics described in section 2 we employ the following procedures:

- To assign individuals to locations, we computed a base location for each of our callers as the tower with which they were most frequently associated. We found our results to be robust to other methods for assigning location, such as computing each individual's Euclidean centroid (the "center of mass" approach).
- For spatial segregation, we calculated homophily at a tower level. Any person who has ever called from a tower was assigned to it. This allowed us to define a community of callers at each of the towers. To calculate w_i, we looked at the fraction of each type that had ever called from that tower or towers within a 10 km radius.
- To compute social segregation, we totaled the number of same and different contacts, separately, that people who called from that tower had. These numbers were averaged to calculate the overall homophily of a tower. Towers with fewer than 10 people assigned to them were filtered out to reduce noise caused by towers with very low activity.
- Since not all individuals have matching metadata, we first pruned the data to exclude communication events where we did not have language information about both the caller and the receiver. We also removed the calls that originated and were directed to the same person - occasionally a person would initiate a call to himself. These were removed as we did not want to include oneself in one's set of friends and distort the amount of contacts of the same type. People with fewer than 3 contacts were excluded.
- To compute segregation, we defined a 10 km radius around the caller. All people assigned a tower within his 10 km radius were included in the spatial segregation calculation. And all people he called who were also based at a tower within that radius were included in the social segregation calculation.
- All data received from the operator is immediately anonymized to remove identifying information. However, as has been noted recently by researchers from several fields, simple anonymization cannot guarantee the privacy of individuals [15, 27]. Thus, all data was handled under a strict Internal Review Board protocol that, among other protections, ensured that only the authors of this work had access to even the anonymized dataset, and that no data was ever copied or removed from a single firewalled server.

4. CASE STUDY RESULTS

4.1 Spatial analysis

Figure 2 provides a high-level perspective on the distribution of spatial segregation across the entire population of phone owners in a single urban area. The bimodal distribution of w in Figure 2 indicates that people in ethnic group A

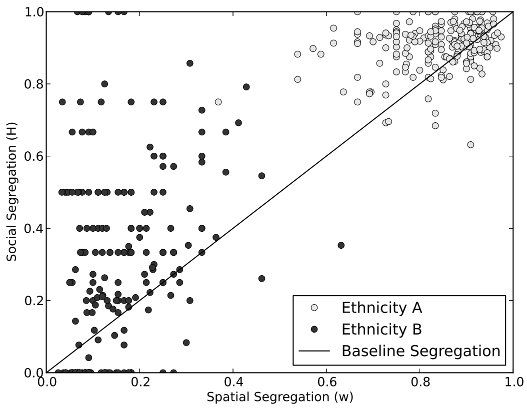

Figure 4: Relationship between spatial segregation (w) and social segregation (H). Each geographic region, defined as the land area covered by a single mobile phone tower, is represented as two points, one grey to reflect the relationship between w and H for ethnic group A, and one black to reflect the relationship for ethnic group B.

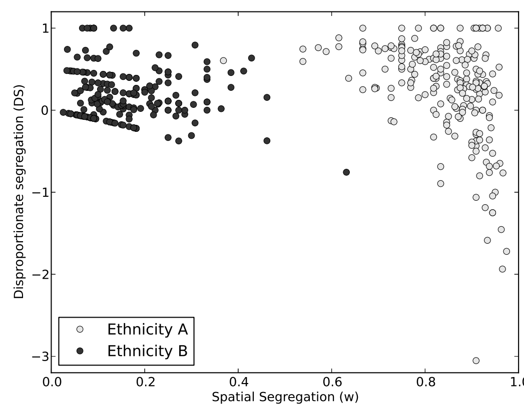

Figure 5: Disproportionate segregation (DS) as a function of spatial segregation w. The mass of observations with values of DS greater than zero are indicative of a general homophilic preference by both ethnicities that is disproportionate to each ethnicity's local population share.

tend to live in places where there is a high fraction of their own ethnicity, whereas people in ethnic group B tend to live in places where there are relatively low proportions of their own ethnicity. While this is presumably driven largely by the fact that ethnicity A represents approximately 85 percent of the total urban population, it also indicates that people of ethnicity B are not clustering themselves in homogeneous pockets of their own ethnic group.

4.2 Social analysis

The distribution of social segregation in the same urban area is shown in Figure 3. As with spatial segregation, levels of social segregation are very high for ethnic group A, and generally low for ethnic group B.

Reflecting on the models of network formation discussed in Section 2.4, for this urban area we can clearly reject the naive model of random attachment, as neither ethnic distribution is tightly unimodal at the points f or $1/f$. Instead, the distribution of social segregation is similar to the distribution of spatial segregation, which suggests the second model of geographic constraints may better fit the data, where individuals form friendships without any preference for their own ethnicity, and simply associate with those nearby.

However, the distributions in Figures 2 and 3 are visibly not identical, and this intuition can be confirmed with a χ^2 test ($p < 0.01$). Indeed, while the distribution of spatial segregation is strictly bimodal, the distribution of social segregation contains large masses at zero for ethnic group B and one for ethnic group A. Intuitively, these masses indicate a disproportionate likelihood for *both* ethnicities to interact with people of ethnicity A.

This is precisely where the power of the large-scale behavioral data and this framework can provide insight that would be unavailable to researchers using traditional datasets. Since the measurement of social segregation requires information on person-to-person interactions, which has historically been very difficult to observe (and nearly impossible to measure at the population scale), scholars have instead typically relied on measures of spatial segregation, which only requires information on where people are located. For this reason, the vast majority of research on segregation uses these spatial measures, which can be derived from census data and other common survey instruments. Implicitly, much of this work assumes that spatial and social segregation can be treated as equivalent, but as noted above this implies a very naive model of interaction. Instead of assuming this equivalency, our framework allows for the hypothesis to be tested empirically.

Figure 4 provides a more direct perspective on the relationship between spatial segregation (on the x-axis) and social segregation (on the y-axis). There is a clear monotonicity in the relationship, and the correlation is statistically significant (p<0.05; $R = 0.24$ for Ethnicity A and $R = 0.30$ for Ethnicity B). This relationship indicates that individuals are more likely to interact with people of their own ethnic group when they are physically surrounded by people of their own ethnic group.

However, there are significant deviations from this average behavior, with some places exhibiting low levels of social segregation despite high levels of spatial segregation, and some behaving the opposite. In particular, the 45-degree line on Figure 4 represents "baseline segregation", with all observations above the line indicating locations of heightened social segregation (relative to the baseline of what would exist under the geographic constraints model of network formation), and all observations below the line indicating locations of better social integration. The disproportionate quantity of points above the line is indicative of disproportionate segregation, a hypothesis we test formally in the following section.

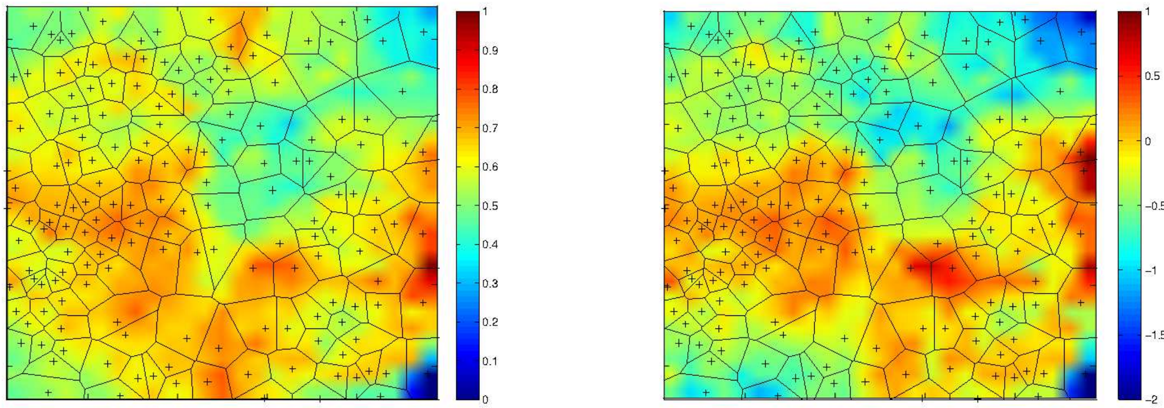

Figure 6: Spatial Segregation w (left) and Social Segregation H (right) in an urban center. Plus (+) signs represent mobile phone towers, with the Voronoi cells indicating approximate coverage areas of each cell. The intensity of segregation is shown in the underlaid heatmap, with highly segregated areas appearing as red and highly integrated areas appearing as blue.

4.3 Joint spatio-social analysis: Disproportionate segregation

The index of disproportionate segregation described in Section 2.2.3 enables us to further untangle the role that spatial segregation plays in determining social segregation. In particular, we can look for evidence of deviations from the naive model of random attachment described above, where people socialize with those nearby, irrespective of ethnicity. In smaller contexts where both social and spatial segregation was elicited through surveys, both [10] and [14] found that a better model for friendship formation in schools involved one where students preferentially associated with others of the same type.

In the data, such deviations from a naive model are indicated by the level of association with co-ethnics above what would be expected under a baseline model of random association. Figure 5 shows the level of this disproportionate segregation as a function of spatial segregation for the South Asian nation. Values of DS greater than zero indicate spatial communities of people who strongly prefer to associate with co-ethnics over non-coethnics. The fact that the vast majority of observations are above the y-axis in Figure 5 thus indicates a disproportionate association of coethnics with coethnics, above what would be expected under a simple model of geographic constraints. Further, the mass of ethnicity B in the top left quadrant exhibit a strong coethnic bias, despite the relative rarity of coethnics in the local population. By contrast, the ethnicity A individuals on the bottom right show a preference for ethnicity B that is surprising given that they are surrounded by people of their own ethnicity (A).

Finally, Figure 6 shows the geographic distribution of both spatial and social segregation in the center of a major urban center. Segregation ranges from zero to one, with extremely segregated areas showing up as red and integrated areas blue. The pockets of high segregation are apparent from the maps, with both ethnicities tending to cluster in areas of extreme segregation. The correlation between the two types of segregation is also visually apparent, though deviations clearly exist with certain areas being prone to high spatial but low social segregation, and vice versa.

5. DISCUSSION

In the previous section, we tested the analytic framework developed in Section 2 using sample data from a large urban area. Even in the rather superficial analysis, we find that different sub-regions of the city are highly geographically segregated, with many pockets of individuals dominated by a single majority ethnicity. We have similarly seen that these patterns of spatial segregation are also reflected in observed dynamics of social segregation, in that individuals are significantly more likely to interact with coethnics than individuals of the opposite ethnicity. Moreover, the observed pattern of social segregation is disproportionate in the sense that individuals are *even more likely* to interact with coethnics than would be expected just given the uneven distribution of ethnicities over space.

Taking these findings at face value (we will discuss reasons to doubt these findings soon), there are several immediate implications of relevance to to the development community. First, the dynamics of disproportionate segregation, wherein the minority ethnicity is observed to cohabitate with the majority ethnicity but the majority ethnicity appears more likely to socially interact with the minority, is consistent with qualitative accounts of interaction in this country. In this society the minority speakers tend to be relatively less educated and of lower socioeconomic status, and a norm exists whereby it is perceived as more appropriate for individuals in ethnicity A to initiate contact with those in ethnicity B, rather than the other way around. While this result is more tantalizing than it is robust, it is suggestive of how this quantitative framework can be used to corroborate qualitative observations, and to measure at a population scale social phenomena documented with more careful ethnographic research.

More generally, the comparison of spatial and social segregation can be used to inform policies in developing countries that are designed to promote integration. While much of our analysis is in the cross-section, i.e. we do not analyze temporal trends in the different types of segregation we observe, in principle the framework is easily adapted to a dynamic context by appropriately discretizing t in the models from Section 2. We see this to be a promising area for future work, as the ability to measure changes over time separately for spatial and social segregation may enable a deeper investigation of the causal relationship of the two forces. In many developing countries, where ethnic conflict can create significant hardship, a deeper understanding of the causes and nature of ethnic segregation can point to potential solutions.

For instance, a focus of much attention in the development community is on building social capital through the removal of frictions that arise from ethnic tensions in the labor force and education system [17, 12]. However, whether reducing spatial segregation - for instance by encouraging interaction in public spaces, requiring pluralities in political representation, or in offering incentives for individuals to colocate in diverse areas - has the desired effect on social interaction and integration is not known. It is, however, an empirical question that can be investigated using the methods outlined in this paper. The reverse hypothesis, that reducing social segregation – for instance by designing technologies to foster interaction across ethnic lines – is equally plausible, and equally amenable to empirical analysis. Of course, this paper stops well short of offering any causal insight, and such conclusions would require a careful research design where changes to one dimension of segregation can be isolated and the effects on another dimension identified.

Beyond the study of ethnic segregation, similar methods could be used to study other types of social mixing, integration, and homophily. For instance, an analogous framework could be used to determine the extent to which different genders are seen to associate in space and time (in some societies it may be expected to see a sharp disconnect between spatial and social segregation by gender), or people of different ages or social status.

5.1 Limitations of the approach

While this analytic framework offers several advantages over traditional methods used to study processes of segregation, it has several significant limitations that merit further discussion.

First, as with many other practices of relying on digital traces as a proxy for underlying social dynamics [25, 34, 31, 20], the external validity of our approach is limited by the fact that ownership and penetration of the digital sensor is often not complete. For our case study, and for many similar applications, adoption of the underlying technology in not uniform, and thus the digital traces used to make social inferences are likely to be biased toward the population of early adopters and heavy users. Such biases are exacerbated in the developing-country context, as several studies have documented differential rates of uptake by different strata of society [21, 7, 8]. Therefore, claims about segregation and integration based on these methods are fundamentally limited to observed behavior among technology users. In limited cases it may be possible to re-weight observations to account for differential rates of uptake [26], but in practice these corrections typically cannot account for all sources of bias.

One type of bias that is particularly difficult to correct occurs when patterns of behavior observed through the trace are not representative of social behavior about which inferences are being made. In certain situations this may lead to random error that increases noise but not bias. For instance, if a satellite were to detect a person's location at random times throughout the day and such observations were used to infer location, the trace would represent a random sampling of activity and should not result in biased estimates of behavior. However, when the observed activity is nonrandom (for instance, if the person's location is only observed when she is at rock concerts and decides to use Twitter), then estimates of activity will only be valid on the biased sample of activities observed (i.e., for inferences about behavior at rock concerts).

In the case study based on mobile phone data, such biases are likely introduced by the fact that we only observe call interactions (the operator did not provide us with SMS records, which are likely to reflect different dynamics of social interaction), and that not all people use their phones for the same purpose. To some extent, we hope to smooth these differences by aggregating data over several months to infer the structure of the social network, and thereby focus only on whether two people have ever communicated instead of the intensity of the communication. Nonetheless, we do not mean to suggest that the data used in the case study provides a perfect reflection of the underlying social network; only that it can provide a high-resolution perspective on segregation within population that is difficult to survey with other methods.

A further limitation of our approach is the manner of determining the ethnicity k of each individual. As discussed earlier, this inference relies on the language selected by the device owner when registering the device for use on the network. Since there is a high correlation between language and ethnicity, we impute ethnicity based on language preference. This approximation of ethnicity with language is common practice in the large literature on ethno-linguistic fractionalization [4, 3]. There are, however, examples when institutional factors (such as legal regulation on language choice) or technological factors (such as default settings on the device) may cause the correlation between ethnicity and language to be far from perfect. For the purposes of the case study in this paper, we found a relatively high correlation ($\rho = 0.845$) between ethnicity and language interface on the phone, which we deemed to be sufficient to provide a strong signal, albeit with some noise, of ethnicity based on language selection.

We thus believe that the locations we observe for each person are relatively representative of that person's location, and we think that the manner by which we infer ethnicity is defensible. Instead, our primary concern is with the fact that we have reason to suspect that the sample of interactions we observe is not entirely representative of underlying social ties between individuals. For while we do observe every phone-based interaction between and across ethnicities, a more acute form of bias can arise when the *nature of use* differs by ethnicity (even given uniform adoption across ethnic groups). Such examples are not uncommon in developing countries, where social norms may dictate technology-mediated interactions that are unrepresentative of traditional modes of interaction. An example we have noted in the country that is the focus of this paper is that speakers of language B sometimes feel it is inappropriate to call their employers (who are disproportionately likely to speak language A), while the reverse is deemed to be social acceptable. While this particular example can be dealt with empirically by defining relationships over a long period of time through undirected rather than directed edges (as we have done above), similar dynamics may lead to more insidious forms of bias that are not so easily detected. Thus, in no way do we mean to discount the importance of this final limitation. Particularly in cases where data of this nature is used to inform policy, rather than as a diagnostic tool as is the case in the current paper, great care must be taken in appropriately scoping the nature of the questions asked and conclusions drawn.

6. CONCLUSION

The dynamics of ethnic segregation and integration play a critical role in determining the social and economic development of many impoverished nations [17, 12]. However, empirical studies of ethnic segregation have been hindered by a lack of reliable data [28, 4]. In particular, social scientists have been limited in their ability to disaggregate patterns of segregation to smaller geographic regions and smaller time intervals, making it difficult to understand the causes and effects of different types of segregation.

Here, we develop a generic computational framework for using spatial network data to develop fine-grained measures of two different types of ethnic segregation: spatial segregation and social segregation. These two types of segregation are often conflated in empirical work on segregation, even though they represent two distinct phenomena with different implications for public policy. Using data from a case study, we provide an example of how these measures might diverge: although the two measures of segregation are highly correlated in the South Asian city, there are statistically significant differences that can be used to distinguish between different stylized models of social preference and network formation. Considerable calibration, contextualization, and cross-validation is needed before this empirical evidence can be useful in a policy context. However, these early results indicate a preference for co-ethnics that is disproportionate to the overall proportion of coethnics in the population.

While the results from our case study are preliminary, our hope is that the analytic framework can be adopted by other researchers studying ethnic and other forms of segregation in developing countries. Through this work, we seek to build a deeper understanding of the key drivers of segregation, and thereby inform policy to foster integration across ethnic lines.

7. REFERENCES

[1] L. Adamic and N. Glance. The political blogosphere and the 2004 U.S. election: Divided they blog. In *In LinkKDD '05: Proceedings of the 3rd international workshop on Link discovery*, pages 36–43, 2005.

[2] A. Alesina, R. Baqir, and W. Easterly. Public goods and ethnic divisions. *The Quarterly Journal of Economics*, 114(4):1243–1284, 1999.

[3] A. Alesina, A. Devleeschauwer, W. Easterly, S. Kurlat, and R. Wacziarg. Fractionalization. *Journal of Economic growth*, 8(2):155–194, 2003.

[4] C. Blattman and E. Miguel. Civil war. *Journal of Economic Literature*, 48(1):3–57, Mar. 2010.

[5] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, Oct. 2008.

[6] J. E. Blumenstock. Inferring patterns of internal migration from mobile phone call records: Evidence from rwanda. *Information Technology for Development*, 18(2):107–125, 2012.

[7] J. E. Blumenstock and N. Eagle. Mobile divides: Gender, socioeconomic status, and mobile phone use in rwanda. *4th International IEEE/ACM Conference on Information and Communications Technologies and Development*, Dec. 2010.

[8] J. E. Blumenstock and N. Eagle. Divided we call: Disparities in access and use of mobile phones in rwanda. *Information Technology and International Development*, 8(2):1–16, 2012.

[9] S. Castles. How nationâĂŘstates respond to immigration and ethnic diversity. *Journal of Ethnic and Migration Studies*, 21(3):293–308, 1995.

[10] S. Cheng and Y. Xie. Structural effect of size on interracial friendship. *Proceedings of the National Academy of Sciences*, Apr. 2013. PMID: 23589848 PMID: 23589848.

[11] J. Coleman. Relational analysis: The study of social organizations with survey methods. *Human Organization*, 17(4):28–36, Dec. 1958.

[12] P. Collier. The political economy of ethnicity. In *Annual World Bank Conference on Development Economics*, pages 387–405, 1998.

[13] M. D. Conover, J. Ratkiewicz, M. Francisco, B. Goncalves, A. Flammini, and F. Menczer. Political polarization on twitter. In *Proc. 5th Intl. Conference on Weblogs and Social Media*, 2011.

[14] S. Currarini, M. O. Jackson, and P. Pin. An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4):1003–1045, 2009.

[15] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3, Mar. 2013.

[16] O. D. Duncan and B. Duncan. A methodological analysis of segregation indexes. *American sociological review*, 20(2):210–217, 1955.

[17] W. Easterly and R. Levine. Africa's growth tragedy: policies and ethnic divisions. *The Quarterly Journal of Economics*, 112(4):1203–1250, 1997.

[18] N. B. Ellison, C. Steinfield, and C. Lampe. The benefits of facebook "friends:" social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.

[19] J. D. Fearon and D. D. Laitin. Ethnicity, insurgency, and civil war. *American political science review*, 97(1):75–90, 2003.

[20] V. Frias-Martinez and J. Virseda. On the relationship between socio-economic factors and cell phone usage. In *Proceedings of the Fifth International Conference on Information and Communication Technologies and Development*, ICTD '12, pages 76–84, New York, NY, USA, 2012. ACM.

[21] A. Gillwald. *Towards an African e-Index: Household and individual ICT access and usage across 10 African countries*. LINK Centre, Wits University, School of Public and Development Management, 2005.

[22] B. Huberman, D. M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1-5), 2009.

[23] M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Nov. 2010.

[24] R. Lambiotte, V. D. Blondel, C. de Kerchove, E. Huens, C. Prieur, Z. Smoreda, and P. Van Dooren. Geographical dispersal of mobile communication networks. *Physica A: Statistical Mechanics and its Applications*, 387(21):5317–5325, Sept. 2008.

[25] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, and M. Van Alstyne. Life in the network: the coming age of computational social science. *Science (New York, N.Y.)*, 323(5915):721–723, Feb. 2009.

[26] C. A. Moser and G. Kalton. Survey methods in social investigation. *Survey methods in social investigation.*, (2nd Edition), 1971.

[27] P. Ohm. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57:1701, 2010.

[28] D. N. Posner. Measuring ethnic fractionalization in africa. *American journal of political science*, 48(4):849–863, 2004.

[29] S. F. Reardon and D. O'Sullivan. Measures of spatial segregation. *Sociological methodology*, 34(1):121–162, 2004.

[30] O. Toomet, S. Silm, E. Saluveer, T. Tammaru, and R. Ahas. Ethnic segregation in residence, work, and free-time: Evidence from mobile communication. *University of Tartu*, 2012.

[31] A. Wesolowski, N. Eagle, A. J. Tatem, D. L. Smith, A. M. Noor, R. W. Snow, and C. O. Buckee. Quantifying the impact of human mobility on malaria. *Science*, 338(6104):267–270, Oct. 2012.

[32] M. J. White. The measurement of spatial segregation. *American journal of sociology*, pages 1008–1018, 1983.

[33] M. J. White. Segregation and diversity measures in population distribution. *Population index*, 52(2):198–221, 1986.

[34] E. Zagheni and I. Weber. You are where you e-mail: using e-mail data to estimate international migration rates. In *Proceedings of the 3rd Annual ACM Web Science Conference*, pages 348–351, 2012.