

Campaign Optimization through Behavioral Modeling and Mobile Network Analysis

Yaniv Altshuler^{1,*}, Erez Shmueli^{2,*}, Guy Zyskind³, Oren Lederman³, Nuria Oliver⁴, Alex (Sandy) Pentland³

Abstract—Optimizing the use of available resources is one of the key challenges in activities that consist of interactions with a large number of “target individuals”, with the ultimate goal of “winning” as many of them as possible, such as in marketing, service provision, political campaigns, or homeland security. Typically, the cost of interactions is monotonically increasing such that a method for maximizing the performance of these campaigns is required. In this paper we propose a mathematical model to compute an optimized campaign by automatically determining the number of interacting units and their type, and how they should be allocated to different geographical regions in order to maximize the campaign’s performance. We validate our proposed model using real world mobility data.

Keywords—Mobile Networks, Network Optimization, Marketing, Behavior Modeling

1 INTRODUCTION

IN a world of limited resources, behavior change campaigns (e.g. marketing, service provision, political or homeland security) can rely on creativity and “coolness” up to a certain point. The success of a campaign can generally be defined as the product of *reach* (portion of the population exposed to the campaign messages) and *value* of a single interaction (the capacity of a message to induce a certain behavior in an exposed audience). Hence, campaign managers typically distribute their budget between content enhancement (to increase the value a single interaction) and wide reach. Yet, to date it seems that the optim trade-off between these two factors is found as a result of “intuition” rather than based on well established analysis.

In this paper, we propose a novel mathematical method that, given the characteristics of the target audience and its ability to be persuaded, generates an optimized campaign strategy in terms of: (a) the quantity of interacting units, also referred to as *insertions* and (b) the monetary allocation to each unit. The model takes into account the population’s mobility in an urban environment as it can be inferred from real data received from a large mobile phone carrier. Even though

different populations located in different environments would be tailored with different campaign strategies, the optimality of each strategy would be maintained.

A major contribution in our optimization model is the use of network analysis methods to approximate the reach of a campaign. More specifically, given the network of mobility between the different geographic locations, and a subset of locations, we use the *Group Betweenness Centrality* (GBC) [54] – a network measure that calculates the percentage of shortest paths among all pairs of network nodes that pass through a pre-defined sub-set of the network’s nodes – to approximate the reach of this subset of locations. We then demonstrate that this function can be approximated using a smooth and easily analyzed *Gompertz* function. This tackles the main limitation of works on campaign optimization hitherto – efficiently estimating the campaign reach as a function of the number of units and their locations.

Finally, we validate our campaign optimization model using a real-world mobility network inferred from CDR data, and demonstrate how GBC based deployment of campaign units outperforms several common alternatives.

The rest of this paper is organized as follows. Related work is discussed in Section 2. A characterization of a campaign model and its target optimization function are presented in Section 3. An analytical optimization of the campaign’s model is shown in Section 4. A validation of the model using real world mobile data is described in Section 5. Concluding remarks appear in Section 6.

2 RELATED WORK

In recent years the social sciences have been undergoing a digital revolution, heralded by the emerging field of “computational social science”. Lazer, Pentland et. al [75] describe the potential of computational social science to increase our knowledge of individuals, groups, and societies, with an unprecedented breadth, depth, and scale. Computational social science combines the leading techniques from network science [19], [88], [116] with new machine learning and pattern recognition tools specialized for the understanding of people’s behavior and social interactions [50].

Marketing campaigns are essential facility in many areas of our lives, and specifically in the virtual medium. One of

¹ Athena Wisdom, yaniv@athenawisdom.com

² Tel-Aviv University, Department of Industrial Engineering, shmueli@tau.ac.il

³ MIT Media Lab, {guyz,orened,sandy}@media.mit.edu

⁴ Telefonica Research, nuriao@tid.es

* Indicates equal contribution

the main thrusts that propels the constant expansions and enhancement of social network based services is its immense impact on the “real world” in a variety of fields such as politics, traditional industry, currency and stock trading and more. This field is becoming increasingly popular [51], [78], due to the possibility of increasing the impact of campaigns by using network related information in order to optimize the allocation of resources in the campaign. This relies on the understanding that a substantial impact of a campaign is achieved through the social influence of people on one another, rather than purely through the interaction of campaign managers with the people that are exposed to the campaign messages directly. A constantly growing portion of commercial and government marketing budgets is being allocated to advertising in social platforms the main goal of which is to spark viral phenomena that by spreading through the social networks would result in global “trends”.

2.1 Coverage optimization – theory and methods

The study of optimal coverage in pre-defined regions, by a distributed system of interacting units, has been the topic of many works in the past couple of decades. The work of [25] considers the problem of locating the minimum number of sensors on the network nodes in order to determine arc flow volumes of the entire network (a variant that considers dynamic environments is discussed in [97]). Influencing the behavior of large consumers population through prices manipulation campaigns by government agencies is discussed in [38]. In [79] a model for optimizing coverage using a multitude of units (in the form of sensors, for traffic surveillance purposes) was discussed.

In general, most of the analytic techniques used for guaranteeing maximal interaction using distributed actors, or “campaign units” use some sort of cellular decomposition of the region to be covered through the campaign. For example, in [30] a decomposition method is being used which is analytically shown to guarantee a complete coverage of an area. Another interesting work is presented in [3], discussing two methods for efficient coverage campaign using mobile units (e.g. cars with posters, or mobile advertising Zeppelins), one probabilistic and the other based on an exact cellular decomposition. Similar results can be found in [12], [13].

2.1.1 Diffusion optimization

Analyzing the spreading of information has been the focus of many social networks studies for the last decade [67] [77]. Researchers have explored both the offline networks structure by asking and incentivizing users to forward real mails and E-mails [45], and online networks by collecting and analyzing data from various sources such as *Twitter* feeds [72].

Researchers believe that such techniques can help understand the inter-influence of individuals in nowadays entangled world, comprising multi-layers of social and media networks [35], and that it can eventually lead to accurate prediction and active optimization and construction for successful and low-cost viral market campaigns, such as the DARPA Challenge [94]. However, the information diffusion process on

social networks is overwhelmingly complicated: the outcome is clearly sensitive to many parameters and model settings that are not entirely well understood and modeled correctly. As a result, accurate trend prediction and influence diffusion optimization are currently among the central research topics in the field.

The dramatic effect of the network topology on the dynamics of information diffusion in communities was demonstrated in works such as [36] [90]. One of the main challenges associated with modeling of behavioral dynamics in social communities stems from the fact that it often involves stochastic generative processes. While simulations on realizations from these models can help explore the properties of networks [63], a theoretical analysis is much more appealing and robust.

The identity and composition of an initial “seed group” in trends analysis has also been the topic of much research. Kempe et al. applied theoretical analysis on the seeds selection problem [70] based on two simple adoption models: *Linear Threshold Model* and *Independent Cascade Model*. Recently, Zaman et al. developed a method to trace rumors back in the topological spreading path to identify sources in a social network [105], and suggest that methods can be used to locate influencers in a network. Some scholars express their doubts and concerns for the influencer-driven viral marketing approach, suggesting that “everyone is an influencer” [18], and companies “should not rely on it” [115]. They argue that the content of the message is also important in determining its spreads, and likely the adoption model we were using is not a good representation for the reality.

2.1.2 Adoption model and social diffusion

A fundamental building block in trends prediction that is not yet entirely clear to scholars is the adoption model, modeling individuals’ behavior based on the social signals they are exposed to. Centola has shown both theoretical and empirically that a complex contagion model is indeed more precise for diffusion [32], [33]. Using social influence relations derived from online domains was discussed for example in [91]. Much research concerning the prediction of users’ behavior based on the dynamics in their community has been carried out in the past, using a variety of approaches such as sociological methods [58], [65], communities-oriented approaches [66], game theory [34] and various machine learning methods [92]. Different adoption models can dramatically alter the model outcome [46]. In fact, a recent work on studying mobile application diffusions using mobile phones demonstrated that in real world the diffusion process is a far more complicated phenomenon, and a more realistic model was proposed in [93].

2.2 Using mobile phones data for social systems modeling

The use of mobile phone data for the mobility and behavioral modeling of large population has become popular in the recent decade. In [74] the behavior and social patterns of 2.5 million mobile phone users, making 810 million phone calls, were analyzed and resulted in efficient mapping of users’ mobility and housing patterns. Similar result appears in [61],

In another example, it was shown that the penetration of cellular phones to the Israeli market is very high, even to lower income households, and specially among individuals in the ages of 10 to 70 (the main focus of travel behavior studies) [23]. This widespread use of cellular phones enables the collection of accurate mobility data that can be used to analyze and optimize coverage and monitoring campaigns.

For example, this data was shown in [23] and [118] to provide a high quality coverage of the network, tracking 94% of the trips (defined as at least 2km in urban areas, and at least 10km in rural areas). The resulting data contained a wealth of traffic properties for a network of over 6,000 nodes, and 15,000 directed links. In addition, the network was accompanied with an Origin Destination (OD) matrix, specifying start and end points of trips.

2.3 Campaign optimization studies

Whenever a company wishes to introduce a new product, increase its market share or merely retain the current one, it needs to engage itself in marketing efforts. In fact, the global marketing spend has been rising fast for several decades, and is currently estimated at 1 trillion dollars a year, which makes it between 1 to 2 percent of global GDP [56]. One important decision with regard to marketing involves finding the optimal balance between cost and effectiveness of the marketing campaign. An appropriate optimization method can help either to obtain more effective marketing results for a given budget or to reduce the marketing cost.

The advertising budgeting problem has been addressed in literature from different perspectives. Early models were relatively simple. [87] focused on the relationship between current marketing spending and future demand. In [112], Vidale and Wolfe represented sales response using three parameters - sales decay, saturation level and a response constant. Building upon the Vidale-Wolfe model, [104] used optimal control theory to obtain an optimal advertising strategy and [43] extended it to include competition in a duopoly.

[81] reviewed aggregate advertising models - functions that show the relationship between product sales and advertising spending for a market as a whole. He stressed that most models at the time often contradicted one another and missed key components, making it difficult to put these models into practical use.

Several works, such as [98] and [52], suggested that at least the short term response to advertising is S-shaped. Meaning that an increase in advertising is typically followed by a period of diminishing returns. In [84], Mesak and Hani provide oligopolistic justification for a pulsing advertising policy where S-shaped response functions are present.

[43] determines the optimal advertising expenditures for a duopoly in an equilibrium. [99] characterized the brands' choice and Nash equilibrium advertising expenditures in an oligopoly. [59] shows a model of oligopolistic competition in which advertising enters into the demand functions of firms, resulting in a positive relationship between product price and the degree of advertising cooperativeness. [89] incorporated random demand for a product to show how demand uncertainty

affects advertising decisions. [68] showed that when the market is saturated, a brand should choose a defensive strategy, in which the goal is preserving the existing customers and sales. [22] explains how to choose between generic advertising and brand advertising strategies in a dynamic duopoly.

Several researchers sought methods of helping marketing managers allocate a given budget and optimize response. [83] focused on assisting marketing managers in optimizing advertising budget. They presented the first comprehensive allocation model that, given a fixed budget, finds the optimal spread over time and market segments. Using simple inputs, the model created by [101] efficiently select advertising schedules for network television.

Instead of looking at advertising as an expense, [40] argues that advertising should be regarded as an investment with long-lived effects. They then lay out an approach for calculating the level of spending that generates optimal return. Following this approach, [41] supplies the formulas required for calculating the level of spending that maximizes ROI. [85] uses a Markov decision process to formulate a stochastic, sequential model that takes into account the maturity and past advertising of the product and determines the optimal advertising spending. [73] considers the advertising investment in a spatial monopoly, contrasting the socially optimal behavior of a benevolent planner against that of a profit seeking monopolist.

The formulas in [41] provide a solution for a single product, single medium and a single stage. [48] addresses the multi-product advertising budgeting problem, in which the cross-effects of the advertised products are taken into consideration, as well the effect of the promoted products on the rest of the products portfolio. [106] and others studied multistage advertising budgeting, showing the short-term and long-term impact on demand.

Additional work has been done by [24], who addressed the multistage multiproduct advertising budgeting problem - optimizing the budget and allocation for multiple products, multiple sale attributes over multiple periods. They tested their model in an actual campaign and observed a clear increase in profits compared to other approaches. In the second part of their work, [24], they introduce a stochastic optimization model and compare it with the deterministic model presented on their early work.

So far, we have reviewed models used for finding optimal expenditures and coverage in marketing campaign. However, a variation of these problems can be found in other domains as well. For example, in [2], the authors discuss two methods for efficient coverage using mobile units, one probabilistic and the other based on an exact cellular decomposition. Similar results and methods can be found in [12], [14]. The work of [26] considers the problem of locating the minimum number of sensors on the network nodes in order to determine arc flow volumes of the entire network. The dynamic environments variant is discussed in [97]. The study of the correlation between topological features of a marketing environment network and the efficiency of a distributed group of (mobile) units is discussed in [16]. In [80] a model for optimizing coverage using a multitude of units (in the form of sensors, for traffic surveillance purposes) was discussed. In general, most of the

analytic techniques used for guaranteeing maximal interaction using distributed actors, or campaign units use some sort of cellular decomposition of the region to be covered through the campaign. For example, in [31] a decomposition method is being used which is analytically shown to guarantee a complete coverage of an area. These methods are the equivalent of the optimization methods used for finding the optimal coverage using various media (e.g. cars with posters, or mobile advertising, Zeppelins) in the world of marketing campaigns.

2.4 Campaign optimization in practice

The author of [82] pointed out that a model that is to be used by a manager should be simple, robust, easy to control, adaptive, as complete as possible and easy to communicate with. In accordance with this recommendation, the models that we survey in this section are simple but realistic enough to be used in the advertising industry.

2.4.1 Estimating effectiveness

In the past, short-term measurement of advertising effectiveness was extremely problematic because of the inherently long-term nature of advertising impact and the very small short-term effects of advertising [1], [29], [110], [111]. Thus, as a practical matter, the media planner often employed a proxy for advertising effectiveness [41]. Popular choices for such a proxy include reach [41], effective reach [4], [39], [86] and average frequency [41]. Reach is the proportion of the target audience exposed to at least one insertion of the advertisement. Effective reach is the proportion of the target audience exposed to at least three insertions of the advertisement. Frequency is the average number of times a person from the reach audience is exposed to an advertisement.

Exposure to an advertisement is often measured in Gross Rating Points (GRPs): the product of reach and frequency. For example, 100 GRPs could mean that 100% of the market is exposed once to an advertisement or that 50% of the market is exposed twice [62]. According to [21], it is usually preferred to measure advertising in GRPs and not in dollars since: (1) most managers evaluate the effectiveness of their campaigns in terms of demand generated per GRP; and (2) it is not clear how much advertising exposure can be purchased for a given budget, and thus GRPs provide a clearer picture of advertising input.

We denote the reach for k insertions as r_k , and for g GRPs as r_g . Similarly, we denote the effective reach for k insertions as er_k , and for g GRPs as er_g .

Traditionally, advertising campaigns are quantitatively described by the exposure distribution (ED), defined as the probability of exposure to none, one, up to all of the ads in the campaign [42], [102]. Denoting the exposure random variable as X , ranging from $0, 1 \dots k$, where k is the total number of insertions in the campaign. It was found in [76] that the most frequently used nonproprietary model for X is the beta-binomial distribution (BBD), with mass function:

$$f_X(x) = \binom{k}{x} \cdot \frac{\Gamma(\alpha + \beta) \cdot \Gamma(\alpha + x)}{\Gamma(\alpha + \beta + k) \cdot \Gamma(\alpha)} \cdot \frac{\Gamma(\beta + k - x)}{\Gamma(\beta)}$$

where Γ is the gamma function and $\alpha, \beta > 0$. The authors of [41] showed that under the BBD model, the reach function can be modeled as:

$$r_k = 1 - \prod_{j=0}^{k-1} \frac{\beta + j}{\alpha + \beta + j}$$

and the number of GRPs can be modeled as:

$$g = 100k \cdot \frac{\alpha}{\alpha + \beta}$$

That is, the number of GRPs g depends linearly on the number of insertions k . At first, this result may seem strange, as the same value of GRPs may have several corresponding insertions values. In reality, however, for a moderately large number of insertions, the GRPs-insertions curve is quite flat. This means that, although there are many possible insertions values for the same GRPs, the fluctuation in GRPs for these combinations are small [109].

In recent years, it has become significantly easier to estimate the effectiveness function directly. We are living the era of Big Data, where companies gather and manage huge databases. By using market response models we can transform this raw marketing information into 'ready to use' information [62]. As a concrete example, [24] models the sales due to advertising as a function of the number of GRPs, denoted as $sales_g$.

Regardless of whether the effectiveness function f_g is estimated directly or via a proxy, in the single product scenario, it corresponds to an increasing concave function which models diminishing returns [62].

The authors of [41] argue that the effectiveness proxies r_g and er_g can be well approximated by a function of the following form:

$$f_g \approx 1 - \gamma g^{-\lambda}$$

where $\gamma, \lambda > 0$ and $g > g_0$ for some lower threshold value of GRPs.

The authors of [62] argue that a better choice for approximating f_g is the so called 'modified exponential' function:

$$f_g \approx \gamma(1 - e^{-\delta g})$$

In the multi product scenario, however, there exist cross elasticities among the products due to relationships of complementarity or substitution [48]. In the case of complementarity (positive elasticity), advertising on one product increases sales of another product and this cross effect can be modeled by an increasing concave function. In the case of substitution (negative elasticity), however, advertising on one product reduces sales of the other product (this cross effect is known as cannibalization [24]). The cannibalization effect can be modeled by a decreasing convex function. If this function is strictly convex, the resulting effectiveness function may not be concave. As is well known, concavity of the objective function is a desirable property in a maximization problem since it guarantees global optimality (assuming a convex feasible domain). Having said that, the cross product effects are usually small relative to the

direct advertising effects, and therefore, in most cases, it is reasonable to assume that the effectiveness function can be modeled by an increasing concave function.

In Section 4, we suggest a method to derive the reach function using real-world data, and propose to model it using a Gompertz function instead of a 'modified exponential' function.

2.4.2 Estimating cost

the authors of [40] have proposed a mathematical relationship between cost and GRPs. The function employed makes two assumptions: (1) more GRPs cost more than fewer GRPs and thus cost is a monotonically increasing function of GRPs; and (2) buying a large number of GRPs can result in discounts and thus the cost curve is concave. The authors further suggest the following flexible functional form to model the cost function c_g .

$$c_g = C \cdot g^\delta$$

where $C > 0$ is a constant and $0 \leq \delta \leq 1$ is a parameter which reflects the expected discounting extent.

In the remainder of this paper we assume that no discounting occurs and thus $\delta = 1$.

$$c_g = C \cdot g$$

2.4.3 Optimal criteria

The authors of [40] and [41] categorize spending criteria into popular ad hoc criteria, and optimal criteria, which are based on modeling and solving an optimization problem. Although criteria that are actively used in practice often tend to be in the ad hoc category, we focus here on the optimal category.

The three primary approaches for optimizing the level of marketing spending are: (1) maximizing advertising profitability, (2) maximizing advertising productivity (efficiency), and (3) maximizing the return on investment of advertising.

Maximizing advertising profitability

This approach, developed by Kaplan and Shocker [69] starts with the assumption that an advertising effectiveness measure exists (perhaps effective reach), and that this measure is directly related to revenue. This seemingly strong assumption is not so unreasonable, because almost all media planners (including the most sophisticated ones) currently use surrogates of advertising effectiveness.

If we denote profitability as E_1 , the measure of advertising effectiveness as f_g , where g is the number of GRPs, and c_g is the cost of buying g GRPs, then the profitability is:

$$E_1 = K \cdot f_g - c_g$$

where K is the dollar value of one unit of effectiveness.

Maximizing advertising productivity (efficiency)

Maximizing productivity or efficiency is another alternative. Economists often emphasize the importance to the economy

of increasing productivity, and management scholars recognize that the greatest potential for gains in productivity are in the knowledge and service sectors of the economy [49].

Using the above notations, productivity (efficiency) is calculated as:

$$E_2 = \frac{K \cdot f_g}{c_g}$$

Maximizing the return on investment

This approach was proposed in [44]. Using the above notations, return on investment (ROI) is calculated as:

$$E_3 = \frac{K \cdot f_g - c_g}{c_g} = \frac{K \cdot f_g}{c_g} - 1$$

3 CHARACTERIZATION OF A CAMPAIGN

Let us define a campaign as an activity confined in time and space with a limited budget whose objective is to send a message or engage with the maximum number of individuals who are located in that space.

More specifically, based on the conventions laid out in the previous section, let us define an *insertion* as some kind of induced intervention that is designed to incentivize certain kinds of inclined behavior in the audience. Such insertions can be for example large billboards, human agents who hand coupons to passing pedestrians, etc. Let us also define a *deployment scheme* as a specific allocation algorithm that for a given number of insertions k outputs a set of k locations for said insertions. Finally, let us define the reach function, r_k , as the number of individuals exposed to at least one insertion (given that exactly k insertions were deployed according to the given deployment scheme).

Selecting profitability as our optimal criteria (see Section 2.4.3), the reach function r_g as the effectiveness measure f_g , and $C \cdot g$ as the cost function c_g (see Section 2.4.2), the *Optimized Campaign Problem* aims at maximizing the following target function:

$$E_g = K \cdot r_g - C \cdot g \quad (1)$$

Given that $k \approx g$ (see Section 2.4.1), the above equation can be rewritten to use the number of insertions k instead of the number of GRPs g .

$$E_k = K \cdot r_k - C \cdot k \quad (2)$$

Next, we are left with efficiently modeling the reach function r_k . Assume that the network of mobility between all possible locations is available (In section 5 we demonstrate such a mobility network which is inferred from CDR data). Given a number of insertions k and a deployment scheme, the reach function r_k can be well approximated by calculating the Group Betweenness Centrality (GBC) [54] of the k locations (i.e. nodes) which are returned by the deployment scheme.

Betweenness Centrality (BC) stands for the ability of an individual node to control the communication flow in a network

and is defined as the total fraction of shortest paths between each pair of vertices that pass through a given node [17], [57]. In recent years Betweenness was extensively applied to analyze various complex networks [20], [108] including social networks [103], [114], computer communication networks [55], [119], and protein interaction networks [27]. Holme [64] has shown that Betweenness is highly correlated with congestion in particle hopping systems. Extensions of the original definition of BC are applicable for directed and weighted networks [28], [117] as well as for multilayer networks where the underlying infrastructure and the origin-destination overlay are explicitly defined [95].

The GBC of a given group ($U \subseteq V$) of vertices accounts for all routes that pass through *at least one* member of the group. Let $\sigma_{s,t}$ be the number of shortest-path routes from s to t , and let $\sigma_{s,t}(U)$ be the number of shortest-path routes from s to t passing through at least one vertex in U :

$$GBC(U) = \sum_{s,t \in V \setminus C | s \neq t} \frac{\sigma_{s,t}(U)}{\sigma_{s,t}} \quad (3)$$

While the ‘‘optimal deployment scheme’’, by definition, would select the set of k locations that yield the maximal group betweenness centrality, our proposed model does not assume any constraint on the deployment scheme and enables the planners to select the optimal number and type of the units, as a function of the deployment scheme used. More specifically, in many cases the optimal deployment scheme might be unfeasible, involve additional costs, or be subject to various regulatory constraints. In such cases campaign managers may choose a different, non-optimal, deployment scheme. Regardless of the deployment scheme chosen, it would of course still provide monotonically increasing reach (and GBC), albeit with a lower incline rate compared to the optimal one.

Finally, we model the approximated reach function, using the well-known *Gompertz function* [60]:

$$r_k = ae^{be^{ck}} \quad (4)$$

The *Gompertz function* is widely used for modeling a great variety of processes, (due to the flexible way it can be controlled using the parameters a , b and c), such as mobile phone uptake [100] or population in a confined space [53]. Its ability to model the progress of optimization process as a function of the available resources can be seen for example in [5]–[7], [15]. In Section 5, we present empirical evidence which illustrates how the reach function r_k approximated by the GBC of three different deployment schemes can be fitted efficiently into a Gompertz function.

Assigning the values of Equation 4 back into Equation 2, results in the following target function:

$$E_k = K \cdot ae^{be^{ck}} - C \cdot k \quad (5)$$

The campaign will be optimized by determining the optimal number of insertions (k) and the optimal cost (C) of each individual insertion that would maximize the campaign’s performance.

4 OPTIMIZED CAMPAIGNS

At this point, we have a clear model for estimating the efficiency of a campaign, that is dependent on the number of insertions and the cost of each insertion.

4.1 Optimizing the Number of Insertions

We now turn our attention to finding the optimal number of insertions k that would maximize the profitability of the campaign. First, we obtain the derivative to find the critical points of the function we seek to optimize:

$$\frac{\partial E_k}{\partial k} = K \cdot \frac{\partial(ae^{be^{ck}})}{\partial k} - C \quad (6)$$

Nullifying Equation 6 results in:

$$\frac{\partial(ae^{be^{ck}})}{\partial k} = \frac{C}{K} \quad (7)$$

In that case, using Equation 7 we obtain:

$$a \cdot b \cdot c \cdot e^{ck} \cdot e^{be^{ck}} = \frac{C}{K}$$

which in turn implies:

$$be^{ck} + ck - \ln \frac{C}{a \cdot b \cdot c \cdot K} = 0 \quad (8)$$

We note that $a, b, c > 0$. Analyzing Equation 8 we can then see that in cases where:

$$\frac{C}{K} \leq -\frac{a \cdot c}{e} \quad (9)$$

and where $W(x)$ is the *Lambert product log*, that can be calculated using the series:

$$W(x) = \sum_{n=1}^{\infty} \frac{(-1)^{n-1} n^{n-2}}{(n-1)!} x^n \quad (10)$$

the optimal value of k would equal:

$$\boxed{\begin{aligned} k_1 &= \frac{\ln\left(\frac{1}{a \cdot b \cdot c} \cdot \frac{C}{K}\right) - W\left(\frac{1}{a \cdot c} \cdot \frac{C}{K}\right)}{c} \\ k_2 &= \frac{\ln\left(\frac{1}{a \cdot b \cdot c} \cdot \frac{C}{K}\right) - W_{-1}\left(\frac{1}{a \cdot c} \cdot \frac{C}{K}\right)}{c} \end{aligned}} \quad (11)$$

Note that $W(x)$ is the *Lambert product log*, and $W_k(x)$ is its analytic continuation over the complex plane (the values of the functions $W(x)$ and $W_{-1}(x)$ in the segment implied by the constraint of Equation 9 are illustrated in Figure 1).

Returning to the optimization of the campaign, we now assign the values of the optimal number of insertions of Equation 11 into the definition of E_k , as follows:

$$E_{k_1} = K \cdot ae^{be^{ck_1}} - k_1 \cdot C$$

$$E_{k_2} = K \cdot ae^{be^{ck_2}} - k_2 \cdot C$$




Fig. 1. The upper and lower charts depict the values of the Lambert functions $W(x)$ and $W_{-1}(x)$ in the segment $[-\frac{1}{e}, 0]$, respectively. The segment is implied by the constraint of Equation 9.

and using the properties of the W function, simplify it into the following form:

$$E_{k_{max}} = \max\{E_{k_1}, E_{k_2}\} \quad \text{where:} \quad (12)$$

$$E_{k_1} = a \cdot b \cdot K \cdot \gamma \cdot \left(W(b \cdot \gamma) + \frac{1}{W(b \cdot \gamma)} - \ln(\gamma) \right)$$

$$E_{k_2} = a \cdot b \cdot K \cdot \gamma \cdot \left(W_{-1}(b \cdot \gamma) + \frac{1}{W_{-1}(b \cdot \gamma)} - \ln(\gamma) \right)$$

where the Campaign Benefit Factor γ is defined as:

$$\gamma = \frac{1}{a \cdot b \cdot c} \cdot \frac{C}{K}$$

From Equation 12 we see that the optimization of a campaign is a function of the Campaign Benefit Factor γ , which takes into account the dollar value of one unit of effectiveness, the total cost, as well as the deployment scheme (characterized by the values of a , b and c). Notice that the value of C does not affect the values of a , b and c (as they are solely derived from the coverage efficiency of the mobility patterns).

4.2 Optimizing the Cost of a Single Insertion

We now proceed to finding the optimal type of units that should be deployed, by optimizing Equation 12 with respect to the cost of various possible kinds of insertions.

Definition 1: Let $Cost_{Base}$ denote the cost of the “most expensive” insertions which we will assume are the units that provide the highest value to the initiators of the campaign,

namely — those insertions that persuade the maximal amount of individuals to take the desired action.

We hereby define K , the dollar value of one unit of effectiveness, as a function dependent on the proportion between the cost of the given type of insertions and the cost of the optimal, but most expensive, insertions:

$$K = f_S \left(\frac{C}{Cost_{Base}} \right) \quad (13)$$

We now look into finding the optimal cost of insertions that would maximize the profitability of a campaign. To do so, we first revise Equation 12 in order to take into account the different types of insertions:

$$E_{k_1} = a \cdot b \cdot f_S \cdot \gamma \cdot \left(W(b \cdot \gamma) + \frac{1}{W(b \cdot \gamma)} - \ln(\gamma) \right) \quad (14)$$

$$E_{k_2} = a \cdot b \cdot f_S \cdot \gamma \cdot \left(W_{-1}(b \cdot \gamma) + \frac{1}{W_{-1}(b \cdot \gamma)} - \ln(\gamma) \right)$$

and where:

$$\gamma = \frac{1}{a \cdot b \cdot c} \cdot \frac{C}{f_S}$$

We then maximize the financial merits of the campaign (namely, $\max\{E_{k_1}, E_{k_2}\}$), by calculating the partial derivatives $\frac{\partial E_{k_1}}{\partial C}$ and $\frac{\partial E_{k_2}}{\partial C}$:

$$\frac{\partial E_{k_1}}{\partial C} = \quad (15)$$

$$\frac{1}{c} \cdot \left(W(b \cdot \gamma) - \ln(\gamma) + \frac{1 - \frac{\partial f_S}{\partial C} \cdot \frac{C}{Cost_{Base} \cdot f_S}}{W(b \cdot \gamma)} \right)$$

and:

$$\frac{\partial E_{k_2}}{\partial C} =$$

$$\frac{1}{c} \cdot \left(W_{-1}(b \cdot \gamma) - \ln(\gamma) + \frac{1 - \frac{\partial f_S}{\partial C} \cdot \frac{C}{Cost_{Base} \cdot f_S}}{W_{-1}(b \cdot \gamma)} \right)$$

Once again, by nullifying the partial derivative and finding the critical points, we obtain the following set of equations:

$$\begin{aligned} & 0 = W(b \cdot \gamma) - \ln(\gamma) + \\ & \frac{1 - \frac{C}{Cost_{Base}} \cdot \frac{\partial f_S}{\partial C} \left[\frac{C}{Cost_{Base}} \right] \cdot \left(f_S \left[\frac{C}{Cost_{Base}} \right] \right)^{-1}}{W(b \cdot \gamma)} \\ & \text{or :} \\ & 0 = W_{-1}(b \cdot \gamma) - \ln(\gamma) + \\ & \frac{1 - \frac{C}{Cost_{Base}} \cdot \frac{\partial f_S}{\partial C} \left[\frac{C}{Cost_{Base}} \right] \cdot \left(f_S \left[\frac{C}{Cost_{Base}} \right] \right)^{-1}}{W_{-1}(b \cdot \gamma)} \\ & \text{where :} \\ & \gamma = \frac{1}{a \cdot b \cdot c} \cdot \frac{C}{K} \cdot \left(f_S \left[\frac{C}{Cost_{Base}} \right] \right)^{-1} \end{aligned} \quad (16)$$

Equation 16 can now be used to calculate the exact optimal cost of a single insertion, for every cost-value relation, and for every deployment scheme!

5 CAMPAIGN OPTIMIZATION FOR A REAL-WORLD MOBILITY NETWORK

In this section we validate our campaign optimization model using a real-world mobility network inferred from CDR data.

5.1 The Dataset

We used Call Data Records (CDR) from a large mobile carrier to create a network that captures the traveling patterns among different urban areas. CDR are readily available today, as they are collected by all carriers and in most (if not all) countries. Furthermore, these records are constantly collected in an automated manner, thus increasing the likelihood of the data being objective and uniform across locations and operators.

More specifically, we denote $G = \langle V, E \rangle$ to be the undirected network graph, where V is the set of vertices representing the cell towers and E is the set of weighted edges representing trips or movements of people between two cell towers. The weight of each edge represents the number of trips that people made between the cell towers connected by that edge.

Each CDR contains an anonymized identifier of the caller/callee, the call time, and the cell tower that the phone was connected to when the call (or SMS) originated. A trip is defined as a change of location by the caller/callee, detected by the existence of two consecutive calls from two different towers or by a change in the cell tower during an existing call.

As shown in figure 2 (left), the weights of the edges in the network seem to follow a power-law distribution. We used the method suggested in [37] and [113] to determine the best $Xmin$ and corresponding γ values that fit a power-law distribution. Since the focus of this paper is not related to the network's topology, our analysis did not include the statistical tests performed in [37]. Figure 2 (right) shows the Probability Density Function (PDF) and the best power-law fit with $Xmin = 3$ and $\gamma = 1.87$.

In order to focus on the network's structure that represents urban mobility patterns of large populations, we retained only edges with weights higher than 10 (arbitrarily chosen threshold), producing a graph with $|V| = 18,315$ and $|E| = 130,313$.

The Complementary Cumulative Distribution Function (CCDF) of node degrees is shown in Figure 3 (left). Once again, we used the method suggested by [37] and [113] to find the parameters that best fit a power-law distribution. Figure 3 (right) shows the PDF and the best power-law fit with $Xmin = 45$ and $\gamma = 5.05$.

5.2 Optimized Deployment Schemes

Acquiring information regarding the mobility patterns of the audience members, and the specific network that is generated

through those patterns can be used to subsequently derive optimized locations for the campaign units, or to the very least, provide a way of measuring the utilization of a set of locations, by calculating the GBC of the set, and comparing it to that of the optimal one.

Several combinatorial optimization techniques can be used to find a group of nodes of given size that has the largest GBC, including greedy approximation [47], a classical *Depth First Branch and Bound* (DFBnB) heuristic search algorithm [71], or the recently proposed *Potential Search* [107]. Both the *DFBnB* and the *Potential* algorithms are anytime search algorithms [120], meaning that their execution can be stopped at any point of time, yielding the best solution found so far.

Similarly to [96], in this section, we examine three methods for finding the group of size k with the highest GBC, and apply them on our real-world mobility network. The first method, Random Deployment, simply selects a random set of k vertices. The second method, BC Deployment, chooses the k vertices with the highest individual Betweenness Centrality scores. The third method, GBC Deployment, uses a greedy algorithm which iteratively finds the vertex which improves the group's GBC score the most and adds it to the group.

Due to the runtime complexity of the GBC calculation and the GBC deployment scheme, we performed two stages of sampling our mobility network. First, a subset of 5000 vertices was randomly selected. Then, we retained only the largest connected component, denoted by G' , a network containing 991 vertices and 3,304 edges. An illustration of the resulting network is shown in figure 4.



Fig. 4. Network G'

It can be seen that this network is a good sampling of the original network, as the median distance from each trimmed node to the closest node in the sampled network is 4, while the diameter of the graph itself is 22. Figure 5 (left) shows the distribution of distances of trimmed vertices from nodes in the sampled network and Figure 5 (right) depicts the PDF and the best power-law fit with $Xmin = 3$ and $\gamma = 2.08$.

We examined the three deployment schemes variants, calcu-




Fig. 2. Edge weights in network G . The left figure shows distribution of edge weights. The right figure depicts the Probability Density Function




Fig. 3. Distribution of the node degrees of G . The left figure depicts the Complementary Cumulative Distribution Function for all degrees. The right figure illustrates the Probability Density Function for all degrees higher than X_{min}

lated their corresponding GBC values, and fitted them onto a *Gompertz function* using regression. Figure 6 illustrates the three deployment schemes for our mobility network. Their fitting yielded the following Gompertz regressions:

$$\begin{aligned}
 \text{Random Deployment : } r_k &= 0.76e^{-3.14e^{-0.02k}} \\
 \text{BC Deployment : } r_k &= 0.92e^{-0.94e^{-0.2k}} \\
 \text{GBC Deployment : } r_k &= 0.96e^{-1.25e^{-0.38k}}
 \end{aligned} \quad (17)$$

The regressions had the following fit quality (in terms of

R^2):

- 1) Random Deployment - 0.6683
- 2) BC Deployment - 0.8035
- 3) GBC Deployment - 0.8531

Consistently with [95], using the GBC Deployment we saw that it is possible to cover the vast majority of the most popular mobility nodes with a few dozen insertions. The BC Deployment produced high quality deployments as well, although based on our findings it required a higher number of insertions. It seems that both schemes significantly outperform




Fig. 5. Distribution of distances between trimmed vertices and nodes in the sampled network G' (left) and Probability Density Function of the distances between trimmed vertices and nodes in the sampled network G' (right)

Random Deployment, which required 100 insertions to reach 50 percent coverage, and a few hundred nodes to guarantee a near-full one.

It is important to note that r_k can significantly change for different networks (modeling different urban environments mobility patterns). With this in mind, we can now proceed to finding the optimal number of insertions, and subsequently — the maximization of the monetary utilization of the investment in the system.

5.3 Optimizing the number of insertions and the cost per insertion

In this section we demonstrate how the proposed method can be used to significantly increase the utilization of a given campaign.

For simplicity, we assume that both K and the cost per unit (C) are continuous, and follow some pre-defined function, known to the campaign managers., For example, we may imagine a function that follows a sub-linear correlation, such as the following:

$$f_{S_1} \left(\frac{C}{Cost_{Base}} \right) = \left(\frac{C}{Cost_{Base}} \right)^2$$

The meaning of this function is that insertions that cost half of the optimal insertions possible, would generate 25% value of the optimal (most expensive) ones. Alternatively, we may imagine environments where the correlation between insertions' cost and value is super-linear, converging to one, such as the function:

$$f_{S_2} \left(\frac{C}{Cost_{Base}} \right) = \sqrt{\frac{C}{Cost_{Base}}}$$

In this analysis we shall use the values of the Gompertz approximation as shown above in Equation 17:

- **Random Deployment** : $a = 0.76, b = 3.14, c = 0.02$
- **BC deployment** : $a = 0.92, b = 0.94, c = 0.2$
- **GBC deployment** : $a = 0.96, b = 1.25, c = 0.38$

In this case, nullifying the partial derivative of Equation 15 for GBC deployment, and assuming f_{S_1} as a measurement of K , would yield:

$$\frac{\partial E_{k_1}}{\partial C} = 0 \quad \rightarrow \quad (18)$$

$$W(1.25 \cdot \gamma) = \ln(\gamma) + \frac{1}{W(1.25 \cdot \gamma)}$$

and:

$$\frac{\partial E_{k_2}}{\partial C} = 0 \quad \rightarrow$$

$$W_{-1}(1.25 \cdot \gamma) = \ln(\gamma) + \frac{1}{W_{-1}(1.25 \cdot \gamma)}$$

subsequently implying:

$$\gamma_{opt} \approx 0.2837$$

(in this example, the optimal value of γ for E_{k_2} has a non-zero imaginary component).

Using this optimal value of γ we would now get:

$$\gamma_{opt} = \frac{1}{a \cdot b \cdot c} \cdot \frac{C}{f_{S_1}} = \frac{2.19 \cdot Cost_{Base}^2}{C} = 0.2837$$

and from this we receive:

$$C_{opt} = \frac{2.19 \cdot Cost_{Base}^2}{0.2837 \cdot K} \approx 7.72 \cdot Cost_{Base}^2 \quad (19)$$




Fig. 6. GBC for three deployment schemes with respect to the number of used units, using mobility patterns extracted from mobile phones data (left to right): (a) Random Deployment, (b) BC Deployment, and (c) GBC Deployment. The appropriate Gompertz fit of the curves is also included.

From Equation 19 we can obtain for each kind of campaign the optimal type of insertions that should be used, in order to maximize Equation 5 — the optimization function that maximizes a campaign’s profitability.

Assigning this back into Equation 11, we can get the optimal number of insertions for each potential campaign (see an illustration in Figure 7):

$$k_1 = \quad (20)$$

$$2.63 \cdot W(2.74 \cdot Cost_{Base}^2) + 0.786 - 5.26 \cdot \ln(Cost_{Base})$$

and:

$$k_2 =$$

$$2.63 \cdot W_1(2.74 \cdot Cost_{Base}^2) + 0.786 - 5.26 \cdot \ln(Cost_{Base})$$

Note that the previous example depends on the assumption of a quadratic relation between the cost and quality of the insertions, and on the assumption regarding the deployment scheme (*i.e.* the parameters of the *Gompertz Model*). However, given any substitute for these assumptions, corresponding solutions to the optimized campaign problem will be generated by the model.

Repeating the above process with the Gompertz fitted values for BC Deployment, as well as Random Deployment, we can obtain the optimized campaigns (in terms of cost per insertion, and number of insertions) for these deployment schemes.

The above analysis examined the profitability of the campaign (see again Section 2.4.3). If we are interested in the return on investment of the campaign, we can divide the profitability of the campaign by its cost (*i.e.* the number of insertions times the cost per insertion). Figure 8 shows the return on investment of the three deployment schemes as a function of their profitability. The figure illustrates the superiority of the GBC method, as well as the performance of the optimized campaigns, compared to non-optimized ones.

6 CONCLUSIONS

In this paper we studied the problem of campaign optimization. We started by formalizing the problem of optimizing a campaign by finding the optimal trade-off between the resources

(cost) allocated to each single unit, and the number of such units. We have presented a novel model to analytically generate an optimized strategy for marketing campaigns, and demonstrated how it can be used with aggregated and anonymized mobility data received from mobile carriers.

Specifically, we have shown a way to analytically calculate the exact optimal cost for units in a campaign, as well as the optimal number of such units, that would guarantee a maximal utilization of the campaign’s budget.

In this work we have discussed the optimization of campaigns resources utilization. However, the exposure results to be guaranteed using such resources and the proposed method were left outside the scope of this work. This aspect however (namely, theoretical lower bounds for *any conceivable* campaign strategy), was discussed in works such as [13] – where a generic analytical bound was developed, under the assumption that the campaign’s target will practice an adversarial strategy that would minimize their exposure, or [8], [9], [13] – where the impact of changes in the topological properties of the environment and the campaign’s theoretical utilization were analyzed.

Finally, it is interesting to note that the problem of finding an optimal (and optionally dynamic) “engagement strategy” is related to other kinds of monitoring problems, such as monitoring for evading land targets by a flock of Unmanned Air Vehicles (UAV). In this problem, however, the fact that the paths of the UAVs is unconstrained (as they are flying in the air) makes the calculation of a near-optimal monitoring strategy fairly easy [10], [11].

REFERENCES

- [1] Magid M Abraham and Leonard M Lodish. Getting the most out of advertising and promotion. *Harvard Business Review*, 68(3):50, 1990.
- [2] Ercan Acar, Yangang Zhang, Howie Choset, Mark Schervish, Albert G Costa, Renata Melamud, David Lean, and Amy Graveline. Path planning for robotic demining and development of a test platform. In *International Conference on Field and Service Robotics*, pages 161–168, 2001.

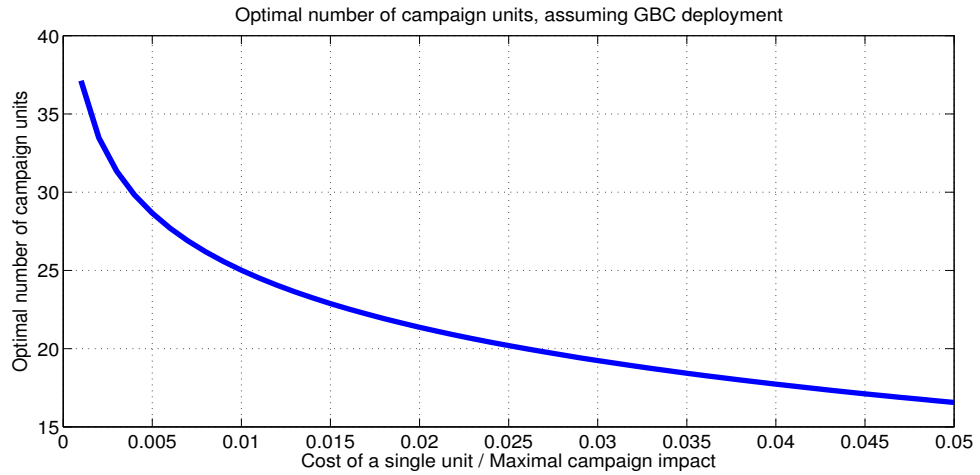


Fig. 7. An illustration of Equation 16, under the GBC Deployment scheme, as approximated by the Gompertz function modeling of Equation 17. For example, when a single unit costs 1% of the maximal campaign impact, the optimal number of campaign units would be 25, whereas if cheaper units are used (such as units that cost merely $\frac{1}{2}\%$ of the maximal campaign impact) the optimal number of units would be 28. Alternatively, for expensive units that cost 5% of the maximal campaign's impact, the optimal number of campaign units would drop to 16.

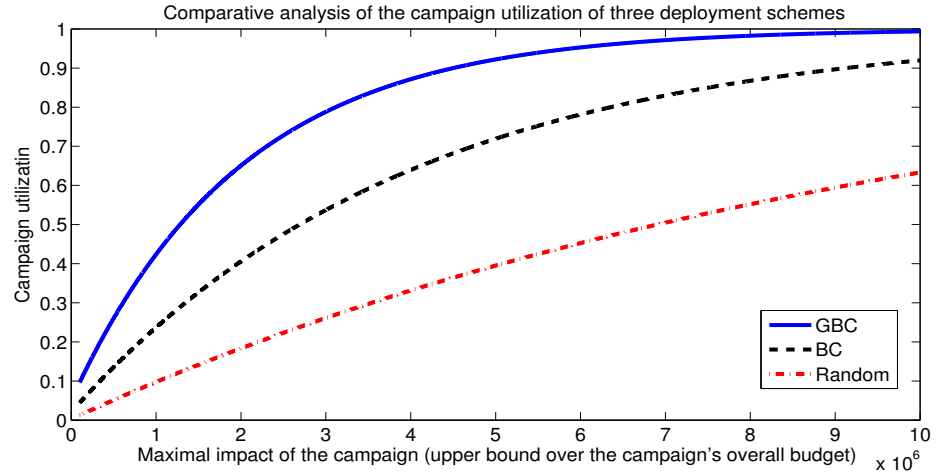


Fig. 8. Comparative analysis of 3 deployment schemes – GBC Deployment, BC Deployment and Random Deployment, in terms of their return on investment, as a function of their profitability. It can be seen that the return on investment of all methods is monotonically increasing. However, it can clearly be seen that whereas Random Deployment achieves only a slow increase in return on investment, the performance of GBC Deployment and BC Deployment are much better.

- [3] E.U. Acar, Y. Zhang, H. Choset, M. Schervish, A.G. Costa, R. Melamud, D.C. Lean, and A. Gravelin. Path planning for robotic demining and development of a test platform. In *International Conference on Field and Service Robotics*, pages 161–168, 2001.
- [4] Alvin A Achenbaum. Effective exposure: A new way of evaluating media. In *Association of National Advertisers Media Workshop, New York, 1977*.
- [5] Y. Altshuler, N. Aharony, M. Fire, Y. Elovici, and A. Pentland. Incremental learning with accuracy prediction of social and individual properties from mobile-phone data. *CoRR*, 2011.
- [6] Y. Altshuler, N. Aharony, A. Pentland, Y. Elovici, and M. Cebrian. Stealing reality: When criminals become data scientists (or vice versa). *Intelligent Systems, IEEE*, 26(6):22–30, nov.-dec. 2011.
- [7] Y. Altshuler, M. Fire, N. Aharony, Y. Elovici, and A. Pentland. How many makes a crowd? on the correlation between groups' size and the accuracy of modeling. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction*, pages 43–52. Springer, 2012.
- [8] Y. Altshuler, I.A. Wagner, and A.M. Bruckstein. Shape factor's effect on a dynamic cleaners swarm. In *Third International Conference on Informatics in Control, Automation and Robotics (ICINCO), the Second International Workshop on Multi-Agent Robotic Systems (MARS)*, pages 13–21, 2006.
- [9] Y. Altshuler, I.A. Wagner, and A.M. Bruckstein. On swarm optimality in dynamic and symmetric environments. volume 7, page 11, 2008.
- [10] Y. Altshuler, V. Yanovski, I.A. Wagner, and A.M. Bruckstein. The

- cooperative hunters - efficient cooperative search for smart targets using uav swarms. In *Second International Conference on Informatics in Control, Automation and Robotics (ICINCO), the First International Workshop on Multi-Agent Robotic Systems (MARS)*, pages 165–170, 2005.
- [11] Y. Altshuler, V. Yanovsky, A.M. Bruckstein, and I.A. Wagner. Efficient cooperative search of smart targets using uav swarms. *ROBOTICA*, 26:551–557, 2008.
- [12] Y. Altshuler, V. Yanovsky, I. Wagner, and A. Bruckstein. Swarm intelligencesearchers, cleaners and hunters. *Swarm Intelligent Systems*, pages 93–132, 2006.
- [13] Yaniv Altshuler and Alfred M. Bruckstein. Static and expanding grid coverage with ant robots: Complexity results. *Theoretical Computer Science*, 412(35):4661–4674, 2011.
- [14] Yaniv Altshuler and Alfred M Bruckstein. Static and expanding grid coverage with ant robots: Complexity results. *Theoretical Computer Science*, 412(35):4661–4674, 2011.
- [15] Yaniv Altshuler, Michael Fire, Nadav Aharoni, Zeev Volkovich, Yuval Elovici, and Alex Sandy Pentland. Trade-offs in social and behavioral modeling in mobile networks. In *Social Computing, Behavioral-Cultural Modeling and Prediction*, pages 412–423. Springer, 2013.
- [16] Yaniv Altshuler, Israel A Wagner, and Alfred M Bruckstein. On swarm optimality in dynamic and symmetric environments. *economics*, 7:11, 2008.
- [17] J. M. Anthonisse. The rush in a directed graph. Technical Report BN 9/71, Stichting Mathematisch Centrum, Amsterdam, 1971.
- [18] E. Bakshy, J.M. Hofman, W.A. Mason, and D.J. Watts. Everyone’s an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 65–74. ACM, 2011.
- [19] Albert-Laszlo Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [20] M. Barthélemy. Betweenness centrality in large complex networks. *The European Physical Journal B – Condensed Matter*, 38(2):163–168, March 2004.
- [21] Frank M Bass, Norris Bruce, Sumit Majumdar, and BPS Murthi. Wearout effects of different advertising themes: A dynamic bayesian model of the advertising-sales relationship. *Marketing Science*, 26(2):179–195, 2007.
- [22] Frank M Bass, Anand Krishnamoorthy, Ashutosh Prasad, and Suresh P Sethi. Generic and brand advertising strategies in a dynamic duopoly. *Marketing Science*, 24(4):556–568, 2005.
- [23] Shlomo Bekhor, Yehoshua Cohen, and Charles Solomon. Evaluating long-distance travel patterns in israel by tracking cellular phone positions. *Journal of Advanced Transportation*, pages n/a–n/a, 2011.
- [24] Cesar Beltran-Royo, H Zhang, LA Blanco, and J Almagro. Multistage multiproduct advertising budgeting. *European Journal of Operational Research*, 2012.
- [25] Confessore G. Gentili M. Bianco, L. Combinatorial aspects of the sensor location problem. *Annals of Operation Research*, 144(1):201–234, 2006.
- [26] Lucio Bianco, Giuseppe Confessore, and Monica Gentili. Combinatorial aspects of the sensor location problem. *Annals of Operations Research*, 144(1):201–234, 2006.
- [27] P. Bork, L. J. Jensen, C. von Mering, A. K. Ramani, I. Lee, and E. M. Marcotte. Protein interaction networks from yeast to human. *Curr. Opin. Struct. Biol.*, 14(3):292–299, 2004.
- [28] U. Brandes. On variants of shortest-path betweenness centrality and their generic computation. *Social Networks*, 30(2):136–145, 2008.
- [29] Simon Broadbent. Point of view-what is a small advertising elasticity. *Journal of Advertising Research*, 29(4):37–39, 1989.
- [30] Z. Butler, A. Rizzi, and R. Hollis. Distributed coverage of rectilinear environments. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, 2001.
- [31] Zack J Butler. *Distributed coverage of rectilinear environments*. PhD thesis, Carnegie Mellon University, 2000.
- [32] D. Centola. The spread of behavior in an online social network experiment. *science*, 329(5996):1194, 2010.
- [33] D. Centola and M. Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702, 2007.
- [34] N. Cesa-Bianchi and G. Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51:239–261, 2003.
- [35] M. Cha, H. Haddadi, F. Benevenuto, and K.P. Gummadi. Measuring user influence in twitter: The million follower fallacy. In *4th International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2010.
- [36] H. Choi, S.H. Kim, and J. Lee. Role of network structure and network effects in diffusion of innovations. *Industrial Marketing Management*, 39(1):170–177, 2010.
- [37] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
- [38] M. C. Cohen and P. Harsha. Designing price incentives in a network with social interactions. *Submitted*, 2013.
- [39] C Samuel Craig and Avijit Ghosh. Using household-level viewing data to maximize effective reach. *Journal of Advertising Research*, 1993.
- [40] Peter J Danaher and Roland T Rust. Determining the optimal level of media spending. *Journal of Advertising Research*, 34(1):28–34, 1994.
- [41] Peter J Danaher and Roland T Rust. Determining the optimal return on investment for an advertising campaign. *European Journal of Operational Research*, 95(3):511–521, 1996.
- [42] PJ Danaher. Parameter estimation and applications for a generalisation of the beta-binomial distribution. *Australian Journal of Statistics*, 30(3):263–275, 1988.
- [43] Kenneth R Deal. Optimizing advertising expenditures in a dynamic duopoly. *Operations Research*, 27(4):682–692, 1979.
- [44] Norman K Dhalla. Assessing the long term value of advertising. *Harvard Business Review*, 56(1):87–95, 1978.
- [45] P.S. Dodds, R. Muhamad, and D.J. Watts. An experimental study of search in global social networks. *Science*, 301(5634):827, 2003.
- [46] P.S. Dodds and D.J. Watts. Universal behavior in a generalized model of contagion. *Physical Review Letters*, 92(21):218701, 2004.
- [47] S. Dolev, Y. Elovici, R. Puzis, and P. Zilberman. Incremental deployment of network monitors based on group betweenness centrality. *Inf. Proc. Letters*, 109:1172–1176, 2009.
- [48] Peter Doyle and John Saunders. Multiproduct advertising budgeting. *Marketing Science*, 9(2):97–113, 1990.
- [49] Peter F Drucker. The new productivity challenge. *Quality in Higher Education*, 37, 1995.
- [50] N. Eagle, A. Pentland, and D. Lazer. Inferring social network structure using mobile phone data. *Proceedings of the National Academy of Sciences (PNAS)*, 106:15274–15278, 2009.
- [51] Nathan Eagle, Michael Macy, and Rob Claxton. Network diversity and economic development. *Science*, 328(5981):1029–1031, 2010.
- [52] Joseph O Eastlack and Ambar G Rao. Modeling response to advertising and pricing changes for v-8 cocktail vegetable juice. *Marketing Science*, 5(3):245–259, 1986.
- [53] G.M. Erickson, P.J. Currie, B.D. Inouye, and A.A. Winn. Tyrannosaur life tables: An example of nonavian dinosaur population biology. *Science*, 313(5784):213–217, 2006.
- [54] M. G. Everett and S. P. Borgatti. The centrality of groups and classes. *Mathematical Sociology*, 23(3):181–201, 1999.
- [55] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. *SIGCOMM Comput. Comm. Rev.*, 29(4):251–262, 1999.
- [56] McKinsey Chief Marketing & Sales Officer Forum. *Big Data, Analytics, and the Future of Marketing & Sales*. 2013.

- [57] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.
- [58] D. Friedman, A. Steed, and M. Slater. Spatial social behavior in second life. In *Proc. Intelligent Virtual Agents LNAI 4722*, pages 252–263, 2007.
- [59] James W Friedman. Advertising and oligopolistic equilibrium. *The Bell Journal of Economics*, pages 464–473, 1983.
- [60] Benjamin Gompertz. On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. *Philosophical Transactions of the Royal Society of London*, 115:513–583, 1825.
- [61] Marta C. Gonzalez, Cesar A. Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 06 2008.
- [62] Dominique M Hanssens, Leonard J Parsons, and Randall L Schultz. *Market response models: Econometric and time series analysis*, volume 12. Springer, 2003.
- [63] C.P. Herrero. Ising model in scale-free networks: A monte carlo simulation. *Physical Review E*, 69(6):067109, 2004.
- [64] P. Holme. Congestion and centrality in traffic flow on complex networks. *Advances in Complex Systems*, 6(2):163–176, 2003.
- [65] C.L. Hsu and H.P. Lu. Why do people play on-line games? an extended tam with social influences and flow experience. *Information and Management*, 41:853–868, 2004.
- [66] C.L. Hsu and H.P. Lu. Consumer behavior in online game communities: A motivational factor perspective. *Computers in Human Behavior*, 23:1642–1659, 2007.
- [67] B.A. Huberman, D.M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1):8, 2009.
- [68] John Philip Jones. *How Much is Enough?: Getting the Most from Your Advertising Dollar*. Lexington Books, 1992.
- [69] Robert S Kaplan and Allan D Shocker. Discount effects on media plans. *Journal of Advertising Research*, 11(3):37–43, 1971.
- [70] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.
- [71] R. E. Korf and W. Zhang. Performance of linear-space search algorithms. *Artificial Intelligence*, 79(2):241–292, 1995.
- [72] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM, 2010.
- [73] Luca Lambertini. Advertising in a dynamic spatial monopoly. *European journal of operational research*, 166(2):547–556, 2005.
- [74] Renaud Lambiotte, Vincent D. Blondel, Cristobald de Kerchove, Etienne Huens, Christophe Prieur, Zbigniew Smoreda, and Paul Van Dooren. Geographical dispersal of mobile communication networks. *Physica A: Statistical Mechanics and its Applications*, 387(21):5317 – 5325, 2008.
- [75] David Lazer, Alex Pentland, Lada Adamic, Sinan Aral, Albert-Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, Tony Jebara, Gary King, Michael Macy, Deb Roy, and Marshall Van Alstyne. Social science: Computational social science. *Science*, 323(5915):721–723, 2009.
- [76] John D Leckenby and Shizue Kishi. How media directors view reach/frequency estimation. *Director*, 9:9–9, 1994.
- [77] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. Citeseer, 2009.
- [78] Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman. The dynamics of viral marketing. *ACM Trans. Web*, 1, May 2007.
- [79] Ouyang Y. Li, X. Reliable sensor deployment for network traffic surveillance. *Transportation Research Part B*, 45:218–231, 2011.
- [80] Xiaopeng Li and Yanfeng Ouyang. Reliable sensor deployment for network traffic surveillance. *Transportation research part B: methodological*, 45(1):218–231, 2011.
- [81] John DC Little. Aggregate advertising models: The state of the art. *Operations Research*, 27(4):629–667, 1979.
- [82] John DC Little. Models and managers: The concept of a decision calculus. *Management science*, 50(12 supplement):1841–1853, 2004.
- [83] John DC Little and Leonard M Lodish. A media planning calculus. *Operations Research*, 17(1):1–35, 1969.
- [84] Hani I Mesak. On the generalizability of advertising pulsation monopoly results to an oligopoly. *European Journal of Operational Research*, 117(3):429–449, 1999.
- [85] George E Monahan. Optimal advertising with stochastic demand. *Management Science*, 29(1):106–117, 1983.
- [86] Michael J Naples. *Effective frequency: the relationship between frequency and advertising effectiveness*. Association of National Advertisers New York, 1979.
- [87] Marc Nerlove and Kenneth J Arrow. Optimal advertising policy under dynamic conditions. *Economica*, pages 129–142, 1962.
- [88] M.E.J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- [89] Dung Nguyen. An analysis of optimal advertising under uncertainty. *Management Science*, 31(5):622–633, 1985.
- [90] V. Nicosia, F. Bagnoli, and V. Latora. Impact of network structure on a model of diffusion and competitive interaction. *EPL (Europhysics Letters)*, 94:68009, 2011.
- [91] Jukka-Pekka Onnela and Felix Reed-Tsochas. Spontaneous emergence of social influence in online systems. *Proc. Natl. Academy of Sciences*, 107(43), 2010.
- [92] J. Orwant. Heterogeneous learning in the doppelganger user modeling system. *User Modeling and User-Adapted Interaction*, 4:107–130, 1994.
- [93] Wei Pan, Nadav Aharony, and Alex Pentland. Composite social network for predicting mobile apps installation. In *Proceedings of the 25th Conference on Artificial Intelligence (AAAI)*, pages 821 – 827, 2011.
- [94] G. Pickard, W. Pan, I. Rahwan, M. Cebrian, R. Crane, A. Madan, and A. Pentland. Time-critical social mobilization. *Science*, 334(6055):509–512, 2011.
- [95] R. Puzis, M. D. Klippel, Y. Elovici, and S. Dolev. Optimization of nids placement for protection of intercommunicating critical infrastructures. In *EuroISI*, 2007.
- [96] Rami Puzis, Yuval Elovici, and Shlomi Dolev. Finding the most prominent group in complex networks. *AI communications*, 20(4):287–296, 2007.
- [97] Vitorino Ramos, Carlos Fernandes, and Agostinho C Rosa. Societal implicit memory and his speed on tracking extrema in dynamic environments using self-regulatory swarms. *Journal of Systems Architecture, Farooq M. and Menezes R.(Eds.), special issue on Nature Inspired Applied Systems, Elsevier, Summer*, 2006.
- [98] Ambar G Rao and Peter Miller. *Advertising–sales Response Functions*. New York University, Graduate School of Business Administration, 1975.
- [99] Ram C Rao. Advertising decisions in oligopoly: An industry equilibrium analysis. *Optimal Control Applications and Methods*, 5(4):331–344, 1984.
- [100] Petri Rouvinen. Diffusion of digital mobile telephony: Are developing countries different? *Telecommunications Policy*, 30(1):46 – 63, 2006.
- [101] Roland T Rust. Selecting network television advertising schedules. *Journal of Business Research*, 13(6):483–494, 1985.
- [102] Roland T Rust. *Advertising media models: A practical guide*. Lexington Books Lexington, MA, 1986.
- [103] J. Scott. *Social Network Analysis: A Handbook*. Sage Publications, London, 2000.

- [104] Suresh P Sethi. Optimal control of the vidale-wolfe advertising model. *Operations Research*, 21(4):998–1013, 1973.
- [105] D. Shah and T. Zaman. Rumors in a network: Who's the culprit? *Arxiv preprint arXiv:0909.4370*, 2009.
- [106] Srinivasaraghavan Sriram and Manohar U Kalwani. Optimal advertising and promotion budgets in dynamic markets with brand equity as a mediating variable. *Management Science*, 53(1):46–60, 2007.
- [107] R. Stern, R. Puzis, and A. Felner. Potential search: a bounded-cost search algorithm. In *AAAI 21st International Conference on Automated Planning and Scheduling (ICAPS)*, 2011.
- [108] S. H. Strogatz. Exploring complex networks. *Nature*, 410:268–276, March 2001.
- [109] Jim Surmanek. *Media planning: a practical guide*, volume 49. NTC Business Books Chicago, IL, 1996.
- [110] Gerard J Tellis. The price elasticity of selective demand: A meta-analysis of econometric models of sales. *Journal of Marketing Research*, pages 331–341, 1988.
- [111] Gerard J Tellis. Interpreting advertising and price elasticities. *Journal of Advertising Research*, 29(4):40–43, 1989.
- [112] ML Vidale and HB Wolfe. An operations-research study of sales response to advertising. *Operations Research*, 5(3):370–381, 1957.
- [113] Yogesh Virkar and Aaron Clauset. Power-law distributions in binned empirical data. *arXiv preprint arXiv:1208.3524*, 2012.
- [114] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*. Cambridge, England: Cambridge University Press., 1994.
- [115] D.J. Watts, J. Peretti, and Harvard Business School. Viral marketing for the real world. 2007.
- [116] D.J. Watts and S.H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.
- [117] D. R. White and S. P. Borgatti. Betweenness centrality measures for directed graphs. *Social Networks*, 16:335–346, 1994.
- [118] Charles Solomon Leonid Kheifits Yehuda J. Gur, Shlomo Bekhor. Intercity person trip tables for nationwide transportation planning in israel obtained from massive cell phone data. *Transportation Research Record: Journal of the Transportation Research Board*, 2121:145–151, 2009.
- [119] S.H. Yook, H. Jeong, and A.-L. Barabasi. Modeling the internet's large-scale topology. *Proceedings of the National Academy of Science*, 99(21):13382–13386, Oct. 2002.
- [120] Shlomo Zilberstein. Using anytime algorithms in intelligent systems. *AI Magazine*, 17(3):73–83, 1996.



Yaniv Altshuler is the CTO and Chief Scientist of Athena Wisdom, and a researcher at MIT (Media Lab). He received his BA (with highest honors) in Computer Science at the Israeli Institute of Technology, and his MSc and PhD (with honors) in Computer Science at the Israeli Institute of Technology. Altshuler spent 3 years as a post-doc at MIT at the Human Dynamics group, headed by Prof. Alex Sandy Pentland. Yaniv specializes in big data, social physics, and network analysis. Altshuler has published over

60 academic papers and filed 15 patent applications.



Erez Shmueli is a senior lecturer at the department of Industrial Engineering at Tel-Aviv University and a research affiliate at the MIT Media Lab. He received his BA degree (with highest honors) in Computer Science from the Open University of Israel, and MSc and PhD degrees in Information Systems Engineering from Ben-Gurion University of the Negev, Israel, under the supervision of Prof. Yuval Elovici. After completing his PhD, Erez spent two years as a post-doctoral associate at the MIT Media Lab, at the Human Dynamics group headed by Prof. Alex Sandy Pentland. His main research interests include Big Data, Complex Networks, Computational Social Science and Information Security and Privacy. His professional experience includes five years as a programmer and a team leader in the Israeli Air-Force and three years as a project manager in Deutsche Telekom Laboratories at Ben-Gurion University of the Negev.



Guy Zyskind is a graduate student in the Human Dynamics group at the MIT Media Lab. He holds a B.Sc degree in Electrical Engineering and Computer Science from Tel Aviv University and is currently pursuing a M.S degree under the supervision of Prof. Alex "Sandy" Pentland. Before joining the Media Lab, Guy has led the development of several start-ups in the Big Data and consumer spaces. His research interests intersect the study of Social Networks and Big Data, Privacy and Distributed Systems.



Oren Lederman is a graduate student at the MIT Media Lab, Human Dynamics group. He holds a B.Sc. in Computer Science and Economics from Tel Aviv University, Israel. His professional experience includes five years as developer and team leader in the Israeli army, founding a mobile social network startup, four years as data infrastructures team leader in an influencers-based marketing startup, and a research engineer position at the Singapore-MIT Alliance for Research and Technology. His research interests include

big data and personal data, group dynamics and studying innovation teams.



Nuria Oliver (PhD, MIT 2000) is currently the Scientific Director and founder of the User, Data and Media Intelligence research areas in Telefonica Research, working on data analytics, machine learning, user modeling and HCI in a variety of domains. Prior to this position, she was a researcher at Microsoft Research in Redmond, WA for over 7 years. She has written over 90 scientific papers in international conferences, journals and book chapters. Her work has been widely recognized by the scientific community

with over 7700 citations. Nuria has over 40 patent applications and granted patents. She is in the organizing and/or program committee of the top conferences in her research areas.

She believes in the power of technology to empower and increase the quality of life of people. She has received a number of awards, including a 10 Year Technical Impact Award (ACM ICMI), a Rising Start Award by the Women's Forum for the Economy and Society (2009), MITs TR100 Young Innovators Award (2004) and the First Spanish Award of EECS graduates (1994). She is senior member of the ACM.

Her work has been widely featured on multiple newspapers, magazines, radio and TV stations both in Spain and the US. She has been featured in EL PAIS Sunday magazine as one of a few 'female directors in technology' (2012), named Rising Talent by the Women's Forum for Economy & Society (October 2009), one of the 'most influential young women in Spain' (MujerHoy Magazine, 2012), one of '100 leaders of the future' by Capital Magazine (May 2009) and one of the 'Generation XXI: 40 Spanish youngsters that will make news in the Third Millenium' by EL PAIS (2000).

She has given two TEDx talks and one WIRED talk. She is also co-organizing the first TEDxBarcelona event devoted to Education.



Alex 'Sandy' Pentland directs MITs Human Dynamics Laboratory and the MIT Media Lab Entrepreneurship Program, co-leads the World Economic Forum Big Data and Personal Data initiatives, and is a Board member for Nissan, Motorola Mobility, Telefonica, and Harvard Business Review. He has previously helped create and direct MITs Media Laboratory, the Media Lab Asia laboratories at the Indian Institutes of Technology, and Strong Hospitals Center for Future Health. In 2012 Forbes named Sandy one of the

'seven most powerful data scientists in the world', along with Google founders and the CTO of the United States, and in 2013 he won the McKinsey Award from Harvard Business Review. He is among the most-cited computational scientists in the world, and a pioneer in computational social science, organizational engineering, wearable computing (Google Glass), image understanding, and modern biometrics. His most recent book is 'Social Physics,' published by Penguin Press.