

Surfing the Blogosphere: Optimal Personalized Strategies for Searching the Web

Stratis Ioannidis
Thomson
Paris, France
stratis.ioannidis@thomson.net

Laurent Massoulié
Thomson
Paris, France
laurent.massoulie@thomson.net

Abstract—We propose a distributed mechanism for finding websurfing strategies that is inspired by the StumbleUpon recommendation engine. Each day, a websurfer visits a sequence of websites recommended by our mechanism, and selects one that matches her daily interests. We formally show that even with this minimal feedback from the surfer—the selected website—our mechanism finds a websurfing strategy that matches the surfer’s interests optimally. The surfer does not need to know—or declare—what her daily interests are before she is presented with content she likes. Moreover, our mechanism is content-agnostic: it is oblivious to the nature of the content the surfer selects.

In addition, we study how the performance of this mechanism can be improved if surfers with similar interests share their feedback. Such surfers can be found indirectly, *e.g.*, if they are all registered as friends in a social networking application. Our analysis characterizes the improvement in the mechanism’s accuracy, based on the size of the group and the degree of similarity between the surfers’ interests. In particular, we show that sharing feedback can significantly accelerate the convergence of our mechanism. Our results are derived analytically using stochastic approximation techniques, but are also validated through a numerical study.

I. INTRODUCTION

An increasing number of people use the Internet, as opposed to more traditional mass media, to learn about daily events. However, the sheer volume, diversity and varying quality of content on the web—ranging from personal blogs to professional news-sites—makes searching for current, interesting information a challenging task. As a result, online recommendation engines (such as, *e.g.*, Delicious¹, Digg², and Reddit³) giving daily suggestions to websurfers have recently proliferated.

The fundamental question the above recommendation engines are trying to answer, which also motivates the work presented in this paper, is the following: *on a given day, which websites should a websurfer visit in order to find content that she is interested in?* A challenge posed by the above problem is that both the interests of websurfers and the content published by websites may change from one day to the next. Moreover, eliciting a websurfer’s interests is not always straightforward: the websurfer may not be able to declare what type of content she is interested in before viewing such content. For this



Fig. 1. The StumbleUpon toolbar. The websurfer is presented with a new website whenever she clicks the “Stumble!” button. If she finds the website shown interesting, she can register her approval by clicking “I like this”.

reason, we would like to address the above question without requiring a-priori knowledge of the websurfer’s interests on a given day.

To address this challenge, we adopt the approach employed by the recommendation engine StumbleUpon⁴. StumbleUpon implements a toolbar, giving websurfers the option between two buttons (see Fig. 1). The first, labeled “Stumble!”, generates a randomly selected website. The second, labeled “I like this” allows surfers to indicate they like the website they see. Websurfers are thus presented with a sequence of different websites until one that matches their interests is found.

The exact algorithm used by StumbleUpon to recommend websites to its users is proprietary. In broad terms [1], StumbleUpon’s server learns through a registered user’s feedback which websites she likes; recommendations are then made by showing the user websites that have been recommended by other registered users with similar preferences. StumbleUpon also leverages *social-networking information* to make recommendations. In particular, registered users can declare their friends. Assuming that friends should have common interests, StumbleUpon exploits this information to recommend websites to users when selected by their friends. Through both of these techniques, StumbleUpon personalizes the search for content to a particular user. Moreover, it does so with minimal feedback; users need only declare which sites they like and, optionally, provide social networking information.

In this paper, we propose a distributed mechanism for determining which websites a surfer should visit that is inspired by StumbleUpon. In particular, our mechanism presents a sequence of websites to the surfer, and all a surfer needs to declare is whether she approves the websites shown to

¹<http://delicious.com>

²<http://digg.com>

³<http://reddit.com>

⁴<http://www.stumbleupon.com>

her. Like StumbleUpon, the websurfer need not know—nor explicitly declare—what her interests are before actually being presented with a website she likes. Our mechanism is *content-agnostic*: it is oblivious to the actual nature of the content shown to the websurfer by the websites she visits. Finally, contrary to StumbleUpon, our mechanism is distributed: no centralized data collection or processing is required.

Our first contribution is to show that, in spite of its simplicity, our mechanism computes a *surfing strategy*, *i.e.*, a strategy for visiting websites, that matches the websurfer’s interests to the content presented by websites in an optimal way. More specifically, the surfing strategy produced by our mechanism maximizes a (non-increasing) function the number of websites the websurfer views before finding interesting content.

Our second contribution is to show that the performance of this mechanism can be significantly improved if feedback is aggregated among surfer communities. In particular, if surfers with similar interests share their feedback, the accuracy of the mechanism and the rate of its convergence can be significantly improved. To the best of our knowledge, this is the first time that the performance of such a mechanism is formally shown to improve through the use of community aggregated feedback.

The remainder of this paper is organized as follows. In Section II we present related work in this area. In Section III we describe our proposed mechanism in detail. In Section IV we state, without proofs, our main results regarding the performance of our mechanism; the full analysis can be found In Section V. Section VI provides a method for reducing the dimension of our problem. Finally, we validate our results through a numerical study in Section VII, and conclude in Section VIII.

II. RELATED WORK

An altogether different approach to searching for content is adopted by search engines like, *e.g.*, Google or Blogscope [2]. Contrary to our scheme, search engines are centralized and content-aware: they require considerable storage and computational resources to store and process the web’s contents. In addition, they respond to query submissions and, as such, cannot provide recommendations when a websurfer is not a priori aware of what content interests her.

Surfers using the Delicious recommendation engine tag blogs with keywords; tagged content then appears on the Delicious website as an aid to daily websurfing. This approach relies on the participation of websurfers through tag submission, which they may not always be able or willing to perform; Digg and Reddit avoid this by letting surfers vote for websites they visit and displaying recent content with high voting scores. Such approaches are centralized but also lack *personalization*: in contrast, our scheme matches surfers’ interests individually.

El-Arini et al. [3] propose a mechanism for recommending a list of websites to surfers that maximizes a personalized coverage function, indicating whether the suggested websites cover topics of interest to the surfer. A similar minimal

feedback scheme as ours is employed, based on the approval and disapproval of suggested sites. However, the mechanism in [3] is not content-agnostic. It requires prior knowledge of the probability that a topic is covered by a proposed website, which, as in Google and Blogscope [2], is obtained by periodically crawling the web and processing the collected data in a centralized manner. Our work also differs by considering aggregation of feedback over communities.

Our mechanism was strong ties to collaborative filtering (see, *e.g.*, [4], [5]), which typically involves clustering users and making recommendations based on the suggestions of other cluster members. Our work is orthogonal: we are not concerned with whether communities are obtained through clustering or, *e.g.*, a social networking application. Irrespective of how the community was formed, we formally characterize the performance gains of aggregating feedback over the community, in terms of its size and the degree of similarity between its members.

III. SYSTEM DESCRIPTION

A. Website Content and Topic Coverage

We model the web as a set of N websites. These websites maintain content that covers M different topics, such as, *e.g.*, sports, politics, movies, health, *etc.* Periodically, the content displayed at each website changes: we denote by $p_{w,f}$ the probability that the content presented in site w covers topic f , where $w = 1, \dots, N$ and $f = 1, \dots, M$. For each website w , we call the vector $\vec{p}_w = [p_{w,1}, \dots, p_{w,M}]$ as the *publishing strategy* of the website. In general, more than one topic (*e.g.*, sports *and* health) can be covered by a certain website at a given point in time; the expected number of topics covered by w will be $\sum_f p_{w,f}$.

We assume that for every topic f there exists a website w such that $p_{w,f} > 0$, *i.e.*, w covers f with positive probability. We also make the assumption that websites have high churn and their content changes frequently (as, *e.g.*, in news-sites like `nytimes.com` or popular blogs like `huffingtonpost.com`). As a result, we will assume that each time a website w is visited its current content is independent from content viewed during previous visits.

B. Surfer Interests and a Minimal Feedback Mechanism

We assume that each day a given websurfer is interested in topic f with some probability d_f , where $\sum_f d_f = 1$. We call $\vec{d} = [d_1, \dots, d_M]$ the *interest profile* of the surfer. The surfer visits websites daily by using a mechanism that works as follows: the mechanism recommends a sequence of websites to the surfer, and the surfer keeps viewing these sites until she finds a website covering the topic that interests her.

Like StumbleUpon, this mechanism can be implemented as a toolbar on the surfer’s web-browser: the surfer would be presented with two different buttons, one called “next” and the other called “approve”. Clicking the first would generate a new website, while clicking the second would indicate that a topic the surfer is interested in is found. However, unlike Stumbleupon, our mechanism is fully decentralized:

the toolbar relies only on the surfer's feedback and there is no need for centralized data collection.

At each click of the “next” button, the website recommended is chosen by the mechanism according to a probability distribution over all sites: we denote by x_w the probability that the mechanism recommends website w , where $\sum_w x_w = 1$. We refer to the vector $\vec{x} = [x_1, \dots, x_N]$ as the *surfing strategy* of the mechanism. As we will see, the mechanism will use the surfer's feedback to update this strategy from one day to the next, in a manner we discuss in detail below.

C. Surfer Optimization

Let Y be the number of sites the surfer visits until it locates a topic it is interested in. Moreover, let $R(Y)$ be a function rating the performance of the system, given that the topic is found within Y steps. We treat R as a system parameter internal to the mechanism, and make the following assumption:

Assumption 1. *The rating function $R : \mathbb{N} \rightarrow \mathbb{R}$ is (a) non-increasing and (b) summable, i.e., $\sum_{k=1}^{\infty} |R(k)| < \infty$.*

Considering non-increasing rating functions is natural: the longer a surfer has to wait, the lower the system's rating should be. The assumption on the summability is technical; note that it holds if R has a finite support and, e.g., visiting more than 30 websites gives a zero rating to the system.

Denote by $\mathbb{E}_{\vec{x}}[R(Y)]$ the expected system rating when the surfing strategy is \vec{x} . Ideally, assuming that the surfer's interest profile \vec{d} and the website publishing strategies \vec{p}_w are fixed, the mechanism recommending sites to the surfer should use a surfing strategy \vec{x} that is a solution to the following problem:

$$\begin{aligned} & \text{SURFER OPTIMIZATION} \\ \text{Maximize: } & F(\vec{x}) = \mathbb{E}_{\vec{x}}[R(Y)], \end{aligned} \quad (1a)$$

$$\text{subject to: } \sum_w x_w = 1, \quad \text{and } x_w \geq \epsilon, \quad \forall w. \quad (1b)$$

for some $0 \leq \epsilon \leq \frac{1}{N}$. Intuitively, given that the rating of the system decreases with the number of visits the surfer makes, the mechanism should choose a surfing strategy that maximizes the expected rating by making the number of visits sufficiently small.

The optimization problem described by (1) is restricted to surfing strategies in which *every site is visited with probability at least ϵ* , i.e., the feasible domain is restricted to

$$D_\epsilon = \{\vec{x} : \sum_w x_w = 1 \text{ and } x_w \geq \epsilon\}. \quad (2)$$

Note that $D_\epsilon \neq \emptyset$ iff $\epsilon \leq \frac{1}{N}$. Moreover, if $\epsilon < \epsilon'$, then $D_{\epsilon'} \subseteq D_\epsilon \subseteq D_0$, where D_0 includes all possible surfing strategies.

The main reason to restrict the domain to D_ϵ , where $\epsilon > 0$, has to do with website discovery—we want to ensure the websurfer has a chance to visit a website she likes. It is also useful for “tracking” [6] changes in \vec{d} or \vec{p}_w : if a website is never visited, the mechanism will never have the opportunity to adapt to the fact that, e.g., it has started covering sports, and thereby may now meet the interests of the surfer. We note that,

if $\epsilon > 0$ is small, an optimal strategy in D_ϵ will not be much worse than the optimal strategy in D_0 (see also Lemma 3).

D. Updating the Surfing Strategy

The surfing strategy \vec{x} can change from one day to the next based only on the minimal feedback given by the surfer; we denote by $\vec{x}(k)$, $k = 1, 2, \dots$, the strategy on day k . Let

$$\Pi_{D_\epsilon}(\vec{x}) = \arg \min_{\vec{y} \in D_\epsilon} \|\vec{y} - \vec{x}\|_2, \quad \vec{x} \in \mathbb{R}^N. \quad (3)$$

be the orthogonal projection of \vec{x} to the domain D_ϵ . Since D_ϵ is closed and convex, Π_{D_ϵ} is well defined. Moreover, there are known algorithms (e.g., [7]) computing Π_{D_ϵ} in $O(N)$ time.

Based on the feedback of the surfer on the k -th day, the strategy $\vec{x}(k)$ is updated according to the following scheme:

$$\vec{x}(k+1) = \Pi_{D_\epsilon}(\vec{x}(k) + \gamma(k) \cdot \vec{g}(k)), \quad k = 1, 2, \dots \quad (4)$$

The term $\gamma(k) \in \mathbb{R}_+$ is a gain factor; if it decreases with k , feedback given later has smaller impact on the surfing strategy. The term $\vec{g}(k) \in \mathbb{R}^N$ is a vector forcing the change on the surfing strategy based the surfer's feedback. In particular, $\vec{g}(k)$ has the following coordinates:

$$g_w(k) = \begin{cases} -\frac{Y(k)\Delta R(Y(k)+1)}{x_w(k)}, & \text{if the surfer selects } w \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

where $Y(k)$, $k = 1, 2, \dots$ is the number of sites a surfer visits on the k -th day until it locates a topic she is interested in, and

$$\Delta R(k) = \begin{cases} R(1), & \text{if } k = 1, \\ R(k) - R(k-1), & \text{if } k > 1. \end{cases} \quad (6)$$

Note that, to compute \vec{g} , all the mechanism needs to know is Y , R and \vec{x} . The rating function R is a system parameter, while \vec{x} is maintained by the mechanism and used to recommend websites. To compute Y , the mechanism needs to keep track of how many websites it has shown to the surfer; the only feedback given by the surfer is the indication that her interests have been met. In particular, the mechanism does not need to know—and the surfer does not need to declare—which is the interesting topic covered by website w . In fact, the mechanism is oblivious to even what the topics $f = 1, \dots, M$ are.

E. Exploiting Surfer Communities

In our analysis, apart from considering the mechanism (4) when the surfer is isolated, we will also consider the case when the surfer is part of a surfer community $C = \{1, \dots, |C|\}$. Each surfer $s \in C$ has its own interest profile \vec{d}_s . Each day the topic it is interested in is chosen independently of the topics that interest other surfers. Furthermore, each surfer executes the minimal feedback mechanism outlined in the previous section. The difference from the single-surfer case will be that these mechanisms will exploit the existence of the surfer community in the manner we describe below.

To begin with, all mechanisms implement a common surfing strategy \vec{x} everywhere in the community. Moreover, at the end of each day, the feedback collected from each surfer is communicated among the mechanisms and is aggregated;

the aggregated feedback is then used to update the common strategy.

More specifically, let $Y_s(k)$, $k = 1, 2, \dots$ be the number of sites surfer $s \in C$ visits on the k -th day and $\vec{g}_s(k)$ be the change vector induced by the feedback of surfer s on day k , given by (5). The latter quantities are computed individually by each mechanism and, at the end of each day, are communicated among all surfers in C .

As in (1a), we denote by $F_s(\vec{x}) = \mathbb{E}_{\vec{x}}[R(Y_s)]$ the expected rating of the mechanism at surfer s . Moreover, we denote by

$$\vec{g}_C = \frac{1}{|C|} \sum_{s=1}^{|C|} \vec{g}_s \quad \text{and} \quad F_C(\vec{x}) = \frac{1}{|C|} \sum_{s=1}^{|C|} F_s(\vec{x}) \quad (7)$$

the average change vector and the average system rating among all surfers in the community.

The common surfing strategy is then updated as follows:

$$\vec{x}(k+1) = \Pi_{D_\epsilon}(\vec{x}(k) + \gamma(k) \cdot \vec{g}_C(k)), \quad k = 1, 2, \dots \quad (8)$$

In other words, the common surfing strategy is updated as in (4), the only difference being that, instead of the individual vectors \vec{g}_s , the community average \vec{g}_C is used.

We will denote by δ_C the maximum l_1 distance between the interest profiles of surfers in the community, *i.e.*,

$$\delta_C = \max_{s, s' \in C} \left\| \vec{d}_s - \vec{d}_{s'} \right\|_1 = \max_{s, s' \in C} \sum_w |d_{s,w} - d_{s',w}|. \quad (9)$$

We will call δ_C the *diameter* of the community. Intuitively, we will be interested in the case where the diameter is small, as surfers in the same community should have similar interests.

F. Discussion on our Assumptions

The case where a surfer is interested in any of k different topics each day can easily be reduced to the single-topic case, by redefining the set of topics to be the set of topic k -tuples. Interestingly, even if our model is extended this way, mechanism (4) would remain unaltered: this is because, as noted, it is entirely oblivious to the nature of the topics f .

The assumption that websites display independent content across visits would not apply to websites that are updated infrequently. In practice, our mechanism should avoid visiting such sites more than once during a search, *e.g.*, through blacklisting them if the surfer does not approve them.

Our mechanism maintains a vector \vec{x} whose dimension is N , the number of websites. Given the number of websites on the Internet, this is clearly impractical. Restricting \vec{x} to a subset of all websites is not satisfactory, as it would limit the surfer's choices. In Section VI, we provide a more appealing solution to this problem by defining surfing strategies over *groups* of websites.

IV. MAIN RESULTS

In this section we state our main results regarding the convergence of mechanism (4) to an optimal surfing strategy and the effect of aggregating community feedback through mechanism (8). We first state our results without proofs; the full analysis can be found in Section V.

A. Convergence to an Optimal Surfing Strategy

Our first main result is to show that the minimal feedback mechanism we propose indeed computes an optimal surfing strategy. This is established for two different cases: one where the gain factors $\gamma(k)$ converge to zero, and one where they remain constant.

To begin with, by appropriately choosing $\gamma(k)$, we can guarantee that the mechanism converges to an optimal strategy, with probability one.

Theorem 1 (Convergence w.p.1). *Let $\vec{x}(k)$ be updated by the mechanism (4), where $\epsilon > 0$, $\gamma(k) = 1/k$ and Assumption 1 holds. Then $\lim_{k \rightarrow \infty} F(\vec{x}(k)) = \sup_{\vec{x} \in D_\epsilon} F(\vec{x})$, w.p.1.*

In other words, if the surfing strategy is updated according to our minimal feedback mechanism with $\gamma(k) = \frac{1}{k}$, the surfing strategy will eventually become optimal, *i.e.*, it will be a solution to the SURFER OPTIMIZATION problem.

Using decreasing gains can make our mechanism non-robust: if either the interest profile \vec{d} of the surfer or the publishing strategies \vec{p}_w change, and $\gamma(k)$ has effectively become zero, our mechanism will not be able to adapt. To be able to “track” [6] such changes, it is interesting to keep $\gamma(k)$ bounded away from zero.

In such a case, convergence to an optimal strategy cannot be guaranteed. However, one expects that if γ is small enough (4) does not deviate too far such a strategy. The following theorem implies that this intuition is indeed correct⁵:

Theorem 2 (Convergence in probability). *Let $\vec{x}(k)$ be updated by the mechanism (4), where $\epsilon > 0$, $\gamma(k) = \gamma > 0$ and Assumption 1 holds. Consider the continuous-time interpolated process $\vec{\xi}(t) = \vec{x}(\lfloor t/\gamma \rfloor)$, $t \in \mathbb{R}_+$. Then,*

$$\lim_{\gamma \rightarrow 0} \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T F(\vec{\xi}(t)) dt = \sup_{\vec{x} \in D_\epsilon} F(\vec{x}), \quad \text{in probability.}$$

In other words, by taking γ sufficiently small, we can make the time-average rating to be arbitrarily close to the one achieved by an optimal surfing strategy.

B. Exploiting Surfer Communities

Having established that our minimal feedback mechanism indeed computes an optimal surfing strategy, we look at its behavior if surfers form a community and implement a common strategy, updated according to (8).

Our first result is to show that, if the surfers have similar interest profiles, the mechanism (8) converges to a strategy that is not too far from each surfer's optimal surfing strategy.

Theorem 3. *Let $\vec{x}(k)$ be updated by the mechanism (8). Then, Theorems 1 and 2 hold if F is replaced by the average rating F_C , given by (7). Moreover, under Assumption 1, if δ_C is the diameter of the community, then, for all $s \in C$,*

$$\left| \sup_{\vec{x} \in D_\epsilon} F_C(\vec{x}) - \sup_{\vec{x} \in D_\epsilon} F_s(\vec{x}) \right| = O(\delta_C). \quad (10)$$

⁵Recall that a sequence of random variables X_n converges to X in probability when $n \rightarrow \infty$, for all $\delta > 0$, $\lim_{n \rightarrow \infty} \mathbf{P}(|X_n - X| > \delta) = 0$.

The theorem states that mechanism (8) maximizes the average rating among all surfers in the community, as opposed to each F_s individually. Most importantly, if surfers have similar interests (*i.e.*, the community diameter is small), the rating of the optimal strategy of an individual surfer will not be too far from the one achieved by algorithm (8).

On the other hand, receiving feedback from the entire community can give (8) a significantly improved performance over (4), as illustrated by the following theorem. Due to space constraints, we omit its derivation—the full proof can be found in our technical report [8].

Theorem 4. *Assume that the objective F_s of some $s \in C$ and the community average F_C have unique maxima at \vec{x}_s^* , \vec{x}_C^* , respectively, and that both are located in the relative interior of D_ϵ . Moreover, assume that $k^2|R(k)|$ is summable, and that the Hessian of F_s at \vec{x}_s^* is invertible.*

If surfer s runs mechanism (4) with $\epsilon > 0$ and $\gamma(k) = \gamma$, and $\vec{\xi}(t)$ is the continuous-time interpolated process $\vec{\xi}(t) = \vec{x}(\lfloor t/\gamma \rfloor)$, $t \in \mathbb{R}_+$, then

$$\lim_{\gamma \rightarrow \infty} \lim_{T \rightarrow \infty} \mathbb{E}[(\|\vec{\xi}(T) - \vec{x}_s^*\|_2)^2/\gamma] = c, \quad (11)$$

where c some positive constant. If, on the other hand, s runs mechanism (8) with the same $\epsilon > 0$ and $\gamma(k) = \gamma$, then, for small enough δ_C ,

$$\lim_{\gamma \rightarrow \infty} \lim_{T \rightarrow \infty} \mathbb{E}[(\|\vec{\xi}(T) - \vec{x}_C^*\|_2)^2/\gamma] = \frac{c}{|C|} + O(\sqrt{\delta_C}), \quad (12)$$

where $|C|$ the size of the community and δ_C the community diameter.

Recall from Theorem 2 that, if $\gamma(k) = \gamma$, (4) does not converge to a maximizer of F_s ; Eq. (11) then states that, in steady state, its expected distance from the maximizer \vec{x}_s^* will be of the order of $\sqrt{c\gamma}$, where c a constant. If, on the other hand, the surfing strategy is updated according to (8), Eq. (12) states that its expected distance from \vec{x}_C^* will be of the order of $\sqrt{\frac{c\gamma}{|C|} + O(\sqrt{\delta_C})}$. Hence, if δ_C is small enough, *sharing feedback with the community decreases the expected distance from the maximizer by a factor of $\sqrt{|C|}$* . Keeping in mind that, from Theorem 3, the $F_s(\vec{x}_s^*)$ is not too far from $F(\vec{x}_C^*)$ (see also Corollary 1) Theorem 4 establishes that, for small enough community diameters, exchanging feedback among surfers in the community can have a significant improvement on the accuracy of the mechanism.

An alternative interpretation of Theorem 4 can be given in terms of speed of convergence. As the distance of (4) from \vec{x}_s^* is of the order of $\sqrt{c\gamma}$, an obvious method for increasing the accuracy is by decreasing the gain factor γ . However, the smaller γ is, the smaller the “jumps” (4) makes and the longer it will take to get close to the maximizer. Theorem 4 therefore implies that, if δ_C is small, (8) can achieve the same accuracy as (4) but much faster, by using a greater γ .

Finally, Theorem 4 suggests an inherent trade-off appearing when the size of the community increases: the larger $|C|$ is, the smaller the first term of eq. (12) will be, indicating that

obtaining feedback from many users will reduce the “noise” in our mechanism. On the other hand, increasing the community can also increase the community diameter δ_C , thus increasing the second term in (12). In this sense, depending on how δ_C relates to $|C|$, a maximal community size might exist, beyond which adding more surfers will not improve performance. We further investigate this in our numerical study in Section VII.

V. ANALYSIS

A. Convexity, Differentiability, and Near-Optimality

Our analysis will exploit the fact that SURFER OPTIMIZATION, described by (1), is a convex optimization problem. In particular, the expected rating is a concave function of the surfing strategy, as indicated by the following lemma.

Lemma 1. *Under Assumption 1, for all $\epsilon \geq 0$, the objective function $F : D_\epsilon \rightarrow \mathbb{R}_+$, given by (1a), is concave.*

Proof: Recall that Y is the number of sites a surfer visits until it locates a topic it is interested in. Moreover, let Y_f the number of websites a surfer visits conditioned on the event that the topic it is interested in is f . Then, Y_f is a geometric random variable, *i.e.*,

$$\mathbf{P}(Y_f = \ell) = \rho_f(1 - \rho_f)^{\ell-1}, \quad \ell \geq 1, \quad (13)$$

where,

$$\rho_f = \sum_w p_{w,f} x_w \quad (14)$$

is the probability that topic f is found in one step. As a result, F can be written as

$$F(\vec{x}) = \sum_{f=1}^M d_f \mathbb{E}_{\vec{x}}[R(Y_f)] \quad (15)$$

where $\mathbb{E}_{\vec{x}}[R(Y_f)] = \sum_{k=1}^{\infty} R(k) \rho_f (1 - \rho_f)^{k-1}$. Observe that Assumption 1 immediately implies that

$$R \text{ is non-negative, bounded, and } \lim_{k \rightarrow \infty} R(k) = 0. \quad (16)$$

If $\mathbb{E}_{\vec{x}}[R(Y_f)]$ is a concave function of $\rho_f \in [0, 1]$ then it is also a concave function of \vec{x} , as a composition of a concave and a linear function. The lemma then follows, as $F(\vec{x})$ would be concave as the sum of concave functions, by (15). If $\rho_f > 0$,

$$\begin{aligned} \mathbb{E}_{\vec{x}}[R(Y_f)] &= \sum_{k=1}^{\infty} R(k) \rho_f (1 - \rho_f)^{k-1} \stackrel{(6)}{=} \sum_{k=1}^{\infty} \sum_{\ell=1}^k \Delta R(\ell) \rho_f (1 - \rho_f)^{k-1} \\ &= \sum_{\ell=1}^{\infty} \sum_{k=\ell}^{\infty} \Delta R(\ell) \rho_f (1 - \rho_f)^{k-1} = \sum_{\ell=1}^{\infty} \Delta R(\ell) (1 - \rho_f)^{\ell-1} \quad (17) \end{aligned}$$

Note that, by Assumption 1, $\Delta R(\ell) \leq 0$ for $\ell > 1$, while $\Delta R(1) = R(1) \geq 0$. As $(1 - \rho_f)^{\ell-1}$ is a convex, decreasing function of $\rho_f \in [0, 1]$ for all $\ell > 1$ and constant for $\ell = 1$, $\mathbb{E}_{\vec{x}}[R(Y_f)]$ is concave for $\rho_f \in (0, 1]$, as the sum of concave functions and the constant $R(1)$; it is also continuous in $(0, 1]$, by the monotone convergence theorem. By the same theorem, $\lim_{\rho_f \rightarrow 0} \sum_{k=1}^{\infty} \mathbb{E}_{\vec{x}}[R(Y_f)] = \sum_{k=1}^{\infty} \Delta R(k) = \lim_{\ell \rightarrow \infty} R(\ell) \stackrel{(16)}{=}$

0. Hence, $\mathbb{E}_{\vec{x}}[R(Y_f)]$ is continuous in all of $[0, 1]$. As it is concave in $(0, 1]$, continuity implies concavity on $[0, 1]$. ■

We can also compute the gradient of F .

Lemma 2. *Under Assumption 1, for all $\epsilon \geq 0$, the objective function $F : D_\epsilon \rightarrow \mathbb{R}_+$ is continuously differentiable and*

$$\frac{\partial F}{\partial x_w} = \sum_f d_f p_{w,f} \cdot \begin{cases} -\frac{1}{\rho_f} \mathbb{E}[Y_f \Delta R(Y_f + 1)], & \text{if } \rho_f > 0 \\ \sum_{k=1}^{\infty} R(k), & \text{if } \rho_f = 0, \end{cases} \quad (18)$$

where $\rho_f = \sum_f p_{w,f} x_w$.

Proof: By (14), $\frac{\partial \mathbb{E}_{\vec{x}}[R(Y_f)]}{\partial x_w} = \frac{\partial \mathbb{E}[R(Y_f)]}{\partial \rho_f} \cdot p_{w,f}$. For ρ_f in $(0, 1]$ we have by (17) that

$$\begin{aligned} \frac{\partial \mathbb{E}_{\vec{x}}[R(Y_f)]}{\partial \rho_f} &= - \sum_{\ell=1}^{\infty} (\ell-1) \Delta R(\ell) (1-\rho_f)^{\ell-2} \\ &\stackrel{k \equiv \ell-1}{=} - \sum_{k=1}^{\infty} k \Delta R(k+1) (1-\rho_f)^{k-1} = - \frac{\mathbb{E}[Y_f \Delta R(Y_f + 1)]}{\rho_f} \end{aligned} \quad (19)$$

This is continuous in $(0, 1]$ by the monotone convergence theorem. Continuity at zero follows by considering the appropriate limits and using again the monotone convergence theorem. ■

An immediate implication is that F is Lipschitz continuous:

Corollary 1. *Under Assumption 1, for all $\epsilon \geq 0$, the objective function $F : D_\epsilon \rightarrow \mathbb{R}_+$ is Lipschitz continuous. In particular,*

$$|F(\vec{x}) - F(\vec{y})| \leq L \|\vec{x} - \vec{y}\|_1, \quad \text{for all } \vec{x}, \vec{y} \in D_0,$$

where $L = \sum_{k=1}^{\infty} R(k) < \infty$. Lipschitz continuity in other norms follows by norm equivalence in \mathbb{R}^N .

The Lipschitz continuity of the objective function F can be used to understand the dependence of the optimal surfing strategy on ϵ . In particular, if ϵ is small enough, one does not lose too much by requiring that all sites are visited with probability at least ϵ . This is stated formally in the following lemma, whose proof can be found in our technical report [8].

Lemma 3. *Under Assumption 1, for any $0 \leq \epsilon \leq 1/N$, $|\sup_{\vec{x} \in D_\epsilon} F(\vec{x}) - \sup_{\vec{x} \in D_0} F(\vec{x})| = O(\epsilon)$.*

B. Estimating the Gradient

Having established the convexity and differentiability of F , we turn our attention to mechanism (4) and how it relates to the SURFER OPTIMIZATION problem. The following result states that the vector \vec{g} , given by (5), capturing the change of the surfing strategy based on the surfer's feedback, is in fact an unbiased estimator of the gradient ∇F . Hence, the mechanism (4) is a *stochastic approximation* algorithm [6], that solves SURFER OPTIMIZATION through a "randomized" gradient descent.

Lemma 4. *Under Assumption 1, given that the surfing strategy of the surfer is $\vec{x}(k)$, the vector $\vec{g} \in \mathbb{R}^N$, can be written as*

$$\vec{g}(k) = \nabla F(\vec{x}(k)) + V(\vec{x}(k)) \quad (20)$$

where $\mathbb{E}[V(\vec{x}(k))]$ is the zero vector and $\mathbb{E}[\|V(\vec{x}(k))\|_2^2] \leq \frac{N}{\epsilon} \sum_{k=1}^{\infty} (k \Delta R(k))^2 < \infty$, for $\epsilon > 0$.

Proof: Let E_f be the site at which the topic is found, given that the topic sought for is f . Then, $\mathbf{P}(E_f = w) = x_w p_{w,f} / \rho_f$, where ρ_f is given by (14). Moreover, E_f and Y_f are independent as

$$\begin{aligned} \mathbf{P}(Y_f = \ell \cap E_f = w) &\stackrel{(13),(14)}{=} (1-\rho_f)^{\ell-1} \cdot x_w p_{w,f} \\ &= \rho_f (1-\rho_f)^{\ell-1} \cdot x_w p_{w,f} / \rho_f = \mathbf{P}(Y_f = \ell) \cdot \mathbf{P}(E_f = w) \end{aligned} \quad (21)$$

By (5), we have that

$$\mathbb{E}[g_w] = - \sum_f d_f \frac{p_{w,f}}{\rho_f} \mathbb{E}[Y_f \Delta R(Y_f + 1) | E_f = w] \stackrel{(21),(18)}{=} \frac{\partial F}{\partial x_w}$$

By (21), $\mathbb{E}[g_w(\vec{x})^2] = - \sum_f \frac{d_f p_{w,f} \mathbb{E}[(Y_f \Delta R(Y_f + 1))^2]}{\rho_f x_w}$ where $\mathbb{E}[(Y_f \Delta R(Y_f + 1))^2] / \rho_f \leq C = \sum_{k=1}^{\infty} (k \Delta R(k))^2 < \infty$, as $\sum_{k=1}^{\infty} k \Delta R(k+1) = \sum_{k=1}^{\infty} R(k) < \infty$. As $p_{w,f} \leq 1$, $\sum d_f = 1$ and $x_w \geq \epsilon$ for all $\vec{x} \in D_\epsilon$, we have $\mathbb{E}[\|g_w(\vec{x})\|_2^2] \leq \frac{C}{\epsilon}$. The lemma thus follows as $\mathbb{E}[\|V(\vec{x})\|_2^2] = \sum_w \mathbb{E}[g_w(\vec{x})^2] - \sum_w \mathbb{E}[g_w(\vec{x})^2] \leq \sum_w \mathbb{E}[g_w(\vec{x})^2]$. ■

C. Proof of Theorem 1

Having established that SURFER OPTIMIZATION is a convex optimization problem, and that the change in the surfing strategy is an unbiased estimator of the gradient ∇F , we now prove Theorem 1. We essentially follow an "ODE method" [6], whereby the random evolution of the surfing strategy, captured by (4), is expressed in terms of the evolution of a deterministic solution to an ODE.

Recall that the feasible domain of (1) is D_ϵ , where $0 \leq \epsilon \leq \frac{1}{N}$, and let $\vec{x} \in D_\epsilon$. Given a vector \vec{y} , we define a mapping $Z : D_\epsilon \times \mathbb{R}^N \rightarrow \mathbb{R}^N$ as follows:

$$Z(\vec{x}, \vec{y}) = \lim_{\delta \rightarrow 0^+} \frac{\Pi_{D_\epsilon}(\vec{x} + \delta \vec{y}) - (\vec{x} + \delta \vec{y})}{\delta}. \quad (22)$$

Intuitively, for small δ , the quantity $\delta \cdot Z(\vec{x}, \vec{y})$ can be seen as the minimal force required to keep $\vec{x} + \delta \vec{y}$ in D_ϵ . The limit in the r.h.s. exists and, moreover, it can be expressed in terms of \vec{x} and \vec{y} in the following manner.

Lemma 5. *Assume that $\vec{x} \in D_\epsilon$ for some $0 \leq \epsilon < \frac{1}{N}$. Let $A(\vec{x}) = \{w : x_w = \epsilon\}$ be the set of active constraints at \vec{x} . Then, there exists a set $B \subseteq A(\vec{x})$, such that $Z(\vec{x}, \vec{y}) = \vec{z}$ where*

$$z_w = \begin{cases} -y_w, & \text{if } w \in B, \\ -\frac{1}{|B^c|} \sum_{w' \in B^c} y_{w'}, & \text{if } w \in B^c. \end{cases} \quad (23)$$

Moreover, for every $w \in B$, $y_w - \frac{1}{|B^c|} \sum_{w' \in B^c} y_{w'} < 0$.

Proof: Our proof will rely on the following property of the orthogonal projection Π_{D_ϵ} . Let $[N] = \{1, \dots, N\}$, and assume that $0 \leq \epsilon < \frac{1}{N}$. Then, for every $\vec{x} \in \mathbb{R}^N$, there exists a set $A \subsetneq [N]$ (that depends on \vec{x}) such that for every $w \in A$, $x_w - \frac{\sum_{w' \in A^c} x_{w'} - 1 + N\epsilon}{|A^c|} < 0$ and $\Pi_{D_\epsilon}(\vec{x}) = \vec{y}$ where

$$y_w = \begin{cases} \epsilon, & \text{if } w \in A, \\ x_w - \frac{\sum_{w' \in A^c} x_{w'} - 1 + N\epsilon}{|A^c|} + \epsilon, & \text{if } w \in A^c. \end{cases} \quad (24a)$$

This property follows from the simple geometry of D_ϵ . Its proof (and an $O(N)$ algorithm by Michelot [7] for computing the set A through repeated projections) is presented in [8].

An immediate implication of the above property of Π_{D_ϵ} is that there exists a set B such that, for all $w \in B$,

$$x_w + \delta y_w - \frac{\sum_{w' \in B^c} (x_{w'} + \delta y_{w'}) - 1 + N\epsilon}{|B^c|} < 0, \quad (25)$$

and, moreover, $\Pi_{D_\epsilon}(\vec{x} + \delta \vec{y}) = \vec{u}$, where

$$u_w = \begin{cases} \epsilon, & \text{if } w \in B, \\ x_w + \delta y_w - \frac{\sum_{w' \in B^c} (x_{w'} + \delta y_{w'}) - 1 + N\epsilon}{|B^c|} + \epsilon, & \text{o.w.} \end{cases} \quad (26)$$

Eq. (26) implies that B is a subset of $A(\vec{u}) = \{w : u_w = \epsilon\}$, the active constraints of \vec{u} . On the other hand, $A(\vec{u})$ is *upper semi-continuous* [6], i.e., there exists $\delta' > 0$ s.t. for any $\vec{u} \in D_\epsilon$ with $\|\vec{x} - \vec{u}\| \leq \delta'$, $A(\vec{u}) \subseteq A(\vec{x})$. The projection on D_ϵ is Lipschitz continuous (in particular, it is non-expansive), so by making δ small enough we can make \vec{u} to be arbitrarily close to \vec{x} . Upper semi-continuity therefore implies that, for δ small enough, $B \subseteq A(\vec{x})$. This, in turn, implies that

$$x_w = \epsilon, \text{ for all } w \in B, \quad \text{and} \quad \sum_{w' \in B^c} x_{w'} = 1 - |B|\epsilon. \quad (27)$$

The first statement of the lemma thus follows, as

$$\frac{1}{\delta} (\vec{u} - (\vec{x} + \delta \vec{y})) \stackrel{(26),(27)}{=} \begin{cases} -y_w, & w \in B \\ -\frac{1}{|B^c|} \sum_{w' \in B^c} y_{w'}, & w \in B^c. \end{cases}$$

In fact, we have proved that (23) holds not only in the limit, but for all small enough δ . The second statement follows from (25) and (27). \blacksquare

Using the above lemma, can define a deterministic process that converges to a solution of SURFER OPTIMIZATION:

Lemma 6. *Consider the following ODE*

$$\dot{\vec{x}} = \nabla F(\vec{x}) + Z(\vec{x}, \nabla F(\vec{x})) \quad (28)$$

and assume that the initial vector is in the feasible domain, i.e., $\vec{x}(0) \in D_\epsilon$. Then, under Assumption 1, $F(\vec{x}(t)) \geq F(\vec{x}(t'))$ for all $t \geq t'$, and $\lim_{t \rightarrow \infty} F(\vec{x}(t)) = \sup_{\vec{x} \in D_\epsilon} F(\vec{x})$.

Proof: To begin with, assuming that $\vec{x}(0) \in D_\epsilon$, $\vec{x}(t) \in D_\epsilon$ for all $t \geq 0$, by the definition of the force \vec{z} . Since, by Lemmas 1 and 2, F is convex, increasing and continuously differentiable, a vector \vec{x}^* will be a solution to the SURFER OPTIMIZATION problem, given by (1), if and only if there exist Lagrange multipliers λ^*, μ^* that satisfy the KKT conditions:

$$\sum_{\alpha} x_{\alpha}^* \leq 1, \quad \lambda^* (\sum_{\alpha} x_{\alpha}^* - 1) = 0, \quad \text{and}, \\ \frac{\partial F}{\partial x_a} \Big|_{\vec{x}^*} - \lambda^* + \mu_a^* = 0, \quad x_a^* \geq \epsilon, \quad \mu_a^* (x_a^* - \epsilon) = 0, \quad \forall a \in [N]$$

where $[N] = \{1, \dots, N\}$. The KKT conditions in turn imply that for an optimal solution \vec{x}^* either $x_a^* = \epsilon$ or $\frac{\partial F}{\partial x_a} \Big|_{\vec{x}^*} = \lambda^*$, where λ^* the Lagrange multiplier of the constraint $\sum_{\alpha} x_{\alpha}^* \leq 1$.

By Lemma 1, F is continuous, hence it attains a (finite) maximum in D_ϵ . Under (28),

$$\begin{aligned} \dot{F} &= \nabla F(\vec{x}) \cdot \dot{\vec{x}} = \nabla F(\vec{x}) \cdot (\nabla F(\vec{x}) + Z(\vec{x}, \nabla F(\vec{x}))) \\ &\stackrel{(23)}{=} \sum_{a \in B^c(\vec{x})} \frac{\partial F}{\partial x_a} \left(\frac{\partial F}{\partial x_a} - \frac{1}{|B^c(\vec{x})|} \sum_{a' \in B^c(\vec{x})} \frac{\partial F}{\partial x_{a'}} \right) \end{aligned}$$

(where $B(\vec{x})$ is some subset of the active constraints at \vec{x}),

$$\begin{aligned} &= \sum_{a \in B^c(\vec{x})} \left(\frac{\partial F}{\partial x_a} \right)^2 - \frac{1}{|B^c(\vec{x})|} \left(\sum_{a' \in B^c(\vec{x})} \frac{\partial F}{\partial x_{a'}} \right)^2 \\ &= \sum_{i \in B^c(\vec{x})} \left(\frac{\partial F}{\partial x_a} - \frac{1}{|B^c(\vec{x})|} \sum_{a' \in B^c(\vec{x})} \frac{\partial F}{\partial x_{a'}} \right)^2 \geq 0 \end{aligned}$$

i.e., F is a Lyapunov function of (28), and this proves that $F(\vec{x}(t))$ is increasing in t . As not all constraints can be active for $\epsilon < \frac{1}{N}$, $B^c(\vec{x}) \neq \emptyset$. Thus, \dot{F} is zero at some \vec{x}^* if and only if $\frac{\partial F}{\partial x_a} \Big|_{\vec{x}^*} = \frac{1}{|B^c(\vec{x}^*)|} \sum_{a' \in B^c(\vec{x}^*)} \frac{\partial F}{\partial x_{a'}} \Big|_{\vec{x}^*}$, for every $a \in B^c(\vec{x}^*)$. Moreover, by Lemma 5, for every $a \in B(\vec{x}^*)$, $x_a^* = \epsilon$ and $\frac{\partial F}{\partial x_a} \Big|_{\vec{x}^*} < \frac{1}{|B^c(\vec{x}^*)|} \sum_{a' \in B^c(\vec{x}^*)} \frac{\partial F}{\partial x_{a'}} \Big|_{\vec{x}^*}$. Any such vector \vec{x}^* will satisfy the KKT conditions with $\lambda^* = \frac{1}{|B^c(\vec{x}^*)|} \sum_{a' \in B^c(\vec{x}^*)} \frac{\partial F}{\partial x_{a'}} \Big|_{\vec{x}^*}$ and $\mu_a^* = \max(\lambda^* - \frac{\partial F}{\partial x_a} \Big|_{\vec{x}^*}, 0)$, $\forall a \in [N]$, and thus is a maximizer of F . \blacksquare

Note that the vector $\vec{x}(t)$ in the ODE (28) essentially follows the gradient $\nabla F(\vec{x})$; the term Z forces the evolution of \vec{x} to be contained in D_ϵ .

The following is an immediate implication of Lemma 6.

Lemma 7. *Assume that $\vec{x}(t)$, $t \in \mathbb{R}_+$, evolves according to the ODE (28), and let $L_{D_\epsilon} = \lim_{t \rightarrow \infty} \bigcup_{\vec{y} \in D_\epsilon} \{\vec{x}(s), s \geq t : \vec{x}(0) = \vec{y}\}$ be the limit set [6] of (28). Then, under Assumption 1, for every $\vec{x} \in L_{D_\epsilon}$, $F(\vec{x}) = \sup_{\vec{y} \in D_\epsilon} F(\vec{y})$.*

Proof: Suppose that $F(\vec{x}) \neq F^* = \sup_{\vec{y} \in D_\epsilon} F(\vec{y})$ for some $\vec{x} \in L_{D_\epsilon}$. Then, there exists a sequence $\{t_n\}_{n \geq 1}$ and a $\vec{y} \in D_\epsilon$ such that $\vec{x}(0) = \vec{y}$ and $\lim_{n \rightarrow \infty} \vec{x}(t_n) = \vec{x}$, under some norm in \mathbb{R}^N . As F is continuous, $\lim_{n \rightarrow \infty} F(\vec{x}(t_n)) = F(\vec{x}) \neq F^*$, which contradicts Lemma 6. \blacksquare

The reason we are interested in (28) is because it is a good ‘‘approximation’’ of the stochastic evolution of our mechanism (4). In particular, the following lemma states that the limit set of (4) will be included in L_{D_ϵ} :

Lemma 8. *Let $\{\vec{x}(k)\}_{k \in \mathbb{N}}$ be the sequence defined by the stochastic approximation algorithm (4), where $\gamma_k = \frac{1}{k}$ and $\epsilon > 0$. Then, under Assumption 1, w.p.1, the limit set of $\{\vec{x}(k)\}_{k \in \mathbb{N}}$ is a subset of the limit set of the ODE (28).*

Proof: We have that $\sup_k \mathbb{E}[\|V(\vec{x}_k)\|^2] < \infty$, by Lemma 4, and $\lim_{k \rightarrow \infty} \gamma_k = 0$, $\sum_{k=1}^{\infty} \gamma_k = \infty$, $\sum_{k=1}^{\infty} \gamma_k^2 = \infty$. Moreover, F is continuously differentiable, by Lemma 2, and constant and equal to $\sup_{\vec{x} \in D_\epsilon} F(\vec{x})$ on all limit points of the ODE (28), by Lemma 7. Hence, the assumptions (A2.1)-(A2.6) of Theorem 2.1 in Chapter 5 of Kushner and Yin [6] are satisfied. As a result, w.p.1, the limit points of $\{\vec{x}_k\}_{k \in \mathbb{N}}$ belong to limit set of the ODE (28). \blacksquare

Theorem 1 follows immediately from Lemmas 8 and 7. \blacksquare

D. Proof of Theorem 2

Let L_{D_ϵ} be the limit points of the ODE (28), and, for some $\delta > 0$, let $N_\delta = \{\vec{x} \in D_\epsilon : \exists \vec{y} \in L_{D_\epsilon} \text{ s.t. } \|\vec{x} - \vec{y}\|_1 \leq \delta\}$ be the δ -neighborhood of L_{D_ϵ} . By Lemma 4, $\sup_k \mathbb{E}[\|V(\vec{x}_k)\|_2^2] < \infty$ and, by Lemma 2, F is continuously differentiable. Hence, assumptions (A1.1) to (A1.4) of Theorem 2.1 in Chapter 8 of Kushner and Yin [6] are satisfied. As a result,

$$\lim_{\gamma \rightarrow 0} \lim_{T \rightarrow \infty} \mathbf{P}\left(\frac{1}{T} \int_0^T \mathbb{1}_{\vec{\xi}(t) \notin N_\delta} dt > \delta'\right) = 0, \text{ for all } \delta, \delta' > 0. \quad (29)$$

I.e., the fraction of time the system spends in the δ -neighborhood of L_{D_ϵ} converges to one, in probability. On the other hand, for $F^* = \sup_{\vec{x} \in D_\epsilon} F(\vec{x})$ and $L = \sum_{k=1}^{\infty} R(k)$,

$$\begin{aligned} \left| F^* - \frac{1}{T} \int_0^T F(\vec{\xi}(t)) dt \right| &\leq \left| F^* - \frac{1}{T} \int_0^T F(\vec{\xi}(t)) \mathbb{1}_{\vec{\xi}(t) \notin N_\delta} dt \right| + \dots \\ &+ \left| \frac{1}{T} \int_0^T F(\vec{\xi}(t)) \mathbb{1}_{\vec{\xi}(t) \notin N_\delta} dt \right| \leq L\delta + \frac{L}{T} \int_0^T \mathbb{1}_{\vec{\xi}(t) \notin N_\delta} dt \end{aligned}$$

by Corollary 1. Hence, for any $\delta', \delta > 0$,

$$\mathbf{P}\left(\left| F^* - \frac{1}{T} \int_0^T F(\vec{\xi}(t)) dt \right| > \delta'\right) \leq \mathbf{P}\left(\frac{1}{T} \int_0^T \mathbb{1}_{\vec{\xi}(t) \notin N_\delta} dt > \frac{\delta' - L\delta}{L}\right)$$

Theorem 2 thus follows from (29) by choosing $\delta < \delta'/L$. \blacksquare

E. Proof of Theorem 3

The first part of Theorem 3 can be proved by replacing F with F_C and \vec{g} with $\vec{g}_C = \frac{1}{|C|} \sum \vec{g}_s$ in the proofs of Theorems 1 and 2. For brevity, we simply give an outline of the proof below. To begin with, F_C is convex and continuously differentiable, as the average of convex and continuously differentiable functions. Lemma 2 implies that $\vec{g}_C(k) = \nabla F_C(\vec{x}(k)) + \vec{V}(\vec{x}(k))$ where $\vec{V}(\vec{x}(k))$ has zero expectation and bounded variance—the same as in Lemma (18), since \vec{g}_r are independent. Moreover, as in Lemma 6, the solutions of the ODE

$$\dot{\vec{x}} = \nabla F_C(\vec{x}) + Z(\vec{x}, \nabla F_C(\vec{x})) \quad (30)$$

where Z is given by (22) converge to maximizers of F_C . Finally, Lemma 8 and Theorem 2 can be similarly extended by relating the evolution of (8) to the ODE (30).

We therefore turn our attention to proving (10). To begin with, observe that for every $\vec{x} \in D_\epsilon$,

$$F_C(\vec{x}) = \frac{1}{|C|} \sum_s \sum_f d_{s,f} \mathbb{E}_{\vec{x}}[R(Y_{s,f})] = \sum_f d_{C,f} \mathbb{E}_{\vec{x}}[R(Y_f)]$$

where $d_{C,f} = \frac{1}{|C|} \sum_s d_{s,f}$. This is because $\mathbb{E}_{\vec{x}}[R(Y_{s,f})]$ does not depend on s : conditioned on the surfing strategy, Y_f has the same distribution for all surfers $s \in C$. For the same reason, for all $\vec{x} \in D_\epsilon$ and all $s \in C$, and for $\bar{R} = R(1)$,

$$|F_C(\vec{x}) - F_s(\vec{x})| = \left| \sum_f (d_{C,f} - d_{s,f}) \mathbb{E}_{\vec{x}}[R(Y_f)] \right| \leq \bar{R} \delta_C. \quad (31)$$

Let \vec{x}^* be a maximizer of F_C and $\vec{\sigma}^*$ a maximizer of F_s . Then $F_C(\vec{\sigma}^*) - F_C(\vec{x}^*) \leq 0$ and $F_s(\vec{x}^*) - F_s(\vec{\sigma}^*) \leq 0$ while, from (31), $F_s(\vec{\sigma}^*) - F_C(\vec{\sigma}^*) \leq \bar{R} \delta_C$ and $F_C(\vec{x}^*) - F_s(\vec{x}^*) \leq \bar{R} \delta_C$. Adding the former inequalities to the latter yields (10). \blacksquare

VI. DIMENSIONALITY REDUCTION

In this section, we outline a method for reducing the dimension of the surfing strategy. Consider a partition $P = \{B_1, B_2, \dots, B_{N'}\}$ of the set $\{1, \dots, N\}$ into N' groups, where $N' \ll N$. Moreover, let $\delta_P = \max_{i=1, \dots, N'} \max_{w, w' \in B_i} \|\vec{p}_w - \vec{p}_{w'}\|_1$ be the maximum l_1 distance between the publishing strategies of two websites belonging to the same group. Intuitively, if δ_P is small, websites in the same group publish content in a similar manner, and B_i can be seen as ‘‘communities’’ of websites.

Instead of implementing a surfing strategy over all sites, the mechanism can be indifferent towards sites from the same group. This can take place as follows. Consider strategies of the form $\vec{v} \in \mathbb{R}^{N'}$ defined over the groups B_i . To recommend a website, the mechanism first chooses a group B_i with probability v_i ; then, it chooses a website uniformly from within this group and displays it to the surfer; as a result, each website in B_i is shown with the same probability, $v_i/|B_i|$. The strategy \vec{v} is then updated according to (4), where websites are replaced by the groups B_i .

The above approach has the obvious advantage that the dimension of the vector \vec{v} is reduced to N' . Let $D'_\epsilon = \{\vec{v} \in \mathbb{R}^{N'} : \sum_i v_i = 1 \text{ and } v_i \geq \epsilon\}$ be the set of strategies \vec{v} defined over the partition P , and denote by $G(\vec{v}) = \mathbb{E}_{\vec{v}}[R(Y)]$ the expected rating of the system when the strategy is \vec{v} .

The following theorem, whose proof is in our technical report [8], states that, if δ_P is not too large—*i.e.*, the websites in the same group indeed have similar publishing strategies, the mechanism (4) will converge to a strategy close to the optimal.

Theorem 5. *Let $\vec{v}(k)$ be updated by the mechanism (4), where \vec{v} is defined as a strategy over the groups in partition P . Then, Theorems 1 and 2 hold if \vec{x} , $F(\vec{x})$ and D_ϵ are replaced by \vec{v} , $G(\vec{v})$, and D'_ϵ , respectively. Moreover, under Assumption 1, $\left| \sup_{\vec{v} \in D'_\epsilon} G(\vec{v}) - \sup_{\vec{x} \in D_\epsilon} F(\vec{x}) \right| = O(\delta_P)$.*

VII. NUMERICAL STUDY

A. Delicious Dataset

We evaluated the performance of our mechanism through a numerical study using traces from Delicious. The trace we used was collected and made publicly available⁶ by Göerlitz *et al.* [9]. It documents the tag assignments made by Delicious users between January 2006 and December 2007; for each tag assignment, the dataset contains the posting date, the user and website IDs and the actual tag.

We use this trace to obtain website publishing strategies and surfer (user) interest profiles as follows. First, due to memory

⁶<http://isweb.uni-koblenz.de/Research/DataSets>

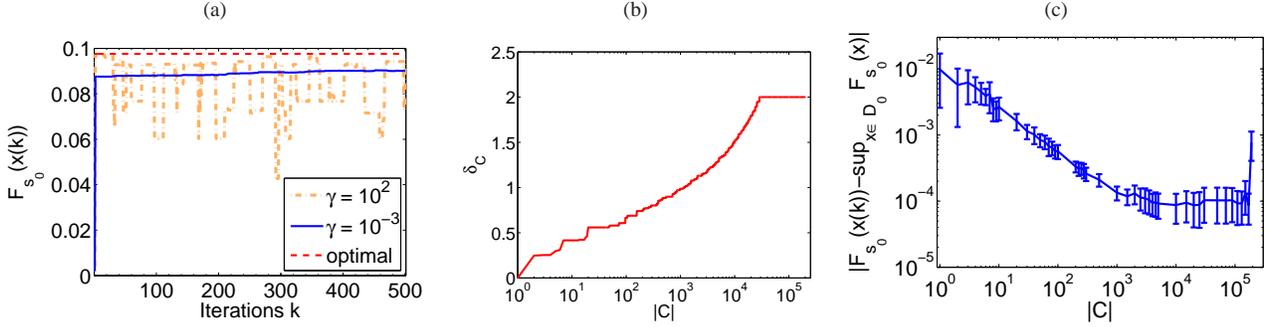


Fig. 2. In figure (a), an execution of mechanism (4) by the central surfer s_0 is illustrated for $\gamma = 10^2$ and $\gamma = 10^{-3}$: within one iteration, the algorithm gets within %1 and %20 of the optimal value, respectively, and the larger value of γ also implies higher variability. In figure (b), the diameter δ_C , given by (9), is plotted as a function of the community size $|C|$. Finally, in (c), we plot the mean distance of $F_{s_0}(x(k))$ from the optimal value $\sup_{D_\epsilon} F_{s_0}(\bar{x})$, for $\gamma = 0.001$, averaged over 10,000 iterations. Each point corresponds to a different community size $|C|$, and the bars show the standard deviation. We observe that, initially, increasing the community size $|C|$ leads to an improvement on the accuracy of the mechanism. The benefit of aggregating feedback among the community disappears at about $|C| = 1500$; above $|C| = 4000$, the trend is reversed and the average distance from the optimal value starts to increase.

constraints, we limit the dataset by focusing on the 20 most popular tags. We then further reduce the dataset by considering only tagging events for the 20 most tagged websites. This gives us a subset of the total trace, containing 20 websites, 19 tags (one was never used on the 20 most tagged websites), 203,084 users and 386,198 tagging events. We deduce the profile \vec{d}_s of user s by taking $d_{s,f}$ to be proportional to the number of times s tagged a website with tag f . Similarly, we deduce the publishing strategy \vec{p}_w of a site w by taking $p_{w,f}$ to be proportional to the number of times w was tagged by a user with tag f .

B. Evaluating the Effect of Community Feedback

To illustrate the behavior of our proposed mechanism, we selected the surfer s_0 that was positioned in the “center” of our dataset, *i.e.*, the average interest profile $\vec{d}_{C_0} = \sum_{C_0} \vec{d}_s / |C_0|$ over the entire dataset C_0 of 203,084 surfers.

Taking $\epsilon = 10^{-4}$ and $R(Y) = Y^{-2}$, we simulated two different executions of mechanism (4) as run by surfer s_0 with $\gamma = 100$ and $\gamma = 0.001$, respectively. These are illustrated in Fig. VI(a): in each case, we plot the expected reward $F_{s_0}(x(k))$ for $k = 1, \dots, 500$. We see that within only one iteration, the mechanism gets within %1 of the optimal value when $\gamma = 100$, and %20 when $\gamma = 0.001$. However, the larger value of γ also translates to higher variability.

To evaluate the effect of aggregating community feedback, we construct a community centered around s_0 as follows. Initially, our community C consists only of s_0 . We then add other surfers sequentially, by taking each time the closest surfer to s_0 that is not already in C . In figure VI(b), we plot the diameter δ_C of the community, *i.e.*, the maximum l_1 distance between two users in C , as a function of the community size $|C|$. Note that the diameter cannot exceed two.

For several different values of the community $|C|$, we ran the mechanism (8), that aggregated community feedback at each iteration, for a total of $k = 10^5$ iterations. In each execution, we took $\gamma = 10^{-3}$, and started from an optimal strategy. In Figure VI(c), we plot the distance of $F_{s_0}(\bar{x}(k))$ from its optimal value in D_ϵ , averaged over all k iterations.

This indicates the accuracy of the mechanism (8) in maximizing the rating as perceived by s_0 .

As predicted by Theorem 4, the accuracy improves drastically with $|C|$ —by two orders of magnitude—as long as the community diameter remains small. However, there exists a critical value, around $|C| = 1,500$, around which the benefits of increasing the community vanish. For $|C|$ larger than 4,000, δ_C becomes so large that the distance from the optimal value starts to increase. The above suggests that the ideal community size for surfer s_0 is between 1,500 and 4,500; based on the profiles of the other surfers extracted from the trace, aggregating feedback from larger communities will only decrease the accuracy of the system, as perceived by surfer s_0 .

VIII. CONCLUSIONS

In this paper, we proposed a model and a mechanism for personalized randomized web-surfing, that operated with only minimal feedback from surfers. In addition, we illustrated the benefit of feedback aggregation over surfer communities: one of our key findings is the identification of a trade-off between the size of a community and its diameter, the impact of which we are able to quantify both analytically and numerically.

REFERENCES

- [1] Stumbleupon’s recommendation technology. [Online]. Available: <http://www.stumbleupon.com/technology/>
- [2] N. Bansal and N. Koudas, “Searching the blogosphere,” in *WebDB*, 2007.
- [3] K. El-Arini, G. V. D. Shahaf, and C. Guestrin, “Turning down the noise in the blogosphere,” in *KDD*, 2009.
- [4] A. Das, M. Datar, A. Garg, and S. Rajaram, “Google news personalization: Scalable online collaborative filtering,” in *WWW*, 2007.
- [5] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl, “Evaluating collaborative filtering recommender systems,” *ACM Transactions on Information Systems*, vol. 22, no. 1, pp. 5–33, 2004.
- [6] H. J. Kushner and G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd ed. Springer, 2003.
- [7] C. Michelot, “A finite algorithm for finding the projection of a point onto the canonical simplex of \mathbb{R}^n ,” *Journal of Optimization Theory and Applications*, vol. 50, no. 1, pp. 195–200, 1986.
- [8] S. Ioannidis and L. Massoulié, “Surfing the blogosphere: Optimal personalized strategies for searching the web,” Thomson, Tech. Rep. CR-PRL-2009-07-0001, 2009.
- [9] O. Görlitz, S. Sizov, and S. Staab, “Pints: Peer-to-peer infrastructure for tagging systems,” in *IPTPS*, 2008.