Physica A 590 (2022) 126752

Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

# Design of biased random walks on a graph with application to collaborative recommendation

Pierre Leleux [a,*], Sylvain Courtain [a], Kevin Françoisse [b], Marco Saerens [c,d]

[a] LouRIM, Université Catholique de Louvain, Belgium
[b] Sagacify Company, Brussels, Belgium
[c] ICTEAM, Université Catholique de Louvain, Belgium
[d] IRIDIA, Université Libre de Bruxelles, Belgium

ARTICLE INFO

*Article history:*
Received 3 September 2020
Received in revised form 24 November 2021
Available online 8 December 2021

*Keywords:*
Network data analysis
Biased random walk
Bag-of-paths model
Collaborative recommendation

ABSTRACT

This work investigates a paths-based statistical physics formalism, inspired from the bag-of-paths framework, for the design of random walks on a graph in which the transition probabilities (the policy) are biased in favor of some node features. More precisely, given a weighted directed graph $G$ and a non-negative cost assigned to each edge, the biased random walk is defined as the policy minimizing the expected cost rate along the walks while maintaining a constant relative entropy rate. As for the standard bag-of-paths and the randomized shortest paths frameworks, the model assigns a Gibbs–Boltzmann distribution to the set of infinite walks and allows to recover known results from the literature, derived here from a different perspective. Examples of quantities of interest are the partition function of the system, the cost rate, the optimal transition probabilities, etc. In addition, the same formalism allows the introduction of capacity constraints on the expected node visit rates and an algorithm for computing the optimal policy subject to such capacity constraints is developed. Simulation results indicate that the proposed procedure can be effectively used in order to define a Markov chain driving the walk towards nodes having some specific properties, like seniority, education level or low node degree (hub-avoiding walk). An application relying on this last property is proposed as a tool for improving serendipity in collaborative recommendation, and is tested on the MovieLens data.

© 2021 Elsevier B.V. All rights reserved.

## 1. Introduction

### 1.1. General introduction

Random walks are used in virtually all areas of science in order to model various real-world stochastic processes [1–6]. In the case of a simple random walk on a discrete structure, the transition probabilities are usually set according to some local properties of the studied process (local affinities, probabilities of interaction, etc; see, e.g., [7–9]). However, in some cases, there is a need for orienting the walks towards regions or nodes of interest according to either some properties of edges and nodes, or to some external information which is made available. One well-known example is the personalization trick used in the famous PageRank algorithm [10].

* Corresponding author.
*E-mail addresses:* p.leleux@uclouvain.be (P. Leleux), sylvain.courtain@uclouvain.be (S. Courtain), kevin@sagacify.com (K. Françoisse), marco.saerens@uclouvain.be (M. Saerens).

https://doi.org/10.1016/j.physa.2021.126752
0378-4371/© 2021 Elsevier B.V. All rights reserved.

This paper describes an extension of the simple random walk on a graph in which walks are optimally biased towards nodes or properties of interest, as already investigated in [11–13]. We therefore assume the availability of a cost defined on each edge of the graph, reflecting the penalty incurred at each movement. In this context, our purpose is to develop a new random walk minimizing the *cost rate* while maintaining a fixed *relative entropy rate*. The resulting random walk is biased towards low-cost walks and the trade-off between exploitation (i.e., favoring low-cost transitions) and exploration (i.e., performing a simple random walk) is monitored by an inverse temperature parameter $\beta$. The policy guiding the walker takes the form of a Markov chain defined by its transition matrix, favoring low-cost transitions, depending on $\beta$. It is obtained by computing the dominant eigenvector of some special matrix, and depends on the whole graph, as already appeared in [14,15], and more recently [11–13]. The present paper therefore builds on this previous work.

This optimal policy is re-derived here in a principled way from a paths-based statistical physics framework [16–18] already used in other applications [19–22], and initially inspired by work in transportation sciences [23]. It assigns a Gibbs–Boltzmann probability distribution to the set of paths in the network, reflecting the probability of following each of these paths. Then, various quantities of interest, such as the visit rate to nodes, the cost rate, the local policy of the random walker, etc, are easily derived from the free energy function of the system. In this formalism, the Gibbs–Boltzmann probability distribution defined on paths and the local policy taking the form of a discrete Markov chain are two equivalent representations of the same process. Notice that the model recovers the results of the maximum entropy random walk (entropy rank, [11,12,24]) and the minimum free energy random walk (free-energy rank, [11,12]). It is also closely related to the method introduced in [13]; see the related work section. It therefore offers a new point of view on the problem which allows to extend its scope in the second part of the paper, by introducing capacity constraints on nodes.

Indeed, thereafter, we consider *capacity constraints* on the expected visit rate of the nodes (the equilibrium distribution) and derive an algorithm for designing transition probabilities minimizing expected cost rate with relative entropy rate regularization satisfying these constraints. This is done by considering the concave dual problem which is optimized through a gradient ascent method, as proposed in [25] for edge flow constraints.

Finally, after some illustrative examples, we show in a concrete case study how to take advantage of these techniques in order to enhance serendipity in collaborative recommendation.

### 1.2. Related work

First of all, the main difference between the model presented in this work and previous work on the **bag-of-paths** (BoP) and the **randomized shortest paths** (RSP) models (see [19–22,26] based on a paths formalism and [27] for a derivation based on edge flows) lies in the fact that, here, we are dealing with a set of paths of infinite length on a strongly connected graph,[1] with the consequence that we have to consider cost rate and relative entropy rate instead of cumulated quantities. In previous work, the paths were essentially finite, although unbounded (they represent killed, or absorbing, random walks). The BoP and RSP frameworks were used, among others, for defining dissimilarity measures and betweenness measures between nodes of a graph [19,22,28].

Besides this, the present paper was mainly inspired by three previous works. The first work adopts a new formalism for deriving a routing policy in transportation networks based on full paths [23,29]; in this field, the model is often called the "Markovian traffic assignment". The main idea is to consider the minimization of the expected path cost subject to entropy constraint (regularization). The routing policy is then expressed as a probability distribution over all the possible paths connecting the source node and the target node. The author shows that this distribution over paths is equivalent to a Markov chain whose transition probability matrix can be computed in closed form [23] (for recent advances in Markovian traffic assignment, see, e.g., [30]). This seminal work led to various extensions in network data analysis, as, e.g., the randomized shortest paths [21,22,31] and the bag-of-paths frameworks [19,20,31], adopting both a statistical physics formalism.[2] The present work extends this framework by considering cost and entropy rate on *regular* Markov chains.

The second work, itself inspired by the work of [14,15], proposes to use the maximally entropic random walks for ranking nodes of an unweighted graph ([11,12], often called Ruelle–Bowen discrete random walk; see also [24,33–38] for other work dedicated to maximal entropy random walks). After having computed a maximally-entropic Markov chain, then, as for PageRank, the stationary distribution is used as a measure of "importance" of the nodes. The authors name this measure *entropy rank*. They further compare the properties of this measure, and some other extensions, to PageRank [10] and HITS [39]. The link between our model and entropy rank is discussed at the end of Section 2.2. The same paper [11] further generalizes this work[3] by considering energy terms on transitions and consequently minimizing the free energy rate, leading to the *free entropy rank* quantity. The paths-based formalism used in the present paper leads to the same quantities and slightly more general results,[4] showing that the two frameworks are equivalent. Recently, this work has

[1] Thus inducing a regular Markov chain.

[2] For related work concerning other frameworks than the BoP/RSP also interpolating between a random walk and an optimal least-cost behavior, see for instance the related works in [19,26,32] and references therein.

[3] Initially introduced for dealing with disconnected graphs in [11] and extended to any energy term in [12].

[4] Besides the fact that the results are derived from a different perspective, the first difference between the work of [11,12] and our work is that we use the relative entropy (or Kullback–Leibler divergence) rate whereas they consider the Shannon entropy. Another difference is that we use a temperature parameter allowing to interpolate between a predefined random and a least cost behavior. In Section 2.2, after Eq. (13), it is shown how to readily derive our results from [11,12].

been extended in [40,41] where the authors develop optimal transportation policies between two sets of nodes along paths of finite length. This appears to be an extension of the so-called continuous state-space Schrödinger bridge problem and its discrete counterpart developed in [42].

The third work [13,43] introduces a method that reduces the computation of the policy of a simplified setting of Markov decision processes (the "environment" is deterministic and costs are defined on nodes) to the extraction of the dominant eigenvector of a matrix. More precisely, Todorov observed that, when injecting a particular form of control into the Bellman-Ford optimality expression, considering costs on states, and adding a Kullback–Leibler (aka relative entropy) term to the immediate cost,[5] this expression simplifies into a linearly solvable form. It turns out that Todorov's model defines a Markov chain whose transition probabilities (the policy) are biased towards lower expected costs, in the same way as in this work. Then the least-cost policy can be obtained at the limit by letting the state costs increase towards infinity. The work [13] had a significant impact in the field of reinforcement learning and control theory (see also [45]). Indeed, many new reinforcement learning procedures were designed based on this pioneering work by adding a Kullback–Leibler term to the immediate costs. They are sometimes called "Kullback–Leibler" or "path integral" control in the literature (see, e.g., [45–52], as well as chapter 7 of [53] and references therein). Moreover, notice also the recent related paper [54] which tackles the problem of synthesizing a policy that maximizes the entropy of a Markov decision process subject to expected reward constraints. This last work assumes an absorbing process while, in this paper, we are interested in regular Markov chains.

The model proposed in the first part of this paper is therefore a reformulation (in terms of the bag-of-paths framework) of Todorov's initial model, putting new light on his technique. Indeed, in the present work, we reformulate the problem as an optimization over paths in the network and deduce the optimal policy from a statistical physics point of view, therefore providing a new interpretation of [13] and extending the scope of the bag-of-paths model.

The second part of this paper focuses on the introduction of capacity constraints on the expected node visit rates. Such constraints are especially important in network flow and routing applications; see for instance [25,55] for related work in the context of the randomized shortest paths framework. Another recent interesting work is [56], studying Markov decision processes subject to linear temporal logic constraints. This looks like a promising way of synthesizing an optimal policy while imposing constraints on the Markov process.

Finally, there has been various propositions for biasing the transition probabilities of a random walk on a graph, usually by using some heuristic rules. For instance, [57] (see also [58] and references therein) assigns link weights according to preferential transition probabilities which define biased random walks that try to avoid hub nodes in the context of stochastic routing of packages in scale-free networks. Other authors bias the transition probabilities based on features available on the nodes in order to achieve a particular goal, for instance predicting new links that will most probably appear in the network, see for instance [59]. Other rely on potential-driven random walks biasing the walk (see the recent paper [60] and references therein). Riascos and Mateos [61] present an overview of different types of random walk strategies with local and non-local transitions on connected networks; see this work for details.

### 1.3. Contributions and organization of the paper

Overall, the main contributions of the paper are

- With respect to the previously introduced BoP and RSP models, their scope is extended here to regular Markov chains by considering the optimization of cost rate and relative entropy rate on a strongly connected graph. The BoP/RSP framework is thus generalized to infinite-length random walks starting from a source node, whereas the standard RSP deals with random walks whose length is finite (although unbounded) between a single source node and a single target node. The standard BoP model is also based on random walks whose length is finite but considers all the paths in the network, between all pairs of source-target nodes. This generalization allows to define biased random walks on a graph in a suitable way while the standard RSP is mainly used to define absorbing random walks from which dissimilarity and betweenness centrality measures can be derived.
- The procedure for computing the policy (the biased random walk), minimizing the objective function, is derived in a principled, generic, way by applying statistical physics concepts [16–18]. The solution is obtained by solving an eigensystem, instead of solving a system of linear equations in the standard RSP or computing a matrix inverse for the BoP. It is also shown that we recover the results of [11,12], which are discussed.
- A new algorithm allowing to introduce capacity constraints on the node visit rate is developed within the same framework, inspired by [25], where flow constraints were considered.
- We apply the algorithms to a collaborative recommendation problem in which we bias the recommendation in favor of infrequent items in order to enhance serendipity.

The paper is mainly divided into three parts. The first part, Section 2, consists in the development of the biased random walk. The second part, Section 3, is devoted to the introduction of capacity constraints on the node visit rate. The third part starts with some preliminary experiments on 2D benchmark lattices for behavior monitoring in Section 4, followed by an application to collaborative recommendation in Section 5. Section 6 concludes the work.

[5] Note that a similar procedure was also introduced in physics [44] where it was called "smart random walkers".

## 2. The biased random walk model

In this section, a biased random walk on a graph $G$ is developed by considering the set of $t$-steps walks starting from a source node $i$. The main idea, close to maximum-entropy arguments [16,62,63], is to define a Gibbs–Boltzmann distribution [16–18] on the set of such paths by minimizing the expected cost rate subject to a fixed relative entropy rate. We then increase the length of the walks by taking the limit for $t \to \infty$. The resulting random walk is biased towards low-cost paths, which have a higher probability of being chosen than high-cost ones. At the limit of infinite walks, this biased random walk defines a first-order regular Markov chain on $G$.

### 2.1. Biased random walk model elaboration

#### 2.1.1. Background and notation

Consider a weighted directed, strongly connected, graph or network $G$ with a set $\mathcal{V}$ of $n$ nodes (or vertices) and a set of $m$ edges $\mathcal{E}$. Its adjacency matrix $\mathbf{A}$ contains nonnegative affinities between nodes while the cost matrix $\mathbf{C}$ contains nonnegative transition costs between nodes. When there is no link between node $k$ and node $l$, we must have $a_{kl} = 0$ and $c_{kl} = \infty$. Otherwise, for existing links, $a_{kl} > 0$ and $c_{kl} \geq 0$ are assumed to be set independently of each other, depending on the application. For instance, typically, a natural random walk on $G$ is defined by setting the likelihood of following a link $(k, l)$ in node $k$ as proportional to its affinity, $a_{kl}$, without taking the costs into account. On the other hand, costs are usually set proportional to some property of the edge or the next node, quantifying the cost-to-go to this node. The biased random walk developed in this work will gradually bias the random walk on $G$ towards low-cost paths. As an example, in a social network marketing application, we could want to bias the random walk in order to favor nodes representing teenagers.

Inspired by [19–23], the developed framework relies on a bag of paths (BoP) containing paths (also called walks) of length[6] $t$ on $G$. More precisely, we assume a bag containing objects,[7] with two properties:

- First, the objects that are drawn are *paths* of length $t > 0$ starting from a given source node.
- Second, each path will be weighted according to its quality, i.e. its *total cost*. The rationale is that the likelihood of picking a low-cost path should be higher than picking a high-cost path — low-cost paths being therefore favored.

Moreover, each path $\wp_{ij}(t)$ of length $t$ is characterized by a starting node $i$ and an ending node $j$, thus connecting $i$ to $j$. Then, the set of all paths (including cycles) of length $t$ linking these two nodes is denoted as $\mathcal{P}_{ij}(t) = \{\wp_{ij}(t)\}$. Thus, $\mathcal{P}_{ij}(t)$ contains all the paths $\wp_{ij}(t)$ allowing to reach node $j$ from node $i$ in *exactly* $t$ steps.

Furthermore, it is assumed that the total cost of a path is *additive*, i.e., $\tilde{c}(\wp_{ij}(t)) = \sum_{\tau=1}^{t} c_{k_{\tau-1}k_\tau}$ where $k_0 = i$ is the starting node and $k_t = j$ is the ending node. In addition, let us define the set (bag) of all $t$-length paths through the graph between node $i$ and any other node as $\mathcal{P}_i(t) = \cup_{j=1}^{n} \mathcal{P}_{ij}(t)$. It contains all paths of length $t$ starting from a fixed node $i$. Note that sets $\mathcal{P}_{ij}(t)$ are mutually exclusive.

#### 2.1.2. Probability distribution on the set of paths

The present section describes how path probabilities are assigned. The probability distribution on the finite set of paths, $\mathcal{P}_i(t)$, representing the probabilities of drawing a path $\wp \in \mathcal{P}_i(t)$ from the bag of paths, is defined as the discrete probability mass $\{\mathrm{P}(\wp)\}_{\wp\in\mathcal{P}_i(t)}$ minimizing the expected cost rate, $\mathbb{E}[\tilde{c}(\wp)/t]$, among all the distributions having a fixed relative entropy rate $J_0$ with respect to a reference distribution — for instance a natural random walk on the graph with transition probabilities $p_{kl}^{\mathrm{ref}} = a_{kl}/\sum_{l'=1}^{n} a_{kl'}$ [20]. The transition matrix containing the transition probabilities of this *reference* random walk will be denoted as $\mathbf{P}_{\mathrm{ref}}$ and is assumed to be regular.

In other words, we are seeking for path probabilities, $\mathrm{P}(\wp)$, $\wp \in \mathcal{P}_i(t)$, minimizing the expected cost rate (average cost along the period $t$) subject to a constant relative entropy rate [36,62,65–67] constraint,

$$\left|\begin{array}{ll} \underset{\{\mathrm{P}(\wp)\}}{\text{minimize}} & \frac{1}{t}\displaystyle\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\tilde{c}(\wp) \\ \text{subject to} & \frac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\log\big(\mathrm{P}(\wp)/\tilde{\pi}^{\mathrm{ref}}(\wp)\big) = J_0 \\ & \sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) = 1 \end{array}\right. \tag{1}$$

where $\tilde{c}(\wp)/t$ is the cost rate along path $\wp$. In this equation, $\tilde{\pi}^{\mathrm{ref}}(\wp)$ represents the likelihood of following the path $\wp$ when walking according to the reference random walk, i.e. using transition probabilities $p_{kl}^{\mathrm{ref}}$. It is simply the product of the transition probabilities along the path $\wp$, $\tilde{\pi}^{\mathrm{ref}}(\wp) = p_{k_0k_1}^{\mathrm{ref}} p_{k_1k_2}^{\mathrm{ref}} \cdots p_{k_{t-1}k_t}^{\mathrm{ref}}$ where the walker sequentially visits nodes $(k_0 = i, k_1, \ldots, k_{t-1}, k_t)$. Note that, because $\sum_{l=1}^{n} p_{kl}^{\mathrm{ref}} = 1$ for each $k$, the path likelihood is properly normalized. Indeed,

$$\sum_{\wp\in\mathcal{P}_i(t)} \tilde{\pi}^{\mathrm{ref}}(\wp) = \sum_{k_1=1}^{n}\sum_{k_2=1}^{n}\cdots\sum_{k_t=1}^{n} p_{ik_1}^{\mathrm{ref}} p_{k_1k_2}^{\mathrm{ref}} \ldots p_{k_{t-1}k_t}^{\mathrm{ref}} = 1$$

[6] The limit $t \to \infty$ will be considered later.

[7] These objects being words in information retrieval. Note that bags of spanning trees and of forests were considered in [32,64].

In Eq. (1), $J_0 > 0$ is provided a priori by the user, according to the desired degree of randomness. It is well-known that optimizing the following Lagrange function

$$\begin{aligned}\mathscr{L}(\mathbb{P}) =& \tfrac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\tilde{c}(\wp) + \mu\left(\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) - 1\right)\\ &+ T\left(\tfrac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\log\frac{\mathrm{P}(\wp)}{\tilde{\pi}^{\mathrm{ref}}(\wp)} - J_0\right)\end{aligned} \tag{2}$$

over the set of path probabilities $\mathbb{P} = \{\mathrm{P}(\wp)\}_{\wp\in\mathcal{P}_i(t)}$ of finite length $t$ starting in a fixed node $i$ [19–21]; see also [11] yields a *Gibbs–Boltzmann* [16,63] probability distribution,

$$\mathrm{P}(\wp) = \frac{\tilde{\pi}^{\mathrm{ref}}(\wp)\exp\left[-\beta\tilde{c}(\wp)\right]}{\sum_{\wp'\in\mathcal{P}_i(t)} \tilde{\pi}^{\mathrm{ref}}(\wp')\exp\left[-\beta\tilde{c}(\wp')\right]}, \text{ for } \wp\in\mathcal{P}_i(t) \tag{3}$$

where the Lagrange parameter $T$ plays the role of a temperature and $\beta = 1/T$ is the inverse temperature related to the relative entropy $J_0$. Thus, as expected, short parts (having a low cost $\tilde{c}(\wp)$) are favored in that they have a large probability of being sampled. The inverse temperature parameter $\beta$ allows to monitor the balance between *exploration* and *exploitation*. Indeed, when $\beta \to 0$, we recover the natural random walk on $G$ (pure exploration) whereas for $\beta \to \infty$, only lowest cost paths are considered (pure exploitation).

In the sequel, it will be more convenient to provide the value of the parameter $\beta$ instead of the relative entropy. Thus, we will assume that $\beta$ is the input to the algorithm, together with the adjacency and the cost matrices.

### 2.1.3. Design of the biased random walk policy

The previous result (Eq. (3)) assigns a probability distribution on all the paths of length $t$ starting in node $i$. Based on this distribution, the objective now is to compute an optimal transition probability matrix gradually favoring transitions towards low-cost nodes when $\beta$ increases, for $t \to \infty$. To this end, the idea will be, first, to find an expression computing the expected passage rate through each edge $(k, l)$ of the graph for bounded walks (of length $t$) and, then, to compute the transition probability matrix by taking the limit for $t \to \infty$, that is, considering walks of infinite length.

To do this, let us further define the *partition function* of the system as the denominator of Eq. (3),

$$\mathcal{Z}_i(t) \triangleq \sum_{\wp\in\mathcal{P}_i(t)} \tilde{\pi}^{\mathrm{ref}}(\wp)\exp\left[-\beta\tilde{c}(\wp)\right] \tag{4}$$

The minimal *free energy* rate of the system of paths is then defined from the partition function as

$$\phi_i(t) \triangleq -\tfrac{1}{t\beta}\log\mathcal{Z}_i(t) \tag{5}$$

It corresponds to the minimized objective function (expected cost rate plus $T$ times relative entropy rate; see Eq. (2)) in which the optimal Gibbs–Boltzmann path probability distribution (3) is injected [31].[8]

From standard statistical physics results [16–18] and proceeding as in [21], the expected *passage rate* through edge $(k, l)$ as well as the *visit rate* to each node, respectively denoted as $\bar{n}_{kl}(t) \triangleq \frac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \eta\big((k,l)\in\wp\big)\,\mathrm{P}(\wp)$ and $\bar{n}_k(t) \triangleq \frac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \eta\big(k\in\wp\big)\,\mathrm{P}(\wp)$, where $\eta\big((k,l)\in\wp\big) = \partial\tilde{c}(\wp)/\partial c_{kl}$ provides the number of times edge $(k, l)$ appears on path $\wp$ and $\eta\big(k\in\wp\big)$ is the number of visits to node $k$ along $\wp$, can be computed from the free energy rate (5) by[9]

$$\begin{cases}\bar{n}_{kl}(t) &= \dfrac{\partial\phi_i(t)}{\partial c_{kl}}\\ \bar{n}_l(t) &= \displaystyle\sum_{k\in\mathcal{P}red(l)} \bar{n}_{kl}(t) + \frac{\delta_{il}}{t}\end{cases} \tag{6}$$

where $\mathcal{P}red(l)$ is the set of predecessor nodes of node $l$. The contribution $\delta_{il}/t$ in the second expression comes from the fact that the initial visit to starting node $i$ at $t = 0$ is not counted in the first term of the same expression. Note, however, that, for a large $t$ ($t \to \infty$), the expected visit rate can simply be computed by $\bar{n}_l(\infty) = \sum_{k\in\mathcal{P}red(l)} \bar{n}_{kl}(\infty)$ or by $\bar{n}_k(\infty) = \sum_{l\in\mathcal{S}ucc(k)} \bar{n}_{kl}(\infty)$ where $\mathcal{S}ucc(k)$ is the set of successor nodes of node $k$.

As for the standard bag-of-paths formalism [19], the optimal policy expressed in terms of path probabilities in Eq. (3), and providing the minimum-cost policy for a constant $J_0$, can be re-expressed in terms of optimal, *local*, *transition probabilities* (i.e., a first order Markov chain — see later) when considering infinite walks and thus stationarity/equilibrium,

$$p^\star_{kl} = \lim_{t\to\infty}\frac{\bar{n}_{kl}(t)}{\sum_{l\in\mathcal{S}ucc(k)}\bar{n}_{kl}(t)} = \lim_{t\to\infty}\frac{\bar{n}_{kl}(t)}{\bar{n}_{k\bullet}(t)} = \frac{\bar{n}_{kl}(\infty)}{\bar{n}_{k\bullet}(\infty)} = \frac{\bar{n}_{kl}(\infty)}{\bar{n}_k(\infty)} \tag{7}$$

[8] See also Eq. (22) for a similar argument.

[9] We omit the specification of the starting node, fixed to node $i$ for simplicity.

which is nothing else than the expected passage rate through edge $(k, l)$ which is normalized in order to sum to one. Therefore, $\bar{n}_l(\infty) = \sum_{k\in\mathcal{P}red(l)} \bar{n}_{kl}(\infty) = \sum_{k\in\mathcal{P}red(l)} \bar{n}_k(\infty) p^\star_{kl}$ which confirms that the $\bar{n}_k(\infty)$ are the elements of the stationary distribution, $\pi_k$, of the biased random walk, up to a scaling factor.[10] A random walker moving according to these transition probabilities minimizes the objective function in (2), that is, the sum of the expected cost rate and $T$ times the relative entropy rate.

Let us now compute these quantities of interest from the free energy rate of the system.

### 2.2. Computation of the main quantities

#### 2.2.1. Free energy rate

It was shown in [19,21,22] that the partition function $\mathcal{Z}_i$ can be computed from the cost matrix, $\mathbf{C}$, and the reference transition probabilities matrix, $\mathbf{P}_{\text{ref}}$, containing the $p^{\text{ref}}_{kl}$. The same is true here when considering cost rate and a strongly connected graph. Let us first define the auxiliary matrix $\mathbf{W}$ as

$$\mathbf{W} = \mathbf{P}_{\text{ref}} \circ \exp[-\beta \mathbf{C}] \tag{8}$$

where the symbol $\circ$ denotes the elementwise matrix product, $\mathbf{P}_{\text{ref}} = \mathbf{D}^{-1}\mathbf{A}$, and $\mathbf{D}$ is the diagonal out-degree matrix. Elementwise, this last Eq. (8) reads $w_{kl} = p^{\text{ref}}_{kl} \exp[-\beta c_{kl}]$. Notice that $\mathbf{W}$ is usually nonsymmetric but, for a random walk on an undirected graph, the matrix is similar to a symmetric matrix $\mathbf{B} = \mathbf{D}^{1/2}\mathbf{W}\mathbf{D}^{-1/2}$. This implies that the eigenvalues and eigenvectors of $\mathbf{W}$ can be chosen as real in the case of a random walk on an undirected graph (reversible Markov chain).

It can be shown (see the Appendix of this paper, or [19–21]) that the partition function (4) for a finite walk of length $t$ starting in $i$ can be computed through

$$\mathcal{Z}_i(t) = \mathbf{e}_i^{\mathrm{T}} \mathbf{W}^t \mathbf{e} \tag{9}$$

where $\mathbf{e}_i$ is a binary indicator vector full of 0s except at position $i$ where the entry is 1, and $\mathbf{e}$ is a vector full of 1s.

Assuming that matrix $\mathbf{W}$ has a complete set of eigenvectors, by considering its spectral decomposition, $\mathbf{W} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^*$ (see, e.g., [68,69]) with $\mathbf{V}^* = \mathbf{U}^{-1}$ and where $*$ denotes the conjugate transpose, we get

$$\mathbf{W}^t = \sum_{k=1}^{n} \mathbf{u}_k \lambda_k^t \mathbf{v}_k^* \tag{10}$$

The $\{\mathbf{u}_k\}_{k=1}^n$ and the $\{\mathbf{v}_k\}_{k=1}^n$ are respectively the columns of $\mathbf{U}$ and $\mathbf{V}$ obtained by the eigensystems $\lambda_k \mathbf{u}_k = \mathbf{W}\mathbf{u}_k$ and $\lambda_k^* \mathbf{v}_k = \mathbf{W}^{\mathrm{T}}\mathbf{v}_k$. Here, we assume the standard normalization $\mathbf{U}^*\mathbf{V} = \mathbf{V}^*\mathbf{U} = \mathbf{I}$. The eigenvalues $\{\lambda_k\}_{k=1}^n$ can be complex and are sorted by decreasing modulus. The nonnegative matrix $\mathbf{W}$ is irreducible and sub-stochastic[11]; therefore its spectral radius is lesser than 1, $\rho(\mathbf{W}) < 1$ [70]. Hence, all its eigenvalues have modulus lesser than 1.

Then, from Eqs. (5), (9) and (10), the free energy rate is

$$\phi_i(t) = -\tfrac{1}{t\beta} \log(\mathbf{e}_i^{\mathrm{T}} \mathbf{W}^t \mathbf{e}) = -\tfrac{1}{t\beta} \log \sum_{k=1}^{n} u_{ik} \lambda_k^t v^*_{\bullet k} \tag{11}$$

where $u_{ij} = [\mathbf{U}]_{ij}$ and $v_{ij} = [\mathbf{V}]_{ij}$, that is, element $i, j$ of the corresponding matrix. As before, the dot $\bullet$ represents summation on the corresponding index. Now, we can pursue the computation of the expected passage rate $\bar{n}_{kl}(t)$ provided by Eq. (6). As explained earlier, we need to take the limit $t \to \infty$ to obtain the expected passage rate for paths of infinite length, that is, when the system reaches its stationary distribution ($n_{kl}(\infty)$),

$$\lim_{t\to\infty} \bar{n}_{kl}(t) = \lim_{t\to\infty} \partial_{c_{kl}} \phi_i(t) = \partial_{c_{kl}} \lim_{t\to\infty} \phi_i(t) \tag{12}$$

From Eq. (11), the limit of the free energy rate can easily be computed

$$\begin{aligned}
\lim_{t\to\infty} \phi_i(t) &= \lim_{t\to\infty} -\tfrac{1}{t\beta} \log\left(\lambda_1^t u_{i1} v^*_{\bullet 1} + \sum_{k=2}^{n} \lambda_k^t u_{ik} v^*_{\bullet k}\right) \\
&= -\tfrac{1}{\beta} \lim_{t\to\infty} \tfrac{1}{t} \log\left(\lambda_1^t \left(u_{i1} v^*_{\bullet 1} + \sum_{k=2}^{n} \left(\frac{\lambda_k}{\lambda_1}\right)^t u_{ik} v^*_{\bullet k}\right)\right) \\
&= -\tfrac{1}{\beta} \lim_{t\to\infty} \left[\tfrac{1}{t} \log \lambda_1^t + \tfrac{1}{t} \log\left(u_{i1} v^*_{\bullet 1} + \sum_{k=2}^{n} \left(\frac{\lambda_k}{\lambda_1}\right)^t u_{ik} v^*_{\bullet k}\right)\right]
\end{aligned}$$

[10] They correspond to the elements of the left dominant eigenvector of the transition probability matrix associated with eigenvalue 1.

[11] Indeed, $\beta > 0$, $c_{ij} \geq 0$ for all $i, j$ and we assume that at least one element of $\mathbf{C}$ is strictly positive.

Now, since the matrix $\mathbf{W}$ is sub-stochastic and irreducible, its dominant eigenvalue (the spectral radius) is real, has multiplicity one, and we have $1 > |\lambda_1| > |\lambda_2| \geq \cdots \geq |\lambda_n|$; thus $\frac{|\lambda_k|}{|\lambda_1|} < 1$ for all $k \neq 1$ [70]. This last equation therefore simplifies into

$$\phi_i(\infty) = \lim_{t\to\infty} \phi_i(t) = -\tfrac{1}{\beta} \log \lambda_1 \tag{13}$$

and we observe that the free energy rate is independent of the starting node $i$ when considering infinite-length walks, as should be.

*2.2.2. Expected visit rate*

The expected visit rate to each node can now be readily computed from the free energy rate. Injecting the last result (13) in Eq. (12) yields

$$\bar{n}_{kl}(\infty) = -\tfrac{1}{\beta} \partial_{c_{kl}} \log \lambda_1 \tag{14}$$

In addition, we know (see, e.g., [71], p. 183, or [72]) that $d\lambda_1 = (\mathbf{v}_1^*(d\mathbf{W})\mathbf{u}_1)/(\mathbf{v}_1^*\mathbf{u}_1)$ which provides, in our case,[12] $\partial\lambda_1/\partial w_{kl} = [\mathbf{v}_1^*]_k[\mathbf{u}_1]_l$ where $\mathbf{u}_1$ and $\mathbf{v}_1$ are the dominant eigenvectors (right and left) of $\mathbf{W}$ and $[\mathbf{v}_1^*]_k$ extracts element $k$ of this (row or column) vector (for node $k$). By Perron–Frobenius theory, these dominant eigenvectors can be chosen as real and non-negative, thus avoiding the additional complication of dealing with complex numbers [70]. Therefore,

$$\partial_{c_{kl}} \log \lambda_1 = \tfrac{1}{\lambda_1} \frac{\partial \lambda_1}{\partial w_{kl}} \frac{\partial w_{kl}}{\partial c_{kl}} = -\beta \frac{[\mathbf{v}_1]_k[\mathbf{u}_1]_l}{\lambda_1} w_{kl} \tag{15}$$

From Eqs. (14) and (15), we finally obtain for the expected passage rate

$$\bar{n}_{kl}(\infty) = \frac{[\mathbf{v}_1]_k[\mathbf{u}_1]_l}{\lambda_1} w_{kl} \tag{16}$$

From (6), summing this quantity over $l$ as well as using $\sum_{l=1}^n w_{kl}[\mathbf{u}_1]_l = \lambda_1[\mathbf{u}_1]_k$ provides the expected visit rate to each node $k$, corresponding to the stationary distribution $\boldsymbol{\pi}$ of the biased random walk (see the discussion after Eq. (7)),

$$\bar{n}_k(\infty) = \pi_k = \sum_{l=1}^n \bar{n}_{kl}(\infty) = [\mathbf{v}_1]_k[\mathbf{u}_1]_k \tag{17}$$

which sums to one as the $\{\mathbf{u}_k\}_{k=1}^n$ and the $\{\mathbf{v}_l\}_{l=1}^n$ are orthonormal ($\mathbf{V}^*\mathbf{U} = \mathbf{I}$).

Here, we recover a result derived in [11] (Section 4, Eq. (7); see also [12]), themselves inspired by the work of [14,15] who already studied related problems. In that work, this quantity was called the free-energy rank. Besides the fact that the result is derived from a bag-of-paths perspective in the present paper, the other differences with this previous work are that (1) the authors use the entropy rate instead of the relative entropy rate and (2) they do not introduce explicitly the temperature parameter $T$ balancing exploration and exploitation. Indeed, for many applications, it is convenient to interpolate between a reference random walk (according to the reference transition probabilities, when $T \to \infty$) and a least cost behavior (when $T \to 0^+$), which is one of the main features of the RSP and the BoP frameworks [22]. Moreover, using the relative entropy instead of simply Shannon entropy allows to introduce reference transition probabilities. This ensures that the fundamental matrix of the bag-of-paths system is well-defined for all values of $T > 0$ [19,20,22,31]. Note, however, that these extensions can easily be derived from [11,12] after redefining the free energy rate by introducing a temperature parameter and rewriting costs according to $c'_{ij} = c_{ij} - T \log p_{ij}^{\text{ref}}$.

*2.2.3. Optimal policy for the biased random walk*

At this stage, the local policy corresponding to the transition probabilities of the biased random walk is easy to compute. Combining Eq. (7) with (16), the optimal transition probabilities, minimizing cost rate subject to constant relative entropy rate in the long run are

$$p_{kl}^\star = \frac{w_{kl}[\mathbf{u}_1]_l}{\sum_{l'=1}^n w_{kl'}[\mathbf{u}_1]_{l'}} = \tfrac{1}{\lambda_1} \frac{w_{kl}[\mathbf{u}_1]_l}{[\mathbf{u}_1]_k} \tag{18}$$

where, as stated previously, $\mathbf{u}_1$ is the right eigenvector of $\mathbf{W}$ associated with the largest eigenvalue in modulus, i.e., $\lambda_1$. Therefore, each row of this transition matrix sums to one as required. Note that this corresponds to the result obtained in [11] (Section 4, Eq. (5)) when deriving the free-energy rank quantity without considering a reference probability (taking $\mathbf{W} = \exp[-\mathbf{C}]$ instead of Eq. (8)). It also appears in Section 2.2, Eq. (24), of Todorov's paper [13], as explained in our related work. An algorithm for computing the optimal policy in matrix form is presented later in Algorithm 1, and is discussed in the text there.

Interestingly, if we adopt an entropy constraint instead of the relative entropy (there is no reference probability $\mathrm{P}_{\text{ref}}$) and the transition costs on the existing edges are made arbitrarily small, then $\mathbf{W} \to \mathbf{A}$ (an unweighted graph with a binary adjacency matrix), and we recover the maximum entropy random walk [11,24]; see also more recently [40].

[12] Recall that we assumed that the eigenvectors are normalized, $\mathbf{v}_1^*\mathbf{u}_1 = 1$.

### 2.3. Other quantities: expected cost and relative entropy rates

#### 2.3.1. Expected cost rate

In addition, following the same basic statistical physics arguments as for the expected visit rate (see Eqs. (4)–(5)), the expected cost rate is provided by

$$\bar{c}(t) \triangleq \frac{1}{t} \sum_{\wp \in \mathcal{P}_i(t)} \mathrm{P}(\wp)\tilde{c}(\wp) = \frac{\partial(\beta\phi_i(t))}{\partial\beta} \tag{19}$$

where $\mathrm{P}(\wp)$ is provided by Eq. (3) [21]. Then, the expected cost rate at equilibrium can be computed from the free energy rate (Eqs. (13) and (15)) thanks to

$$\begin{aligned}\bar{c}(\infty) &= \lim_{t\to\infty} \bar{c}(t) = \partial_\beta(\beta\phi_i(\infty)) = \partial_\beta(-\log\lambda_1)\\ &= \sum_{k,l=1}^{n} \frac{\partial(-\log\lambda_1)}{\partial w_{kl}} \frac{\partial w_{kl}}{\partial\beta} = \sum_{k,l=1}^{n} \frac{[\mathbf{v}_1]_k[\mathbf{u}_1]_l}{\lambda_1} w_{kl}c_{kl}\\ &= \sum_{k,l=1}^{n} \bar{n}_{kl}(\infty)\, c_{kl}\end{aligned} \tag{20}$$

which looks natural: we compute the sum of the flows (expected number of passage rate computed in Eq. (16)) on the edges multiplied by the corresponding costs.

Alternatively, let us follow another, more direct, way for getting the result [73]. By considering $\eta\big((k,l) \in \wp\big)$, the number of times edge $(k,l)$ is present on path $\wp$, the expression for computing the cost rate becomes

$$\begin{aligned}\bar{c}(t) &= \frac{1}{t} \sum_{\wp\in\mathcal{P}_i(t)} \tilde{c}(\wp)\,\mathrm{P}(\wp) = \frac{1}{t} \sum_{\wp\in\mathcal{P}_i(t)} \left(\sum_{k,l=1}^{n} \eta\big((k,l)\in\wp\big)\, c_{kl}\right)\mathrm{P}(\wp)\\ &= \sum_{k,l=1}^{n} \underbrace{\left(\frac{1}{t} \sum_{\wp\in\mathcal{P}_i(t)} \eta\big((k,l)\in\wp\big)\,\mathrm{P}(\wp)\right)}_{\bar{n}_{kl}(t)} c_{kl} = \sum_{k,l=1}^{n} \bar{n}_{kl}(t)\, c_{kl}\end{aligned} \tag{21}$$

and by taking the limit $t \to \infty$, we obtain the result (20).

#### 2.3.2. Relative entropy rate

Let us now compute the relative entropy rate from its definition (1) and from the Gibbs–Boltzmann distribution (3),

$$\begin{aligned}\bar{J}(t) &= \frac{1}{t} \sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) \log\left(\frac{\mathrm{P}(\wp)}{\tilde{\pi}^{\,\mathrm{ref}}(\wp)}\right)\\ &= \frac{1}{t} \sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) \log\left(\frac{\exp\left[-\beta\tilde{c}(\wp)\right]}{\mathcal{Z}_i(t)}\right)\\ &= \frac{1}{t} \sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\left(-\beta\tilde{c}(\wp) - \log\mathcal{Z}_i(t)\right)\\ &= -\beta\bar{c}(t) - \frac{1}{t}\log\mathcal{Z}_i(t) = -\beta\bar{c}(t) + \beta\phi_i(t)\end{aligned} \tag{22}$$

where we used Eq. (5). Again, taking the limit provides

$$\bar{J}(\infty) = \beta\big(\phi_i(\infty) - \bar{c}(\infty)\big) \tag{23}$$

Let us now consider yet another way of biasing random walks by introducing capacity constraints on the node visit rate.

## 3. Capacity constraints on node visit rate

Another interesting option, investigated in this section, is to restrict the expected rate of visits to each node in order to soften the hubness effect. Indeed, this section is largely inspired by recent work [25] showing how to introduce capacity constraints on edges in the standard randomized shortest paths formalism. This idea is exploited here in order to introduce constraints on the node visit rate, and thus on the stationary distribution, instead of the edge flow in [25]. The main objective is thus to constrain the visit rate of some nodes, belonging to a set of *constrained nodes* $\mathcal{C}$, to stay below a given threshold,

$$\bar{n}_k(t) \leq \sigma_k \quad \text{for nodes } k \in \mathcal{C} \tag{24}$$

which ensures that the rates are inside the interval $[0, \sigma_k]$. Of course, we must assume that the provided constraints are feasible, which could be difficult to verify a priori. The procedure developed in this section will minimize the objective function (expected cost rate with relative entropy rate regularization) and provide an optimal policy while satisfying the capacity constraints (24) on node visits.

### 3.1. The Lagrange function integrating capacity constraints

Let us first study the form of the Lagrange function including capacity constraints. By using standard nonlinear optimization theory [74–76], the Lagrange function (see Eq. (2)) can be rewritten as

$$\begin{aligned}\mathscr{L}(\mathbb{P}, \boldsymbol{\lambda}) = &\tfrac{1}{t} \sum_{\wp \in \mathcal{P}_i(t)} \mathrm{P}(\wp)\tilde{c}(\wp) + \tfrac{T}{t} \sum_{\wp \in \mathcal{P}_i(t)} \mathrm{P}(\wp) \log\left(\frac{\mathrm{P}(\wp)}{\tilde{\pi}^{\mathrm{ref}}(\wp)}\right) \\ &+ \mu\left(\sum_{\wp \in \mathcal{P}_i(t)} \mathrm{P}(\wp) - 1\right) + \sum_{l \in \mathcal{C}} \lambda_l\big(\bar{n}_l(t) - \sigma_l\big)\end{aligned} \tag{25}$$

where the Lagrange parameters $\{\lambda_k\}_{k\in\mathcal{C}}$ are all nonnegative in the case of inequality constraints and there is a Lagrange parameter associated with each constrained node. Now, recall that $\bar{n}_l(t) = \frac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \eta(l \in \wp)\,\mathrm{P}(\wp)$ where $\eta(l \in \wp)$ represents the number of times node $k$ appears on path $\wp$ of length $t$. Therefore, note that the objective function to be minimized with respect to the discrete probabilities is convex and the inequality constraints are all linear.

Let us develop the last term of this expression. From Eq. (6) and $\bar{n}_{kl}(t) = \frac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \eta\big((k,l)\in\wp\big)\,\mathrm{P}(\wp)$, we obtain

$$\begin{aligned}\sum_{l\in\mathcal{C}} \lambda_l\big(\bar{n}_l(t) - \sigma_l\big) &= \sum_{l\in\mathcal{C}} \lambda_l\left(\sum_{k\in\mathcal{P}red(l)} \bar{n}_{kl}(t) + \tfrac{\delta_{il}}{t} - \sigma_l\right) \\ &= \sum_{(k,l)\in\mathcal{E}} \bar{n}_{kl}(t)\,\delta(l\in\mathcal{C})\lambda_l + \sum_{l\in\mathcal{C}} (\tfrac{\delta_{il}}{t} - \sigma_l)\lambda_l \\ &= \tfrac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) \sum_{(k,l)\in\mathcal{E}} \eta\big((k,l)\in\wp\big)\,\delta(l\in\mathcal{C})\lambda_l \\ &\quad + \sum_{l\in\mathcal{C}} (\tfrac{\delta_{il}}{t} - \sigma_l)\lambda_l\end{aligned} \tag{26}$$

and $\delta(k \in \mathcal{C})$ is defined as 1 when node $k \in \mathcal{C}$ and 0 otherwise. Moreover, in this last equation, $\delta_{il}$ is a constant and thus the term $\delta_{il}/t$ disappears when $t \to \infty$. Because we will be interested in infinite-length walks, it can be neglected.

Injecting this last result in the Lagrange function of Eq. (25), using $\tilde{c}(\wp) = \sum_{(k,l)\in\mathcal{E}} \eta\big((k,l)\in\wp\big)\,c_{kl}$, and rearranging the terms yields

$$\begin{aligned}\mathscr{L}(\mathbb{P}, \boldsymbol{\lambda}) = &\tfrac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) \underbrace{\left(\tilde{c}(\wp) + \sum_{(k,l)\in\mathcal{E}} \eta\big((k,l)\in\wp\big)\,\delta(l\in\mathcal{C})\lambda_l\right)}_{\text{augmented cost } \tilde{c}'(\wp) \text{ cumulated on path } \wp} \\ &+ \tfrac{T}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\log\left(\frac{\mathrm{P}(\wp)}{\tilde{\pi}^{\mathrm{ref}}(\wp)}\right) + \mu\left(\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) - 1\right) - \sum_{l\in\mathcal{C}} \sigma_l\,\lambda_l \\ = &\tfrac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) \sum_{(k,l)\in\mathcal{E}} \eta\big((k,l)\in\wp\big) \underbrace{\big(c_{kl} + \delta(l\in\mathcal{C})\,\lambda_l\big)}_{\text{augmented costs } c'_{kl}} \\ &+ \tfrac{T}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\log\left(\frac{\mathrm{P}(\wp)}{\tilde{\pi}^{\mathrm{ref}}(\wp)}\right) + \mu\left(\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) - 1\right) - \sum_{l\in\mathcal{C}} \sigma_l\,\lambda_l \\ = &\underbrace{\tfrac{1}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\tilde{c}'(\wp) + \tfrac{T}{t}\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp)\log\left(\frac{\mathrm{P}(\wp)}{\tilde{\pi}^{\mathrm{ref}}(\wp)}\right)}_{\text{objective function based on augmented costs}} \\ &+ \mu\left(\sum_{\wp\in\mathcal{P}_i(t)} \mathrm{P}(\wp) - 1\right) - \sum_{l\in\mathcal{C}} \sigma_l\,\lambda_l\end{aligned} \tag{27}$$

Note that in this reasoning, the costs $c_{kl}$ were redefined into *augmented costs*,

$$c'_{kl} \triangleq \begin{cases} c_{kl} + \lambda_l & \text{when node } l \in \mathcal{C} \\ c_{kl} & \text{otherwise} \end{cases} \tag{28}$$

**Algorithm 1** Optimal policy for the biased random walk with capacity constraints on node visit rate.

**Input:**
- A weighted directed, strongly connected, graph $G$ containing $n$ nodes.
- The $n \times n$ transition probability matrix $\mathbf{P}_{\text{ref}}$ of the natural random walk on $G$. The transition matrix must be stochastic and regular.
- The $n \times n$ cost matrix $\mathbf{C}$ associated with $G$ defining non-negative costs of transitions.
- The set of constrained nodes $\mathcal{C}$.
- The set of non-negative capacity constraints on node visit rate, $\{\sigma_k\}$, defined on the constrained nodes set $\mathcal{C}$.
- The inverse temperature parameter $\beta > 0$.
- The gradient ascent step $\alpha > 0$.

**Output:**
- The $n \times n$ policy of the biased random walk with capacity constraints, defined by the transition matrix $\mathbf{P}^\star$.

$\boldsymbol{\lambda} \leftarrow \mathbf{0}$ ▷ initialize a $n \times 1$ Lagrange parameters column vector
$\mathbf{C}' \leftarrow \mathbf{C}$ ▷ initialize the augmented costs matrix
**repeat** ▷ main iteration loop
  $\mathbf{W} \leftarrow \mathbf{P}_{\text{ref}} \circ \exp[-\beta \mathbf{C}']$ ▷ recompute $\mathbf{W}$ matrix from augmented costs (elementwise exponential and multiplication $\circ$)
  $(\mathbf{u}_1, \lambda_1) \leftarrow \text{abs}(\text{eigs}(\mathbf{W}, 1))$ ▷ extract the right, dominant, real eigenvector and eigenvalue of $\mathbf{W}$
  $(\mathbf{v}_1, \lambda_1) \leftarrow \text{abs}(\text{eigs}(\mathbf{W}^{\mathrm{T}}, 1))$ ▷ extract the left, dominant, real eigenvector and eigenvalue of $\mathbf{W}$
  $\mathbf{u}_1 \leftarrow \mathbf{u}_1 / \sqrt{\mathbf{u}_1^{\mathrm{T}} \mathbf{v}_1}$; $\mathbf{v}_1 \leftarrow \mathbf{v}_1 / \sqrt{\mathbf{u}_1^{\mathrm{T}} \mathbf{v}_1}$ ▷ to ensure that the two eigenvectors are properly normalized
  $\bar{\mathbf{n}} \leftarrow \mathbf{u}_1 \circ \mathbf{v}_1$ ▷ visit rate to each node (= stationary distribution)
  **for all** $l \in \mathcal{C}$ **do** ▷ gradient ascent: update Lagrange parameters associated with constrained nodes
    $\lambda_l \leftarrow \max(\lambda_l + \alpha(\bar{n}_l - \sigma_l), 0)$
  **end for**
  $\mathbf{C}' \leftarrow \mathbf{C}' + \mathbf{e}\boldsymbol{\lambda}^{\mathrm{T}}$ ▷ recompute augmented costs. $\mathbf{e}$ is a column vector full of 1's
**until** convergence
$\bar{\mathbf{N}} \leftarrow \mathbf{W} \circ (\mathbf{v}_1 \mathbf{u}_1^{\mathrm{T}})/\lambda_1$ ▷ compute the passage rate through each edge
$\mathbf{P}^\star \leftarrow \mathbf{Diag}\big(\mathbf{e} \div (\bar{\mathbf{N}}\mathbf{e})\big)\, \bar{\mathbf{N}}$ ▷ compute optimal policy; $\div$ is the elementwise division between the two vectors
**return** $\mathbf{P}^\star$

and $\mathbf{C}' = (c'_{kl})$ will be the matrix containing these augmented costs. As frequently observed, the Lagrange parameters have a heuristic interpretation similar to the dual variables in linear programming, the cost to pay on each node in order to satisfy the constraints [76]. This commonly appears in standard network flow problems [55].

As in [25], the important point here is that the resulting Lagrange function (27) has *exactly the same form* as Eq. (2), except the last extra term which does not depend on the probability distribution. Therefore, the computation of the path probabilities can be done in the same way as for the standard procedure (without inequality constraints), except the fact that costs are replaced by the *augmented costs* in Eq. (28). This property will be exploited in the next section for computing the optimal policy satisfying the constraints.

### 3.2. The optimal policy satisfying capacity constraints

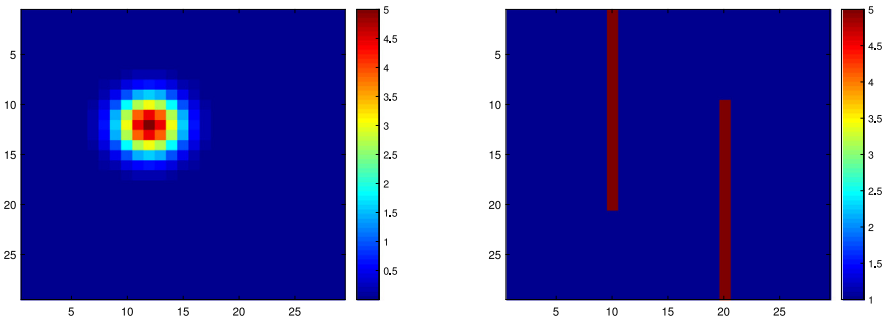
The procedure for computing the optimal policy with capacity constraints on node visits is inspired by [25] and standard convex optimization results (see, e.g., [74–76]). As in this work, we take advantage of the fact that, in our bag-of-paths formulation of the problem, the Lagrange dual function and its gradient are easy to compute (see [77] for a similar derivation in the context of supervised classification). We therefore apply a common procedure (closely related to the Arrow–Hurwicz–Uzawa algorithm [78]) aiming at (1) computing the dual Lagrange function which depends on the Lagrange parameters, and (2) performing a gradient ascent on the dual function (see [25] for details). This leads to a two-steps algorithm [25], which is summarized now. After initializing the augmented costs to the edge costs, it iterates the two following steps:

1. Compute the expected visit rate from the augmented costs through Eqs. (28), (8) and (17) while considering the Lagrange parameters as fixed. The expected visit rate is then used in the second step.
2. Update the Lagrange parameters through a gradient ascent-like technique involving the expected visit rate (see Eq. (29) below),

until convergence. The optimal policy is finally obtained from Eq. (18). Before developing the final algorithm, let us provide a few useful information about these two steps.

#### 3.2.1. Computing the expected visit rate

This step simply aims at finding the expected visit rate to constrained nodes, $\bar{n}_k(\infty)$, by evaluating Eq. (17) in which the costs $c_{ij}$ are replaced by the augmented costs $c'_{ij}$ provided by Eq. (28). Therefore, the problem studied in Section 2 must be solved iteratively.



(a) Gaussian obstacle (b) Double wall

**Fig. 1.** Two examples of cost matrix associated with the lattice.

*3.2.2. Updating the Lagrange parameters*

We maximize the dual function by a simple gradient ascent-based technique developed by Rockafellar [79]. Indeed, it turns out that the partial derivative of the dual function with respect to $\lambda_k$ is simply $(\bar{n}_k(\infty) - \sigma_k)$ [25]. Hence, the updating rule (with gradient step $\alpha$) for the Lagrange parameters associated to constrained nodes is

$$\lambda_k \leftarrow \max(\lambda_k + \alpha(\bar{n}_k(\infty) - \sigma_k), 0) \text{ for all } k \in \mathcal{C} \tag{29}$$

which converges in the convex case as far as the chosen $\alpha$ parameter is positive, not too large, and the problem is feasible (see [79] for details). More efficient convex optimization techniques could be used as well, but this simple method was sufficient for our application, although slow and requiring some parameter tweaking on $\alpha$.

*3.3. The final algorithm*

The two previous steps are iterated until convergence. The resulting algorithm computing the quantities in matrix form is presented in Algorithm 1. From Eqs. (17) and (18), the transition probability matrix of the biased random walk, $\mathbf{P}^\star$, is computed from $\bar{\mathbf{N}}$ by normalizing each of its rows so that they all sum to 1 (line 15 in the algorithm).

## 4. Some illustrative examples

A series of simulations has been performed in order to investigate the behavior of the method. The first series illustrates the influence of a cost matrix $\mathbf{C}$ on walks on a graph $G$ as well as the influence of the diffusion coefficient $\beta$. We then show the impact of a visit constraint on the stationary distribution.

For this experiment, $G$ will be referring to a two-dimensional undirected square lattice forming a Markov chain. Each node is associated with a cost, represented in the cost matrix $\mathbf{C}$. We consider that each node of $G$ has four neighbors, except for the border nodes (3 neighbors) and the corner nodes (2). This lattice is then transformed into an affinity matrix $\mathbf{A}$ respecting the topology of the lattice $G$ before the biased random walk is run to obtain a biased transition probability matrix $\mathbf{P}^\star$ (see Eq. (18)). Taking its left eigenvector corresponding to the dominant eigenvalue yields the stationary probability distribution on the set of nodes of $G$. As a result, it will be observed from the stationary distribution that the random walker indeed avoids visits to nodes associated with high costs.

In this context, two examples are considered, one simple uncentered Gaussian obstacle and another with two walls of higher cost. The cost matrices associated with these two graphs are represented in Fig. 1. For these two examples, we show the resulting stationary distribution without any visit constraint in the Figs. 2 for $\beta = 0.01$ and 3 for $\beta = 0.1$. Fig. 4 illustrates the stationary distribution when using $\beta = 0.1$ (as in Fig. 3) but now with a uniform visit constraint forcing the expected visit rate to each node to be below $2/n$, where $n$ represents the total number of nodes. Note that in all these heatplots, the scaling is relative to the difference between the largest and the lowest values that are represented.

## 5. Collaborative recommendation application

The previous section visually illustrates how the bias (and the node visit constraints) influences the behavior of the random walk on two small benchmark grids. Let us now further observe the impact of this bias in a practical application of a random walk. More precisely, we consider ItemRank [80–83], an algorithm designed for collaborative recommendation that relies on a random walk on items.

Please note that our objective is not to suggest a new recommendation algorithm that would outperform all pre-existing ones, but rather to provide an example of a concrete application with an explicit objective, rather than a purely theoretical illustration like in the previous section. The aim is therefore to observe how the addition of the bias can affect the resulting recommendations with respect to the objective.

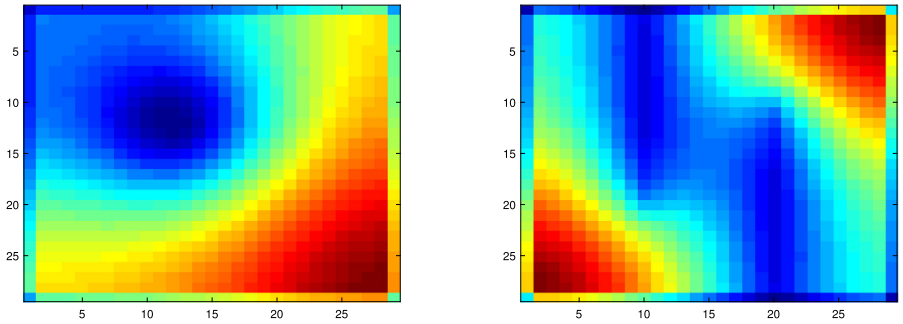


**Fig. 2.** Stationary distribution ($\beta = 0.01$) for the examples of Fig. 1.

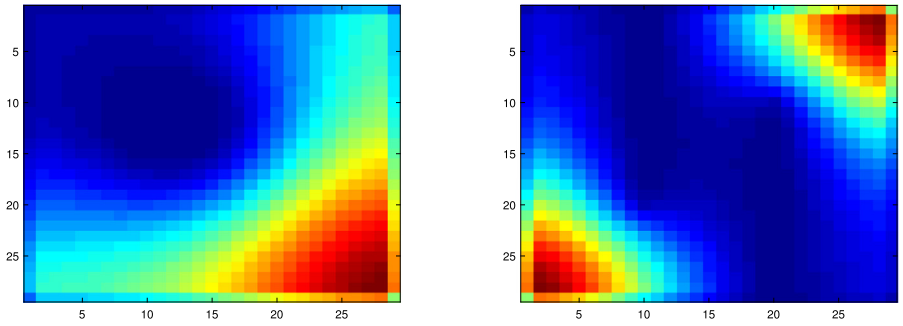


**Fig. 3.** Stationary distribution ($\beta = 0.10$) for the examples of Fig. 1.

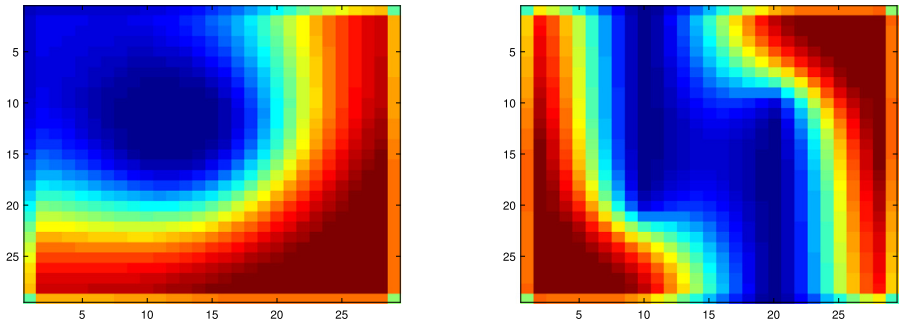


**Fig. 4.** Stationary distribution ($\beta = 0.10$) and uniform visit rate constraint $2/n$, for the examples of Fig. 1.

### 5.1. Context and objective

In order to bias the transition probabilities, the first step is to decide towards which nodes the transitions should be favored. Our approach is to bias the random walk transition probabilities towards items with a lower degree (less known and less popular), with the goal of improving the serendipity of the recommendations. Users are usually not interested in suggestions recommending highly popular items because they already know them (like blockbusters in movie recommendation).

The serendipity is a complex concept and represents a "good surprise" in recommendation. Although there is no consensus on the precise definition of serendipity, it is often described as composed of three components: relevance, novelty and unexpectedness [84].

Additionally, a second experiment that considers the introduction of node visit constraints is performed, with the goal of improving diversity in recommendations. All these concepts are further developed in Section 5.3.

### 5.2. Application of the biased random walk with restart

In the field of collaborative recommendation, the dataset often takes the form of a list of links between users and items, corresponding to the product purchases (transactions), and can therefore be represented in the form of a bipartite

graph (i.e. whose nodes can be split into two mutually exclusive subsets such that there does not exist any link between two nodes belonging to the same subset [32]). This is achieved by splitting the nodes into a set $\mathcal{U}$ of users and a set $\mathcal{I}$ of items that are available to the users. For each user $u$, we define $\mathcal{I}_u$ the subset of items that $u$ consumed and conversely for the items, $\mathcal{U}_i$ the subset of users that consumed an item $i$.

From $\mathbf{B}$, the $(|\mathcal{U}| \times |\mathcal{I}|)$ biadjacency matrix containing the links between users and items (binary values), a workable adjacency matrix between items can be computed using a projection [85] with stochastic complementation [86], as $\mathbf{A}_{\mathcal{I}} = \mathbf{B}^{\mathrm{T}}\mathbf{D}_{\mathcal{U}}^{-1}\mathbf{B}$ where $\mathbf{D}_{\mathcal{U}}$ is a diagonal matrix containing the degrees of the users.

#### 5.2.1. The standard ItemRank

Introduced by Gori and Pucci in the field of collaborative recommendation [80], the ItemRank model consists of a random walk with restart on a graph of products (adjacency matrix $\mathbf{A}_{\mathcal{I}}$) (see also [83,87]). The random walker starts on a product selected at random among those consumed by the current user who is waiting for some recommendations, then, at each time step, it has a probability $\alpha$ to continue the random walk and $(1-\alpha)$ to be teleported back onto one of the products consumed by the user. The nodes that are visited the most often among those that were not already consumed by the user are the ones that will be recommended.

Therefore, the random walk takes the following form for a user $u$

$$\mathbf{x}_u(t+1) = \alpha\mathbf{P}^{\mathrm{T}}\mathbf{x}_u(t) + (1-\alpha)\mathbf{v}_u \tag{30}$$

where $\mathbf{x}_u(t)$ represents the probability for the random walker to be located on each product at time $t$, $\mathbf{P}$ is the transition probability matrix computed from $\mathbf{A}_{\mathcal{I}}$, the adjacency matrix between items, and $\mathbf{v}_u$ is a personalization vector containing $|\mathcal{I}_u|^{-1}$ for each item rated by the user and 0 elsewhere. Eq. (30) yields the steady-state solution

$$\mathbf{x}_u = (1-\alpha)(\mathbf{I} - \alpha\mathbf{P}^{\mathrm{T}})^{-1}\mathbf{v}_u \tag{31}$$

where $\mathbf{I}$ is the identity matrix and the vector $\mathbf{x}_u$ contains the score of each item for user $u$.

#### 5.2.2. The unconstrained biased ItemRank

With the aim to improve the serendipity of the recommendation against the standard ItemRank, our model biases the random walk towards less consumed items. This is achieved by defining the cost matrix $\mathbf{C}$ in function of the degree of the nodes. The transition cost from any item to a specific item $j$ will be computed from a monotonic increasing transformation of the degree of node $j$. The steady-state solution of the biased ItemRank equation for a user $u$ then becomes Eq. (31) with $\mathbf{P}$ replaced by $\mathbf{P}^{\star}$, the optimal transition probability matrix computed from Eq. (18).

### 5.3. Experimental setup

#### 5.3.1. Dataset

As part of this paper, the experiments were conducted on the 100k MovieLens (see https://grouplens.org) benchmark dataset (1998). It contains a total of 100,000 links between approximately 1000 users and 1700 movies, forming an unweighted undirected bipartite graph (ratings are not considered). The users composing this dataset have been selected so that each of them rated at least 20 movies.

#### 5.3.2. Performance measures

The performance of the algorithms will be compared through a 10-fold cross-validation, using measures quantifying relevance, novelty, unexpectedness, and diversity. The scores for these measures are computed for each fold independently and averaged across the 10 folds afterwards.

*Relevance measures.* Here are the two relevance measures, quantifying the quality of the recommendations, that will be reported.

Recall. This score represents, for any user $u$, the proportion of the links in the `test set` associated to $u$ that are retrieved among the top 20 (recall at 20) movies from the ranked list computed from the `training set`.

Degree of Agreement. When links of a user are hidden, one would reasonably expect to find these links well positioned in the ranked list. The degree of agreement consists in a variant of Somer's $D$ [88]. It compares the ranked list returned by the algorithm to the optimal list in which all the movies in the `test set` (that were actually watched) are ranked before the rest. It scales from 0 (complete disagreement) to 1 (complete agreement). The score is therefore computed for the entire ranked list, instead of only considering the top $k$ items like the recall. For more details, see [89].

*Novelty.* In this work, we consider an item to be novel to a given user if it is unknown to him/her [90–93]. It is therefore necessary to estimate the probability that a user has seen or heard about an item. The novelty can be assessed by considering that the probability that an item $i$ has been noticed reflects a factor of popularity [84,94]

$$\mathrm{nov}(i) = 1 - \frac{|\mathcal{U}_i|}{|\mathcal{U}|} \tag{32}$$

where $|\mathcal{U}_i|$ is the number of users having purchased item $i$.

*Unexpectedness.* The definition of unexpectedness fluctuates in literature [84]. We consider the approach suggested in [90,95,96] that proposes to evaluate the unexpectedness of recommendations by comparing them with those from a benchmark, primitive, recommendation model that makes expected recommendations. Let $\mathcal{I}_u^R$ be the set of items recommended for a user $u$ and $\mathcal{I}_u^P$ the set of items recommended by the primitive model for this same user. The unexpectedness of the set of recommendations $\mathcal{I}_u^R$ is computed as

$$\text{unexp}(\mathcal{I}_u^R) = \frac{|\mathcal{I}_u^R \setminus \mathcal{I}_u^P|}{|\mathcal{I}_u^R|} \tag{33}$$

In the present work, we consider, as primitive model that makes expected recommendations, the algorithm that recommends the items the most similar to those already consumed by the user, that is, the top-$k$ items maximizing

$$\arg\max_{i \in \mathcal{I} \setminus \mathcal{I}_u} \frac{1}{|\mathcal{I}_u|} \sum_{j \in \mathcal{I}_u} \text{sim}(i,j) \tag{34}$$

where $\text{sim}(i,j)$ is the cosine similarity.[13]

*Diversity.* Diversity differs from serendipity in the sense that recommended items are compared between each other rather than with respect to the user profile [94]. In this work, we consider diversity in terms of catalog coverage [90,91], which corresponds to the proportion of items that appears in the top-20 recommendations of at least one user. It is thus computed as

$$\text{coverage} = \frac{|\bigcup_{u \in \mathcal{U}} \mathcal{I}_u^R|}{|\mathcal{I}|} \tag{35}$$

For a given fold, except for the coverage, which is global, all the measures are computed user-wise, then averaged across all the users.

### 5.4. First experiment: Unconstrained biased ItemRank

#### 5.4.1. Results

The parameter $\alpha$ in Eqs. (30) and (31) represents the probability, at any time step, to continue the random walk. Its value has been arbitrarily fixed at $\alpha = 0.1$ to put emphasis on consumed items.

The two remaining inputs of the biased ItemRank are the cost matrix $\mathbf{C}$ and the parameter $\beta$ controlling the exploration. The higher the value of $\beta$, the larger the bias of the random walk in favor of less popular items (in terms of degree). As already discussed, the transition costs are computed from the degree of the destination node. More specifically, three case are considered: $c_{ij} = d_j$, $c_{ij} = \sqrt{d_j}$, and $c_{ij} = \log(d_j + 1)$.

For each of these cost matrices, two figures depict the evolution of the different metrics in function of $\beta$. The first figure uses an exponential scale to identify key values for $\beta$ and the second one provides smooth curves around these key values using a linear scale.

The results for the degree, square root of the degree and logarithm of the degree are displayed in Fig. 5, Fig. 6, and Fig. 7 respectively.

#### 5.4.2. Discussion of the results

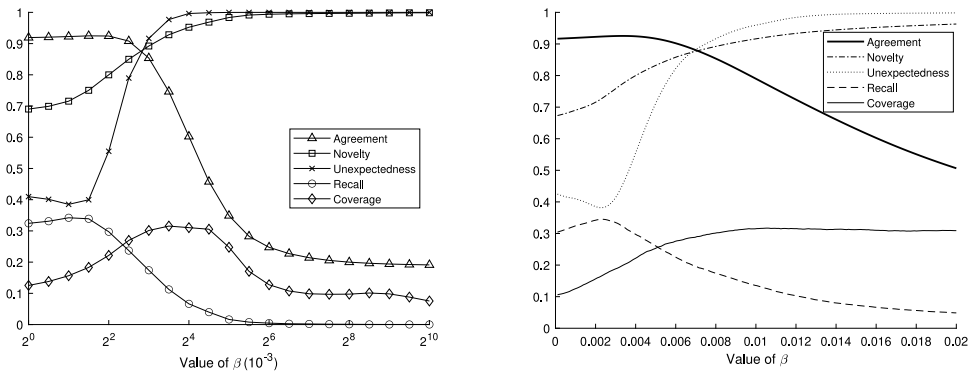
It can be observed from Figs. 5, 6, and 7 that the curves for the different measures tend to follow similar patterns for the three cost matrices considered. As one would reasonably expect, the relevance (recall and degree of agreement) tends to drop and the unexpectedness and novelty to increase when the random walk gets strongly biased towards unpopular items. The balance between these measures is controlled by the intensity of the parameter $\beta$.

Interestingly enough, the biased ItemRank led to an increase in the value of the recall (and agreement) before an effective decline, as can be observed in all three cases. The maximum value that could be achieved are 0.3405, 0.3547, and 0.3434 for the degree, square root of the degree, and logarithm of the degree, respectively. Even though it is not the objective of the biased ItemRank, a notable improvement with respect to the initial recall (0.3007) from the standard ItemRank can be observed.

A plausible explanation for this last result, that can seem counter-intuitive at first sight, comes from the mechanics of the ItemRank and, more generally, of random walks on graphs. It is known that hubs on graphs tend to naturally attract random walkers. In a collaborative recommendation context, it results in hub items being recommended, not because they are relevant to the user profile, but because these items attract the random walker. It has been shown that this "hubness effect" can be detrimental in some situations [97,98]. In our experiment, the significant increase in recall when a carefully chosen bias in favor of lower degree nodes is applied, is probably due to the harmfulness potential of this hubness effect.

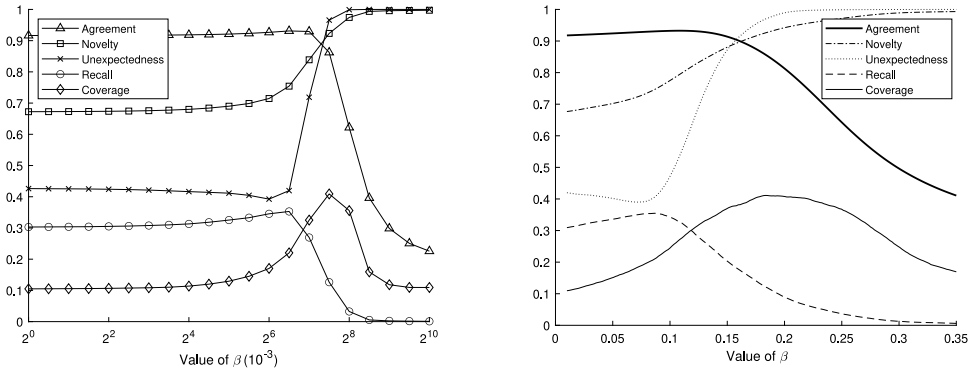
As regards diversity, the catalog coverage behaves comparably to relevance measures with an evolution split into a growing phase followed by a decline. This effect appears natural with respect to the biased ItemRank mechanics,

[13] The similarity is computed between the columns of the biadjacency matrix.



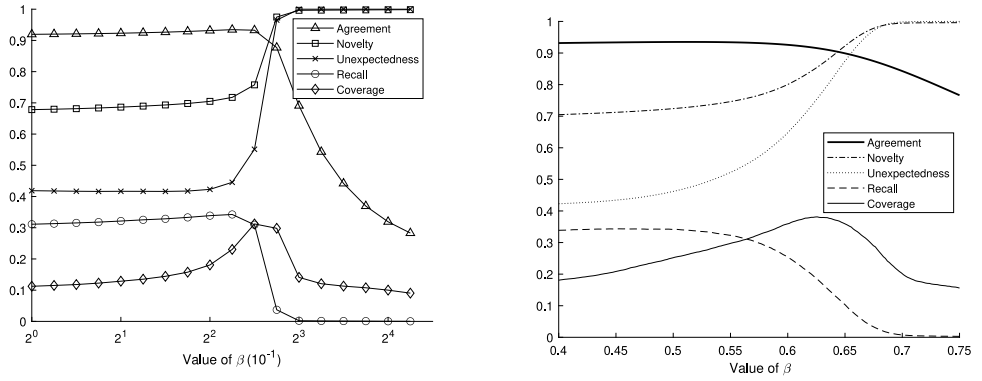

(a) Exponential $\beta$ scale. (b) Linear $\beta$ scale.

**Fig. 5.** Average (over a 10-fold cross-validation) of the recall, degree of agreement, novelty, unexpectedness, and diversity in function of $\beta$ by using the degree as cost and $\alpha = 0.1$.




(a) Exponential $\beta$ scale. (b) Linear $\beta$ scale.

**Fig. 6.** Average (over a 10-fold cross-validation) of the recall, degree of agreement, novelty, unexpectedness, and diversity in function of $\beta$ by using the square root of the degree as cost and $\alpha = 0.1$.

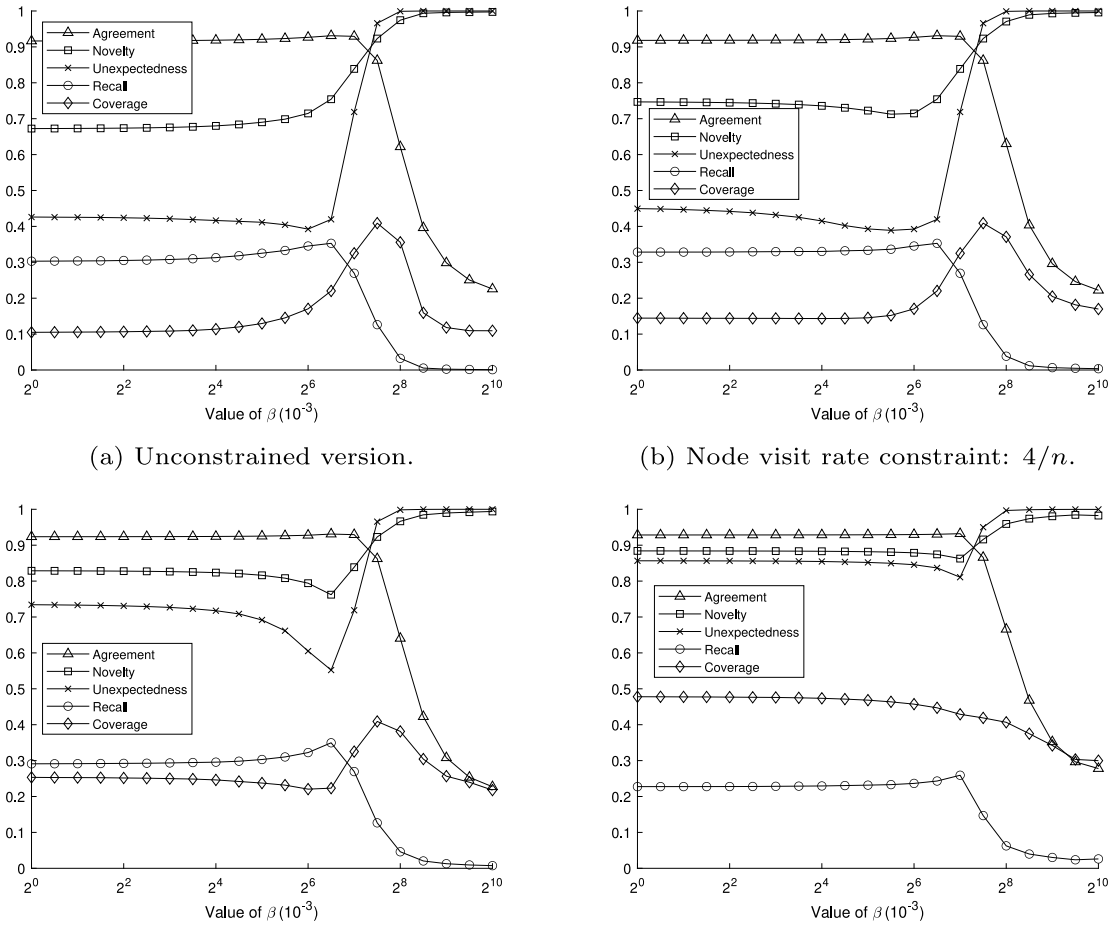



(a) Exponential $\beta$ scale. Mind the difference in the scale with respect to Figures 5a and 6a.

(b) Linear $\beta$ scale.

**Fig. 7.** Average (over a 10-fold cross-validation) of the recall, degree of agreement, novelty, unexpectedness, and diversity in function of $\beta$ by using the logarithm of the degree as cost and $\alpha = 0.1$.

since alleviating the hub-effect improves the diversity by avoiding systematic recommendations of popular items. On the contrary, a bias too strong harms diversity by limiting recommendations to extremely unpopular items. There is, however, a clear offset between the peak in recall and in diversity.




(a) Unconstrained version.

(b) Node visit rate constraint: $4/n$.

(c) Node visit rate constraint: $2.5/n$.

(d) Node visit rate constraint: $1.5/n$.

**Fig. 8.** Average (over a 10-fold cross-validation) of the recall, degree of agreement, novelty, unexpectedness, and diversity in function of $\beta$ by using the square root of the degree as cost and $\alpha = 0.1$, using a constrained random walk (see the text for details).

### *5.5. Second experiment: Constrained biased ItemRank*

This second experiment aims to improve the diversity of the recommendations by making use of the node visit rate constraint. The objective is thus to use the constraint on node visit rate to force the random walker to spread its visits across the graph in order to increase the catalog coverage.

We opted for a uniform threshold for each node, constraining the stationary distribution to each node to $\frac{a}{n}$ with $a > 1$ where $n$ is the number of nodes. Figs. 8(b), 8(c), and 8(d) depict the evolution of the different measures for $a = 4$, 2.5, and 1.5 respectively, with transition costs chosen as the square root of the degree. Fig. 8(a) is a duplicate of Fig. 6(a) for ease of comparison.

The addition of a node visit constraint effectively increases the coverage but comes nonetheless at the expense of recall. However, the constraint seems to have limited or no notable impact on the degree of agreement, which indicates that the global product order of the ranked list is still pertinent.

As regards the novelty and unexpectedness of the recommendations, a notable increase can be observed when the bias is small. Recall that, in the present paper, the novelty quantifies the unpopularity and the unexpectedness indirectly evaluates the dissimilarity with respect to the products already consumed by the user. With a small bias, since the random walker is most likely to visit hubs and items close to those already consumed by the user, the addition of capacity constraint mostly affects these items, therefore effectively increasing these two measures.

When a strong bias is used, apart from a better coverage and a small alleviation in the recall drop, the addition of a node visit rate constraint does not seem to have any significant impact.

### 5.6. Experiments summation

The biased ItemRank globally behaves as expected, by progressively tilting the recommendations towards items with a lower degree. This is reflected by the increase in novelty and unexpectedness along with a decrease in the recall. On the other hand, as regards the degree of agreement, Figs. 5(b), 6(b), and 7(b) show that it is possible to pick a bias that provides high novelty and unexpectedness, while keeping a good degree of agreement, indicating that the overall item order within the ranked lists remains relevant.

This application offers the advantage of having a tunable parameter $\beta$ allowing for customizable balance between relevance, novelty and unexpectedness depending on one's needs, up to extremely novel recommendations. It comes nonetheless with the disadvantage of the sensitivity of this parameter $\beta$ to the cost matrix.

With regard to the constrained version, the experiments have highlighted that a small bias combined with a node visit rate constraint seem to provide a good balance. The coverage, novelty, and unexpectedness can be improved significantly, while having limited impact on the recall and almost no impact on the degree of agreement.

However, the constrained approach makes use of the Algorithm 1, which requires to compute the dominant eigenvector and eigenvalue iteratively until convergence of the biased transition probability matrix, before proceeding to the computation of the ItemRank scores. Therefore, the constrained version of the biased random walk is significantly more computationally heavy than its unconstrained counterpart, which only requires to compute the dominant eigenvector once.

Finally, it must be stressed that all these results were obtained on a single dataset and should therefore be confirmed on additional data sources; this will be part of our future work.

## 6. Conclusion

Our main objective in this paper was to study two different ways of designing a biased random walk on a graph. The first approach is based on a bag-of-paths model considering walks of increasing length biased towards low-cost paths. This model allows to define random walks avoiding nodes having some undesirable properties by setting the cost depending on these properties, and then minimizing expected cost rate under relative entropy rate constraint. It is shown that all quantities of interest, already derived from a different perspective in [11,12,14,15], can be obtained from the free energy rate function of the system, which can easily be computed in closed form from the adjacency and the custom cost matrices by standard matrix operations. The balance between exploration and exploitation is controlled by a temperature parameter biasing gradually the walk towards low-cost walks. A concrete application of this model would be to define a cost-sensitive ItemRank based on the biased random walk.

The second approach, also based on the bag-of-paths framework, considers capacity constraints on the node visit rates. The derived algorithm, based on Lagrangian duality, assigns additional virtual costs to constrained nodes in order to enforce the capacity constraints and is iterative, solving the policy problem until convergence. The result is a new biased random walk minimizing expected cost rate regularized by relative entropy rate (as in the previous approach) while, at the same time, satisfying the provided capacity constraints.

The illustrative application in collaborative recommendation, in which the ItemRank is biased in favor of low-degree items, shows than this bias can be used to obtain a desired balance between relevance, novelty, and unexpectedness of recommendations.

With the unintended increase in recall, our first experiment also suggests that our method might be used to counteract the hubness effect. Therefore, future work will study the application of the introduced techniques to problems where the hubness effect is detrimental. Indeed it has been shown that, in some situations, hubness can alter significantly the performance of graph data analysis, and more generally machine learning, techniques [97,99–101], as probably observed in the present experiment.

Moreover, following [102], we will further investigate the relationships between the discrete bag-of-paths/randomized shortest paths frameworks and their continuous space–time counterparts; especially models involving either Brownian bridges [103,104], Schrödinger bridges (which were extended on discrete structures in [40–42]), or importance sampling-related techniques for dealing with rare events [105–109], usually involving Doob's transform [110]. Indeed, the discrete probability of a randomized shortest path $\wp$ can also be understood as the probability of following this path $\wp$ given that the random walker survives and thus reaches the target node during a "killed" random walk on the graph with transient, sub-stochastic, probability matrix $\mathbf{W}$ [19,32]. This interpretation is related to both Brownian bridges, Schrödinger bridges, and importance sampling techniques.

In addition, as suggested by one of the reviewers, another interesting topic for future work would be to go beyond the hypothesis of additive costs along paths (see Section 2.1.1).

## CRediT authorship contribution statement

**Pierre Leleux:** Methodology, Software, Investigation, Writing – original draft, Writing – review & editing. **Sylvain Courtain:** Methodology, Writing – original draft. **Kevin Françoisse:** Conceptualization, Writing – original draft. **Marco Saerens:** Conceptualization, Methodology, Software, Writing – original draft, Writing – review & editing.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

First, many thanks to the reviewers of this paper whose remarks allowed us to improve significantly the paper. We also thank Prof. Jean-Charles Delvenne, Prof. François Fouss, Prof. Frédéric Vrins, and Dr. Ilkka Kivimäki for reading our manuscript and providing helpful references and remarks. This work was partially supported by projects funded by InnovIris (Brussels region) as well as by the Walloon region, Belgium. We thank these institutions for giving us the opportunity to conduct both fundamental and applied research.

## Appendix

In this appendix, the expression for computing the partition function is derived in detail. Let us start from Eq. (4) and use the definition of $\mathcal{P}_i(t) = \cup_{k=1}^{n} \mathcal{P}_{ik}(t)$. Recall that $\mathcal{P}_{ik}(t)$ is the set of paths connecting $i$ to $k$ in exactly $t$ steps or hops, with the $\mathcal{P}_{ik}(t)$ sets being mutually exclusive. First, observe that, for a path length $t = 1$, only one transition from node $i$ is allowed and in this case $\mathcal{Z}_i(1) = \sum_{k=1}^{n} p_{ik}^{\text{ref}} \exp[-\beta c_{ik}] = \sum_{k=1}^{n} w_{ik} = \mathbf{e}_i^{\mathrm{T}}\mathbf{W}\mathbf{e}$.

For $t > 1$, we split the sum over all paths of length $t$ from $i$ to $k$ into a length-one path (a single transition) from $i$ to a successor node $j \in \mathcal{S}ucc(i)$ and a $(t-1)$-length path from $j$ to $k$,

$$
\begin{aligned}
\mathcal{Z}_i(t) &= \sum_{\wp(t)\in\mathcal{P}_i(t)} \tilde{\pi}^{\text{ref}}(\wp(t)) \exp\left[-\beta\tilde{c}(\wp(t))\right] \\
&= \sum_{k=1}^{n} \sum_{\wp_{ik}(t)\in\mathcal{P}_{ik}(t)} \tilde{\pi}^{\text{ref}}(\wp_{ik}(t)) \exp\left[-\beta\tilde{c}(\wp_{ik}(t))\right] \\
&= \sum_{k=1}^{n} \sum_{j=1}^{n} \underbrace{p_{ij}^{\text{ref}} \exp\left[-\beta c_{ij}\right]}_{w_{ij}} \times \\
&\quad \sum_{\wp_{jk}(t-1)\in\mathcal{P}_{jk}(t-1)} \tilde{\pi}^{\text{ref}}(\wp_{jk}(t-1)) \exp\left[-\beta\tilde{c}(\wp_{jk}(t-1))\right] \\
&= \sum_{j=1}^{n} w_{ij} \underbrace{\left( \sum_{k=1}^{n} \sum_{\wp_{jk}(t-1)\in\mathcal{P}_{jk}(t-1)} \tilde{\pi}^{\text{ref}}(\wp_{jk}(t-1)) \exp\left[-\beta\tilde{c}(\wp_{jk}(t-1))\right] \right)}_{\mathcal{Z}_j(t-1)} \\
&= \sum_{j=1}^{n} w_{ij} \mathcal{Z}_j(t-1)
\end{aligned}
$$

and we obtain $\boldsymbol{\mathcal{z}}(t) = \mathbf{W}\boldsymbol{\mathcal{z}}(t-1)$ where $\boldsymbol{\mathcal{z}}(t)$ is a column vector containing the $\mathcal{Z}_i(t)$. By induction, Eq. (9) follows,

$$
\mathcal{Z}_i(t) = \mathbf{e}_i^{\mathrm{T}} \underbrace{\mathbf{W}\mathbf{W}\cdots\mathbf{W}}_{t \text{ times}} \mathbf{e} = \mathbf{e}_i^{\mathrm{T}}(\mathbf{W}^t)\mathbf{e}
$$

## References

[1] O. Ibe, Elements of Random Walk and Diffusion Processes, Wiley, 2013.
[2] N. Masuda, M. Porter, R. Lambiotte, Random walks and diffusion on networks, Phys. Rep. 716–717 (2017) 1–58.
[3] J. Rudnick, G. Gaspari, Elements of the Random Walk, Cambridge University Press, 2004.
[4] P. Sarkar, A. Moore, Random walks in social networks and their applications: a survey, in: C. Aggarwal (Ed.), Social Network Data Analytics, Springer, 2011, pp. 43–77.
[5] Z. Wang, L. Guo, S. Wang, L. Chen, H. Wang, Review of random walk in image processing, Arch. Comput. Methods Eng. 26 (1) (2017) 17–34.
[6] F. Xia, J. Liu, H. Nie, Y. Fu, L. Wan, X. Kong, Random walks: A review of algorithms and applications, IEEE Trans. Emerg. Top. Comput. Intell. (2019) 1–13.
[7] J.R. Norris, Markov Chains, Cambridge University Press, 1997.
[8] S. Ross, Introduction To Probability Models, tenth ed., Academic Press, 2010.
[9] H.M. Taylor, S. Karlin, An Introduction to Stochastic Modeling, third ed., Academic Press, 1998.
[10] L. Page, S. Brin, R. Motwani, T. Winograd, The Pagerank Citation Ranking: bringing Order to the Web, Technical report 1999-66, Computer Science Department, Stanford InfoLab, Stanford University, USA, 1998.
[11] J.-C. Delvenne, A.-S. Libert, Centrality measures and thermodynamic formalism for complex networks, Phys. Rev. E 83 (4) (2011) 046117.
[12] R. Lambiotte, J.-C. Delvenne, M. Barahona, Random walks, Markov processes and the multiscale modular organization of complex networks, IEEE Trans. Netw. Sci. Eng. 1 (2) (2014) 76–90.

[13] E. Todorov, Linearly-solvable Markov decision problems, in: Advances in Neural Information Processing Systems 19, NIPS 2006, MIT Press, 2007, pp. 1369–1375.
[14] W. Parry, Intrinsic Markov chains, Trans. Amer. Math. Soc. 112 (1) (1964) 55–66.
[15] D. Ruelle, Thermodynamic Formalism, second Ed., Addison-Wesley, 1978.
[16] E.T. Jaynes, Information theory and statistical mechanics, Phys. Rev. 106 (1957) 620–630.
[17] M. Kardar, Statistical Physics of Particles, Cambridge University Press, 2007.
[18] L.E. Reichl, A Modern Course in Statistical Physics, second ed., Wiley, 1998.
[19] K. Francoisse, I. Kivimaki, A. Mantrach, F. Rossi, M. Saerens, A bag-of-paths framework for network data analysis, Neural Netw. 90 (2017) 90–111.
[20] A. Mantrach, L. Yen, J. Callut, K. Francoise, M. Shimbo, M. Saerens, The sum-over-paths covariance kernel: A novel covariance between nodes of a directed graph, IEEE Trans. Pattern Anal. Mach. Intell. 32 (6) (2010) 1112–1126.
[21] M. Saerens, Y. Achbany, F. Fouss, L. Yen, Randomized shortest-path problems: Two related models, Neural Comput. 21 (8) (2009) 2363–2404.
[22] L. Yen, A. Mantrach, M. Shimbo, M. Saerens, A family of dissimilarity measures between nodes generalizing both the shortest-path and the commute-time distances, in: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '08, 2008, pp. 785–793.
[23] T. Akamatsu, Cyclic flows, Markov process and stochastic traffic assignment, Transp. Res. B 30 (5) (1996) 369–386.
[24] Z. Burda, J. Duda, J. Luck, B. Waclaw, Localization of the maximal entropy random walk, Phys. Rev. Lett. 102 (16) (2009) 160602.
[25] S. Courtain, P. Leleux, I. Kivimaki, G. Guex, M. Saerens, Randomized shortest paths with net flows and capacity constraints, Inform. Sci. 556 (2021) 341–360.
[26] G. Guex, I. Kivimäki, M. Saerens, Randomized optimal transport on a graph: framework and new distance measures, Netw. Sci. 7 (1) (2019) 88—122.
[27] F. Bavaud, G. Guex, Interpolating between random walks and shortest paths: A path functional approach, in: K. Aberer, A. Flache, W. Jager, L. Liu, J. Tang, C. Guéret (Eds.), Proceedings of the 4th International Conference on Social Informatics, SocInfo '12, in: Lecture Notes in Computer Science, vol. 7710, Springer, 2012, pp. 68–81.
[28] I. Kivimäki, B. Lebichot, J. Saramäki, M. Saerens, Two betweenness centrality measures based on randomized shortest paths, Scientific Reports 6 (2016) srep19668.
[29] R. Dial, A probabilistic multipath assignment model that obviates path enumeration, Transp. Res. 5 (1971) 83–111.
[30] Y. Oyama, E. Hato, Prism-based path set restriction for solving Markovian traffic assignment problem, Transp. Res. B 122 (2019) 528–546.
[31] I. Kivimäki, M. Shimbo, M. Saerens, Developments in the theory of randomized shortest paths with a comparison of graph node distances, Physica A 393 (2014) 600–616.
[32] F. Fouss, M. Saerens, M. Shimbo, Algorithms and Models for Network Data and Link Analysis, Cambridge University Press, 2016.
[33] L. Arnold, V.M. Gundlach, L. Demetrius, Evolutionary formalism for products of positive random matrices, Ann. Appl. Probab. 4 (3) (1994) 859–901.
[34] F. Biondi, A. Legay, B.F. Nielsen, A. Wasowski, Maximizing entropy over Markov processes, J. Log. Algebraic Methods Program. 83 (5) (2014) 384–399.
[35] Y. Chen, T.T. Georgiou, M. Pavon, Ruelle-Bowen continuous-time random walk, in: Proceedings of the 23rd International Symposium on Mathematical Theory of Networks and Systems, 2018, pp. 496–499.
[36] L. Demetrius, T. Manke, Robustness and network evolution – an entropic principle, Physica A 346 (3) (2005) 682–696.
[37] J. Justesen, T. Hoholdt, Maxentropic Markov chains, IEEE Trans. Inform. Theory 30 (4) (1984) 665–667.
[38] R. Sinatra, J. Gómez-Gardeñes, R. Lambiotte, V. Nicosia, V. Latora, Maximal-entropy random walks in complex networks with limited information, Phys. Rev. E 83 (2011) 030103.
[39] J.M. Kleinberg, Authoritative sources in a hyperlinked environment, J. Assoc. Comput. Mach. 46 (5) (1999) 604–632.
[40] Y. Chen, T.T. Georgiou, M. Pavon, A. Tannenbaum, Robust transport over networks, IEEE Trans. Automat. Control 62 (9) (2017) 4675–4682.
[41] Y. Chen, T.T. Georgiou, M. Pavon, A. Tannenbaum, Efficient robust routing for single commodity network flows, IEEE Trans. Automat. Control 63 (7) (2017) 2287–2294.
[42] M. Pavon, F. Ticozzi, Discrete-time classical and quantum Markovian evolutions: Maximum entropy problems on path space, J. Math. Phys. 51 (4) (2010) 042104.
[43] E. Todorov, Efficient computation of optimal actions, Proc. Natl. Acad. Sci. 106 (28) (2009) 11478–11483.
[44] J. Perotti, O. Billoni, Smart random walkers: The cost of knowing the path, Phys. Rev. E 86 (2012) 011120.
[45] H.J. Kappen, Path integrals and symmetry breaking for optimal control theory, J. Stat. Mech. Theory Exp. 2005 (11) (2005) P11011.
[46] A. Busic, S. Meyn, Action-constrained Markov decision processes with Kullback-Leibler cost, in: Proceedings of the 31st Conference on Learning Theory, COLT, PMLR 75, 2018, pp. 1431–1444.
[47] R. Fox, A. Pakman, N. Tishby, G-learning: taming the noise in reinforcement learning via soft updates, in: Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence, UAI 2016, 2001, pp. 202–211.
[48] H.J. Kappen, V. Gómez, M. Opper, Optimal control as a graphical model inference problem, Mach. Learn. 87 (2) (2012) 159–182.
[49] G. Neu, A. Jonsson, V. Gómez, A unified view of entropy-regularized Markov decision processes, 2017, arXiv preprint arXiv:1705.07798.
[50] J. Rubin, O. Shamir, N. Tishby, Trading value and information in MDPs, in: Decision Making with Imperfect Decision Makers, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 57–74.
[51] E.A. Theodorou, D. Krishnamurthy, E. Todorov, From information theoretic dualities to path integral and Kullback-Leibler control: continuous and discrete time formulations, in: The Sixteenth Yale Workshop on Adaptive and Learning Systems, 2013.
[52] E.A. Theodorou, E. Todorov, Relative entropy and free energy dualities: Connections to path integral and KL control, in: Proceedings of the 51st IEEE Conference on Decision and Control, CDC 2012, IEEE, 2012, pp. 1466–1473.
[53] P. Winder, Reinforcement Learning. Industrial Applications of Intelligent Agents, O'Reilly, 2020.
[54] Y. Savas, M. Ornik, M. Cubuktepe, U. Topcu, Entropy maximization for constrained Markov decision processes, in: Proceedings of the IEEE 56th Annual Allerton Conference on Communication, Control, and Computing, ICML '98, 2018, pp. 515–521.
[55] R.K. Ahuja, T.L. Magnanti, J.B. Orlin, Network Flows: Theory, Algorithms, and Applications, Prentice Hall, 1993.
[56] X. Ding, S.L. Smith, C. Belta, D. Rus, Optimal control of Markov decision processes with linear temporal logic constraints, IEEE Trans. Automat. Control 59 (5) (2014) 1244–1257.
[57] A. Fronczak, P. Fronczak, Biased random walks in complex networks: The role of local navigation rules, Phys. Rev. E 80 (2009) 016107.
[58] M. Bonaventura, V. Nicosia, V. Latora, Characteristic times of biased random walks on complex networks, Phys. Rev. E 89 (2014) 012803, http://dx.doi.org/10.1103/PhysRevE.89.012803.
[59] L. Backstrom, J. Leskovec, Supervised random walks: predicting and recommending links in social networks, in: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM '11, ACM, 2011, pp. 635–644.
[60] B. Benigni, R. Gallotti, M. De Domenico, Potential-driven random walks on interconnected systems, Phys. Rev. E 104 (2) (2021) 024120.
[61] A. Riascos, L. Mateos, Random walks on weighted networks: a survey of local and non-local dynamics, J. Complex Netw. 9 (5) (2021) 1–39.

[62] T.M. Cover, J.A. Thomas, Elements of Information Theory, second ed., Wiley, 2006.
[63] J.N. Kapur, Maximum-Entropy Models in Science and Engineering, New Age International, 1989.
[64] M. Senelle, S. García-Díez, A. Mantrach, M. Shimbo, M. Saerens, F. Fouss, The sum-over-forests density index: Identifying dense regions in a graph, IEEE Trans. Pattern Anal. Mach. Intell. 36 (6) (2014) 1268–1274.
[65] L. Ekroot, T. Cover, The entropy of Markov trajectories, IEEE Trans. Inform. Theory 39 (4) (1993) 1418–1421.
[66] J. Gómez-Gardenes, V. Latora, Entropy rate of diffusion processes on complex networks, Phys. Rev. E 78 (6) (2008) 065102.
[67] M. Kafsi, M. Grossglauser, P. Thiran, The entropy of conditional Markov trajectories, IEEE Trans. Inform. Theory 59 (9) (2013) 5577–5583.
[68] A. Basilevsky, Applied Matrix Algebra in the Statistical Sciences, North-Holland, 1983.
[69] G. Seber, A Matrix Handbook for Statisticians, Wiley, 2008.
[70] C.D. Meyer, Matrix Analysis and Applied Linear Algebra, Society for Industrial and Applied Mathematics, 2000.
[71] J. Magnus, H. Neudecker, Matrix Differential Calculus with Applications in Statistics and Econometrics, third ed., Wiley, 2007.
[72] J. Franklin, Matrix Theory, Prence-Hall, 1968.
[73] S. García-Díez, F. Fouss, M. Shimbo, M. Saerens, A sum-over-paths extension of edit distances accounting for all sequence alignments, Pattern Recognit. 44 (6) (2011) 1172–1182.
[74] D.P. Bertsekas, Nonlinear Programming, second ed., Athena Scientific, 1999.
[75] S. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004.
[76] I. Griva, S. Nash, A. Sofer, Linear and Nonlinear Optimization, second ed., SIAM, 2008.
[77] T. Jebara, Machine Learning: Discriminative and Generative, Kluwer Academic Publishers, 2004.
[78] K. Arrow, L. Hurwicz, H. Uzawa, Studies in Linear and Non-Linear Programming, Stanford University Press, 1958.
[79] R. Rockafellar, The multiplier method of hestenes and Powell appplied to convex programming, J. Optim. Theory Appl. 12 (6) (1973) 555–562.
[80] M. Gori, A. Pucci, ItemRank: A random-walk based scoring algorithm for recommender engines, in: Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI '07, 2007, pp. 2766–2771.
[81] T. Haveliwala, Topic-sensitive PageRank, in: Proceedings of the 11th International World Wide Web Conference, WWW '02, 2002, pp. 517–526.
[82] S. Park, W. Lee, B. Choe, S.-G. Lee, A survey on personalized PageRank computation algorithms, IEEE Access 7 (2019) 163049–163062.
[83] H. Tong, C. Faloutsos, J.-Y. Pan, Fast random walk with restart and its applications, in: Proceedings of the 6th IEEE International Conference on Data Mining, ICDM '06, 2006, pp. 613–622.
[84] D. Kotkov, S. Wang, J. Veijalainen, A survey of serendipity in recommender systems, Knowl.-Based Syst. 111 (2016) 180–192, http://dx.doi.org/10.1016/j.knosys.2016.08.014.
[85] M. Latapy, C. Magnien, N. Del Vecchio, Basic notions for the analysis of large two-mode networks, Social Networks 30 (1) (2008) 31–48.
[86] C.D. Meyer, Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems, SIAM Rev. 31 (2) (1989) 240–272.
[87] H. Tong, C. Faloutsos, J.-Y. Pan, Random walk with restart: Fast solutions and applications, Knowl. Inf. Syst. 14 (3) (2008) 327–346.
[88] S. Siegel, J. Castellan, Nonparametric Statistics for the Behavioral Sciences, second ed., McGraw-Hill, 1988.
[89] F. Fouss, A. Pirotte, J.-M. Renders, M. Saerens, Random-walk computation of similarities between nodes of a graph, with application to collaborative recommendation, IEEE Trans. Knowl. Data Eng. 19 (3) (2007) 355–369.
[90] M. Ge, C. Delgado, D. Jannach, Beyond accuracy: evaluating recommender systems by coverage and serendipity, in: Proceedings of the 4th ACM Conference on Recommender Systems, RecSys, 2010, pp. 257–260.
[91] J. Herlocker, J. Konstan, L. Terveen, J. Riedl, Evaluating collaborative filtering recommender systems, ACM Trans. Inf. Syst. 22 (1) (2004) 5–53.
[92] L. Iaquinta, M. De Gemmis, P. Lops, G. Semeraro, M. Filannino, P. Molino, Introducing serendipity in a content-based recommender system, in: 2008 Eighth International Conference on Hybrid Intelligent Systems, IEEE, 2008, pp. 168–173.
[93] K. Kapoor, V. Kumar, L. Terveen, J.A. Konstan, P. Schrater, I like to explore sometimes: Adapting to dynamic user novelty preferences, in: Proceedings of the 9th ACM Conference on Recommender Systems, ACM, 2015, pp. 19–26.
[94] S. Vargas, P. Castells, Rank and relevance in novelty and diversity metrics for recommender systems, in: Proceedings of the Fifth ACM Conference on Recommender Systems, 2011, pp. 109–116.
[95] P. Adamopoulos, A. Tuzhilin, On unexpectedness in recommender systems: Or how to better expect the unexpected, ACM Trans. Intell. Syst. Technol. (TIST) 5 (4) (2015) 54.
[96] T. Murakami, K. Mori, R. Orihara, Metrics for evaluating the serendipity of recommendation lists, in: Annual Conference of the Japanese Society for Artificial Intelligence, Springer, 2007, pp. 40–46.
[97] M. Radovanović, A. Nanopoulos, M. Ivanović, Hubs in space: Popular nearest neighbors in high-dimensional data, J. Mach. Learn. Res. 11 (2010) 2487–2531.
[98] M. Radovanović, A. Nanopoulos, M. Ivanović, On the existence of obstinate results in vector space models, in: Proceedings of the 33rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10, 2010, pp. 186–193.
[99] I. Suzuki, K. Hara, M. Shimbo, Y. Matsumoto, M. Saerens, Investigating the effectiveness of Laplacian-based kernels in hub reduction, in: Proceedings of the 26th AAAI Conference on Artificial Intelligence, AAAI '12, 2012, pp. 1112–1118.
[100] N. Tomasev, M. Radovanovic, D. Mladenic, M. Ivanovic, The role of hubness in clustering high-dimensional data, IEEE Trans. Knowl. Data Eng. 26 (3) (2014) 739–751.
[101] K. Hara, I. Suzuki, M. Shimbo, K. Kobayashi, K. Fukumizu, M. Radovanovic, Localized centering: Reducing hubness in large-sample data, in: Proceedings of the 29th AAAI Conference on Artificial Intelligence, AAAI '15, 2015, pp. 2645–2651.
[102] S. García-Díez, E. Vandenbussche, M. Saerens, A continuous-state version of discrete randomized shortest-paths, in: Proceedings of the 50th IEEE International Conference on Decision and Control, CDC '11, 2011, pp. 6570–6577.
[103] R. Chetrite, H. Touchette, Nonequilibrium Markov processes conditioned on large deviations, Ann. Henri Poincaré 16 (9) (2014) 2005–2057.
[104] S.N. Majumdar, H. Orland, Effective langevin equations for constrained stochastic processes, J. Stat. Mech. Theory Exp. 2015 (6) (2015) P06039.
[105] P.G. Bolhuis, D. Chandler, C. Dellago, P.L. Geissler, Transition path sampling: Throwing ropes over rough mountain passes, in the dark, Annu. Rev. Phys. Chem. 53 (1) (2002) 291–318.
[106] A. Das, D.C. Rose, J.P. Garrahan, D.T. Limmer, Reinforcement learning of rare diffusive dynamics, 2021, arXiv preprint arXiv:2105.04321.
[107] H.J. Kappen, H.C. Ruiz, Adaptive importance sampling for control and inference, J. Stat. Phys. 162 (5) (2016) 1244–1266.
[108] U. Ray, G.K.-L. Chan, D.T. Limmer, Importance sampling large deviations in nonequilibrium steady states. I, J. Chem. Phys. 148 (12) (2018) 124120.
[109] G.M. Torrie, J.P. Valleau, Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling, J. Comput. Phys. 23 (2) (1977) 187–199.
[110] J.L. Doob, Conditional Brownian motion and the boundary limits of harmonic functions, Bull. Soc. Math. France 85 (1957) 431–458.