

Learning to Represent Human Motives for Goal-directed Web Browsing

Jyun-Yu Jiang*
jyunyu@cs.ucla.edu
University of California, Los Angeles
USA

Chia-Jung Lee*
cjlee@amazon.com
Amazon Inc.
USA

Longqi Yang
loy@microsoft.com
Microsoft
USA

Bahareh Sarrafzadeh
bahar.sarrafzadeh@microsoft.com
Microsoft
USA

Brent Hecht
brhecht@microsoft.com
Microsoft
USA

Jaime Teevan
teevan@microsoft.com
Microsoft
USA

Abstract

Motives or goals are recognized in psychology literature as the most fundamental drive that explains and predicts why people do what they do, including when they browse the web. Although providing enormous value, these higher-ordered goals are often unobserved, and little is known about how to leverage such goals to assist people’s browsing activities. This paper proposes to take a new approach to address this problem, which is fulfilled through a novel neural framework, *Goal-directed Web Browsing* (GoWeB). We adopt a psychologically-sound taxonomy of higher-ordered goals and learn to build their representations in a structure-preserving manner. Then we incorporate the resulting representations for enhancing the experiences of common activities people perform on the web. Experiments on large-scale data from Microsoft Edge web browser show that GoWeB significantly outperforms competitive baselines for in-session web page recommendation, re-visitation classification, and goal-based web page grouping. A follow-up analysis further characterizes how the variety of human motives can affect the difference observed in human behavioral patterns.

Keywords

User Behavior; User Goals; Web Browser Session Modeling; Goal Representation Learning

ACM Reference Format:

Jyun-Yu Jiang, Chia-Jung Lee, Longqi Yang, Bahareh Sarrafzadeh, Brent Hecht, and Jaime Teevan. 2021. Learning to Represent Human Motives for Goal-directed Web Browsing. In *Proceedings of The 15th ACM Conference on Recommender Systems (RecSys ’21), September 27–October 1, 2021, Amsterdam, Netherlands*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3460231.3474260>

*Work done while at Microsoft.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys ’21, September 27–October 1, 2021, Amsterdam, Netherlands

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8458-2/21/09.

<https://doi.org/10.1145/3460231.3474260>

1 Introduction

Constructs such as motives and goals are recognized in the psychology literature as the fundamental forces that guide human behavior. While a variety of interpretations of goals exist, there is consensus that these constructs regulate controlled cognitive processes such as planning and resource allocation, and direct behavioral sequences intended to enact specific performances [10]. When people browse the web, which is one important behavior common in modern lives, goals also underlie why they do what they do. For example, they may buy products to work towards their fitness goals, plan trips to stay connected with friends and family, conduct business to pursue career success, or do research or seek health advice in support of their well-being. Understanding peoples’ goals not only helps answer the “why” questions about their activities [18, 56], but also signifies where the opportunities sit for potential improvements.

Although goals and motives are fundamental to browsing behaviors, they stay unobserved and cannot be easily inferred due to the variety and complexity of activities prevalent on modern web browsers. Moreover, the question of how goals and motives can be utilized to assist browsing remains unanswered. To bridge the gap, we propose to adopt the well-established bodies of psychological theories [5, 11, 44] regarding human motives to the need for browser-centric activities. For instance, the motive of *staying physically healthy* may energize a person to consider *purchasing a rowing machine* or *to learn nutritious ingredients* via a number of web interactions. As a result, a better understanding of human motives can guide how we might enhance browsing experiences pertaining to people’s needs.

Previous research on supporting user goals on the web has primarily sought to understand and categorize a person’s intent manifested through search queries. In their seminal studies, Broder [8] and Rose and Levinson [54] classified the goal of a user query into one of three categories: navigational, informational, and transactional. Successive work refined intent categorization by introducing purchase, sell or job search intents [14, 36], as well as examined the variability of intents across different users [62]. The main objective of research in this field is often to inform and improve the search process [8, 9, 35, 54], and the characteristics of user goals tend to be task-oriented or situation-specific. In contrast to that, we emphasize incorporating the fundamental human goals into the context of web browsing, since they are recognized to be key to the core

values in people’s lives [5] (e.g. health, well-being, sustainability, and learning).

In this paper, we focus on tackling two research agendas toward our goal-directed vision, namely, how to represent influential human goals¹ and how to leverage the learned representations to provide assistance in browsing sessions. We propose a unified neural framework, Goal-directed Web Browsing (GoWeB), for both objectives laid out in two phases. The first phase concerns human motives. Psychology research has developed decades of knowledge informing which human goals are considered influential through theoretical and empirical evidence. Inspired by the findings, GoWeB uses these goals as the backbone and builds on the top to learn intrinsic goal representations following a distance-based reconstruction objective to retain inherent hierarchical structures². GoWeB further devises a neural goal estimator, which learns to transfer the knowledge entailed in these goal representations to estimating people’s motives for visiting web pages. The second phase focuses on enhancing web experiences with what we have learned. Specifically, GoWeB integrates the goal estimator with modern language models [15] to assist three common interactions or needs people have on the web. We consider an in-session web page recommender that recommends web pages to help people advance their goals, a web page re-visitation classifier that predicts if a person will revisit a web page in the future due to recurring or unfinished goals, and a web page grouping method that clusters in-session web pages according to underlying goals. The overall framework is illustrated in Figure 4.

We evaluate GoWeB using anonymized browsing logs obtained from the Microsoft Edge browser. The experimental results show that GoWeB consistently outperforms competitive baselines in all three browser-centric applications, suggesting that capturing fundamental human motives can empirically improve the intelligence of web browsers. To understand the effectiveness, we find that the learned intrinsic goal representations largely preserve the human-curated hierarchical relations. Meanwhile, the goal estimator can effectively predict the goals behind web page visits evaluated based on quantitative and qualitative exercises. Our follow-up analysis further characterizes the browsing patterns when people engage in pursuing different goals. The results suggest that people may explore multiple directions simultaneously when the categories of session’s focal goals are broad (e.g., *Ethics & Idealism*), while they may revisit web resources in shorter intervals for goals concerning social interactions (e.g., *Friendship*).

2 Related Work

Studying goals is an important, long-lasting topic across multiple research disciplines. We examine the different notions of goals that have been investigated in the literature.

Task-oriented Goals in Search and Web Browsing. Prior work has a strong focus on studying search goals as they contribute to a critical segment of web browser usage. The mainstream work in this space follows Broder [8], and Rose and Levinson [54] who

classified the goal of user queries into navigational, informational, and/or transactional. Lee et al. [35] identified that 60% of search queries can be associated with informational or navigational goals, and proposed using clicks and anchors to predict those. Caruccio et al. [9] introduced a lightweight taxonomy that added a handful of sub-classes to informational and transactional queries. Casting as a classification problem, they proposed content and behavioral features to predict which class in the taxonomy a query should belong to using data collected from a custom browser extension. For more complex search needs, Jones and Klinkner [26] leveraged human annotations and built predictive models that were trained to segment sequences of user queries into same or different search goals. Among a set of hand-crafted features, they showed that lexical features such as words and characters were identified as the most useful. Similarly, Law and Zhang [32] showed that the ability to decompose a complex search goal into sub-problems (i.e., a set of queries and corresponding search results) can better support a user’s information need.

General web browsing behavioral patterns tend to involve higher flexibility and complexity than search. Kumar and Tomkins [31] analyzed large-scale commercial logs and showed that page visits related to search comprise only one-sixth of the log sample. In addition to search, they identified that content consumption (e.g., news, portals, games, verticals, multimedia) comprise about half of online page views, while communication (e.g., email, social networking, forums, blogs, chat) represent about one-third of those. As a result, general web browsing patterns require different types of support. For example, past studies have shown that people are often unable to focus on or return to their goals when using a web browser due to the prevalence of distractions on the web [1, 63]. Research on Cyberloafing [20, 37] and task resumption on the web [19, 43, 48] have shown that individuals experience difficulty when resuming their main task after external- or self-interruptions [3, 13].

Faaborg and Lieberman [18] introduced a program-by-example goal-oriented web browser that allows people to customize how they organize and interact with the web. A person who visits a recipe website, for example, may want to know the nutritional information. By explicitly demonstrating how to complete nutrition extraction once, the system could potentially evoke the same human-specified macros on similar websites. Dix et al. [16] attempted to understand the goals behind user behavior by connecting actions and user-created personal ontology structures.

Higher-ordered Human Goals. Understanding human goals and motives has long been a central area of research in the psychology literature, as the goals of individuals largely direct the behavior in which they engage [29]. Earlier attempts in creating goal taxonomies relied on theoretical viewpoints. McDougall [42] presented a list of 13 instincts while Murray [45] articulated 44 variables of personality as forces determining behavior. Chulef et al. [11] took an empirical approach and recruited participants of diverse backgrounds to delve into developing a hierarchical human goal taxonomy based on similarity between goals, providing a concrete and comprehensive structure. More recently, Talevich et al. [60] iterated on the taxonomy derived in [11] with several added classes of human motives.

¹Unless otherwise mentioned, we use human motives and goals interchangeably in this paper.

²For example, in Chulef et al. [11]’s human goal taxonomy, the goals *be affectionate* and *share feelings* are both decedents of a parent goal *friendship*.

As our work incorporates long-term goals in browsing sessions to provide goal-directed assistance, we review the classic and state-of-the-art technical methods applicable in the three browser-centric tasks, and conclude with a brief summary comparing previous approaches and ours.

Recommendation. The task of web page recommendation [21] aims to predict the information needs of users and provide them with recommendations to facilitate their navigation. Classic approaches to recommendation relies on mining user-item relationships through collaborative filtering [58] and matrix factorization [23, 52]. Session-based approaches [61] take into account session-level user actions as a sequential stream of events, based on which recommendations are made to users. In recent years, large-scale deep neural networks [68] have shown top performing effectiveness in recommender systems, including leveraging GRU [24], self-attention blocks [27] and BERT-based architectures [59] for session-based recommendation.

Revisitation. Web page revisitation behavior is prevalent on the web. Obendorf et al. [47] suggested that people may revisit web pages to access the same resource again (e.g., for unfinished goals) or may re-access a resource as they expect changed content (e.g., for recurring goals, such as new headlines on a news site). Adar et al. [2] combined quantitative and qualitative approaches and showed that people revisit web pages with varying speed. To help resume previously encountered information, Dontcheva et al. [17] presented a web extension that allows people to extract entities of interest and creates an interactive summary of extracted entities (e.g., hotels that a person considers for a trip).

Grouping. Automatic managing and grouping the vast amount of unorganized web pages can help reduce the cognitive load required when navigating the web. A common strategy is to classify web pages [6, 51] by the encapsulated topics; popular class choices include the Open Directory Project, the Yahoo! Directory, categorized URLs from the Wikipedia, etc. To relax the requirement of pre-defined classes, clustering techniques have been studied to mine related web pages based on semantics or graph partitioning [50]. Different from prior studies, we adopt the notion of goals as the basis to associate web pages for web page grouping.

In this work, we incorporate long-term goals to facilitate effective web browsing experiences, including recommendation, revisitation and goal-based grouping, which differ significantly from previous perspectives and approaches. We leverage Poincaré embeddings [46] to derive the goal representations, which are combined with state-of-the-art transformer-based architectures [15] to model users’ web session behavior.

3 Representation Learning for Human Motives

This section presents our approach to learning distributed intrinsic goal representations (iGoalRep r_g) for human motives, which are grounded by an expert-curated goal taxonomy in the psychology literature. Moreover, we propose a goal estimator to embed any web page visit p in browsing sessions into the same goal embedding space as r_g , and we denote them as visit goal representation (vGoalRep) r_p . To learn the goal estimator, we collect and rely on a weak supervision dataset through the associations between related

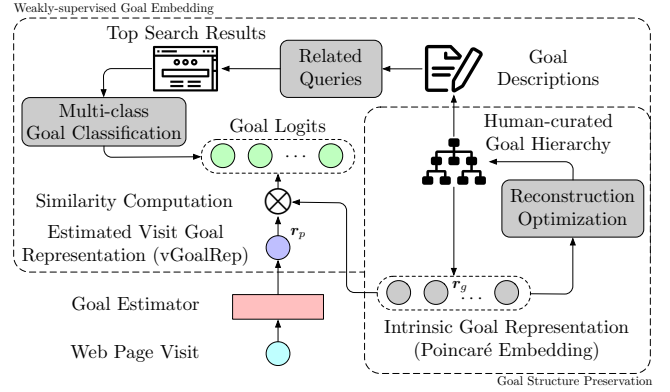


Figure 1: Framework for learning intrinsic goal representations (iGoalRep) and training the goal estimator to predict visit goal representations (vGoalRep).




Figure 2: Illustration of the constructed 3-layer goal taxonomy with 166 nodes.

queries and web pages determined by the Bing search engine. Figure 1 summarizes the flow of the steps, and we describe the details of each step next.

3.1 Learning Intrinsic Goal Representations

Taxonomy of Human Goals. We leverage the hierarchical human goal taxonomy curated by Chulef et al. [11] to form the basis of the goal space. Specifically, the taxonomy leaf nodes were curated based on an aggregation of psychology literature [41, 45, 53, 65] and the responses from participants of diverse background, followed by a similarity-based sorting process to create the hierarchy relationship. The resulting set consists of 135 leaf goal nodes descended from 30 parent categorical nodes. We further introduce a “non-goal” node as the root that covers all the categorical nodes. Conceptually, the root represents prosaic behavior that might not be tied with specific goals; examples could include random distractions on the web or one-time matters. Figure 2 illustrates the final adopted human goal taxonomy consisting of 166 nodes in three layers. Formally, G denotes the set of goals in the goal taxonomy and H denotes the hierarchical structure.

Representation Learning in Hyperbolic Space. To induce the structural bias H , we propose to learn intrinsic goal representations (iGoalRep) in a hyperbolic space. Hyperbolic geometry brings the advantage of learning compact representations that capture both hierarchy and similarity, which, for our case, is desirable as we aim to preserve the hierarchical properties of the goal taxonomy.




Figure 3: Illustration of learned 2-dimensional Poincaré goal representations with the 3-layer goal taxonomy from [11]. The center red node indicates the concept of “non-goal,” while the blue nodes represent goal categories that cover more specific goals presented as orange nodes.

Concretely, we adopt Poincaré embeddings [46] to derive d_G -dimensional iGoalRep r_g for each $g \in G$. Let $\mathcal{B}^{d_G} = \{\mathbf{x} \in \mathbb{R}^{d_G} \mid \|\mathbf{x}\| < 1\}$ be an open Poincaré ball, where $\|\cdot\|$ indicates the Euclidean norm. In contrast to Euclidean distance, the hyperbolic distance between any two points $\mathbf{u}, \mathbf{v} \in \mathcal{B}^{d_G}$ is given as:

$$d(\mathbf{u}, \mathbf{v}) = \operatorname{arcosh} \left(1 + 2 \frac{\|\mathbf{u} - \mathbf{v}\|^2}{(1 - \|\mathbf{u}\|^2)(1 - \|\mathbf{v}\|^2)} \right).$$

Subsequent representation learning is then guided by this distance measure where a similar pair of goals should be closer to each other than a remote pair.

Reconstruction Optimization. To start learning, we treat the edges in the goal taxonomy as transitive closure and reconstruct the relationship by minimizing the Poincaré distances between the embeddings of their endpoints. Following prior work [46], we define the loss function \mathcal{L}_r as:

$$\mathcal{L}_r = \sum_{(g_u, g_v) \in H} \log \frac{e^{-d(r_{g_u}, r_{g_v})}}{\sum_{g_{v'} \in \mathcal{N}(g_u)} e^{-d(r_{g_u}, r_{g_{v'}})}},$$

where $\mathcal{N}(g_u) = \{g_{v'} \mid (g_u, g_{v'}) \notin H\} \cup \{g_u\}$ denotes a set of negative examples³ for g_u and $H = \{(g_u, g_v)\} \cup \{(g_v, g_u) \mid g_u, g_v \in G\}$ contains symmetric relationship. We adopt RSGD [7], a stochastic Riemannian optimization method, to learn goal embeddings within the Poincaré ball as in [46]. To do so, we need to rely on the Riemannian manifold structure of the Poincaré Ball when minimizing \mathcal{L}_r . Since Euclidean gradient is not directly applicable in hyperbolic space, rescaling gradients depends on the Riemannian metric tensor t_x equipped in the Poincaré ball: $t_x = \left(\frac{2}{1 - \|\mathbf{x}\|^2}\right)^2 t^E$, where $\mathbf{x} \in \mathcal{B}^{d_G}$; t^E denotes the Euclidean metric tensor.

³In practice, we randomly sample 50 negative examples for each positive example.

Figure 3 visualizes the resulting iGoalRep r_g with two dimension⁴ learned from the goal taxonomy. The results demonstrate that the learned embeddings preserve the desired hierarchical properties and align with the original goal taxonomy curated by human.

3.2 Estimating Visit Goal Representations (vGoalRep)

When a person visits a web page, we aim to predict a visit goal representation (vGoalRep) r_p such that r_p can best reflect which goals may have driven the visit. To do so, we first build a weak supervision data collection that provide the association information of what the most probable goals are when web pages are being visited. Then, a parameterized goal estimator is introduced to predict r_p based on the weak labels according to a multi-class classification objective.

Weak Supervision Data Collection. Inspired by past work [57] that demonstrated search engine queries can contain explicit user objectives (e.g., get rid of belly fat), we first reify every goal in the goal taxonomy with a set of seed queries. For example, we manually generate queries such as “how to be charismatic” and “how to meet new friends” to elicit the high-ordered goal of *being likeable, making friends, drawing others near*. Then we expand the seed query set with related queries by requesting publicly available Microsoft Bing search API. In total, 300 seed queries and 1,932 related queries represent 165 defined human goals. The last step involves querying the same API using the enlarged query set together with the original goals; we then keep the top 5 returned web pages as the weakly-positive instances for the corresponding goal. For the concept of “non-goal,” i.e., the root of the goal taxonomy, we randomly select 1,000 web page visits from general web browsing logs as “negative examples”. Here we denote this data collection as D_{weak} .

Goal Estimator Construction. Next, we aim to devise a goal estimator that can predict vGoalRep r_p by projecting each web page p to the same goal embedding space in which r_g exists. A web page visit p is commonly characterized by the semantics conveyed via its textual content c_p , as well as the website h_p where it is hosted (e.g., pages on the same site often demonstrate higher associations). To build an effective goal estimator, we take as input the two sources of information, $p = (h_p, c_p)$, and estimate vGoalRep $r_p \in \mathbb{R}^{d_G}$ by:

$$r_p = \mathcal{F}_G([\operatorname{emb}_{\text{host}}^G(h_p); \operatorname{BERT}^G(c_p)]),$$

where $\operatorname{emb}_{\text{host}}^G(h_p)$ projects the host to a d_h -dimensional embedding space; $\operatorname{BERT}^G(\cdot)$ is a contextualized content encoder, such as BERT [15] and RoBERTa [38]; $\mathcal{F}_G(\cdot)$ as a fully-connected hidden layer derives ultimate d -dimensional estimated goal embeddings. The output of the goal estimator, r_p , can be regarded as a continuous representation of certain underlying motives for visiting a particular web page p .

Multi-class Goal Classification. While the architecture of goal estimator is the same as web page encoder, we learn the parameters using a multi-class goal classification objective guided by D_{weak} . The multi-class classification loss function is defined by categorical cross-entropy [22] as $\mathcal{L}_c: \mathcal{L}_c(p) = -\sum_{g_u \in G} \mathbb{1}(\hat{g}_p =$

⁴We select dim=2 for visualization convenience here; in the actual experiments, dim=64 is used as described in Section 5.

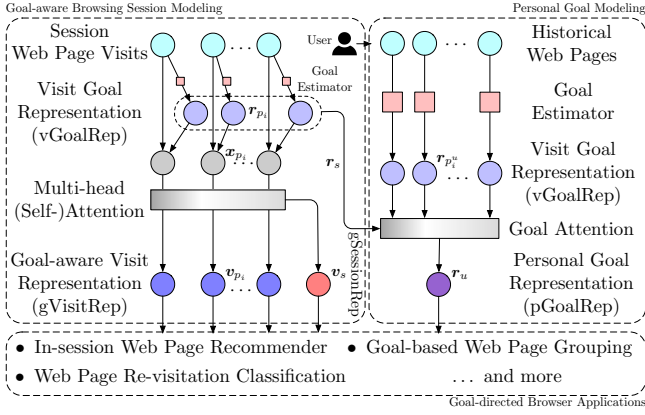


Figure 4: Overview of GoWeB. The goal estimator produces goal representations r_p , which are then used to derive goal-aware visit v_p and session v_s representations, as well as personal goal representation r_u . These representations facilitate goal-directed experiences in browser applications. The details of the goal estimator are described in Section 3.

$g_u) \log(\sigma_{g_u}(r_p, G))$ where $\mathbb{1}(\cdot)$ is an indicator function; \hat{g}_p is the labeled goal of the web page p ; $\sigma_{g_u}(r_p, G)$ indicates the softmax function calculating the probabilistic distribution over goals in the goal taxonomy as: $\sigma_{g_u}(r_p, G) = \frac{r_p \otimes r_{g_u}}{\sum_{g \in G} r_p \otimes r_g}$. Here $r_p \otimes r_g$ denotes the similarity score between the web page p and the goal g in the Poincaré embedding space. In other words, by treating the Poincaré similarity scores as classification logits, r_p is trained to reflect its ground-truth goal embedding r_g . Formally: $r_p \otimes r_g = \|r_p\| \cdot \|r_g\| \cdot \cos(\theta_{r_p}, \theta_{r_g})$ where θ_{r_p} and θ_{r_g} indicate the angles of embeddings in the hyperbolic coordinate system.

A final remark to note is that we could derive the most probable goal g_p associated with a page p based on the learned r_p , by simply calculating the similarity between r_p and r_g : $g_p = \operatorname{argmax}_{g \in G} r_p \otimes r_g$.

When designing GoWeB in Section 4, we choose to incorporate r_p rather than g_p because a continuous representation can encode more information and reduce sparsity compared to a discrete one.

4 Goal-directed Web Browsing (GoWeB)

Based on the goal estimator and visit goal representations introduced in Section 3, we present a unified framework, Goal-directed Web Browsing (GoWeB), to assist common activities that people perform when browsing the web. We first architect a generic neural framework that can encapsulate information incurred in browsing using goal-aware representations. Second, we employ this framework to assist people to advance their goals through in-session recommendation, to pick up their goals through re-visitation prediction, and to focus on certain goals through goal-based clustering. Figure 4 further illustrates an overview of GoWeB.

4.1 Goal-aware Browsing Session Modeling

A web browsing session involves three key actors: individual web pages that a person visits, the sequence of visits in the same session, and the person who performs these visits. To support goal-directed

browsing, GoWeB needs to derive representations that effectively incorporate goal awareness for each of these. More concretely, suppose a web browsing session s consists of n web page visits $s = \{p_1, \dots, p_n\}$ performed by a person u . In the below, we describe the steps to form individual *goal-aware visit representation* (gVisitRep) as v_{p_i} , *goal-aware session representation* (gSessionRep) as v_s , and *personal goal representation* (pGoalRep) as r_u .

Goal-aware Visit Representation (gVisitRep). To represent a web page visit, we consider two types of essential signals, namely the lexical content expressed inside the page as well as the probable goals accounting for the visit. Specifically, we derive content-based embeddings by initializing a separate set of weights using the same encoder architecture described in Section 3.2. For each visit p_i , the content embedding $w_{p_i} \in \mathbb{R}^{d_v}$ is determined by its underlying content c_{p_i} and the host of the page h_{p_i} . Formally, we compute content embedding by $w_{p_i} = \mathcal{F}_V([\text{emb}_{\text{host}}^V(h_{p_i}), \text{BERT}^V(c_{p_i})])$, where $\text{emb}_{\text{host}}^V(h_{p_i})$ projects the host to a d_h -dimensional embedding space; $\text{BERT}^V(\cdot)$ is a contextualized language model; $\mathcal{F}_V(\cdot)$ as a fully-connected hidden layer derives the final d_v -dimensional content embedding. While the content-based embeddings are important, we take a step further to incorporate probable goals for the act of visiting, by concatenating the vGoalRep r_{p_i} and the content embedding w_{p_i} as $x_{p_i} = [r_{p_i}, w_{p_i}]$. Thus, visits in the same session s can be denoted as $X \in \mathbb{R}^{n \times d_{\text{model}}}$, where n is the number of pages and $d_{\text{model}} = d_G + d_v$.

To make page visits sensitive to the session context in which they appear, we propose to apply the multi-head attention mechanism [64] to derive contextualized visit representations. Formally, for each web page visit, the gVisitRep v_{p_i} can be computed as:

$$[v_{p_1}; \dots; v_{p_n}] = \text{Concat}(\text{head}_1, \dots, \text{head}_k)W^O,$$

where $\text{head}_i = \text{Attention}(XW_i^Q, XW_i^K, XW_i^V)$; k is the number of attention heads; $W_i^Q, W_i^K \in \mathbb{R}^{d_{\text{model}} \times d_k}$, $W_i^V \in \mathbb{R}^{d_{\text{model}} \times d_v}$, $W^O \in \mathbb{R}^{hd_v \times d_{\text{model}}}$; d_k and d_v are the dimension numbers of attention keys and values; $\text{Attention}(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$.

Goal-aware Session Representation (gSessionRep). To represent the overall browsing session, we learn a context vector $c_i \in \mathbb{R}^{1 \times d_k}$ for each head to dynamically estimate the importance of each visit for a specific application. The gSessionRep v_s of the browsing session can then be derived as $v_s = \text{Concat}(\text{head}_1^s, \dots, \text{head}_k^s)W^{Os}$, where $\text{head}_i^s = \text{Attention}(c_i, XW_i^{Ks}, XW_i^{Vs})$; $W_i^{Ks} \in \mathbb{R}^{d_{\text{model}} \times d_k}$, $W_i^{Vs} \in \mathbb{R}^{d_{\text{model}} \times d_v}$, $W^{Os} \in \mathbb{R}^{hd_v \times d_{\text{model}}}$.

Personal Goal Representation (pGoalRep). As an individual's past goals could be predictive of their future goals or behavior, conventionally it is a common practice to rely on an individual's past browsing activities to derive their personal representation [25]. We argue that it is equally important to account for the activities in a currently active session, according to Shah et al. [55] that the activation of a given focal goal tends to result in an inhibition of alternative goals. To implement this idea, GoWeB constructs pGoalRep r_u for an individual u by aggregating the current session activities as a query to attend to past activities of the same individual. Suppose $\{r_{p_i}\}$ denotes the vGoalReps for all page visits in a current

session; $Z_u = \{p_i^u\}$ is the set of web pages visited in the past by a user u . We adopt the Luong’s attention mechanism [39] to aggregate current session activities as:

$$\mathbf{r}_s = \sum_i \alpha_i \cdot \mathbf{r}_{p_i}, \alpha_i = \frac{e^{c_s \cdot \mathcal{F}_s(\mathbf{r}_{p_i})}}{\sum_j e^{c_s \cdot \mathcal{F}_s(\mathbf{r}_{p_j})}},$$

where α_i is the weight for \mathbf{r}_{p_i} ; c_s is the context vector to estimate the importance of each visit; $\mathcal{F}_s(\cdot)$ is a fully-connected hidden layer. The pGoalRep \mathbf{r}_u is computed by using \mathbf{r}_s as a query to discover and aggregate past goals related to current session:

$$\mathbf{r}_u = \sum_{p_i^u \in Z_u} \beta_i \cdot \mathbf{r}_{p_i^u}, \beta_i = \frac{e^{\mathbf{r}_s \cdot \mathbf{r}_{p_i^u}}}{\sum_{p_j^u \in Z_u} e^{\mathbf{r}_s \cdot \mathbf{r}_{p_j^u}}},$$

where β_i is the weight for each historical visit.

To summarize, we have described how to create goal-aware representations for a page visit (\mathbf{v}_{p_i}), a browsing session (\mathbf{v}_s), and an individual acting in the session (\mathbf{r}_u). These representations can, when designed properly, be used for directing people towards their goals while they browse the web. In particular, we consider three opportunities for goal-aware assistance described as follows.

4.2 Task 1: In-session Web Page Recommender

Making progress towards goals when navigating the web can be challenging, which could be due to unfamiliarity with subject matter or other external factors. To address this, GoWeB aims to recommend most goal-related, unseen resources that can help people advance their goals. Concretely, GoWeB formulates an in-session web page recommender. Given n preceding web pages $s = \{p_1, \dots, p_n\}$ visited in a session, GoWeB ranks web pages from a candidate set C according to their likelihood of being visited later in the same session. We follow prior work [12] and cast recommendation as a classification task. Based on the information available in s , we combine the gSessionRep \mathbf{v}_s and pGoalRep \mathbf{r}_u as a feature vector. The ranking scores $\mathbf{y}_{\text{rec}} \in \mathbb{R}^{|C|}$ of each web page in C can then be calculated as classification logits: $\mathbf{y}_{\text{rec}} = \mathcal{F}_{\text{rec}}(\mathcal{F}_{\text{hidden}}([\mathbf{v}_s, \mathbf{r}_u]))$, where $\mathcal{F}_{\text{hidden}}(\cdot)$ is a fully-connected hidden layer; $\mathcal{F}_{\text{rec}}(\cdot)$ projects the hidden state to the ranking scores of candidate web pages. Finally, the ranked list of recommendations can be generated by retrieving candidates with top ranking scores.

4.3 Task 2: Re-visitation Prediction

Web page re-visitation is prevalent as people may visit the same resources for unfinished or recurring goals. The ability to forecast potential future re-visitations can be useful for supporting people resuming their goals when predicted correctly. For instance, the predictions can be stored in the browser backend and resurface to users when they start a new session as a reminder. To support this, GoWeB aims to predict whether any of web pages in a current session will be revisited in a future session. For each visit p in the session s , we first construct its feature vector by concatenating the gVisitRep \mathbf{v}_p and the pGoalRep \mathbf{r}_u . The probability of the web page being re-visited can then be estimated as:

$$P(\text{revisit} = \text{True} \mid p, s) = \sigma(\mathcal{F}_{\text{rev}}(\mathcal{F}_{\text{hidden}}([\mathbf{v}_p, \mathbf{r}_u]))),$$

where $\mathcal{F}_{\text{hidden}}(\cdot)$ is a fully-connected hidden layer; $\mathcal{F}_{\text{rev}}(\cdot)$ projects the hidden state to a logit so that the sigmoid function $\sigma(\cdot)$ can

	# of Users	# Sessions	Avg. Session Length
Training	57,715	836,073	21.689 ± 11.922
Test (Warm)	6,701	38,737	22.257 ± 12.124
Test (Cold)	6,062	75,339	22.076 ± 12.127

Table 1: Statistics of web session experimental datasets.

derive a probabilistic score. We note that for the first and second tasks, GoWeB is trained end-to-end to directly optimize the target objective.

4.4 Task 3: Goal-based Web Page Grouping

Another common challenge with web browsing is being able to rationalize with goals and focus on subsets of web pages by goals. Considering that people often visit multiple, potentially diverse web pages in a session, this may create high cognitive load when people switch in-between goals. For this, we employ GoWeB to group in-session web pages by goals, such that the resulting groups of pages can be used to categorize ongoing flows and help people concentrate on certain goals of choice. Specifically, we cast this as a clustering task where the vGoalRep \mathbf{r}_p of each visit p is used as features, based on which subsequent feature-based clustering algorithms can be applied.

5 Experiments

This section examines the effect of incorporating high-ordered goals into the context of web browsing. We compare each of the browser-centric tasks with competitive baselines and demonstrate how GoWeB can enhance multiple browsing experiences.

5.1 Experimental Setup

Experimental Datasets. Our main experimental dataset, denoted as D_{web} , consists of web browsing sessions that are used for evaluating the three goal-based applications supported by GoWeB. D_{web} was constructed by randomly sampling the anonymized logs of Microsoft Edge web browser⁵. A session is composed of a sequence of web page visits of a user, where a boundary is found between two consecutive visits that are at least 30 minutes apart. The logs contain records of web pages visits assembled by sessions, where each page comes with a host, a title and a timestamp when the page was visited; the titles are used as the web page content. The data was sampled from June 2020, where the training and test sets were respectively gathered from the periods of June 1st to 23rd and June 24th to 30th. To avoid tail behavior, a web page is discarded if it appears less than 10 times in the sampling period, and short sessions with less than 10 page visits are disregarded as the associated goals are likely to be simpler. In total, our dataset contains web pages originating from 79,695 unique hosts (websites).

We prepare two disjoint test sets to study the effect of incorporating the notion of personal goals. The warm-start sample includes users whose behavior can be found in the training period, while the cold-start sample draws from users new to the system. The statistics of D_{web} used for the experiments are shown in Table 1. For cold-start users, since we do not have access to their past behavior,

⁵To the best of our knowledge, unfortunately, there is no publicly available user web browsing dataset that is suitable for evaluating our hypothesis and framework, but we will release our implementation to facilitate the community development.

we exclude the personal goal representations from the framework and denote this variant as **GoWeB (NP)**.

Implementation Details. We implement GoWeB and baseline methods with the PyTorch framework [49]. The dimensions of goals and hosts (i.e., d_G and d_h) are set to 64, while the number of hidden units in fully-connected layers and the dimension of the remaining embeddings are set to 128. The number of attention heads k is set to 8. We use a learning rate of 0.3 for RSGD [7] in the Riemannian optimizer, while a learning rate 10^{-5} and $(\beta_1, \beta_2) = (0.9, 0.999)$ are used in the Adam optimizer [28] for optimizing the supervised goal-based applications and the goal estimator. For contextualized language modeling (i.e., the BERT functions in Section 4 and 3), we use the pre-trained RoBERTa [38] models provided by Hugging-Face [66]. We will open-source our implementation.

5.2 Web Page Recommendation Performance

Recall that the in-session recommender predicts goal-related web pages that are likely to be visited later in the same session. Specifically, for a session, we take the first half of page visits as observed input information and use the second half as ground-truth to evaluate our model. The ranking candidates of web pages are generated by removing the top 10 popular web pages (e.g., `msn.com` and `google.com`) and selecting the top 50,000 frequent web pages⁶ from the remaining set. For both test sets, we assert that the ground-truth (i.e., the second half) of every session contains at least one web page from the candidate set. We treat the task as a retrieval problem and evaluate it by conventional metrics, including $MRR@10$, $HR@\{1, 5, 10\}$, and $NDCG@\{5, 10\}$. We compare our method with state-of-the-art top- K session-based recommenders, including popularity (Pop), BPRMF [52], NCF [23], Caser [61], GRU4Rec [24], SASRec [27], and BERT4Rec [59]. We further consider a modern content-based modeling method that incorporates web page semantics using BERT and contextualization. We implement this by reusing the GoWeB neural architecture but with all the goal-awareness removed (i.e. excluding r_p and any representations derived based on r_p), which is denoted as SemRec.

Table 2 demonstrates the recommendation performance in two test sets. Among the baseline methods, SemRec performs better than conventional session-based recommenders that disregard web page content. It advocates for the importance of modeling the semantics of web pages, as the open-domain nature of the web can make recommendation more challenging than domain-specific applications (e.g. movies). Without personalization, GoWeB (NP) performs the best and outperforms SemRec by 18.52% and 16.09% in $MRR@10$ for two test sets. This empirical finding sheds light on the opportunity of modeling human motives and higher-level goals carried by the goal estimator without accessing historical data, since only the in-session page visits are considered. For the warm-start test set where access to past page visits is available, GoWeB can further provide more satisfactory recommendations. Overall, the results suggest that incorporating people’s motives can lead to better recommendation in web browsing sessions.

⁶The number was selected in accordance with the prior study [2] that analyzed large-scale web data.

5.3 Re-visitation Prediction Performance

For re-visitation prediction, given a browsing session, the task is formulated as a binary classification problem to predict if each web page visit in the session will be re-visited by the same user in any future session. Similarly, we consider two families of baseline methods. The conventional sequence modeling methods learn an embedding vector for each unique web page without using web page content, including CNN, RNN, SAS [27], BERT [15, 38]. To account for page content, we likewise build a strong semantic classifier, SemCLS, based on the sequence modeling component embedded in GoWeB without using any goal-related representations. We utilize common binary classification metrics, including F1-score, precision, recall, and accuracy, for evaluation.

Table 3 shows the classification performance of methods in two test sets. Consistent with the results in Table 2, SemCLS that incorporates web page content is the best-performing baseline method. Compared to SemCLS, GoWeB (NP) respectively provides 2.32% and 3.30% gains in F1 scores for the two test sets. Similar to in-session recommender, GoWeB introduces further gains when using the focal goals to attend to past goals. These results indicate that the GoWeB framework can provide goal-aware assistance in supervised ranking and classification tasks.

5.4 Goal-based Grouping Performance

The objective of the web page grouping task is to create coherent clusters of web pages inside a session in an unsupervised fashion. We treat the goal representations r_p derived in GoWeB as features and apply K -means++ algorithm [4] to cluster page visits into different groups. We consider two popular text modeling methods, `doc2vec` [34] and BERT [15], as the baselines where the corresponding embedding vectors are used for clustering.

To create ground-truth for evaluation, we leverage the outcome of a proprietary hierarchical topical classifier as references that different methods can compare to. The proprietary classifier predicts the topics of web pages according to a 3-layer taxonomy consisting of a root node, 22 category nodes, and 288 leaf nodes. This setup resembles the traditional web page classification task [6, 51] where a topical organization of the web was found to be useful. Once we obtain the outcome topics, the number of ground-truth clusters in each session is also determined accordingly. Note that, following previous studies [33, 67], our focus is to evaluate the clustering quality using different embedding features. We acknowledge past work on determining the optimal number of clusters [30] as a different research problem.

Table 4 shows the resulting grouping performance on NMI and AMI metrics [40]; here we combine the two test sets together since this is an unsupervised task. BERT, as is powered by a decent pretrained contextualized language model, performs better than `doc2vec`. Meanwhile, the goal representations generated by goal estimator outperforms both baselines, suggesting its efficacy for tackling unsupervised tasks. Figure 5 further shows the percentage of clusters over the number of unique ground-truth topics included in a cluster. An ideal cluster should contain exactly one topic.

The results show that GoWeB tends to derive purer clusters with fewer unique classes compared to baselines. It may suggest that the goal-aware representations are also more topically coherent.

		Pop	BPRMF	NCF	Caser	GRU4Rec	SASRec	BERT4Rec	SemRec	GoWeB (NP)	GoWeB
Test (Warm)	MRR@10	0.0386	0.0484	0.0537	0.0506	0.0533	0.0693	0.0727	0.0817	0.0968 (+18.52%)	0.1080 (+32.22%)
	HR@1	0.0169	0.0243	0.0259	0.0216	0.0269	0.0365	0.0382	0.0472	0.0581 (+23.24%)	0.0674 (+42.96%)
	HR@5	0.0618	0.0751	0.0864	0.0870	0.0837	0.1093	0.1150	0.1252	0.1445 (+15.38%)	0.1585 (+26.53%)
	HR@10	0.1123	0.1243	0.1365	0.1329	0.1445	0.1681	0.1748	0.1801	0.2090 (+16.05%)	0.2276 (+26.38%)
	NDCG@5	0.0202	0.0276	0.0295	0.0278	0.0303	0.0398	0.0418	0.0476	0.0563 (+18.13%)	0.0626 (+31.34%)
	NDCG@10	0.0292	0.0354	0.0373	0.0363	0.0406	0.0503	0.0521	0.0571	0.0676 (+18.33%)	0.0751 (+31.43%)
Test (Cold)	MRR@10	0.0454	0.0513	0.0528	0.0508	0.0559	0.0675	0.0683	0.0856	0.0993 (+16.09%)	N/A
	HR@1	0.0215	0.0234	0.0268	0.0251	0.0267	0.0353	0.0337	0.0499	0.0597 (+19.66%)	N/A
	HR@5	0.0752	0.0882	0.0884	0.0829	0.0937	0.1034	0.1094	0.1350	0.1496 (+10.83%)	N/A
	HR@10	0.1206	0.1355	0.1269	0.1340	0.1493	0.1696	0.1761	0.1874	0.2150 (+14.68%)	N/A
	NDCG@5	0.0252	0.0318	0.0316	0.0280	0.0335	0.0392	0.0405	0.0518	0.0598 (+15.30%)	N/A
	NDCG@10	0.0341	0.0398	0.0383	0.0375	0.0435	0.0505	0.0520	0.0611	0.0712 (+16.50%)	N/A

Table 2: Performance of methods in session-based web page recommendation. Note that cold-start users do not have personalized historical goal embeddings from training data, so GoWeB with personalization is not available for the cold-start testing dataset. GoWeB (NP) denotes the non-personalized version of GoWeB.

		F1	Precision	Recall	Accuracy
Test (Warm)	CNN	0.5249	0.4818	0.5765	0.6066
	RNN	0.5387	0.6269	0.4722	0.6951
	SAS	0.5302	0.6376	0.4537	0.6969
	BERT	0.5948	0.5479	0.6505	0.6659
	SemCLS	0.6718	0.6155	0.7394	0.7277
	GoWeB (NP)	0.6874	0.6247	0.7642	0.7380
	GoWeB	0.7010	0.6234	0.8007	0.7425
Test (Cold)	CNN	0.5454	0.5271	0.5650	0.6075
	RNN	0.5369	0.6826	0.4425	0.6819
	SAS	0.5290	0.6956	0.4268	0.6833
	BERT	0.6127	0.6093	0.6160	0.6754
	SemCLS	0.6824	0.6829	0.6820	0.7355
	GoWeB (NP)	0.7050	0.6984	0.7117	0.7518

Table 3: Performance of methods in re-visitation classification. Similar to Table 2, GoWeB with personalization is N/A for cold-start users. GoWeB (NP) denotes the non-personalized version of GoWeB.

	doc2vec	BERT	GoWeB (Goal Estimator)
NMI	0.7290	0.7612	0.7683 (+0.93%)
AMI	0.4320	0.4929	0.5155 (+4.59%)

Table 4: Performance of methods in goal-based clustering.




Figure 5: The percentage of clusters over different numbers of incorporated unique topics.

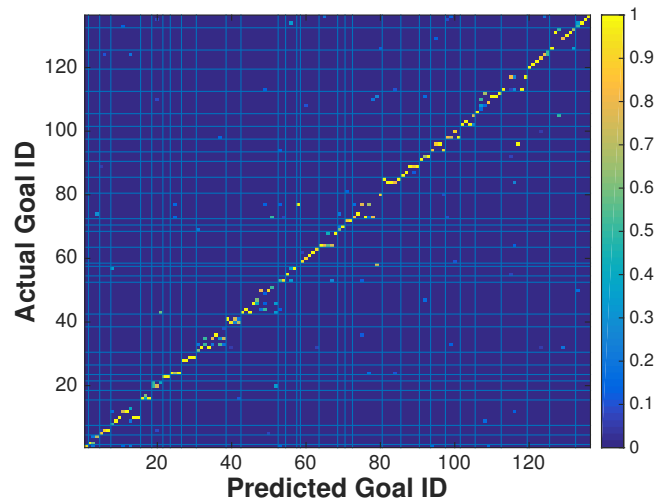


Figure 6: The probabilistic confusion matrix for our goal estimator in goal classification. The grid partitions goals into categories. The “non-goal” class has goal ID 1.

6 Analysis and Discussions

In section 5, we demonstrate empirically the effectiveness of GoWeB when applying to browser applications. Now we turn our attention to deepen our understanding for the effectiveness of the goal estimator, and how the variety of goals may affect the differences observed in people’s behavioral patterns.

Effectiveness of Goal Estimator. To uncover the details, we scrutinize the goal estimator as it powers the prediction for goal representations when people make visits to web pages. Recall that the goal estimator is learned with a multi-class objective, where a 136-way classification is carried out to predict among the 135 leaf goals plus a non-goal class. We use 90% of weakly-labeled instances from D_{weak} for training, and evaluate on the remaining 10%. As a result, we achieve 55.27% and 83.01% F1 scores for predicting individual goals and goal categories. Figure 6 depicts the confusion matrix for our goal estimator, where the grid partitions leaf goals into goal categories. We can see that most of the predictions are

Goal Category	Top 1	Top 2	Top 3
Appearance	alphacute.com (✓)	interesticle.com (✓)	trend-chaser.com (✓)
Career	indeed.com (✓)	careerbuilder.com (✓)	linkedin.com (✓)
Entertainment	en.wikipedia.org	gocomics.com (✓)	nytimes.com
Finances	cnbc.com (✓)	zillow.com (✓)	realtor.com (✓)
Order	microsoft.com (✓)	indeed.com	lowes.com (✓)
Physical Health	healthygeorge.com (✓)	myfitnesspal.com (✓)	financialadvisorheroes.com
Religion	biblegateway.com (✓)	kingjamesbibleonline.org (✓)	bibletrivia.com (✓)

Table 5: Top 3 hosts whose web pages are more likely to belong to some goal categories. Hosts with (✓) are labeled as relevant websites to the corresponding categories. The goal “Order” means “(Keep Things in) Order.”

aligned with the diagonal as correct predictions; among other cases, most of the misclassifications are in the same grid block, indicating that errors are bounded within the same goal category.

We further conduct a lightweight qualitative exercise to analyze the relations between hosts and predicted goals. We first identify top hosts that cover a higher proportion of web pages being predicted closest to each goal category. Two annotators were then asked to judge whether these hosts are topically relevant to the corresponding categories. Table 5 shows the assessment results with 100% agreement from the two annotators. The assessments suggest that the goal estimator can make reasonable predictions in the majority of cases.

Characterization of In-session Goals. To start, we note that people tend to involve in multiple goals in a single session; for example, the average number of goals can be more than 5 for sessions containing 10 page visits. We also find that the increase in the number of goals pursued is sub-linear to the session length; for example, sessions of 60 page visits are associated with fewer than 14 goals on average. Given that people may seek multiple goals in a session, to which degree each goal category could be considered as the session’s focal goal? Figure 7a illustrates the percentage of being the only goal in a single session. On one hand, for sessions that contain a *Finances* goal, about 16% of the times people concentrate solely on this goal category. We conjecture that it might be related to the time-dependent sensitivity required in the process of financial decision making; further studies need to be conducted in the future. On the other hand, we find that more nebulous goals such as *Ethics & Idealism* are seldom being pursued singularly (i.e., only 3% of the times). The results suggest people behave differently according to the types of goals, and fathoming the underlying objectives is key to providing delightful web browsing experiences. For example, enabling a focus mode on detecting goals that require high concentration, or proactively recommending new resource for more exploratory goals could be new designs to consider.

Characterization of Revisitation Patterns by Goals. People’s goals that motivate visits to certain web pages may also have an impact on how these pages are revisited. We analyze temporal cross-session revisitation patterns. A revisit is identified by two consecutive visits of the same web page, where the elapsed time between the former visit and the latter is considered as the duration. Following Adar et al. [2], we group all collected revisitation patterns into 14 unequal-size buckets according to the length of duration. Figure 7b illustrates the distributions of revisits in different buckets for different goal categories. Interestingly, at a high level, all the

Rank	Duration Scale		
	Hours	Days	Weeks
1	Friendship	Romance	(Keep Things in) Order
2	Appearance	Social Qualities	Receiving from Others
3	Entertainment	Career	Career
4	Social Recognition	Self-sufficiency	Stability & Safety
5	Marriage	Religion	Social Qualities

Table 6: Top 5 goal categories with different duration scales in re-visitations.

30 goal categories exhibit similar revisiting patterns. Most of the revisits (37%) occur roughly after a day or within few days, while between 23% to 36% of revisits occur within 7 hours. We observe far less revisiting after a week has passed. This could be due to the limitation of our data being sampled from a month period.

Beyond the distributional similarity at a macro level, we observe distinct patterns in the top 5 goal categories for three revisitation duration scales in Table 6. The results suggest that goal categories concerning social interactions (*Friendship*, *Social Recognition*, and *Marriage*), addictability (*Entertainment*), and timeliness (*Appearance*) can lead to relatively quicker revisits within hours. It may imply that inter-personal relationships tend to urge people to voluntarily stay up-to-date, potentially due to the desire for intimacy or attractive looks, or the fear of missing out. On the contrary, goal categories (*Career* development, *Self-sufficiency*, and *Social Qualities*) that demand continuous investment associate more frequently with slower revisits. More abstract, higher-ordered goals, such as *Order*, *Receiving from Others* and *Stability & Safety* may result in the slowest revisitation compared to more concrete goals such as *Entertainment*. This highlights the opportunities for supporting people browsing the web, where an ambient reminder for a potential revisit can reduce the overhead of manual retrieval.

7 Conclusion

In this paper, we highlight the importance of modeling human motives grounded by the long history of psychology literature. A unified neural framework, GoWeB, is presented to fulfill this vision. We build on the top of existing taxonomy and concertize these goals with structure-preserving representation learning in hyperbolic space, based on which a goal estimator is introduced to tighten the loop of how goals could be employed for enhancing browsing experiences. We showcase the generality of GoWeB and adopt it in three browser-centric applications. On real-world data, GoWeB consistently outperforms competitive baselines for both warm-start and cold-start users, and demonstrates additional gains when using the focal goals to attend to past goals. Our follow-up analysis reveals the effectiveness of the goal estimator via quantitative and qualitative exercises, and characterizes the similarities and differences found in behavioral patterns when people pursue different goals.

Our work brings new perspectives in multiple ways. We present a promising paradigm where we capture the fundamental motives that drive people in their actions and reflect those in digital applications. Broadly speaking, we introduce and transfer the knowledge from psychology findings to modeling browsing sessions on the web, while keeping the framework flexible such that future research could incorporate other types of taxonomies. It is our hope that

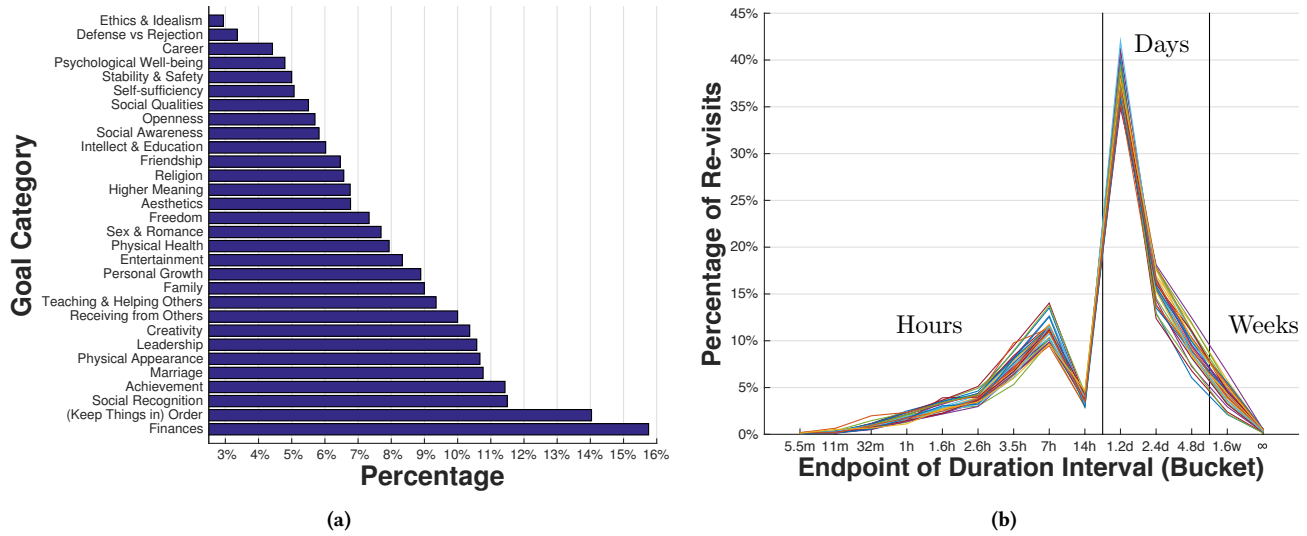


Figure 7: (a) The percentage of being in a web browsing session with a single goal for each goal category. (b) Percentage of re-visits within different duration interval buckets [2] and three duration scales for goal categories. Note that each line represents the statistics of goals in a certain category.

these findings can lift the burden of understanding and characterizing complex human goals for ubiquitous web browsing applications.

References

- [1] Jesper Aagaard. 2015. Drawn to distraction: A qualitative study of off-task use of educational technology. *Computers & Education* 87 (2015), 90–97.
- [2] Eytan Adar, Jaime Teevan, and Susan T. Dumais. 2008. Large Scale Analysis of Web Revisitation Patterns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, 1197–1206.
- [3] Rachel F Adler and Raquel Benbunan-Fich. 2013. Self-interruptions in discretionary multitasking. *Computers in Human Behavior* 29, 4 (2013), 1441–1449.
- [4] David Arthur and Sergei Vassilvitskii. 2006. *k-means++: The advantages of careful seeding*. Technical Report. Stanford.
- [5] Debjanee Barua, Judy Kay, Bob Kummerfeld, and Cécile Paris. 2014. Modelling Long Term Goals. In *User Modeling, Adaptation, and Personalization*. Springer International Publishing, 1–12.
- [6] E. Baykan, M. Henzinger, L. Marian, and Ingmar Weber. 2011. A Comprehensive Study of Features and Algorithms for URL-Based Topic Classification. *ACM Trans. Web* 5 (2011), 15:1–15:29.
- [7] Silvere Bonnabel. 2013. Stochastic gradient descent on Riemannian manifolds. *IEEE Trans. Automat. Control* 58, 9 (2013), 2217–2229.
- [8] Andrei Broder. 2002. A Taxonomy of Web Search. *SIGIR Forum* 36, 2 (2002), 3–10.
- [9] Loredana Caruccio, Vincenzo Deufemia, and Giuseppe Polese. 2015. Understanding user intent on the web through interaction mining. *Journal of Visual Languages & Computing* 31 (2015), 230–236. Special Issue on DMS2015.
- [10] Ada S. Chulef. 1993. A Hierarchical Taxonomy of Human Goals. *Dissertation* (1993).
- [11] Ada S. Chulef, S. Read, and D. Walsh. 2001. A Hierarchical Taxonomy of Human Goals. *Motivation and Emotion* 25 (2001), 191–232.
- [12] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
- [13] Mary Czerwinski, Eric Horvitz, and Susan Wilhite. 2004. A diary study of task switching and interruptions. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 175–182.
- [14] Honghua Dai, Lingzhi Zhao, Zaiqing Nie, Ji-Rong Wen, Lee Wang, and Ying Li. 2006. Detecting online commercial intention (OCI). In *Proceedings of the 15th international conference on World Wide Web*. 829–837.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 4171–4186.
- [16] Alan Dix, Giorgos Lepouras, Akrivi Katifori, Costas Vassilakis, Tiziana Catarci, Antonella Poggi, Yannis Ioannidis, Miguel Mora, Ilias Daradimos, Nazihah Md Akim, et al. 2010. From the web of data to a world of action. *Journal of Web Semantics* 8, 4 (2010), 394–408.
- [17] Mira Dontcheva, Steven M. Drucker, Geraldine Wade, David Salesin, and Michael F. Cohen. 2006. Summarizing Personal Web Browsing Sessions (*UIST '06*). 115–124.
- [18] A. Faaborg and H. Lieberman. 2006. A goal-oriented web browser. In *CHI*.
- [19] Tony Gillie and Donald Broadbent. 1989. What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological research* 50, 4 (1989), 243–250.
- [20] Jeremy Glassman, Marilyn Prosch, and Benjamin BM Shao. 2015. To monitor or not to monitor: Effectiveness of a cyberloafing countermeasure. *Information & Management* 52, 2 (2015), 170–182.
- [21] Murat Göksedef and Sule Gündüz Ögüdücü. 2010. Combination of Web page recommender systems. *Expert Syst. Appl.* 37 (2010), 2911–2922.
- [22] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*. Vol. 1. MIT press Cambridge.
- [23] X. He, Lizi Liao, Hanwang Zhang, L. Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. *Proceedings of the 26th International Conference on World Wide Web* (2017).
- [24] Balázs Hidasi, Alexandros Karatzoglou, L. Baltrunas, and D. Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. *CoRR* abs/1511.06939 (2016).
- [25] Jyun-Yu Jiang, Tao Wu, Georgios Roumpos, Heng-Tze Cheng, Xinyang Yi, Ed Chi, Harish Ganapathy, Nitin Jindal, Pei Cao, and Wei Wang. 2020. End-to-End Deep Attentive Personalized Item Retrieval for Online Content-sharing Platforms. In *Proceedings of The Web Conference 2020*. 2870–2877.
- [26] Rosie Jones and Kristina Lisa Klinkner. 2008. Beyond the Session Timeout: Automatic Hierarchical Segmentation of Search Topics in Query Logs. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management (Napa Valley, California, USA) (CIKM '08)*. 699–708.
- [27] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. *2018 IEEE International Conference on Data Mining (ICDM)* (2018), 197–206.
- [28] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [29] Eric Klingner. 1987. The Interview Questionnaire technique: Reliability and validity of a mixed idiographic-nomothetic measure of motivation. *J. N. Butcher & C. D. Spielberger (Eds.), Advances in personality assessment* 6 (1987), 31–48.
- [30] Trupti M Kodinariya and Prashant R Makwana. 2013. Review on determining number of Cluster in K-Means Clustering. *International Journal* 1, 6 (2013), 90–95.
- [31] Ravi Kumar and Andrew Tomkins. 2010. A Characterization of Online Browsing Behavior. In *Proceedings of the 19th International Conference on World Wide Web*

- (WWW '10). 561–570.
- [32] Edith Law and Haoqi Zhang. 2011. Towards Large-Scale Collaborative Planning: Answering High-Level Search Queries Using Human Computation. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI '11)*. 1210–1215.
- [33] Marc T Law, Raquel Urtasun, and Richard S Zemel. 2017. Deep spectral clustering learning. In *International Conference on Machine Learning*. 1985–1994.
- [34] Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*. 1188–1196.
- [35] Uichin Lee, Zhenyu Liu, and Junghoo Cho. 2005. Automatic Identification of User Goals in Web Search. In *Proceedings of the 14th International Conference on World Wide Web (WWW '05)*. 391–400.
- [36] Xiao Li, Ye-Yi Wang, and Alex Acero. 2008. Learning query intent from regularized click graphs. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*. 339–346.
- [37] Vivien KG Lim. 2002. The IT way of loafing on the job: Cyberloafing, neutralizing and organizational justice. *Journal of organizational behavior: the international journal of industrial, occupational and Organizational Psychology and Behavior* 23, 5 (2002), 675–694.
- [38] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [39] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 1412–1421.
- [40] Christopher D Manning, Hinrich Schütze, and Prabhakar Raghavan. 2008. *Introduction to information retrieval*. Cambridge university press.
- [41] Abraham H. Maslow. 1970. Motivation and personality (2nd ed.). *New York: Harper & Row* (1970).
- [42] William McDougall. 1933. *The Energies of Men*. *New York: Scribner's* (1933).
- [43] Daniel C McFarlane and Kara A Latorella. 2002. The scope and importance of human interruption in human-computer interaction design. *Human-Computer Interaction* 17, 1 (2002), 1–61.
- [44] Gordon B. Moskowitz and Heidi Grant. 2009. The psychology of goals.
- [45] H. A. Murray. 1938. Explorations in personality : a clinical and experimental study of fifty men of college age. (1938).
- [46] Maximilian Nickel and Douwe Kiela. 2017. Poincaré Embeddings for Learning Hierarchical Representations (*NIPS'17*). 6341–6350.
- [47] Hartmut Obendorf, H. Weinreich, E. Herder, and Matthias Mayer. 2007. Web page revisitation revisited: implications of a long-term click-stream study of browser usage. In *CHI*.
- [48] Brid O'Conaill and David Frohlich. 1995. Timespace in the workplace: Dealing with interruptions. In *Conference companion on Human factors in computing systems*. 262–263.
- [49] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*. 8026–8037.
- [50] Dipak Patel and Mukesh Zaveri. 2011. A Review on Web Pages Clustering Techniques.
- [51] Xiaoguang Qi and Brian D. Davison. 2009. Web page classification: Features and algorithms. *ACM Comput. Surv.* 41 (2009), 12:1–12:31.
- [52] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI '09)*. 452–461.
- [53] Milton Rokeach. 1973. *The nature of human values*. *New York: Free Press* (1973).
- [54] Daniel E. Rose and Danny Levinson. 2004. Understanding User Goals in Web Search. In *Proceedings of the 13th International Conference on World Wide Web (New York, NY, USA) (WWW '04)*. 13–19.
- [55] James Y Shah, Ron Friedman, and Arie W Kruglanski. 2002. Forgetting all else: on the antecedents and consequences of goal shielding. *Journal of personality and social psychology* 83, 6 (2002), 1261.
- [56] Dustin A Smith and Henry Lieberman. 2010. The Why UI: using goal networks to improve user interfaces. In *Proceedings of the 15th international conference on Intelligent user interfaces*. 377–380.
- [57] Markus Strohmaier and Mark Kröll. 2012. Acquiring knowledge about human goals from Search Query Logs. *Information Processing & Management* 48, 1 (2012), 63 – 82.
- [58] Xiaoyuan Su and Taghi M. Khoshgoftaar. 2009. A Survey of Collaborative Filtering Techniques. (2009).
- [59] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer (*CIKM '19*). 1441–1450.
- [60] Jennifer Talevich, Stephen J. Read, David A. Walsh, Ravi Iyer, and Gurveen Chopra. 2017. Toward a comprehensive taxonomy of human motives. *PLoS ONE* 12 (2017).
- [61] Jiayi Tang and Ke Wang. 2018. Personalized Top-N Sequential Recommendation via Convolutional Sequence Embedding. *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (2018).
- [62] Jaime Teevan, Susan T Dumais, and Daniel J Liebling. 2008. To personalize or not to personalize: modeling queries with variation in user intent. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*. 163–170.
- [63] Vincent W-S Tseng, Matthew L Lee, Laurent Denoue, and Daniel Avrahami. 2019. Overcoming distractions during transitions from break to work using a conversational website-blocking system. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*. 5998–6008.
- [65] Frank Wicker, Frank Lambert, Frank Richardson, and Joseph Kahler. 1984. Categorical goal hierarchies and classification of human motives. *Journal of Personality* 52 (1984), 285 – 305.
- [66] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *ArXiv* (2019), arXiv–1910.
- [67] Xu Yang, Cheng Deng, Feng Zheng, Junchi Yan, and Wei Liu. 2019. Deep spectral clustering using dual autoencoder network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4066–4075.
- [68] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep Learning Based Recommender System: A Survey and New Perspectives. 52, 1 (2019).