Predicting Intent Using Activity Logs: How Goal Specificity and Temporal Range Affect User Behavior

Justin Cheng*
Stanford University
jcccf@cs.stanford.edu

Caroline Lo*
Stanford University
clo@cs.stanford.edu

Jure Leskovec
Pinterest, Stanford University
jure@cs.stanford.edu

ABSTRACT

People have different intents in using online platforms. They may be trying to accomplish specific, short-term goals, or less well-defined, longer-term goals. While understanding user intent is fundamental to the design and personalization of online platforms, little is known about how intent varies across individuals, or how it relates to their behavior. Here, we develop a framework for understanding intent in terms of goal specificity and temporal range. Our methodology combines survey-based methodology with an observational analysis of user activity. Applying this framework to Pinterest, we surveyed nearly 6000 users to quantify their intent, and then studied their subsequent behavior on the web site. We find that goal specificity is bimodal – users tend to be either strongly goal-specific or goal-nonspecific. Goal-specific users search more and consume less content in greater detail than goal-nonspecific users: they spend more time using Pinterest, but are less likely to return in the near future. Users with short-term goals are also more focused and more likely to refer to past saved content than users with long-term goals, but less likely to save content for the future. Further, intent can vary by demographic, and with the topic of interest. Last, we show that user's intent and activity are intimately related by building a model that can predict a user's intent for using Pinterest after observing their activity for only two minutes. Altogether, this work shows how intent can be predicted from user behavior.

Keywords

Goal-directed behavior; goal setting; user motivations; Pinterest

1. INTRODUCTION

People have different intents in using online platforms – they may search for information using Google, connect with others using Facebook, or buy products on Amazon. These intents, or goals that people aim to achieve, can also differ for the same service – visitors to Google may intend to navigate to a desired site or learn about a particular topic [5]; Facebook users may aim to maintain social connections or simply consume content [20]; users may visit Amazon to purchase a particular product, or explore what is available [19]. More generally, a user's intent strongly influences their behavior [1], and affects the actions that they take.

*Research partly done while at Pinterest.

WWW'17 Companion, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4914-7/17/04.
http://dx.doi.org/10.1145/3041021.3054198

An understanding of user intent can help designers surface different interaction modes [38], provide better contextual help [41], and personalize search results and recommendations [43]. For example, a user who started working on a long-term weekend do-it-yourself project may appreciate ways to track their progress and reminders to continue working on it the following week; a user looking for recipes for dinner tonight may prefer being shown only main course recipes that are easy to make and that have familiar ingredients, but when looking for potluck ideas for next week may instead prefer being shown recipes that are generally uncommon.

Intent has been studied both broadly in terms of goal-setting [4, 16], and specifically within different domains [5, 23]. The former line of work has measured the effect of intent on general outcomes such as performance [25], but has tended to focus on high-level differences in behavior rather than the relation of intent to more granular user activity. Further, this work has mostly relied on smaller-scale surveys [17] and controlled laboratory studies [27]. The latter line of work has performed large-scale studies and proposed models for predicting specific actions (e.g., clicking on a search result [36, 43], or purchasing a product [24, 33]), but these models are hard to generalize to other settings. What is missing is a large-scale methodology that can link a wide range of intents with specific user behaviors, and that can be used in a broad range of online settings. By identifying user intent generally and connecting it to specific behavioral patterns, we can develop better models of user behavior and build platforms that are more personalized to an individual's needs.

The present work: a conceptual framework for intent. In this work, we focus on two key dimensions of intent: goal specificity and temporal range. We present a generalizable methodology for studying these dimensions of intent and how they relate to user behavior. Goal specificity relates to the extent to which a person visits a web site for a specific purpose (e.g., to look for gluten-free chocolate-chip cookie recipes, to find gardening tips, or just to kill time), while temporal range relates to the time horizon for goal completion (e.g., recipe ideas for dinner tonight, for the weekend, or for next month). We study intent in the context of content discovery (i.e., browsing, searching for, and consuming online media), and in particular, apply our methodology to Pinterest, a content sharing platform with over 150 million monthly active users [40]. To elicit intent in terms of goal specificity and temporal range, we conducted a survey involving nearly 6000 users on Pinterest; to understand how these users behaved on the platform, we then analyzed complete traces of these users' activity. By connecting this survey data with actual usage patterns, we can study how intent informs behavior, as well as predict intent by observing behavior of a user for just a few minutes.

Summary of results. Our methodology reveals that goal specificity is strongly bimodal – users are either very goal-specific, or goal-nonspecific, suggesting two distinct modes of site use, corroborating prior qualitative work [23]. We find that goal-specific users are more focused. They engage in more task-oriented activities such as searching, spend more time browsing specific categories of content in detail, and are also more likely to reference past saved content. But while goal-specific users spend more time browsing the site, they are also less likely to return in the next seven days.

Studying temporal range, we find that users with short-term goals, like goal-specific users, also focus on specific categories of content in detail. But while these users are similarly more likely to reference past saved content, they are less likely to save new content for the future. Temporal range also correlates with goal specificity: goal-specific users are likely to act in the short-term, while goal-nonspecific users are more unsure of taking action.

Gender, age, and categorical differences in intent also exist. Female users are more likely to be goal-specific on Pinterest, while male users are more unsure of taking action. Older users tend to have short-term goals, while younger users are more unsure of taking action. Users interested in food and drink or DIY are more likely to be goal-specific and act in the short-term, while users interested in travel and entertainment tend to be goal-nonspecific and are likely to act in the long-term. An analysis of recipe-saving (a popular use of Pinterest [18]) shows differences in the types of recipes users look up depending on their intent.

Our methodology further connects user behavior back to user intent. In particular, we show that a user's intent can be reliably predicted based on their actions. Notably, intent can be predicted early in a user session. We develop a model that can predict both goal specificity (AUC=0.76) and temporal range (AUC=0.70) within two minutes of a user logging in, with performance improving as more of a user's session is observed. We find that activity in the current session alone is sufficient in predicting intent, but lacking those signals, historical activity and demographics can also be used to infer intent. As user intent can vary with each visit, quickly identifying broad intent can help support relevant changes to the interface and content presented to the user on-the-fly.

In summary, we:

- propose goal specificity and temporal range as a way to study user intent,
- demonstrate how intent translates into subsequent behavior on an online content discovery platform, Pinterest, and
- develop a model that accurately and quickly predicts intent from user behavior.

2. BACKGROUND

To begin, we review literature on intent, and focus on the dimensions of goal specificity and temporal range (i.e., when a goal will be achieved), as well as work on predicting user behavior.

Intent and Goals. Intent precedes any behavior, and is one of the strongest predictors of future behavior [1]. In this work, we study intent in the context of goal-setting. Prior work studying goals and goal-setting behavior has typically focused on establishing the relationship between a goal's dimensions [4] (e.g., difficulty or specificity) and high-level behavioral outcomes (e.g., performance [47] or productivity [25]). A majority of this work tends to fall into one of two categories – those that are interview or survey-based and where behavior is self-reported [17, 37], and those that do measure behavior, but at small-scale and typically in laboratory settings [10, 27]. Other work studied intent along with an individual's beliefs, attitudes, and norms [2] in the context of behavior change (e.g., technology adoption [15]), but these were also mainly survey-based.

In the current work, we focus on two key dimensions of intent: goal specificity and temporal range. Together, these dimensions cover a wide variety of motivations seen in prior work [39] (e.g., on Pinterest [23]). In contrast to prior work, we propose a methodology that allows us to, at larger scales, survey user intent, and then objectively measure their behavior.

Goal Specificity. Goals vary in their *specificity* [16, 25], and may be more specific and quantitative, or less specific and qualitative. Specific goals lead to more focused behavior [26], while less well-defined goals lead to greater variability in behavior, and thus, performance [27], as many outcomes can be consistent with vague goals [16]. While specificity can covary with task difficulty [29], we focus on the former, and see consideration of the latter as future work.

Applying goal specificity to particular domains, research has characterized differences between navigational and exploratory web searching [28]. Similarly, online shopping has generally been dichotomized as either being goal-directed or experiential [44]. In the specific case of Pinterest, research that studied motivations identified two primary modes of use – "casual browsing with no particular goal in mind", or "responding to a specific task" (e.g., finding a new hairstyle) [23]. Drawing on this rich literature, we broadly characterize goal specificty on a range – from being *goal-nonspecific* (or having less defined, abstract goals) to being *goal-specific* (or having specific, concrete goals).

Temporal Range. *Temporal range*, or when a goal will be achieved, is another important dimension of intent [4, 14]. Goals can be oriented towards either the short-term, or the long-term; past work has argued that because people are future-oriented, motivations must include a temporal aspect [34]. Temporal range can also indicate the level of abstraction of a goal – shorter-term goals tend to be more explicit (e.g., hunger satisfaction), while longer-term goals tend to be more abstract (e.g., future aspirations) [35].

People also set goals depending on their future time perspective, or their perceptions of whether time is limited or expansive. For this reason, older people tend to pursue more emotionally meaningful goals, and spend less time than younger people gathering new information and expanding horizons [7].

In case of Pinterest, a recent survey showed that Pinterest is used to achieve short-term goals such as preparing daily meals and long-term goals such as planning vacations [18], suggesting that it may be a good platform for studying temporal range.

Predicting User Behavior. By understanding user intent, we can better understand, and even predict user behavior [2]. For example, goal-oriented shoppers are more likely to browse specific products and directly search for them, while "experiential" shoppers are more likely to browse more products [9]. Recommender systems can take advantage of understanding temporal range of intent in order to provide better contextual recommendations [45]. But in addition to being mainly interested in purchasing intent, many online shopping studies are largely survey-based [33, 39] or conducted in laboratory settings [9], and do not seek to predict intent.

Large-scale studies on predicting future behavior or intent have been relatively domain-specific, for example, focusing on information search intent by modeling user interests [43], the likelihood of clicking on search results [46], or purchase intent [9, 24].

Altogether, prior work on intent and goal-setting has tended to be rely on self-reported behavior or be limited to small-scale laboratory settings, and has mainly focused on relating intent to behavior at a high level. When prediction have been involved, it has typically been of domain-specific actions, which makes generalization difficult. In this work, we propose a scalable methodology that relates intent, characterized in terms of both goal specificity and temporal range, to an individual's specific behavioral traces. We use a qualitative survey to elicit intent, and perform quantitative data analysis of user activity to measure behavior. This then allows us to develop predictive models of intent.

3. METHOD

In this paper, we present a conceptual framework for studying user intent on online platforms. To elicit intent explicitly, we first design and conduct a user survey, which builds on fundamental work in modeling user intent [2, 15, 25]. To then see how intent influences subsequent behavior, we also measure these surveyed users' behavior as they continue to use the platform.

Survey. The first component of our methodology is a user survey. Here, we conducted two identical surveys of English-speaking Pinterest users residing in the United States during the month of July 2016: one on female users (N=5369, mean age of 33.5), and a smaller one on male users (N=564, mean age of 38.6). Users navigating to the Pinterest web site were invited to complete the survey via a modal popup, which could be dismissed. Each survey consisted of four questions. Two questions asked about goal specificity and temporal range ("are you visiting Pinterest today with a goal in mind?" and "when are you planning to act on what you're looking at today?"). To further substantiate our findings, and drawing from motivations identified in prior work specific to social curation web services [23], a third question asked users about their Pinterest-specific motivations (e.g., to look for ideas or inspiration). We discuss these in relation to goal specificity and temporal range where relevant. The last question surveyed users about which categories they were interested in ("what are you looking for on Pinterest today?"). These categories were derived from the categorization that Pinterest uses for categorizing its pins. Finally, these survey responses were then matched to server logs of user activity.

To minimize coverage bias, the surveys were conducted using probability sampling (i.e., all users visiting Pinterest were equally likely to receive the survey). To test for participation bias, we compared the activity of users who completed the survey to those who saw the modal popup but did not complete the survey. Overall, surveyed users tended have been using Pinterest longer and have saved more pins, suggesting that engaged users were more likely to complete the survey. And as users completed the survey at the beginning of their browsing session, we can only measure predetermined intent, but not changes in intent over time.

Observational data: User behavior on Pinterest. The second component of our methodology involves capturing behavioral data on an online platform. In the case of Pinterest, users can view, as well as save *pins* (i.e., pieces of content) to *boards* (i.e., collections), which are typically organized by category or purpose. These pins are organized into 33 different categories (e.g., food and drink, DIY, travel). Further, the user may look at *closeups* of pins, which provide additional information, or alternatively *click through* to visit the original web site that content originated from. In this work, we focus on user activity in the browsing session immediately following the survey. We primarily analyze the first ten minutes of this session as this approximately corresponds to peak performance in predicting intent from user behavior, suggesting that this time window is most informative. 83% of users have sessions longer than 10 minutes. Notwithstanding, we also consider shorter, as well as longer periods of observations. Measuring activity in a fixed amount of time also allows our methodology to more fairly identify behavioral differences across users. Within these ten minutes, we observed over 850 thousand individual behavioral events across 5933 users (e.g., views, closeups, searches). We additionally analyzed past and future user activity, from the time a user first created an account to a week following survey completion.

All survey and observational data was de-identified and analyzed in aggregate. In the rest of the paper, we report the results on both male and female users in aggregate, and note any significant differences between male and female users appropriately. As we make multiple comparisons, we report Holm-corrected p-values. Error bars on plots represent the standard error of the mean or 95% confidence interval of a proportion.

4. DIMENSIONS OF INTENT

Goal specificity and temporal range can together describe a wide range of intent. Goal specificity can tell us how purposeful a user's visit is, while temporal range can tell us when a user plans to take action based on their visit. In this section, we study both dimensions of intent and how they influence a user's activity on Pinterest.

4.1 Goal Specificity

Goal specificity relates to how purposeful or goal-specific a user's intent is. User intent may be strongly specified or *goal-specific* (e.g., finding a recipe for chicken pot pie), or instead less well-defined or *goal-nonspecific* (e.g., to kill time). Survey participants rated how goal-specific their current visit to Pinterest was on a seven-point Likert scale. We define lower values (1 to 3) to be goal-nonspecific (i.e., just browsing), and higher values (5 to 7) to be goal-specific (i.e., looking for something specific). Considering only the extremes of the scale lead to qualitatively similar findings, and accentuates the differences we observe. As later described, goal specificity and temporal range are correlated, so we performed additional regression analysis (not reported for brevity) to control for temporal range.

Goal specificity is bimodal. As Figure 1a shows, goals vary along a spectrum from being non-specific, to being moderately specific, to being very specific. Nonetheless, surveyed ratings were more likely to be at the two extremes of the scale – 29% of users were strongly goal-specific (i.e., they picked the maximum value of 7 on the Likert scale), while 31% are strongly goal-nonspecific (i.e., they picked the minimum value of 1). In other words, most users are either visiting Pinterest with a very specific goal in mind, or no particular goal in mind. Examining gender, female users are more likely to be goal-specific (49% vs. 40%, χ^2=12.0, p<0.01) than male users.

Studying users' more specific motivations, we find that goal-nonspecific users tended to use Pinterest to fill time. More goal-nonspecific users cited boredom as their motivation for visiting than goal-specific users (25% vs. 2%, χ^2=646, p<10^{-3}). Instead, goal-specific users were more likely to be visiting Pinterest to make something (26% vs. 5% for goal-nonspecific users, χ^2=275, p<10^{-3}).

This bimodality in goal specificity suggests two primary modes of using Pinterest – goal-specific use and goal-nonspecific use, and points towards an observation from a prior interview study that found that most users were either "responding to a specific task" or "casually browsing with no particular goal in mind" [23]. Nonetheless, that a substantial proportion of users report moderate goal specificity suggests that future work could study it at greater granularity.

Goal-specific users are task-focused, searching more while consuming less content in greater detail. Defined goals direct attention to relevant tasks and away from irrelevant tasks [26], and goal-specific users do tend to engage in more focused behavior. Goal-specific users searched more than goal-nonspecific users (a

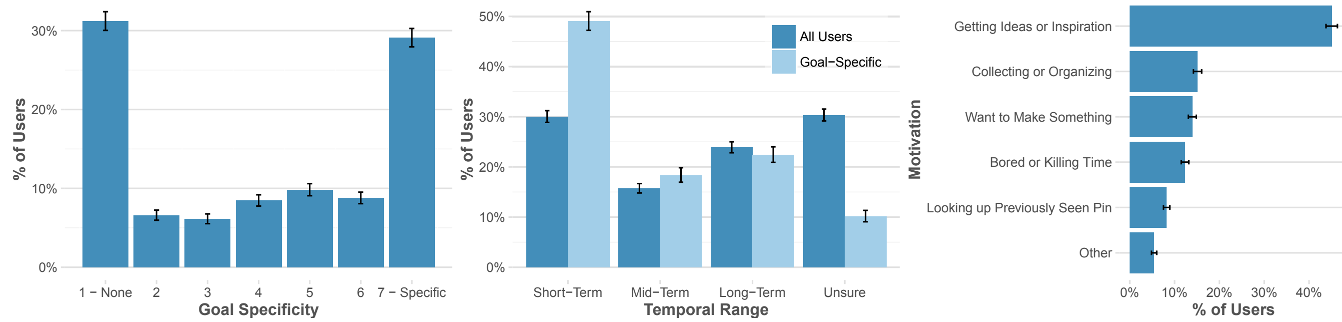

(a) Goal Specificity

(b) Temporal Range

(c) Motivations

Figure 1: In two surveys of nearly 6000 Pinterest users on their intentions for their current visit, we find that (a) goal specificity is largely bimodal, or that users are either goal-specific (e.g., learning how to make something) or goal-nonspecific (e.g., killing time). (b) Further, while temporal range is more varied, goal-specific users tend to act in the short-term (e.g., finding recipes for dinner tonight). (c) Looking at specific motivations for using Pinterest, we find that many use the service for ideas and inspiration.

mean of 1.1 vs. 0.4 searches; using an unequal variances t-test, $t(5162)=15.5$, $p<10^{-3}$, Cohen's $d=0.42$). Goal-specific searches also tended to be more complex – among search queries issued by either group, goal-specific users used more words per query than goal-nonspecific users (3.0 vs. 2.7 words, $t(858)=4.6$, $p<10^{-3}$, $d=0.24$). And because goal-specific users know what they want, they also started searching more quickly than goal-nonspecific users (1.9 vs. 3.0 minutes from signing in, $t(719)=7.3$, $p<10^{-3}$, $d=0.42$).

Goal-specific users also browsed less content, but in greater detail. They looked at less content than goal-nonspecific users (144 vs. 156 viewed pins, $t(5135)=2.8$, $p<0.05$, $d=0.08$) in the same amount of time. Further, goal-specific users also click through to a pin's original source more often (0.9 vs. 0.7 click-throughs, $t(5423)=5.6$, $p<10^{-3}$, $d=0.15$), suggesting that they are more likely to examine content in detail. These results corroborate and extend prior work that found that goal-specific purchasers spent more time per product viewed than experiential purchasers, but not a difference in the number of products browsed [9].

Further, goal-specific users were more focused in terms of subject matter. Controlling for the number of pins viewed, we find that goal-specific users viewed pins belonging to significantly fewer categories than goal-nonspecific users (8.7 vs. 10.4 categories; using a Wilcoxon signed-rank test, $V>10^6$, $p<10^{-3}$, effect size $r=0.35$). Goal-specific users were also less likely to be browsing entertainment-related content (e.g., film, music and books or humor) ($t>2.1$, $p<0.05$, $d>0.05$).

Taken together, these findings suggest that goal-specific users are more task-focused. Rather than browsing Pinterest generally, goal-specific users are more particular about the content they examine, and rely more on search to look for what they want. And though goal-specific users view less content overall, they instead spend their time on the content they do select in greater detail.

Goal-specific users reference past saved content. Many content discovery platforms allow users to save content that they are interested in. On Pinterest, users can save content in the form of pins to curated boards. As such, differences may exist in how goal-specific and goal-nonspecific users split their time between browsing for new content and engaging with previously saved content. Referencing past saved content could be more goal-specific, if a user is looking up something in particular, or less well-defined, if a user is generally reminiscing about their past activity. Here, we find evidence for the former. Goal-specific users were more likely to cite looking up a pin they previously saw as their motivation for visiting (22% vs. 6%, $\chi^2=274.8$, $p<10^{-3}$). In fact, 75% of users who viewed a closeup of a pin or board previously saved were goal-specific. Of users who viewed at least one closeup or board, the proportion of viewed closeups or boards that were saved or created previously was greater for goal-specific than goal-nonpsecific users (19% vs. 7% for closeups, $t(3589)=12.8$, $p<10^{-3}$, $d=0.40$; 74% vs. 52% for boards, $t(740)=7.64$, $p<10^{-3}$, $d=0.47$).

However, are all goal-specific users looking at more past saved content, or is a larger proportion of goal-specific users only looking at past saved content? The distribution of the proportion of closeups that are of previously saved pins seems to indicate the latter (Figure 2a). That is, there may be two distinct types of goal-specific use: either to look up new content, or to refer to past saved content. Now, we divide goal-specific users into those where less than half of the closeups they looked at was of past saved content (i.e., users looking up new content), and those where at least half was of past saved content (i.e., users referencing past saved content). Corroborating this dichotomy, we find that the latter group of users tends to have more boards for organizing content (39 vs. 33 boards, $t(545)=2.7$, $p<0.05$, $d=0.13$), and is less likely to be searching for additional content (0.3 vs. 1.3 searches, $t(1535)=16.6$, $p<10^{-3}$, $d=0.59$).

Goal-specific users spend more time using the service, but are less likely to return soon. Experiential intent, rather than goal-directed intent, correlates with achieving flow, which can lead to time distortion (i.e., time appearing to pass quickly) [33]. Similarly, we might expect that goal-nonspecific users may be more likely to spend greater amounts of time simply browsing the service, as opposed to goal-specific users who are there for a singular purpose. However, we instead find that goal-specific users tend to spend more time on Pinterest than goal-nonspecific users. 48% of goal-specific users spent more than half an hour using the service, compared to only 42% of goal-nonspecific users ($\chi^2=20.8$, $p<10^{-3}$).

On one hand, we might hypothesize that goal-specific users, having satisfied their initial goals, may switch to casually browsing Pinterest like goal-nonspecific users. On the other hand, goal-specific users, having a more specific reason for visiting Pinterest, may end up spending more time on their topic of interest. Previously, we found that goal-specific users tended to view pins in fewer categories than goal-nonspecific users. Here, we test whether this continues to be the case over time. Thus, we compare the number of categories viewed pins belonged to in the first five minutes of activity with the subsequent five minutes, and control for the total number of viewed pins. In the first five minutes, as expected, goal-specific users viewed pins that belonged to fewer categories than goal-nonspecific users (8.2 vs. 9.8 categories, $t(3150)=10.2$, $p<10^{-3}$, $d=0.37$). In the next

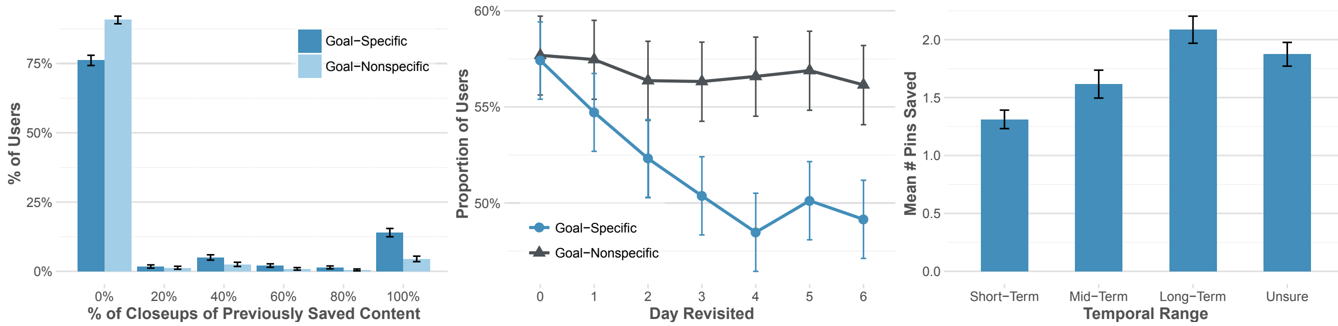

(a) Closeups of Previously Saved Content

(b) Likelihood of Revisiting Pinterest

(c) Amount of Content Saved

Figure 2: (a) The distribution of the proportion of closeups that were of content a user previously saved is bimodal for goal-specific users, suggesting that these users are either primarily looking for new content, or referring to past saved content. (b) Goal-specific users are less likely to revisit Pinterest on subsequent days. (c) Users with short-term goals save the least pins, while those with long-term goals save the most.

five minutes however, while both groups viewed fewer pins overall, the number of categories goal-specific users viewed decreased more than for goal-nonspecific users (6.2 vs. 8.3 , $t(3138)=12.0$, $p<10^{-3}$, $d=0.43$), suggesting increased specificity in what goal-specific users are looking for. Further, despite viewing fewer pins, goal-specific users did not click through on pins significantly less (0.58 vs. 0.53 clickthroughs, n.s.), while goal-nonspecific users did (0.52 vs. 0.42 clickthroughs, $t(3053)=3.4$, $p<0.01$, $d=0.12$). In other words, rather than switching to casual browsing, goal-specific users appear to be more strongly focusing on their goals over time.

But while goal-specific users spend more time on Pinterest, which may indicate greater engagement, are they more likely to return in the near future? As Figure 2b shows, initially, the likelihood of using Pinterest again on the same day is similarly high (71%) for both goal-specific and goal-nonspecific users. However, on subsequent days, goal-specific users are less likely to visit Pinterest (e.g., 50% vs. 56% on the third day, $\chi^2=16.0$, $p<10^{-3}$), suggesting that goal-specific users' visits may be driven by a particular need, while goal-nonspecific users' visits may be more habitual. Comparing gender, goal-specific female users were more likely to return in the near future than goal-specific male users (e.g., 52% vs. 40% on the third day, $\chi^2=7.8$, $p<0.05$).

Past behavior suggests future goal specificity. Past behavior can predict future intention [3]. While intent can vary from session to session, an individual's goals may persist over longer periods of time. Looking at activity from the past 24 hours prior to the current session, we find that while goal-specific users are not any more likely to have visited Pinterest, they tended to have been browsing less content (311 vs. 263 views, $t(5432)=2.4$, $p<0.05$, $d=0.06$) and viewing pins in fewer categories (controlling for the number of viewed pins, 7.9 vs. 8.4 categories, $V>10^6$, $p<10^{-3}$, $r=0.14$). Thus, goal-specific behavior can be reflected in past sessions, and can inform behavior in future sessions.

4.2 Temporal Range

Temporal range corresponds to when a user anticipates a goal will be accomplished. Goals may be oriented towards the shorter-term future (e.g., tomorrow), or the longer-term future (e.g., the end of the week, perhaps on an indefinite timescale) [14]. Understanding temporal range allows us to understand the urgency of a visit – users looking for recipes to make right away may behave differently from users who are looking for recipes to use sometime in the future, who in turn are likely to behave differently from users who save recipes with little intention of making them. Thus, we asked Pinterest users if they planned to take action on what they were doing on Pinterest in the short-term (defined as within two days), the medium-term (within three to seven days), the long-term (a week or more), or if they were unsure of taking (or not intending to take) action. In this section, we focus on comparing users with short-term goals, long-term goals and those unsure of taking action; observations for users with medium-term goals tend to fall between users with short-term and long-term goals.

Temporal range varies significantly. In contrast to goal specificity's bimodality, temporal range is relatively varied on Pinterest. A third of all users had short-term goals, and another third was unsure of taking action (Figure 1b). Male users were significantly more likely than female users to be unsure of taking action (46% vs. 28%, $\chi^2=55.6$, $p<10^{-3}$). Relating temporal range to specific motivations, a majority of users with long-term goals were looking for ideas or inspiration (56% of users with long-term goals), while users with short-term goals were either looking for ideas or wanting to make something (41% and 26% respectively).

Temporal range correlates with goal specificity. Long-term goals tend to be more abstract and less specific, while shorter-term goals tend to be more concrete and more specific [35]. 49% of goal-specific users planned to act in the short-term (Figure 1b), while 52% of goal-nonspecific users were unsure of taking action. In fact, goal specificity correlates positively with having short-term goals (Pearson's $r=0.42$, $t(5931)=35.8$, $p<10^{-3}$), and negatively with both having long-term goals ($r=-0.31$, $t(5931)=2.4$, $p<0.05$) and being unsure about taking action ($r=-0.45$, $t(5931)=39.9$, $p<10^{-3}$).

Nonetheless, being goal-specific does not necessarily imply taking action in the short-term (e.g., a user may be looking for a coffee table for their new home), and conversely, having short-term goals does not imply being goal-specific (e.g., a user may be looking for something to do during the weekend, but not decided on exactly what). Thus, to isolate the effect of temporal range, we have to disentangle the effects of goal specificity. As such, while this and the previous subsections report differences in one dimension, we also performed regression analysis using both temporal range and goal specificity to ensure that any observations reported are not due to interactions between them.

Users with short-term goals also have greater task-focus, and look at more past saved content but save less new content. Where greater goal specificity suggests that users know what they want, shorter temporal range suggests that users want to do something soon. Given the implication of greater urgency in the latter case, users with short-term goals may also exhibit greater task focus in their activity on Pinterest. Though temporal range does not have a significant effect on either the number of searches or the average search query length, users with short-term goals were more likely to click through to a pin than other users (0.96 vs. 0.70 click-throughs, $t(3183)>6.9$, $p<10^{-3}$, $d=0.20$), even after controlling for goal specificity. Users with short-term goals also viewed pins in fewer categories overall, after controlling for the number of pins viewed (8.6 vs. 9.1 categories, $V>10^5$, $p<10^{-3}$, $r=0.19$). Thus, users with short-term goals are more discriminative about the content they examine but examine it in greater detail.

Similarly to our analysis of goal specificity in the previous section, we also study the effect of temporal range on how users save content and reference past saved content. When a user looks up information they previously saved, they likely intend to use that information right away. As such, we might expect that users with short-term goals are more likely to reference past saved content. At the same time, given that short-term goals indicate a sense of immediacy, users may also be less likely to save new content, being less likely to be thinking about the long-term future.

In addition to users with short-term goals being most likely to report looking up previously seen pins as their motivation for visiting, we find that they are more likely to view closeups of pins they previously saved (20% of closeups are of previously saved pins), or boards that they own (79% of boards viewed are their own) than users with long-term goals (10% and 57% respectively, $t>6.3$, $p<10^{-3}$, $d>0.30$). Like goal-specific users, users with short-term goals can also be divided into those looking for new content or referring to past saved content. Users unsure of taking action are least likely to look at closeups of previously saved pins and their own boards (6% for closeups, 50% for boards, $t>3.6$, $p<0.05$, $d>0.16$).

Contrasting with goal specificity which does not significantly influence how much users save, users with short-term goals save fewer pins than users with long-term goals or who are unsure about taking action (1.3 vs. 2.1 and 1.9 pins saved respectively, $t>4.3$, $p<10^{-3}$, $d>0.14$). Notably, users with long-term goals, being the most future-oriented, saved the most. Users with long-term goals were also more likely than users with short-term goals to state collecting or organizing as their motivation for visiting Pinterest (16% vs. 10%, $\chi^2=24.9$, $p<10^{-3}$).

In sum, users with short-term goals tend to look up information to use immediately, while those with long-term goals or who are unsure of taking action appear to be more future-oriented, and more likely to save what they find.

Users with short-term goals spend more time using the service. One might expect that an action taken in the short term suggests shorter deadlines, and thus a more rapid work pace [6, 22], which would suggest that short-term goals lead to shorter sessions. In contrast, we find that a greater proportion of users with short-term goals spend more than half an hour using Pinterest than those with long-term goals (48% vs. 42%, $\chi^2=13.7$, $p<10^{-3}$). Importantly, these differences remain significant even when controlling for goal specificity. We also might expect that users with long-term goals may be more likely to return at a future date to continue working towards their goals. However, unlike goal specificity, temporal range has no significant effect on return visits in the next seven days.

Older people realize goals in the short-term; younger people are less certain. Prior research also suggests that people adjust their time horizons with increasing age, as they increasingly perceive their future time as being more limited, and that conversely, younger people are more likely to expend time exploring their options [7]. Indeed, age is positively correlated with having short-term goals ($r=0.03$, $p<0.05$), and negatively correlated with being unsure about taking action ($r=0.05$, $p<10^{-3}$). In other words, older users are more likely to focus on accomplishing short-term goals, and younger users

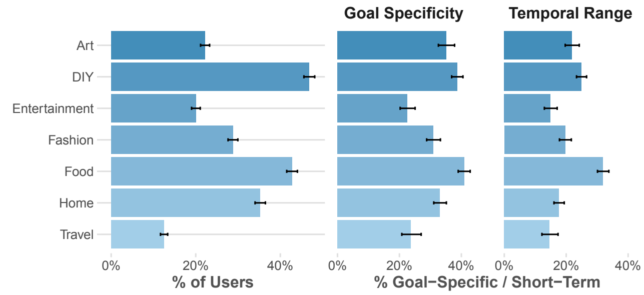

(a) Category Distribution (b) Intent by Category

Figure 3: (a) Surveyed users were most interested in food and drink or DIY. (b) Intent differs by category. For example, users interested in food and drink tend to be the most goal-specific and most likely to have short-term goals.

think less about a goal's timeframe for completion. We note that there are no significant differences with respect to goal specificity.

5. CATEGORY AND INTENT

In this section, we analyze how intent and behavior varies by category. As a case study, we consider recipe-finding on Pinterest, which lets us study in greater detail how intent may further influence category-specific behaviors (e.g., the type of food users look for).

5.1 Overall Categorical Differences

Food and DIY are the most popular categories on Pinterest. As Figure 3a shows, food and drink and DIY were the two most popular categories, corroborating prior market studies (e.g., [18]). Looking deeper at the specific motivations users have in each category, users interested in these categories are also more likely to be planning to make something (17% and 14% respectively), than if they were interested in other categories. Across all categories however, looking for ideas and inspiration was still the most common motivation ($\geq$38%), and this was most pronounced among users interested in home and decor – over half (52%) were looking for ideas in this category. Boredom as a motivating factor was cited most commonly among users looking for entertainment (27%). And where surveyed male users were more interested in art and design, female users were more interested in food and drink, DIY, home decor, and fashion ($\chi^2>34.3$, $p<10^{-3}$).

Category moderates goal specificity and temporal range. Depending on the category of interest, goals may be more actionable and shorter term (e.g., finding recipes), or less actionable and longer-term (e.g., planning a vacation). Examining goal specificity, users interested in food and drink or DIY were most likely to be goal-specific (40% and 39% of users are goal-specific respectively, Figure 3b), and those interested in travel or entertainment less goal-specific (24% and 23% respectively).

For temporal range, users interested in food were most likely to have short-term goals (32%), and least likely to be unsure about taking action (30%). At the other end, users interested in home and decor or travel tended to have long-term goals (29% and 27%); users interested in entertainment were most likely to be unsure about taking action (52%). In other words, users looking for food or DIY-related content tended to be looking for something to make right away, while users interested in home and decor or travel were more likely to be looking for ideas and planning for the longer term.

Intent accentuates behavior differently in different categories. While many of our prior results hold within individual categories, intent affects specific behaviors differently in different categories.

Studying goal specificity, among users interested in DIY, those that were goal-specific made just over twice as many searches as those that were goal-nonspecific (1.0 vs. 0.4, $t(1902)=10.0$, $p<10^{-3}$, $d=0.38$). In contrast, among users interested in entertainment, goal-specific users made over three times as many searches (1.0 vs. 0.3, $t(360)=6.3$, $p<0.001$, $d=0.53$). Among users interested in food and drink, those that were goal-specific were more likely to reference past saved content than those that were goal-nonspecific (21% vs. 7% of closeups were of past content, $t>4.7$, $p<10^{-3}$, $d>0.42$), but this was not the case for users interested in fashion (11% vs. 8% for closeups, n.s.).

Differences also exist for temporal range. Among users interested in food and drink, those with long-term goals pinned almost twice as much as those with short-term goals (2.4 vs. 1.3, $t(840)=4.4$, $p<10^{-3}$, $d=0.27$). In contrast, among users interested in fashion, those with long-term goals did not pin significantly more (2.1 vs. 1.9, n.s.). Among users interested in food and drink or DIY, those with short-term goals were over twice as likely to be viewing closeups of past saved content compared to those with long-term goals (food and drink: 22% vs. 10% of closeups were of past content, DIY: 18% vs. 8%, $t>5.0$, $p<10^{-3}$, $d>0.32$). However, this was not the case for users interested in travel (9% vs. 9%, n.s.).

In summary, by examining individual categories of interest, we can discover subtle differences in the specific behaviors users engage in. In the case of the food and drink category, users may be likely to be looking up recipes saved in the past. In the case of travel or home decor, users are instead more likely to be engaging in more exploratory idea-finding and longer-term planning.

5.2 Recipes on Pinterest

One of the most common uses of Pinterest is to find recipes [18]. To study how intent influences recipe-finding behavior, we consider the subset of users who viewed a closeup of at least one recipe. We find that goal-specific users view more closeups of recipes than goal-nonspecific users (1.6 vs. 0.9 recipes, $d=0.50$, $t(762)=7.3$, $p<10^{-3}$), indicating that goal-specific users may be more interested in how to make the depicted food item. Users with short-term goals also view more recipe closeups than users with long-term goals (1.8 vs. 0.9, $d=0.30$, $t(230)=2.9$, $p<0.01$).

Intent may also affect the specific types of food that users look for. An examination of a recipe's ingredients reveals that the proportion of recipes that users view closeups of containing meat or seafood-related ingredients is highest among users with short-term goals (42% vs. 27% for users with long-term goals, $t(189)=3.2$, $p<0.01$, $d=0.34$). On the other hand, the proportion of recipes containing sugar is higher for goal-nonspecific users than for goal-specific users (39% vs. 27%, $t(295)=2.9$, $p<0.01$, $d=0.28$).

Together, these findings suggest that users with short-term goals are more likely to be looking for main courses to make, perhaps for dinner, and that users who are casually browsing are more likely to be looking for desserts to admire. Future work here may involve studying differences in recipe complexity and nutritional value. With regards to the latter, we observed a trend that users with short-term goals may view recipes with more sugar and salt than those with long-term goals. While these effects were not significant, they are suggestive of time discounting [13], or that users looking to make food in the short-term may be undervaluing the future health benefits of food with less salt or sugar.

6. PREDICTING INTENT

Thus far, we have described how goal specificity and temporal range affect user behavior. However, is it possible to use these behavioral signals to recover intent? If we can predict user intent near the beginning of their current session, we can alter the user interface and content to better serve that visitor's needs. In this section, we construct predictive models of intent, and study how performance and feature importance changes with the observation window and category, for both goal specificity and temporal range.

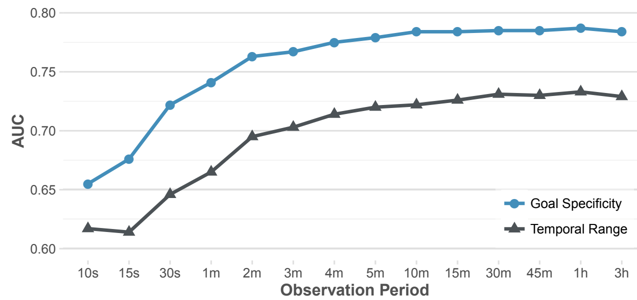

Figure 4: Observing just the first two minutes of a user's session results in robust prediction performance for goal specificity and temporal range, with prediction performance increasing the longer a session is observed.

Feature Set	Goal Specificity	Temporal Range
Demographics	0.56	0.54
+ Historical Activity	0.67 (0.66)	0.62 (0.61)
+ Current Activity	0.78 (0.77)	0.72 (0.71)

Table 1: Intent is best predicted by what a user is currently doing (i.e., current activity in the first ten minutes of a user session), but can still be predicted even before a user logs on (using demographics and historical activity). Shown are the performance improvements from incrementally adding these features, with AUC reported. Individual feature set performance is in parentheses.

Challenges in predicting intent. Several challenges exist in accurately predicting intent. First, intent is provisional and may change over time; models of planned behavior from prior work only explain up to 38% of the variance in observed behavior [42]. Further, if one seeks to predict intent in the minutes following a user logging on, there are few behavioral signals, if any at all. In the first two minutes of a user session on Pinterest, the median number of pins, closeups, content click-throughs, and searches are all zero. The impreciseness of how intent is defined coupled with data sparsity suggests that high performance is difficult to achieve.

Features. Based on our findings in our previous sections, we considered three broad classes of features:

- **Demographics**. Demographic factors such as gender, age, and location can affect intent [10]. For example, we found that women on Pinterest are more likely to be goal-specific, and other work also found gender effects on behavior on Pinterest [8]. Age can also influence users' future time perspective [21], and hence a goal's temporal range.
- **Current activity**. Behavior is directly influenced by intent, and affects how much time users spend on individual pieces of content, how much they search, and what categories of content they browse. Thus, we consider factors relating to the different actions users may take on Pinterest (e.g., searches, views, pins, closeups, and click-throughs), the time of day and day of week of the current session, as well as the categories (of which there are 33) within which these actions are taken.
- **Historical activity**. Past behavior may be indicative of future intent. For example, users who viewed less content in the past tended to be more goal-specific in the current session. Thus,

we measured user activity in the 24 hours prior to the current session, in addition to other longer-term features such as the days since a user signed up and the total number of pins a user saved over their lifetime.

Prediction tasks. With these features in mind, we considered two prediction tasks. After observing a user's behavior for a period of time, can we (a) predict whether their intent was goal-specific or not, and (b) whether they planned to take action in the short-term or the long-term? For the former prediction task, we used a balanced dataset of goal-specific or goal-nonspecific users, noting that the original dataset is already fairly balanced, and that users who reported being neither comprise a relatively small fraction of users. As exactly half of users are goal-specific, random guessing achieves classification accuracy of 50%. For the latter prediction task, we instead use a balanced dataset of users with either short-term or long-term goals. We also consider a multi-class setting of this prediction task which includes having mid-term goals or being unsure of taking action as possible outcomes. Using a random forest classifier, we performed ten-fold cross-validation, and primarily report the area under the ROC curve (AUC). All features were standardized.

Overall performance is robust. We obtain strong performance in predicting both goal specificity (AUC=0.78, F1=0.70) and temporal range (AUC=0.72, F1=0.67) after observing the first ten minutes from when a user logs in (Table 1). A logistic regression classifier gives empirically similar results. In the multi-class version of predicting temporal range, we also obtain comparatively robust performance (weighted F1=0.38, as compared to 0.14 when simply predicting the majority class).

Current activity is most predictive of intent. Given that user activity in the current session resulted directly from their stated intent, we expect that current activity alone would strongly predict intent. In fact, with current activity alone, we achieve performance almost equal to that of using all features (AUC=0.77, 0.71 respectively). In comparison, while demographic and historical activity are less predictive, they remain useful – demographics can provide a baseline estimate of what a new user is likely to do during their first visit, while historical activity can be used to estimate what a current user is likely to do the next time they visit the web site.

For both goal specificity and temporal range, search most strongly indicated intent. The mean number of words in search queries (AUC=0.66 and 0.59) and the number of search queries (0.65 and 0.58) were two of most individually predictive features, i.e., a greater number of more complex queries corresponds to a higher likelihood of being goal-specific or having short-term goals. Other measures of task focus also played a significant role, as did content category – viewed pins belonging to fewer categories (0.61) was also predictive of greater goal specificity, while viewing pins related to home and decor (0.62) was most predictive of having long-term goals.

Intent can be predicted quickly. As Figure 4 shows, while intent becomes easier to discern the longer a user's behavior is observed, performance remains relatively robust even when predictions are based on shorter durations of time. In just the first two minutes, performance for both goal specificity and temporal range are already substantial (AUC=0.76 and 0.70 respectively). Prediction remains possible even with shorter amounts of time (in 30 seconds, 0.72 and 0.65). Predictions can also be made before a user does anything by using only demographics and historical activity features (0.67 for goal specificity, 0.61 for temporal range).

Thus, not only can we guess a user's intent before they even log on, but we can quickly improve on our guess within minutes, and immediately adjust a user's experience to match their intent.

Predictibility varies by category. In the context of consumer research, prior work found that segmentation helped improve sales forecasts based on purchasing intent [31]. Similarly, if we know what category a user is interested in, we may be able to make better predictions about their intent. Here, our results are mixed. Training classifiers on individual categories, we find that intent is most predictable for food and drink (AUC=0.80 for goal specificity, 0.75 for temporal range), but least predictable for fashion (0.71 and 0.62).

7. DISCUSSION AND CONCLUSION

In this work, we presented a framework for characterizing the relationship between a person's intent and their behavior. Through a survey designed to clarify a user's intent that was followed by an observational study of subsequent behavior, we discovered significant differences in how users behaved depending on whether they were goal-specific or goal-nonspecific, or if they were planning to take action in the short-term, long-term, or take no action at all. Users differed in how focused their activity was, what content they looked at and at what level of detail, how long they spent on the site, and whether they would return soon. Intent also varied with gender and age, and by category. Finally, we used these behavioral signals to recover a user's intent.

Design implications. How may we apply these insights to the design of online platforms? First, our findings (e.g., on task focus) may be directly useful in similar content discovery and sharing services (e.g., Flickr or Netflix). Next, as goal-specific and goal-nonspecific users view content differently, recommender systems could prioritize showing specific, targeted content to goal-specific users, and more diverse content to goal-nonspecific users. Recommendations can also be tailored to specific categories – depending on whether a user looking for restaurants has short-term goals or long-term goals, a system might suggest either restaurants currently open nearby, or ones with higher ratings that accept reservations further away. Further, while goal-specific users may not return to a web site as often, they stay longer whenever they do, so providing specific goals (e.g., learning how to write a simple computer game like Pong) may encourage these users to visit more. When users have longer-term goals, sites could offer feedback or track progress towards the goal [11], for example, through providing checklists or email reminders. And as intent can be predicted quickly, content and interface changes can be made in real-time, with these predictions improving as a user continues to use the site.

Limitations and future work. Several limitations of this analysis exist. Measuring intent may influence a user's subsequent behavior [30]. Intent can also change [12] during a user session, but we partially mitigate this by primarily considering only the first ten minutes of a user session, and surveying intent right before the beginning of the session. Explicitly modeling how intent changes over time may improve predictions over longer user sessions. Surveyed users also tend to be more engaged or invested. Further, while we sought to present a broad overview of aggregate behavior on Pinterest, our results suggest that category-specific behaviors exist (e.g., in recipes) – their detailed study remains future work.

Our study in this work is limited to understanding goal specificity and temporal range on Pinterest, but we see our methods generalizing to other online settings. On social networks (e.g., Facebook), we could survey users' social support and self-presentation motivations [32], then observe and subsequently predict their posting and commenting behavior. More generally, by considering other aspects of intent such as difficulty and commitment, we may also predict if a user is likely to succeed in their goals [25].

8. REFERENCES

[1] I. Ajzen. From intentions to actions: A theory of planned behavior. In *Action Control*. 1985.

[2] I. Ajzen and T. J. Madden. Prediction of goal-directed behavior: Attitudes, intentions, and perceived behavioral control. *J Exp. Soc. Psychol.*, 1986.

[3] D. Albarracin and R. S. Wyer Jr. The cognitive impact of past behavior: influences on beliefs, attitudes, and future behavioral decisions. *J. Pers. Soc. Psychol.*, 2000.

[4] J. T. Austin and J. B. Vancouver. Goal constructs in psychology: Structure, process, and content. *Psychol. Bull.*, 1996.

[5] A. Broder. A taxonomy of web search. In *SIGIR Forum*, 2002.

[6] J. F. Bryan and E. A. Locke. Parkinson's law as a goal-setting phenomenon. *Organ. Behav. Hum. Perf.*, 1967.

[7] L. L. Carstensen. The influence of a sense of time on human development. *Science*, 2006.

[8] S. Chang, V. Kumar, E. Gilbert, and L. G. Terveen. Specialization, homophily, and gender in a social curation site: findings from Pinterest. In *Proc. CSCW*, 2014.

[9] J.-S. Chiou and C.-C. Ting. Will you spend more money and time on internet shopping when the product and situation are right? *Comput. Hum. Behav.*, 2011.

[10] A. J. Elliot and J. M. Harackiewicz. Goal setting, achievement orientation, and intrinsic motivation: A mediational analysis. *J. Pers. Soc. Psychol.*, 1994.

[11] M. Erez. Feedback: A necessary condition for the goal setting-performance relationship. *J. Appl. Psychol.*, 1977.

[12] M. Fishbein and I. Ajzen. Belief, attitude, intention, and behavior: An introduction to theory and research. 1977.

[13] S. Frederick, G. Loewenstein, and T. O'Donoghue. Time discounting and time preference: A critical review. *J. Econ. Lit.*, 2002.

[14] M. Frese and D. Zapf. Action as the core of work psychology: A German approach. *Handbook of industrial and organizational psychology*, 1994.

[15] J. F. George. The theory of planned behavior and internet purchasing. *Internet research*, 2004.

[16] J. R. Hollenbeck and H. J. Klein. Goal commitment and the goal-setting process: Problems, prospects, and proposals for future research. *J. Appl. Psychol.*, 1987.

[17] S. Hutchison and M. L. Garstka. Sources of perceived organizational support: goal setting and feedback. *J. Appl. Psychol.*, 1996.

[18] Ipsos. App overload? a cheat sheet for the best ways to stay connected and discover new things in 2016. http://bit.ly/2eCMXR0, 2016.

[19] J. L. Joines, C. W. Scherer, and D. A. Scheufele. Exploring motivations for consumer web use and their implications for e-commerce. *J. Consum. Mark.*, 2003.

[20] A. N. Joinson. Looking at, looking up or keeping up with people?: motives and use of facebook. In *Proc. CHI*, 2008.

[21] F. R. Lang and L. L. Carstensen. Time counts: future time perspective, goals, and social relationships. *Psychol. Aging*, 2002.

[22] G. P. Latham and E. A. Locke. Increasing productivity and decreasing time limits: A field replication of parkinson's law. *J. Appl. Psychol.*, 1975.

[23] R. Linder, C. Snodgrass, and A. Kerne. Everyday ideation: All of my ideas are on pinterest. In *Proc. CHI*, 2014.

[24] C. Lo, D. Frankowski, and J. Leskovec. Understanding behaviors that lead to purchasing: A case study of pinterest. In *Proc. KDD*, 2016.

[25] E. A. Locke. Toward a theory of task motivation and incentives. *Organ. Behav. Hum. Perf.*, 1968.

[26] E. A. Locke and J. F. Bryan. The directing function of goals in task performance. *Organ. Behav. Hum. Perf.*, 1969.

[27] E. A. Locke, D.-O. Chah, S. Harrison, and N. Lustgarten. Separating the effects of goal specificity from goal level. *Organ. Behav. Hum. Dec.*, 1989.

[28] G. Marchionini. Exploratory search: from finding to understanding. *CACM*, 2006.

[29] A. J. Mento, R. P. Steel, and R. J. Karren. A meta-analytic study of the effects of goal setting on task performance: 1966–1984. *Organ. Behav. Hum. Dec.*, 1987.

[30] V. G. Morwitz, E. Johnson, and D. Schmittlein. Does measuring intent change behavior? *J. Consum. Res.*, 1993.

[31] V. G. Morwitz and D. Schmittlein. Using segmentation to improve sales forecasts based on purchase intent: Which" intenders" actually buy? *J. Marketing Res.*, 1992.

[32] A. Nadkarni and S. G. Hofmann. Why do people use Facebook? *Pers. Indiv. Differ.*, 2012.

[33] T. P. Novak, D. L. Hoffman, and A. Duhachek. The influence of goal-directed and experiential activities on online flow experiences. *J. Consum. Psychol.*, 2003.

[34] J. R. Nuttin. The future time perspective in human motivation and learning. *Acta Psychol.*, 1964.

[35] W. T. Powers. *Behavior: The control of perception*. 1973.

[36] M. Richardson, E. Dominowska, and R. Ragno. Predicting clicks: estimating the click-through rate for new ads. In *Proc. WWW*, 2007.

[37] W. W. Ronan, G. P. Latham, and S. Kinne. Effects of goal setting and supervision on worker behavior in an industrial situation. *J. Appl. Psychol.*, 1973.

[38] K. Sherwin. User intent affects filter design. http://bit.ly/2evdOjs, 2016.

[39] H.-P. Shih and B.-H. Jin. Driving goal-directed and experiential online shopping. *J. Org. Comp. Elect. Com.*, 2011.

[40] B. Silbermann. 150 million people finding ideas on pinterest. http://bit.ly/2eQVc0L, 2016.

[41] S. Stumpf, X. Bao, A. Dragunov, T. G. Dietterich, J. Herlocker, K. Johnsrude, L. Li, and J. Shen. Predicting user tasks: I know what you're doing. In *Proc. AAAI Workshop on Human Comprehensible Machine Learning*, 2005.

[42] S. Sutton. Predicting and explaining intentions and behavior: How well are we doing? *J. Appl. Psychol.*, 1998.

[43] J. Teevan, S. T. Dumais, and D. J. Liebling. To personalize or not to personalize: modeling queries with variation in user intent. In *Proc. SIGIR*, 2008.

[44] M. Wolfinbarger and M. C. Gilly. Shopping online for freedom, control, and fun. *Calif. Manage. Rev.*, 2001.

[45] L. Xiang, Q. Yuan, S. Zhao, L. Chen, X. Zhang, Q. Yang, and J. Sun. Temporal recommendation on graphs via long-and short-term preference fusion. In *Proc. KDD*, 2010.

[46] Y. Zhang, W. Chen, D. Wang, and Q. Yang. User-click modeling for understanding and predicting search-behavior. In *Proc. CIKM*, 2011.

[47] B. J. Zimmerman, A. Bandura, and M. Martinez-Pons. Self-motivation for academic attainment: The role of self-efficacy beliefs and personal goal setting. *Am. Educ. Res. J.*, 1992.