

much recent work on exploiting the web as a repository of human beliefs, we show how to mine very large libraries of human activities from the web, instead of analyzing sensor data. In particular, we use step-by-step human descriptions of activities (such as recipes and how-tos) from appropriate websites to obtain the structure of our models, and supplement these with web-wide measures of similarity between concepts and terms to obtain the parameter values for our models. The technique is feasible with very lightweight natural language tools and has enabled our system to mine a library of models for roughly twenty thousand mundane activities, covering many aspects of human life.

At the heart of our technique is a breakthrough in sensing technology. Advances in miniaturization and manufacturing have yielded postage-stamp sized, forty-cent radio transceivers called Radio Frequency Identification (RFID) tags that can be attached unobtrusively to objects as small as a toothbrush. The tags are wireless and battery free. When queried by radio, the tag responds with a globally unique identifier using power scavenged from the querying signal. When combined with special wearable and ambient tag readers and a database mapping identifiers to names, RFID technology can reliably name the objects with which a person is interacting. Because of the sensors' accuracy and specificity we model activities in a novel way: we define an activity in terms of the probability and sequence of the objects that are physically involved in the activity. We are thus able to view models as probabilistic translations from terms that represent an activity name (such as "drinking tea") to terms that represent the objects involved (such as "teacup" and "teabag").

To show that the resulting models (and the recognition system as a whole) can be used for detecting actual activities, we present the results of analyzing activities of daily living (ADLs) data from eight subjects in a real home. ADLs are regarded as indicators of patient wellness, and professional caregivers are often required by law to record them. Our users performed subsets of fourteen classes of activities comprising 66 activities in all. To explore the usefulness of the remaining thousands of models, and to evaluate the components of our model extractor more thoroughly, we also present results based on simulated traces and cross-corpus comparisons.

In this paper, we make three main contributions:

1. We formulate the problem of extracting activity models as that of generating translations between natural language terms.
2. We describe a simple set of techniques to mine these translations automatically from the web.
3. We validate these techniques via a combination of real-world experiments and simulations.

To the best of our knowledge, this paper is the first to show how to mine useful models of human physical activity from the web.

The system for mining models is part of a larger activity recognition system called the Proactive Activity Toolkit (PROACT)[15]. To place the model extractor in context, section 2 sketches the structure and usage model for PROACT. Section 3 describes in detail how we mine models. Section 4 evaluates the model extractor. Section 5 presents related work. Section 6 summarizes the paper and presents future work.

2. ACTIVITY INFERENCING SYSTEM

Below, we describe how the PROACT activity recognition system, which uses the mined models, is intended to be used. We then describe the high-level structure of PROACT, place the model miner in context.

2.1 Usage model

PROACT assumes that "interesting" objects in the environment contain RFID tags. These can be purchased off the shelf, cost roughly \$0.40 each, have the form factor of postage stamps (including adhesive backing), and can withstand day-to-day use for years. PROACT deployment involves tagging tens to hundreds of objects in their environment. This can be done incrementally. As the number of tags increase more accurate and detailed recognition becomes possible. Tagging an object involves sticking an RFID tag on it, and making a database entry mapping the tag ID to a name. Current trends indicate that within a few years, many household objects may be RFID-tagged before purchase, thus eliminating the overhead of tagging.

Users employ RFID tag readers to track tag objects they interact with. They may wear tag-detecting bracelets or gloves, place medium-range readers in corners of rooms, or run robots, vacuum cleaners, or janitorial carts, with mounted long-range readers. As users go about their daily activities, the readers detect tags that (a) users touch, (b) are close to them, or (c) are moved by them, and thereby deduce which objects are currently involved in an activity. PROACT uses the sequence and timing of object involvement to deduce what activity is happening.

An application can query PROACT at any time for the likelihood of various activities being tracked or details of those activities (e.g. objects involved or durations), or subscribe for event notification when activities occur with a specified degree of certainty. Programmers name activities using plain English phrases (e.g. "paying bills"). The phrase used can either be chosen from a pre-mined list provided by PROACT, or can be a new one provided by the programmer. For a new phrase, the programmer can define the activity by providing a text document containing an English description of the steps involved in the activity (much like a recipe). The model extractor converts text into activity definitions.

2.2 System Overview

Figure 1 presents the main components of PROACT. It is centered on an inference engine which, given models for activities, and sequences of sensor readings, returns the likelihood of current activities. The models are produced by the model extractor, which extracts them automatically from text documents, including but not limited to websites. The sensor readings are produced while the end-user performs activities. For debugging, PROACT provides an activity viewer, which provides programmers with a real-time view of activities in progress, the sensor data seen, and an indication of how belief in each activity changes with the data. A more detailed description of the system is in [15]. We elaborate on some of the relevant components below.

we will generically call “directions” below) include how-tos (e.g. those at ehow.com), recipes (e.g. from epicurious.com), training manuals, experimental protocols, and facility/device use manuals. For example, Figure 4 shows directions for making tea.

| |
|--|
| <p>Making Tea:</p> <ol style="list-style-type: none"> 1. Fill a teapot from the faucet. Place kettle on the stove and boil. 2. Pour hot water into a cup, filling $\frac{3}{4}$ of the cup. Immerse teabag in cup for two minutes and dispose of teabag. 3. Add milk and sugar to taste. |
|--|

Figure 4: Directions for Making Tea

3.1 Syntactic structure of directions

A key observation in enabling this translation is that, in many cases, the structure of natural language directions closely parallels the structure of our formal model. Typically, each direction consists of:

1. A title t for the activity, and
2. A textual list r_1, \dots, r_m of steps. Each step r_i has:
 - Possibly a special keyword delimiting duration d_i , and
 - A natural-language description nld_i , typically a paragraph, of what to do during the step, typically mentioning some subset of the objects involved in the step, and often also constraints on the duration of the step.

The example of Figure 4 has title $t = \text{“Making Tea”}$, number of steps $m = 3$, no formally stated duration for each step, a two-minute minimum constraint on the duration of step 2. The objects mentioned in each step are highlighted in bold font.

In practice, each corpus of directions we wish to mine has its own concrete syntax (keywords, numbering scheme, indentation, etc), which needs to be parsed into the above abstract syntax. For each corpus, therefore, we require a user to provide a front-end that performs this parsing.

3.2 Converting directions to activity models

The similarity of structure between the formal models and the directions suggests the following scheme for converting the latter into the former. Intuitively, we produce one step in the model for each step in the directions; the objects mentioned in the paragraph for the step are those involved with the step in the model. We name key steps in bold, and provide acronyms for them when we need to use them in the future.

1. **Labeling.** Set label l of the mined model to title t of the directions.
2. **Parsing Steps.** For each step r_i in list r_1, \dots, r_m generate step s_i as follows:
 - a. If r_i has keyword-delimited duration d_i , set the mean duration μ_i for the step to d_i , and standard deviation σ_i to $S(d_i, i, l)$.
 - b. Let $O_i = O(nld_i)$ be the set of terms that represents objects in the descriptive paragraph nld_i . For each term o_{ij} in O_i , calculate the probability $P(o_{ij}, i, l)$ that an object named o_{ij} is touched in step i of activity l .

3. **Tagged Object Filtering.** Given the set $O_{deployed}$ of tagged named objects in the space where the activity is happening, we remove from our model all observations related to objects not in $O_{deployed}$. This step exploits the fact that in most actual deployments of the system, the set of objects the system can possibly see is often much smaller than the set of objects named in the models.

The scheme depends on three helper functions:

- Function $S(d, i, l)$ computes the standard deviation for the duration of step i of activity l (which has mean duration d). A very simple definition is $S(d, i, l) = N \text{ seconds}$, where N is an integer fixed for each corpus being mined.
- Function $O(nld)$ computes the set of terms representing objects in paragraph nld . This happens as follows:
 - a. **Object Extraction.** Compute the set T of terms in nld that represent objects, as per an ontology over the terms. We currently use the WordNet [18] ontology, and include all terms that have either “object” or “substance” as hypernyms.
 - b. **Noun-Phrase Extraction.** Compute the subset of terms in T that are used as nouns in the original paragraph, nld . We currently use the QTag tagger [10] on the incoming paragraph for part-of-speech tagging. We run the tagged paragraph through a customized regular-expression based noun-phrase extractor, which extracts the set of maximal phrases that contains only nouns.
- Function $P(o, i, l)$ computes the probability that object o_i is involved in step i of the activity labeled l . We consider two alternative approximations for P , neither of which makes an effort to distinguish between two different steps of a model i.e. they ignore parameter i of function P .
 - a. **Fixed Probabilities.** $P(o, i, l) = P_{obj}$. where $0 < P_{obj} < 1$ is fixed for each corpus being mined. In our experiments below we use $P_{obj} = 0.5$.
 - b. **Google Conditional Probabilities (GCP).** $P(o, i, l) = \text{GoogleCount}("l"+o)/\text{GoogleCount}(l)$. $\text{GoogleCount}(s)$ is the number of pages on the web matching string s as reported by Google [6], and $s+s'$ is the concatenation of string s and s' . For example, if the phrase “making tea” has 24,200 matches, and the phrase “making tea” cup has 7,340 matches, we conclude that the conditional probability of a cup being involved in (any step of) making tea is $7340/24,200 = 0.3$.

| object→ ↓activity | Cup | Diaper | remote control | Stapler | keyboard | wrench |
|------------------------|------|--------|-------------------|---------|----------|--------|
| making tea | 0.30 | 0.01 | 0.01 | 0.00 | 0.03 | 0.01 |
| changing baby | 0.09 | 0.14 | 0.08 | 0.00 | 0.07 | 0.00 |
| watching television | 0.07 | 0.01 | 0.06 | 0.00 | 0.02 | 0.00 |
| copying paper | 0.03 | 0.01 | 0.01 | 0.01 | 0.05 | 0.00 |
| sending e- mail | 0.01 | 0.00 | 0.01 | 0.00 | 0.04 | 0.00 |
| auto repair | 0.09 | 0.06 | 0.02 | 0.00 | 0.04 | 0.05 |

Table 1: GCP for some common activities and objects

The GCP is predicated on what we call the *mirror assumption*: the probability of an object being involved in an activity in the real world is reflected by the probability of the phrases describing the two co-occurring in human discourse (and therefore on the web). In practice, since we interpret the results of our inference engines as likelihoods rather than probabilities, it is sufficient for our purposes that the probabilities are consistent; if the actual probability of involvement of one object-activity pair is lower than that of the other, the GCP of the one should be lower than that of the other.

Table 1 presents the GCPs for six activities and six objects. We pick the objects such that the i 'th object is intuitively one that would be involved in activity i , but not in any other. The i 'th row and i 'th column of the table contains the GCP of the i 'th object being involved in the j 'th activity.

In most cases, as expected, the table has its highest values along the diagonal. Pairings that intuitively seem unlikely (such as “diaper” and “making tea”) have substantially lower value than more plausible ones (e.g. “cup” and “making tea”). Although the relative values of the probabilities are on the whole sensible, the absolute values are lower than one might expect in most cases. For instance one expects to use a keyboard more than 4% of the time when sending e-mail! It is therefore important that any scheme that uses GCP require at most consistency, but not absolute accuracy, of probabilities.

3.3 Example

To show how the model extractor works, we apply the model-mining steps of the previous subsection to the directions for making tea:

1. **[Labeling]** As per step 1, the label for the new activity is “Making Tea”.
2. **[Parsing Steps]** As per step 2, the new activity has $m = 3$ steps.
 - a. Since none of the steps in the figure have a keyword-delineated duration, we do not ascribe durations to any of the steps.
 - b. Figure 5 shows the result of applying the helper functions to each of the paragraphs 1, 2 and 3 of the incoming directions.
 - i. Sets T_1 through T_3 show terms describing objects, as extracted by the OT pass for each of the three paragraphs. Note that the terms include “filling”, which though a hyponym of “substance” in WordNet, is as a verb here.
 - ii. Sets O_1 through O_3 are the subsets of T_1 through T_3 , as produced by the POST pass that are used as nouns in the directions. Since “filling” is used as a verb, it is eliminated here.
 - iii. Sets s_1 through s_3 are the result of adding to each object the probability that it is involved in the activity as a whole, using GCP to get probabilities.

Note that this resulting model has structure very similar to that of the model in Figure 3, except that the probabilities are different and that the new model has no timing constraints.

3. **[Tagged Object Filtering]** Finally, suppose that we use the model to track activities in a space where the set O_{deployed} of

names used for tagged objects is {kettle, stove, cup, teabag, milk}. Sets t_1 through t_3 result from removing from sets s_1 through s_3 objects that are not in this set.

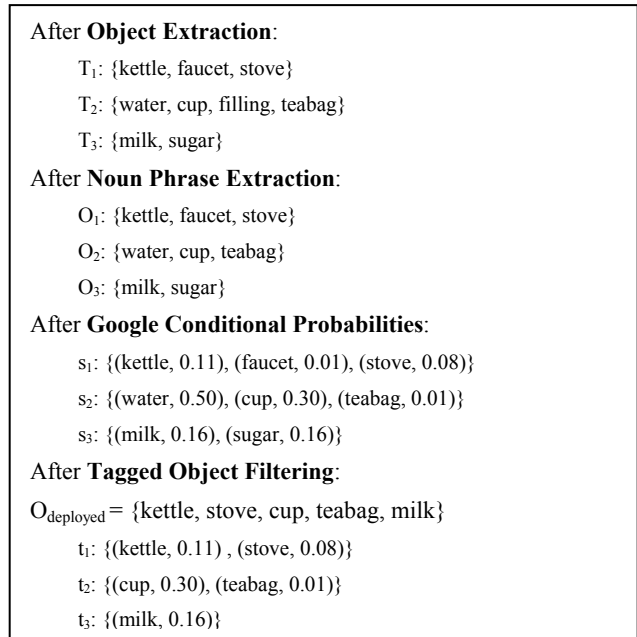


Figure 5: Steps in Mining the Directions for Making Tea

4. EVALUATION

We have used the techniques described in the previous section to mine roughly 21,300 models from two websites that provide directions for activities. We mined ehow.com for roughly 2300 directions on performing domestic tasks (from “boiling water in the microwave” to “change your air filter”), and fpts.com and epicurious.com for a further 400 and 18,600 recipes respectively. In the rest of this section, we evaluate these models. None of the models we mined for this set of experiments have keyword-delimited timing constraints; we have used a default of five seconds per activity step.

We are interested in answering two main questions through our evaluation:

1. **Are the models good enough?** We wish to evaluate whether the models we mined contain sufficient information to enable correct recognition of object traces gathered from actual humans performing activities. The subjects may, of course, perform the activities in a wide variety of ways, most of which will hopefully be captured by the model.
2. **How effective are the various steps that comprise the model extractor?** We wish to evaluate the impact of the various steps (and of alternate design options) of the mining scheme.

We assume for the evaluation that the inference engine is fixed to be the Monte-Carlo based solver sketched in section 2.2.3.

A comprehensive but impractical strategy to determine sufficiency and necessity would be to collect, from real-world activities, traces that exercise all models we have mined. These traces would represent the large variety of ways in which people perform activities. By comparing to ground truth the results reported by the inference engine on these traces, we could

evaluate sufficiency of all the models. By comparing accuracy of results with various (combinations of) mining techniques turned off, we could evaluate the efficacy of those techniques.

Unfortunately, it is impractical to collect actual usage data for thousands of activities. The activities are time consuming, often complex and sometimes unpleasant; getting subjects to perform all of them in various combinations is clearly an impractical task. Instead, we use the following three strategies to approximate comprehensive evaluation. The three options apply to increasingly larger classes of activities, but provide increasingly indirect evidence about the quality of our system:

1. **Human activity-trace recognition (HTR).** We instrumented the home of one of the researchers with 108 RFID tags. Over a period of six weeks, we collected traces (consisting of time-stamped RFID tag numbers) from 14 subjects, while each performed a randomly chosen 12 of 14 Activities of Daily Living (ADLs). Our subjects all wore the glove-based RFID reader (Figure 2) to allow tracking of touched objects. We picked our ADLs from state-mandated ADL lists. Each ADL (e.g. housework) corresponds to many ordinary activities (e.g. washing dishes, making beds, dusting, vacuuming). In all, our subjects picked from roughly 50 activities of the latter kind. Subjects recorded activities as they performed them, to provide us with ground truth.

To measure the quality of our models, we select 66 of the models mined from ehow.com. These correspond to the fifty activities being performed by the user (some user activities are described by more than one model). We measure the accuracy with which our inference engine infers activities given these models. We define accuracy below.

2. **Inter-corpus consistency (ICC).** In the absence of actual activity traces, we cannot measure the ability of our models to represent real activities directly. However, we can still indirectly measure the ability of our models to represent the diverse ways in which activities may be performed. Specifically, if an activity model mined from a particular corpus is to represent most instances of the corresponding activity, it must, in particular be compatible with another representation of the same activity (since the latter description presumably represents some person’s way of performing the activity).

Based on the above insight, we use two different corpora (epicurious.com and ffts.com) containing the same activities to test the quality of our models. Each corpus contains roughly 120 recipes for making cookies. We first generate models for all activities in the two corpora. Next, for each model derived from one corpus, we generate traces compatible with model: for each state in the model, in order, we include in the trace a subset of the objects corresponding to that state. Each object is picked with probability equal to its conditional probability p_{ij} . We measure the accuracy with which we infer activities on these traces given the models generated from the second corpus.

3. **Intra-corpus distinguishability (ICD).** Unfortunately, even inter-corpus comparisons are difficult to perform extensively due to lack of data. Most activities in ehow.com, for instance, are not easily available from other sources. We therefore adopt a third strategy. Once again, we resort to an indirect technique for measuring if the mined models are sufficient for recognizing traces from real activities. We

observe that if the models are good enough to recognize their corresponding activities, then in particular each model should have a stronger match than any of the other models on activities that conform perfectly to it. In other words, the models must contain enough information to distinguish themselves from each other.

To keep the number of activities inferred simultaneously manageable, we restrict ourselves to testing distinguishability between models from the same activity domain. To this end, we select seven domains (ADLs, automobiles, food, grooming, parties and the two cookie datasets described above). As a measure of distinguishability, we measure the accuracy with which traces generated from models in a given domain are detected, given the set of all models in the domain. Since the traces generated are “perfect” with respect to the model being matched against, the distinguishability metric is in a sense a “limit” on how well the models can do given our inference engine.

The three strategies above rely on measuring the accuracy with which activity traces can be detected. To determine accuracy over a trace, we split the trace into five-second windows. Accuracy then refers to the fraction of windows where our inference engine gives the same result as ground truth.

In the following subsections, we present the limit measurements from our ICD study, followed by the more direct quality measures from the HTR and ICC measurements. Next, we analyze the contribution of the component techniques of the model extractor towards accuracy. Finally, we measure the effect of some of these techniques in producing more compact models.

4.1 Intra-Corpus Distinguishability

Figure 6 shows the distinguishability for each of the seven domains mentioned above. The distinguishability of a domain is the accuracy across the result of concatenating all traces for that domain i.e. the fraction of all 5-second windows across all activities in the domain that we labeled correctly. The numbers range from 67% for autos to 95% for the second cookie domain. The implication is that if the sensor trace from an activity conforms closely to that expected by the model for that activity, the models we mined for the domains are distinct enough, and our activity inference engine sensitive enough, that we can correctly identify the activity much of the time.

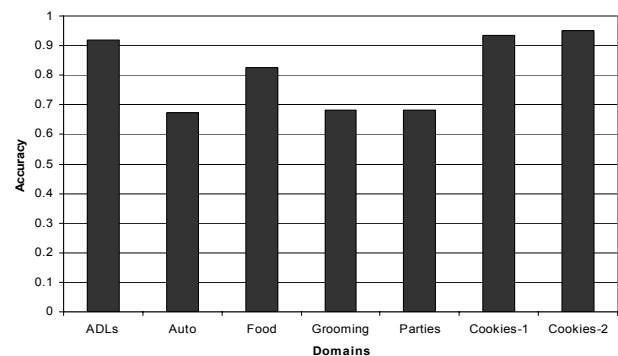


Figure 6: Distinguishability Within Activity Domains

4.2 Human and inter-corpus trace recognition

Figure 7 presents the accuracy achieved on the HTR measurements (testing against measured human traces, labeled ADLs), and the ICC measurements (testing inter-corpus consistency, labeled Cookie-1 and Cookie-2).

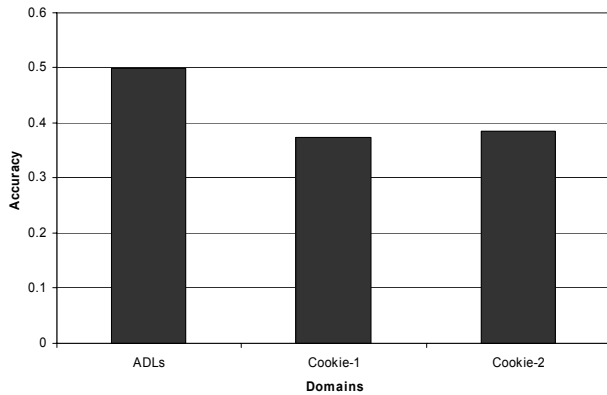


Figure 7: Accuracy of Recognition for Three Datasets

The bar labeled “ADLs” represents the accuracy across the unified trace that results from concatenating the traces from all 14 subjects across all 12 activities that they performed. It therefore says that for 50% of all 5-second timeslots, we are able to correctly infer the ongoing activity. Although the number may seem modest, we point out for comparison that a random assignment of activities to slots, given that there are 66 possible activities being modeled, will yield only 1.5% accuracy. We also note that we are unaware of any previous work that recognizes such a large variety of activities being performed by such a large number of subjects in an unmodified home (even with handcrafted activity models) with at least this level of accuracy.

The fact that recognition rates for ADLs on real traces is significantly lower than the distinguishability implies that the traces seen in practice deviate from models. There are two main reasons for this deviation. First, the traces violate some of the assumptions underlying our models. For instance, many objects touched and were not tagged, tagged objects were sometimes missed, and objects from different activities were interleaved. Second, the models were not fully general representations of the corresponding activities. For instance, they imposed extraneous ordering constraints, omitted possible objects or included optional objects.

The bar labeled Cookie-1 represents the accuracy of Cookie-1 models across the unified trace that results from concatenating the individual traces generated from the models of Cookie-2, and vice versa for the bar labeled Cookie-2. As explained previously, these bars address the question of how generally valid the models mined from each of these corpora are. In comparison with the 95% distinguishability rates, the accuracy is roughly 40% in both cases. There are two reasons for this discrepancy. One, of course, is that the identical recipe can have quite different structure in the two corpora. The other is that for some of the recipes, there is no (unique) counterpart in the other corpus. Our results may have improved had we removed these.

4.3 Impact of techniques on accuracy

We wish to measure the relative impact of our various techniques on recognition accuracy. Accordingly, we use the full system,

with all techniques in use, as a baseline and remove parts of the system one at a time. Figure 8(a) and (b) present the efficacy of the various techniques that comprise our model extractor on the ADLs and the two cookie datasets respectively. In each chart, the bar labeled “All” uses all available techniques. The other bars represent our models generated without Google probabilities (using a fixed probability of 0.5 for all objects), without performing tagged object filtering, and without performing part-of-speech tagging and noun phrase extraction, respectively. The error bars show a 95% confidence interval around the results.

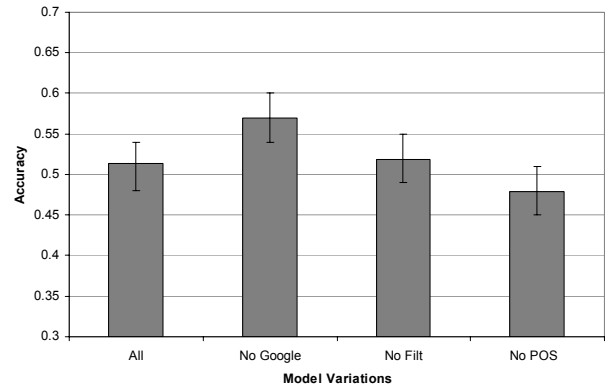


Figure 8(a): Impact of Selectively Disabling Mining Techniques with ADL Dataset

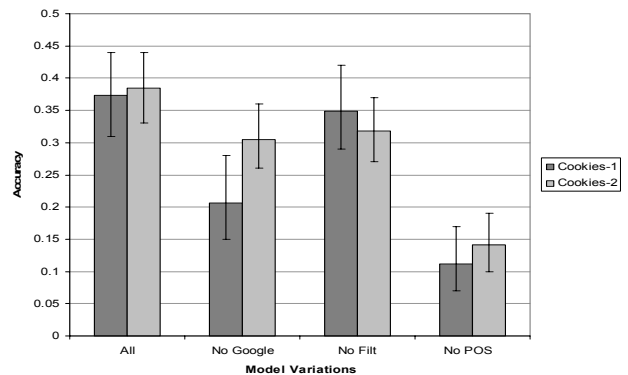


Figure 8(b): Impact of Selectively Disabling Mining Techniques with Cookies Dataset

The markedly different shapes of the two charts make clear that the domains of ADLs and Cookies are quite different. In the ADL domain, using Google probabilities actually seems to hurt accuracy. We therefore made an additional comparison between using none of our techniques and using only Google probabilities. This comparison is shown in Figure 9, with the “All techniques” accuracy for comparison. Clearly, Google probabilities alone do have a positive effect. We believe that these various techniques interact, so that doing them all together does not necessarily help. The Google probabilities provide some information, but, as we saw in **Error! Reference source not found.**, the probabilities are noisy; once we have done part-of-speech tagging and object filtering, Google probabilities provide no additional information, but the inherent noise hurts overall accuracy. This provides strong motivation for future work in improving these probabilities.

In comparison to the ADL domain, the Cookie domain involves many more activities and many more objects, and our best

performance is significantly worse. In this domain, each technique contributes some additional information, without as much apparent interaction. As a result, we see the highest accuracy when using all techniques. These differences between the ADL and Cookie domains suggest that choosing the best technique for a particular problem domain may depend highly on the structure of that domain.

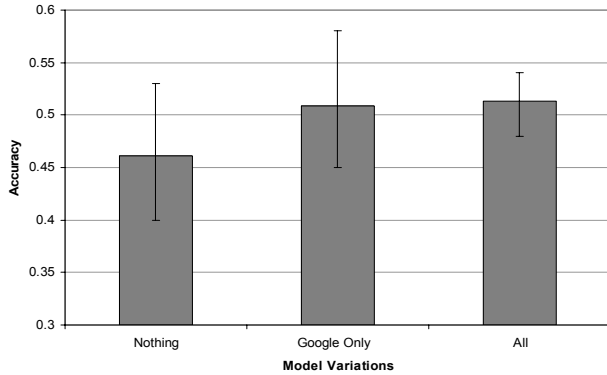


Figure 9: Impact of Google probabilities with ADL Dataset

4.4 Impact of techniques on compactness

The various techniques we apply in the extractor have significant impact on the resulting models. Figure 10 shows how the average number of objects per activity varies as we apply part-of-speech tagging and tagged object filtering to models in the ADL and Cookie domains. As discussed in Section 4.3, these two domains differ significantly in their structure. ADL activity models are only somewhat reduced when we filter for noun phrases, but drastically reduced by tagged object filtering. The resulting models are significantly smaller. In fact, approximately 30% of activity steps become empty; when we remove these empty nodes, we change the structure of the resulting activities. On the other hand, the Cookie activities are barely affected by tagged object filtering, but are significantly reduced by part-of-speech tagging. These models contain many more objects, and these techniques do not create any significant change in structure.

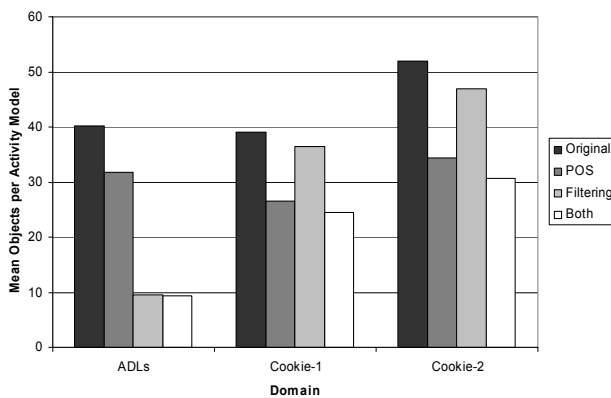


Figure 10: Impact of Mining Techniques on Model Size

5. RELATED WORK

To the best of our knowledge our work is the first to mine models of human physical activities directly from textual or web-based

descriptions. Related work falls into two main categories, and we view our work as being complementary to both.

In the past few years, much work has gone into extracting models of human beliefs (other than activities) such as shared interests [17], preferences [3], reputations [5], concept definitions [9] and spam [16] by mining virtual communities such as the web, chat rooms and newsgroups. In many cases, these efforts have revolved around a mixture of techniques similar to those we use: lightweight natural language processing (NLP), document co-occurrence analysis for detecting semantic proximity, mining topic and community-specific corpora. Even the notion of Google Conditional Probability (GCP) has in essence been used for extracting conditional probability of beliefs [7], although our application of Google to real-world phenomena (via the mirror assumption) is novel. These projects also employ a variety of techniques (such as link-structure analysis and more advanced NLP techniques) that could profitably be used in our system. In some cases, the mined models are even used by virtual sensors to recognize new data; for instance, spam classification software uses mined models to classify incoming data. None of this work, however, makes the connection between mining models from data and recognizing activities in the physical world.

An extensive literature exists on activity recognition, most of it using computer vision. In most cases [1][11][12], the focus is on identifying features, model structure and mining algorithms that enable robust detection of small subsets of activities. The models used in these systems are constructed and labeled by hand. In some of these cases, the parameters of the model are learned on the fly, whereas the structure and labels are provided by clients of the system. Similarly, some systems [12] learn patterns of user behavior, essentially modeling their likely future actions, but do not recognize or model activities per se.

A smaller set [2][8] of systems learn quantitative models of classes of activities; they observe video footage depicting the movement of pre-defined robust features and extract probability density functions over possible sequences of feature configurations as “interesting” models. The models so extracted are sufficient for their intended purpose (e.g. enabling computers to mimic facial movements and gestures, and detecting periodic phenomena in footage of people’s lives). However, in order to use them as part of a library of activities, the extracted models still have to be labeled. Viewing the quantitative models and labeling them is a challenging task for humans.

A final small class of systems [4] extracts models automatically from video, but provides qualitative descriptions of the footage. For instance, Fernyhough et al. are able to detect one car overtaking another in video footage and report the phenomena as a sequence of three operators: traveling *behind-right* in the *same* direction, traveling *right* in the *same* direction, and traveling *infront-right* in the *same* direction, rather than as a probabilistic graphical model over low-level features. The resulting description models are, of course, much easier for humans to parse and then label. However, it is conceivable that more models with more steps will get more difficult to comprehend or specify even in this qualitative fashion. Further, the system designer has to build in the primitives (such as *infront-right*) and the grammar to combine them.

We view our techniques as being complementary to the ones discussed above. In particular, our models are suited to modeling activities that are characterized by the identity of objects involved. In many cases, however, activities are better

characterized by geometry, color and texture information; the above techniques are more suited for such activities.

6. SUMMARY AND FUTURE WORK

We show how the use of novel sensing technology (RFID tags) allows us to view human activity models as probabilistic translations from natural language phrases to words describing objects. Given this view of models, we provide a suite of techniques that allow conventional natural language descriptions of activities (such as how-tos and recipes) to be completely automatically converted into formal models of activities. We show, through a combination of measured data of human activities and simulations, that the models we generate are useful. We also provide an analysis of the contribution of our mining techniques to improving the accuracy and compactness of our models. Our techniques have allowed us to construct many orders of magnitude more models for human activities than described previously.

This paper is essentially an introduction to the idea of mining activity detection from the web. We have identified an extensive research agenda that follows up on these ideas.

Our first order of business for the future is to perform a more comprehensive evaluation of our models. We intend to collect human traces on many more activities, with many more tagged objects in the activity space (a few thousand as opposed to the current hundred or so), and on more challenging, but common, patterns of activity such as multi-person and interleaved activities. We would like to find corpora that provide alternate descriptions for many of our activities, both for performing inter-corpus consistency comparisons, and more importantly to examine techniques for combining multiple descriptions for a given activity. Finally, we have not yet attempted to mine timing constraints on activity steps on any large scale; we hope to locate corpora where timing constraints can be extracted easily and to test the usefulness of the timed models that result.

We also plan to explore techniques for improving the effectiveness of our mined models. We are examining how to mine observables relevant to sensors other than RFID tags. In particular, we intend to include location in our models, extracting location information from textual descriptions. For example, the activity representing “baking cookies” might be annotated with the location “kitchen”. The location observables in the model are then triggered by location sensors as opposed to RFID sensors.

Close examination of how we fail in recognizing certain activities has led us to believe that it will be profitable to model activities as partial orders, as opposed to lists. We are considering simple natural-language processing techniques to determine dependencies between two steps of a list of directions. In addition, human activities are often described in terms of each other, as when the recipe “basic pie crust” is given as a component of the recipe “apple pie”. We intend to augment our model of activities to allow references to other activities, and to then mine these references based on syntactic cues. We expect that, in many cases, clients of our system may be interested in activities for which no step-by-step description is known (e.g. “paying bills”). In these cases, we suspect that we can still create simple naïve-Bayesian models by simply identifying object names that are “semantically close” to the phrase describing activities. Our early forays in this direction have been quite promising.

A difficulty in connecting mined activities with tagged objects is that the activity models may refer to objects synonymously. For example, we might have an object in our kitchen tagged “skillet” while our pancake recipe calls for a “frying pan”. One possible solution to the problem is to use synsets to represent objects. Synsets – collections of synonymous words – can be extracted from WordNet. Another solution would be to classify objects with an object ontology – again, an ontology can be constructed from WordNet – and recognize objects that are close in the ontology can sometimes be used interchangeably in activities.

In this paper, we assumed that the spaces in which activities are performed have been extensively tagged with RFID tags. In practice, however, the person tagging the space may have only a limited number of tags available, and may value feedback on which objects are most profitable to tag, given the set of activities to be detected. We call this the “eigenobject” problem, and have preliminary ideas on solving it.

7. ACKNOWLEDGMENTS

We thank Suzi Soroczak and Meliha Yetisgen-Yildiz for performing the noun phrase extraction. We thank Henry Kautz and Tanzeem Choudhury for feedback and discussions.

8. REFERENCES

- [1] Bobick, A. F. and Ivanov, Y.A. Action Recognition Using Probabilistic Parsing. CVPR 1998: 196-202.
- [2] Clarkson, B. Life Patterns: Structure from Wearable Sensors. PhD Thesis, MIT, 2002.
- [3] Dave, K., Lawrence, S., and Pennock, D. M. Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Product Reviews. WWW 2003: 519-528
- [4] Fernyhough, J., Cohn, A. G. and Hogg, D. Constructing Qualitative Event Models Automatically from Video Input. Image and Vision Computing, 18, 2000, pages 81-103.
- [5] Gibson, D. Communities and Reputation on the Web. PhD Thesis, UC Berkeley, 2002.
- [6] Google. <http://www.google.com>
- [7] Heylighen, F. Mining Associative Meanings from the Web: From Word Disambiguation to the Global Brain. Proceedings of Trends in Special Language and Language Technology, R. Temmerman (ed.), Standaard Publishers, Brussels, 2001.
- [8] Jebara, T. and Pentland, A. Action Reaction Learning: Automatic Visual Analysis and Synthesis of Interactive Behaviour. ICVS 1999: 273-292
- [9] Liu, B., Chin, C. W., and Ng, H.T. Mining Topic-specific Concepts and Definitions on the Web. WWW 2003: 251-260.
- [10] Mason, O. QTag home page. <http://web.bham.ac.uk/o.mason/software/tagger/index.html>
- [11] Moore, D.J., Essa, I. A., and Hayes, M. H. III. Exploiting Human Actions and Object Context for Recognition Tasks. In IEEE International Conference on Computer Vision, volume 1, pages 80-86, Corfu, Greece, September 1999.
- [12] Mozer, M. C. (1998). The Neural Network House: An Environment that Adapts to its Inhabitants. In M. Coen

(Ed.), *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligent Environments* (pp. 110-114). Menlo, Park, CA: AAAI Press.

- [13] Oliver, N., Horvitz, E., and Garg, A. Layered Representations for Human Activity Recognition. In Fourth IEEE Int. Conf. on Multimodal Interfaces, pages 3--8, 2002.
- [14] Patterson, D., Liao, L., Fox, D. and Kautz, H. Inferring High-Level Behavior from Low-Level Sensors. UbiComp 2003.
- [15] Philipose, M., Fishkin, K., Perkowitz, M., Patterson, D., and Haehnel, D. The Probabilistic Activity Toolkit: Towards Enabling Activity Aware Computer Interfaces. Intel Research Seattle Technical Memo IRS-TR-03-013, December 2003.
- [16] Sahami, M., Dumais, S., Heckerman, D. and Horvitz, E. A Bayesian Approach to Filtering Junk E-mail. In AAAI-98 Workshop on Learning for Text Categorization, 1998.
- [17] Schwartz, M. F. and Wood, D. C. M. Discovering Shared Interests using Graph Analysis. *Communications of the ACM*, 36(8):78--89, 1993.
- [18] WordNet. www.cogsci.princeton.edu/~wn