

Using Topic Models to Mine Everyday Object Usage Routines Through Connected IoT Sensors

Yanxia Zhang

FX Palo Alto Laboratory
Palo Alto, USA
yzhang@fxpal.com

Hayley Hung

Delft University of Technology
Delft, Netherlands
h.hung@tudelft.nl

ABSTRACT

With the tremendous progress in sensing and IoT infrastructure, it is foreseeable that IoT systems will soon be available for commercial markets, such as in people's homes. In this paper, we present a deployment study using sensors attached to household objects to capture the resourcefulness of three individuals. The concept of resourcefulness highlights the ability of humans to repurpose objects spontaneously for a different use case than was initially intended. It is a crucial element for human health and wellbeing, which is of great interest for various aspects of HCI and design research. Traditionally, resourcefulness is captured through ethnographic practice. Ethnography can only provide sparse and often short duration observations of human experience, often relying on participants being aware of and remembering behaviours or thoughts they need to report on. Our hypothesis is that resourcefulness can also be captured through continuously monitoring objects being used in everyday life. We developed a system that can record object movement continuously and deployed them in homes of three elderly people for over two weeks. We explored the use of probabilistic topic models to analyze the collected data and identify common patterns.

ACM Classification Keywords

1.5.2. PATTERN RECOGNITION: Design MethodologyPattern analysis

Author Keywords

Internet of Things (IoT); Deployment; Machine Learning; Sensing

INTRODUCTION

Internet of Things (IoT) systems are emerging for commercial markets, which serve as an opportunity for resourcefulness research. This paper addresses the research question of whether resourcefulness can be captured through sensor data attached to objects being used in everyday life. Using pervasive sensors and computational tools provides a novel approach for

us to understand how objects are used in everyday life. Traditional methods are based on human ethnography interviews, surveys, self-report and logging. These methods can suffer from subjective bias. In addition, they can only sample limited information from experiences recalled from the users' perspective. We propose that it would be desirable to uncover patterns of behavior that would not be easy to observe through conventional ethnographic methods. By equipping continuous passive sensing to household objects, we hope to ultimately be able to discover routines such as daily or nightly, weekday or weekend patterns. This information is valuable in understanding the usage patterns of objects across practices which can be useful for suggesting new design practices.

To address our research questions, we propose to use machine learning and leverage sensor data to capture resourcefulness, in parallel with ethnography. Specifically, our project aims to understand resourcefulness of third-age adults (people who have generally just retired, and are still fit and able). Previous works have identified that resourcefulness emerges in daily practices [3]. Hence, we focus on identifying situations where objects have been used in a creative and unexpected situation. For example, it can be when objects have been used together or at unexpected times during the day. To record events of object usage with sensors, we use their movement as a proxy of objects being handled by the participants.

Our first step is to examine raw events recorded by the sensor. These are a set of time series data that correspond to events triggered by object movement in chronological order, coming from each object and individual sensing modality. At this step, we are interested in whether the detected events from the sensors are consistent with observations during ethnographic fieldwork.

We further explore interactions between objects to uncover patterns of objects used together and their temporal context. To achieve this, we created feature representations to encode object movements for the entire group of objects per household. We extract common usage patterns by applying a topic model. Since topic models are probabilistic, they also generate quantitative metrics to rank the saliency of these patterns. We are interested in whether machine learning can generate meaningful patterns that are related to resourcefulness without human annotation.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

IOT '18 October 15–18, 2018, Santa Barbara, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6564-2/18/10...\$15.00

DOI: [10.1145/3277593.3277634](https://doi.org/10.1145/3277593.3277634)

Table 1. Summary of Participants and objects to track in each household

ID	Age	Housing	Objects to track
1	69	Terraced house	Remote Control, Spider Stick, Garden Door, Fridge, Breakfast Chair, Tray
2	76	Apartment	Chair Pillow, Remote Control, Rope on Stairs, Kitchen Drawer, Fridge
3	74	Terraced house	Kitchen Chair, Fridge, Remote Control, Kitchen Cabinet Door, Knitting Needle, Tablet

DATA COLLECTION

A set of six TI SensorTag (CC2650STK) ¹ were deployed at the lower most hierarchy, which collects the object usage data from the Inertial Measurement Unit (IMU) and environment sensors. The data is streamed to a Raspberry Pi, which acts as a data collection gateway at the site and regularly records the data on a remote server (IBM Cloud). The software that was loaded on the Raspberry Pi is located at the resourceful-gateway repository ². The software on the sensors is found at the resourceful-sensortag repository ³.

We initially conducted a deployment study in the homes of five elderly people for over two weeks. The initial deployment for the first two households were not successful due to technical issues such as sensor disconnection and running out of battery. Subsequently, we refined the system development and managed to collect data continuously for over two weeks in the remaining three households. The characteristics of the household data and tracked objects are summarized in Table 1. The environment data (air pressure, humidity, temperature and light) is sampled every minute. The movement data from the IMU is logged using a sleep and wake up strategy, logging continuously for 20 seconds whenever the amplitude of movements is above a certain threshold. The threshold was determined through a pilot test with two researchers explicitly manipulating the sensor tags to ensure that noisy jitters are not recorded.

COMPUTATIONAL MODEL

We first extracted the movement data and then applied Latent Dirichlet Allocation (LDA) [1] to discover hidden patterns. Our first step is to convert raw movement data of objects to word tokens and build a vocabulary. We consider the data from one day to be a document and each movement instance as a word. Then, we form the document collections as features represented with word distributions over the built vocabulary. Each corpus consists of data from one household during the field deployment. Finally, we use LDA to identify topics in the documents, in other words, daily usage patterns.

LDA for identifying object interaction routine

Topic models (i.e., LDA [1]) have been successfully applied in text mining community to extract summaries of large and unstructured collection of documents. LDA treats words in documents as a generative probabilistic process, in which hidden variables describe the structure of how topics are composed of mixture of words and how documents are composed of mixture of topics. Figure 1 illustrates the generative process

of the LDA model. Assuming we have a corpus D consisting of M documents and N words per document, LDA learns the topic proportions of documents by representing each topic as a mixture of word distribution and each document as a mixture of topic-document distribution. The probability of each topic θ is initialized with a Dirichlet prior with hyperparameter α . The word probability over each topic is denoted as $p(w|z_n, \beta)$ with parameter β .




Figure 1. Graphical model representation of LDA model

The idea of using topic models to discover object usage routines is that we consider the combination of object movements are a mixture of events related to activities from daily practices. If we consider a breakfast routine, this would involve the handling of different objects involving the kitchen drawer, fridge and mug in the morning. On the other hand, a leisure routine could be composed of a mixture of using remote control, fridge and sofa chair at late evening. If we treat each of these routines as a topic that consists of commonly co-occurring object usage instances, we can consider each day to be a mixture of social interaction routines in the same way that a document is formed by a mixture of topics. Typically for data clustering, popular methods include the K-means clustering which hard assigns a data sample to one particular cluster, or Gaussian mixture models (GMM) which assigns a soft membership score for each cluster assuming a Gaussian density [2]. The benefit of LDA is that it provides a richer representation for topics that are composed of probabilistic description of words, in our case, the object movement events.

Feature extraction and word representation

We used a similar approach that have successfully applied LDA to mine human interaction patterns from infrared data [4]. The code of our data analysis is available at the resourceful-topicmodels repository ⁴. Figure 2 illustrates the workflow of feature extraction and word representation from the movement event data. Our first step is to convert raw infrared log data to word tokens and build a vocabulary. We consider the recorded data from one day to be a document and each object movement event as a word. Then, we form the document collections as features represented with word distributions over the built vocabulary. Finally, we use LDA to identify topics in the documents.

¹http://www.ti.com/ww/en/wireless_connectivity/sensortag/

²<https://github.com/resourceful-ageing/resourceful-gateway>

³<https://github.com/resourceful-ageing/resourceful-sensortag>

⁴<https://github.com/resourceful-ageing/resourceful-topicmodels>

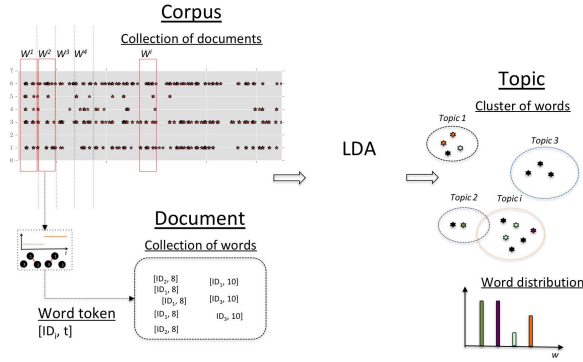


Figure 2. An illustration of the workflow

We constructed a corpus consisting of documents $D = \{w^1, \dots, w^M\}$ to represent the interaction data during M days, where one day w is a document in the corpus. Each document w consists of N words $w = \{w_1, \dots, w_N\}$ that represent interaction events within that day. The number of words N can vary daily depending on the number of times an object is moved during that day.

Word tokens

As the IMU sensors were sampled at 20Hz, this generated time series of discrete observations when movement occurred at every 0.05s. If the object moves continuously, data related to movement amplitude, the ID of the object and timestamp will be continuously recorded at the predefined resolution. We encoded the word tokens with both the timestamp and ID information of the object only when movement events were detected. We used a non-overlapping windows of $\Delta_t = 1$ hour to extract movement instances and construct words w . The word w for object ID_i is represented as:

$$w = [ID_i, t] \quad (1)$$

where ID_i is the current object and t is the temporal context that describes the hour of the day in a 24 hour format. In this way, we implicitly encode the movement duration as more words will be generated when more movement events observed within a period of Δ_t .

Choice of Number of Topics

The choice of the optimal number of topics is usually selected via a cross validation procedure (see [4] for more details). The idea is to choose a cost metric (i.e., perplexity) and select the number that best fits the training data and also generalizes well to unseen data. This is an important step because, as with any unsupervised method, metrics are required to validate the quality of the clustering in some way. Unfortunately, due to the limited number of days available in our dataset there are not enough data samples to learn optimal model parameters. Therefore, we opted for a heuristic approach. The 'rule of thumb' is that the number of topics should not exceed the number of documents (number of days). Since we only have

two weeks of data, we choose to set the number of topics as ten in this work.

RESULT

To illustrate the identified patterns, we plot the extracted 10 topics from the data of household 3 in Figure 3 where each color corresponds to the specific object. Each topic is a distribution over words and corresponds to an object usage routine, in other words, a mixture of objects being used together. Each row is a collection of interaction event words indicated as circles. Larger circle indicates a higher probability of those interaction events occurring in that topic. For clarity of display, Figure 3 shows only the words with probability greater than 0.01. The ordering of the circles within each time window is arbitrary. Therefore, for a given hour of the day, multiple circles appearing in that window can only be considered as having been moved during the same hour. We refer to this in the remainder of the text as co-use. We can see that topic 0 and topic 8 are identical. Duplicated topics appearing suggests that the number of routines is less than the parameter 10 we set in this data.

Object co-use in different routines

If we examine each individual topic, each of them appears to associate with different routine behaviours. For instance, topic 0 consists of a repeated pattern of two combinations of object usages including 'kitchen Cabinet Door' with 'Fridge' and 'Remote Control' with 'Fridge'. This topic is associated with routines around 'Fridge' which indicates its significance in this household's daily activities. The combination of 'Remote Control' with 'Fridge' usage at around 9:00-10:00 could be a daily morning ritual such as the participant watching TV while eating breakfast. The same combination appears at around 19:00-20:00 indicating the same routine of watching while eating also appearing in the evening. On the other hand, the use of the 'kitchen Cabinet Door' with 'Fridge' occurs at around 16:00-17:00 which is most likely cooking behaviour, which is less surprising.

Salient combination of object usage

If we examine words occurring at the same 1 hour period, we can observe a variety of configurations of object co-use combinations. It is very common that one or two objects are used. These topics could also be related to resourcefulness (though we can only verify this with the participants themselves). For instance, topic 1 is mostly related to usage of digital device and shows a less unique routine that the tablet are used throughout the day. We also observe a few cases where three or more are co-used. These cases are particularly interesting as the higher number of objects involved might correspond to a complex situation where something creative or resourceful is happening with the co-use of the objects. On the other hand, it could be a random and coincidental anomaly that has no meaning for the participant. Besides looking at the number of co-used objects, we also observe unexpected combinations of object usage. For instance, topic 2 shows a combination of five objects including 'knitting Needle', 'tablet' and other objects in the kitchen. The combination is preceded by using the kitchen door and followed by two other objects (fridge and kitchen chair) again




Figure 3. Extracted topics for one household. Each row is one topic and its associated word events. The color indicates the object and the size corresponds to the probability of the movement events. The x axis shows the hour of the day.

in the kitchen. One possibility could be that the participant uses the tablet to search for food menus. However, the usage of knitting needle in that situation is uncommon and is a case that might be interesting to probe the participant further during field interviews.

This section used two examples to discuss ways to interpret the patterns by looking at shared objects and combination of objects and associate those with situations where resourcefulness might have occurred. Nonetheless, our interpretation has limitations as they are entirely based on the continuous digital traces collected from the sensors attached to the household objects.

CONCLUSION

This paper explored the use of machine learning, specifically topic models, to identify resourcefulness from IoT sensors data. We conceptualize resourcefulness as the creative usage of a combination of objects. We use the movement of objects as a proxy for object usage. The topic models were proposed to extract patterns of co-usage the raw movement data. Our results show that machine learning can extract higher level routine behaviors but alone is not sufficient to identify resourcefulness. It is important to verify the associated patterns identified by the machine learning algorithm with semantic meanings in order to generate meaningful insights to understand resourcefulness. In the future, we will explore the dynamics of usage routines and their patterns at various temporal and spatial resolutions on a larger dataset.

Acknowledgment

This research was supported by the project 'Resourceful Ageing' funded by NWO/STW under the Research through Design program (2015/16734/STW).

REFERENCES

- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022. <http://dl.acm.org/citation.cfm?id=944919.944937>
- M Bishop Christopher. 2016. *Pattern Recognition and Machine Learning*. Springer-Verlag New York.
- Lenneke Kuijer, Iohanna Nicenboim, and Elisa Giaccardi. 2017. Conceptualising Resourcefulness As a Dispersed Practice. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. ACM, New York, NY, USA, 15–27. DOI: <http://dx.doi.org/10.1145/3064663.3064698>
- Yanxia Zhang, Jeffrey Olenick, Chu-Hsiang Chang, Steve W. J. Kozlowski, and Hayley Hung. 2018. The I in Team: Mining Personal Social Interaction Routine with Topic Models from Long-Term Team Data. In *23rd International Conference on Intelligent User Interfaces (IUI '18)*. ACM, New York, NY, USA, 421–426. DOI: <http://dx.doi.org/10.1145/3172944.3172997>