Routine Pattern Learning and Anomaly Detection Applied to Lone Workers Through Topic Modeling

Ana Cravidão Pereira[1], Marília Barandas[1] and Hugo Gamboa[1,2]

[1]*Fraunhofer Portugal AICOS, Rua Alfredo Allen 455/461, 4200-135 Porto, Portugal*

[2]*Laboratório de Instrumentação, Engenharia Biomédica e Física da Radiação (LIBPhys-UNL), Departamento de Física, Faculdade de Ciências e Tecnologia da Universidade Nova de Lisboa, Monte da Caparica, 2829-516 Caparica, Portugal*

Keywords: Pattern Learning, Anomaly Detection, Topic Modeling, Incremental Learning, Lone Workers.

Abstract: Learning routines and identifying anomalous behaviors play a critical role in worker safety. Identifying deviations from normal patterns helps prevent accidents, ensuring enhanced safety in complex environments. Topic modeling is frequently used to discover hidden semantics patterns and is well-suited to the complexity of routines in human behavior. However, its utility in complex time-series analysis and as a baseline for anomaly detection has not been widely explored. This work proposes a novel solution to accurately model complex routines using topic modeling, enabling the identification of anomalies through a statistical approach. A dataset of human movement recordings was collected over up to seven consecutive months, capturing the routines of three lone workers, with each accumulating between 414 and 955 hours of recording time. This dataset served as the basis for a comprehensive analysis of the results, showing strong alignment between visually observed patterns in routines and the outcomes of the proposed method. Additionally, detecting anomalies across models with varying training days confirms that online learning enhances the accuracy and adaptability of routine modeling. Topic modeling allows for in-depth learning of routines, capturing latent patterns undetectable to humans. This capability prevents abnormal events, leading to a proactive approach to predictive risk assessment.

1 INTRODUCTION

Working alone is a common practice across various sectors, spanning from transport and logistics to healthcare and the manufacturing industry. With the intensification of remote working worldwide, it has also become frequent in other fields. According to the UK Ministry of Defence (2008), a lone worker is defined as someone who cannot be immediately assisted by another person during their working day, whether for a short or extended period. The market for lone worker safety solutions is growing as more organizations recognize the need for comprehensive management policies and safety systems to ensure employee well-being. Current strategies primarily focus on detecting falls or prolonged inactivity. However, lone workers may also face other hazards, such as assault or hostage situations, which are inadequately addressed by existing safety reactive measures.

To address these challenges, key methodologies for a more proactive risk prediction and categorization for lone workers in real-time include routine pattern learning, anomaly detection tailored to individual worker routines, and online learning to incorporate new patterns and respond to previously unseen scenarios.

Contributions. This study proposes a framework that learns and adapts to each worker's unique routine patterns, allowing it to identify deviations that may indicate critical anomalies. The proposed solution was evaluated in a real-world setting, involving three lone workers from different sectors, ranging from office work to the manufacturing industry. These workers recorded their daily routines over 45 days across five to seven consecutive months. Results demonstrated the framework's ability to learn routines and detect anomalies in those routines. These contributions resulted from a research path that addressed the following research questions:

- **RQ1.** How can different worker routines be learned in an unsupervised manner?
- **RQ2.** Are anomaly detection methods effective in identifying deviations from an individual worker routine?
- **RQ3.** How do different online learning techniques impact routine pattern learning?

Pereira, A. C., Barandas, M. and Gamboa, H.
Routine Pattern Learning and Anomaly Detection Applied to Lone Workers Through Topic Modeling.
DOI: 10.5220/0013236600003911

In *Proceedings of the 18th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2025) - Volume 1*, pages 935-942
ISBN: 978-989-758-731-3; ISSN: 2184-4305

2 RELATED WORK

Over recent years, various strategies have been employed for pattern learning on human movement, utilizing activity and spatial routines on daily or weekly cycles. In the context of pattern learning, topic modeling emerges as a relevant technique. Huynh et al. (2008) contributed to the early application of topic modeling for recognizing daily routines using time-series activity data. Their approach begins by clustering features relevant to routines, followed by applying Latent Dirichlet Allocation (LDA) to uncover hidden topics within these routines. Similarly, a related study from Tang (2021) explored the use of LDA to identify location patterns in user trip behaviors using Global Positioning System (GPS) coordinates. However, this study did not consider the temporal factor of patterns. Steinhauer et al. (2019) adopted a similar approach in their study of telecommunication networks, interpreting run-time variable readings as words in LDA to obtain representative topics. Furthermore, Seiter et al. (2014) employed LDA to discover activity routines of stroke rehabilitation patients, subsequently applying KNN to map the discovered topics to specific routines. Sun et al. (2021) used a two-dimensional probability distribution, a variation of LDA, to discover routine patterns in individual travel behavior. This work does not consider the order of occurrences, essential for routine comprehension.

In the context of anomaly detection in human routines, Sun et al. (2021) also explored the detection of anomalies by using perplexity as a scoring function to assess whether new behaviors were predictable (thus normal) or anomalous. Steinhauer et al. (2019) also aimed to detect anomalies following topic modeling by establishing a ”normal model” and relying on an expert's visual comparison to confirm the presence and nature of anomalies. In their study on detecting patterns in anomalous clusters, Abraham and Nair (2018) applied LDA to define the normal topics within the data, subsequently using a statistical test to detect anomalies. Similarly, Thornton et al. (2020) used LDA to identify normal patterns in avionics networks and employed a combination of variational inference and statistical analysis to detect anomalies.

Based on studies on routine pattern learning, topic modeling has shown promising results for in-depth analysis of human movement routines. This work will therefore explore its potential in complex time series analysis, focusing on a critical use case: lone worker safety. Human routines are complex, and a few days or weeks may be insufficient to capture them accurately. To the best of our knowledge, no study has addressed long-term routine learning in human movement using both activity patterns and GPS data and only a few have focused on anomaly detection within individual routines, which is the aim of this study.

3 METHODS

The primary aim of this work is to develop a generalized approach to predictive risk assessment based on deviations from a user's routine within the time-series domain. To achieve this, a lone worker use case was selected, and a general approach for routine learning, anomaly detection, and model updating was developed. Figure 1 provides an overview of the proposed approach. For the routine learning stage, a topic modeling technique was applied, consisting of a clustering step, document creation, LDA model training, and topic activation. For the anomaly detection stage, a statistical approach was developed. Since the solution is designed for real-life implementation, routine evolution was also considered. Therefore, model updates were explored, using both retraining and Incremental Learning (IL) approaches.

3.1 Datasets and Preprocessing

For this study's dataset, three lone workers volunteered to collect data on their daily work routines during the first semester of 2024. The typical daily routine of LW1 includes a commuting period via subway or car and a work period in an office. LW2 commutes by car and has work periods that include being in an office and walking outdoors within a specific region. Lastly, LW3 is a manufacturing worker who has a long commuting period by car and spends part of the workday walking through facilities.

The recording protocol involved installing a logger Android application on the workers' smartphones, which was activated by the user at the beginning and end of each workday. During the collection period, the lone workers did not record data every weekday due to holidays, application issues, battery drain, and forgetfulness in activating the application. As a result, this led to a total of 45, 46, and 81 days of usable data for routine learning for LW1, LW2, and LW3, respectively. Additionally, the volunteers were asked to annotate any days on which activities and/or locations occurred that were not part of their routine. This request was made to prevent the contamination of routine learning with anomalous patterns. Therefore, the test days were selected based on these annotations, corresponding to a percentage between 23% and 34% of all available days. Table 1 summarizes the dataset details for the data collected from each lone worker.

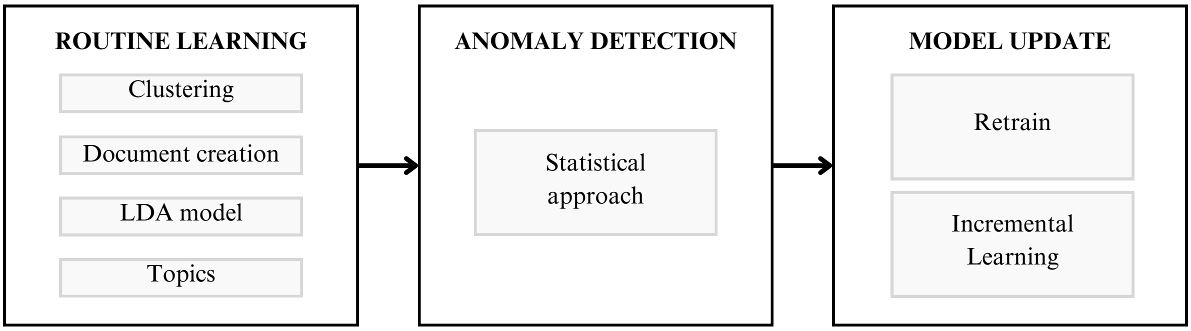

Figure 1: Overview of the main stages of the proposed approach.

Table 1: Summary of data for lone workers. Mean day duration and total duration are presented in hours

	LW 1	**LW 2**	**LW 3**
Number of train days	30	35	62
Number of test days	15	11	19
Mean day duration	9,22	11,03	11,79
Total duration	414,69	507,43	955,37

The application recorded GPS information (latitude, longitude, and accuracy), human activity data (i.e., a set of pre-defined, automatically recognized activities with associated confidence levels), and temporal records. The activity data was obtained with the Activity Recognition API from Google (2021), which detects user activities through mobile device sensors. The detected activities are *in vehicle*, *on bicycle*, *on foot*, *running*, *still*, *tilting*, *unknown*, and *walking*.

For preprocessing the human activity data, low-confidence activity labels were filtered out using a threshold of 80%. For GPS information, only locations with an accuracy of up to 30 meters were considered. Finally, to align the time intervals for both activities and locations, we applied linear interpolation with a one-minute step, resulting in a data structure where each minute contains location data and a confidence value for each activity label. The time of day, weekday, and day of the month were added as features in the final data structure.

3.2 Routine Learning

This section is divided into two subsections: the first provides a brief overview of topic modeling, and the second presents our proposed approach using topic modeling for routine learning.

3.2.1 Topic Modeling

Topic modeling is a methodology mainly used in text processing that allows the discovery of hidden patterns and structures, i.e., topics, in collections of data. In the case of a text document, topic modeling intents to find common patterns across the document and to understand how its words are associated to make up the existing topics. As long as the elements that compose the document can be processed into sequences of characters, latent patterns can be discovered. Thus, this methodology can also be applied to routine identification. The most common approach for topic modeling is LDA, a generative probabilistic model developed by Blei et al. (2003). A graphical model representation of LDA is presented in Figure 2, illustrating how LDA transforms a collection of documents into a probabilistic distribution of topics. Topic modeling assumes documents are formed by combining the words from the existing topics in the correct proportions - both the proportions of the words within the topics and of the topics within the document. This document is the input to train the model.

3.2.2 Proposed Approach

Before applying the LDA model, the input data structure is normalized, followed by the K-means clustering algorithm. The output clusters represent crucial combinations of data elements that form the routine. The elbow method was used to determine the optimal number of clusters, having been chosen 7 clusters for LW1 and LW3 and 8 clusters for LW2. Additionally, one drawback of K-means is that it does not handle points far from the centroids as outliers, assigning them to one of the known clusters. To address this issue, we adapted K-means by creating an ”outlier cluster” for anomalous events. During the training phase, no points are assigned to this outlier cluster. However, in the testing phase, if the distance between the data point and the cluster centroids exceeds three times the

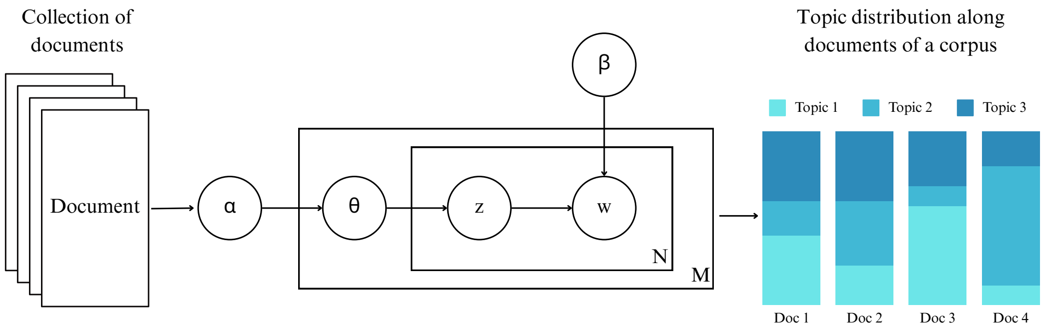

Figure 2: Graphical model representation of LDA. The outer plate, M, represents documents and the inner plate, N, represents the words and topics within a document. Each word, w, is assigned to a topic, whose association is represented by z. Parameter α defines the topic distribution within the corpus, θ defines the topic distribution within a document and β defines the word distribution within topics. The diagram is presented in the work of Blei et al. (2003).

standard deviation of the training data, is assigned to the outlier cluster.

Before creating the document, distances must be transformed into probabilities so that shorter distances to a centroid correspond to higher association probabilities with its cluster. This conversion was done according to the following equation presented in the work of Huynh et al. (2008):

$$\omega_i = \frac{e^{-d_i/\sigma}}{\sum_{j=1}^{K} e^{-d_j/\sigma}}, \tag{1}$$

where ω_i represents the weight, i.e., probability of a cluster, σ denotes the standard deviation of the distances of all instances, d_i is the distance of the instance to the centroid of the cluster to which it belongs, d_j is the distance of the instance to the centroid of cluster j and K is the number of clusters.

Since an event spans several seconds rather than a single point in time, a sliding window of ten minutes was defined. This window captures information over intervals by summing the probabilities associated with each instance within that time frame and normalizing the result. The window size was selected considering the duration of the events targeted. For each sliding window, it was produced a list whose elements were the clusters, with their quantities proportional to the probabilities obtained. The collection of these lists forms the document used to train the model.

The LDA model trained outputs an activation function that indicates the composition of the topics, i.e., the contribution of each cluster to each topic, which is illustrated by the following example:

$$\begin{aligned} \text{Topic}_1 &= \text{Cluster}_1 \times a_{1,1} + \ldots + \text{Cluster}_K \times a_{1,K} \\ &\ldots \\ \text{Topic}_N &= \text{Cluster}_1 \times a_{N,1} + \ldots + \text{Cluster}_K \times a_{N,K} \end{aligned} \tag{2}$$

where N represents the number of topics, K the number of clusters and $a_{N,K}$ corresponds to the activation value of Cluster_K in Topic_N. This function is used to decode the latent structures in the test data, which undergoes the same processing as the training data, by multiplying the cluster probabilities by the activation function factors.

Despite the creation of an outlier cluster, the number of topics remains unchanged, as topics are influenced only by the clusters identified during training. Therefore, this approach does not generate a topic that detects anomalies. Instead, as highlighted by Equation 2, it minimizes the values of all other topics, which might be an indicator of an anomaly.

3.3 Anomaly Detection

To the best of our knowledge, very few studies have applied anomaly detection after pattern learning in time series and no standard method exists. Therefore, we implemented a statistical approach over the pattern learning results to identify potential deviations. This approach was applied to compare the behavior of the topics in the training data with the activation of the topics in the test data. For this purpose, the mean value and corresponding standard deviation were calculated for each topic and each timestamp across all training days.

Let z be a set composed of N topics, $z = \{z_1, z_2, \ldots, z_N\}$, which form a user's routine. Let $z_{x,i}$ denote the topic z_x at time instant i, and the mean and standard deviation across all training days for $z_{x,i}$ be represented by $\mu_{z_{x,i}}$ and $\sigma_{z_{x,i}}$, respectively. Given a test instance where $\tilde{z}_{x,i}$ represents the activation of $z_{x,i}$, an anomaly is detected through the following rule:

$$\sum_{x=1}^{N} \mathbb{1}\left\{\left|\tilde{z}_{x,i} - \mu_{z_{x,i}}\right| > \sigma_{z_{x,i}}\right\} > \frac{N}{2} \tag{3}$$

where $\mathbb{1}\{.\}$ is an indicator function that takes the value 1 (anomaly) if the expression inside is true, and 0 (normal) otherwise. An anomaly is detected if more

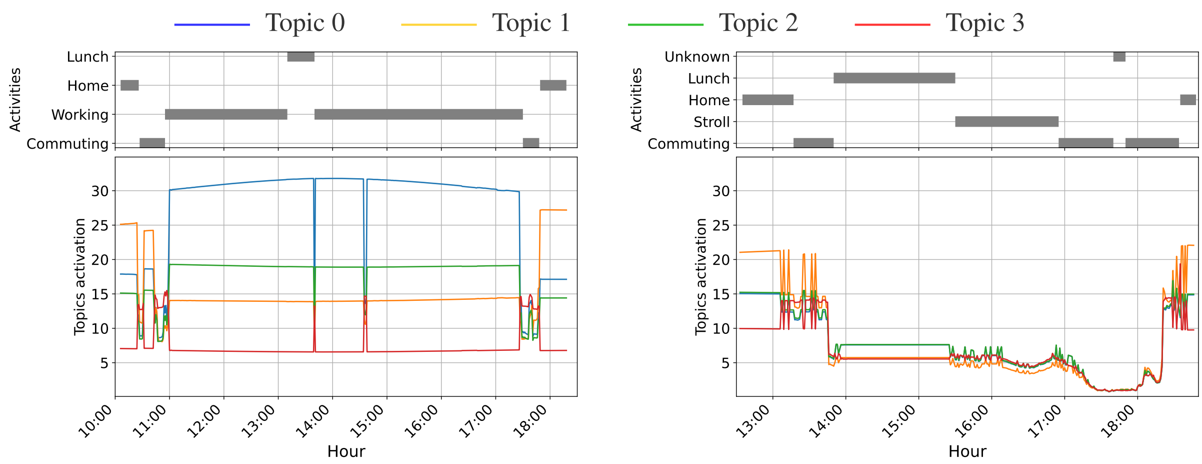

(a) Normal day. (b) Anomalous day.

Figure 3: Top: Annotation of activities performed along the complete workday of a normal and anomalous day. The annotated activities are not exhaustively descriptive. For instance, *Commuting* includes walking to and from the train station, subway travel, and car transportation. Bottom: Topic activations from the model selected, applied to the same test day.

than half of the topics at time instant i deviate from their expected range, defined as $\mu_{z_{x,i}} \pm \sigma_{z_{x,i}}$.

A post-processing technique was implemented to require a minimum streak of samples flagged as anomalies to confirm the detection, and a minimum streak of normal samples to revert the status to normal. This safeguards the algorithm against inaccuracies and insignificant variations.

3.4 Model Update

Anomalies might be detected based on a pre-established number of days. However, routine originally identified might have evolved. Therefore, it is essential to maintain an updated routine baseline.

For the purpose of model updating, two approaches were studied: (a) Retraining the models from scratch and (b) Applying Incremental Learning (IL). While IL is computationally lighter but may result in catastrophic forgetting, model retraining considers all data equally. In the case of the LDA model, the implementation used supports incremental updates and, therefore, both approaches were tested. However, the approach of topic modeling includes a preliminary step: the clustering algorithm before the LDA model. In the IL approach, the clusters are not updated with the new data, as it would demand producing a new document for the LDA model and, therefore, require retraining it instead of incrementally updating it. Hence, it was assumed that the features (i.e. clusters) are well-defined based on the initial training data.

4 RESULTS AND DISCUSSION

Regarding routine learning, it is difficult, if not impossible, to access all factors that compose a user's routine. Without a ground-truth routine, quantitative evaluation is compromised, requiring relative comparisons with other methods and the visualization of semi-annotated data. Therefore, to evaluate the model, we visualized topic activations across test days, comparing them with the semi-annotated data.

The evaluation of anomaly detection was divided into two phases. Firstly, it was hypothesized that the number of detected anomalies decreases with the number of training days, reflecting a stabilization of the learned routine. Secondly, it was evaluated whether the expected anomalies were being detected. To achieve this, LW1 was asked to annotate perceived anomalies in their routine, allowing for the application of quantitative measures on the annotated data. However, it is important to note that users' perception of an anomaly is often limited compared to the model's considerations. As a result, these annotations may be biased. For this reason, the evaluation focuses exclusively on the limited set of annotated anomalies, calculating the true positive rate for assessment purposes.

Finally, for online learning evaluation, a comparison was made between retraining the model from scratch and applying incremental updates over different training days, using the number of detected anomalies as a proxy for performance.

In summary, the performance evaluation comprised the following steps to address the previously

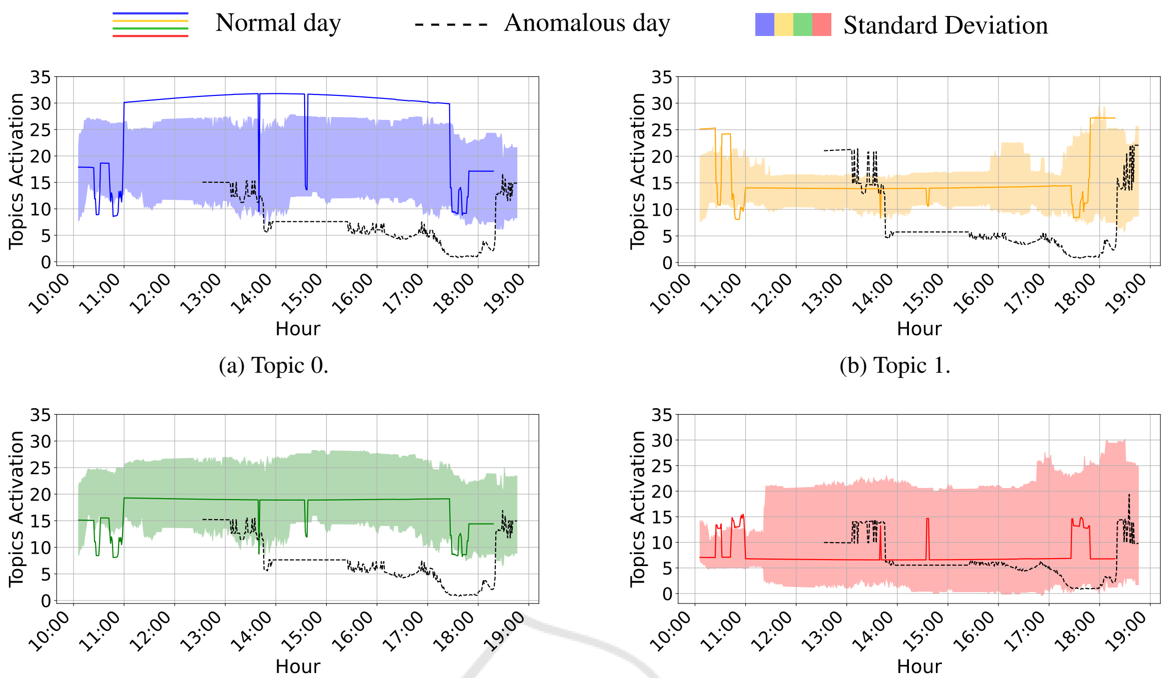

(a) Topic 0. (b) Topic 1. (c) Topic 2. (d) Topic 3.

Figure 4: Comparison of activated topics on a normal and an anomalous day, overlaid with the normality range obtained from training days.

identified research questions:

1. Visual analysis of different test days compared to the learned routine (**RQ1**).
2. Evaluation of the number of detected anomalies as a function of the number of training days used for routine learning (**RQ2**).
3. Comparison of the percentage of anomalies correctly detected (true positive rate) across the different tested models (**RQ2**).
4. Comparison of model updating strategies using the number of detected anomalies as a proxy for performance evaluation (**RQ3**).

4.1 Routine Learning: Visual Analysis

LW1 semi-annotated one typical and one anomalous day from their routine to assess how well the model learned the routine. Figure 3 illustrates the activation of topics throughout the workday, along with semi-annotated labels for both days. On the normal day (Figure 3a), it is observed that the topics activated by the model closely correspond to the annotations shown on top. For example, the periods labeled *Home*, *Commuting* and *Working* display a consistent topic configuration within their respective pairs. This alignment suggests that the model effectively extracts meaningful information from the data, organizing it into logical categories. On the anomalous day (Figure 3b), the topics do not align with the annotated activities. While both *Home* periods display a consistent topic configuration, all other periods exhibit topics indicative of anomalous behavior. Throughout most of the day, changes in annotated activities are indistinguishable based on the activated topics, suggesting that these activities are not recognized due to their absence in the training data.

Another visualization uses the proposed approach for anomaly detection. It is assumed that topics activated during the testing days should fall within each topic's mean $\pm$ standard deviation (referred to as the normality range), to be considered normal. Thus, each activated topic on both the normal and anomalous days is shown in Figure 4, along with the corresponding normality range obtained during training. For most of the normal day, Topics 1, 2 and 3 fall within the normality range while Topic 0 does not. It is important to note that even on training days, not all topic activations consistently fall within the normality range, emphasizing the importance of the combination of topics in routine formation. In contrast, on the anomalous day, all topics except Topic 3 fall outside the normality range for most of the day, confirming its anomalous nature.

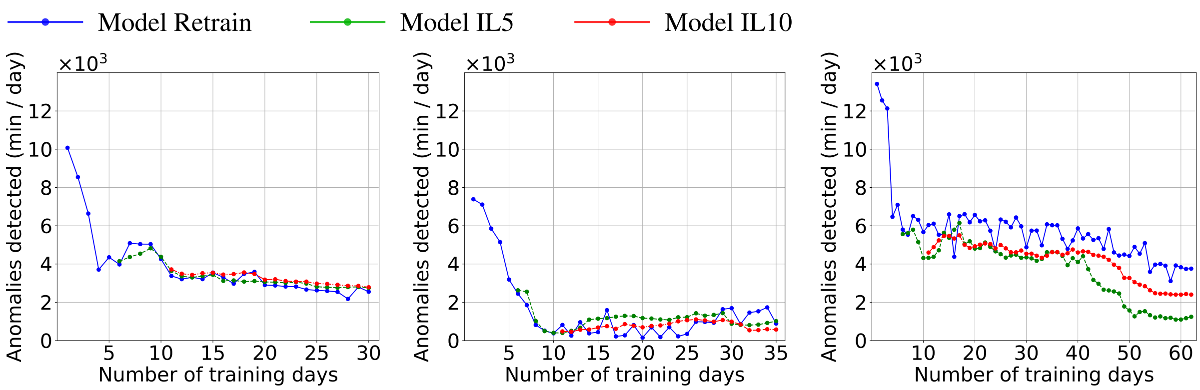

(a) LW1 (b) LW2 (c) LW3

Figure 5: Number of anomalies detected with the increased number of training days for the three lone workers and three tested approaches: Model Retrain, IL5, and IL10.

4.2 Anomaly Detection and Model Updates

To evaluate anomaly detection performance and address research questions **RQ2** and **RQ3**, the topic modeling technique was tested in the context of model updating using the following approaches:

1. **Model Retrain.** Topic model retrained each time new data is added to the training set.
2. **Model IL5.** Topic model trained with IL, starting from a base model of five days and updating each new day individually.
3. **Model IL10.** Topic model trained with IL, starting from a base model of 10 days and updating each new day individually.

As the number of training days grows, it is expected that the number of detected anomalies decreases, indicating that the routine is being better learned and that instances previously considered anomalous are now recognized by the model. This hypothesis is confirmed in Figure 5. For all lone workers' routines, it is observed a decreasing trend in detected anomalies across all models evaluated, with a more pronounced drop during the first 10 days. These results suggest that ten consecutive days are sufficient for the model to learn a user's routine. However, a gradual decrease in detected anomalies is still evident for LW1 and LW3, indicating that new data continues to provide learnable information. In the case of LW2, this decreasing pattern appears to plateau earlier, with the Model Retrain approach even showing an increase in detected anomalies toward the final days. It is observed a more stable pattern in IL models, whereas Model Retrain shows more fluctuation, with each new day having a greater impact. Comparing the three approaches, no clear conclusions can be drawn from this analysis, as the differences in anomalies detected between models do not appear to be significant across users. More information about routine data would be needed to determine which model performs best.

Regarding the correctly detected anomalies evaluation, a rough annotation of anomalies was performed for the test set of LW1 to serve as ground truth. It was established that for an anomaly to be correctly detected, the model must identify at least one instance within the annotated anomaly. In Figure 6, the true positive rate of annotated anomalies was evaluated across the different models. This analysis reveals that Model Retrain tends to outperformed both IL5 and IL10 models with more training days, but there is not a model that consistently produces a higher true positive rate. Additionally, there does not appear to be a significant difference between the performances of Models IL5 and IL10, likely because the routine learning stabilizes by day 10, as indicated in Figure 5.

5 CONCLUSIONS

In the context of lone workers, detecting anomalies in their routines can help prevent risky situations and promote preventive actions.

To discover the patterns that form the routines, a topic modeling approach was implemented. This method reveals hidden patterns, thereby identifying relationships that, although coherent with annotations, are not readily discernible to humans due to the complexity of routines. Through visual inspection, it was found that the LDA model identified routines, and the derived topics accurately represented the human activities performed. Moreover, it was concluded that the routine does not correspond to a single topic, but rather to the combination of all topics activated.

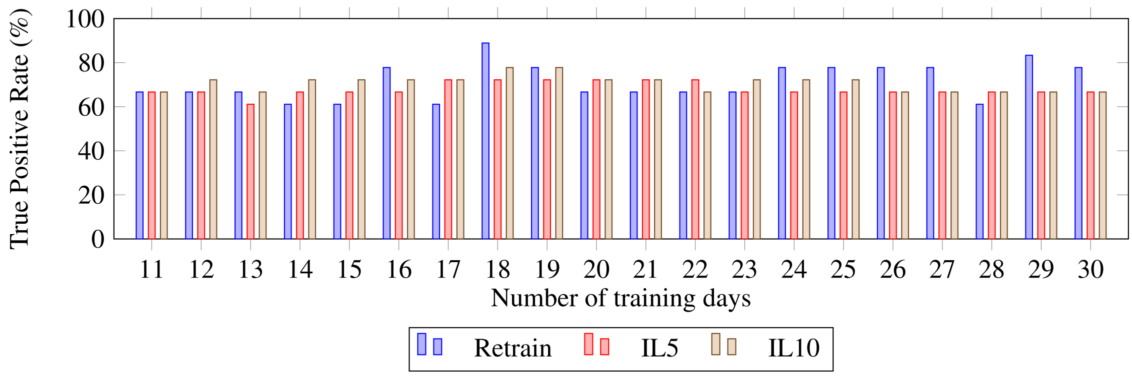

Figure 6: Comparison of correctly detected anomalies across different days for all models on data from Lone Worker 1.

The patterns identified were subsequently used as the ground truth for defining the routine and detecting any deviations from it. The anomaly detection results were satisfactory, demonstrating effective performance in detecting anomalies that had been annotated. However, this evaluation is not straightforward, as the annotated anomalies do not fully correspond to the ground truth, due to their complex nature.

Finally, the routine's continuous evolution was addressed using IL techniques. A comparison of the results between models updated with IL and those fully retrained shows that the latter detected a higher number of anomalies when provided a higher number of training days. However, further studies are necessary to draw more definitive conclusions.

While the results presented are promising, there are some limitations and future work to consider. Firstly, parameters such as the number of topics need further testing and optimization. Additionally, the statistical approach to anomaly detection does not consider potential relationships between activated topics, although the results indicate a strong inter-topic relationship. Thus, alternative anomaly detection methods should be explored, and more annotated data should be collected to enhance the robustness and generalizability of the methodologies. This would also facilitate a more thorough comparison between retraining and incremental learning.

ACKNOWLEDGMENTS

This work was supported by European funds through the Recovery and Resilience Plan, project "Center for Responsible AI", project number C645008882-00000055.

REFERENCES

Abraham, B. L. and Nair, A. P. (2018). Anomalous topic discovery based on topic modeling from document cluster. *International Research Journal of Engineering and Technology (IRJET)*, 5(2):966–972.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(null):993–1022.

Google (2021). Google Activity Recognition API. https://developers.google.com/location-context/activity-recognition. [Accessed: November 4, 2024].

Huynh, T., Fritz, M., and Schiele, B. (2008). Discovery of activity patterns using topic models. pages 10–19.

Seiter, J., Derungs, A., Schuster-Amft, C., Amft, O., and Tröster, G. (2014). Activity routine discovery in stroke rehabilitation patients without data annotation. *Proceedings - REHAB 2014*, pages 270–273.

Steinhauer, H., Helldin, T., Mathiason, G., and Karlsson, A. (2019). Topic modeling for anomaly detection in telecommunication networks. *Journal of Ambient Intelligence and Humanized Computing*, 14.

Sun, L., Chen, X., He, Z., and Miranda-Moreno, L. (2021). Routine pattern discovery and anomaly detection in individual travel behavior. *Networks and Spatial Economics*, 23.

Tang, T. (2021). Characterizing shared mobility operator and user behavior using big data analytics and machine learning. Master's thesis, University of Virginia, School of Engineering and Applied Science.

Thornton, A., Meiners, B., and Poole, D. (2020). Latent dirichlet allocation (lda) for anomaly detection in avionics networks. In *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, pages 1–5.

UK Ministry of Defence (2008). *JSP 375 Volume 1: Management of Health and Safety in Defence: Arrangements*. Last updated December 7, 2023.