

“Chatty Devices” and Edge-based Activity Classification

Mike Lakoju

Cardiff University

Amir Javed (✉ javeda7@cardiff.ac.uk)

Cardiff University

Omer Rana

Cardiff University

Peter Burnap

Cardiff University

Samuelson Atiba

Cardiff University

Soumaya Cherkaoui

Université de Sherbrooke

Research Article

Keywords: Product Use Activity classification, IoT, Chatty Factories, Federated Learning, Edge computing.

Posted Date: November 20th, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-109382/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Discover Internet of Things on February 24th, 2021. See the published version at <https://doi.org/10.1007/s43926-021-00004-9>.

“Chatty Devices” and Edge-based Activity Classification

Mike Lakoju · Amir Javed ·
Omer Rana · Pete Burnap · Samuelson
T. Atiba · Soumaya Cherkaoui ¹

Received: date / Accepted: date

Abstract With increasing automation of manufacturing processes (focusing on technologies such as robotics and human-robot interaction), there is a realisation that the manufacturing process and the artefacts/products it produces can be better connected post-production. Built on this requirement, a “chatty” factory involves creating products which are able to send data back to the manufacturing/ production environment as they are used, whilst still ensuring user privacy. The intended use of a product during design phase may differ significantly from actual usage. Understanding how this data can be used to support continuous product refinement, and how the manufacturing process can be dynamically adapted based on the availability of this data provides a number of opportunities. We describe how data collected on product use be used to: (i) classify product use; (ii) associate a label with product use using unsupervised learning – making use of edge-based analytics; (iii) transmission of this data to a cloud environment where labels can be compared across different products of the same type. Federated learning strategies are used on edge devices to ensure that any data captured from a product can be analysed locally (ensuring data privacy). Using a 6th gen. Apple iPad as a “chatty device” (with acceleration, orientation, angular velocity and magnetic field sensors) we demonstrate how product use activities can achieve a classification accuracy of 99.35%. A comparison is also undertaken with the Human Activity Recognition (HAR) data set, achieving an accuracy of 98%. Our approach demonstrates how semantic activity labels can be associated with product use, and subsequently used to improve product design.

Keywords Product Use Activity Classification · IoT · Chatty Factories · Federated Learning · Edge computing.

School of Computer Science & Informatics, Cardiff University, 5 The Parade, Cardiff CF24 3AA, UK Email: Lakojum@cardiff.ac.uk · ¹ Interlab Lab., Dept. of Electical & Comp. Engineering, University of Sherbrooke, Canada

1 Introduction

Increasing availability and reduction in size of sensors and actuators has enabled integration of these devices within a number of physical objects. The use of sensors within built environments has existed for a number of years, enabling control systems (e.g. Building Energy Management Systems) to process and undertake control actions based on data obtained from such sensors. The term “Industry 4.0” has drawn significant attention over recent years, providing an ecosystem of sensing, actuation and analytics to support product development and process automation. A key theme is the transformation of a production environment from manual (but machine supported) design and manufacturing to digital design and manufacturing [39] – often also the basis for significant recent interest in developing “digital twins”. Within this paradigm, interconnected computer systems, intelligent devices and smart materials communicate with one another while interacting with the environment with little to no human involvement [15]. The manufacturing industry, in particular, has been significantly impacted by the use of such sensing and actuation technologies [37]. In the typical manufacturing system, taking a product to market requires multiple discrete and highly specialised activities across research, design and manufacturing disciplines. The *Chatty Factories* vision for the manufacturing factory of the future is to take these traditionally discrete activities and collapse them into one seamless process that is capable of continuous real-time product refinement at an industrial scale. The Industrial Internet of Things (IIoT) offers the opportunity to draw new insights through real-time data harvested from sensors embedded in products. “Chatty Factories” seeks to leverage product use data from “Chatty” (smart devices that are able to communicate their usage back to the factory) devices and equip design engineers and manufacturers within the Industry 4.0 paradigm with the ability to understand how the products are being used *in the wild*, providing the basis for new forms of creative design and leading to tailor-built manufacturing processes. A key challenge, as outlined in this work, is to ensure that product use data is collected whilst still preserving user privacy. Placing IoT-enabled data driven systems at the core of design and manufacturing processes provides a number of opportunities for refining product design and aligning this more closely with its actual (vs. estimated/ expected) use. We focus on collecting data from IoT-enabled sensors embedded in products during real-time use by consumers, to understand how that data might be immediately transferred into usable information to inform design. We also consider what characteristics of the manufacturing environment might optimise the response to such data in a privacy-preserving manner. However, achieving continuous product refinement in response to real-time data on product use requires: data driven systems that provide an auditable, secure flow of information between all operations inside and outside the factory. By implication, this raises concerns around privacy, security and scale of data. Recent approaches to achieve this include Federated Learning, which can be used to train a globally shared model by exploiting a massive amount of user-generated data samples on sen-

sorized/ chatty devices while preventing data leakage [49]. A *surrogate* model is developed on each device, which can then be combined at a cloud data centre. Product use activity classification is carried out on data at each device at the network edge, using federated learning approaches. This is achieved through the use of an Autoencoder at the edge to identify unknown product use activities, and then uses unsupervised machine learning methods to identify repeated patterns/ anomalous activities, to distinguishing between noise and actual product use. Furthermore, we use entangled ethnography to create labels for these newly identified activities. The new use activity labels created are pushed to the cloud which aids federated learning and model correction from a collective of chatty devices while maintaining appropriate privacy and trust. This paper contributes to the broader literature on edge computing and federated learning by:

- Proposing a novel system that identifies *product-use* activities at the edge, thereby ensuring user privacy while providing for model correction;
- Demonstrating the system’s capability to detect unknown *product-use* activities on chatty devices at the edge;
- Recognising previously “seen” activities in the presence of noise;
- Rigorously testing the novel system against two independent datasets (a chatty dataset collected by us and a publicly available dataset on human activity recognition).

2 The Chatty Factories Vision

The Chatty Factories project explores the transformative potential of placing IoT-enabled data driven systems at the core of design and manufacturing processes. The aims of the research project can be summarised as [8]: (i) Enabling new forms of agile engineering product development via “chatty” products that continue to relay their experiences from consumers back to designers and product engineers—through the mediation provided by embedded sensors, IoT and data driven design tools; (ii) Enabling the manufacturing environment to dynamically respond to changes in physical configuration and ethical re-skilling of production elements such as robots and humans by utilising theory from expopedagogy and interpretable data analytics; (iii) Confronting the challenges of data volume, privacy and cybersecurity to develop an access-controlled manufacturing ecosystem in which product use data will be collated, analysed and disseminated across the factory floor, merging Operational Technology (OT) on the factory floor with Information Technology (IT) in the business. Figure 1 illustrates the high-level vision of Chatty Factories. With the “Chatty” concept to understand how products are used in the wild starts by using sensors retrofitted on products in “the wild”. The chatty process flow is initiated at the “Product Data from the wild” phase as illustrated in Figure 1. Acquired product use labels are subsequently transferred for data annotation and pattern detection at the “Data Annotation” phase. As denoted by the “Product Use Models” phase, the successful labelling of newly identified product use

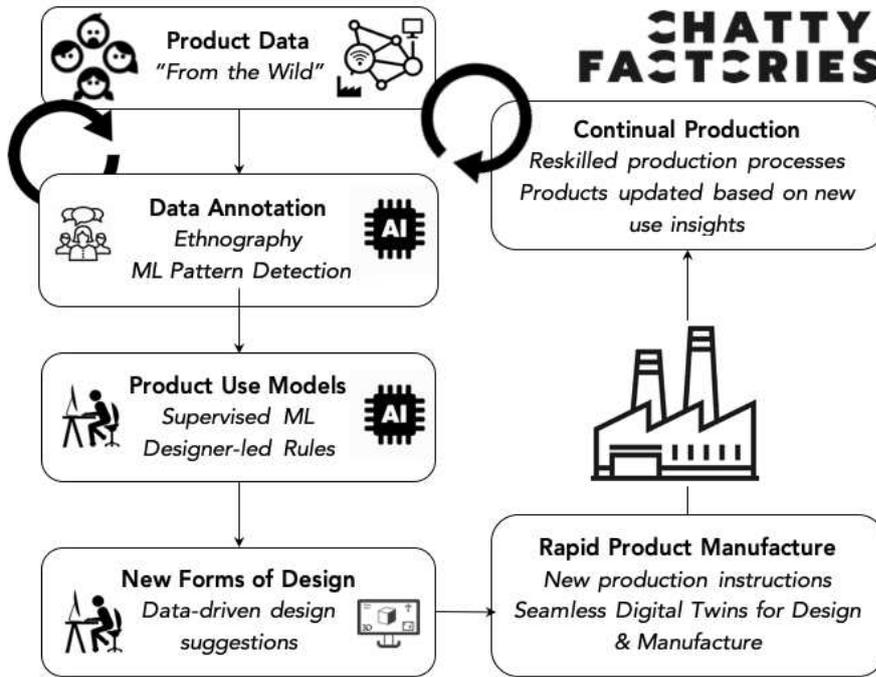


Fig. 1 The Chatty Factories Vision

activities is followed by the creation of product use models through supervised machine learning algorithms assisted by designer led rules. After various analyses, new forms of data-driven design suggestions are put forward to the design team at the “New Forms of Design” phase. Valuable insights gathered at this stage inspire the creation of enhancements to existing products and/or the creation of a totally new product. As illustrated by the “Rapid Product Manufacture” phase, this design and manufacturing process are rapidly sped up by the interaction between the design digital twin and the manufacturing digital twin. The entire process runs constantly and iteratively. This is highlighted by the “Continual Production” phase in Figure 1. The approach proposed in this work automatically produces new activity labels (based on users behaviour) on Edge devices while protecting the privacy of the user. Furthermore, the proposed system can be adapted by analysing the *social chatter* between devices.

3 Related Work

3.1 Activity Classification

While research on product use activity recognition is still in its infancy, there has been substantial research dedicated to activity recognition and classifica-

tion especially in human activity classification, smart homes, activities of daily living (ADL) and medical domains. Advancements in sensor technology have fuelled the growing interest in context-aware systems for different domain-specific applications [18] [48] [42]. This provides a better understanding of sensing and reasoning to establish a context for people’s actions. Researchers have particularly considered how data from sensors can be used to investigate fall and activities of daily living (ADL) [9],[34],[23]. General research on activity recognition with sensor technology can be categorised into three general themes: wearable sensors, vision-based detection and environmental sensors [23],[21], [36], [4].

There is also interest in vision-based detection sensors for identifying patterns using image processing techniques. However, a notable limitation lies within its coverage area and the obvious inability to provide awareness about users out of visibility range. Ambient sensor-based systems have the capability of tracking movements, and can be used for analysis of gait and trajectory of the subjects/ users [14], [12] movement. The system uses sensor readings of pressure, sound, vibrations of the floor and infrared imaging to determine and characterise movements. These approaches are still constrained by the perimeter (radius of the area) of sensor coverage.

Wearable sensors (continuously reducing in cost) are a hybrid of different types of environment sensors. Wearable devices are commonly used in sports, health applications, entertainment etc. [3][30]. The availability of this category of sensors is inspiring more innovative research applications around activity recognition along with efficient processing techniques [38][16]. Sensors such as an accelerometer, a gyroscope, a magnetometer, a camera etc are now not only wearable but can also be retrofitted into high-value products, thus enabling an ecosystem of sophisticated systems which can aid the analysis of any information captured by them, without a restriction of being indoors or outdoors.

Li et al. [28] carried out experiments on fall detection using a sample thresholding method between a gyroscope and an accelerometer. Lee and Carlisle [27] proposed a two-thresholding method to identify movements and the different simulated falls using a smartphone and accelerometers. Human behaviour research has gained significant attention due to its relevance in pervasive and mobile computing, health, ambient assistive living, surveillance-based security and context-aware computing [48]. For example, Tabia et al. [47] identified human activities such as running, jogging and running by using Naive Bayes and Support Vector Machine (SVM) classifiers. Jayasinghe et al. [22] conducted experiments using a combination of clothing mounted sensors and wearable sensors for movement analysis and activity classification with a focus on running, walking, sitting and riding in a bus. By calculating correlation coefficients for each sensor pair, they suggest that even though the two data streams have some notable differences, results indicate high classification accuracy.

Moussa et al. [35] approached human activity classification from a spatial features-based method. They utilised the Scale Invariant Feature Transform (SIFT) algorithm to identify key points of interest. Furthermore, they created

a Bag of Words using the K-means algorithm, which aided the assignment of each descriptor generated from the key points of interest to the nearest visual word. They also calculated the frequency histogram of the visual words. Lastly, an SVM classifier was used to determine the human activity class.

By using a wearable Shimmer device, Kerdjidj et al. [23] proposed a system for fall detection and detection of activities of daily living (ADL). Their experiment employed 17 subjects performing a set of movements. They focused on three distinct systems: first detects the presence or absence of a fall; the second detects static or dynamic movements including the fall, and the last recognizes the fall and six other ADL activities. Furthermore, in order to reduce the size of transmitted data and also energy consumption, they applied a compressive sensing (CS) method, reporting an accuracy of 99.9% from their experiment. Similarly, Hui and Zhongmin [20] proposed the use of a compressed sensing method to identify activities in which a combination of acceleration data and phone placement information are used. Their experiment focused on three placements of mobile phones: trouser pocket, handbag and the hand.

3.2 Federated Learning

Federated Learning (FL) is a unique implementation of a distributed machine learning approach that enables training on decentralised data residing on devices like mobile phones, tablets or any chatty device [32],[6]. FL is one instance of *pushing code to the data*, rather than data being pushed to the code which therefore supports user privacy, ownership and locality of data. A more detailed background of Federated Learning has been provided by McMahan & Ramage [31]. Likewise, Konecny et al. [26] and McMahan et al. [33] have elaborated on its theory. Federated Learning infrastructure implementation allows for either asynchronous or synchronous training algorithms. Even though literature reveals substantial successful work on deep learning using asynchronous training, e.g Dean et al. [13], more recent studies have tilted towards synchronous large batch training even in the data centers [45],[17]. The McMahan et al. [32] – a Federated Averaging algorithm also follows a similar approach. Similarly, several studies including differential privacy by McMahan et al. [33] and Secure Aggregation by Bonawitz et al. [7] also use FL to enhance privacy. They operate on the notion of synchronization on a fixed set of devices, which allows the server-side of the learning algorithm to only consume a simple aggregate of the updates from multiple users.

4 Systems Design and Implementation

In this section we describe how devices on the edge equipped with a baseline autoencoder model can be used to identify known product use activities. Using a threshold-based detection mechanism, if the results of the activity cannot be classified into a known class, an unsupervised machine learning model is

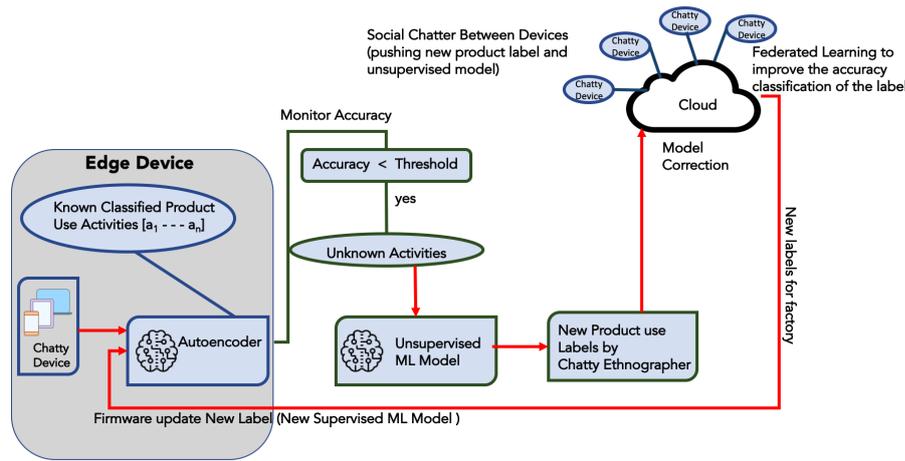


Fig. 2 Flow of work

automatically triggered to cluster such an activity. If the outcome of the unsupervised model identifies a particular cluster which persists over time, we posit that the cluster is a new activity and not just random noise that the autoencoder is unfamiliar with. At this point, an ethnographer creates new labels for the this activity cluster – enabling these clusters to to be aggregated based in similarity using federated learning mechanisms on a cloud platform. Cloud aggregation also enables model correction, enabling this new activity detection to be deployed on the edge. This flow of work is illustrated in Figure 2.

4.1 Using Federated Learning

The use of federated learning is particularly suited to “Social Chatter” for two reasons. First, the data generated by a device is directly related to how the device is being used. This essentially leads to non-independent and identically distributed (non-iid) and unbalanced distributions. Second, the chatty device’s data can also happen to be privacy sensitive. In a general federated learning approach, a global model is sent to several devices. This model is then trained at each device using local data. Information about the locally trained models (e.g. gradients) are then sent from the devices back to the central server (e.g. at the cloud) where they are aggregated to form a combined enhanced global model which is sent back to the devices for subsequent rounds of training. For the edge-based activity classification, the baseline autoencoder model is used to identify known and unknown activities and an unsupervised machine learning model is used to cluster unknown activities.

4.2 Date Exchange Protocol & Frequency

In each communication round, the server residing at the cloud sends up-to-date global autoencoder to a set of devices. The devices run the autoencoder during a training period with the unlabeled data. When the autoencoder detects activities with an accuracy below a threshold, this triggers the use of an unsupervised machine learning method to cluster the new activity. At the end of the communication round, the devices send the gradients of the resulting model to the server. The server uses the Federated Averaging Algorithm (FedAvg) to calculate the resulting model which will be sent in the next round. If the unsupervised model identifies a cluster which continues over time a new label is generated. The server also updates the global classifier using the new label and send the resulting classifier to the selected devices with the global autoencoder.

The number of scheduled devices is limited by the available bandwidth that can be allocated to clients at each round. Also, a communication round cannot be considered valid unless a minimum number of updates is collected from devices. The frequency of the communication and aggregation can be reduced with additional local new activity learning rounds. Therefore, for an optimal learning round duration, bandwidth resource use should be adapted based on the number of local iterations of unknown activities at each device, and the number of global learning rounds.

5 Experiment, Data Collection and Annotation

In this study, two (2) datasets were used: one a publicly available Human Activity Recognition (HAR) dataset and a second one which was captured from devices in the Chatty Factories project.

5.0.1 Experimental Set up

A 6th generation Apple iPad was selected as the “Chatty device” for the experiment. Modern mobile iOS and Android devices are equipped with a variety of sensors, providing a number of potential mechanisms to collect data for research. iOS devices typically utilise a wide range of inbuilt sensors eg. proximity sensors, accelerometers, gyroscope, etc. [11]. Using this “Chatty device” a dataset was created while performing common product use activities, defined as: walking, sitting, standing with the device, dropping, picking up the device, placing the device stationary on a surface and a vibrating surface. The data was collected from the following four in-built sensors.

- Acceleration: 3-axes acceleration data
- Orientation: Azimuth, Pitch, and Roll
- Angular Velocity: 3-axes gyroscopes for rotational motion
- Magnetic Field: 3-axes magnetic field

5.0.2 Data Acquisition

-

Chatty Dataset

The MATLAB Mobile application [1] was used to stream data from the in-built sensors of the Chatty-device. A sampling rate of 100Hz was selected for logging the data. The product use Data was streamed into a MathWorks Cloud for analysis. Two subjects (a male and female) were selected to carry out each activity for a continuous period of two minutes each. A total of four minutes of sensor data readings were collected for every activity, producing a total of 148,778 data points. The steps highlighted below were followed:

1. The MATLAB Mobile was installed and configured on the Chatty device through the Apple App store.
2. A registered user account was used and the Chatty device was connected to MathWorks Cloud.
3. Using the MATLAB Mobile interface, the sensors were turned on.
4. A sampling rate of 100Hz was selected and the log button was used to initiate recording of the sensor data readings.
5. After the expected duration, the logged results were saved to the cloud and given file names which corresponded to the activities carried out.

Human Activity Recognition Dataset

The publicly available Human Activity Recognition (HAR) data set [46] covers acceleration, GPS, gyroscope, light, magnetic field, and sound level data of the activities. Where the activities that were performed by fifteen participants (age 31.9 ± 12.4 , height 173.1 ± 6.9 , weight 74.1 ± 13.8 , eight males and seven females) were, climbing stairs down and up, jumping, lying, standing, sitting, running/jogging, and walking. For each activity, a log file was generated that simultaneously recorded the acceleration of the body based on chest, forearm, head, shin, thigh, upper arm, and waist positions during each activity. Each activity was performed for a period of 10 minutes except for jumping due to physical exertion (1.7 minutes).

5.0.3 Data Pre-Processing

All generated product use data were stored locally on the device and automatically uploaded to the Mathworks cloud account. The files were then downloaded and imported into MATLAB desktop application and processed. Each sensor reading was timestamped, pre-processed and then labelled for each activity. In the end a master labelled data log file was generated that contained all sensor readings.

A summary of the activities from the datasets used in this work is given in Table 1 The Chatty dataset was used to train and test the model. The model was further tested against an unseen publicly available dataset. For training of the model, out of the six activities defined in Chatty dataset, four were used. While two activities, namely: Standing and “Drop and Pick up” were used to evaluate the new activity detecting capability of the system.

Table 1 Summary of Datasets

Activities	Chatty	Public HAR
Known	99001	61606
Walking	25267	30734
Sitting	24563	30871
Vibration	24558	-
Stationary on Table	24613	-
Unknown	49776	26264
Standing	25214	-
Drop and Pickup	24562	-
Jumping	-	4325
Climbing down stairs	-	21939

5.1 Autoencoders

Unlabelled data is captured at the edge to identify known product use activities and flag any unknown activity. To accomplish this problem we propose the use of the Autoencoder method for new activity detection. Autoencoder is a deep learning architecture that has the capability to identify outliers in a dataset [41],[19],[44]. Thus it does not need labelled datasets which contain both known and unknown product use activities for training purposes. Considering the data will be collected in real time, it will not be free from noise, a stacked denoising autoencoder was chosen to detect product use activities. A stacked denoising autoencoder, shown in Fig. 3, is a neural network that learns a reconstruction function that produces a denoised vector from a corrupted, noisy input [5]. It is a feed-forward non-recurrent net with the input layer connected to the output layer through several hidden layers which contract from the input to the latent code at the bottleneck layer and then expand to the output layer [2].

The basic autoencoder learns to find a parameter vector θ which minimizes a reconstruction function of the form

$$J_{AE}(\theta) = \sum_t L(x^{(t)}, g_{\theta}(f_{\theta}(x^{(t)}))) \quad (1)$$

where $x^{(t)}$ represents a training sample of the known product use activities from the Chatty dataset. However, the stacked denoising autoencoder adds a corruption term $E_{q(\tilde{x}|x^{(t)})}[\cdot]$ such that the reconstruction function becomes

$$J_{AE}(\theta) = \sum_t E_{q(\tilde{x}|x^{(t)})}[L(x^{(t)}, g_{\theta}(f_{\theta}(x^{(t)})))] \quad (2)$$

where the notation $E_{q(\tilde{x}|x^{(t)})}[\cdot] = \sum_{x^{(t)}} q(\tilde{x}|x^{(t)}) \times [\cdot]$

The corruption term so introduced reduces the sensitivity of the network to noise in the input and thus increases the ability of the network to generalize to previously unseen input data.

Experiment settings and Hyperparameters: a stacked denoising Autoencoder was implemented in Python 3.8.5 with Keras 2.4.0 and TensorFlow

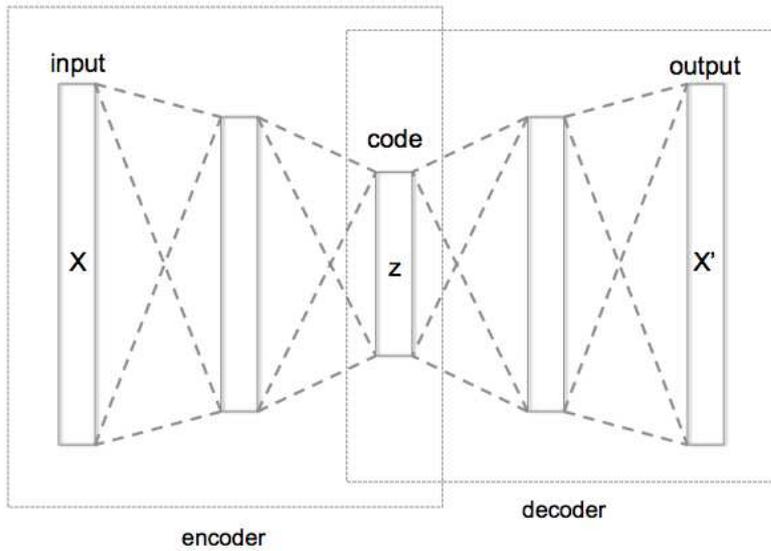


Fig. 3 Stacked Autoencoder Architecture

2.3.0-rc0 with GPU support. The experiment was carried out on a Macbook Pro with six (6) core Intel Core i7 processor running at 2.6GHz, 16GB RAM and a Radeon Pro 555X GPU. Table 2 summarizes the Autoencoder hyperparameter configuration. Linear activation was used on the output and hidden layers to preserve correspondence between the input and output domains since no normalization was used on the input data. For optimization, the Adam optimizer was used since it achieves faster convergence and better performance compared to RMSProp, Adadelta and vanilla SGD [43], [24]. The number of training epochs was fixed to 64 since the training loss flattens out considerably at about 55 epochs. To combat overfitting, an L1_L2 activity regularizer was used on the bottleneck layer to introduce sparsity on that layer. A 0.6 standard deviation Gaussian Noise layer was also used after the input layer to reduce sensitivity of the model to noise in the input data and enhance the models ability to generalize to previously unseen data samples. The noise layer was only active during training, it was inactive at inference time.

5.2 Unsupervised Machine Learning

While the focus of this paper is not to explore various unsupervised methods for clustering, it plays a crucial role in helping to distinguish between an actual product activity from random noise. If the clustering method repeatedly

Table 2 Autoencoder Hyperparameters

Hyperparameters	Description	Values
N_l	number of hidden layers	3
N_h	number of units on each hidden layer	5, 3, 5
N_{epoch}	number of training epochs	64
α	learning rate	0.1

identifies a specific cluster of data points, then we infer that those data points belong to an actual product activity and not just random noise that may occur infrequently. To this end, it is critically important to determine an optimal number of clusters that can be fed to the clustering algorithm given that the product use activities are being generated during actual use, and as such is unknown.

Various methods can be used to find the optimal number of clusters. The elbow method is one such methods [25] and is selected to achieve this task in this study. In this algorithm, the centre of clusters is determined by using K-means clustering algorithm (to execute the task of clustering), and the *Within-Cluster-Sum-of-Squares* (WCSS) is estimated for each identified cluster. The WCSS is referred to as the sum of squares of the distances for every data point in all clusters to their corresponding centroids [40],[10].

To demonstrate the clustering of data points in this study, *K-means* algorithm was used. It is characterized by its distance-based technique which executes in an iterative manner.

Let $X = \{x_i\}$, $i = 1, 2, \dots, n$. be the n objects to be clustered, $S = S_1, S_2, \dots, S_k$ is the set of clusters. Let μ be the mean of cluster S_i . The squared-error between μ_i and the objects in cluster S_i is defined as $WCSS(S_i) = \sum_{X_j \in S_i} \|X_j - \mu_i\|^2$. The K-means algorithm aims to minimize the sum of the squared error over all k clusters, that is $min(WCSS(S)) = argmin_s \sum_{i=1}^k \sum_{X_j \in S_i} \|X_j - \mu_i\|^2$, where WCSS denotes the sum of the squared error in the inner-cluster. A good point to note is that one drawback to the K-means algorithm is that it requires the specification of the number of "k" (clusters) beforehand and also the inappropriateness for discovering non-convex clusters [29]

6 Results and Analysis

Once the auto-encoder was configured and deployed, the Chatty dataset was divided into train and test datasets using a 75%-25% split for the train and test datasets respectively. For the train dataset only known product user activities were included. However, the dataset that was used to test the model contained both known and unknown product use activities.

Table 3 RECONSTRUCTION LOSS DISTRIBUTION

	Train (Chatty)	Known Activity (Chatty)	Unknown Activity (Chatty)	Known Activity (Public) Sitting	Unknown Activity (Public) Climb down
Mean	0.129162	0.129993	0.462168	1.054878	1.222696
Std	0.173400	0.174832	0.511599	0.027329	0.385470
Min	0.012407	0.012829	0.018665	0.768975	0.331072
25%	0.028761	0.028780	0.068573	1.047389	0.975942
50%	0.032796	0.033920	0.385282	1.052245	1.123178
75%	0.151830	0.152089	0.773422	1.058461	1.359943
Max	1.363831	1.375441	11.554137	2.316129	3.493574

Though Autoencoder have a deep learning architectures they differ from neural network based on what they produce as an output. The target output of the autoencoder is not classes but the input itself, primarily to learn the representation of the data or to remember the pattern of known product use activities. During training the autoencoder only known product activities were used to build the model. The rationale behind this was that the model was able to recognise only known product activities and was capable to successfully able to reconstruct the error.

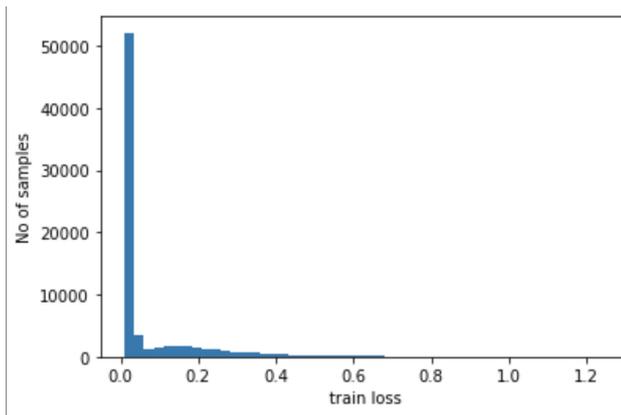


Fig. 4 Training Reconstruction Loss (Chatty Device)

While training the model on known product use activities the model was able to reconstruct the train & test set data with an accuracy of 99.35% and mean reconstruction error of 0.129 (see table 3 and figures 4, 5). Similarly, the model was tested against known product use activities (walking & sitting) from a publicly available dataset and the performance was consistent with results obtained with the Chatty dataset. The model recorded a performance of 98.03% with a mean of 1.05 and max of 2.3 for the reconstruction error (see table 3 and figures 6).

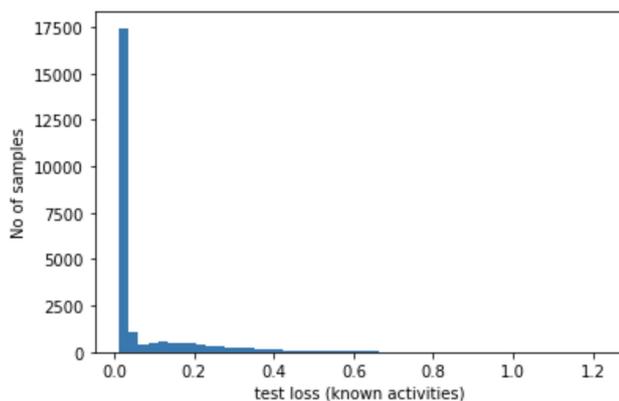


Fig. 5 Test Loss Known Activity (Chatty Device)

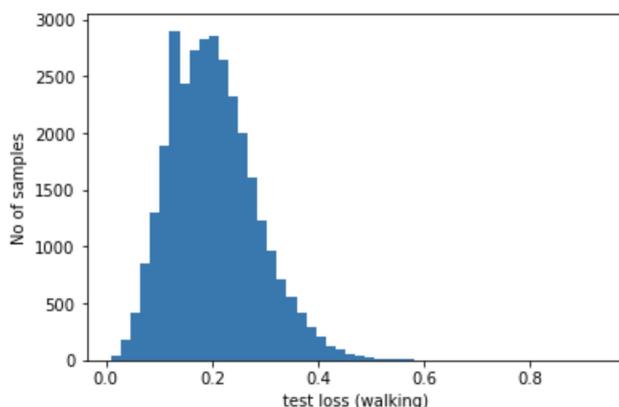


Fig. 6 Known Activity Reconstruction Loss (Public Dataset)

However, when the model was tested on unknown Chatty product use activities (Standing and Drop & Pickup) the accuracy dropped to 89.51% and the maximum reconstruction error increased by around 743%. This indicated that the model was unable to recognise the data, thus an increase in the reconstruction error was observed. Similarly, we checked the model against unknown product use activities (jumping and climbing down stairs) from a publicly available data set. Again we observed a drop in performance to 83.59% and a rise in reconstruction error (see figure 8, 9). In figure 7 we can see how the reconstruction error changes as the model encounters unknown product use activity resulting in a drop of overall performance.

Table 4 presents the performance of the model on known and unknown product use activities from the chatty dataset and the publicly available HAR dataset. The results demonstrate that the model is able to generalize to previously unseen datasets of known activities with an accuracy of 98.06%. As

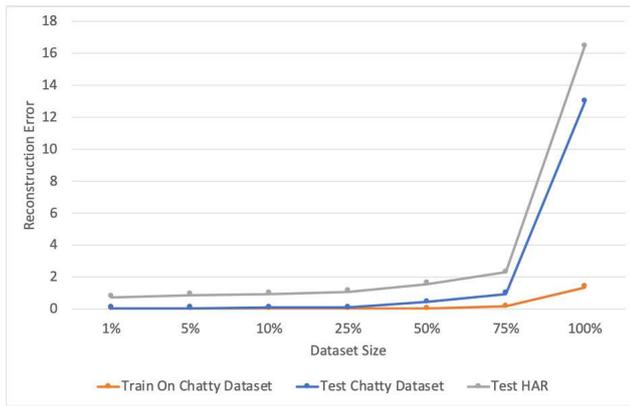


Fig. 7 Reconstruction Error observed while training and testing the model

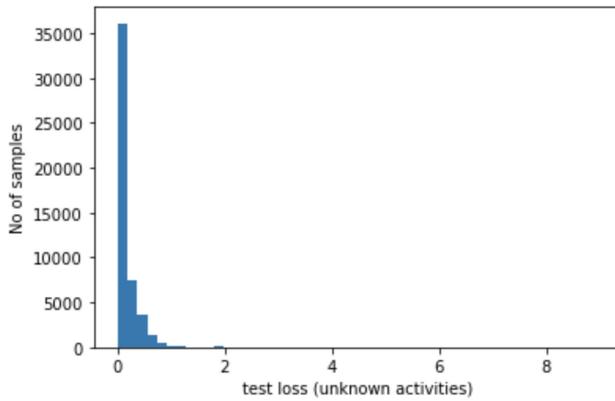


Fig. 8 Unknown Activity Reconstruction Loss (Chatty Device)

Table 4 Anomaly Detection Model Performance

Dataset	Accuracy
Chatty Train set (Known)	99.35%
Chatty Test set (Known)	99.35%
Chatty Test set (Unknown)	89.51%
Public Test set (Known)	98.03%
Public Test set (Unknown)	83.59%

expected, the accuracy drops significantly for unknown activities from the held out test data from the Chatty device and the publicly available HAR dataset.

In order to create a trigger for the model to flag unknown product use activities, a threshold value was to be decided. From the observed results, an accuracy of 93.7% was selected as the threshold for minimum acceptable accuracy for known activities. This threshold is 1.5 standard deviations lower than the minimum observed accuracy for known activities across the chatty

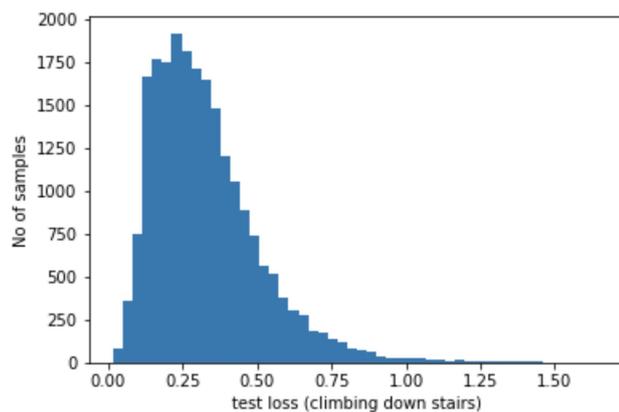


Fig. 9 Unknown Activity Reconstruction Loss (Public Dataset)

datasets and the publicly available HAR dataset. The selected threshold allows the model to maximally discriminate between known and unknown activities while preserving the ability of the model to generalize to previously unseen datasets.

The data that had been categorised as unknown product use activity, was further processed to check if it was actually an activity data or it was just noise. In order to check that the data was passed through unsupervised machine learning model to cluster similar data points together. While building the unsupervised clustering model, Elbow method was used to determine the number of clusters and subsequently, the corresponding graph is plotted with respect to the cluster numbers and WCSS (as reflected in Figure 10).

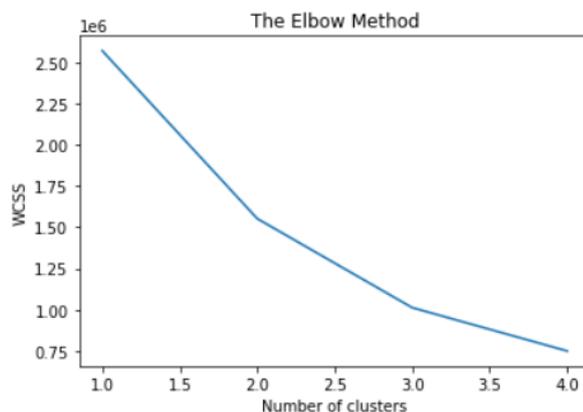


Fig. 10 Elbow method to select number of Clusters

The optimal number of clusters can be understood as the point at which adding to the number of clusters yields no further improvement to the clustering index considerably. Figure 10 demonstrates clusters for the unknown product usage activities discovered by the autoencoder. From Figure 10, we can infer that the optimal number of clusters selected is 2.

Acceleration 3D-Plot for Clustered Unknown Activities

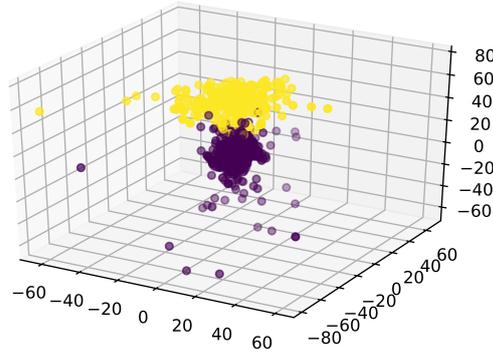


Fig. 11 Acceleration 3D Plot for Clustered for Unknown Product Use Activity

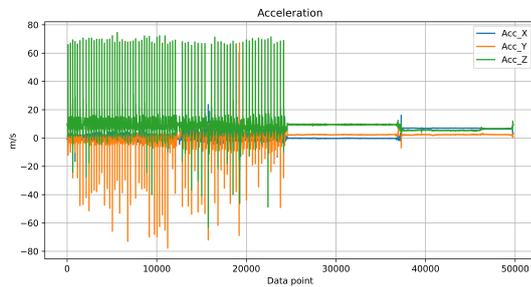


Fig. 12 Acceleration Line Plot for Clustered for Unknown Product Use Activity

Figure 11 presents the empirical results from the K-means cluster analysis that was carried out on the unknown activity data points. Two clear cluster points were formed in the visualisation as already indicated from the results from the WCSS. Figure 12 illustrates a line plot of the triaxial accelerometer signals for the unknown activities. Thus, the data flagged as unknown product use data corresponds to user activity and not to noise.

7 Conclusion

A novel system that can distinguish between known and unknown product use activities, while providing model correction using federated learning is described. The proposed system is designed in such a way that user data is processed at the edge and any correction needed to the model based on other user's data are undertaken on a cloud environment. Thus ensuring that no user related data has to be transferred in proximity to the device, and maintaining user privacy at every stage. To identify between known and unknown product use activities a stacked denoising autoencoder was chosen because of its ability to reduce sensitivity to noise in the input dataset and thus increase the ability to generalise to previously unseen input dataset. The autoencoder was trained on dataset collected from devices in the Chatty Factories project. It was tested on part of unseen data collected from *chatty devices* (i.e. devices that have in-built sensors and where such data can be used to associate particular use of the device from this data) and from a publicly available dataset, the Human Activity Recognition (HAR) dataset. The model was successful in identifying known product use activities with an accuracy of 99.35% while training the model and that of 99.35% (Chatty dataset) & 98.03% (HAR dataset) while testing the model. In order to identify unknown product use activities a performance threshold of 97.5% was defined and a drop in performance below this level would indicate that the model had encountered unknown product use activities. Based on this threshold the model was successful in identifying unknown product use activities in both Chatty and HAR dataset. This led to the creation of a new label representing unknown product use activities, which the manufacture could use to understand how the product was *actually* being used, and how the data could be used to improve product design. The methodology and approach outlined in this work can lead to a new generation of privacy-aware data-driven adaptive manufacturing process.

Acknowledgements This research was funded by the Engineering and Physical Sciences Research Council (EPSRC) - project EP/R021031/1 - Chatty Factories.

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. Makers of matlab and simulink. URL <https://uk.mathworks.com/>
2. Autoencoder with recurrent neural networks for video forgery detection. *Electronic Imaging* **2017**(7) (2017)
3. Abbasi, Q.H., Rehman, M.U., Qaraqe, K., Alomainy, A.: *Advances in Body-Centric Wireless Communication: Applications and State-of-the-Art*, 1st edn. Institution of Engineering and Technology (2016)

4. Alhimale, L., Zedan, H., Al-Bayatti, A.: The implementation of an intelligent and video-based fall detection system using a neural network. *Applied Soft Computing Journal* **18**, 59–69 (2014). DOI 10.1016/j.asoc.2014.01.024
5. Bengio, Y., Courville, A., Vincent, P.: Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**(8), 1798–1828 (2013). DOI 10.1109/TPAMI.2013.50
6. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H.B., Van Overveldt, T., Petrou, D., Ramage, D., Roselander, J.: Towards Federated Learning at Scale: System Design (2019). URL <http://arxiv.org/abs/1902.01046>
7. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K.: Practical secure aggregation for privacy-preserving machine learning. In: *Proceedings of the ACM Conference on Computer and Communications Security*, pp. 1175–1191. Association for Computing Machinery, New York, USA (2017). DOI 10.1145/3133956.3133982. URL <http://dl.acm.org/citation.cfm?doi=3133956.3133982>
8. Burnap, P., Branson, D., Murray-Rust, D., Preston, J., Richards, D., Burnett, D., Edwards, N., Firth, R., Gorkovenko, K., Khanesar, M., Lakoju, M., Smith, T., Thorp, J.: Chatty factories: a vision for the future of product design and manufacture with IoT. In: *Living in the Internet of Things (IoT 2019)*, pp. 4 (6 pp.)–4 (6 pp.). Institution of Engineering and Technology (2019). DOI 10.1049/cp.2019.0129. URL <https://digital-library.theiet.org/content/conferences/10.1049/cp.2019.0129>
9. Casilari, E., Santoyo-Ramón, J.A., Cano-García, J.M.: Analysis of Public Datasets for Wearable Fall Detection Systems. *Sensors* **17**(7), 1513 (2017). DOI 10.3390/s17071513. URL <http://www.mdpi.com/1424-8220/17/7/1513>
10. Chakrabarty, N., Rana, S., Chowdhury, S., Maitra, R.: Rbm based joke recommendation system and joke reader segmentation. In: *International Conference on Pattern Recognition and Machine Intelligence*, pp. 229–239. Springer (2019)
11. Costello, S.: The Sensors That Make the iPhone So Cool (2019). URL <https://www.lifewire.com/sensors-that-make-iphone-so-cool-2000370>
12. Daher, M., Diab, A., El Badaoui El Najjar, M., Ali Khalil, M., Charpillet, F.: Elder Tracking and Fall Detection System Using Smart Tiles. *IEEE Sensors Journal* **17**(2), 469–479 (2017). DOI 10.1109/JSEN.2016.2625099
13. Dean, J., Corrado, G.S., Monga, R., Chen, K., Devin, M., Le, Q.V., Mao, M.Z., Ranzato, A., Senior, A., Tucker, P., Yang, K., Ng, A.Y.: Large Scale Distributed Deep Networks. Tech. rep. (2012)
14. Feng, G., Mai, J., Ban, Z., Guo, X., Wang, G.: Floor Pressure Imaging for Fall Detection with Fiber-Optic Sensors. *IEEE Pervasive Computing* **15**(2), 40–47 (2016). DOI 10.1109/MPRV.2016.27
15. Ghobakhloo, M.: Industry 4.0, digitization, and opportunities for sustainability (2020). DOI 10.1016/j.jclepro.2019.119869
16. Gibson, R.M., Amira, A., Ramzan, N., Casaseca-De-La-Higuera, P., Pervez, Z.: Multiple comparator classifier framework for accelerometer-based fall detection and diagnostic. *Applied Soft Computing Journal* **39**, 94–103 (2016). DOI 10.1016/j.asoc.2015.10.062
17. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour (2017). URL <http://arxiv.org/abs/1706.02677>
18. Guidoux, R., Duclos, M., Fleury, G., Lacomme, P., Lamaudière, N., Manenq, P.H., Paris, L., Ren, L., Rousset, S.: A smartphone-driven methodology for estimating physical activities and energy expenditure in free living conditions. *Journal of Biomedical Informatics* **52**, 271–278 (2014). DOI 10.1016/j.jbi.2014.07.009
19. Hao, Y., Sheng, Y., Wang, J.: Variant gated recurrent units with encoders to preprocess packets for payload-aware intrusion detection. *IEEE Access* **7**, 49985–49998 (2019)
20. Hui, S., Zhongmin, W.: Compressed sensing method for human activity recognition using tri-axis accelerometer on mobile phone. *Journal of China Universities of Posts and Telecommunications* **24**(2), 31,71–37 (2017). DOI 10.1016/S1005-8885(17)60196-1
21. Igual, R., Medrano, C., Plaza, I.: Challenges, issues and trends in fall detection systems (2013). DOI 10.1186/1475-925X-12-66. URL <http://biomedical-engineering-online.biomedcentral.com/articles/10.1186/1475-925X-12-66>

22. Jayasinghe, U., Harwin, W.S., Hwang, F.: Comparing Clothing-Mounted Sensors with Wearable Sensors for Movement Analysis and Activity Classification. *Sensors* **20**(1), 82 (2019). DOI 10.3390/s20010082. URL <https://www.mdpi.com/1424-8220/20/1/82>
23. Kerdjadj, O., Ramzan, N., Ghanem, K., Amira, A., Chouireb, F.: Fall detection and human activity classification using wearable sensors and compressed sensing. *Journal of Ambient Intelligence and Humanized Computing* **11**(1), 349–361 (2020). DOI 10.1007/s12652-019-01214-4
24. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)
25. Kodinariya, T.M., Makwana, P.R.: Review on determining number of cluster in k-means clustering. *International Journal* **1**(6), 90–95 (2013)
26. Konečný, J., McMahan, H.B., Ramage, D., Richtárik, P.: Federated Optimization: Distributed Machine Learning for On-Device Intelligence (2016). URL <http://arxiv.org/abs/1610.02527>
27. Lee, R.Y.W., Carlisle, A.J.: Detection of falls using accelerometers and mobile phone technology. *Age and Ageing* **40**(6), 690–696 (2011). DOI 10.1093/ageing/afr050. URL <https://doi.org/10.1093/ageing/afr050>
28. Li, Q., Stankovic, J.A., Hanson, M.A., Barth, A.T., Lach, J., Zhou, G.: Accurate, fast fall detection using gyroscopes and accelerometer-derived posture information. In: *Proceedings - 2009 6th International Workshop on Wearable and Implantable Body Sensor Networks, BSN 2009*, pp. 138–143 (2009). DOI 10.1109/BSN.2009.46
29. Likas, A., Vlassis, N., J. Verbeek, J.: The global k-means clustering algorithm. *Pattern Recognition* **36**(2), 451–461 (2003). DOI 10.1016/S0031-3203(02)00060-2
30. Makhoulf, A., Boudouane, I., Saadia, N., Ramdane Cherif, A.: Ambient assistance service for fall and heart problem detection. *Journal of Ambient Intelligence and Humanized Computing* **10**(4), 1527–1546 (2019). DOI 10.1007/s12652-018-0724-4
31. McMahan, B., Ramage, D.: Federated Learning: Collaborative Machine Learning without Centralized Training Data (2017). URL <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>
32. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., Arcas, B.A.y.: Communication-Efficient Learning of Deep Networks from Decentralized Data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017* (2016). URL <http://arxiv.org/abs/1602.05629>
33. McMahan, H.B., Ramage, D., Talwar, K., Zhang, L.: Learning Differentially Private Recurrent Language Models. *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings* (2018). URL <http://arxiv.org/abs/1710.06963>
34. Micucci, D., Mobilio, M., Napolitano, P., Tisato, F.: Falls as anomalies? An experimental evaluation using smartphone accelerometer data. *Journal of Ambient Intelligence and Humanized Computing* **8**(1), 87–99 (2017). DOI 10.1007/s12652-015-0337-0
35. Moussa, M.M., Hamayed, E., Fayek, M.B., El Nemr, H.A.: An enhanced method for human action recognition. *Journal of Advanced Research* **6**(2), 163–169 (2015). DOI 10.1016/j.jare.2013.11.007
36. Mubashir, M., Shao, L., Seed, L.: A survey on fall detection: Principles and approaches. *Neurocomputing* **100**, 144–152 (2013). DOI 10.1016/j.neucom.2011.09.037
37. Muscio, A., Ciffolilli, A.: What drives the capacity to integrate Industry 4.0 technologies? Evidence from European R&D projects. *Economics of Innovation and New Technology* **29**(2), 169–183 (2020). DOI 10.1080/10438599.2019.1597413. URL <https://www.tandfonline.com/doi/full/10.1080/10438599.2019.1597413>
38. Ozcan, K., Velipasalar, S., Varshney, P.K.: Autonomous Fall Detection with Wearable Cameras by Using Relative Entropy Distance Measure. *IEEE Transactions on Human-Machine Systems* **47**(1), 31–39 (2017). DOI 10.1109/THMS.2016.2620904
39. Oztemel, E., Gursev, S.: Literature review of Industry 4.0 and related technologies (2020). DOI 10.1007/s10845-018-1433-8
40. Pollard, D.: Strong consistency of k-means clustering. *The Annals of Statistics* pp. 135–140 (1981)
41. Pratomo, B.A., Burnap, P., Theodorakopoulos, G.: Unsupervised approach for detecting low rate attacks on network traffic with autoencoder. In: *2018 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, pp. 1–8 (2018). DOI 10.1109/CyberSecPODS.2018.8560678

42. Preece, S.J., Goulermas, J.Y., Kenney, L.P.J., Howard, D., Meijer, K., Crompton, R.: Activity identification using body-mounted sensors a review of classification techniques. *Physiological Measurement* **30**(4), R1—R33 (2009). DOI 10.1088/0967-3334/30/4/r01. URL <https://doi.org/10.1088/0967-3334/30/4/r01>
43. Ruder, S.: An overview of gradient descent optimization algorithms. CoRR **abs/1609.04747** (2016). URL <http://arxiv.org/abs/1609.04747>
44. Schneider, P., Böttinger, K.: High-performance unsupervised anomaly detection for cyber-physical system networks. In: *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy*, pp. 1–12 (2018)
45. Smith, S.L., Kindermans, P.J., Ying, C., V.Le, Q.: Don’t decay the learning rate, increase the batch size. In: *International Conference on Learning Representations (ICLR)*, 2018 (2018)
46. Sztyler, T., Stuckenschmidt, H.: On-body localization of wearable devices: An investigation of position-aware activity recognition. In: *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pp. 1–9. IEEE Computer Society (2016). DOI 10.1109/PERCOM.2016.7456521. [Http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7456521](http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7456521)
47. Tabia, H., Gouiffes, M., Lacassagne, L.: Motion histogram quantification for human action recognition. In: *21st International Conference on Pattern Recognition (ICPR 2012)*, *Icpr*, pp. 2404–2407 (2012)
48. Uddin, M.Z.: A wearable sensor-based activity prediction system to facilitate edge computing in smart healthcare system. *Journal of Parallel and Distributed Computing* **123**, 46–53 (2019). DOI 10.1016/j.jpdc.2018.08.010
49. Wu, Q., He, K., Chen, X.: Personalized Federated Learning for Intelligent IoT Applications: A Cloud-Edge based Framework (2020). URL <http://arxiv.org/abs/2002.10671>

Author, Book title, page numbers. Publisher, place (year)

Figures

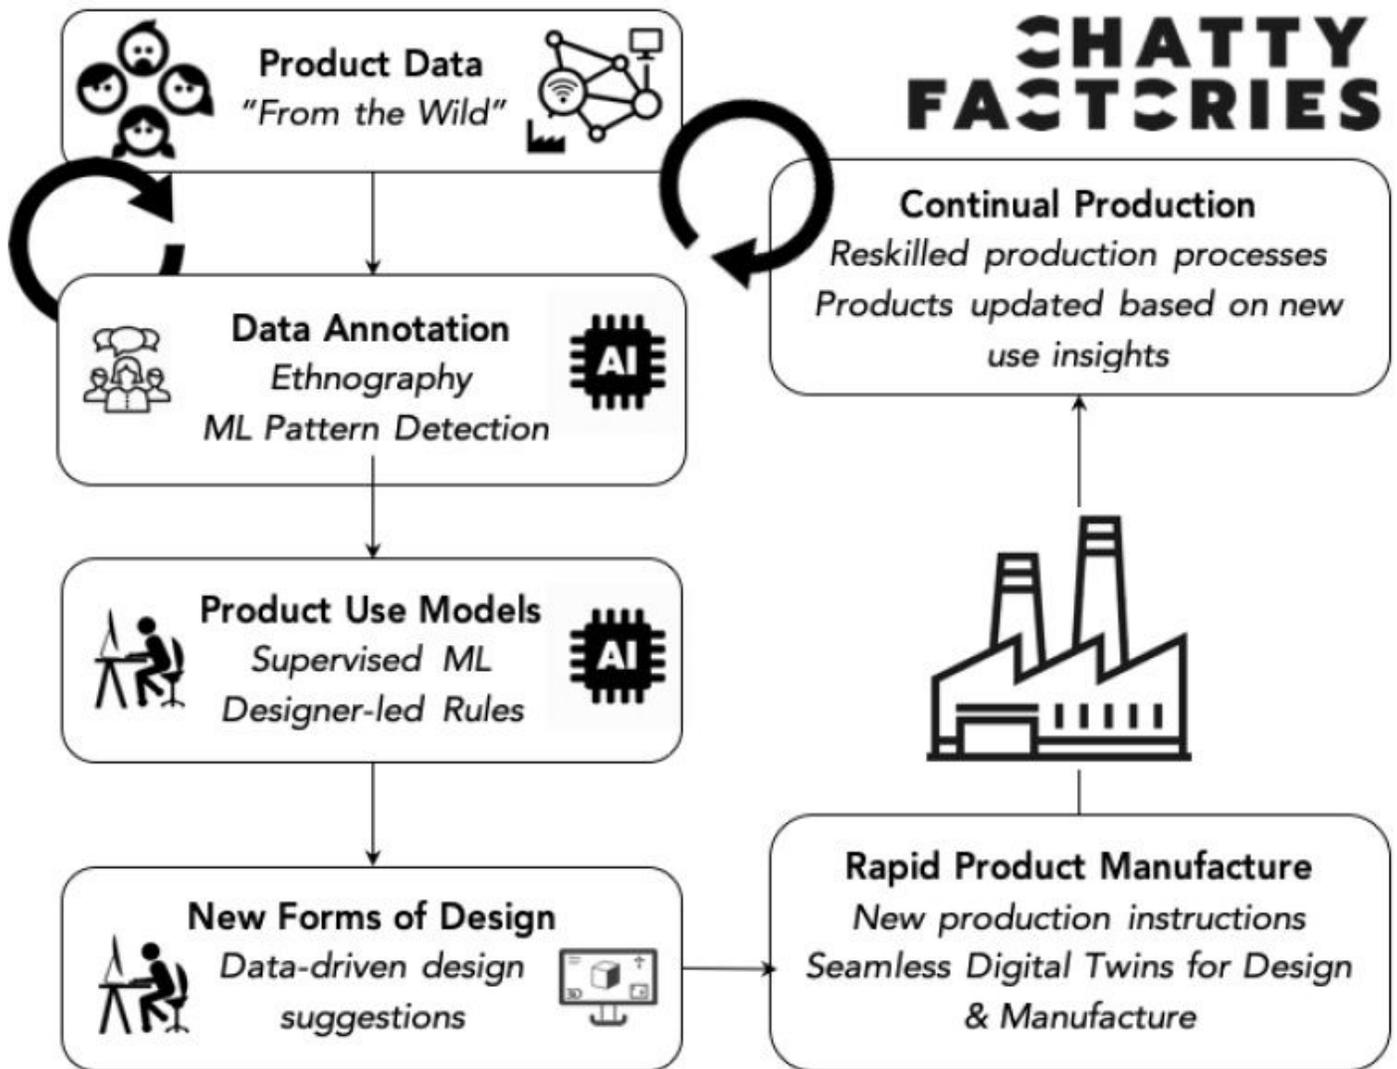


Figure 1

The Chatty Factories Vision

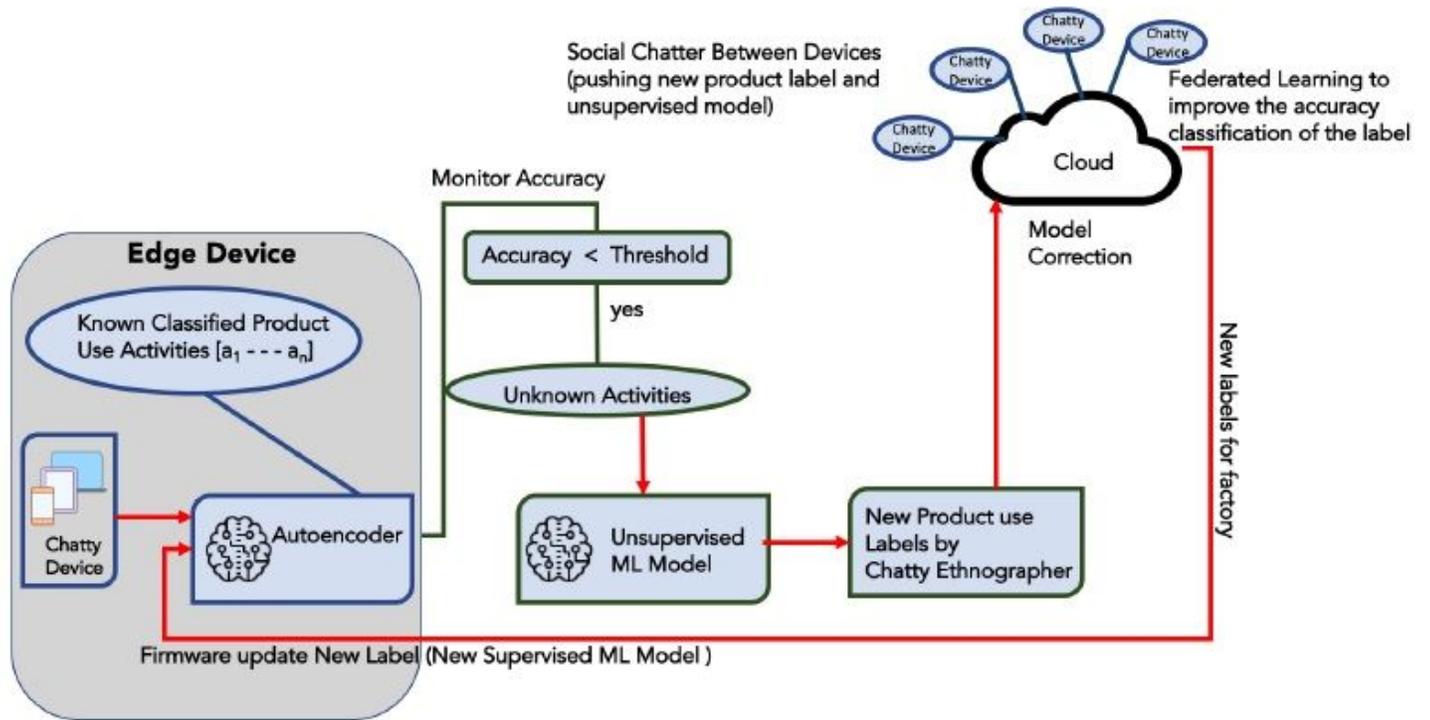


Figure 2

Flow of work

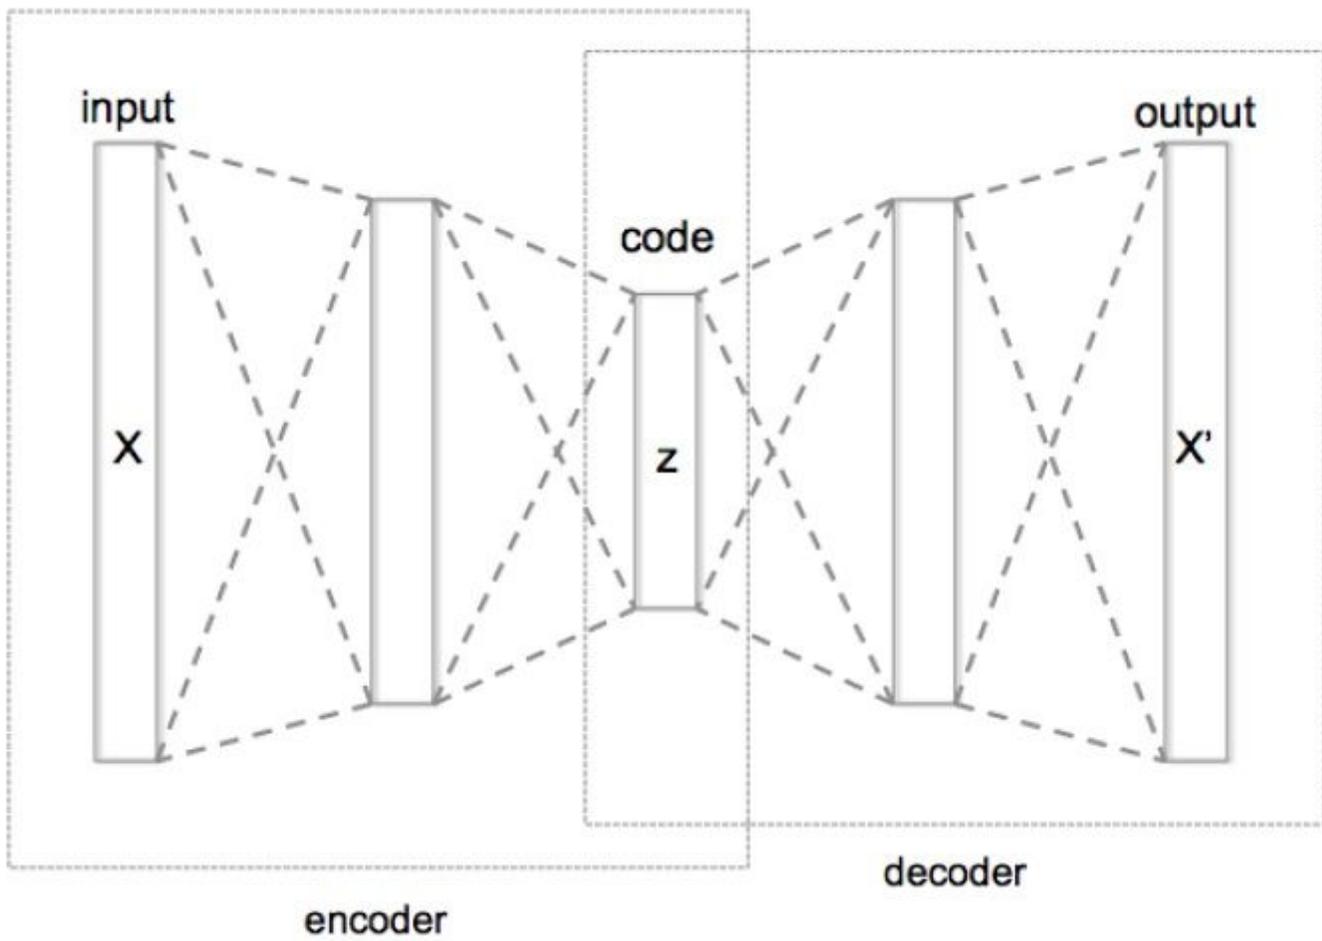


Figure 3

Stacked Autoencoder Architecture

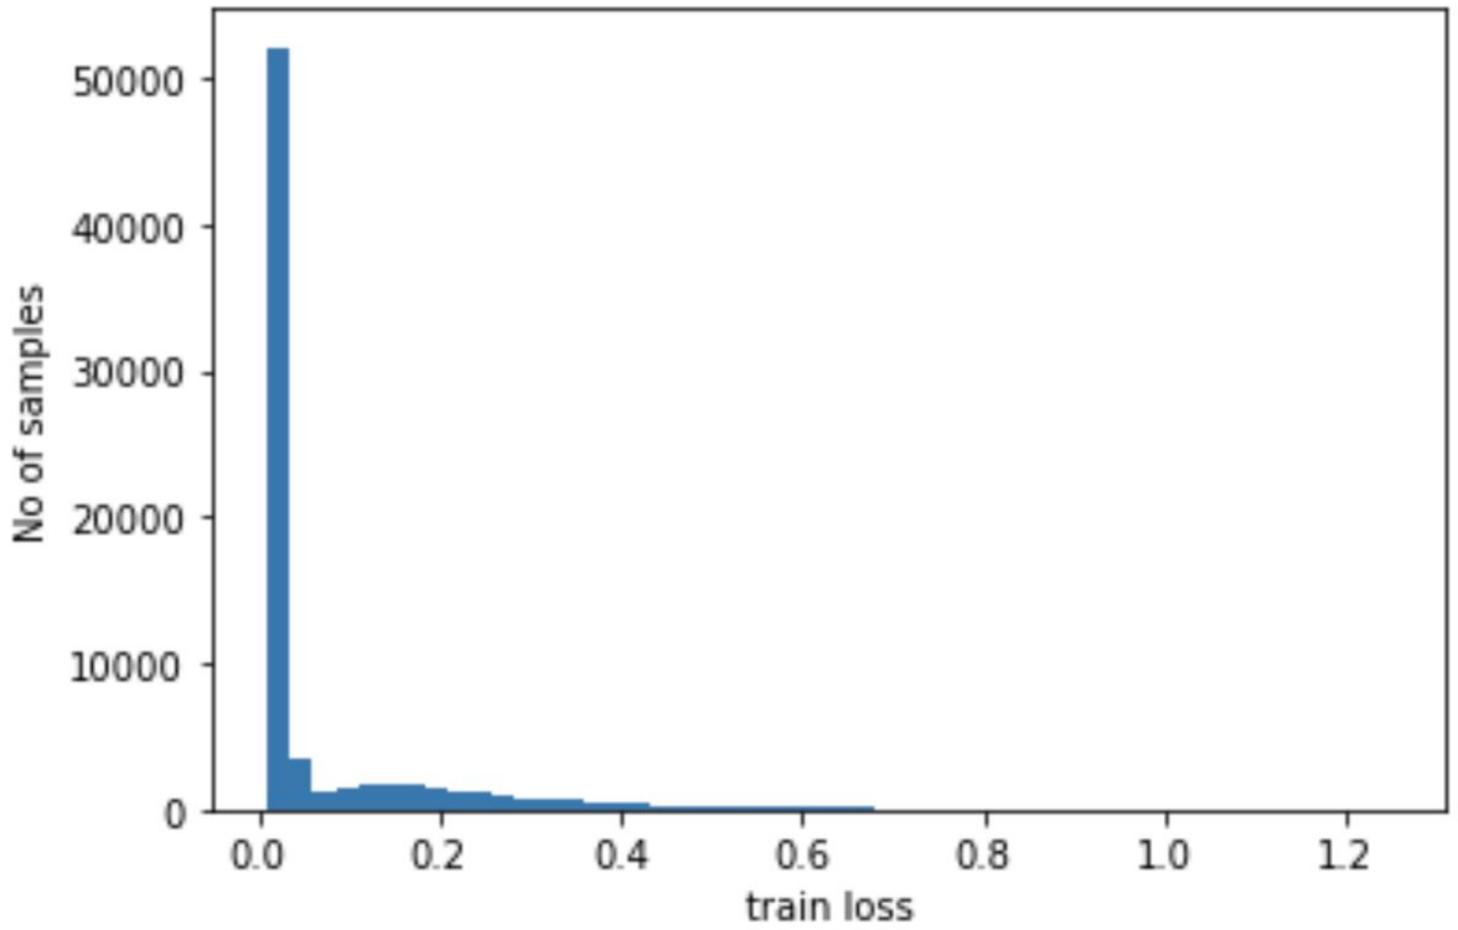


Figure 4

Training Reconstruction Loss (Chatty Device)

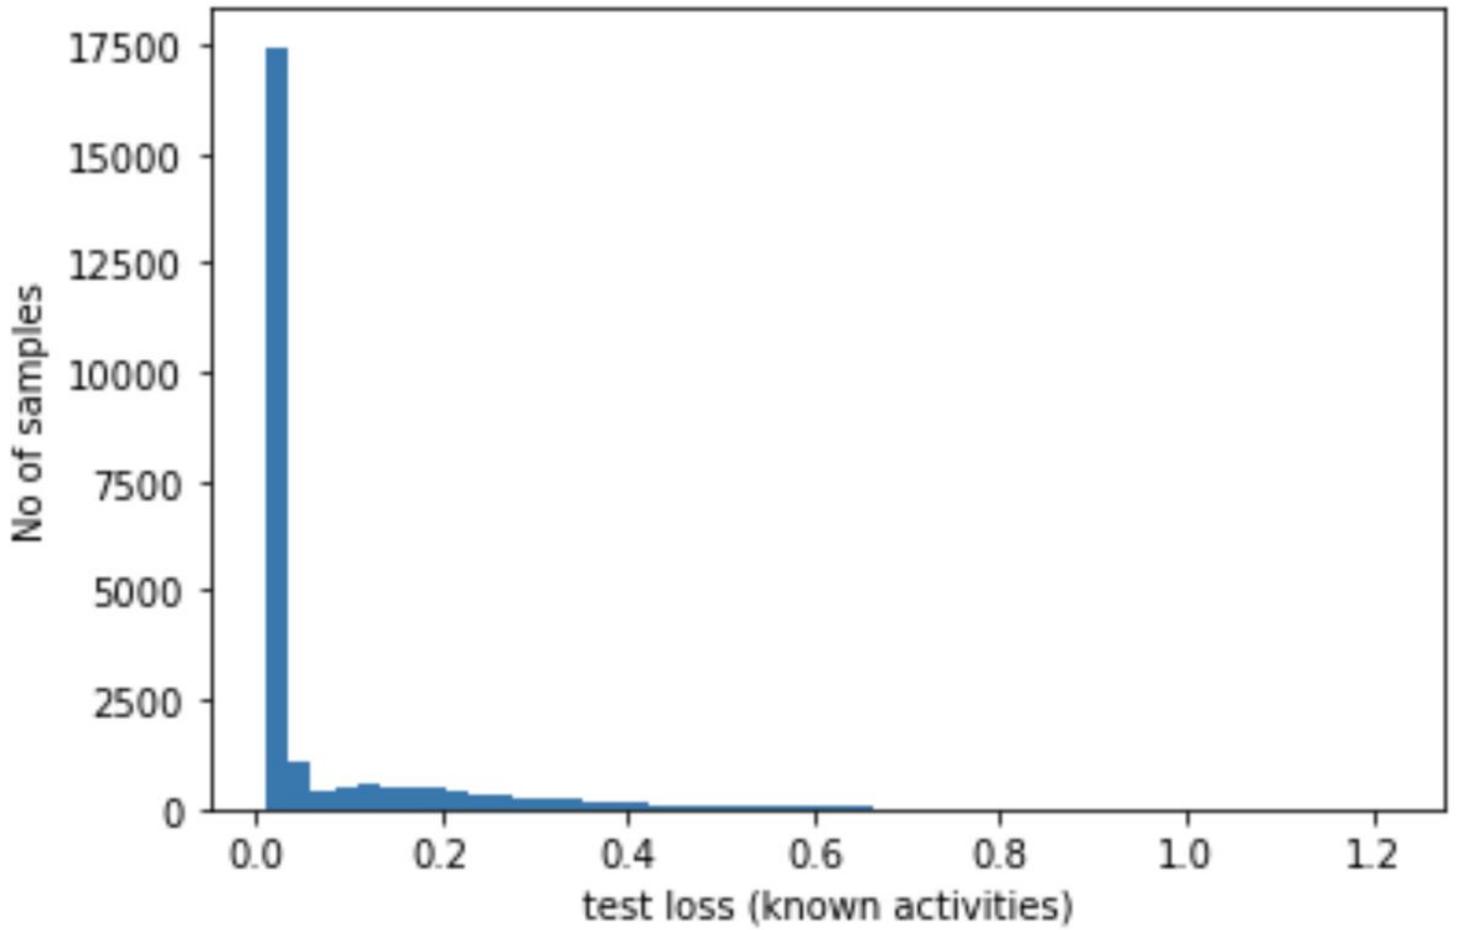


Figure 5

Test Loss Known Activity (Chatty Device)

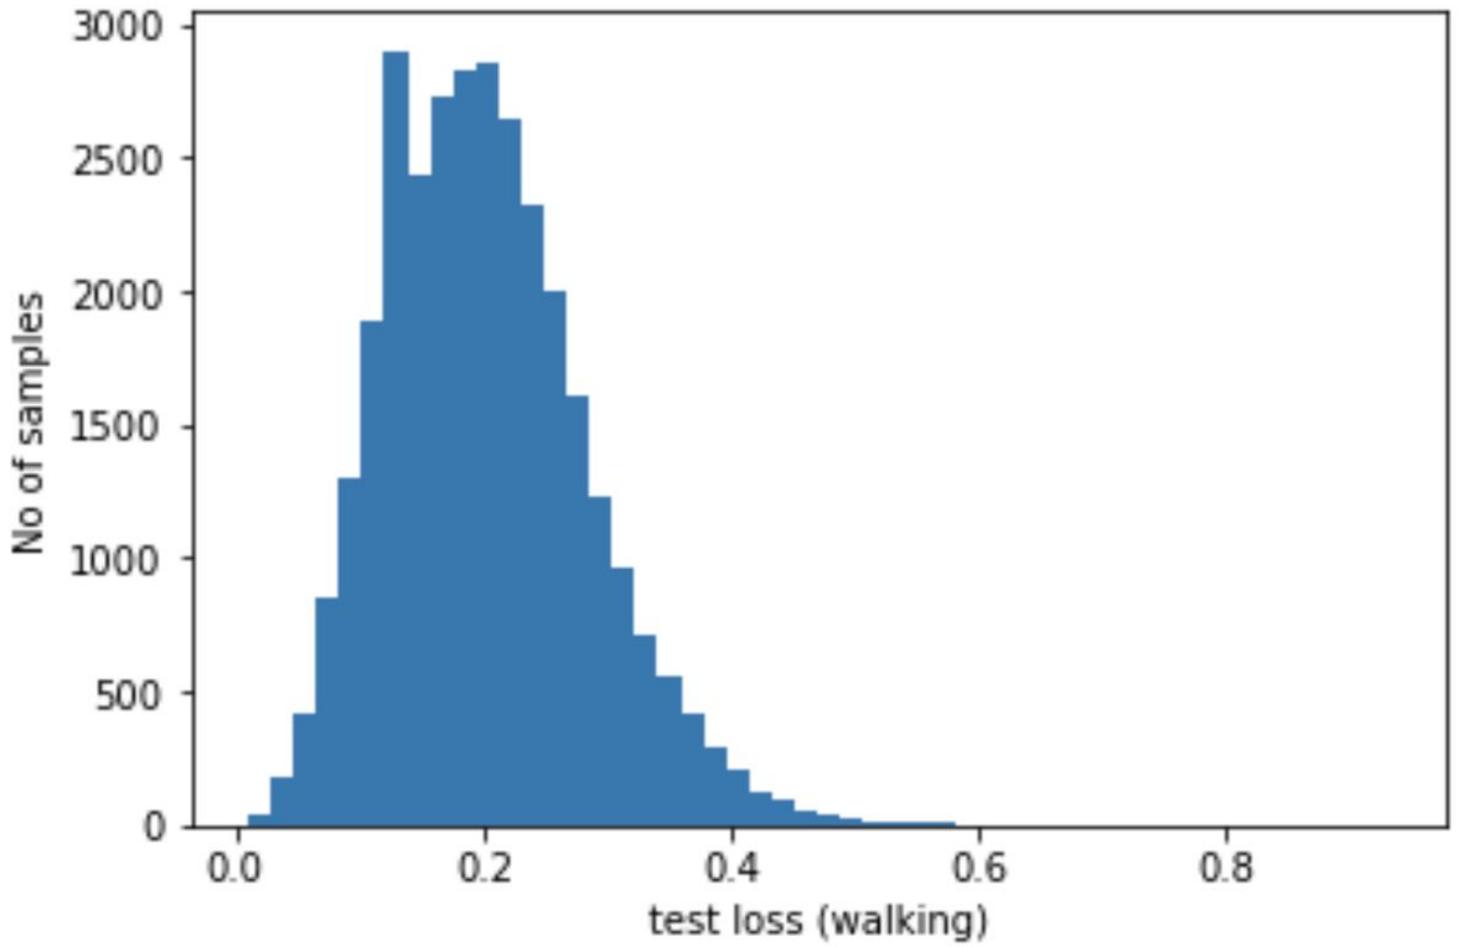


Figure 6

Known Activity Reconstruction Loss (Public Dataset)

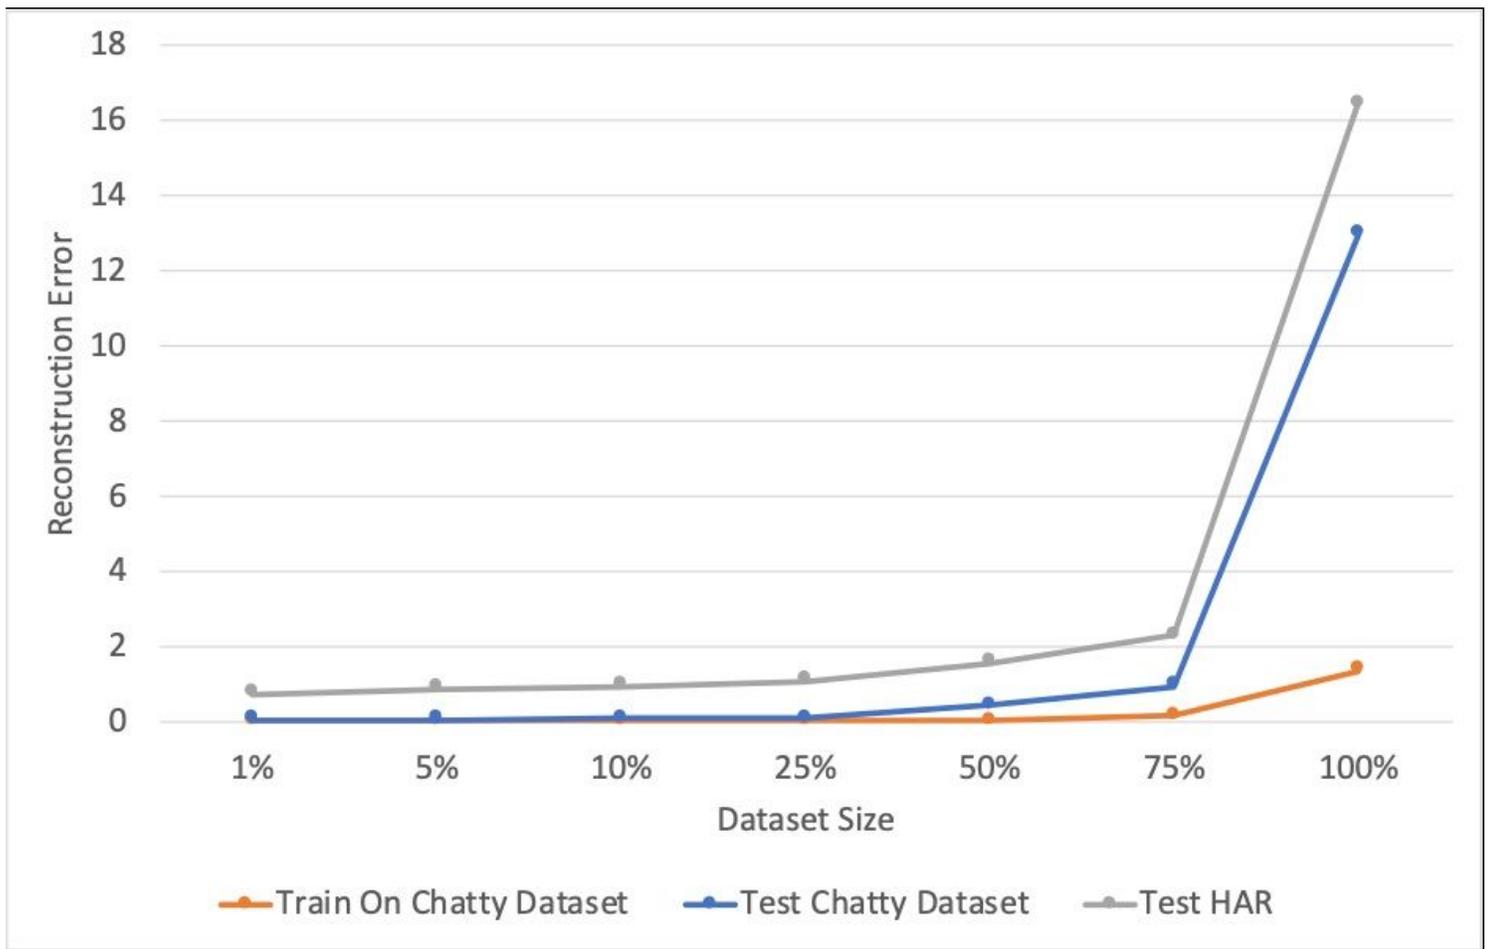


Figure 7

Reconstruction Error observed while training and testing the model

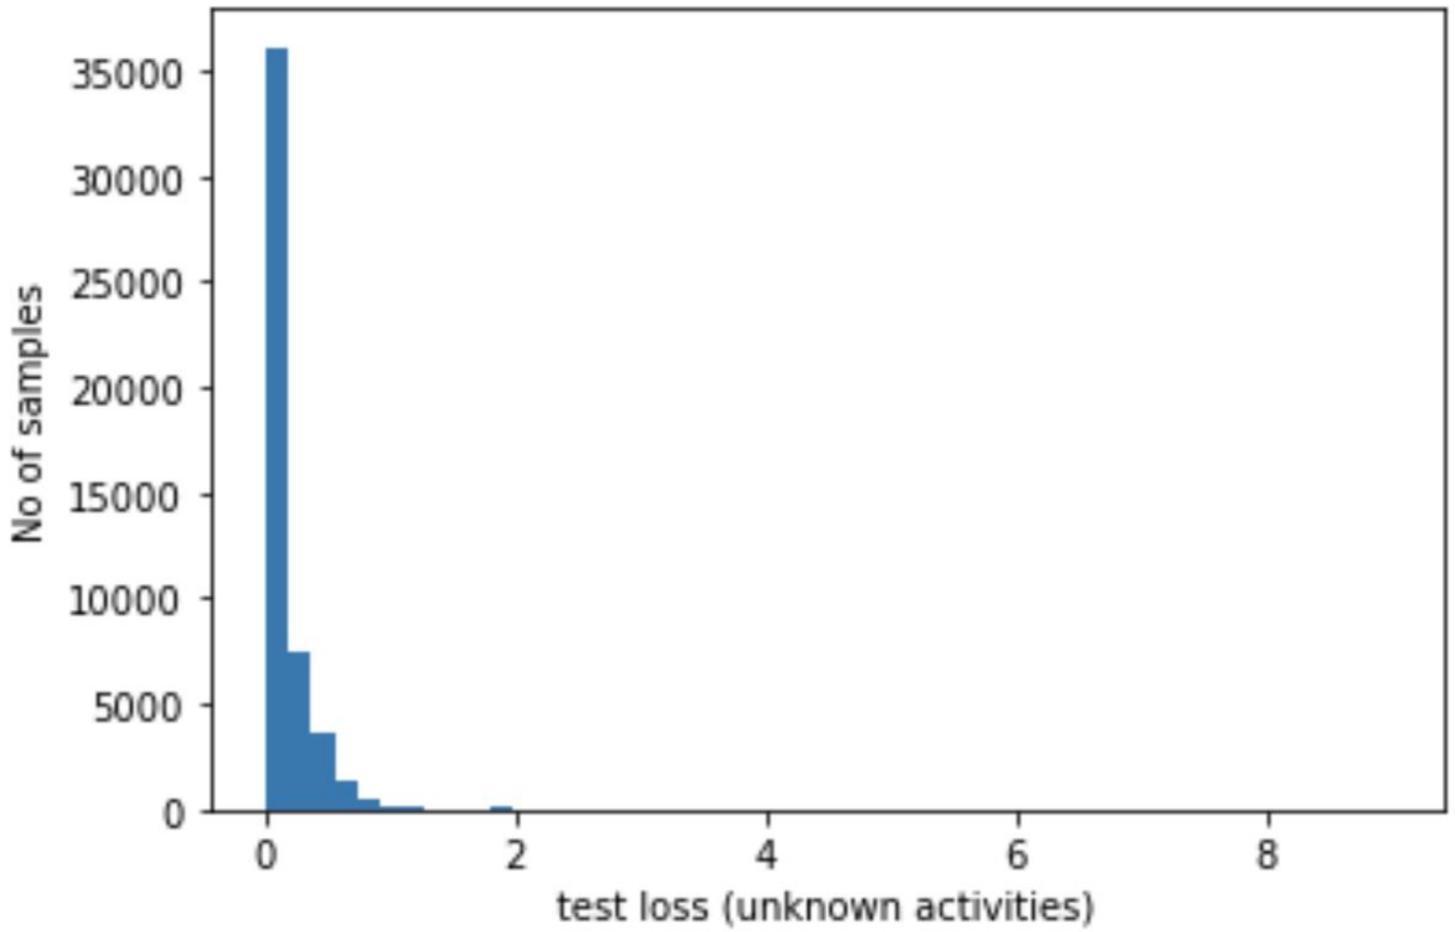


Figure 8

Unknown Activity Reconstruction Loss (Chatty Device)

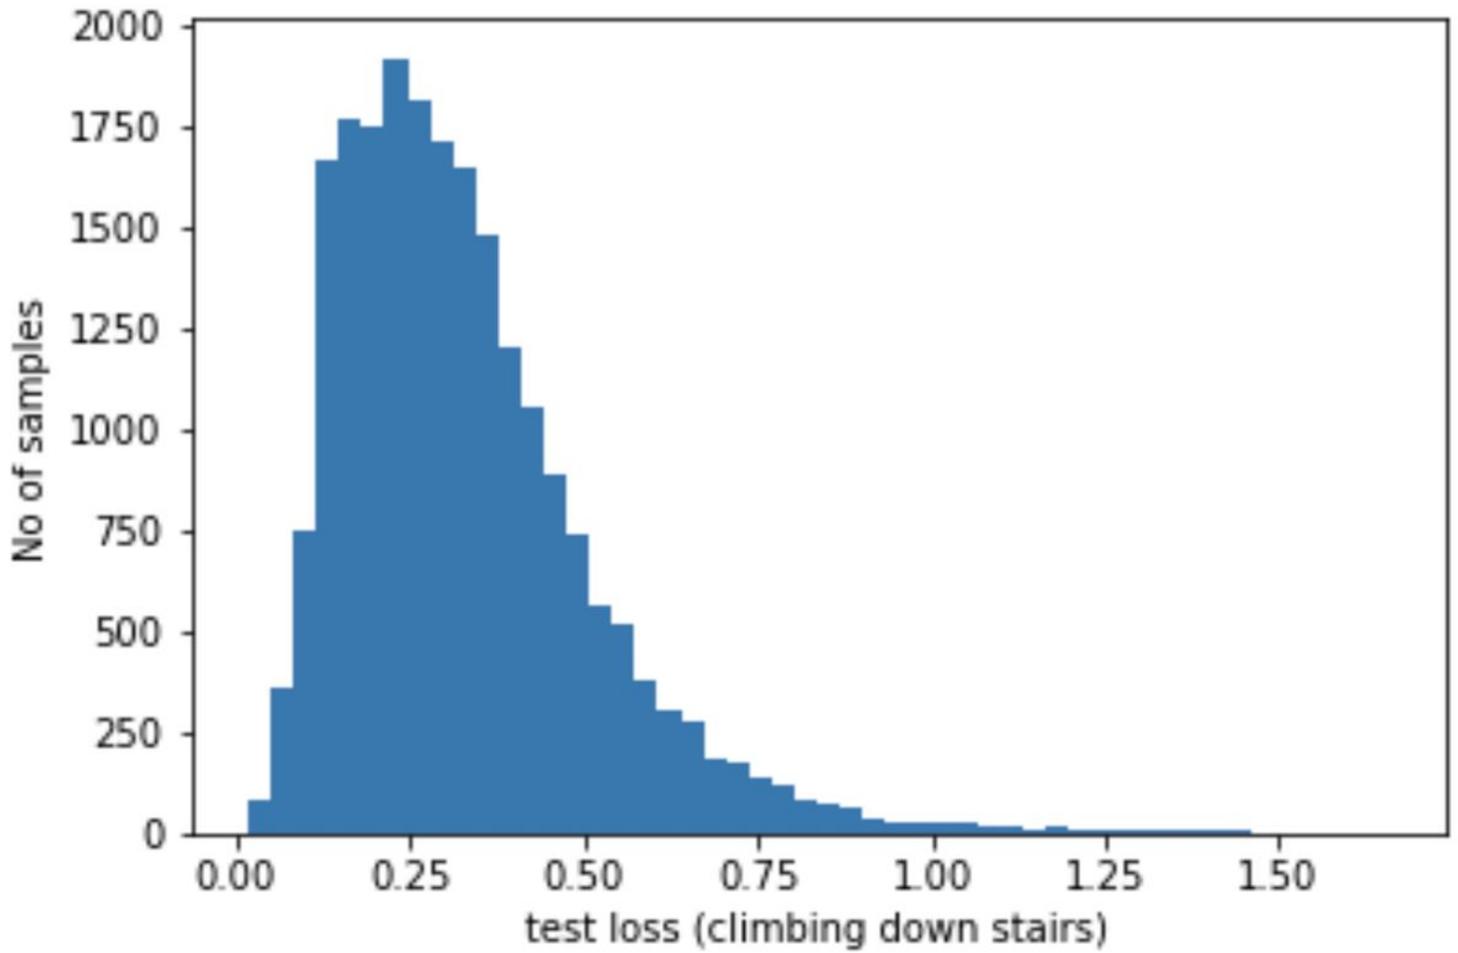


Figure 9

Unknown Activity Reconstruction Loss (Public Dataset)

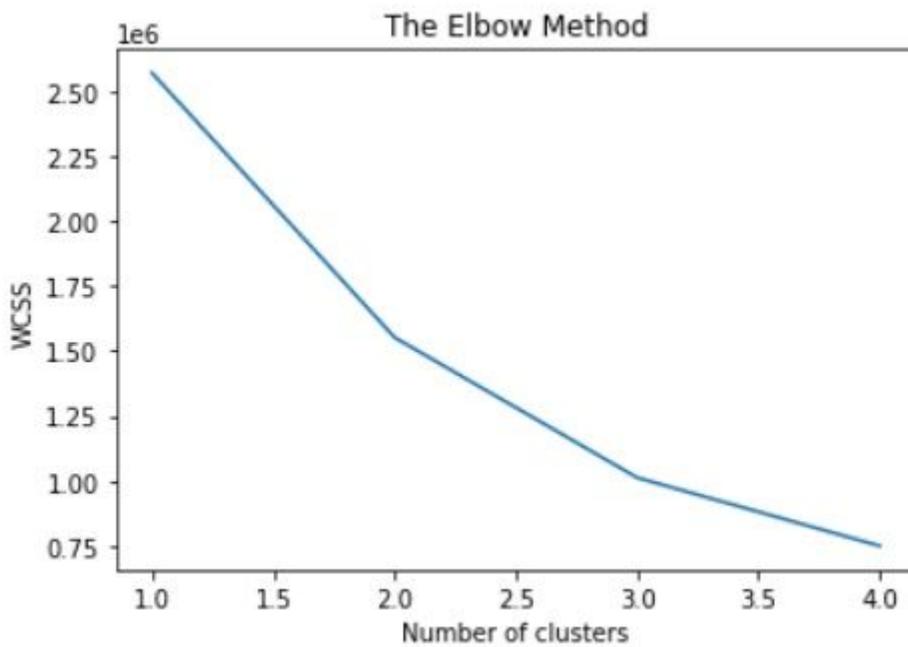


Figure 10

Elbow method to select number of Clusters

Acceleration 3D-Plot for Clustered Unknown Activities

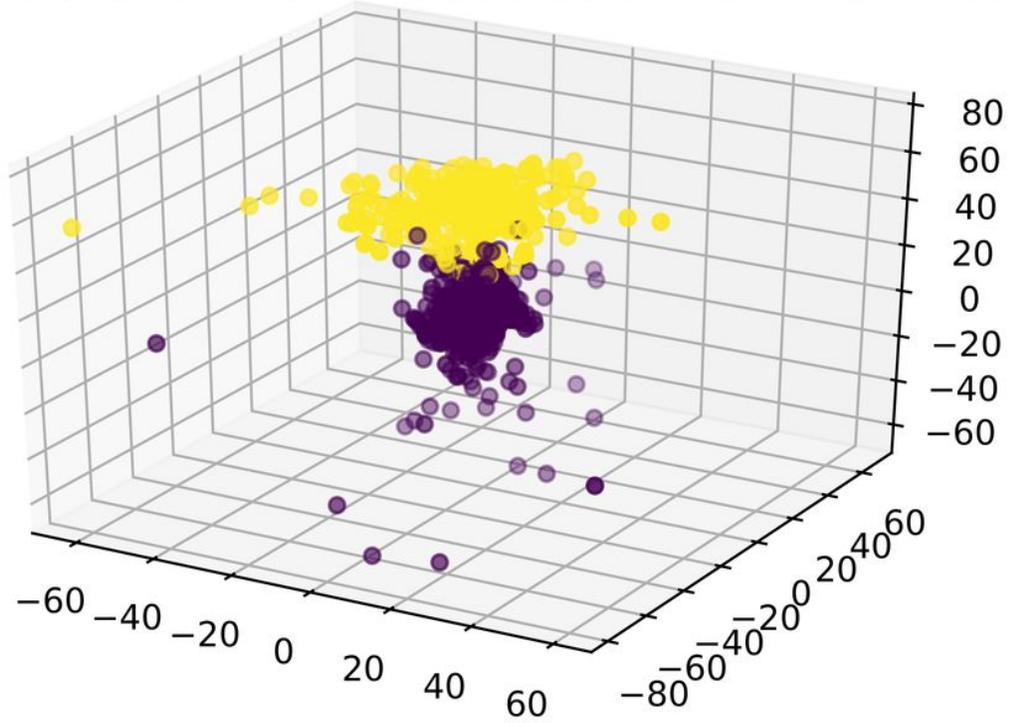


Figure 11

Acceleration 3D Plot for Clustered for Unknown Product Use Activity

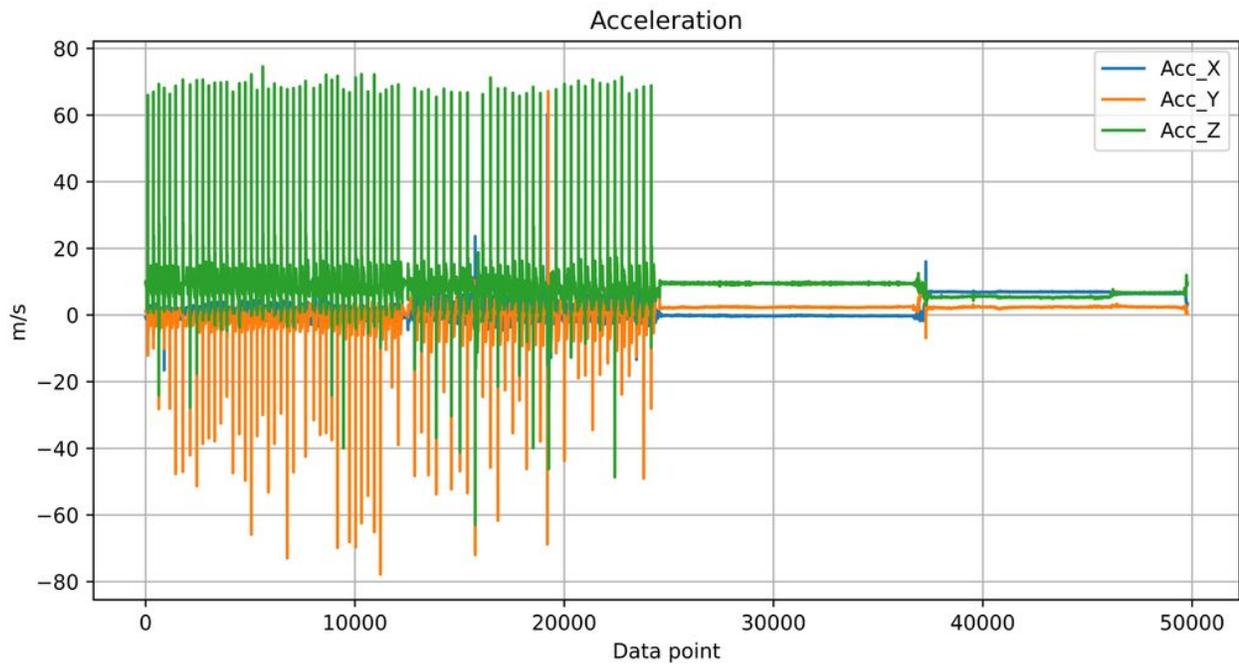


Figure 12

Acceleration Line Plot for Clustered for Unknown Product Use Activity