

mmFall: Fall Detection using 4D MmWave Radar and a Hybrid Variational RNN AutoEncoder

Feng Jin, *Student Member, IEEE*, Arindam Sengupta, *Student Member, IEEE*, and Siyang Cao, *Senior Member, IEEE*

Abstract—Elderly fall prevention and detection becomes extremely crucial with the fast aging population globally. In this paper, we propose *mmFall* - a novel fall detection system, which comprises of (i) the emerging millimeter-wave (mmWave) radar sensor to collect the human body’s point cloud along with the body centroid, and (ii) a Hybrid Variational RNN AutoEncoder (HVRAE) to compute the anomaly level of the body motion based on the acquired point cloud. A fall is detected when the spike in anomaly level and the drop in centroid height occur simultaneously. The mmWave radar sensor offers privacy-compliance and high sensitivity to motion, over the traditional sensing modalities. However, (i) randomness in radar point cloud and (ii) difficulties in fall collection/labeling in the traditional supervised fall detection approaches are the two major challenges. To overcome the randomness in radar data, the proposed HVRAE uses variational inference, a generative approach rather than a discriminative approach, to infer the posterior probability of the body’s latent motion state every frame, followed by a recurrent neural network (RNN) to summarize the temporal features over multiple frames. Moreover, to circumvent the difficulties in fall data collection/labeling, the HVRAE is built upon an autoencoder architecture in a semi-supervised approach, which is only trained on the normal activities of daily living (ADL). In the inference stage, the HVRAE will generate a spike in the anomaly level once an abnormal motion, such as fall, occurs. During the experiment¹, we implemented the HVRAE along with two other baselines, and tested on the dataset collected in an apartment. The receiver operating characteristic (ROC) curve indicates that our proposed model outperforms baselines and achieves 98% detection out of 50 falls at the expense of just 2 false alarms.

Note to Practitioners—Traditional non-wearable fall detection approaches typically make use of a vision-based sensor, such as camera, to monitor and detect fall using a classifier that is trained in a supervised fashion on the collected fall and non-fall data. However, several problems render these methods impractical. Firstly, camera-based monitoring may trigger privacy concerns. Secondly, fall data collection using human subjects is difficult and costly, not to mention the impossible ask of the elderly repeating simulated falls for data collection. In this paper, we propose a new fall detection approach to overcome these problems by (i) using a palm-size mmWave radar sensor to monitor the elderly, that is highly sensitive to motion while protecting privacy; and (ii) using a semi-supervised anomaly detection approach to circumvent the fall data collection. Further hardware engineering and more training data from people with different body figures could make the proposed fall detection solution even more practical.

Index Terms—Fall detection, millimeter wave radar, variational autoencoder, recurrent autoencoder, semi-supervised learning, anomaly detection.

I. INTRODUCTION

GLOBALLY, the elderly aged 65 or over make up the fastest-growing age group [1]. Approximately 28-35% of the elderly fall every year [2], making it the second leading unintentional injury death after road traffic injuries [3]. Moreover, elderly falls are cost intensive with the total 2015 direct cost of fall among the elderly, adjusted for inflation, being 31.9 billion USD in the United States alone [4]. Therefore, researchers seek to detect the fall right after it occurs, along with an immediate alert trigger so a timely treatment can be implemented [5]. Based on the choice of sensor, fall-related research can be broadly divided into wearables, non-wearables and fusion domains [5].

In this paper, we are focusing on non-wearable fall detection using the emerging millimeter-wave (mmWave) radar sensor [6]. In short, mmWave radar represent moving objects in a scene as a point cloud in which each point contains the 3D position in space and a 1-D Doppler (radial velocity component) information, thereby resulting in a 4D mmWave radar as referred to in the paper title.

MmWave radar sensor can offer several advantages over the other traditional sensing technologies, viz. (i) non-intrusive and convenience over the wearable solutions [7]–[9] that also need frequent battery recharging; (ii) privacy-compliance over camera [10]; (iii) high-sensitivity to motion and operationally robust to occlusions, when compared to depth sensors [11], especially in a complex living environment; (iv) more informative than typical ambient sensors [12]–[14] which suffer interference from the external environment [15]; and (v) low-cost, compact and high resolution over the traditional radar counterparts [16].

The World Health Organization (WHO) [2] defines fall as “*inadvertently coming to rest on the ground, floor or other lower level, excluding intentional change in position to rest in furniture, wall or other objects.*” Therefore, we propose the *mmFall*, in which a generative recurrent autoencoder measures the motion inadvertence or anomaly level based on the mmWave radar point cloud of the body, and the drop of centroid height, which is estimated from the point cloud, indicates the motion of coming to rest on a lower level. Moreover, such a semi-supervised approach can circumvent the difficulties of real-world elderly fall data collection.

The rest of this paper is organized as follows. Section II discusses related radar-based fall detection research and

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Feng Jin, Arindam Sengupta and Siyang Cao are with Department of Electrical and Computer Engineering, *the University of Arizona*, Tucson, AZ 85719, USA. (e-mail: {fengjin, sengupta, caos}@email.arizona.edu)

¹All the codes and datasets are shared on GitHub (<https://github.com/radar-lab/mmfall>).

semi-supervised learning approaches. Section III introduces all the components that constitute our proposed *mmFall* system, including the principles of mmWave radar sensor, variational inference, variational autoencoder, and recurrent autoencoder. Section IV presents the overall system architecture, a novel data oversampling method and a custom loss function for model training. Section V shows the experimental evaluation of *mmFall*, compares the performance with two baseline architectures, and discusses the limitations of current research and future work. Finally, Section VI concludes the paper.

II. RELATED WORK

Traditionally in radar-based fall detection research, researchers mainly focus on extracting the micro-Doppler [17] features, i.e., Doppler distribution over time, and then train a classifier that can distinguish fall from non-fall data [18]–[23]. However, micro-Doppler features have no spatial information (range and angle), presence of multiple people, other motion sources, and similar motions as fall (such as sitting), can lead to inaccuracies. Jokanović [24] fused information from both the micro-Doppler and range domains to reduce the false alarm rate by training a logistic regression classifier. Similar research can be found in [25]–[27]. Tian *et al.* [28] used a 3D convolutional neural network (CNN) to learn different activities by exploiting the range-angle heatmap in both azimuth and elevation over time. However, there are two problems in incorporating the spatial information, viz. i) achieving high angular resolution using low-frequency radars requires a bulky antenna (see Fig. 1 of a 4 GHz radar in [18]); ii) and a low signal bandwidth also limits the range resolution.

The use of high-bandwidth mmWave radar looks very promising to overcome these limitations than its traditional counterparts, and is an emerging trend. In our previous research [29]–[31] we adopted such a palm-size mmWave radar sensor that first segregates multiple people based on the spatial information, and then uses a CNN to classify each person’s activity, including fall, based on the Doppler pattern separately, and even reconstruct their skeletal pose. Sun *et al.* [32] also used a mmWave radar and long short-term memory (LSTM) to detect fall based on the range-angle heatmap over time. The advantage of using mmWave radars is furthered if we can take advantage of all the information available from it, such as range, azimuth angle, elevation angle, and Doppler.

On the other hand, a vast majority of these radar-based fall detection research adopts a supervised approach. Researchers manually label the collected fall and non-fall data, manually or automatically extract features over time, and then train a classifier that can distinguish fall from non-fall data. The challenge with these supervised approaches is that the rare and non-continuous fall event is very difficult to collect, not to mention the impossible ask of the elderly repeating falls for data collection. Furthermore, the manual extraction and labelling of short portions of fall event from

the long duration data is very expensive, time-consuming and inefficient.

To overcome these problems, we leverage the semi-supervised anomaly detection (SSAD) approach. Anomaly detection refers to the problem of finding patterns in data that do not conform to expected behavior [33]. In our case, we can use SSAD to train a model only on the normal activities of daily living (ADL), such as walking/sitting/crouching, etc., such that the model will ‘recognize’ the normal ADL, while a fall event will ‘surprise’ the model as an anomaly.

The commonly used SSAD methods include one-class support vector machine (SVM), autoencoders, etc. [34], [35]. SSAD has been applied to detect fall using other sensor modalities [36]–[39]. However, we found little research on SSAD in radar-based fall detection systems. Diraco *et al.* [40] introduced one-class SVM and the K-means based approach using micro-Doppler features obtained from a 4.3 GHz radar. As normal ADL data is generally easy to collect and imminently develop, we prefer to use autoencoders [41], which, unlike other methods, can be incorporated into a neural network to learn large-scale datasets.

Particularly, we propose a Hybrid Variational RNN AutoEncoder (HVRAE) that adopts two autoencoder substructures, viz. i) the variational (inference) autoencoder (VAE) [42], a generative model rather than a discriminative model, to learn the radar data per frame; ii) a recurrent autoencoder (RAE) to learn temporal features over multiple frames to model fall as a sequence of events. Combining the VAE and RAE has been widely studied in computer vision (CV) and natural language processing (NLP) areas. Fabius *et al.* [43] developed the Variational Recurrent Autoencoder (VRAE) that first uses RAE to summarize the temporal features over multiple frames and then uses VAE to learn the distribution of the summarized features. Chung *et al.* [44] proposed a more profound structure called Variational RNN (VRNN) that applies VAE every frame to learn the distribution but conditioned on previous frame. Our HVRAE model, in which VAE performs on every frame independently, can be viewed as an adaptation of VRNN to simplify the temporal learning. Similar models have been proposed to detect anomaly in other applications as outlined in [45].

In summary, our major contributions include: i) the first method to detect fall based on the 4D radar point cloud of a human in a semi-supervised approach; ii) introducing a variational inference into radar point cloud distribution learning.

III. PRELIMINARIES

In this section, we introduce the background of all the components that constitute the proposed *mmFall* system detailed in the next section.

A. 4D mmWave FMCW Radar Sensor

The carrier frequency of mmWave frequency-modulated continuous-wave (FMCW) radar sensor, or mmWave radar sensor for short, ranges from 57 GHz to 85 GHz according

to various applications. For example, 76-81 GHz is primarily used for automotive applications such as objects' dynamics measurement [46], and 57-64 GHz can be used for short-range interactive motion sensing such as in Google's Soli project [47]. Coming along with the high carrier frequency, a high bandwidth up to 4 GHz is available, and the physical size of hardware components, including antennas, shrinks. This eventually makes the mmWave radar sensor more compact and higher resolution than the traditional low-frequency band radars.

There are no significant differences in signal modulation and processing of mmWave radar sensor than that of conventional FMCW radars described in [48]. Generally, the mmWave radar sensor transmits multiple linear FMCW signals over multiple antenna channels in both azimuth and elevation. After the stretch processing and digitalization, a raw multidimensional radar data cube is obtained. Followed by a series of fast Fourier transform (FFT), the parameters of each reflection point in a scene, i.e., range r , azimuth angle θ_{AZ} , elevation angle θ_{EL} , and Doppler DP , are estimated. In addition, during this process the constant false alarm rate (CFAR) is incorporated to detect the points with signal-to-noise ratio (SNR) greater than an adaptive threshold, and the moving target indication (MTI) is applied to distinguish the moving points from the static background. Eventually, a set of moving points, also called radar point cloud, is obtained.

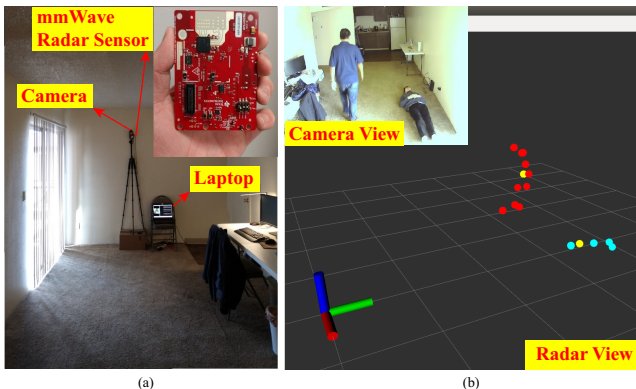


Fig. 1: MmWave radar sensor and radar point cloud. (a) The mmWave radar sensor is set up in an apartment, the camera provides a view for reference, and the laptop is used for data acquisition. The same setup is also used in the experiment in Section V. (b) Radar point cloud in a two-person scenario (lying down on the floor and walking). For the points, different color indicates different person while the yellow point indicates the centroid. For the coordinates, red is the cross-radar direction, green is the forward direction, and blue is the height direction. The original radar measurement of each point is a vector of $(r, \theta_{AZ}, \theta_{EL}, DP)$, along with the estimated centroid of (x_c, y_c, z_c) .

If multiple moving targets are present in a scene, the obtained point cloud is a collection of such points from all targets. Thus, a clustering method, such as the Density-Based Spatial Clustering of Applications with Noise (DBSCAN), has to be applied to segregate multiple targets. Meanwhile, the target's centroid can be estimated from the point subset associated with it. Followed by a tracking

algorithm, such as Kalman filtering, the trajectory of each target will be recorded with an association of a unique target ID. Particularly, a joint clustering/tracking algorithm called Group Tracking [49] can be used as well. Fig. 1 shows an example of the mmWave radar sensor and the radar point cloud. With the help of target ID, the motion history of each person can be gathered separately, such that we are able to analyze each person's motion individually. For simplicity but without loss of generality, we will only discuss the single-person scenario thereafter.

B. Radar Point Cloud Distribution for Human Body Motion

From Fig. 1 (b), a straightforward fall detection approach could be to analyze the height of the body centroid. For instance, a fall can be detected when there is a sudden drop in the body centroid. However, this approach may easily cause a false alarm when the person is crouching or sitting.

Considering the randomness in radar measurement, we now start to view the radar point cloud of the human body as a probabilistic distribution. From the observation in Fig. 1 (b), the distribution of the point cloud of the lying-down person is different than that of the walking person. Specifically, the covariance of the distribution is related to the human pose, and the mean is related to the human body centroid's location. Therefore, a distribution point of view has a physical significance, in a way that it represents the human pose and location of a person. Moreover, a motion, such as walking/fall, is a change of human pose/location over time, and we therefore call the pose/location as "motion state" for short. A depiction of motions is shown in Fig. 2.

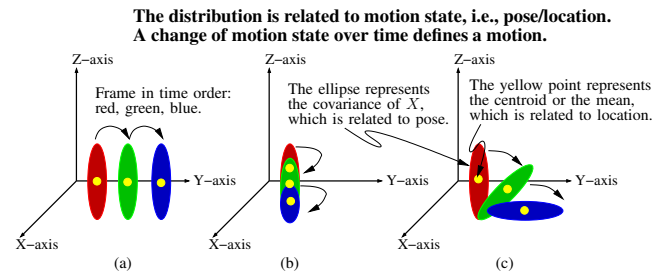


Fig. 2: A depiction of motion pattern. Compare this figure with Fig. 1 (b). (a) Walking; (b) Crouching; (c) Fall.

From the discussion above, now we make an assumption

Assumption 1. Let \mathbf{X} denote the radar point cloud of the human body. Let \mathbf{z} denote the body's motion state representing the pose and location. The assumption is, given a \mathbf{z} , the distribution of \mathbf{X} , i.e., the likelihood $p(\mathbf{X}|\mathbf{z})$, follows a particular multivariate Gaussian distribution. And a change of \mathbf{z} over multiple frames defines a motion, such as walking or fall, etc., we therefore need to infer $p(\mathbf{z}|\mathbf{X})$ at every frame and learn the change of \mathbf{z} over multiple frames.

Although Assumption 1 might not hold true as we never know the true physical generation process of radar data from a human body, we at least believe that this assumption is

enough for our purpose, i.e., distinguish different human motion. Therefore, we propose to intuitively detect fall through ‘learning’ the uniqueness of such motion patterns.

The overview of following subsections is, we propose to (i) learn the distribution at each frame through variational inference, (ii) learn the distribution change over multiple frames through recurrent neural network (RNN), (iii) and discuss, overall, in the framework of autoencoder for semi-supervised learning approach.

C. Variational Inference

More formally, at each frame we obtain a N -point radar point cloud $\mathbf{X}=\{\mathbf{x}_n\}_{n=1}^N$. The original radar measurement of each point \mathbf{x}_n is a four-dimensional vector of $(r, \theta_{AZ}, \theta_{EL}, DP)$. After coordinates transformation to the Cartesian coordinate, \mathbf{x}_n goes to be (x, y, z, DP) . We view the points in \mathbf{X} are independently drawn from the likelihood $p(\mathbf{X}|\mathbf{z})$, given a latent motion state \mathbf{z} which is a D -dimensional continuous vector. According to Assumption 1, $p(\mathbf{X}|\mathbf{z})$ follows multivariate Gaussian distribution. The Bayes’ theorem shows

$$\underbrace{p(\mathbf{z}|\mathbf{X})}_{\text{posterior}} = \frac{\overbrace{p(\mathbf{X}|\mathbf{z})p(\mathbf{z})}^{\text{likelihood prior}}}{\underbrace{p(\mathbf{X})}_{\text{evidence}}} = \frac{p(\mathbf{X}|\mathbf{z})p(\mathbf{z})}{\int p(\mathbf{X}|\mathbf{z})p(\mathbf{z})d\mathbf{z}}. \quad (1)$$

We expect to infer the motion state \mathbf{z} based on the observation \mathbf{X} . This is equivalent to infer the posterior $p(\mathbf{z}|\mathbf{X})$ of \mathbf{z} . Due to the difficulties in solving $p(\mathbf{z}|\mathbf{X})$ analytically as the evidence $p(\mathbf{X})$ is usually intractable, two major approximation approaches, i.e Markov Chain Monte Carlo (MCMC) and variational inference (VI), are mostly used.

Generally, the MCMC approach [50] uses a sampling method to draw enough samples from a tractable proposal distribution which is eventually approximate to the target distribution $p(\mathbf{z}|\mathbf{X})$. The most commonly used MCMC algorithm iteratively samples a data \mathbf{z}^t from an arbitrary tractable proposal distribution $q(\mathbf{z}^t|\mathbf{z}^{(t-1)})$ at step t , and then accept it with a probability of

$$\begin{aligned} & \min\left\{1, \frac{p(\mathbf{z}^t|\mathbf{X}) * q(\mathbf{z}^t|\mathbf{z}^{(t-1)})}{p(\mathbf{z}^{(t-1)}|\mathbf{X}) * q(\mathbf{z}^{(t-1)}|\mathbf{z}^t)}\right\} \\ & = \min\left\{1, \frac{p(\mathbf{X}|\mathbf{z}^t)p(\mathbf{z}^t) * q(\mathbf{z}^t|\mathbf{z}^{(t-1)})}{p(\mathbf{X}|\mathbf{z}^{(t-1)})p(\mathbf{z}^{(t-1)}) * q(\mathbf{z}^{(t-1)}|\mathbf{z}^t)}\right\}, \quad (2) \end{aligned}$$

where the difficult calculation of $p(\mathbf{X})$ has been circumvented. And it has been proven that this approach constructs a Markov chain whose equilibrium distribution equals to $p(\mathbf{z}|\mathbf{X})$ and is independent to the initial choice of $q(\mathbf{z}^0)$. One of the disadvantages in the MCMC approach is that the chain needs a long and indeterminable burn-in period to approximately reach the equilibrium distribution. This makes the MCMC not suitable for learning on large-scale dataset.

On the other hand, the VI approach [51] uses a family of tractable probability distribution $Q\{q(\mathbf{z})\}$ to approximate

the true $p(\mathbf{z}|\mathbf{X})$ instead of solving it analytically. The VI approach changes the inference problem to an optimization problem as

$$q^*(\mathbf{z}) = \arg \min_{q(\mathbf{z}) \in Q} \text{KLD}\{q(\mathbf{z})||p(\mathbf{z}|\mathbf{X})\}, \quad (3)$$

where KLD is the Kullback–Leibler divergence that measures the distance between two probability distributions. And by definition we have

$$\begin{aligned} & \text{KLD}\{q(\mathbf{z})||p(\mathbf{z}|\mathbf{X})\} \\ & := \int q(\mathbf{z}) \log \frac{q(\mathbf{z})}{p(\mathbf{z}|\mathbf{X})} d\mathbf{z} \\ & = \int q(\mathbf{z}) \log q(\mathbf{z}) d\mathbf{z} - \int q(\mathbf{z}) \log p(\mathbf{z}|\mathbf{X}) d\mathbf{z} \\ & = \mathbb{E}_q[\log q(\mathbf{z})] - \mathbb{E}_q[\log p(\mathbf{z}|\mathbf{X})] \\ & = \mathbb{E}_q[\log q(\mathbf{z})] - \mathbb{E}_q\{\log [p(\mathbf{X}|\mathbf{z})p(\mathbf{z})]\} + \mathbb{E}_q[\log p(\mathbf{X})] \\ & = \underbrace{\mathbb{E}_q[\log q(\mathbf{z})] - \mathbb{E}_q\{\log [p(\mathbf{X}|\mathbf{z})p(\mathbf{z})]\}}_{\mathcal{L}(q)} + \log p(\mathbf{X}), \quad (4) \end{aligned}$$

where $\mathbb{E}_q[*]$ is the statistical expectation operator of function $*$ whose variable follows $q(\mathbf{z})$, and \mathcal{L} is called the evidence low bound (ELBO). As the term $\log p(\mathbf{X})$ is constant with respect to \mathbf{z} , the optimization in Eq. (3) is simplified to be

$$q^*(\mathbf{z}) = \arg \min_{q(\mathbf{z}) \in Q} \mathcal{L}(q). \quad (5)$$

Here, the difficult computation of $p(\mathbf{X})$ is also circumvented. This optimization approach leads to one of the advantages of VI, that it can be integrated into a neural network framework and optimized through the back-propagation algorithm.

It is critical to choose the variational distribution $Q\{q(\mathbf{z})\}$ such that it is not only flexible enough to closely approximate the $p(\mathbf{z}|\mathbf{X})$, but also simple enough for efficient optimization. The most commonly used option is the factorized Gaussian family

$$q(\mathbf{z}) = \prod_{d=1}^D q(\mathbf{z}[d]) = \prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]), \quad (6)$$

where $(\boldsymbol{\mu}_q, \boldsymbol{\sigma}_q)$ are mean and covariance of the distribution of latent variable \mathbf{z} with a predetermined length of D , and the components in \mathbf{z} are mutually independent.

D. Variational Autoencoder

As we briefly state previously, we adopt the semi-supervised anomaly detection approach to train model only on normal ADL such that the model will be surprised by the ‘unseen’ fall data. The common approach is autoencoder, whose basic architecture is shown in Fig. 3 (a). The autoencoder consists of two parts, i.e., encoder and decoder. In most cases, the decoder is simply a mirror of the encoder. The encoder compresses the input data \mathbf{X} to a latent feature vector \mathbf{z} with fewer dimensions, and reversely the decoder reconstructs \mathbf{X}' to be as close to \mathbf{X} as possible, based on the latent feature vector \mathbf{z} . Generally, the Multilayer Perceptrons

(MLP) are used to model the non-linear mapping function between \mathbf{X} and \mathbf{z} , as the MLP is a powerful universal function approximator [52]. Besides a predetermined non-linear activation function, such as sigmoid/tanh, the MLP is characterized by its weights and biases. The training objective is to minimize the loss function between \mathbf{X} and \mathbf{X}' with respect to the weights and biases of encoder MLP and decoder MLP. The loss function could be cross-entropy for a categorical classification problem or mean square error (MSE) for a regression problem.




Fig. 3: Autoencoder architecture. (a) Vanilla autoencoder architecture. (b) Variational autoencoder architecture with factorized Gaussian parametrized by (μ_q, σ_q) .

In this way, the autoencoder squeezes the dimensionality to reduce the redundancy of input data. So it learns a compressed yet informative latent feature vector in \mathbf{X} . Therefore, the autoencoder will result in a close reconstruction \mathbf{X}' from the input data similar to \mathbf{X} , with a low reconstruction loss. However, whenever an ‘unseen’ data passes through, the autoencoder will erroneously squeeze it and be unable to reconstruct it well. This will lead to a loss spike from which an anomaly can be detected.

Similarly, in VAE [42], [53] in Fig. 3 (b), the encoder learns $q(\mathbf{z})$ that aims to approximate $p(\mathbf{z}|\mathbf{X})$ from the input data \mathbf{X} using VI approach, and the decoder reconstructs the $p(\mathbf{X}|\mathbf{z})$ based on \mathbf{z} sampled from the learned $q(\mathbf{z})$. The VAE training objective is as in Eq. 5. From Eq. 4, the loss function is

$$\mathcal{L}_{\text{VAE}} = \mathcal{L}(q) = \mathbb{E}_q[\log q(\mathbf{z})] - \mathbb{E}_q[\log p(\mathbf{z})] - \mathbb{E}_q[\log p(\mathbf{X}|\mathbf{z})] \\ = \text{KLD}\{q(\mathbf{z})||p(\mathbf{z})\} - \mathbb{E}_q[\log p(\mathbf{X}|\mathbf{z})]. \quad (7)$$

For the variational distribution $q(\mathbf{z})$, the factorized Gaussian in Eq. 6 is used, and for the prior $p(\mathbf{z})$, a common choice of Gaussian $\mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I})$ is used as we do not have a strong assumption on it. Therefore, the first term in \mathcal{L}_{VAE} in Eq. (7) is reduced to

$$\text{KLD}\{q(\mathbf{z})||p(\mathbf{z})\} \\ = -\frac{1}{2} \sum_{d=1}^D \{1 + \log \sigma_q[d]^2 - \mu_q[d]^2 - \sigma_q[d]^2\}, \quad (8)$$

where (μ_q, σ_q) is the mean and variance of the factorized Gaussian $q(\mathbf{z})$ with D -dimensional latent vector \mathbf{z} . See Appendix A for detailed derivation.

For the second term in \mathcal{L}_{VAE} in Eq. (7), it is reduced to

$$\mathbb{E}_q[\log p(\mathbf{X}|\mathbf{z})]$$

$$= \int q(\mathbf{z}) \log p(\mathbf{X}|\mathbf{z}) d\mathbf{z} \\ \text{(Using the single-data Monte Carlo estimation, the single-data } \mathbf{z} \text{ is sampled from } q(\mathbf{z}), \text{ then)} \\ \approx \log p(\mathbf{X}|\mathbf{z}) \quad (9a) \\ \text{(Using the Assumption 1 that the likelihood is Gaussian)} \\ = \log \mathcal{N}(\mathbf{X}|\mu_p, \sigma_p) \\ = \log \prod_{n=1}^N \mathcal{N}(\mathbf{x}_n|\mu_p, \sigma_p) \\ = \sum_{n=1}^N \sum_{k=1}^K \log \mathcal{N}(\mathbf{x}_n[k]|\mu_p[k], \sigma_p[k]) \\ \text{(Ignoring the constant } \log \sqrt{2\pi} \text{ in an optimization problem)} \\ \approx -\frac{1}{2} \sum_{n=1}^N \sum_{k=1}^K \left\{ \frac{(\mathbf{x}_n[k] - \mu_p[k])^2}{\sigma_p[k]^2} + \log \sigma_p[k]^2 \right\}, \quad (9)$$

where $\mathbf{X} = \{\mathbf{x}_n\}_{n=1}^N$ is the input point cloud, each point \mathbf{x}_n is a K -dimensional vector, and (μ_p, σ_p) is the mean and variance of the likelihood $p(\mathbf{X}|\mathbf{z})$.

In Eq. (9a), a single sample of \mathbf{z} is needed. Instead of drawing from $q(\mathbf{z})$ directly, the reparameterization trick [42], [54] is used as

$$\mathbf{z} = \mu_p + \sigma_p \odot \epsilon, \quad (10)$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, \odot is element-wise product. The trick is first to draw a sample ϵ from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, and then compute \mathbf{z} .

By viewing Eq. (8), the VAE encoder becomes clear as

$$(\mu_q, \log \sigma_q^2) = \text{EncoderMLP}_\phi\{\mathbf{X}\}, \quad (11)$$

where the weights and biases of encoder MLP are denoted as ϕ . In other words, the EncoderMLP_ϕ estimates the parameters of $q(\mathbf{z})$ from the input \mathbf{X} .

Similarly, from Eq. (9), the VAE decoder becomes clear

$$(\mu_p, \log \sigma_p^2) = \text{DecoderMLP}_\theta\{\mathbf{z}\}, \quad (12)$$

where the weights and biases of decoder MLP are denoted as θ . In other words, the DecoderMLP_θ estimates the parameters of $p(\mathbf{X}|\mathbf{z})$ from the \mathbf{z} sampled from Eq. (10).

Then the VAE architecture shown in Fig. 3 (b) becomes clear by combining EncoderMLP_ϕ and DecoderMLP_θ together, where these two parts are bridged through the sampling of \mathbf{z} . And the VAE training objective is to minimize the loss function with respect to the network parameters (ϕ, θ) . According to Eq. (7-9), the overall VAE loss function is

$$\mathcal{L}_{\text{VAE}} = \text{KLD}\{q(\mathbf{z})||p(\mathbf{z})\} - \mathbb{E}_q[\log p(\mathbf{X}|\mathbf{z})] \\ = \frac{1}{2} \sum_{n=1}^N \sum_{k=1}^K \left\{ \frac{(\mathbf{x}_n[k] - \mu_p[k])^2}{\sigma_p[k]^2} + \log \sigma_p[k]^2 \right\} - \\ \frac{1}{2} \sum_{d=1}^D \{1 + \log \sigma_q[d]^2 - \mu_q[d]^2 - \sigma_q[d]^2\}. \quad (13)$$

E. Recurrent Autoencoder

While we use the VI approach to learn the radar point cloud distribution at each frame, we also need a sequence-to-sequence modeling approach to learn distribution changes over multiple frames, as stated in Section III-B previously.

The recurrent neural network (RNN) is such a basic sequence-to-sequence model for temporal applications. At every frame l , an RNN accepts two inputs, input from the sequence at the l -th frame x_l and its previous hidden state h_{l-1} , to output a new hidden state h_l , calculated as:

$$h_l = \tanh(W * h_{l-1} + U * x_l) \quad \forall l = 1, 2, \dots, L \quad (14)$$

where W and U are learnable weights (including the bias term, omitted for brevity), and L is the length of the sequence. Note that at $l=1$, h_0 is defined as the initial RNN state that is either initialized as zeros, or randomly initialized. Also, note that the hidden state h_l acts as an accumulated memory state as it continuously computed and updated with new information in the sequence. Based on the basic RNN, the Long-Short-Term-Memory (LSTM) and Gated-Recurrent-Units (GRUs) [55], [56] has been developed to solve the vanishing/exploding gradients issue in modeling long term dependencies [57] in RNN. However, in our case, as a fall motion may last for about one second, that is ten frames for the radar data rate of ten frames per second, the long term dependency is not an issue here. Only the basic RNN is used for light computation load consideration.

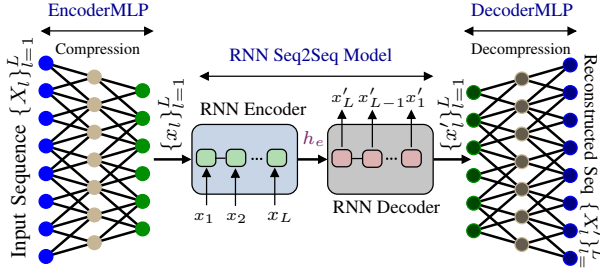


Fig. 4: A depiction of a Recurrent Autoencoder (RAE). The input sequence $\{X_i\}_{i=1}^L$ is first compressed to an embedded feature sequence $\{x_i\}_{i=1}^L$ on per frame basis through the EncoderMLP. An RNN Encoder iteratively processes the data over L frames and the final hidden state h_e is passed on to the RNN Decoder that outputs the reconstructed embedded feature sequence $\{x'_i\}_{i=1}^L$ in reverse. Finally, $\{x'_i\}_{i=1}^L$ are decompressed to reconstruct the sequence $\{X'_i\}_{i=1}^L$ through the DecoderMLP. The output sequence $\{X'_i\}_{i=1}^L$ is compared with the input sequence $\{X_i\}_{i=1}^L$ to compute the reconstruction loss, which is desired to be low for an autoencoder.

The RNN-based autoencoder [58] [59], or RAE as shown in Fig. 4, is built upon the vanilla autoencoder architecture in Fig. 3 (a). As the input is a time sequence of feature vectors, it has two dimensions, i.e., feature dimension and time dimension. In RAE, the EncoderMLP/DecoderMLP is for compressing and reconstructing the feature vector on per frame basis, and the RNN-Encoder/Decoder is for compressing and reconstructing the time sequence over multiple frames. Overall, the RAE reduces redundancy in both feature and time dimension.

IV. PROPOSED SYSTEM

To effectively learn the motion pattern of human body, which is formed by a sequence of radar point cloud, for fall detection in a semi-supervised approach, we propose a Hybrid Variational RNN AutoEncoder (HVRAE) which has two autoencoder substructures, i.e., VAE for learning radar point cloud distribution on per frame basis and RAE for learning the change of distribution over multiple frames. The HVRAE is trained only on normal ADL, such that an ‘unseen’ fall will cause a spike in the loss or anomaly level. If the height of body centroid, which is estimated from the point cloud, drops suddenly at the same time, a fall is detected. The proposed system, called *mmFall*, including both hardware and software, is presented in Fig. 5.

A. Data Preprocessing

With a proper mmWave radar sensor, we are able to collect the radar point cloud, as shown in Fig. 1 (b). In Fig. 5, the radar sensor could be mounted on the wall in a room with a height of h over the head of people, and could also be rotated with an angle θ_{tilt} so that it has a better coverage of the room. The radar sensor can detect multiple moving persons simultaneously, each person has a unique target ID as a result of the clustering/tracking algorithms. With the multiple frame data with the same target ID, we can analyze the motion of the person associated with this target ID. In other words, each person’s motion analysis can be processed separately based on the target ID. Afterwards, we will only discuss the single-person scenario for brevity.

We then propose a data preprocessing flow denoted in Fig. 5 for the following reasons.

The original measurement for each point in the radar point cloud is in the radar spherical coordinates. We need to transfer it to the radar Cartesian coordinates, and then to the ground Cartesian coordinates on the basis of the tilt angle and height. Therefore, we have a transformation matrix as

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{\text{tilt}} & \sin \theta_{\text{tilt}} \\ 0 & -\sin \theta_{\text{tilt}} & \cos \theta_{\text{tilt}} \end{bmatrix} \begin{bmatrix} r \cos \theta_{\text{EL}} \sin \theta_{\text{AZ}} \\ r \cos \theta_{\text{EL}} \cos \theta_{\text{AZ}} \\ r \sin \theta_{\text{EL}} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ h \end{bmatrix}, \quad (15)$$

where $(r, \theta_{\text{AZ}}, \theta_{\text{EL}})$ is range, azimuth angle and elevation angle in the radar spherical coordinates, θ_{tilt} is radar tilt angle, h is the radar platform height, and $[x, y, z]^T$ is the result in the ground Cartesian coordinates.

After coordinate transformation, at each frame we obtain a radar point cloud, in which each point is a vector of (x, y, z, DP) where DP is the Doppler from the original radar measurement. And we also have the centroid (x_c, y_c, z_c) as a result of the clustering/tracking algorithms in the radar.

We use a sliding window to accumulate every consecutive L frames as a motion pattern. The value of L equals to the radar frame rate in frames per second (fps) multiplied by the predetermined detection window in seconds. For each motion pattern with L frames, we subtract the x and y value of each point in each frame from the x_c and y_c value of

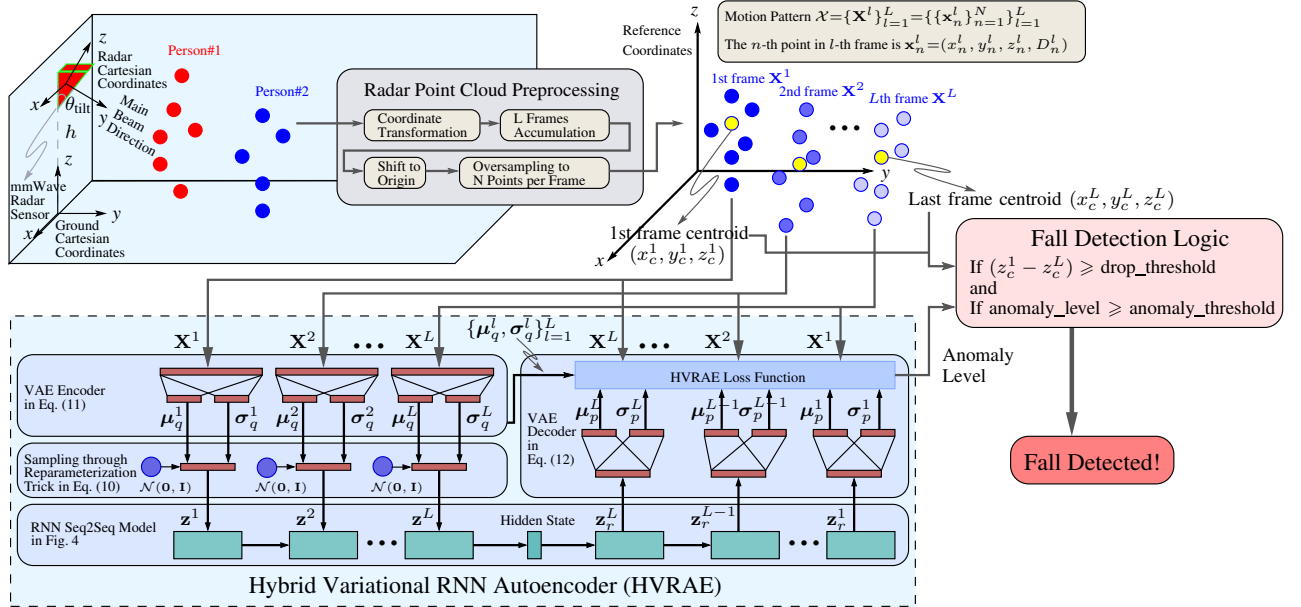


Fig. 5: An overview of the proposed *mmFall* System. At each frame, we obtain the point cloud of a human body along with its centroid from mmWave radar sensor. We use a sliding window to accumulate. After the preprocessing stage, we get a motion pattern $\mathcal{X} = \{\mathbf{X}^l\}_{l=1}^L$ in the reference coordinates. For each l -th frame, we use the VAE Encoder to model the mean μ_q^l and variance σ_q^l of the factorized Gaussian family $q(\mathbf{z}^l) = \prod_{d=1}^D \mathcal{N}(z_d^l | \mu_d^l, \sigma_d^l)$ that aims to approximate the true posterior $p(\mathbf{z}^l | \mathbf{X}^l)$ of the latent motion state \mathbf{z}^l , where D is the predetermined length of \mathbf{z} . Then we use the reparameterization trick to sample \mathbf{z}^l from $q(\mathbf{z}^l)$. After we have a sequence of latent motion states $\mathcal{Z} = \{\mathbf{z}^1, \dots, \mathbf{z}^L\}$, we use the RAE to compress and then reconstruct it as $\mathcal{Z}_r = \{\mathbf{z}_r^1, \dots, \mathbf{z}_r^L\}$. Based on \mathcal{Z}_r , we use the VAE Decoder to model the mean μ_p^l and variance σ_p^l of the likelihood $p(\mathbf{X}^l | \mathbf{z}^l)$. With (μ_q, σ_q) , (μ_p, σ_p) and \mathcal{X} , we are able to compute the HVRAE loss defined in Eq. 17 as an indication of anomaly level. In the fall detection logic, if a sudden drop of centroid height is detected at the same time when the HVRAE outputs an anomaly spike, we claim a fall detection.

Algorithm 1: Data Oversampling Method

Input: Input dataset $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^M$ with a length of M , M is a random number, each data sample \mathbf{x}_i is a vector. N , target length after oversampling. N is always $\geq M$.

Output: $\mathbf{X}' = \{\mathbf{x}'_i\}_{i=1}^N$ with a length of N .

```

1  $\hat{\mu} = \frac{1}{M} \sum_{i=1}^M \mathbf{x}_i$  // Get the estimated mean
2 for  $i = 1$  to  $N$  do
3   if  $i \leq M$  then // Rescale and shift
4      $\mathbf{x}'_i = \sqrt{\frac{N}{M}} \mathbf{x}_i + \hat{\mu} - \sqrt{\frac{N}{M}} \hat{\mu}$ ;
5   else // Pad with  $\hat{\mu}$ 
6      $\mathbf{x}'_i = \hat{\mu}$ ;
7   end
8 end
```

centroid in the first frame, respectively. In this way, we shift the motion pattern to the origin of a reference coordinates.

At each frame, the number of points in the radar point cloud is random due to the nature of radar measurement. We need a data oversampling method to meet the fixed input of the HVRAE model. The traditional oversampling method in deep learning is such as zero-padding or random oversampling. Zero-padding simply adds more zeros into the original data and random sampling simply duplicates some original data. Using both these two oversampling methods,

the distribution of the input may be changed. However, our purpose is to learning the distribution of radar point cloud and changing the distribution is definitely not what we want. Therefore, we propose a novel data oversampling Algorithm 1 that extends the original point cloud to a fixed number while keeping its distribution (mean and covariance) the same. The proof of this algorithm is in Appendix B.

Finally, we obtain a motion pattern \mathcal{X} in the reference coordinates,

$$\mathcal{X} = \{\mathbf{X}^l\}_{l=1}^L = \{\{\mathbf{x}_n^l\}_{n=1}^N\}_{l=1}^L = \{\{(x_n^l, y_n^l, z_n^l, DP_n^l)\}_{n=1}^N\}_{l=1}^L, \quad (16)$$

where L is the number of frames in the motion pattern; N is the number of points at each frame; \mathbf{X}^l is the l -th frame point cloud; \mathbf{x}_n^l is the n -th point in l -th frame, that is also a 4D vector of $(x_n^l, y_n^l, z_n^l, DP_n^l)$. We also have the centroid $\{(x_c^l, y_c^l, z_c^l)\}_{l=1}^L$ over L frames. Afterwards, we use the superscript l to denote the frame index.

B. HVRAE Model

The HVRAE architecture is shown in Fig. 5 and detailed in the caption. The HVRAE model is a combination of VAE and RAE, discussed in the previous section. The HVRAE loss $\mathcal{L}_{\text{HVRAE}}$ is the VAE loss \mathcal{L}_{VAE} in Eq. (13) over all the L frames. Then, we have

$$\mathcal{L}_{\text{HVRAE}} = \sum_{l=1}^L \left\{ \text{KLD}\{q(\mathbf{z}^l) || p(\mathbf{z}^l)\} - \mathbb{E}_q[\log p(\mathbf{X}^l | \mathbf{z}^l)] \right\}$$

$$\begin{aligned}
&= \sum_{l=1}^L \left\{ \frac{1}{2} \sum_{n=1}^N \sum_{k=1}^K \left\{ \frac{(\mathbf{x}_n^l[k] - \boldsymbol{\mu}_p^l[k])^2}{\sigma_p^l[k]^2} + \log \sigma_p^l[k]^2 \right\} \right. \\
&\quad \left. - \frac{1}{2} \sum_{d=1}^D \left\{ 1 + \log \sigma_q^l[d]^2 - \boldsymbol{\mu}_q^l[d]^2 - \sigma_q^l[d]^2 \right\} \right\}
\end{aligned} \tag{17}$$

where L , N and \mathbf{x}_n^l are from the motion pattern in Equ (16); K is the length of point vector, in our case $K=4$ as each point is a 4D vector; D is the length of latent motion state \mathbf{z} ; $(\boldsymbol{\mu}_q, \sigma_q)$ and $(\boldsymbol{\mu}_p, \sigma_p)$ are parameters of factorized Gaussian $q(\mathbf{z})$ and likelihood $p(\mathbf{X}|\mathbf{z})$, respectively, both are modeled through the architecture in Fig. 5.

For HVRAE training, the objective is to minimize $\mathcal{L}_{\text{HVRAE}}$ with respect to the network parameters. The standard stochastic gradient descent algorithm Adam [53] is used.

It is noted that, for the implementation of VAE Encoder/Decoder in HVRAE, only a dense layer or fully-connected layer is used, as the model should be invariant to the order of point cloud at each frame.

C. Fall Detection Logic

In a semi-supervised learning approach, we train this HVRAE model only on normal ADL, which are easy to collect compared to falls. For normal ADL, the HVRAE will output a low $\mathcal{L}_{\text{HVRAE}}$ as this is the training objective. In the inference stage, the model will generate a high loss $\mathcal{L}_{\text{HVRAE}}$ when an ‘unseen’ motion happens, such as fall occurs. Therefore, we denote the HVRAE loss $\mathcal{L}_{\text{HVRAE}}$ as an anomaly level measure of human body motion.

Along with the body centroid height $\{z_c^l\}_{l=1}^L$ over L frames, we can calculate the drop of centroid height as $z_c^1 - z_c^L$ during this motion. Then we propose a fall detection logic as in Fig. 5, that is if the centroid height drop is greater than a threshold at the same time when the anomaly level is greater than a threshold, we claim a fall detection.

According to the fall definition from WHO as in the Section I, in the proposed *mmFall* system, the HVRAE measures the inadvertence or anomaly level of the motion, while the centroid height drop indicates the motion of coming to rest on a lower level.

V. EXPERIMENTAL RESULTS AND DISCUSSION

To verify the effectiveness of the proposed system, we used a mmWave radar sensor to collect experimental data and implemented the proposed *mmFall* system along with two baselines for performance evaluation and comparison.

A. Hardware Configuration and Experiment Setup

We adopt the Texas Instrument (TI) AWR1843BOOST mmWave FMCW radar evaluation board [60] for radar point cloud acquisition. This radar sensor has three transmitting antenna channels and four receiving antenna channels, as shown in Fig. 1 (a). The middle transmitting channel is displaced above the other two by a distance of half a wavelength. Through the direction-of-angle (DOA) algorithm

using multiple-input and multiple-output (MIMO), it can achieve 2x4 MIMO in azimuth and 2x1 MIMO in elevation. Thus, we have 3D positional measurement of each point. Plus the 1D Doppler, we finally have a 4D radar point cloud. Based on a demo project from TI [61], we configure the radar sensor with the parameters listed in Table I.

TABLE I: mmWave FMCW radar parameter configuration. Refer to [62] for waveform details. f_s , FMCW starting frequency. BW, FMCW bandwidth. r_{Chirp} , FMCW chirp rate. f_{ADC} , ADC sampling rate. N_{Fast} , ADC samples per chirp. CPI, coherent processing interval. N_{Slow} , chirps per CPI per transmitting channel. T_{Frame} , duration of one frame. ΔR , range resolution. R_{max} , maximum unambiguous range. ΔD , Doppler resolution. D_{max} , maximum unambiguous Doppler. $\Delta\theta_{\text{AZ}}$, azimuth angle resolution. $\Delta\theta_{\text{EL}}$, elevation angle resolution. r_{Frame} , frame rate in frames per second.

Parameter	Value	Unit	Parameter	Value	Unit
f_s	77	GHz	ΔR	0.078	m
BW	1.92	GHz	R_{max}	9.99	m
r_{Chirp}	30	MHz/us	ΔD	0.079	m/s
f_{ADC}	2	MHz	D_{max}	± 2.542	m/s
N_{Fast}	128		MIMO	2x4/2x1	AZ/EL
CPI	24.2	ms	$\Delta\theta_{\text{AZ}}$	15	deg
N_{Slow}	64		$\Delta\theta_{\text{EL}}$	57	deg
T_{Frame}	100	ms	r_{Frame}	10	fps

Based on the Robotic Operating System (ROS) on an Ubuntu laptop, we developed an interface program to connect the TI AWR1843BOOST and collect the radar point cloud over the USB port. Then we set up the equipment in the living room (2.7m*8.2m*2.7m) in an apartment, as shown in Fig. 1 (a). There are two large desks in the living room and most area is relatively empty. More occlusion discussion can be found in Section V-E. The radar sensor was put on top of a tripod with a height of 2 meters, and rotated with a tilt angle of 10 degrees for better area coverage. Later, we processed the collected data offline using a Jupyter Notebook that you can find in the GitHub repository.

B. Data Collection

During the experiment, the first two authors, as shown in Fig. 1 (b), collected three datasets in Table II together. Firstly, we collected the DS_0 dataset which contains about two hours of normal ADL without any labels for training. Secondly, in the DS_1 dataset, we collected randomly walking along with one sample of each other motion, including fall, etc. We showed the motion pattern for every motion in Fig. 6 for visualization purposes. Lastly, we collected a comprehensive inference dataset DS_2 and manually labeled the frame index when a fall happens as the ground truth, and it is used for overall inference performance evaluation. It is noted that in DS_1 and DS_2 , both the fall and jump are anomalies that can not be found in DS_0 . We expect that HVRAE will output an anomaly level spike for both fall and jump, but the fall detection logic involving the centroid height drop will reject the jump but detect the fall.

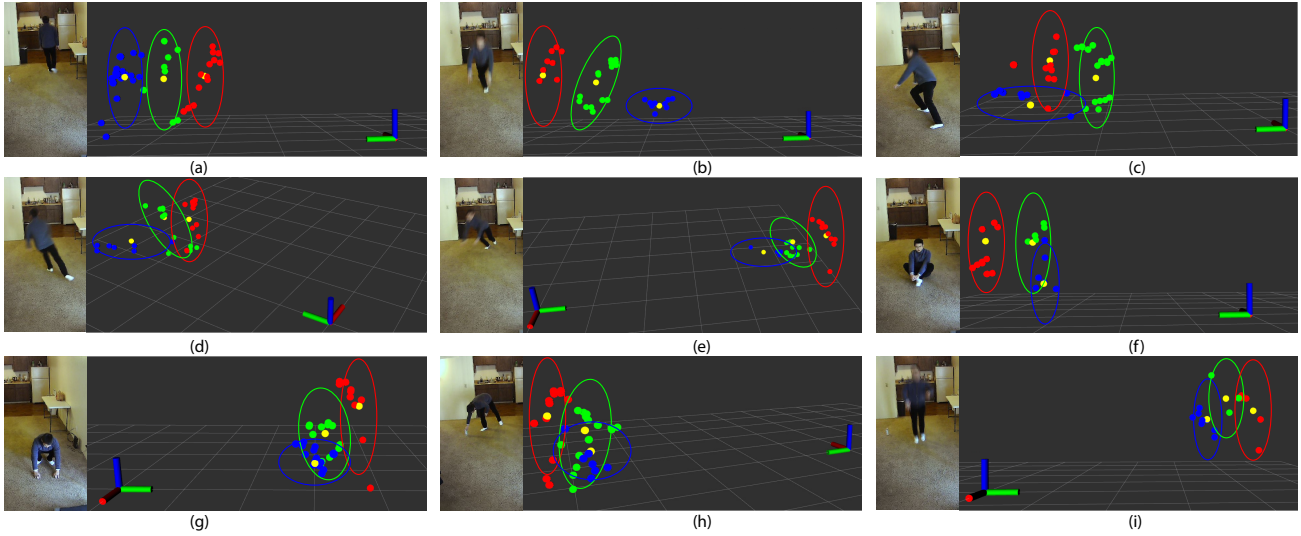


Fig. 6: Motion patterns in dataset DS_1 along with the associated camera view. Only the ellipse was manually added for depicting the distribution of point cloud. For the points, different color indicates the frame in time order: red, green, and then blue, while the yellow point indicates the centroid estimated by the mmWave radar sensor. For simplicity, we showed the frames with the increment of five frames. Each frame is 0.1 seconds. Please compare this figure with Fig. 2. For the coordinates, red is the cross-radar direction, green is the forward direction, and blue is the height direction. (a) Randomly walking; (b) Forward fall; (c) Backward fall; (d) Left fall; (e) Right fall; (f) Sitting down on the floor; (g) Crouching; (h) Bending; (i) jump.

TABLE II: Collected Dataset.

Name	Description
DS_0	Two hours of normal ADL, including randomly walking, sitting on the floor, crouching, bending, etc. No labeling.
DS_1	Randomly walking with one forward fall, one backward fall, one left fall, one right fall, one sitting on the floor, one crouching, one bending, and one jump.
DS_2	Randomly walking with 15 forward falls, 15 backward falls, 10 left falls, 10 right falls, 50 sitting on the floor, 50 crouching, 50 bending, and 50 jump. Labeling fall as ground truth.

C. Model Implementation and Two Baselines

We first implemented the proposed *mmFall* system in Fig. 5 on Keras (Tensorflow backend), with loss function in Eq. (17). In this implementation, we set the number of frames, L , equal to 10 for a one-second detection window with 10 fps radar data rate; the number of points each frame N equal to 64 for data oversampling. Thus, the motion pattern \mathcal{X} , i.e., the model input, is $10*64*4$. We set the length of latent motion state z , D , equal to 16. For performance comparison purposes, we also implemented two other baselines. All the three models are listed in Table III.

The baseline HVRAE_SL is the same as the proposed *mmFall* system except for using a simplified loss function in Eq. (18). The simplified loss function Eq. (18) is based on a weak assumption on likelihood, that is $p(\mathbf{X}|\mathbf{z})$ follows a Gaussian with identity covariance, i.e., $\mathcal{N}(\boldsymbol{\mu}_p, \mathbf{I})$. This leads to that the σ_p term in Eq. (17) is ignored, or

$$\mathcal{L}_{\text{HVRAE_SL}} = \sum_{l=1}^L \left\{ \frac{1}{2} \sum_{n=1}^N \sum_{k=1}^K \{ (\mathbf{x}_n^l[k] - \boldsymbol{\mu}_p^l[k])^2 \} - \frac{1}{2} * \sum_{d=1}^D \{ 1 + \log \sigma_q^l[d]^2 - \boldsymbol{\mu}_q^l[d]^2 - \sigma_q^l[d]^2 \} \right\}. \quad (18)$$

To compare HVRAE with HVRAE_SL, we will verify that the concept that the covariance represents the pose contributes to the radar point cloud learning for human motion inference, as discussed in Section III-B.

Another baseline is RAE with MSE loss in Fig. 4, which uses MLP in the feature dimension instead of VI approach in HVRAE every frame. To compare HVRAE with RAE, we will show that the VI approach for motion state inference based on the distribution of radar point cloud makes more sense than the vanilla MLP feature compression.

TABLE III: Implemented Models.

Name	Description
HVRAE	The proposed HVRAE and fall detection logic in Fig. 5 with loss function in Eq. (17).
HVRAE_SL	The proposed HVRAE and fall detection logic in Fig. 5 with simplified loss in Eq. (18).
RAE	The vanilla RAE in Fig. 4 with MSE loss function and fall detection logic in Fig. 5.

D. Training and Inference

First, we trained these three models on the normal dataset DS_0 , and then tested on dataset DS_1 in which there are some normal motions as in DS_0 and two different ‘unseen’ motions, i.e., fall and jump, that do not appear in DS_0 . The anomaly level outputted by these three models on DS_1 is shown in Fig. 7. The proposed HVRAE model can generate significant anomaly level for fall and jump while keeping low for normal motions. Along with the fall detection logic involving body centroid drop, the jump will be rejected, and only fall will be detected. As a comparison, the HVRAE_SL model suffers great noise during normal motions that easily

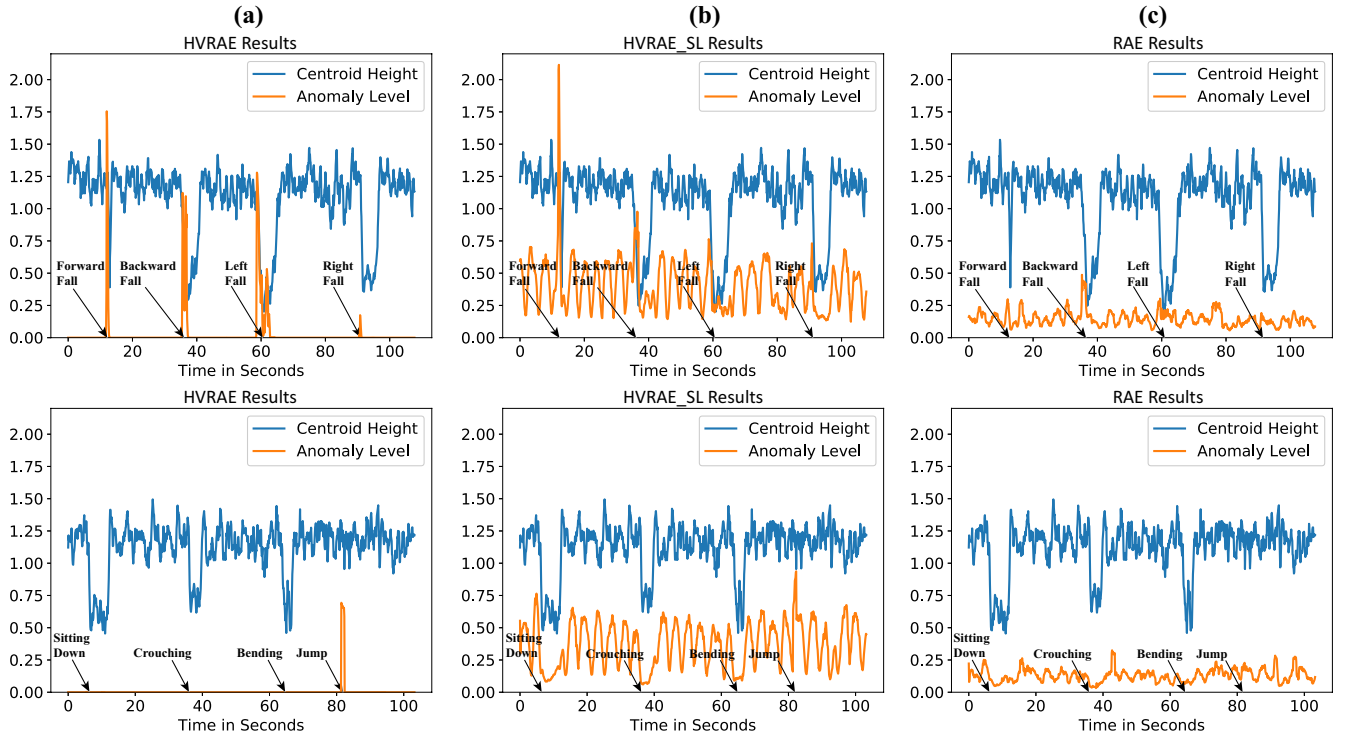


Fig. 7: Inference results of the models listed in Tab III on the dataset DS_1 described in Tab II. In each figure, the blue line represents the body’s centroid height, and the orange line represents the model’s loss output, or anomaly level. Only the black text and arrows were manually added as the ground truth when a motion happens. Except for the motion indicated by the black text, the rest of time are always randomly walking. (a) HVRAE inference results: The HVRAE model can clearly generate a spike in anomaly level when fall/jump happens while keeping low anomaly level for normal motions. Jump is another abnormal motion that does not appear in the training dataset DS_0 , but the fall detection logic involving the body centroid drop at the same time will reject jump. On the other hand, without the help of anomaly level it is difficult to distinguish fall from other motions if only the change of centroid height is considered; (b) HVRAE_SL inference results: The HVRAE_SL can also have anomaly level spike generation for fall/jump but suffer significant noise during normal motion occurrence. For example, the ‘Sitting Down’ and the ‘Right Fall’ have almost the same anomaly level output. As a result, either the ‘Sitting Down’ causes a false alarm, or ‘Right Fall’ causes a missed detection, depending on the threshold; (c) Vanilla RAE inference results: The vanilla RAE model can not effectively learn the anomaly level for ‘unseen’ motions.

leads to false alarm, and the vanilla RAE model can not learn the anomaly level effectively.

Finally, we tested these three well-trained models on the dataset DS_2 . In DS_2 , there are 50 falls with manually labeled ‘ground truth fall frame index’ when a fall happens, along with many other different motions without labeling. The fall detection logic will detect the frame index when a fall happens. We allow a flexible detection, i.e., if the ‘detected fall frame index’ falls into the 1-second detection window centered at one ‘ground truth fall frame index’, we treat it as true positive. In this experiment, for the fall detection logic we fixed the threshold of centroid height drop as 0.6 meters. By varying the anomaly level threshold, we got the Receiver Operating Characteristic (ROC) curves as shown in Fig. 8. The false alarm will cause waste of caring resources, and high true fall detection rate guarantees the elderly safety. Therefore, we want to achieve as high fall detection rate as possible at the expense of a few false alarms. From the ROC, we clearly see that HVRAE outperformed the other two baselines. Specifically, at the expense of two alarms, our HVRAE model can achieve 98% fall detection rate out of 50 falls, while the HVRAE_SL

can only achieve around 60% and the vanilla RAE can only achieve around 38%.




Fig. 8: ROC curves for all the three models.

E. Limitations of Current Research and Future Work

In this research, we did the experiment in a relatively empty apartment where occlusion is not a problem. To have more comprehensive results, an experiment in a complex

living environment is necessary. Also, the radar sensor is essentially robust to occlusion due to the nature of radio frequency (RF) signal. Basically, the Signal-to-Noise Ratio (SNR), that is related to the radar hardware's noise figure and transmitting power, determines the occlusion performance. A powerful radar sensor can even 'see' through the wall [63]. In addition, the current radar device in our experiment is not a optimized design and it may not acquire enough points from the human body occasionally, which may cause a false negative or missed detection. To obtain a more practical validation and robust result, in the future we aim to incorporate necessary hardware engineering to improve the SNR of the current radar sensor, and demonstrate the performance in an apartment with lots of furniture.

The human subjects in the experiment are the first two authors who have a very similar body figure ($\sim 175cm/75kg$). As we view the point cloud of human body as a distribution, typically, if there is a huge difference of body figure than the subjects' in this experiment, the distribution should also be quite different. Therefore, the model trained in this experiment can not be directly applied to the person with a significantly different body figure, for example $\sim 190cm/110kg$. In the future, we will collect more training data from multiple human subjects with a wide range of build/height, to make the model be able to cover more cases. On the other hand, in the experiment, we mainly focused on a quick fall event. However, in practice, the elderly may fall in a slow motion, which means our 1-second detection window may not be enough.

VI. CONCLUSION

In this study, we used a mmWave radar sensor for fall detection on the basis of its advantages such as privacy-compliant, non-wearable, sensitive to motions, etc. We made an assumption that the radar point cloud for the human body can be viewed as a multivariate Gaussian distribution, and the distribution change over multiple frames has a unique pattern for different motions. And then, we proposed a Hybrid Variational RNN AutoEncoder to effectively learn the anomaly level of 'unseen' motion, such as fall, that does not appear in the normal training dataset. We also involved a fall detection logic that checks the body centroid drop to further confirm the anomaly motion is fall. In this way, we detected the fall in a semi-supervised learning approach that does not require the difficult fall data collection and labeling. The experiment results showed our proposed system can achieve 98% detection rate out of 50 falls at the expense of just two false alarms, and outperformed the other two baselines. In the future, we will have necessary hardware engineering to improve the SNR and demonstrate the occlusion performance of the mmWave radar sensor in a complex living environment and also collect more training data from people with different body figures.

APPENDIX A PROOF OF EQ. 8

$$\begin{aligned}
& \text{KLD}\{q(\mathbf{z})||p(\mathbf{z})\} \\
&= \text{KLD} \left\{ \prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) || \prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|0, 1) \right\} \\
&:= \int \dots \int \left\{ \prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) * \right. \\
&\quad \left. \log \frac{\prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d])}{\prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|0, 1)} \right\} d\mathbf{z}[1] \dots d\mathbf{z}[D] \\
&= \int \dots \int \left\{ \prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) * \right. \\
&\quad \left. \sum_{d=1}^D \log \frac{\mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d])}{\mathcal{N}(\mathbf{z}[d]|0, 1)} \right\} d\mathbf{z}[1] \dots d\mathbf{z}[D] \\
&= \sum_{d=1}^D \left\{ \int \dots \int \prod_{d=1}^D \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) * \right. \\
&\quad \left. \log \frac{\mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d])}{\mathcal{N}(\mathbf{z}[d]|0, 1)} d\mathbf{z}[1] \dots d\mathbf{z}[D] \right\} \\
&= \sum_{d=1}^D \left\{ \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) \log \frac{\mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d])}{\mathcal{N}(\mathbf{z}[d]|0, 1)} d\mathbf{z}[d] \right\} \\
&= \sum_{d=1}^D \left\{ \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) \right. \\
&\quad \left. \left\{ \log \frac{1}{\sqrt{\boldsymbol{\sigma}_q[d]^2}} - \frac{(\mathbf{z}[d] - \boldsymbol{\mu}_q[d])^2}{2\boldsymbol{\sigma}_q[d]^2} + \frac{\mathbf{z}[d]^2}{2} \right\} d\mathbf{z}[d] \right\} \\
&= -\frac{1}{2} \sum_{d=1}^D \left\{ \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) \log \boldsymbol{\sigma}_q[d]^2 d\mathbf{z}[d] + \right. \\
&\quad \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) \frac{(\mathbf{z}[d] - \boldsymbol{\mu}_q[d])^2}{\boldsymbol{\sigma}_q[d]^2} d\mathbf{z}[d] - \\
&\quad \left. \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) \mathbf{z}[d]^2 d\mathbf{z}[d] \right\} \\
&= -\frac{1}{2} \sum_{d=1}^D \left\{ \log \boldsymbol{\sigma}_q[d]^2 * \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) d\mathbf{z}[d] + \right. \\
&\quad \frac{1}{\boldsymbol{\sigma}_q[d]^2} * \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) (\mathbf{z}[d] - \boldsymbol{\mu}_q[d])^2 d\mathbf{z}[d] - \\
&\quad \left. \int \mathcal{N}(\mathbf{z}[d]|\boldsymbol{\mu}_q[d], \boldsymbol{\sigma}_q[d]) \mathbf{z}[d]^2 d\mathbf{z}[d] \right\} \\
&\quad (\text{As } \mathbb{E}[(x - \mu)^2] = \sigma^2 \text{ and } \mathbb{E}[x^2] = \mu^2 + \sigma^2, \text{ then}) \\
&= -\frac{1}{2} \sum_{d=1}^D \left\{ 1 + \log \boldsymbol{\sigma}_q[d]^2 - \boldsymbol{\mu}_q[d]^2 - \boldsymbol{\sigma}_q[d]^2 \right\}, \quad (19)
\end{aligned}$$

where $(\boldsymbol{\mu}_q, \boldsymbol{\sigma}_q)$ is the mean and variance of of the factorized Gaussian $q(\mathbf{z})$ with D -dimensional latent vector \mathbf{z} .

APPENDIX B
PROOF OF ALGORITHM 1

Given a set of statistically independent and identically distributed (i.i.d.) data $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^M$ drawn from a multivariate Gaussian random variable $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma})$, thus the maximum likelihood (ML) estimator of its mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\sigma}$ is,

$$\hat{\boldsymbol{\mu}} = \frac{1}{M} \sum_{i=1}^M \mathbf{x}_i, \quad \hat{\boldsymbol{\sigma}} = \frac{1}{M} \sum_{i=1}^M \{\mathbf{x}_i - \hat{\boldsymbol{\mu}}\}^2. \quad (20)$$

For the output dataset $\mathbf{X}' = \{\mathbf{x}'_i\}_{i=1}^N$, its first M elements are modified from the input dataset according to Step 4 in Algorithm (1), and its last $(N - M)$ elements are simply the mean of the input dataset according to Step 6 in Algorithm (1). Thus, its ML estimator of its mean $\boldsymbol{\mu}'$ and covariance $\boldsymbol{\sigma}'$ is

$$\begin{aligned} \hat{\boldsymbol{\mu}}' &= \frac{1}{N} \sum_{i=1}^N \mathbf{x}'_i = \frac{1}{N} \left\{ \sum_{i=1}^M \left(\sqrt{\frac{N}{M}} \mathbf{x}_i + \hat{\boldsymbol{\mu}} - \sqrt{\frac{N}{M}} \hat{\boldsymbol{\mu}} \right) + \sum_{i=M+1}^N \hat{\boldsymbol{\mu}} \right\} \\ &= \frac{1}{N} \left\{ \sqrt{\frac{N}{M}} \sum_{i=1}^M \mathbf{x}_i + M \hat{\boldsymbol{\mu}} - M \sqrt{\frac{N}{M}} \hat{\boldsymbol{\mu}} + (N - M) \hat{\boldsymbol{\mu}} \right\} \\ &= \frac{1}{N} \left\{ \sqrt{\frac{N}{M}} M \hat{\boldsymbol{\mu}} + M \hat{\boldsymbol{\mu}} - M \sqrt{\frac{N}{M}} \hat{\boldsymbol{\mu}} + (N - M) \hat{\boldsymbol{\mu}} \right\} \\ &= \frac{1}{N} \{N \hat{\boldsymbol{\mu}}\} = \hat{\boldsymbol{\mu}}, \end{aligned} \quad (21)$$

and

$$\begin{aligned} \hat{\boldsymbol{\sigma}}' &= \frac{1}{N} \sum_{i=1}^N \{\mathbf{x}'_i - \hat{\boldsymbol{\mu}}'\}^2 = \frac{1}{N} \sum_{i=1}^N \{\mathbf{x}'_i - \hat{\boldsymbol{\mu}}\}^2 \\ &= \frac{1}{N} \sum_{i=1}^M \left\{ \sqrt{\frac{N}{M}} \mathbf{x}_i + \hat{\boldsymbol{\mu}} - \sqrt{\frac{N}{M}} \hat{\boldsymbol{\mu}} - \hat{\boldsymbol{\mu}} \right\}^2 + \frac{1}{N} \sum_{i=(M+1)}^N \{\hat{\boldsymbol{\mu}} - \hat{\boldsymbol{\mu}}\}^2 \\ &= \frac{1}{N} \sum_{i=1}^M \left\{ \sqrt{\frac{N}{M}} \mathbf{x}_i - \sqrt{\frac{N}{M}} \hat{\boldsymbol{\mu}} \right\}^2 \\ &= \frac{1}{M} \sum_{i=1}^M \{\mathbf{x}_i - \hat{\boldsymbol{\mu}}\}^2 = \hat{\boldsymbol{\sigma}}. \end{aligned} \quad (22)$$

Therefore, the proposed algorithm oversamples the original input dataset to a fixed number while keeping the ML estimation of mean and variance the same.

REFERENCES

- [1] *World Population Prospects 2019: Highlights (ST/ESA/SER.A/423)*, Department of Economic and Social Affairs, Population Division, United Nations, 2019. [Online]. Available: https://population.un.org/wpp/Publications/Files/WPP2019_Highlights.pdf
- [2] *WHO Global Report on Falls Prevention in Older Age*, World Health Organization, 2008. [Online]. Available: <https://extranet.who.int/agefriendlyworld/wp-content/uploads/2014/06/WHO-Global-report-on-falls-prevention-in-older-age.pdf>
- [3] (2018, Jan.) Falls. World Health Organization. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/falls>
- [4] E. R. Burns, J. A. Stevens, and R. Lee, "The direct costs of fatal and non-fatal falls among older adults—United States," *J. Safety Res.*, vol. 58, pp. 99–103, 2016.
- [5] K. Chaccour *et al.*, "From fall detection to fall prevention: A generic classification of fall-related systems," *IEEE Sensors J.*, vol. 17, no. 3, pp. 812–822, Feb 2017.
- [6] *mmWave radar sensors in robotics applications*, Texas Instruments, 2017. [Online]. Available: <http://www.ti.com/lit/wp/spr311/spr311.pdf>
- [7] J. K. Lee, S. N. Robinovitch, and E. J. Park, "Inertial sensing-based pre-impact detection of falls involving near-fall scenarios," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 23, no. 2, pp. 258–266, March 2015.
- [8] J. Liu and T. E. Lockhart, "Development and evaluation of a prior-to-impact fall event detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 7, pp. 2135–2140, July 2014.
- [9] J. Sun *et al.*, "A plantar inclinometer based approach to fall detection in open environments," in *Emerg. Trends Adv. Technol. Comput. Intell.* Springer, 2016, pp. 1–13.
- [10] B. Mirmahboub *et al.*, "Automatic monocular system for human fall detection based on variations in silhouette area," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 2, pp. 427–436, Feb 2013.
- [11] Z.-P. Bian *et al.*, "Fall detection based on body part tracking using a depth camera," *IEEE J. Biomed. Health Informat.*, vol. 19, no. 2, pp. 430–439, 2014.
- [12] Y. Li, K. Ho, and M. Popescu, "A microphone array system for automatic fall detection," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 5, pp. 1291–1301, 2012.
- [13] K. Chaccour *et al.*, "Smart carpet using differential piezoresistive pressure sensors for elderly fall detection," in *Proc. IEEE 11th Int. Conf. Wireless Mobile Comput. Netw. Commun. (WiMob)*, 2015, pp. 225–229.
- [14] X. Fan *et al.*, "Robust unobtrusive fall detection using infrared array sensors," in *Proc. IEEE Int. Conf. Multisensor Fusion Integr. Intell. Syst. (MFI)*, 2017, pp. 194–199.
- [15] L. Ren and Y. Peng, "Research of fall detection and fall prevention technologies: A systematic review," *IEEE Access*, vol. 7, pp. 77 702–77 722, 2019.
- [16] M. G. Amin *et al.*, "Radar signal processing for elderly fall detection: The future for in-home monitoring," *IEEE Signal Process. Mag.*, vol. 33, no. 2, pp. 71–80, March 2016.
- [17] S. Z. Gurbuz and M. G. Amin, "Radar-based human-motion recognition with deep learning: Promising applications for indoor monitoring," *IEEE Signal Process. Mag.*, vol. 36, no. 4, pp. 16–28, July 2019.
- [18] M. S. Seyfioğlu, A. M. Özbayoğlu, and S. Z. Gürbüz, "Deep convolutional autoencoder for radar-based classification of similar aided and unaided human activities," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 54, no. 4, pp. 1709–1723, Aug 2018.
- [19] B. Y. Su *et al.*, "Doppler radar fall activity detection using the wavelet transform," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 3, pp. 865–875, 2014.
- [20] H. Sadreazami, M. Bolic, and S. Rajan, "Capsfall: Fall detection using ultra-wideband radar and capsule network," *IEEE Access*, vol. 7, pp. 55 336–55 343, 2019.
- [21] H. Yoshino, V. G. Moshnyaga, and K. Hashimoto, "Fall detection on a single doppler radar sensor by using convolutional neural networks," in *2019 IEEE Int. Conf. on Syst., Man, Cybern. (SMC)*, 2019, pp. 2889–2892.
- [22] A. Seifert, M. G. Amin, and A. M. Zoubir, "Toward unobtrusive in-home gait analysis based on radar micro-doppler signatures," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 9, pp. 2629–2640, 2019.
- [23] Y. Shankar, S. Hazra, and A. Santra, "Radar-based non-intrusive fall motion recognition using deformable convolutional neural network," in *2019 18th IEEE Int. Conf. Mach. Learn. Appl. (ICMLA)*, 2019, pp. 1717–1724.
- [24] B. Jakanović and M. Amin, "Fall detection using deep learning in range-doppler radars," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 54, no. 1, pp. 180–189, 2017.
- [25] A. Bhattacharya and R. Vaughan, "Deep learning radar design for breathing and fall detection," *IEEE Sensors J.*, vol. 20, no. 9, pp. 5072–5085, 2020.
- [26] C. Ding *et al.*, "Fall detection with multi-domain features by a portable fmcw radar," in *2019 IEEE MTT-S Int. Wireless Symp. (IWS)*, 2019, pp. 1–3.
- [27] B. Erol and M. G. Amin, "Radar data cube processing for human activity recognition using multisubspace learning," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 55, no. 6, pp. 3617–3628, 2019.

- [28] Y. Tian *et al.*, “Rf-based fall monitoring using convolutional neural networks,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 2, no. 3, Sep. 2018.
- [29] F. Jin *et al.*, “Multiple patients behavior detection in real-time using mmwave radar and deep CNNs,” in *2019 IEEE Radar Conf. (RadarConf)*, 2019, pp. 1–6.
- [30] R. Zhang and S. Cao, “Real-time human motion behavior detection via cnn using mmwave radar,” *IEEE Sens. Lett.*, vol. 3, no. 2, pp. 1–4, 2019.
- [31] A. Sengupta *et al.*, “mm-Pose: Real-time human skeletal posture estimation using mmwave radars and cnns,” *IEEE Sensors J.*, pp. 1–12, 2020.
- [32] Y. Sun *et al.*, “Privacy-preserving fall detection with deep learning on mmWave radar signal,” in *2019 IEEE Vis. Commun. Image Process. (VCIP)*, 2019, pp. 1–4.
- [33] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Comput. Surv.*, vol. 41, no. 3, Jul. 2009.
- [34] M. Braei and S. Wagner, “Anomaly detection in univariate time-series: A survey on the state-of-the-art,” 2020. [Online]. Available: arXiv:2004.00433
- [35] R. Chalapathy and S. Chawla, “Deep learning for anomaly detection: A survey,” 2019. [Online]. Available: arXiv:1901.03407
- [36] J. Nogas, S. S. Khan, and A. Mihailidis, “Deepfall: Non-invasive fall detection with deep spatio-temporal convolutional autoencoders,” *J. Healthc. Inform. Res.*, vol. 4, p. 50–70, 2020.
- [37] K. Makantasis *et al.*, “3d measures exploitation for a monocular semi-supervised fall detection system,” *Multimed. Tools Appl.*, vol. 75, p. 15017–15049, 2016.
- [38] C. Liu *et al.*, “Detection of human fall using floor vibration and multi-features semi-supervised svm,” *Sensors*, vol. 19, no. 17, pp. 3720–3740, 2019.
- [39] E. A. Kringle *et al.*, “Iterative processes: a review of semi-supervised machine learning in rehabilitation science,” *Disability Rehabil. Assistive Technol.*, vol. 15, no. 5, pp. 515–520, 2020.
- [40] G. Diraco, A. Leone, and P. Siciliano, “A fall detector based on ultra-wideband radar sensing,” in *Sensors*, B. Andò *et al.*, Eds. Cham: Springer International Publishing, 2018, pp. 373–382.
- [41] D. Charte *et al.*, “A practical tutorial on autoencoders for nonlinear feature fusion: Taxonomy, models, software and guidelines,” *Inf. Fusion*, vol. 44, pp. 78 – 96, 2018.
- [42] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” 2013. [Online]. Available: arXiv:1312.6114
- [43] O. Fabius and J. R. van Amersfoort, “Variational recurrent auto-encoders,” 2014. [Online]. Available: arXiv:1412.6581
- [44] J. Chung *et al.*, “A recurrent latent variable model for sequential data,” 2015. [Online]. Available: arXiv:1506.02216
- [45] D. Park, Y. Hoshi, and C. C. Kemp, “A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder,” *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 1544–1551, 2018.
- [46] *Operation of Radar Services in the 76-81 GHz Band*, Federal Communications Commission, Washington, D.C. [Online]. Available: <https://docs.fcc.gov/public/attachments/FCC-15-16A1.pdf>
- [47] *Google LLC Request for Waiver of Part 15 for Project Soli*, Federal Communications Commission, Washington, D.C. [Online]. Available: <https://docs.fcc.gov/public/attachments/DA-18-1308A1.pdf>
- [48] G. Hakobyan and B. Yang, “High-performance automotive radar: A review of signal processing algorithms and modulation schemes,” *IEEE Signal Process. Mag.*, vol. 36, no. 5, pp. 32–44, 2019.
- [49] S. Blackman, *Multiple-target Tracking with Radar Applications*, ser. Radar Library. Dedham, MA: Artech House, 1986, ch. 11.
- [50] C. M. Bishop, *Pattern recognition and machine learning*, ser. Information science and statistics. New York, NY: Springer, 2006, ch. 11.
- [51] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *J. Am. Stat. Assoc.*, vol. 112, no. 518, pp. 859–877, 2017.
- [52] S. Haykin, *Neural Networks and Learning Machines*, 3rd ed. Upper Saddle River, NJ: Pearson Higher Ed, 2011, ch. 4.
- [53] D. P. Kingma and M. Welling, “An introduction to variational autoencoders,” *Foundations and Trends in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019.
- [54] D. J. Rezende, S. Mohamed, and D. Wierstra, “Stochastic backpropagation and approximate inference in deep generative models,” in *Proc. 31st Int. Conf. Machine Learning (ICML)*, 2014, pp. II–1278–II–1286.
- [55] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [56] K. Cho *et al.*, “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2014, pp. 1724–1734.
- [57] Y. Bengio, P. Simard, and P. Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE Trans. Neural Netw.*, vol. 5, no. 2, pp. 157–166, 1994.
- [58] A. M. Dai and Q. V. Le, “Semi-supervised sequence learning,” in *Advances Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 3079–3087.
- [59] T. Kieu *et al.*, “Outlier detection for time series with recurrent autoencoder ensembles,” in *Proc. 28th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2019, pp. 2725–2732.
- [60] *xWR1843 Evaluation Module (xWR1843BOOST) Single-Chip mmWave Sensing Solution*, Texas Instruments, 2019. [Online]. Available: <http://www.ti.com/lit/ug/spruim4a/spruim4a.pdf>
- [61] (2020, Jan.) Overview of traffic monitoring for 18xx or 68xx. Texas Instruments. [Online]. Available: http://dev.ti.com/tirex/explore/node?node=A1v9f4cIpJmNkdgYh8SGsw_VLyFKFf_LATEST
- [62] *Programming Chirp Parameters in TI Radar Devices*, Texas Instruments, 2020. [Online]. Available: <http://www.ti.com/lit/an/swra553a/swra553a.pdf>
- [63] F. Adib *et al.*, “Capturing the human figure through a wall,” *ACM Trans. Graph.*, vol. 34, no. 6, Oct. 2015.




Feng Jin (S'19) received the B.S. and M.S. degrees in electrical engineering from Beihang University, Beijing, China, in 2011 and 2014, respectively. He is currently working toward the Ph.D. degree at University of Arizona, Tucson, Arizona, USA. His research interests include automotive radar signal processing and machine learning on mmWave radar sensor for object classification.



Arindam Sengupta (S'20) is a PhD student at the University of Arizona. He received his Bachelor of Engineering degree in ECE from Birla Institute of Technology in India, in 2012. He received his MS in ECE from the University of Akron in 2014 where his primary research was on multi-dimensional digital signal processing architectures. Under the supervision of Dr. Siyang Cao, he is currently researching methods to improve existing target tracking and localization methods using mmWave radars and applied machine learning.

His research interests include Signal Processing, Radars, Machine Learning and Autonomous Vehicles.



Siyang Cao (S'11-M'15-SM'20) received the Ph.D. degrees in Electrical Engineering from The Ohio State University in Columbus, Ohio in 2014. He received the B.S. degree in electronic and information engineering from Xidian University in Shanxi, China in 2007, and the M.S. degree in circuits and system from South China University of Technology in Guangdong, China in 2010. Since August 2015, he has been an Assistant Professor in the Electrical and Computer Engineering Department of The University of Arizona

in Tucson, AZ, USA. His research interests include radar signal processing, radar imaging, sensor fusion, electronically scanned radar system and machine learning with an emphasis on sensors.