

Anomaly Detection in Video Data Based on Probabilistic Latent Space Models

Giulia Slavic*, Damian Campo*, Mohamad Baydoun*, Pablo Marin†, David Martin†,
Lucio Marcenaro* and Carlo Regazzoni*

*DITEN, University of Genova, Italy

Contact email: slavic.giulia@gmail.com, damian.campo@edu.unige.it and mohamad.baydoun@edu.unige.it

†Intelligent systems lab, University Carlos III de Madrid, Spain

Contact email: pamarinp@ing.uc3m.es and dmgomez@ing.uc3m.es

Abstract—This paper proposes a method for detecting anomalies in video data. A Variational Autoencoder (VAE) is used for reducing the dimensionality of video frames, generating latent space information that is comparable to low-dimensional sensory data (e.g., positioning, steering angle), making feasible the development of a consistent multi-modal architecture for autonomous vehicles. An Adapted Markov Jump Particle Filter defined by discrete and continuous inference levels is employed to predict the following frames and detecting anomalies in new video sequences. Our method is evaluated on different video scenarios where a semi-autonomous vehicle performs a set of tasks in a closed environment.

Index Terms—Variational autoencoder, anomaly detection, particle filtering, Kalman filtering

I. INTRODUCTION

The detection of anomalies on video data is currently one of the most relevant topics in signal processing and computer science fields due to its multiple applications such as machine automation [1]–[5], estimation of future instances [6]–[8] and improvements in surveillance systems [9]–[11]. Currently, video analysis and computer vision are important fields that attract large research and industrial interest. Furthermore, the automatic detection of anomalies in video information is a key element for generating robust autonomous systems that can adapt themselves to unknown situations/experiences and incrementally learn predictive models from them.

Over the last years, different deep learning algorithms have demonstrated their capabilities for solving several problems, e.g., image classification with a human-like performance [12]–[14]. Since we are moving closer and closer towards the ultimate aim of human-like vision for machines [15], more complex problems such as recognizing contextual information incrementally and understanding/adapt to new scenarios are current challenges to be automatically solved by machines. For generating autonomous systems, it is fundamental to provide machines with models that can handle the dynamic properties of real-world situations, e.g., dealing with streams of data that have not been seen before and including uncertainty in contextual representations and model predictions.

Motivated by the necessity of generating *active* artificial agents that are able to interact with the real-world in often uncontrolled or detrimental conditions [16], this paper proposes a method for detecting anomalies in video data that facilitates

the identification of unusual situations (previously unseen data) as they are experienced. Our method uses a probabilistic structure that facilitates the potential insertion/learning of new models as they are detected as abnormal, which would undoubtedly increase the adaptability of autonomous systems to unknown scenarios by continuously evolving their current models through the incorporation of new concepts.

As is well known, in signal processing and surveillance systems, the detection of anomalies is an essential topic of large research [11], [17]–[21] and commercial interest [22]–[24]. Nonetheless, in computer vision, due to the large dimensionality of images, few attempts for recognizing and understanding anomalies have been made. Accordingly, existing methods based on deep learning for detecting anomalies in video sequences [8], [25]–[27] usually do not allow to perform learning of models in a probabilistic framework that can be compatible with models relying on low dimensional information, e.g., 2-Dimensional position data. Novel research has tried to include the concept of *lifelong learning* to deep neural networks (DNNs) [28], which aim at allowing DNNs to acquire, fine-tune, and transfer knowledge through time continually. Nonetheless, advances in DNN lifelong learning remain in large dimensional representations that do not contemplate the possibility of including multisensory data for making robust models that could rely on either low or high dimensional information.

The proposed method is based on hierarchical probabilistic models that facilitate inferring future instances of video sequences and detecting anomalies. A similar approach was used for making inferences in low dimensional data, namely positional [29] and control [30] information coming from a moving vehicle. Accordingly, this paper proposes a method for inferring/estimating video sequences similarly as [29] and [30] do through a Markov Jump Particle Filter (MJPF).

Our method is based on the latent representation of a Variational Autoencoder (VAE), which is employed for obtaining a low-dimensional state of the video at each time instant in a probabilistic fashion. Clusters of similar latent spaces are identified, which facilitates obtaining a semantic representation of video states. In each cluster, a fully connected neural network (NN) is employed to learn a non-linear dynamical model that allows estimating a future image given the content

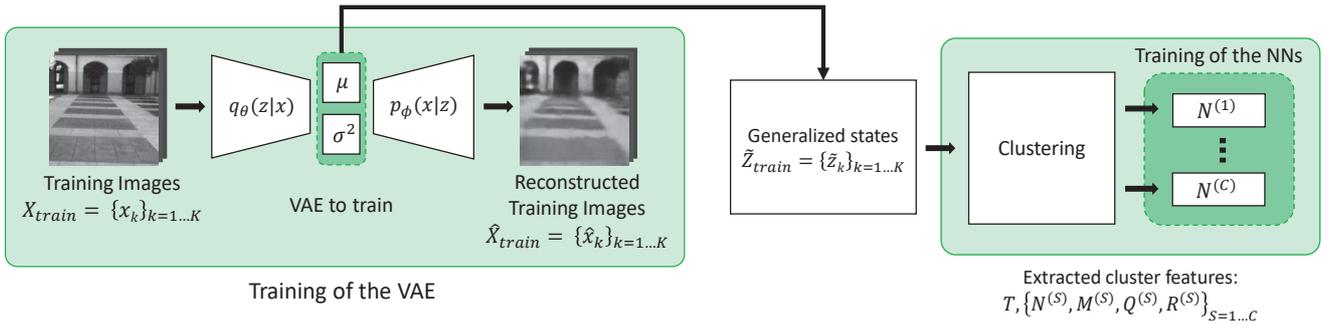


Fig. 1: Training phase: the VAE is trained to reconstruct a set of training images. The bottleneck features are then extracted and the GSs are derived and used for clustering. A Neural Network is trained for each cluster.

of the current one. Accordingly, two representation levels, i.e., discrete and continuous, are learned for making inferences in video sequences and detecting anomalies. For testing learned models, a particle filter coupled with a set of Unscented Kalman Filters (UKFs) [31] are employed for predicting at discrete and continuous levels, respectively, and detecting anomalies.

As mentioned before, the proposed method is based on research previously done on low-dimensional data [29], [30]. Nonetheless, the novel contributions of this paper are: *i)* A full probabilistic method that represents video sequences into latent spaces so that video predictions can be made at continuous and discrete hierarchical levels. *ii)* A method that models non-linearly the dynamics of latent space information through a set of NNs. *iii)* Finally, an approach that is compatible with any low and large dimensional data, which allows us to detect anomalies in multisensory data and potentially use them for incremental learning.

The rest of the paper is structured as follows: Section II presents the proposed method for detecting anomalies in video sequences. Section III introduces the dataset employed to evaluate our method. Section IV shows and discusses obtained results. Finally, Section V concludes the paper and presents some insights about possible future work.

II. PROPOSED METHOD

The proposed method is composed of two basic steps, namely *training* and *testing phases*; during the former (section II-A and Fig. 1), algorithms are trained based on observed data; whereas the latter (section II-B and Fig. 2) uses learned algorithms to detect anomalies on new data.

A. Training phase

1) *Variational Autoencoder*: As a first step of the training phase, a VAE is used for describing the images in a latent space that has significantly reduced dimension with respect to the original image size. A VAE has been used instead of a normal Autoencoder because the former facilitates to represent images in the latent state probabilistically by using a mean μ and variance σ^2 to approximate each latent variable. This enables probabilistic reasoning and inference.

It is well known that a VAE is composed of two parts: an encoder $q_\theta(z|x)$ and a decoder $p_\phi(x|z)$. The latent state z sampled from $\mathcal{N}(\mu, \sigma^2)$, returns an approximate reconstruction of the observation x . Through θ and ϕ , we define the parameters of the encoder and decoder, respectively. In order to optimize them, the VAE maximizes the sum of the lower bound on the marginal likelihood of each observation x of the dataset D , as described in [32], [33]:

$$\mathcal{L}_{\phi, \theta}(D) = \sum_{x \in D} \mathcal{L}_{\phi, \theta}(x), \quad (1)$$

being $\mathcal{L}(\theta, \phi; x)$ defined as:

$$\begin{aligned} \mathcal{L}_{\phi, \theta}(x) = & -D_{KL}(q_\theta(z|x) || p_\phi(z)) + \\ & + E_{q_\theta(z|x)}[\log p_\phi(x|z)], \end{aligned} \quad (2)$$

where D_{KL} defines the Kullback-Leibler divergence. Therefore, the first term measures the difference between the encoder's distribution $q_\theta(z|x)$ and the prior $p_\phi(z)$; being the prior typically a standard normal distribution $\mathcal{N}(0, 1)$. The second term is the expected log-likelihood of the observation x and forces the VAE to reconstruct the input data.

This work uses the ability of the VAE to encode the input information in a significant lower-dimensional space that exhibits probabilistic properties exploitable to detect anomalies at the latent feature level. Additionally, as we are interested in producing predictive models that can work for multisensory data regardless of their dimensionality, the VAE turns out to be an excellent choice for representing and treating video sequences as small-dimensional data, enabling a more homogeneous way of making algorithms for data fusion with multimodal information.

We first train the VAE by using a set of training images X_{train} . Then, we input again X_{train} to the VAE and obtain a set of latent features described by μ_{train} and σ_{train}^2 .

2) *Generalized states*: Starting from the set of μ_{train} , considered as the state of training images, we build a set of Generalized States (GSs) containing several time-order derivatives. This work only uses the first time-order derivative since no abrupt dynamics are considered.

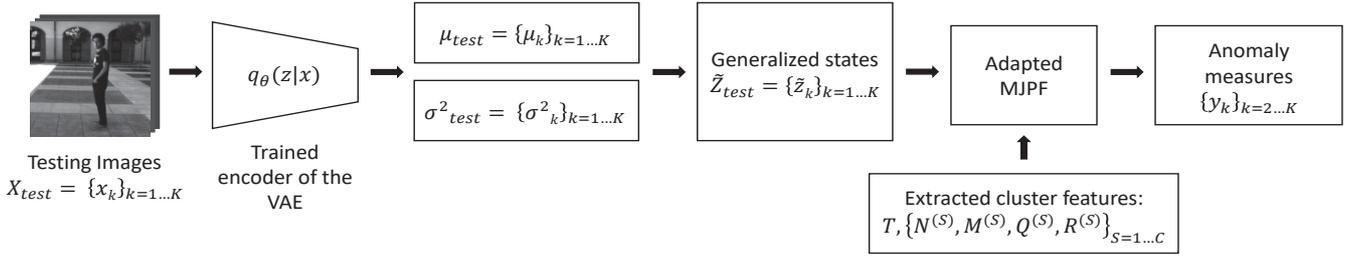


Fig. 2: Testing phase: the encoder of the VAE is used to extract the bottleneck features of the testing images. The GSs are derived and given as input to the Adapted MJPF, which detects anomalies by using information coming from clusters.

Let μ_k be the value of μ for the image x_k at time k , its first time-order derivative can be approximated by $\dot{\mu}_k \sim \frac{\mu_k - \mu_{k-1}}{\Delta k}$, where $\Delta k = 1$, which assumes a normalized regular sampling of images. The GS at time k can thus be written as $\tilde{z}_k = [\mu_k \ \dot{\mu}_k]^\top$. Repeating this for each consecutive couple of training images, we obtain a set of GSs for the training set, defined by:

$$\tilde{Z}_{train} = [\mu_{train} \ \dot{\mu}_{train}]^\top. \quad (3)$$

3) *Clustering and neural networks*: After obtaining GSs related to training video sequences, we use a traditional k-means algorithm to cluster GSs into groups that carry similar information. Since we use the values of μ and $\dot{\mu}$ to perform the clustering process, obtained clusters take into consideration the encoded image and also its dynamics w.r.t. the next frame. This facilitates recognizing and cluster different ways of moving, e.g., the vehicle crossing the same zone at different speeds.

Once the clustering is performed, a transition matrix T encodes the transition probabilities from each cluster to the others. Additionally, the following features are extracted from each cluster indexed as S : *i*) cluster's centroid $M^{(S)}$, *ii*) cluster's covariance $Q^{(S)}$ and *iii*) cluster's radius of acceptance $R^{(S)}$. Finally, a fully connected neural network $N^{(S)}$ defining the dynamics of GSs, i.e., continuous predictive model, is learned for each cluster. Assuming that a total number of C clusters have been identified, it is possible to write $S = \{1, \dots, C\}$. For training each $N^{(S)}$, the value of every μ_k is taken as input and the corresponding $\dot{\mu}_{k+1}$ as output, where $[\mu_k, \dot{\mu}_k]^\top \in S$. Moreover, to include the uncertainty of the Gaussian latent spaces encoded in σ^2 , $2L$ additional inputs and outputs are used, where L is the dimension of the latent state. Such $2L$ points, together with the initial mean $\mu_k^0 = \mu_k$, permit to completely capture and define the Gaussian $\mathcal{N}(\mu_k, \sigma_k^2)$, as described in [31]:

$$\begin{aligned} \mu_k^i &= \mu_k + (\sqrt{(L+\lambda)\Sigma_k})_i & \mathbf{if} \ i &= 1 \dots L \\ \mu_k^i &= \mu_k - (\sqrt{(L+\lambda)\Sigma_k})_{i-L} & \mathbf{if} \ i &= L+1 \dots 2L, \end{aligned} \quad (4)$$

where λ is a scaling parameter and $\Sigma_k \sim I_L \sigma_k^2$, being I_L the identity matrix of dimension L .

The μ_k^i values calculated in Eq. (4) are the so-called sigma points associated with (μ_k, σ_k^2) . A corresponding group of sigma points can be defined in the same way for $(\mu_{k+1}, \sigma_{k+1}^2)$ and each value of $\mu_{k+1}^i - \mu_k^i$ is given as additional output for the training of the NNs.

To summarize, each $N^{(S)}$ performs the following approximation:

$$\dot{\mu}_{k+1}^i \sim N^{(S)}(\mu_k^i) + w_k^i,$$

where $\dot{\mu}_{k+1}^i$ and μ_k^i are calculated based on $[\mu_k, \dot{\mu}_k]^\top \in S$ and w_k^i is the residual error after the convergence of the network.

Each $N^{(S)}$ learns a sort of *quasi-semantic* information based on a particular image appearance and motion detected by the cluster S , facilitating the estimation of future latent spaces, i.e., predicting following frames. Such a feature can be employed to detect whether new observations are similar to previously learned situations encoded in the set of NNs. In case predictions from NNs are not compliant with observations, an anomaly should be detected, and models should be adapted to learn new situations, generating new semantic information.

B. Testing Phase

During the testing phase, each testing image is processed by the VAE, and GSs are calculated. Then, an adapted version of the MJPF is used to detect anomalies in video sequences.

1) *Adapted Markov Jump Particle Filter*: A MJPF can be described as a Probabilistic Switching Graphical Model [34], [35] for prediction and anomaly detection purposes by using a bank of Kalman Filters (KFs) at the continuous state level and a Particle Filter at the discrete state level [29]. This work tackles a problem that requires a non-linear model for prediction purposes and a non-linear observation model, solved by a set of NNs (each of them associated with a detected cluster) and a VAE, respectively. Accordingly, a bank of standard KFs at the state level cannot be used due to the nonlinearities described above. This work uses a bank of modified KFs whose predictions follow the same logic of the UKFs and employ the encoded information of the VAE for updating purposes.

The DBN associated with the Adapted MJPF (A-MJPF) is showed in Fig. 3. A detailed description of the MJPF can be found in [29]. This paper will only provide a brief description

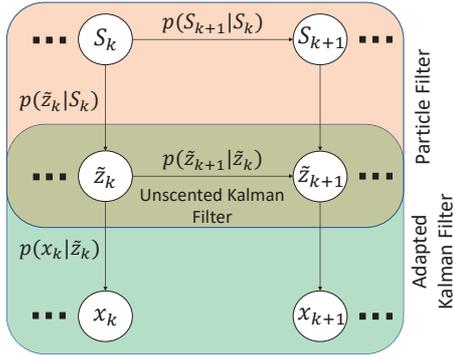


Fig. 3: DBN associated with the A-MJPF.

of the overall logic of the employed probabilistic architecture, which is identical to the one used in [29]. Instead, it will focus on the parts that have been modified in the A-MJPF w.r.t. the standard MJPF.

For both architectures (MPJF and A-MPJF), at each time instant k , two stages are performed: *prediction* and *update*. During prediction, the next cluster S_{k+1} (discrete level) and the next GS \tilde{z}_{k+1} (continuous level) are estimated for each particle, i.e., $p(S_{k+1}|S_k)$ and $p(\tilde{z}_{k+1}|\tilde{z}_k)$, respectively. Similarly to the standard MPJF, predictions at the discrete level in the A-MJPF are performed by using the transition matrix T in each particle. On the other hand, predictions at the continuous level in the A-MPJF are performed by the neural network $N^{(S_k)}$ associated with the selected discrete state S_k . Since non-linear models are considered for predicting continuous level information, an UKF performs the estimations as described in [31] by taking $2L$ additional sigma points as already done in section II-A3. Therefore, the prediction for each sigma point follows the equation below:

$$\tilde{z}_{k+1}^i = f(\tilde{z}_k^i) = A\tilde{z}_k^i + BN^{(S)}(\mu_k^i) + w_k^i, \quad (5)$$

where A and B are two matrices used to map the previous state \tilde{z}_k^i and the new velocity computed by $N^{(S)}(\mu_k^i)$ on the new state \tilde{z}_{k+1}^i , such that as $A = [A_1 A_2]$ with $A_1 = [I_L 0_{L,L}]^\top$, $A_2 = 0_{2L,L}$ and as $B = [I_L I_L]^\top$. The mean and covariance of the predicted state are then calculated using the UKF formulas for the propagation of a Gaussian random variable through a non-linear model.

The update phase is performed when a new measurement (image) is observed. At the discrete level, particles are re-sampled based on a measure of the anomaly (see section II-B2). At the state level, a modified version of the KF update is performed. This update takes into consideration the fact that μ_k and σ_k^2 from the VAE for image x_k can be used as the mapped observation on the state space at time k . Consistently, σ_k^2 can approximate the covariance matrix, such that $\Sigma_k \sim I_L \sigma_k^2$, representing the uncertainty while encoding images. Algorithm 1 describes the employed KF's steps.

2) *Anomaly measurement*: After the update phase, at each time instant k , the predicted value of $\mu_k^{l,p}$ related to latent

Algorithm 1 Equations of the prediction and update phase of the Adapted Kalman Filter.

▷ **PREDICTION**:

- 1: Calculation of the sigma points $\tilde{z}_{k|k}^i$ and of their respective weights $\tilde{W}^{i,m}$ and $\tilde{W}^{i,c}$ as described in [31].
- 2: $\tilde{z}_{k+1|k}^i = f(\tilde{z}_{k|k}^i)$
- 3: $\tilde{z}_{k+1|k} = \sum_{i=0}^{2L} \tilde{W}^{i,m} \tilde{z}_{k+1|k}^i$
- 4: $P_{k+1|k} = \sum_{i=0}^{2L} \tilde{W}^{i,c} \{\tilde{z}_{k+1|k}^i - \tilde{z}_{k+1|k}\} \{\tilde{z}_{k+1|k}^i - \tilde{z}_{k+1|k}\}^\top$
- 5: $P_{k+1|k}^L = P_{k+1|k} \Big|_{\{row:1\dots L, col:1\dots L\}}$

▷ **UPDATE**:

- 6: $K_{k+1} = [P_{k+1|k}^L; I_L] (P_{k+1|k}^L + \Sigma_{k+1})^{-1}$
- 7: $\tilde{z}_{k+1|k+1} = \tilde{z}_{k+1|k} + K_{k+1} (\mu_{k+1} - \mu_{k+1|k})$
- 8: $P_{k+1|k+1} = P_{k+1|k} - K_{k+1} (P_{k+1|k}^L + \Sigma_{k+1}) K_{k+1}^\top$

state component l and particle p is compared with the actual updated value, outputting a measure of innovation defined as:

$$y_k = \min_p \frac{\sum_{l=1}^L |\mu_{k|k}^{l,p} - \mu_{k|k-1}^{l,p}|}{L}. \quad (6)$$

The anomaly values of training video sequences are used to set an anomaly threshold defined as:

$$thresh = \bar{y}_{train} + 3std(y_{train}), \quad (7)$$

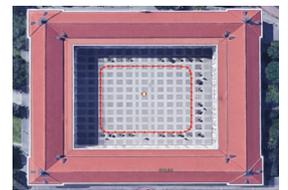
being \bar{y}_{train} and $std(y_{train})$ the mean value and standard deviation of anomalies from the training data, respectively. When using algorithm 1 on testing data, video sequences that produce anomaly signals above the threshold in Eq.(7) are considered as potential anomalies. Moreover, to filter out spurious peaks, also a temporal criterion based on a window of 3 frames is used, such that anomaly signals that are above the threshold $thresh$ but last less than 3 frames are not considered as actual anomalies.

III. EMPLOYED DATASET

A real vehicle called “iCab” [36], see Fig. 4a, is used to collect the video dataset. A human drives the iCab performing different tasks in a closed environment displayed in Fig. 4b. The proposed dataset was captured from an onboard front camera while the vehicle executes 4 different tasks.



(a) Autonomous vehicle “iCab”



(b) Closed environment “iCab”

We aim at detecting dynamics that have not been seen previously in a normal situation (Scenario I), which is used for learning purposes. Scenarios II, III and IV include unseen

maneuvers caused by the presence of pedestrians while the vehicle performs a perimeter monitoring task. Accordingly, four scenarios (see Fig. 5) are considered in this work:

Scenario I (perimeter monitoring): the vehicle follows a rectangular trajectory along with a closed building. The temporal evolution of the perimeter monitoring maneuver from a first-person perspective is shown in Fig. 6.

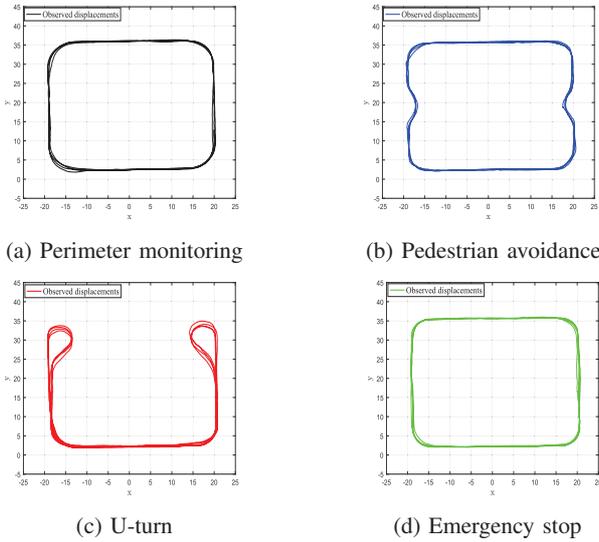


Fig. 5: Vehicle tasks used to evaluate the proposed method. Perimeter monitoring is utilized in the training phase whereas the other three tasks are employed for testing purposes

Scenario II (emergency stop maneuver): the vehicle executes the perimeter monitoring task and encounters two pedestrians crossing its path at each lap. The vehicle performs an emergency stop and then continues the perimeter monitoring task when pedestrians exit its field of view. A vehicle’s first-person perspective of the temporal evolution of the stop maneuver is provided in Fig. 7.

Scenario III (pedestrian avoidance maneuver): two obstacles (stationary pedestrians) in different locations interfere with the perimeter monitoring task of Scenario I. The vehicle performs an avoidance maneuver and continues the perimeter monitoring. Fig. 8 shows the temporal evolution of the avoidance maneuver from a first-person perspective.

Scenario IV (U-turn maneuver): while the vehicle executes a perimeter monitoring, it encounters two static pedestrians located in different locations. In this scenario, the vehicle performs a U-turn motion and then continues the perimeter monitoring in the opposite direction w.r.t. training data. Fig. 9 shows the U-turn maneuver from the front camera viewpoint.

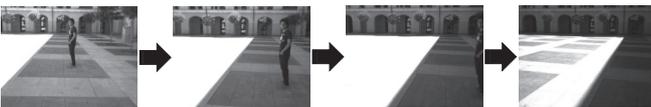


Fig. 8: Scenario III (avoidance maneuver).

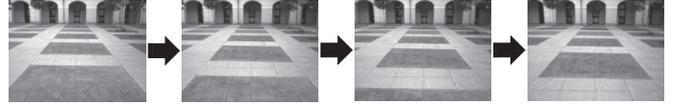


Fig. 6: Scenario I (perimeter monitoring).

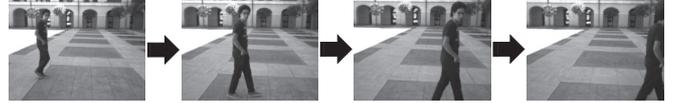


Fig. 7: Scenario II (stop maneuver).

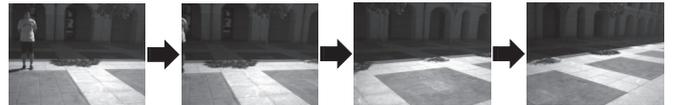


Fig. 9: Scenario IV (U-turn maneuver).

IV. EXPERIMENTAL RESULTS

A. Definition of normality

As a first step, it is necessary to train our model to learn normal patterns. The Perimeter Monitoring frames are therefore used as X_{train} data to perform the training phase described in section II-A. Consequently, a VAE is trained (II-A1), from which a set of clusters and their corresponding NNs (II-A3) are learned based on frames where the iCab moves in a *pedestrian-less* environment. We considered several cluster experiments where the total number of clusters C was varied. Consequently, we selected the case of $C = 6$ empirically as it describes the normal scene accurately with a relatively low number of clusters. Video frames are clustered based on images’ appearance and dynamics (changes in consecutive frames) at each time instant.

The threshold described in section II-B2 is then obtained by performing a testing procedure using the training data. This threshold will be used in the actual testing phases described in the following sections with the objective of detecting abnormal behaviors in new scenarios.

B. Emergency stop maneuver

In this modality, the vehicle (artificial agent) performs an emergency stop maneuver that allows a pedestrian to cross in front of it. Fig. 10a displays the anomaly signal obtained from a single vehicle’s lap around the courtyard where it encounters/interacts with two pedestrians. The signal is normalized based on the threshold calculated from the perimeter monitoring experience; see Eq.(7). Consistently, when the anomaly signal goes above the threshold (displayed as a dotted black line in Fig. 10(a)), a potential abnormal situation is detected. The blue color background indicates the presence of the pedestrian, such as from the moment it enters into the camera’s field of view, see (1) and (4); until the instant where it leaves it, see (3) and (6). The anomaly signal grows rapidly as soon as the pedestrian completely enters the field of view of the camera, see (2) and (5).

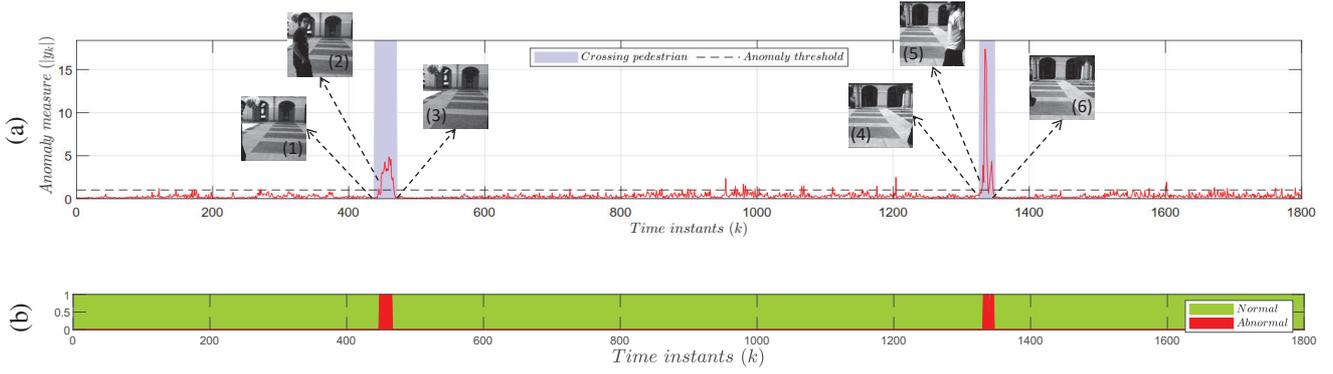


Fig. 10: Testing phase on the emergency stop task: (a) Anomaly signal. (b) Color-coded final anomaly.

Fig. 10(b) displays the final color-coded anomalies w.r.t. the perimeter monitoring task (training data), normal and abnormal frames are colored in green and red, respectively. It can be seen how the proposed method enables the detection of anomalies due to moving pedestrians that have not seen before in the training data.

C. Pedestrian avoidance

In this modality, the vehicle avoids a static pedestrian. Fig. 11a shows the resulting anomaly signal. The blue zones refer to video frames that contain the static pedestrian, and yellow zones encode the avoidance maneuvers. As can be seen from Fig. 5b, at each lap, the vehicle encounters two different static pedestrians in the environment. They wear t-shirts of different colors (black in the first case and white in the second one), which make them “camouflage” with the environment in some particular configurations due to changeable illumination conditions. This factor influences the different anomaly values for both pedestrians, with the second one generating a higher anomaly.

At each pedestrian encounter, the anomaly signal is composed of three zones with high values: a first one due to the pedestrian presence beginning, see (1) and (5); and other two due to the avoidance maneuver, see (2)-(4) and (6)-(8).

Between the latter two peaks, a zone with a low anomaly is present, see (3), or (6); this is due to the execution of similar behaviors already observed in the training set.

D. Pedestrian avoidance through U-turn

In this modality, the vehicle avoids a static pedestrian by performing a U-turn maneuver. In this case, the vehicle will move in the opposite direction w.r.t. its previous motion towards the pedestrian, see Fig. 5c, which introduces later on new situations, e.g., curving in the opposite direction, and other already known behaviors, e.g., moving straight in regions that present similar symmetries to those in the training set. Additionally, the U-turn presents some cases containing both normal and abnormal information, e.g., moving straight in regions that are similar to video sequences already seen but showing some differences related to structures that are illuminated differently in the courtyard.

Fig. 12a shows the anomaly signal for this case. Three regions have been highlighted describing the main abnormal situations appearing in this task: *i*) the presence of pedestrians colored in blue, *ii*) the U-turn maneuver represented in yellow and *iii*) the curves performed in the opposite direction w.r.t. the training set, which is coded in purple. Regarding the first situation, it must be noted that the blue zone is again related

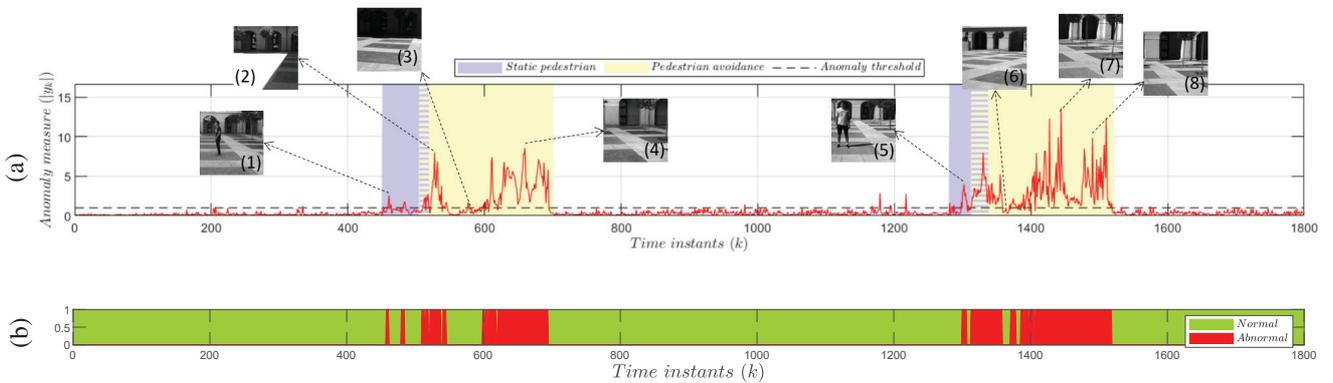


Fig. 11: Testing phase on the pedestrian avoidance task: (a) Anomaly signal. (b) Color-coded final anomaly.

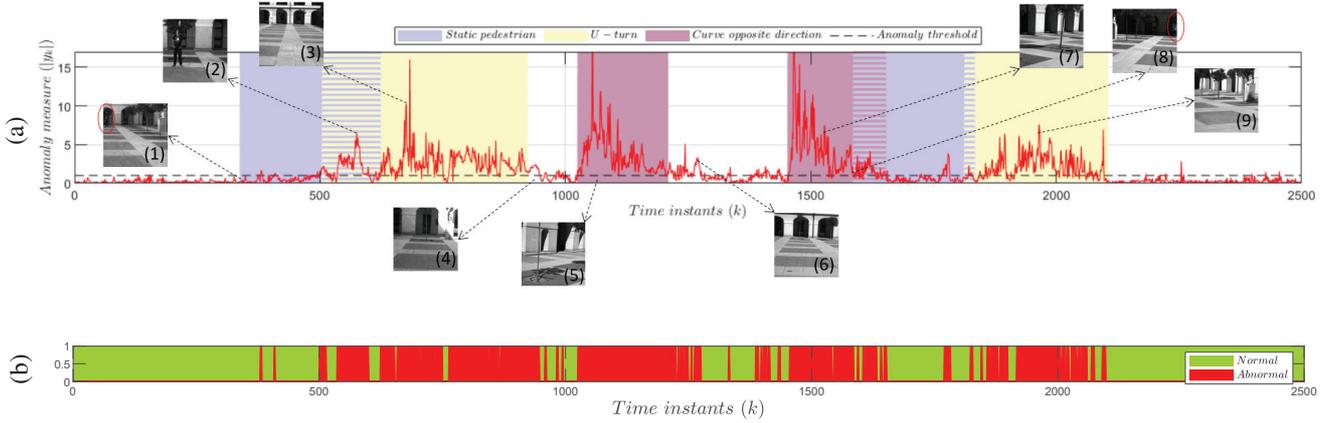


Fig. 12: Testing phase on pedestrian avoidance through U-turn task: (a) Anomaly signal. (b) Color-coded final anomaly.

to the entire sequence of frames in which the pedestrian is present. In this task, as the pedestrian first appears on the opposite corner of the courtyard w.r.t. where the vehicle is moving, it occupies only a small amount of pixels, which makes its recognition difficult even for the human eye; see (1) and (8) in Fig. 12a. Moreover, as pedestrians in this task are located on a background that has a similar color shade to their clothing, even when the vehicle is moving closer, low anomaly levels are detected. In (2), the anomaly is particularly high because the vehicle starts moving right to perform the U-turn. The other two anomaly zones regarding the U-turn maneuver, see (3) and (9), and the curves in the opposite direction (5) and (7) are both detected effectively by the proposed method as anomalies.

In addition to the three anomaly cases mentioned above, other smaller anomaly peaks can be observed in zones where the vehicle moves straight but in the opposite direction to what is experienced in the training set. Those anomalies due to different causes: In (6), they are caused by differences in the background (e.g., the right side is abnormal due to the presence of a tree and the lack of shadows). In (4), anomalies are found due to differences in the background and velocities (e.g., the vehicle is in the middle of the courtyard, where it is expected the maximum velocity based on the training dataset, but due to the U-turn maneuver, the vehicle decelerates at that section of the courtyard).

E. Results Discussion

It can be observed how the proposed method is able to determine whether a situation is normal or abnormal, with a high accuracy level in different scenarios. The detection of anomalies was particularly good in three cases: *i)* visual data that differs substantially from training samples, e.g., when performing a curve in the opposite direction w.r.t. the training video sequences or during the central part of the U-turn movement. *ii)* new image motions related to visual information that has already been seen, e.g., the beginning of the U-turn maneuver (due to the changing of direction) and after exiting from it (due to previously unseen velocity changes). *iii)* the

presence of moving or static objects, e.g., pedestrians, that are not considerably far away from the camera. On the other hand, the accuracy is much lower in cases where anomalies are not completely defined, e.g., pedestrians placed far away from the camera or that “camouflage” with the background.

We have therefore observed how our method successfully recognizes anomalies based on the appearance (new images) and the dynamics (abnormal motions) of video data.

V. CONCLUSION AND FUTURE WORK

This paper presented a method for the detection of anomalies from visual data. By using a VAE, we were able to bring high-dimensional data acquired from a camera to a low-dimensionality that is compatible with other sensor data acquired from the vehicle and already examined in previous contributions (e.g., position, steering angle). An A-MJPF has been introduced and used to detect anomalies both at the observation and prediction levels.

This work aims at generating a predictive model for video sequences. Our method can be inserted in a multi-modal architecture for autonomous vehicles. Future work is oriented to examine the incremental learning of the new scenarios and adaptation to them. Anomaly detection constitutes a fundamental aspect for this: when abnormal situations are detected, the corresponding input images can be used for building/refining new and already existing models. As discussed in the paper, two main abnormal situations can be distinguished: *i)* Video sequences containing images that were never seen before; *ii)* Video sequences containing known images but new video dynamics/motions. This paper mainly focused on the first case, as also done by [8], [25]. In the future, we will consider both cases by exploiting the incremental learning of new situations by taking abnormal data and use it for building/refining predictive models.

The insertion of observations coming from other sensory data, e.g., positional and control information, into the proposed probabilistic framework constitutes another future path for our work. By allowing algorithms to handle multi-modal data when making predictions and detecting anomalies, it is

possible to generate more realistic algorithms that associate multiple heterogeneous signals to familiar concepts as humans and other animals do. Such will certainly enable a more robust inference process, which leads to more efficient decision-making in autonomous systems.

REFERENCES

- [1] G. I. Parisi and S. Wermter, "A neurocognitive robot assistant for robust event detection," in *Trends in Ambient Intelligent Systems*. Springer, 2016, pp. 1–27.
- [2] F. Rezazadegan, S. Shirazi, B. Uproft, and M. Milford, "Action recognition: From static datasets to moving robots," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 3185–3191.
- [3] H. Lu, Y. Li, S. Mu, D. Wang, H. Kim, and S. Serikawa, "Motor anomaly detection for unmanned aerial vehicles using reinforcement learning," *IEEE Internet of Things Journal*, vol. 5, no. 4, pp. 2315–2322, Aug 2018.
- [4] M.-K. Ko, "Autonomous vehicle driving support system and autonomous driving method performed by the same," Jun. 18 2015, uS Patent App. 14/132,249.
- [5] D. Feng, L. Rosenbaum, and K. Dietmayer, "Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, Nov 2018, pp. 3266–3273.
- [6] A. R. M. Forkan, I. Khalil, Z. Tari, S. Foufou, and A. Bouras, "A context-aware approach for long-term behavioural change detection and abnormality prediction in ambient assisted living," *Pattern Recognition*, vol. 48, no. 3, pp. 628–641, 2015.
- [7] A. Bera, S. Kim, and D. Manocha, "Realtime anomaly detection using trajectory-level crowd behavior learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016, pp. 50–57.
- [8] M. Ravanbakhsh, M. Baydoun, D. Campo, P. Marin, D. Martin, L. Marcenaro, and C. S. Regazzoni, "Learning multi-modal self-awareness models for autonomous vehicles from human driving," in *2018 21st International Conference on Information Fusion (FUSION)*. IEEE, 2018, pp. 1866–1873.
- [9] S. Coşar, G. Donatiello, V. Bogorny, C. Garate, L. O. Alvares, and F. Brémond, "Toward abnormal trajectory and event detection in video surveillance," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 27, no. 3, pp. 683–695, 2016.
- [10] V. Bastani, L. Marcenaro, and C. S. Regazzoni, "Online nonparametric bayesian activity mining and analysis from surveillance video," *IEEE Transactions on Image Processing*, vol. 25, no. 5, pp. 2089–2102, 2016.
- [11] M. Ravanbakhsh, E. Sangineto, M. Nabi, and N. Sebe, "Training adversarial discriminators for cross-channel abnormal event detection in crowds," in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 1896–1904.
- [12] B. M. Lake, R. R. Salakhutdinov, and J. Tenenbaum, "One-shot learning by inverting a compositional causal process," in *Advances in neural information processing systems*, 2013, pp. 2526–2534.
- [13] S. Dodge and L. Karam, "A study and comparison of human and deep learning recognition performance under visual distortions," in *2017 26th international conference on computer communication and networks (ICCCN)*. IEEE, 2017, pp. 1–7.
- [14] R. Geirhos, D. H. Janssen, H. H. Schütt, J. Rauber, M. Bethge, and F. A. Wichmann, "Comparing deep neural networks against humans: object recognition when the signal gets weaker," *arXiv preprint arXiv:1706.06969*, 2017.
- [15] U. Nadeem, S. Shah, F. Sohel, R. Togneri, and M. Bennamoun, "Deep learning for scene understanding," *Smart Innovation, Systems and Technologies*, vol. 136, pp. 21–51, 2019.
- [16] N. Sünderhauf, O. Brock, W. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Uproft, P. Abbeel, W. Burgard, M. Milford *et al.*, "The limits and potentials of deep learning for robotics," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 405–420, 2018.
- [17] W. Lin, Y. Zhang, J. Lu, B. Zhou, J. Wang, and Y. Zhou, "Summarizing surveillance videos with local-patch-learning-based abnormality detection, blob sequence optimization, and type-based synopsis," *Neurocomputing*, vol. 155, pp. 84–98, 2015.
- [18] T. Wang, M. Qiao, Y. Deng, Y. Zhou, H. Wang, Q. Lyu, and H. Snoussi, "Abnormal event detection based on analysis of movement information of video sequence," *Optik*, vol. 152, pp. 50–60, 2018.
- [19] A. B. Mabrouk and E. Zagrouba, "Abnormal behavior recognition for intelligent video surveillance systems: A review," *Expert Systems with Applications*, vol. 91, pp. 480–491, 2018.
- [20] J. Sun, J. Shao, and C. He, "Abnormal event detection for video surveillance using deep one-class learning," *Multimedia Tools and Applications*, vol. 78, no. 3, pp. 3633–3647, 2019.
- [21] M. Ravanbakhsh, M. Baydoun, D. Campo, P. Marin, L. Marcenaro, D. Martin, and C. Regazzoni, "Learning self-awareness for autonomous vehicles: Exploring multisensory incremental models," *Transactions on Intelligent Transportation Systems*, 2020.
- [22] D. G. Aviv, "Abnormality detection and surveillance system," Sep. 13 2011, uS Patent App. 12/466,340.
- [23] N. Kiryati, T. Riklin-Raviv, Y. Ivanchenko, S. Rochel, I. Dvir, and D. Harari, "Apparatus and methods for the detection of abnormal motion in a video stream," May 13 2014, uS Patent 8,724,891.
- [24] S. Ishikawa and Y. Umemoto, "Abnormality detection program, abnormality detection apparatus and abnormality detection method," Oct. 4 2018, uS Patent App. 15/899,433.
- [25] M. Ravanbakhsh, M. Baydoun, D. Campo, P. Marin, D. Martin, L. Marcenaro, and C. S. Regazzoni, "Hierarchy of gans for learning embodied self-awareness model," in *2018 25th IEEE International Conference on Image Processing (ICIP)*, Oct 2018, pp. 1987–1991.
- [26] H. Wei, Y. Xiao, R. Li, and X. Liu, "Crowd abnormal detection using two-stream fully convolutional neural networks," in *2018 10th International Conference on Measuring Technology and Mechatronics Automation (ICMTMA)*, Feb 2018, pp. 332–336.
- [27] N. Li and F. Chang, "Video anomaly detection and localization via multivariate gaussian fully convolution adversarial autoencoder," *Neurocomputing*, vol. 369, pp. 92–105, 2019.
- [28] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Networks*, 2019.
- [29] M. Baydoun, D. Campo, V. Sanguineti, L. Marcenaro, A. Cavallaro, and C. Regazzoni, "Learning switching models for abnormality detection for autonomous driving," in *2018 21st International Conference on Information Fusion (FUSION)*, July 2018, pp. 2606–2613.
- [30] D. Kanapram, D. Campo, M. Baydoun, L. Marcenaro, E. L. Bodanese, C. Regazzoni, and M. Marchese, "Dynamic bayesian approach for decision-making in ego-things," in *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*. IEEE, 2019, pp. 909–914.
- [31] E. Wan and R. van der Merwe, "The unscented kalman filter for nonlinear estimation," in *Symposium on Adaptive Systems for Signal Processing, Communication and Control*. IEEE, 2000.
- [32] D. P. Kingma and M. Welling, "Auto-encoding variational bayes." in *ICLR*, Y. Bengio and Y. LeCun, Eds., 2014.
- [33] —, "An introduction to variational autoencoders," *Foundations and Trends in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019.
- [34] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, ser. Adaptive computation and machine learning. MIT Press, 2009. [Online]. Available: <https://books.google.it/books?id=7dzhCHzNQ4C>
- [35] L. Sucar, *Probabilistic Graphical Models: Principles and Applications*, ser. Advances in Computer Vision and Pattern Recognition. Springer London, 2015.
- [36] P. Marin-Plaza, J. Beltrán, A. Hussein, B. Musleh, D. Martin, A. de la Escalera, and J. M. Armingol, "Stereo vision-based local occupancy grid map for autonomous navigation in ros," *Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISI-GRAPP)*, 2016.