



HAL
open science

Improving Chairlift Security with Deep Learning

Kevin Bascol, Rémi Emonet, Elisa Fromont, Raluca Debusschere

► **To cite this version:**

Kevin Bascol, Rémi Emonet, Elisa Fromont, Raluca Debusschere. Improving Chairlift Security with Deep Learning. International Symposium on Intelligent Data Analysis (IDA 2017), Oct 2017, Londres, United Kingdom. hal-01581392

HAL Id: hal-01581392

<https://hal.science/hal-01581392v1>

Submitted on 6 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improving Chairlift Security with Deep Learning

Kevin Bascol^{1,2} & Rémi Emonet¹ & Elisa Fromont¹ & Raluca Debusschere²

¹ Univ Lyon, UJM, Lab Hubert Curien, CNRS UMR 5516, F-42000 St-Etienne, Fr

² Bluecime inc., 445 rue Lavoisier, 38330 Montbonnot Saint Martin, France

Abstract. This paper shows how state-of-the-art deep learning methods can be combined to successfully tackle a new classification task related to chairlift security using visual information. In particular, we show that with an effective architecture and some domain adaptation components, we can learn an end-to-end model that could be deployed in ski resorts to improve the security of chairlift passengers. Our experiments show that our method gives better results than already deployed hand-tuned systems when using all the available data and very promising results on new unseen chairlifts.

Keywords: Deep learning, Convolutional neural networks, Image Classification, Domain adaptation.

1 Introduction

Ski resorts are all equipped with different and numerous chair and ski-lifts. To ensure the security at each boarding terminal, one person (at least) is in charge of continuously monitoring the lift traffic and ensuring safe boarding for all its passengers. Possible hazardous situations include: improper seating, restraining bar not pulled down, loss of a gliding equipment, unaccompanied child, etc. On average, all these situations lead each year to 20 serious injuries in France alone, out of which more than a third occur during or just after boarding the lift. Thus the chairlift administrators are in search of viable solutions to prevent the possible accidents. All the chairs could be equipped with sensors that could monitor the number of persons on a given chair, whether they are well seated, whether the position of the restraining bar (or security railing) is correct a few seconds after leaving the boarding station (which is very important to prevent falls), whether a child is not alone on the chair, whether the skiers have not lost a ski or a snowboard etc. Such sensors should trigger an alarm that would stop or slow down the chairlift if a particular anomaly is detected. However, in the extreme weather conditions that ski-resorts can face, small mechanical sensors are subject to wear and often prone to failure, thus other more robust solutions are required.

The work presented in this article is made in collaboration with a start-up called Bluecime, which has installed video cameras filming the boarding platforms in several ski-resorts. Each of these cameras is linked to a computer that processes the video stream in real time. The company has developed a software

(called SIVAO) that takes as input the video of a running chairlift, focuses on one moving chair and automatically (using signal processing techniques) detects if the security railing is up or down a few seconds after boarding. Since its creation, the company has successfully shown that video cameras represent a reliable solution which needs almost no maintenance. Besides, they have proved that the real-time image processing of the captured images leads to better results for the cited tasks than mechanical sensors, even in extreme weather conditions (e.g. the artifacts created by the presence of snow flakes are compensated by various pre-processing methods applied to the input images). If the results obtained by the company are already very good, they are currently limited to one particular task (triggering an alarm if the restraining bar position is not as expected). Moreover, their system needs to be calibrated on a per-lift basis, which is time consuming and makes it more difficult to scale the system to a very large number of lifts.

We believe that intelligent data analysis techniques and, in particular, deep learning, can be used to solve this problem in a more precise and adaptable way. We explain in this article how state-of-the-art deep learning architectures designed for general image classification tasks as well as dedicated domain adaptation methods can be used to improve and generalize the results obtained by the company.

2 Background on Neural Networks and Related Work

Deep learning is nowadays almost a synonym for learning with deep (more than 2 layers) convolutional neural networks (CNN, see [4] for some insights on neural networks). This learning technique has become tremendously popular since 2012 when a deep architecture, called ALEXNET, proposed by A. Krizhevsky et al. [7] was able to win the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) with an outstanding improvement of the classification results over the existing systems: the error rate was 15% compared to 25% for the next ranked team. Since then, countless different architectures have shown excellent performance in many domains such as computer vision, natural language processing or speech recognition [4].

The amount of labeled data available, the systematic use of convolutions, the better optimization techniques and the advances in manufacturing graphical processing units (GPU) have contributed to this success. However, **deep** networks were supposed to suffer from (at least) three curses. 1) The deeper the network, the more difficult it is to update the weights of the first layers (the ones closer to the input): deep networks are subject to *vanishing and exploding gradient* that leads to convergence problems. 2) The bigger the network, the more weights need to be learned. In statistical machine learning, it is well known that more complex models require more training examples to avoid overfitting phenomena and guarantee relatively good test accuracy. 3) Neural network training aims at minimizing a loss which is a measure of the difference between the computed output of the network and the target. The computed outputs depend on (i) the

inputs, (ii) all the weights of the network and (iii) the non linear activation functions that are applied at each layer of the network. The function to minimize is thus high-dimensional and non-convex. As the minimization process is usually achieved using stochastic gradient descent and back-propagation, it can easily get trapped in local optimum.

In this paper, we are interested in deep learning approaches developed for image classification. In this context, the use of convolutions (i.e., a reduced number and a better configuration of the weights) and the huge size of the current image datasets ([2,8]) partially solved the second and third curses mentioned above. In the following, we will show other architectures and learning “tricks” that practically alleviate all the remaining problems with deep learning for image classification.

Learning tricks. The first improvement concerns the *initialization of the weights*. It was shown [1] that image classification tasks could all benefit from a pre-training step on a large dataset such as ImageNet [2]. The ImageNet dataset contains images grouped in 1000 different classes. Practically, the output of a network trained on ImageNet consists in a 1000-D vector which gives the probability for an input image to be of each of the 1000 classes. To use such a network on a new task (with much fewer and different target classes), the idea is to “cut” the last layer(s) of the pre-trained network. Only the convolutional part is kept and its output is considered as a generic image descriptor. A new network can be built from this pre-trained generic feature extractor by training new additional “classification” layers (usually fully connected). The whole network can then be further *fine-tuned*, i.e., trained end-to-end on the new problem of interest. This procedure is called “transfer learning”.

The second improvement is the *batch normalization* [6]. This simple yet very effective idea consists in normalizing the activation outputs of each layer of the network according to the current input batch. Its main benefit is to reduce the internal covariate shift, which corresponds to the amplification of small changes in the input distribution after each layer, and so creates high perturbations in the inputs of the deepest layers.

CNN architectures. K. Simonyan et al. [9] created VGGNET 16 and VGGNET 19 composed of respectively 13 and 16 convolutional layers followed by 3 fully-connected layers. As for the previously mentioned ALEXNET, the spatial dimensions were reduced with max-pooling layers, the inputs of the convolutions being typically padded so as to keep the spatial size. However, each time the spatial dimensions were divided by two, the number of filters in the next layers was doubled to compensate for the information loss. With this method, Simonyan et al. managed to obtain an error rate of 7% at ILSVRC 2014.

Following these results, several other techniques have been designed to be able to train deeper networks. For instance, He et. al. [5] introduced the concept of *residual mapping* which allowed them to create a network with 18 to 152 convolutional layers called a Residual Network (ResNet). Their architecture is divided into blocks composed of 2 or 3 convolutional layers (depending on the

total depth of the network). At the end of a block, its input is added to the output of the layers. This sum of an identity mapping with a “residual mapping” has proven effective at overcoming the vanishing gradient phenomenon during the back-propagation phase and allows to train very deep networks. An example of such blocks is shown in Figure 2. Using residual blocks, the network is learned faster and with better performance. At ILSVRC 2015, ResNet 152 showed the best performance with less than 4% of error rate. They also used *batch normalization* [6] on the first (or first two) layer(s) of each residual block. We will base our system on this architecture.

Domain adaptation. Some applications require to be able to learn a model on some data (e.g. images of cats) called the *source domain* and deploy it on similar but different data (images of tigers) called the *target domain*. This well-known machine learning problem is called *Domain Adaptation*. Many interesting approaches have been proposed in deep learning to tackle this problem such as the one from Ganin et al. [3]. Their idea (as shown in Figure 1) is to train two networks that share the same first (convolutional) layers called the “feature extractor”. The first network is dedicated to the classification task on one particular domain. The second network aims at predicting the domain of an input example. Note that to train this second network, the examples of the target domain do not need to be labeled (and are not in [3]). The only information needed to train the second network is whether an example belongs to the source or to the target domain. The two networks are trained in an adversarial way according to the shared layers (using a mechanism called gradient reversal on the second network optimization). As a consequence, the shared features of the networks are discriminative for the classification task as well as domain invariant.

3 Proposed Deep ResNet with Domain Adaptation

Based on the state-of-the art study presented in Section 2, we propose an image classification architecture using domain adaptation and a convolutional residual network pre-trained on ImageNet. This architecture is shown in Figure 1 and is divided into three parts: 1) a feature extractor which learns a new image representation; 2) a classifier which predicts the class of an image, and 3) a domain discriminator which ensures that the feature extractor is domain invariant, to improve classification accuracy on unseen data.

3.1 Network Architecture

Our network inputs are RGB images of size 224×224 that can be viewed as three-dimensional tensors (two spatial dimensions and one dimension for the RGB channels). After some trial and error, we decided to use the ResNet architecture presented in the previous section with 50 layers and pre-trained on the ImageNet dataset (see the “learning tricks”). Networks with more layers gave a slightly better accuracy but were longer to train and the biggest ones (with

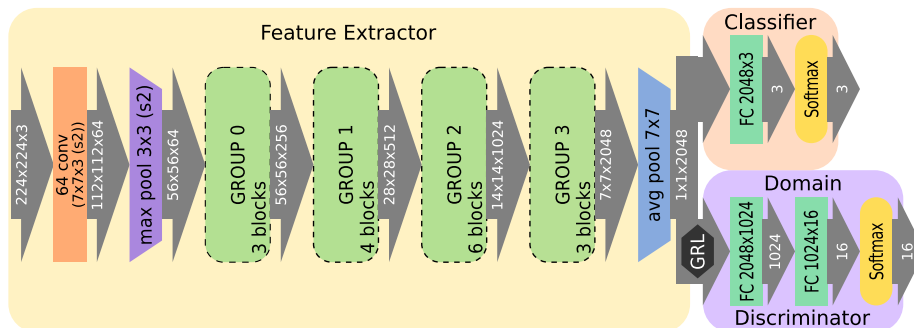


Fig. 1. The proposed ResNet architecture with domain adaptation, for a 3-class classification problem with 16 domains (chairlifts).

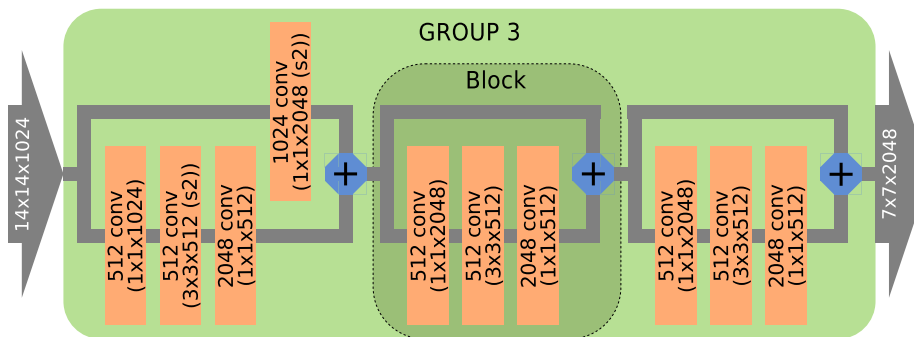


Fig. 2. Details of the last group of layers in our architecture (see Fig. 1), which is composed of three residual blocks and different convolution filters.

more than 101 layers) did not allow us to process test images in real time on a CPU. The 50 layers architecture (shown in Fig. 1) gave us a good trade-off between computational efficiency and accuracy. This network is composed of 49 convolutional layers and only one fully connected layer at the output of the network. This last layer is preceded by an average pooling layer which reduces the size of the feature maps to a 1D-vector. The last layer of the original ResNet 50 architecture was changed to fit our 3-classes classification problem (instead of the 1000 classes of the ImageNet classification problem). We used our chairlift dataset to train this last layer and fine-tune the entire network.

In the feature extractor, each group is composed of a set of blocks. Each block is composed of three layers: a 1×1 convolution that acts as a learnable dimensionality reduction step, a 3×3 convolution that extracts some features and a 1×1 convolution that restores the dimensionality. As explained in Section 2, we use the residual connections between each block to ease the training phase.

The first block of each group contains a 3×3 convolution with a stride of 2 which reduces the spatial size by a factor two. To compute the block output with the sum operator, the input needs to have the same dimensions as the output.

Therefore, a 1×1 convolutional layer with a stride of 2 is added to reduce the spatial size and also to match the number of feature maps which changes according to the group. In Figure 2, we show in more details the last group of layers in the feature extractor.

3.2 Objective Function and Training

Different losses are classically used in deep learning. The best known for image classification is the cross-entropy [4]. After experimenting with different losses, we decided to use the more robust multi-class hinge loss (eq. 1) commonly used in SVM:

$$loss(W, X, t) = \frac{1}{b} \sum_{i=1}^b \frac{1}{|C|} \sum_{c \in C \setminus \{t\}} max(0, o_c^i - o_t^i + m) \quad (1)$$

where W is the set of all the parameters in the network, X is the subset of the input dataset (mini-batch) given to the network during a forward pass, b is the mini-batch size, C is the set of classes, and t is the ground-truth label corresponding to each example of the mini-batch. The margin m can be seen as the minimal “distance” required between the computed probability of the prediction of the target class (o_t^i) and the other classes (o_c^i), since the loss value is equal to 0 when $o_t^i \geq o_c^i + m$. This loss presents several advantages: (i) it is defined on $p(x) = 0$ which makes it more robust than the cross-entropy; (ii) it does not penalize the weights on well classified examples which could speed up the convergence of the network; (iii) the slope and the margin of the function can be easily changed to give a custom weight to the different type of errors (false positive, false negative). This last advantage will be explored in our future work.

To update the weights of the network, a stochastic gradient descent algorithm is applied on the loss function. The weight update rule is given by:

$$W(\tau + 1) = W(\tau) - \eta \cdot \frac{\partial loss(W(\tau), X, t)}{\partial W(\tau)} \quad (2)$$

where τ is the current iteration and η is *the learning rate* which needs to be tuned.

4 Experiments

SIVAO, the current system deployed by Bluecime (introduced in Section 1), can be decomposed into two parts: a *vehicle detector* and a *situation classifier*. The purpose of the vehicle detector is to determine the vehicle position, scale, and orientation so as to track it each time it passes in front of the camera. The situation classifier detects different vehicle parts with respect to the position given by the vehicle detector. Its purpose is to detect whether the vehicle is empty or not and also to infer the position of the restraining bar, so that it




Fig. 3. Training images from three different chairlifts (labeled from left to right *Empty*, *Safe*, and *Unsafe*)

can classify a frame into *Safe* or *Unsafe*. A detection zone is pre-configured and, at the end of the zone, the system computes a final decision based on a configured number of frame-level decisions made during the tracking of the vehicle. Depending on this decision (predicted label), an alarm is triggered or not.

Our proposed network could replace the situation classifier to enhance the classification performance while keeping the efficient vehicle detector designed by Bluecime. In the following, we evaluate the performance of our deep learning approach applied on images provided by the vision-based vehicle detector.

4.1 Chairlift Dataset

Bluecime has already equipped 16 chairlifts (here named A, B, ...) from 3 different ski-resorts with the system presented in Section 1. To build the dataset, each time a vehicle (chair) passes in front of the camera, we use SIVAO to record 2 images that are approximately centered on the chair: one far from and one close to the camera. Examples of such images are given in Figure 3. A total of 50 000 such images have been recorded and manually labeled in 3 categories: *Empty*, *Safe* and *Unsafe*. In the first case (*Empty*), the vehicle does not carry any passengers. The system should not trigger an alarm in this situation. In the second case (*Safe*), the vehicle carries passengers who closed the restraining bar completely. In the last case (*Unsafe*), the vehicle carries passengers and the restraining bar is slightly or completely open.

The dataset is unbalanced with respect to the classes: around 28 000 images are labeled *Empty*, 18 000 are labeled *Safe*, and 4 000 are labeled *Unsafe*. The dataset is also unbalanced with respect to the different chairlifts (which will be called *domains* in the following): the least represented chairlift has 1 800 examples, whereas the most represented one has 6 200 examples.

As they were designed by different manufacturers at different times, all the 16 chairs in the dataset are different, though some are similar. There are even some unique cases, for instance a chairlift having a glass bubble (chairlift D) as a second protection, or another whose vehicles do not have a complete frame (see Fig. 3, left).

4.2 Evaluation Procedure

SIVAO is a commercialized product that aims at improving the work of chairlift operators. We use its current performance as an indication of the minimal result requirements for an acceptable quality of service that we aim to surpass with our automatic system. To compare SIVAO and our method, the results of our multiclass approach are cast into a binary classification problem. To do that, we merge the classes *Empty* and *Safe* as Negative, and the class *Unsafe* becomes Positive.

To study the behavior of our approach, we considered six different experimental settings:

1. **OOC** ("Only One Chairlift"). 16 independent experiments are averaged: each chairlift is considered independently (as it is done by SIVAO), thus the training (resp. the test) set contains 85% (resp. 15%) of the images of a single chairlift. Obviously, with only one chairlift per experiment, the domain adaptation component is not used on this setting.
2. **ALL** One experiment is performed using only the features extractor and the classifier, with 85% of all the available images (from all the chairlifts) for training and 15% for testing. The same experiment is done 16 times with 16 randomly chosen training/test sets and the results are averaged in table 1.
3. **ALL DA** ("Domain Adaptation") same as ALL but using the domain adaptation component described in Section 2.
4. **LOCO** ("Leave One Chairlift Out") 16 experiments are averaged: in each, we use only the feature extractor and the classifier, with the images of all the chairlifts but one (thus 15) mixed in the training set and all the images of the remaining chairlift as test set.
5. **LOCO DA-** Same as LOCO but with domain adaptation using *no examples* from the target chairlift.
6. **LOCO DA** Same as LOCO but with classical domain adaptation where some *unlabeled* examples of the target chairlift are used by the domain adaptation component.

In setting 1 (OOC), due to time constraints and to be closer to a realistic setting, we do not tune the hyper-parameters for each domain but we use one single hyper parameter setting (given in the next section) for all the chairlifts. In this setting, we expect the network to quickly overfit our data and also to be penalized by the lack of examples especially for the least represented chairlifts.

In settings 2 and 3 (ALL and ALL DA), we train our network using all the training data available. We only build one model for all the chairlifts which makes this setting easier to deploy in practice. However, the hyper-parameters of the system are also global which may harm the final performance. These settings could be used with the current cameras installed by the company but do not evaluate the real ability of our system to work on new chairlifts.

Settings 4 to 6 (LOCO, LOCO DA-, and LOCO DA) really show the potential of our approach. Ideally, our method should show good enough performance on these settings to allow Bluecime to deploy their system equipped with this

model on any new chairlift with no manual labeling. In these settings, we expect a performance drop compared to the OOC or ALL settings because of the variability of the different domains.

To evaluate their system, the company relies on numerous measures. Among these, 4 statistical measures assess the overall performance of the system: recall, precision, F-measure, and accuracy. *Unsafe* examples are considered positive (the alarm has to be triggered), and safe examples negative (no alarm needed), thus the classes *Empty* and *Safe* are both considered negative. The *recall* gives the proportion of examples correctly detected positive among all the examples expected positive. In our case, that is the ability of the system to trigger an alarm in unsafe situations. The *precision* gives the proportion of examples correctly detected positive among all the examples detected positive. Thus it indicates the ability of the system to avoid useless alarms. The *F-measure* (harmonic mean between precision and recall) and the accuracy (percentage of correct predictions) give a more global view on the performance of the system. In the following, we report these 4 measures.

4.3 Training Details

In our experiments, we use mini-batches of 84 images (guided by GPU memory limitations), enforcing a balance over classes and domains: we randomly select a chairlift, then randomly select a class, then randomly select an example. We tune our hyper-parameters using a grid search. In all the reported experiments, the learning rate is set to 10^{-5} (with an ADAM optimizer), the hinge loss margin to 0.33 for the domain adaptation component (with the reversal layer) and to 0.01 for the domain discriminator. The gradient reversal layer used for the domain adaptation induces a hyper-parameter γ set to 10 as defined by Ganin et al. [3].

Our dataset is composed of images of size 237×237 . When an image is loaded, we randomly select a 224×224 crop in the image to fit the size of the original ResNet architecture. To maximally exploit the small number of examples available in the positive class, we decided not to use a validation set and to stop the learning process after a constant (high enough) number of iterations set to 12 000. Training our model for 12 000 iterations takes about 4 to 5 hours. The classification of one image is done in approximately 15ms which allows our method to be used in real time.

4.4 Experimental Results

In Table 1, we present the results on all the datasets for all the experimental settings. The first line of the table gives the performance of the current hand-tuned method developed by Bluecime. In the OOC setting (the closest to what the company is currently doing), the proposed network brings a 7 point improvement in precision, recall and F-measure and +1pt in accuracy. This validates the relevance for the company of the proposed machine learning-based method. With the ALL setting, a 5pts gain in F-measure and 1pt gain in accuracy are

Experiment	F-measure	Precision	Recall	Accuracy
Bluecime	84.72	87.88	81.78	97.40
OOC	91.70	94.64	88.94	98.68
ALL	89.47	93.49	85.78	98.33
ALL DA	89.71	93.88	85.89	98.36
LOCO	76.23	76.71	75.75	96.07
LOCO DA-	72.36	70.84	73.93	95.30
LOCO DA	84.24	82.76	85.77	97.33

Table 1. Performance results of our deep learning system for the three main group of settings (OOC, ALL, LOCO) compared to the hand-tuned existing system of the Bluecime company.

observed. While the gain is smaller than with OOC, the *one-model-fits-all* nature of ALL makes it a very attractive solution from an industrial standpoint. Adding domain adaptation (ALL DA) has a very limited impact which can be explained by the fact that all the domains are already in the training set.

As expected, during the LOCO experiment, performance losses occur for all the measures, with -8pts of F-measure and -1pts of accuracy. In this setting, performance is too low for industrial deployment. However, adding domain adaptation (LOCO DA) brings back the performance to a level comparable with the Bluecime system, with -0.5pts of F-measure and a similar accuracy. These results emphasize that domain adaptation is relevant and allows to create models that have competitive results even on a chairlift with non-labeled examples.

In the case when we do not have any examples of a specific chairlift (e.g., launching the system immediately after its installation) adding domain adaptation (LOCO DA-) causes a 4pts loss of F-measure (compared to plain LOCO). This behavior may be unexpected but it has an intuitive explanation: domain adaptation encourages the learning of features that are domain-invariant but also specializes these features for the domains it has seen, causing a detrimental effect on an unseen chairlifts. These results show that in a LOCO setting, domain adaptation should be used only if we can retrain the model using new unlabeled images (setting LOCO DA), and not if the goal is to produce an off-the-shelf model.

In the Table 2, we focus on 4 different chairlifts in order to illustrate how performance varies in function of the specificity of each chair model.

Chairlift A – In both settings, results are better than the average with +4pts (OOC) and +9pts (LOCO) of F-measure. This gain in performance is possible, for OOC, thanks to a good balance in the classes of the images of this chairlift, and for LOCO, thanks to the chairlift configuration similar to several other ones.

Chairlift B – OOC shows a slightly lower performance with -3pts of F-measure but +0.5pts of accuracy. This is explained by the over-representation of the *Safe* class (4 600 images versus 340 for the *Unsafe* class). With the LOCO setting, we

Chairlift	Size	Experiment	F-measure	Precision	Recall	Accuracy
A	2842	OOC	95.68	97.08	94.33	97.06
		LOCO DA	92.98	93.56	92.40	95.21
B	6306	OOC	88.89	91.67	86.27	98.83
		LOCO DA	72.19	66.50	78.95	96.70
C	1822	OOC	76.92	78.95	75.00	96.79
		LOCO DA	77.78	83.05	73.13	96.93
D	2736	OOC	85.71	100.00	75.00	99.76
		LOCO DA	82.76	85.71	80.00	99.63
Overall Average (Table 1)		OOC	91.70	94.64	88.94	98.68
		LOCO DA	84.24	82.76	85.77	97.33

Table 2. Comparison of OOC and LOCO experiments on 4 specific chairlifts

observe a considerable performance drop with -12pts of F-measure. This poor performance is mostly due to the sun which casts different shadows depending of the time of the day. Besides, a high number of images present flares decreasing the image quality.

Chairlift C – OOC has a F-measure of 15pts below the overall result. It is even outperformed by LOCO (+1pts of F-measure and +0.1 of accuracy). Chairlift C is the domain with the fewest number of examples (around 1 800). This number is probably not sufficient to train our deep architecture. In the LOCO experiment, the poor performance is mostly due to the shadow cast by the chairlift tower in the afternoon.

Chairlift D – This is the only chairlift with a glass bubble, and also the only one which has the restraining bar (and also the bubble) closed on *Empty* images. This implies that, even if its data is unbalanced (2 000 *Empty*, 670 *Safe*, and 30 *Unsafe*), on the OOC setting, the network has perfect precision and retrieves 3 of the 4 positives examples, as only 15% of the 30 *Unsafe* are used for testing. The F-measure is 6pts lower than the average result, yet it is highly impacted by this lack of positive testing examples. In LOCO, we could have expected an important performance drop considering the chairlift unique features (the bubble), but when the restraining bar is open, the bubble is open too, and so the configuration in that case is not dramatically different to other chairlifts. This explains why the loss of F-measure is limited to 2pts.

These results underline the different factors limiting the performance of the proposed system. When training and testing is done on the same domain, we need a consequent amount of data to reach good performances. When testing on a new domain, if this domain is very different from the others, we can expect worse results with DA since the domain adaptation will not be able to bridge the domain gap.

For Bluecime, collecting data is easy but annotating them is not because of human resource limitations. Besides, several new chairlifts are equipped each

season. Thus the LOCO setting with DA seems the most promising one in a first deployment phase, since it avoids the need for numerous annotations for the new chairlift. This setting also provides a basis for an active learning system that would limit the number of needed annotations. In a second phase, with enough annotated data, the OOC setting could be used for each resort.

5 Conclusion

We have presented an end-to-end deep learning system that can improve the security on chairlifts using visual information. Leveraging domain adaptation techniques, the model is able to achieve competitive performance even with (new) chairlifts for which no labeled data is available. Overall, this ability to generalize, its competitive accuracy and its operation in real-time make this system well suited for industrial deployment.

Acknowledgment

The authors acknowledge the support from the *SoLStiCe* project ANR-13-BS02-0002-01.

References

1. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: Delving deep into convolutional nets. In: BMVC (2014) **3**
2. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR (2009) **3**
3. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. JMLR (2016) **4, 9**
4. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016), <http://www.deeplearningbook.org> **2, 6**
5. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Computer Vision and Pattern Recognition (CVPR) (2016) **3**
6. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML. pp. 448–456 (2015) **3, 4**
7. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012) **2**
8. Lin, T., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: 13th European Conference in Computer Vision - ECCV (2014) **3**
9. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) **3**