



**HAL**  
open science

## Linear classification of chairlift images for presence analysis

Julien Muzeau, Patricia Ladret, Pascal Bertolino

► **To cite this version:**

Julien Muzeau, Patricia Ladret, Pascal Bertolino. Linear classification of chairlift images for presence analysis. QCAV 2019 - 15th International Conference on Quality Control by Artificial Vision, May 2019, Mulhouse, France. pp.55, 10.1117/12.2521686 . hal-01980423v2

**HAL Id: hal-01980423**

**<https://hal.science/hal-01980423v2>**

Submitted on 24 Oct 2019

**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.

# Linear classification of chairlift images for presence analysis\*

Julien Muzeau<sup>a</sup>, Patricia Ladret<sup>a</sup> and Pascal Bertolino<sup>a</sup>

<sup>a</sup>Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab  
38000 Grenoble, France

## ABSTRACT

In the recent past years, innumerable techniques more complex than the others have emerged in computer vision. They have been applied to many fields and, thanks to the tremendous computational power one has access to nowadays, have made possible more and more elaborate applications. In this article, we propose a classification tool, using hand-crafted interpretable (statistical and digital imaging) features, in order to confirm or invalidate the presence of passengers on skilift vehicles in mountain ranges. More precisely, Linear Discriminant Analysis, which is a dimensionality reduction alongside classification technique, and its less restrictive variant Quadratic Discriminant Analysis, are applied. One of the paper's objectives consists in illustrating the famous law of parsimony, also known as Occam's razor, in the sense that the simpler solutions should be considered first and more complex models should be built afterwards if needed.

**Keywords:** Binary classification, Linear Discriminant Analysis, chairlift images, presence analysis

## 1. INTRODUCTION

Computer vision is widely used nowadays, from quality control of processes or products in industry to human surveillance and navigation. A digital camera is indeed rather easy and cheap to set up. Further algorithms allow us to deduce information about the observed scene. Numerous techniques have been developed throughout the years in order to provide more and more accurate results. However, it naturally implies an increase in complexity.

A question we can ask ourselves is the next: does a gain in assessment justify some raise in complexity? In this paper, we illustrate Occam's razor (or law of parsimony) which can be summarized as following: simpler theories should be preferred to more complex ones, not because they are more realistic or give better understanding, but because they can be tested in an easy manner.

We apply this idea to mountain ranges and more specifically to the detection of passengers on skilift vehicles. We propose a simple machine learning process, namely linear classification, and compare it to a more evolved technique which makes use of deep learning. The results are naturally different but have to be compared with respect to the required resources.

The remainder of this article is organized as follows: section 2 explains the context of our application. The proposed method is detailed in section 3 and results are given in section 4.

## 2. CONTEXT

Ensuring security in mountain ranges has become a concerning and challenging task for the last decades. Most of the accidents in this field are indeed mainly due to human behaviour. It is specifically the case for skilifts: people often forget to lower the security railing, the chair (generally referred to as *vehicle*) can be overcharged, that means it hosts more persons than the number of available seats, a passenger loses one of his/her ski, *etc.* All those cases can lead to a hazardous situation.

---

Further author information: (Send correspondence to J.M.)

J.M.: E-mail: julien.muzeau@gipsa-lab.grenoble-inp.fr, Telephone: +33 (0)4 76 82 71 34

\* This work is funded by the Auvergne-Rhône-Alpes region





Figure 1: Variability demonstrated through 3 images taken from each of 4 different skilifts (top-left, top-right, bottom-left and bottom-right). Blur is intentional for privacy reasons.

To prevent accidents, a computer vision system has been developed by a French company from Grenoble in the Alps: a camera, positioned around the first skilift pylon, monitors each vehicle departure, analyzes several features of the scene (*i.e.* a sequence of 30 to 100 frames) and concludes to a dangerous or a safe situation with a certain reliability. An audible alarm is triggered whenever such a sequence is detected: that way, the skilift operator is aware that a potential danger can occur and thus pays more attention, while the passengers secure the vehicle (put the safety rail down for instance).

One of the main scene features obviously is the presence of passengers: as a matter of fact, if the vehicle does not transport any skier, further post-processing is not needed. This paper provides a simple albeit effective method which aims at establishing whether the chair is occupied or not. We show this classification can be achieved without implementing complex and computationally expensive techniques such as deep learning approaches.

### 3. PROPOSED METHOD

#### 3.1 Details on the dataset

As explained in section 2, a digital Red-Green-Blue camera observes the boarding from the first pylon of the skilift. Examples of camera frames are given in figure 1: it depicts 3 images before desaturation taken at various times for each of 4 skilifts. One can clearly notice how difficult our problem is: the different variabilities appear, either the intra-skilift ones (weather, lightning, passengers, season, *etc*) or the ones highlighting the differences between all skilifts (vehicle size, skilift geometry, front or back view, *etc*).

The dataset is divided into 19 skilifts in 3 different resorts. The names are kept secret for privacy reasons. Table 1 summarizes the dataset. One can notice how unbalanced it is: the number of available images are oscillating between 10k to 140k and the percentage of each class can be quite different from 50%.

#### 3.2 Linear Discriminant Analysis

The proposed method makes use of Fisher’s famous Linear Discriminant Analysis (LDA).<sup>1</sup> Its different steps are summarized here.

Let us first consider a set of  $N$  examples  $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$ , where  $\mathbf{x}_j$  is the  $d$ -dimensional feature vector and  $y_j$  its associated label, *i.e.* the class  $\mathbf{x}_j$  belongs to. We consider a  $C$  classes problem. Each class  $i = 1 \dots C$  contains  $N_i$  data points. The mean of the whole dataset is  $\mu$ .

In real terms, the Linear Discriminant Analysis can be broken down into six steps.

Table 1: Details on the dataset. The first column is the identifier of each skilift. n\_images represents the number of available images and n\_empty\_images the number of images with empty vehicles. The same information is given as a percentage of n\_images in the last column.

Skilift	n_images	n_empty_images	%_empty
Skilift 1	87143	53735	61.7
Skilift 2	11161	3032	27.2
Skilift 3	14330	7575	52.9
Skilift 4	22315	12953	58.1
Skilift 5	36493	17567	48.1
Skilift 6	33803	4792	14.2
Skilift 7	139058	61324	44.1
Skilift 8	102605	58293	56.8
Skilift 9	20718	10574	51.0
Skilift 10	39596	23329	58.9
Skilift 11	29957	9078	30.3
Skilift 12	9862	7146	72.5
Skilift 13	97572	56770	58.2
Skilift 14	93357	40646	43.5
Skilift 15	26669	22688	85.1
Skilift 16	100492	86811	86.4
Skilift 17	17399	10207	58.7
Skilift 18	67771	42636	62.9
Skilift 19	25013	19129	76.5

1. For each class  $i = 1 \dots C$ , compute the mean vector  $\boldsymbol{\mu}_i \in \mathbb{R}^d$  and the scatter matrix

$$\mathbf{S}_i = \sum_{\mathbf{x}_j \text{ s.t. } y_j=i} (\mathbf{x}_j - \boldsymbol{\mu}_i)(\mathbf{x}_j - \boldsymbol{\mu}_i)^T \quad (1)$$

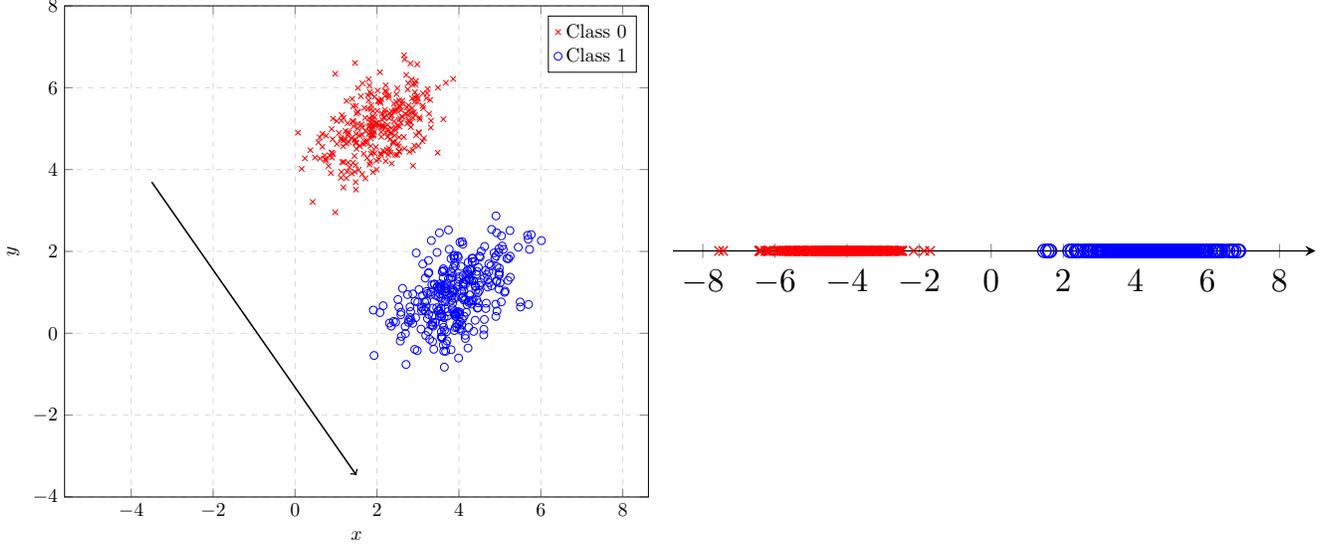
2. Compute the within-class scatter matrix

$$\mathbf{S}_w = \sum_{i=1}^C \mathbf{S}_i \quad (2)$$

3. Compute the between-class scatter matrix

$$\mathbf{S}_b = \sum_{i=1}^C N_i (\boldsymbol{\mu} - \boldsymbol{\mu}_i)(\boldsymbol{\mu} - \boldsymbol{\mu}_i)^T \quad (3)$$

4. Obtain the eigenvalues  $\lambda$  and eigenvectors  $\mathbf{v}$  of the matrix  $\mathbf{S}_w^{-1} \mathbf{S}_b$ . This operation aims at determining the axis, represented by the eigenvectors  $\mathbf{v}$ , which lead to the minimum intra-class variance, that means how gathered the data points of each class are, and the maximum inter-class variance, *i.e.* how far each class lies from one another.
5. Sort the eigenvalues in descending order and keep the first  $k$  corresponding eigenvectors in a matrix  $\mathbf{P}$ . The choice of  $k$  is usually motivated by the explained variance ratio.
6. The matrix  $\mathbf{P}$ , of size  $d \times k$ , allows to project the original  $d$ -dimensional data onto a reduced space of dimension  $k \leq d$ .



(a) The data points which belong to class 0 are represented by red crosses, while the ones from class 1 are depicted with blue circles. The axis plotted in black is the Linear Discriminant Analysis output.

(b) Projection of the original bidimensional data points onto unidimensional axis. The two classes are subsequently easily separable.

Figure 2: Example of a Linear Discriminant Analysis applied to a bidimensional dataset with 2 classes.

Eventually, the computed projection allows the best discrimination between data from different classes. The whole process is illustrated on a trivial example in figure 2. The bidimensional dataset, made up of 2 classes, can be separated by the projection onto the LDA axis, depicted in black in both subfigures. To sum up, LDA consists in dimensionality reduction alongside classification.

### 3.3 Chosen features

The main point when applying any machine learning method is the choice of the variables (or features). In this paper, after converting the input color images (like those depicted in figure 1) to gray levels, the idea is to mix statistical indicators and image processing metrics. Thus, considering an image  $\mathbf{I} = (I_{ij})$  ( $i = 1 \dots h, j = 1 \dots w$ ), the chosen features are:

- Mean

$$\mu = \frac{1}{hw} \sum_{i=1}^h \sum_{j=1}^w I_{ij} \tag{4}$$

- Maximum

$$max = \max_{i=1 \dots h, j=1 \dots w} I_{ij} \tag{5}$$

- Minimum

$$min = \min_{i=1 \dots h, j=1 \dots w} I_{ij} \tag{6}$$

- Standard-deviation

$$\sigma = \frac{1}{hw} \sum_{i=1}^h \sum_{j=1}^w (I_{ij} - \mu)^2 \tag{7}$$

- **Skewness**

$$sk = \frac{1}{hw} \sum_{i=1}^h \sum_{j=1}^w \left( \frac{I_{ij} - \mu}{\sigma} \right)^3 \quad (8)$$

- **Kurtosis**

$$kt = \frac{1}{hw} \sum_{i=1}^h \sum_{j=1}^w \left( \frac{I_{ij} - \mu}{\sigma} \right)^4 \quad (9)$$

- **Michelson contrast** (or visibility)<sup>2</sup>

$$c = \frac{max - min}{max + min} \quad (10)$$

- **Blur level**<sup>3</sup>
- **Number of connected components**
- **Contours variance**
- **Contours length**

## 4. EXPERIMENTS

This section describes the conducted experiments and gives results.

As explained in subsection 3.1, many variabilities affect the content of the images. In order to thwart them, at least the inter-skilift ones, we decided to work on each skilift independently and to build different models.

For each skilift, we proceed as following. The images are divided into two subsets of equivalent size, each of those has the same empty/non-empty ratio as the whole skilift dataset. This way one guarantees that the learned model is representative of every kind of situation. The aforementioned model is trained on the first part then tested on the second.

### 4.1 Metrics

In order to assess the quality of the model, as we measure binary classification results, we naturally make use of True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN) counters. Those are reminded in table 2. For instance, FN represents the number of data points which are predicted as members of class 0 while actually belonging to class 1.

Table 2: Definition of True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN) counters.

	<b>0</b>	<b>1</b>	Groundtruth
<b>0</b>	TN	FN	
<b>1</b>	FP	TP	
Prediction			

Consequently, the metrics used are Precision, Accuracy, Recall and F-measure:

$$Precision = \frac{TP}{TP + FP} \quad (11)$$

$$Accuracy = \frac{TN + TP}{TP + TN + FP + FN} \quad (12)$$

$$Recall = \frac{TP}{TP + FN} \quad (13)$$

$$F\text{-measure} = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall} \quad (14)$$

The F-measure is in fact the harmonic mean between Precision and Recall. This criterion is needed and meaningful in our case as the classes are unbalanced. Class 1 represents images with empty vehicles while images with passengers belong to class 0.

## 4.2 Results

Table 3 presents an excerpt of the results obtained for the 19 different chairlifts. The whole process is repeated 30 times for each skilift, with the train and test subsets chosen randomly at each iteration, the mean and standard deviation are displayed for each of the aforementioned metrics. The last row of the table shows the averaged Precision, Recall, Accuracy and F-measure over the whole dataset.

Table 3: Linear Discriminant Analysis applied to skilift images. Precision, recall, accuracy and F-measure are given for each skilift (mean and standard-deviation, as the process is repeated 30 times). The last row gives the average measures over the whole dataset.

Skilift	Precision	Recall	Accuracy	F-measure
Skilift 1	0.805 ± 0.025	0.832 ± 0.013	0.839 ± 0.009	<b>0.818 ± 0.013</b>
Skilift 2	0.8 ± 0.016	0.82 ± 0.018	0.781 ± 0.011	<b>0.81 ± 0.01</b>
Skilift 3	0.857 ± 0.003	0.968 ± 0.014	0.835 ± 0.009	<b>0.909 ± 0.006</b>
Skilift 4	0.737 ± 0.019	0.755 ± 0.036	0.737 ± 0.015	<b>0.745 ± 0.018</b>
Skilift 5	0.677 ± 0.011	0.813 ± 0.025	0.661 ± 0.013	<b>0.739 ± 0.012</b>
Skilift 6	0.804 ± 0.025	0.732 ± 0.022	0.864 ± 0.008	<b>0.766 ± 0.014</b>
Skilift 7	0.878 ± 0.004	0.986 ± 0.004	0.869 ± 0.004	<b>0.929 ± 0.002</b>
Skilift 8	0.736 ± 0.011	0.83 ± 0.015	0.712 ± 0.007	<b>0.78 ± 0.006</b>
Skilift 9	0.944 ± 0.011	0.982 ± 0.007	0.955 ± 0.006	<b>0.962 ± 0.005</b>
Skilift 10	0.725 ± 0.012	0.821 ± 0.018	0.692 ± 0.012	<b>0.77 ± 0.012</b>
Skilift 11	0.841 ± 0.029	0.829 ± 0.024	0.91 ± 0.008	<b>0.834 ± 0.013</b>
Skilift 12	0.943 ± 0.018	0.985 ± 0.019	0.946 ± 0.017	<b>0.963 ± 0.012</b>
Skilift 13	0.894 ± 0.021	0.933 ± 0.013	0.906 ± 0.012	<b>0.913 ± 0.01</b>
Skilift 14	0.767 ± 0.013	0.869 ± 0.02	0.77 ± 0.008	<b>0.814 ± 0.007</b>
Skilift 15	0.753 ± 0.01	0.809 ± 0.018	0.734 ± 0.007	<b>0.779 ± 0.007</b>
Skilift 16	0.928 ± 0.015	0.944 ± 0.009	0.901 ± 0.013	<b>0.936 ± 0.008</b>
Skilift 17	0.893 ± 0.014	0.976 ± 0.007	0.932 ± 0.008	<b>0.933 ± 0.007</b>
Skilift 18	0.886 ± 0.034	0.919 ± 0.025	0.971 ± 0.006	<b>0.902 ± 0.02</b>
Skilift 19	0.829 ± 0.01	0.824 ± 0.011	0.847 ± 0.005	<b>0.826 ± 0.006</b>
Average (Ours)	0.822	0.863	0.831	0.841

## 4.3 Improving results

The linear classification detailed in this paper can also be approached from a probabilistic point of view. One can indeed consider the probabilities of a new example  $\mathbf{x}$  belonging to class 0  $p(\mathbf{x}|y = 0)$  and class 1  $p(\mathbf{x}|y = 1)$ . These probabilities are subsequently modeled by multivariate normal distributions  $\mathcal{N}(\boldsymbol{\mu}_0; \boldsymbol{\Sigma}_0)$  and  $\mathcal{N}(\boldsymbol{\mu}_1; \boldsymbol{\Sigma}_1)$ , where  $\boldsymbol{\mu}_j$  and  $\boldsymbol{\Sigma}_j$  respectively are the mean and covariance matrix of all data points that belongs to class  $j$ .

Thus, comparing these values leads to:

$$\begin{aligned}
p(\mathbf{x}|y = 0) &\stackrel{1}{\underset{0}{\leq}} p(\mathbf{x}|y = 1) \\
\iff \frac{1}{(2\pi)^{d/2}|\boldsymbol{\Sigma}_0|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}_0^{-1}(\mathbf{x} - \boldsymbol{\mu}_0)\right) &\stackrel{1}{\underset{0}{\leq}} \frac{1}{(2\pi)^{d/2}|\boldsymbol{\Sigma}_1|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_1^{-1}(\mathbf{x} - \boldsymbol{\mu}_1)\right) \\
\iff -\frac{1}{2} \ln |\boldsymbol{\Sigma}_0| - \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}_0^{-1}(\mathbf{x} - \boldsymbol{\mu}_0) &\stackrel{1}{\underset{0}{\leq}} -\frac{1}{2} \ln |\boldsymbol{\Sigma}_1| - \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_1^{-1}(\mathbf{x} - \boldsymbol{\mu}_1) \\
\iff \ln |\boldsymbol{\Sigma}_0| + (\mathbf{x} - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}_0^{-1}(\mathbf{x} - \boldsymbol{\mu}_0) &\stackrel{0}{\underset{1}{\leq}} \ln |\boldsymbol{\Sigma}_1| + (\mathbf{x} - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_1^{-1}(\mathbf{x} - \boldsymbol{\mu}_1)
\end{aligned} \tag{15}$$

For enhancing further processing, the Linear Discriminant Analysis makes use of the homoscedasticity hypothesis, that means the covariance matrices are assumed to be equal ( $\boldsymbol{\Sigma}_0 = \boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}$ ). Applying this assumption to equation 15, one can derive:

$$\mathbf{x}^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1) \stackrel{1}{\underset{0}{\leq}} \frac{1}{2}(\boldsymbol{\mu}_0^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_0 - \boldsymbol{\mu}_1^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_1) \tag{16}$$

We can notice how equation 16 depends linearly on  $\mathbf{x}$ , hence the LDA name. However, this homoscedasticity assumption is not needed: we then refer to this method (see equation 15) as Quadratic Discriminant Analysis (QDA).<sup>4</sup> Therefore, for our application, we combine Linear and Quadratic Discriminant Analyzes. Results are given in table 4. Furthermore, we display in the last row the scores obtained by deep learning techniques.<sup>5</sup>

Table 4: Averaged metrics over the whole dataset when LDA only is applied and when LDA is combined with QDA. One can see the improvement brought by the use of both methods. The last row also gives the average measures over the whole dataset for the deep learning approach.<sup>5</sup>

	<b>Precision</b>	<b>Recall</b>	<b>Accuracy</b>	<b>F-measure</b>
Average (LDA only)	0.822	0.863	0.831	0.841
Average (LDA + QDA)	0.840	0.881	0.846	0.859
Average <sup>5</sup>	0.946	0.889	0.987	0.917

Obviously, the results produced by our linear classification do not compete with the ones obtained with neural networks. However they have to be put in contrast to the complexity of the model (more than 20 million parameters for<sup>5</sup>) and the duration of the training phase (4 to 5 hours versus less than an hour in our case).

## 5. CONCLUSION & PERSPECTIVES

We presented in this paper a classification tool that, applied to skilift images, allows to distinguish between the ones representing empty vehicles and the ones with passengers onboard. We deliberately chose one of the simplest model available, *e.g.* Linear Discriminant Analysis. Although not ideal, the results are satisfactory, especially when considering the simplicity and the speed of the model and the fact that interpretable features are used.

Perspectives are numerous in order to improve our results. First, it is crucial to investigate other features. One can think about spatial ones such as Signal-to-Noise Ratio, Total Variation, local contrast, histogram (gray level and color). It can also be interesting to consider the frequential content of each image (using Gabor filters for instance). Another perspective would be the work on local patches instead of considering the image as whole as the emptiness decision can be enhanced when looking at the heads or feet of the passengers. Finally, the results can be improved by making use of more elaborate classification tools such as Support Vector Machines<sup>6</sup> or Random Forests.<sup>7</sup>

## ACKNOWLEDGMENTS

The authors of this article would like to thank Yulu WANG for his insight into this subject and the Auvergne-

Rhône-Alpes region  for the funding of this work.

## REFERENCES

- [1] Fisher, R. A., “The use of multiple measurements in taxonomic problems,” *Annals of Eugenics* **7**(2), 179–188 (1936).
- [2] Michelson, A. A., [*Studies in optics*], The University of Chicago Press (1927).
- [3] Crete-Roffet, F., Dolmiere, T., Ladret, P., and Nicolas, M., “The blur effect: Perception and estimation with a new no-reference perceptual blur metric,” *Proceedings of SPIE - The International Society for Optical Engineering* **6492** (Feb. 2007).
- [4] Hastie, T., Tibshirani, R., and Friedman, J., [*The elements of statistical learning*], Springer (2009).
- [5] Bascol, K., Emonet, R., Fromont, E., and Debusschere, R., “Improving chairlift security with deep learning,” in [*International Symposium on Intelligent Data Analysis*], (2017).
- [6] Cortes, C. and Vapnik, V., “Support-vector networks,” *Machine learning* **20**, 273–297 (Sept. 1995).
- [7] Breiman, L., “Random forests,” *Machine learning* **45**, 5–32 (Oct. 2001).