RECPAD 2023
27th October 2023

INSTITUTO DE SISTEMAS E ROBÓTICA
UNIVERSIDADE DE COIMBRA

Dealing with Overfitting in the Context of Liveness Detection using FeatherNets with RGB images

Miguel Leão - miguel.leao@isr.uc.pt
Nuno Gonçalves - nunogon@deec.uc.pt

Liveness Detection or Face Anti-Spoofing has been developed in tandem with facial recognition technology

The interest of liveness detection:

- To provide security to facial recognition applications;
- To be used in the day-to-day applications that already use facial recognition;

Problems:

- These applications do not have access to the grade of equipment used in the development of the current state of the art methods.
- The simplicity of the loss function results in overfitting;
- The imbalance of the available datasets;

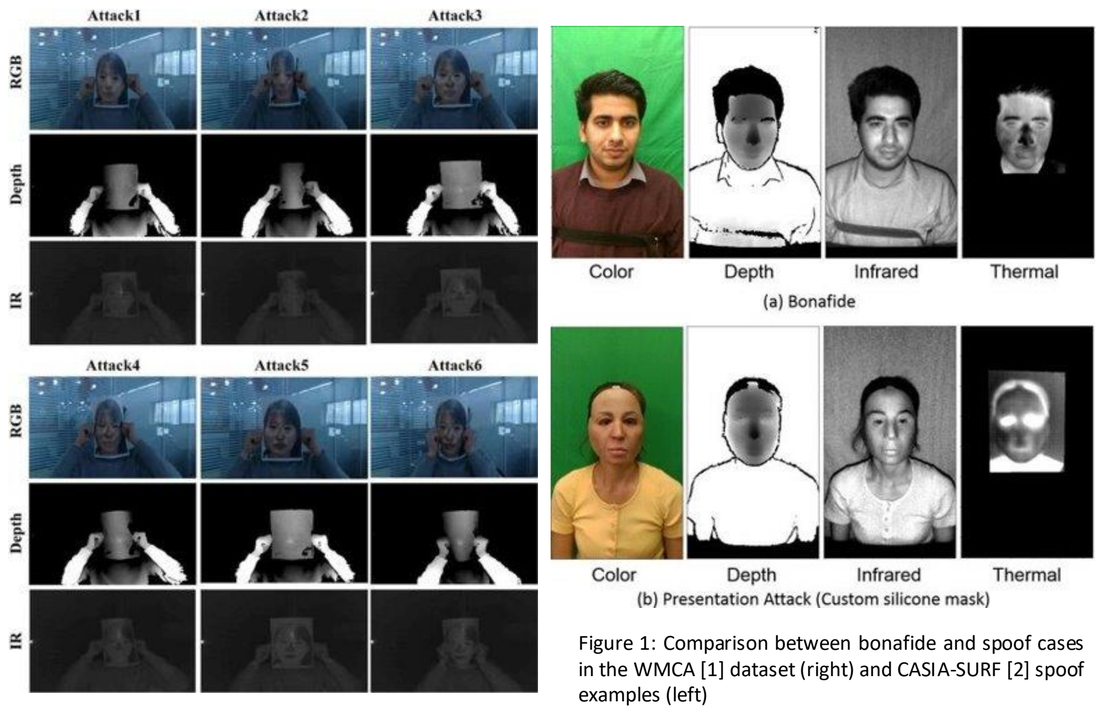

Figure 1: Comparison between bonafide and spoof cases in the WMCA [1] dataset (right) and CASIA-SURF [2] spoof examples (left)

FeatherNets

- Deep Learning (DL) approaches to liveness detection tend to be heavy in both computational requirements and data storage;
- The objective is to make these approaches function in any scenario, independent of system capabilities;
- [3] achieves ACER of 0.00168, with only 0.35 million parameters and 83 million flops down from CASIA-SURF's baseline using ResNet18 [4] with an ACER of 0.05 with 11.18 million parameters and 1800 million flops.

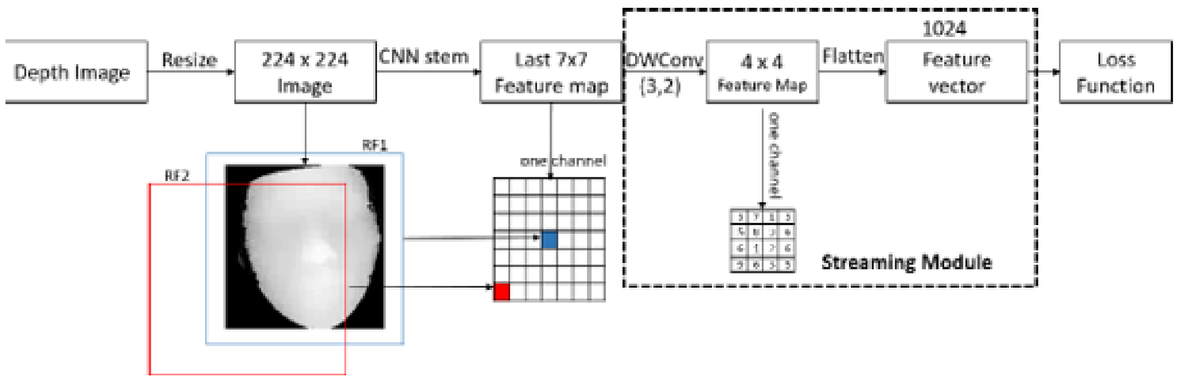

Figure 2: FeatherNet strucuture

- Binary cross-entropy loss, often used in liveness detection, results in overfitting due to the its simplicity;
- The focal loss function [5] extends the base function with a weight parameter to each of the labels present in the dataset (alpha) and a modulating factor that down weights easier examples, with a tunable focus parameter (gamma).

$$FocalLoss(p_t) = -\alpha_t(1-p_t)^{\gamma}log(p_t)$$

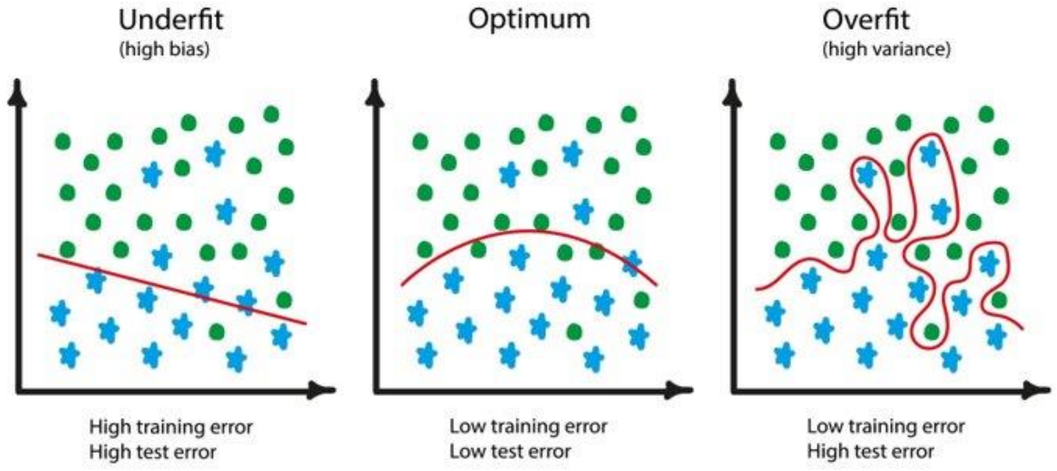

Figure 3: Visualization of how a model underfits, overfits and the ideal result.

Results obtained when running FeatherNets with RGB images

Model	Dataset	γ	Best Epoch	EER	Accuracy	APCER	BPCER	ACER
FeatherNet A	CASIA-SURF	2	1	0.093	91.996	0.039	0.172	0.105
		3	9	0.081	89.748	0.129	0.044	0.087
		5	20	0.080	90.466	0.117	0.048	0.082
FeatherNet B	CASIA-SURF	2	12	0.093	89.675	0.117	0.073	0.095
		3	3	0.093	91.674	0.068	0.117	0.092
		5	19	0.067	92.038	0.093	0.049	0.071
FeatherNet A	WMCA	2	16	0.0005	99.972	0.0001	0.001	0.0005
		3	69	0.043	96.988	0.024	0.051	0.038
		5	160	0.004	99.696	0.001	0.008	0.005
FeatherNet B	WMCA	2	81	0.0005	99.986	0.000	0.005	0.0003
		3	63	0.026	98.529	0.009	0.033	0.021
		5	122	0.0003	99.993	0.000	0.0003	0.0001

Model	Dataset	γ	Avg. Acc.	Std.	APCER Avg.	BPCER Avg.
FeatherNet A	CASIA-SURF	2	52.306	0.888	0.692	0.002
		3	53.179	0.877	0.679	0.002
		5	60.866	1.422	0.566	0.004
FeatherNet B	CASIA-SURF	2	56.582	1.275	0.630	0.002
		3	58.762	1.170	0.598	0.001
		5	57.667	1.328	0.614	0.002
FeatherNet A	WMCA	2	99.392	0.061	0.005	0.010
		3	95.984	0.165	0.032	0.067
		5	99.449	0.085	0.005	0.008
FeatherNet B	WMCA	2	99.868	0.073	0.001	0.001
		3	97.479	0.298	0.015	0.060
		5	99.917	0.089	0.001	0.0003

Results obtained after adjusting the decision threshold according to a PR curve

Model	Dataset	γ	Best Epoch	EER	Accuracy	APCER	BPCER	ACER
FeatherNet A	CASIA-SURF	3	36	0.117	89.373	0.049	0.232	0.141
FeatherNet A	GRAFTSET - 10% Both	3	59	0.110	90.377	0.046	0.217	0.131
FeatherNet A	WMCA	3	189	0.018	91.165	0	0.385	0.193

Model	Dataset	γ	Avg. Acc.	Std.	APCER Avg.	BPCER Avg.
FeatherNet A	CASIA-SURF	3	85.746	1.003	0.160	0.105
FeatherNet A	GRAFTSET - 10% Both	3	87.478	0.401	0.113	0.155
FeatherNet A	WMCA	3	88.446	1.012	0	0.504

With this work it was possible to:

- Conclude on the importance of a well-rounded dataset;
- Mitigating overfitting, reducing the difference from the best epoch to the average of the last epochs from 36.57% to 3.63%.

In the future:

- Apply these conclusions in the making of a new model, keeping it as simple as possible;
- Developing a dataset with as much variety as possible, be it in types of spoof, individuals and capture condition.

References

[1] A. George, Z. Mostaani, D. Geissenbuhler, O. Nikisins, A. Anjos, and S. Marcel, "Biometric face presentation attack detection with multi-channel convolutional neural network," IEEE Transactions on Information Forensics and Security, vol. 15, 2020.

[2] S. Zhang, X. Wang, A. Liu, C. Zhao, J. Wan, S. Escalera, H. Shi, Z. Wang, and S. Z. Li, "A dataset and benchmark for large-scale multi-modal face anti-spoofing," in Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2019-June, 2019.

[3] P. Zhang et al., "FeatherNets: Convolutional Neural Networks as Light as Feather for Face Anti-Spoofing," 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Long Beach, CA, USA, 2019, pp. 1574-1583, doi: 10.1109/CVPRW.2019.00199.

[4] K. He, X. Zhang, S. Ren and J. Sun, "Deep Residual Learning for Image Recognition," 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770-778, doi: 10.1109/CVPR.2016.90.

[5] T. Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 42, 2020.