SD® D&LEMC

September 23-26, 2019 Santa Clara, CA

New Perspective on Machine Learning Predictions Under Uncertainty

Rahul Vishwakarma, Jayanth Reddy

Agenda

September 23-26, 2019 Santa Clara, CA SD® D¢LLEMC

Understanding Trustworthiness of Prediction

Quantifying Uncertainty

Application in risk-sensitive system

Why should we care?

2019 Storage Developer Conference. © Dell EMC. All Rights Reserved.

Classical ML Approach

September 23-26, 2019 Santa Clara, CA

- Classification (Binary)
 - Input: $(x_1, y_1), (x_2, y_2), (x_n, y_n)... (x_{n+1}, y_{n+1} = ?)$
 - Task: Predict *label* of (y_{n+1}) new data point
 - Prediction model : $f(x_i) \rightarrow (y_i)$
- Model performance
 - $f(x_i) \to (y_i)$ accuracy is 98% for $(x_1, y_1), (x_2, y_2), (x_n, y_n)$

Can we trust the prediction accuracy of (y_{n+1}) is also 98% Are we really confident about the specific prediction

Motive

September 23-26, 2019 Santa Clara, CA

DELLEMC

- Prediction models only output bare predictions but not the confidence in those predictions
- Obtain a metrics which explains:
 - Confidence of prediction for new label
 - Informativeness of each new data points

How do we quantify prediction uncertainty

How do we know when to trust our results?

September 23-26, 2019

Sources of Uncertainty in ML



Figure 1: Sources of uncertainty in machine learning

SD® DØLLEMC

Uncertainty Estimation

September 23-26, 2019 Santa Clara, CA

- Algorithmic Randomness
 - Algorithmic information theory provides universal measures of confidence but these are, unfortunately, **non-computable**
 - Obtain practicable approximations to universal measures of confidence

Conformal Prediction



Learning Framework

September 23-26, 2019 Santa Clara, CA

- Generalizable for any machine leaning algorithm
- Framework
 - Algorithmic randomness¹
 - problem of assigning confidences to predictions is closely connected to the problem of defining random sequences
 - Hypothesis testing

¹Algorithmic Learning in a Random World by Vovk, Gammerman and Shafer. Springer, 2005.

SD (®

D%LLEMC

Assumption

September 23-26, 2019 Santa Clara, CA

Exchangeability

 $P\{(z_1, z_2, \dots) \in Z^{\infty} : (z_1, \dots, z_n) \in E\} = P\{(z_1, z_2, \dots) \in Z^{\infty} : (z_{\pi(1)}, \dots, z_{\pi(n)}) \in E\}$



SD® Deellemc

Conformal Prediction

September 23-26, 2019 Santa Clara, CA

SD® D&LLEMC



Problem statement

Predict the label (Y) of new data point



September 23-26, 2019 Santa Clara, CA

- Given
 - Disk Drives dataset
- Predict
 - Label of new Disk
 - Assign Confidence and Credibility
- Model
 - k-Nearest Neighbor (k =1)

SD[®]

DELLEMC

September 23–26, 201 Santa Clara, CA

Nonconformity score

Assuming the predicted label is a Failed disk: Calculate non conformity score





SDC D¢LLEMC

2019 Storage Developer Conference. © Dell EMC. All Rights Reserved.

September 23-26, 20 Santa Clara, CA

Nonconformity score

Assuming the predicted label is a Normal disk: Calculate non conformity score





SDC DØLLEMC

p-value and Prediction



$$p(\alpha_{n+1}^{y_p}) = \frac{\#\{i: \alpha_i^{y_p} \ge \alpha_{n+1}^{y_p}\}}{n+1}$$

Higher non-conformity : Lower p-value

Credibility = p_{max} Confidence = $1 - p_{max}^{2nd}$



SD© D¢llemc

Variants of Conformal Prediction

- Inductive Conformal Prediction
 - Computationally efficient
- Aggregated Conformal Prediction
- Cross Conformal Prediction

2019 Storage Developer Conference. © Dell EMC. All Rights Reserved.

SD (®

D%LLEMC



Applications



Applications

September 23-26, 2019 Santa Clara, CA SD® DØLLEMC

- Hardware failure prediction
 - Disk drive
- Anomaly detection
 - Time Series
- Feature selection
- Prediction quality assessment

September 23-26, 201

Disk Drive Failure Prediction

- Dataset: Black Blaze Q2_2019
- Classification using Conformal Framework
 - K- Nearest Neighbors
 Want to try out Conformal Classification package?
 https://pypi.org/project/ConfClr/ (pip install ConfClr)
- Interpretation of results
 - Confidence
 - Credibility
- Application
 - Ability to Rank the Labels

https://f001.backblazeb2.com/file/Backblaze-Hard-Drive-Data/data_Q2_2019.zip

18

SD (®

D%LLEMC

Dataset

Databases 🗗	X Gri	id view	Form vie	w												
Filter by name							100 Tel (5)	The state								
 drive_stats (SQLite 3) 	22	• •				9 1 2 2 (8)	AL 45 22	Hiter data	IT • T	otal rows loaded: 45842						
✓ I Tables (1)				and a surface				faith and	and C and							
✓ III drive_stats		Gate		enal_number		node	capacity_bytes	laiure	smart_p_raw	smart_12_raw	smart_to/_raw	smart_100_rav	smart_105_raw	smart_190_raw	smart_130_raw	
Columns (12)	1	date	s	erial_number		model	capacity_bytes	failure	smart_1_normalized	smart_1_raw	smart_2_normalized	smart_2_raw	smart_3_normalized	smart_3_raw	smart_4_normalized	
III date	2	2019-06	-24 2	305B2QN		ST4000DM000	4000787030016	1	1 114	72570896	134	97	91	(100	ð
U serial_number	3	2019-06	-24 2	JV0XJQ4		ST12000NM0007	12000138625024	. (0 80	99542112	134	0	98	I I	10	Ó
	4	2019-06	-24 7	UV0XJQ3		ST12000NM0007	12000138625024	. (0 83	215409040	132	0	99	j r	10	D
failure	5	2019-06	-24 2	0DLX0AL		ST12000NM0007	12000138625024		1 78	56561808	132	97	93	i r	10	0
<pre>smart_5_raw</pre>	6	2019-06	-24 F	L1331LAHG1	4H	HGST HMS5C4040ALE640	4000787030016	(0 100) 0	136	96	100	43	5 10	D
smart_12_raw	7	2019-06	-24 2	A16NQJR		ST8000NM0055	8001563222016	1	1 83	223352360	133	96	91	ſ	10	0
smart_187_raw	8	2019-06	-24 Z	JV02XWG		ST12000NM0007	12000138625024		0 73	22090816	134	94	89) (10	D
smart_189_raw	9	2019-06	-24 2	JV1CSVX		ST12000NM0007	12000138625024		1 83	207001192	135	98	98	i r	10	D
III smart_190_raw	10	2019-06	-24 7	JV02XWA		ST12000NM0007	12000138625024		1 78	66102672	134	106	97	i r	10	0
III smart_198_raw																
P Triggers																

Disk Drive Dataset

In [2]:	M	<pre>import pandas as pd data = pd.read_csv('HDD.csv') data.head()</pre>
---------	---	--

Out[2]:

	0.419355		3.2	4./	4.7.1	1.4	1.4.1	NORMAL
0	0.419355	6.4	3.2	4.5	4.5	1.5	1.5	NORMAL
1	0.225806	6.9	3.1	4.9	4.9	1.5	1.5	NORMAL
2	0.290323	5.5	2.3	4.0	4.0	1.3	1.3	NORMAL
3	0.290323	6.5	2.8	4.6	4.6	1.5	1.5	NORMAL
4	0.451613	5.7	2.8	4.5	4.5	1.3	1.3	NORMAL

https://f001.backblazeb2.com/file/Backblaze-Hard-Drive-Data/data Q2 2019.zip

2019 Storage Developer Conference. © Dell EMC. All Rights Reserved.

Package ConfClr

September 23-26, 2019 Santa Clara, CA



In [1]: ▶ !pip install ConfClr

Collecting ConfClr Using cached https://files.pythonhosted.org/packages/ee/83/6948bdb787bc048020410e50a10ae1ca34ce90c9f55da81fe31b84b44ab3/Co nfClr-0.1.4-py3-none-any.whl Installing collected packages: ConfClr Successfully installed ConfClr-0.1.4

SD® D¢LLEMC

2019 Storage Developer Conference. © Dell EMC. All Rights Reserved.

Prediction with Explainable Reliability

Compute p-values for each label and calculate Confidence and Credibility

In [4]: ▶ import ConfClr

CC.conformalScore(x,y)

Out[4]:

	p-FAILED	p-NORMAL	Status	Credibility	Confidence
0	0.360656	0.147541	Failed	0.360656	0.639344
1	0.213115	0.147541	Failed	0.213115	0.786885
2	0.180328	0.508197	Normal	0.508197	0.819672
3	0.229508	0.147541	Failed	0.229508	0.770492
4	0.163934	0.131148	Failed	0.163934	0.836066
5	0.180328	0.147541	Failed	0.180328	0.819672
6	0.147541	0.721311	Normal	0.721311	0.852459
7	0.163934	0.081967	Failed	0.163934	0.836066
8	0.655738	0.147541	Failed	0.655738	0.344262
9	0.180328	0.721311	Normal	0.721311	0.819672
10	0.163934	0.737705	Normal	0.737705	0.836066
11	0.065574	0.655738	Normal	0.655738	0.934426
12	0.163934	0.852459	Normal	0.852459	0.836066
13	0.622951	0.147541	Failed	0.622951	0.377049
14	0.721311	0.147541	Failed	0.721311	0.278689
15	0.770492	0.147541	Failed	0.770492	0.229508
16	0.163934	0.147541	Failed	0.163934	0.836066
17	0.163934	0.803279	Normal	0.803279	0.836066
18	0.852459	0.147541	Failed	0.852459	0.147541
19	0.409836	0.147541	Failed	0.409836	0.590164

- Confidence level of 95% means that in 95% the predicted class is also a true class
- Credibility = p_{max}
- Confidence = $1 p_{max}^{2nd}$

Interpretation

September 23-26, 2019 Santa Clara, CA

Plot Confidence and Credibility of each new Label





Figure 4: Interpreting quality of prediction

Framework Application

September 23-26, 2019 Santa Olara, CA

Disk Serial Number	Status	Confidence
Z305B2QN	FAILED	0.950820
ZA16NQJR	FAILED	0.918003
ZJV1CSVX	FAILED	0.868852
ZA18CEBF	FAILED	0.737705
ZJV02XWA	FAILED	0.672131
ZJV0XJQ0	FAILED	0.426230
ZJV02XWV	FAILED	0.114754
PL2331LAG9TEEJ	FAILED	0.098361

Figure 4: Ranking the failed drive (Confidence)

- Priority based insightful action
- Efficient Inventory management

SD® DØLLEMC

September 23-26, 201

Sample Interface - Application

							Q ☆			:
									-	
					Q Search					
SYSTEMS	roduction									Ĩ
	🛢 Cap	pacity								
Tota	al Capacity	76.8 TB	Disk Failure F	precast in next 5 da	ays	Storage Usage				
Savi Provi Used	Used Free Junconfigured Drives Ings Island	55.3 TB 21.5 TB 0 B 724 TB 55.3 TB	Disk SN P9J81144W K4Kf012L YVKW267K ZSEUIQ4K LM7RQWD9 X037U92 MJKUIYT5 PRTL385K 02R645JK	Confidence 0.923 0.881 0.853 0.761 0.752 0.751 0.749 0.727 0.714	Credibility 0.777 0.538 0.962 0.973 0.965 0.978 0.538 0.874	Snapshots LUNs File Systems VMware	55 _{TB} Used	26.7 TB 6.6 TB 7.2 TB 14.5 TB	A characteristic state in the	

SDC DEELEMC

2019 Storage Developer Conference. © Dell EMC. All Rights Reserved.

Conclusion and Discussion



Challenges

September 23-26, 2019 Santa Clara, CA SD® DØLLEMC

- Designing nonconformity score
- Application in Time Series

Summary

September 23-26, 2019 Santa Clara, CA

- Standard prediction models are based on bare predictions
 - Classification: discrete classes
 - Regression: real-valued point prediction
- Model with high accuracy does not guarantee reliable forecasting
- Conformal prediction can be used with any machine learning
- Conformal Framework
 - Confidence of each single prediction
 - Credibility tells about the quality of data point

SD® DØLLEMC

Why should we care?

September 23-26, 2019 Santa Clara, CA

It matters how much we can rely upon a given prediction in risk-sensitive applications.



Horizon 2020 European Union Funding for Research & Innovation

Drug design



InsuranceInvestment



- Development of new drugs
- Precision medicine & Genomics

Seven scientific experts in Italy were convicted of manslaughter and sentenced to six years in prison for failing to give warning before the April 2009 earthquake that killed 309 people.

Diacu F. Is failure to predict a crime? October 2012. New York Times; October 2012. https://www.nytimes.com/2012/10/27/opinion/a-failed-earthquake-prediction-a-crime.html

SDY (9

DELLEMC

References

September 23-26, 2019 Santa Clara, CA

- V. Vovk, A. Gammerman, and G. Shafer, Algorithmic learning in a random world. Springer, 2005.
- G. Shafer and V. Vovk, "A tutorial on conformal prediction," The Journal of Machine Learning Research, vol. 9, pp.
- Balasubramanian, V., Ho, S. and Vovk, V. (2014). Conformal Prediction for Reliable Machine Learning.

SD[®]

D%LLEMC

September 23-26, 2019 Santa Clara, CA

"As far as the laws of mathematics refer to reality, they are not certain; and as far as they are certain, they do not refer to reality." – Albert Einstein

SD©

