New Directions In Automated Traffic Analysis

github.com/nprint/

Jordan Holland, Paul Schmitt, Nick Feamster, Prateek Mittal

ML + Networking

- Device fingerprinting
- OS detection
- Website fingerprinting
- Application identification
- Protocol fingerprinting
- Anomaly detection
- …

Classic ML Pipeline

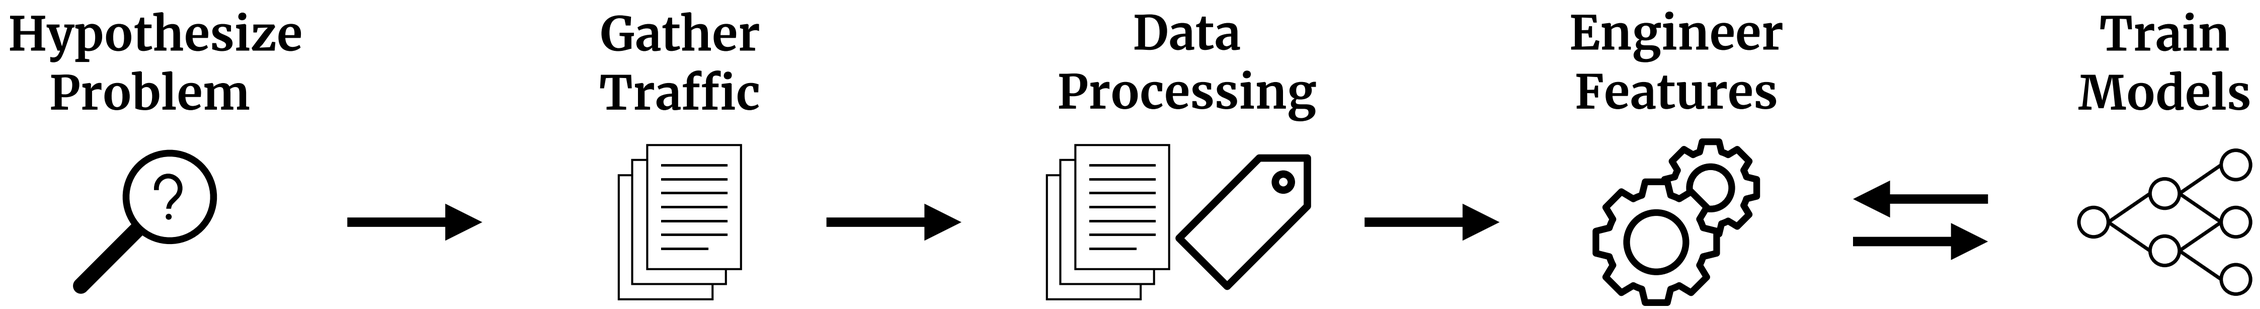

Bespoke Solutions

Application Identification

Hypothesize Problem → Gather Traffic → Data Processing → Engineer Features ⇄ Train Models

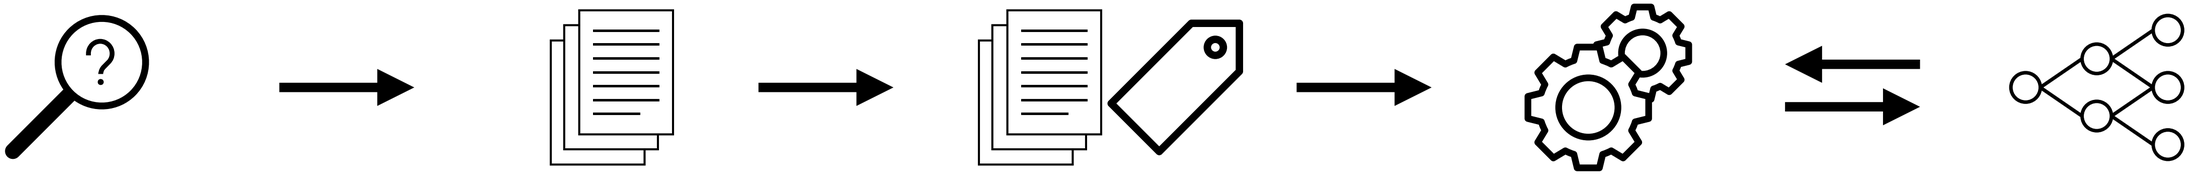

Anomaly Detection

Hypothesize Problem → Gather Traffic → Data Processing → Engineer Features ⇄ Train Models

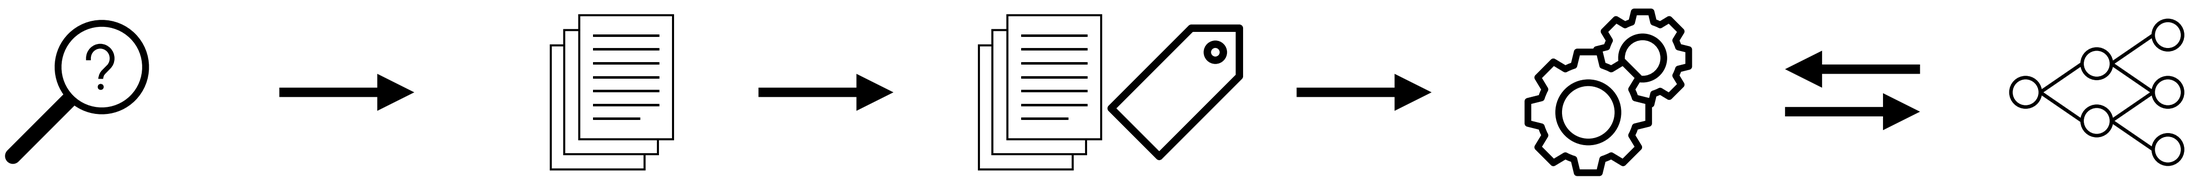

Generalizable Solutions?

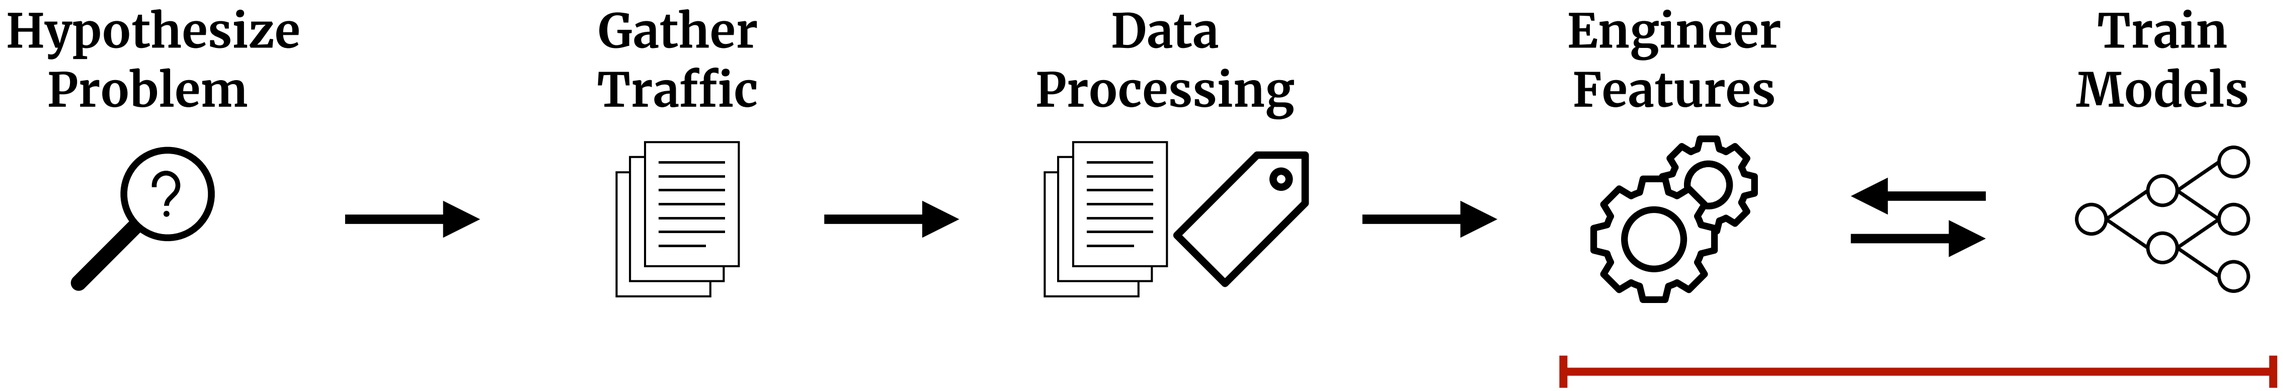

Are We Working Too Hard?

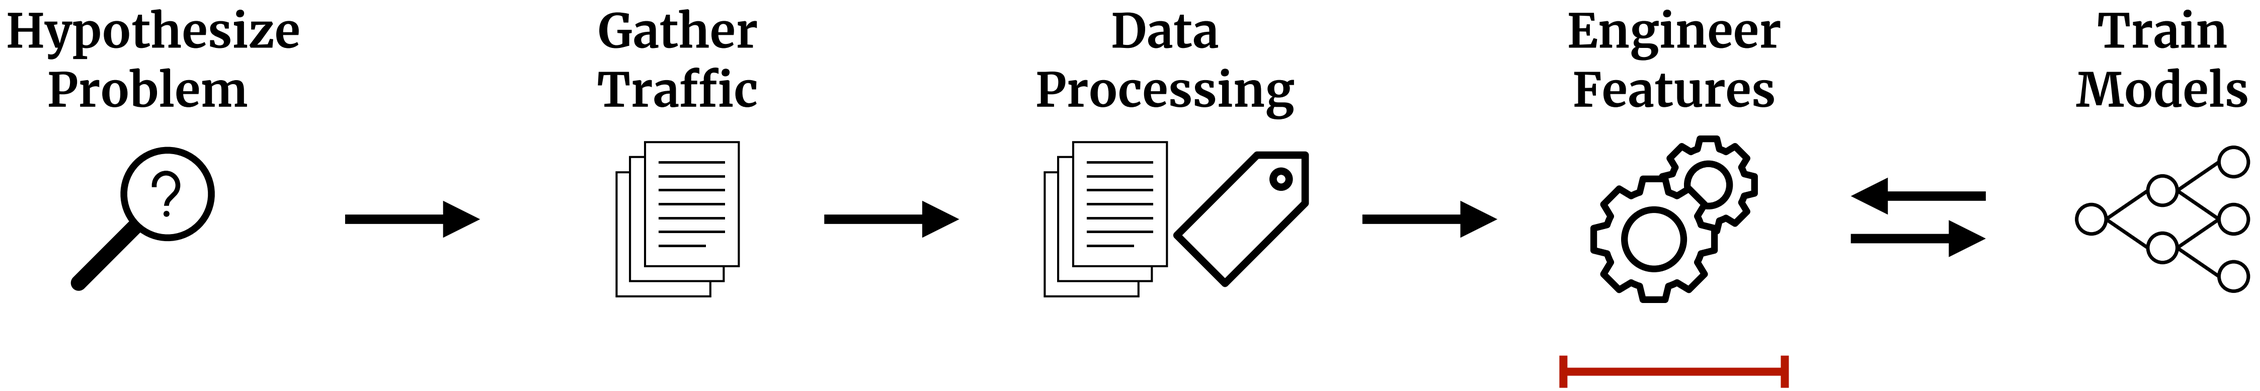

nPrint

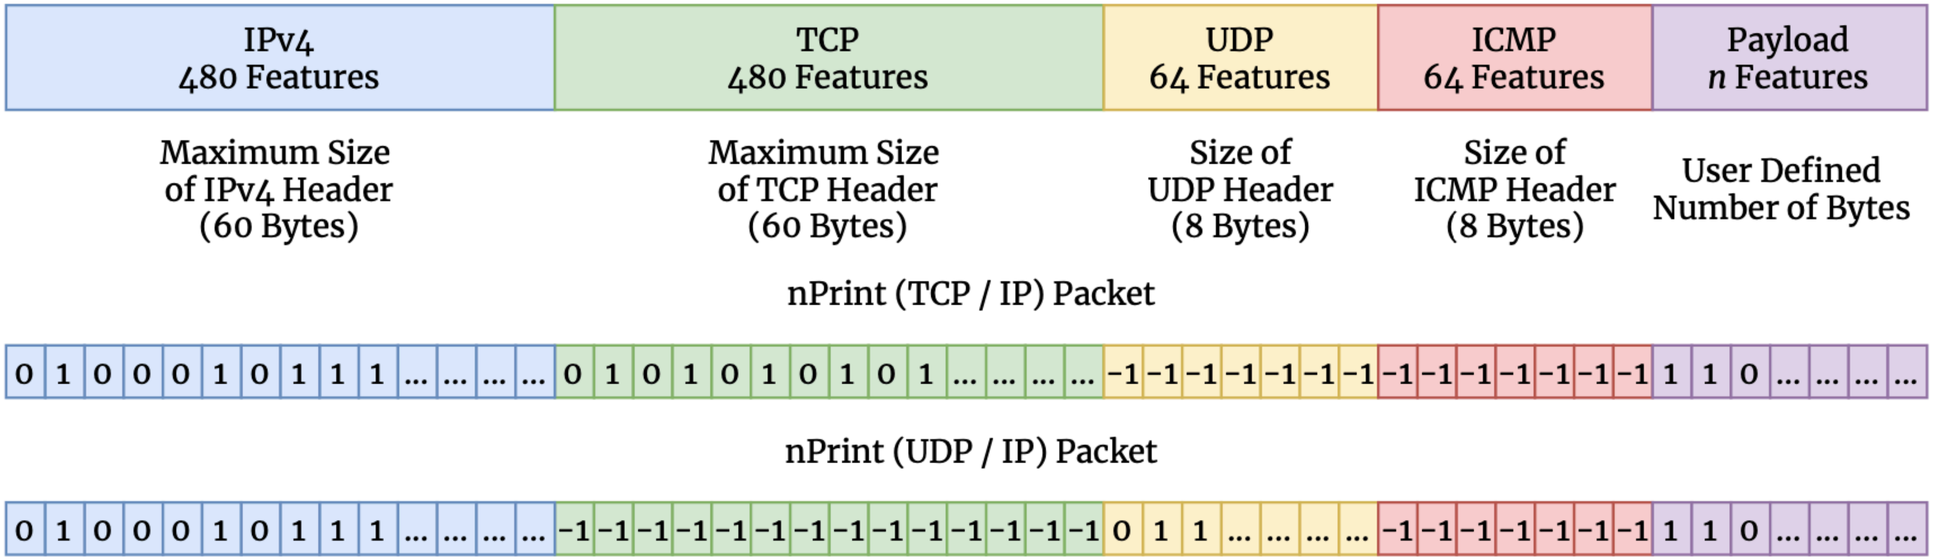
IPv4
480 Features
TCP
480 Features
UDP
64 Features
ICMP
64 Features
Payload
n Features
Maximum Size of IPv4 Header (60 Bytes)
Maximum Size of TCP Header (60 Bytes)
Size of UDP Header (8 Bytes)
Size of ICMP Header (8 Bytes)
User Defined Number of Bytes
nPrint (TCP / IP) Packet
nPrint (UDP / IP) Packet

Automate This Step?

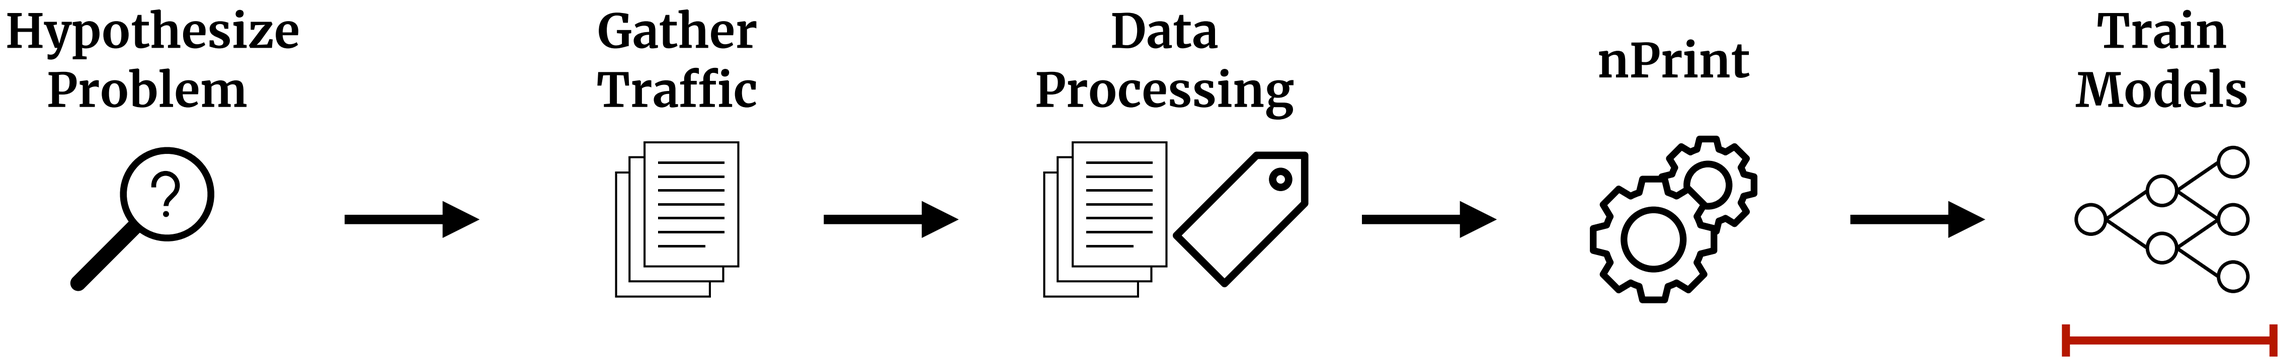

Automated Machine Learning

- Model selection
- Feature selection
- Hyperparameter search

nPrintML

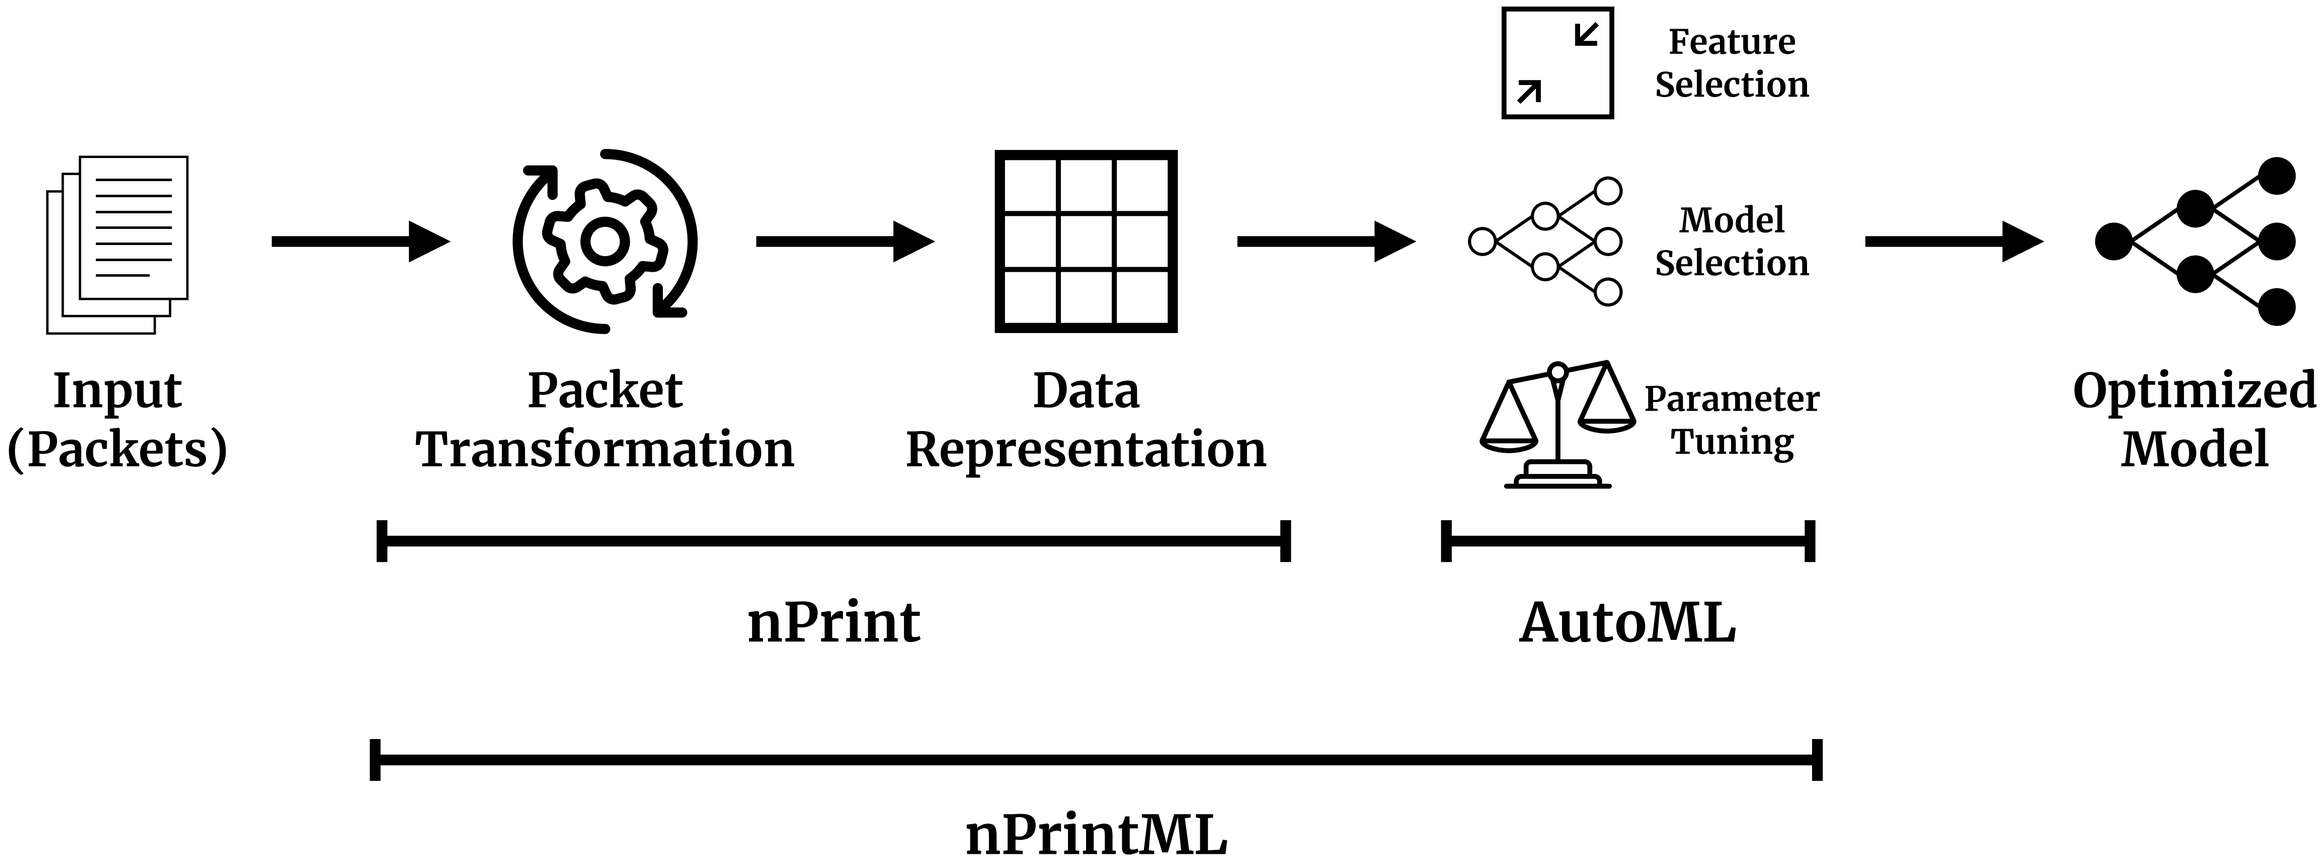
Feature Selection
Model Selection
Parameter Tuning
Input (Packets)
Packet Transformation
Data Representation
Optimized Model
nPrint
AutoML
nPrintML

8 Discrete Case Studies

Problem Overview			nPrintML					Comparison	
Description	Dataset	# Classes	Configuration eAppendix A.4)	Sample Size (# Packets)	Balanced Accuracy	ROC AUC	Macro F1	Score	Source
Active Device Fingerprinting (§5.1)	Network Device Dataset [22]	15	-4 -t -i	21	95.4	99.7	95.5	92.9 (Macro-F1)	ML-Enhanced Nmap [31]
Passive OS Detection (§5.2)	CICIDS 2017 [48]	3	-4 -t	1	99.5	99.9	99.5	81.3 (Macro-F1)	p0f [40]
				10	99.9	100	99.9		
				100	99.9	100	99.9		
		13		100	77.1	97.5	76.9	No Previous Work	
Application Identification via DTLS Handshakes (§5.3)	DTLS Handshakes [32]	7	-4	43	99.8	96.9	99.7	99.8 (Average Accuracy)	Hand-Curated Features [32]
			-u		99.9	99.7	99.5		
			-p 10		95.0	78.8	77.4		
			-p 25		99.9	99.7	99.7		
			-p 100		99.9	99.7	99.7		
			-4 -u -p 10		99.8	99.9	99.8		
Malware Detection for IoT Traces (§5.4.1)	netML IoT [6, 28]	2	-4 -t -u	10	92.4	99.5	93.2	99.9 (True Positive Rate)	NetML Challenge Leaderboard [37]
		19			86.1	96.9	84.1	39.7 (Balanced F1)	
Type of Traffic in Capture (§5.4.1)	netML Non-VPN [6, 12]	7	-4 -t -u -p 10	10	81.9	98.0	79.5	67.3 (Balanced F1)	
					76.1	94.2	75.8		
		18	-4 -t -u		66.2	91.3	63.7	42.1 (Balanced F1)	
		31			60.9	92.2	57.6	34.9 (Balanced F1)	
Intrusion Detection (§5.4.1)	netML CICIDS 2017 [6, 48]	2	-4 -t -u	5	99.9	99.9	99.9	98.9 (True Positive Rate)	
		8			99.9	99.9	99.9	99.2 (Balanced F1)	
Determine Country of Origin for Android & iOS Application Traces (§5.4.2)	Cross Platform [44]	3	-4 -t -u -p 50	25	96.8	90.2	90.4	No Previous Work	
Identify streaming video (DASH) service via device SYN packets (§5.4.3)	Streaming Video Providers [10]	4	-4 -t -u -R	10	77.9	96.0	78.9	No Previous Work	
				25	90.2	98.6	90.4		
				50	98.4	99.9	98.6		

nPrintML Results

Outperforming Hand-engineered Solutions

Problem Overview			nPrintML					Comparison	
Description	Dataset	# Classes	Configuration eAppendix A.4)	Sample Size (# Packets)	Balanced Accuracy	ROC AUC	Macro F1	Score	Source
Active Device Fingerprinting (§5.1)	Network Device Dataset [22]	15	-4 -t -i	21	95.4	99.7	95.5	92.9 (Macro-F1)	ML-Enhanced Nmap [31]
Passive OS Detection (§5.2)	CICIDS 2017 [48]	3	-4 -t	1	99.5	99.9	99.5	81.3 (Macro-F1)	p0f [40]
				10	99.9	100	99.9		
				100	99.9	100	99.9		
		13		100	77.1	97.5	76.9	No Previous Work	
Application Identification via DTLS Handshakes (§5.3)	DTLS Handshakes [32]	7	-4	43	99.8	96.9	99.7	99.8 (Average Accuracy)	Hand-Curated Features [32]
			-u		99.9	99.7	99.5		
			-p 10		95.0	78.8	77.4		
			-p 25		99.9	99.7	99.7		
			-p 100		99.9	99.7	99.7		
			-4 -u -p 10		99.8	99.9	99.8		
Malware Detection for IoT Traces (§5.4.1)	netML IoT [6, 28]	2	-4 -t -u	10	92.4	99.5	93.2	99.9 (True Positive Rate)	NetML Challenge Leaderboard [37]
		19			86.1	96.9	84.1	39.7 (Balanced F1)	
Type of Traffic in Capture (§5.4.1)	netML Non-VPN [6, 12]	7	-4 -t -u -p 10	10	81.9	98.0	79.5	67.3 (Balanced F1)	
					76.1	94.2	75.8		
		18	-4 -t -u		66.2	91.3	63.7	42.1 (Balanced F1)	
		31			60.9	92.2	57.6	34.9 (Balanced F1)	
Intrusion Detection (§5.4.1)	netML CICIDS 2017 [6, 48]	2	-4 -t -u	5	99.9	99.9	99.9	98.9 (True Positive Rate)	
		8			99.9	99.9	99.9	99.2 (Balanced F1)	
Determine Country of Origin for Android & iOS Application Traces (§5.4.2)	Cross Platform [44]	3	-4 -t -u -p 50	25	96.8	90.2	90.4	No Previous Work	
Identify streaming video (DASH) service via device SYN packets (§5.4.3)	Streaming Video Providers [10]	4	-4 -t -u -R	10	77.9	96.0	78.9	No Previous Work	
				25	90.2	98.6	90.4		
				50	98.4	99.9	98.6		

nPrintML Results

Thank You!

github.com/nprint/