EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

MATHEMATISCH-
NATURWISSENSCHAFTLICHE FAKULTÄT

Detection of Malicious PDF Files Based on Hierarchical Document Structure

Nedim Šrndić and Pavel Laskov

Network and Distributed System Security Symposium (*NDSS*),
February 24-27, 2013; San Diego, California, USA

The Security of PDF

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- PDF is complex
  - over 7,000 pages of documentation
  - JavaScript, Flash, fonts, images, forms
- Adobe (Acrobat) Reader popular and vulnerable
  - CVE-2013-0641, CVE-2013-0640, CVE-2013-0627, CVE-2013-0626, CVE-2013-0624, CVE-2013-0623, CVE-2013-0622, CVE-2013-0621, CVE-2013-0620, CVE-2013-0619, CVE-2013-0618, CVE-2013-0617, CVE-2013-0616, CVE-2013-0615, CVE-2013-0614, CVE-2013-0613, CVE-2013-0612, CVE-2013-0611, CVE-2013-0610, CVE-2013-0609, CVE-2013-0608, CVE-2013-0607, CVE-2013-0606, CVE-2013-0605, CVE-2013-0604, CVE-2013-0603, CVE-2013-0602, CVE-2013-0601, CVE-2012-4162, CVE-2012-4161, CVE-2012-4160, CVE-2012-4159, CVE-2012-4158, CVE-2012-4157, CVE-2012-4156, CVE-2012-4155, CVE-2012-4154, CVE-2012-4153, CVE-2012-4152, CVE-2012-4151, CVE-2012-4150, CVE-2012-4149, CVE-2012-4148, CVE-2012-4147, CVE-2012-2051, CVE-2012-2050, CVE-2012-2049, CVE-2012-1530, CVE-2012-1525, CVE-2012-0777, CVE-2012-0775, CVE-2012-0774, CVE-2011-4373, CVE-2011-4372, CVE-2011-4371, CVE-2011-4370, CVE-2011-4369, CVE-2011-2462, CVE-2011-2442, CVE-2011-2441, CVE-2011-2440, CVE-2011-2439, CVE-2011-2438, CVE-2011-2437, CVE-2011-2436, CVE-2011-2435, CVE-2011-2434, CVE-2011-2433, CVE-2011-2432, CVE-2011-2431, CVE-2011-2106, CVE-2011-2105, CVE-2011-2104, CVE-2011-2103, CVE-2011-2102, CVE-2011-2101, CVE-2011-2100, CVE-2011-2099, CVE-2011-2098, CVE-2011-2097, CVE-2011-2096, CVE-2011-2095, CVE-2011-2094, CVE-2011-0611, CVE-2011-0610, CVE-2011-0609, CVE-2011-0606, CVE-2011-0605, CVE-2011-0604, CVE-2011-0603, CVE-2011-0602, CVE-2011-0600, CVE-2011-0599, CVE-2011-0598, CVE-2011-0596, CVE-2011-0595, CVE-2011-0594, CVE-2011-0593, CVE-2011-0592, CVE-2011-0591, CVE-2011-0590, CVE-2011-0589, CVE-2011-0588, CVE-2011-0587, CVE-2011-0586, CVE-2011-0585, CVE-2011-0570, CVE-2011-0568, CVE-2011-0567, CVE-2011-0566, CVE-2011-0565, CVE-2011-0564, CVE-2011-0563, CVE-2011-0562, CVE-2010-4091, CVE-2010-3658, CVE-2010-3657, CVE-2010-3656, CVE-2010-3654, CVE-2010-3632, CVE-2010-3631, CVE-2010-3630, CVE-2010-3629, CVE-2010-3628, CVE-2010-3627, CVE-2010-3626, CVE-2010-3625, CVE-2010-3624, CVE-2010-3623, CVE-2010-3622, CVE-2010-3621, CVE-2010-3620, CVE-2010-3619, CVE-2010-2890, CVE-2010-2889, CVE-2010-2888, CVE-2010-2887, CVE-2010-2884, CVE-2010-2883, CVE-2010-2862, CVE-2010-2212, CVE-2010-2211, CVE-2010-2210, CVE-2010-2209, CVE-2010-2208, CVE-2010-2207, CVE-2010-2206, CVE-2010-2205, CVE-2010-2204, CVE-2010-2203, CVE-2010-2202, CVE-2010-2201, CVE-2010-2168, CVE-2010-1297, CVE-2010-1295, CVE-2010-1285, CVE-2010-1278, CVE-2010-1241, CVE-2010-1240, CVE-2010-0204, CVE-2010-0203, CVE-2010-0202, CVE-2010-0201, CVE-2010-0199, CVE-2010-0198, CVE-2010-0197, CVE-2010-0196, CVE-2010-0195, CVE-2010-0194, CVE-2010-0193, CVE-2010-0192, CVE-2010-0191, CVE-2010-0190, CVE-2010-0188, CVE-2010-0186, CVE-2009-4324, CVE-2009-3959, CVE-2009-3958, CVE-2009-3957, CVE-2009-3956, CVE-2009-3955, CVE-2009-3954, CVE-2009-3953, CVE-2009-3462, CVE-2009-3459, CVE-2009-3458, CVE-2009-3431, CVE-2009-2998, CVE-2009-2997, CVE-2009-2996, CVE-2009-2994, CVE-2009-2993, CVE-2009-2992, CVE-2009-2991, CVE-2009-2990, CVE-2009-2988, CVE-2009-2987, CVE-2009-2986, CVE-2009-2985, CVE-2009-2983, CVE-2009-2982, CVE-2009-2981, CVE-2009-2980, CVE-2009-2979, CVE-2009-2028, CVE-2009-1862, CVE-2009-1861, CVE-2009-1859, CVE-2009-1858, CVE-2009-1857, CVE-2009-1856, CVE-2009-1855, CVE-2009-1600, CVE-2009-1599, CVE-2009-1598, CVE-2009-1597, CVE-2009-1492, CVE-2009-1062, CVE-2009-1061, CVE-2009-0928, CVE-2009-0927, CVE-2009-0889, CVE-2009-0888, CVE-2009-0658, CVE-2009-0512, CVE-2009-0511, CVE-2009-0510, CVE-2009-0509, CVE-2009-0198, CVE-2009-0193, CVE-2008-4817, CVE-2008-4815, CVE-2008-4814, CVE-2008-4813, CVE-2008-4812, CVE-2008-2992, CVE-2008-2641, CVE-2008-2549, CVE-2008-0883, CVE-2008-0726, CVE-2008-0667, CVE-2008-0655, CVE-2007-5666, CVE-2007-5663, CVE-2007-5659, CVE-2007-5020, CVE-2007-3896, CVE-2007-1199, CVE-2007-0048, CVE-2007-0047, CVE-2007-0046, CVE-2007-0045, CVE-2007-0044, CVE-2006-6236, CVE-2006-6027, CVE-2006-5857, CVE-2006-3452, CVE-2006-3093, CVE-2005-2470, CVE-2005-1625, CVE-2005-1347, CVE-2005-1306, CVE-2005-0492, CVE-2005-0035, CVE-2004-1598, CVE-2004-1153, CVE-2004-1152, CVE-2004-0631, CVE-2004-0630, CVE-2004-0629, CVE-2004-0194, CVE-2003-0508, CVE-2003-0142, CVE-2002-1764, CVE-2002-1016, CVE-2002-0030, CVE-2000-0713, CVE-1999-1576
  - 230 CVEs since 2009
  - 28 CVEs in 2013

Related Work

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

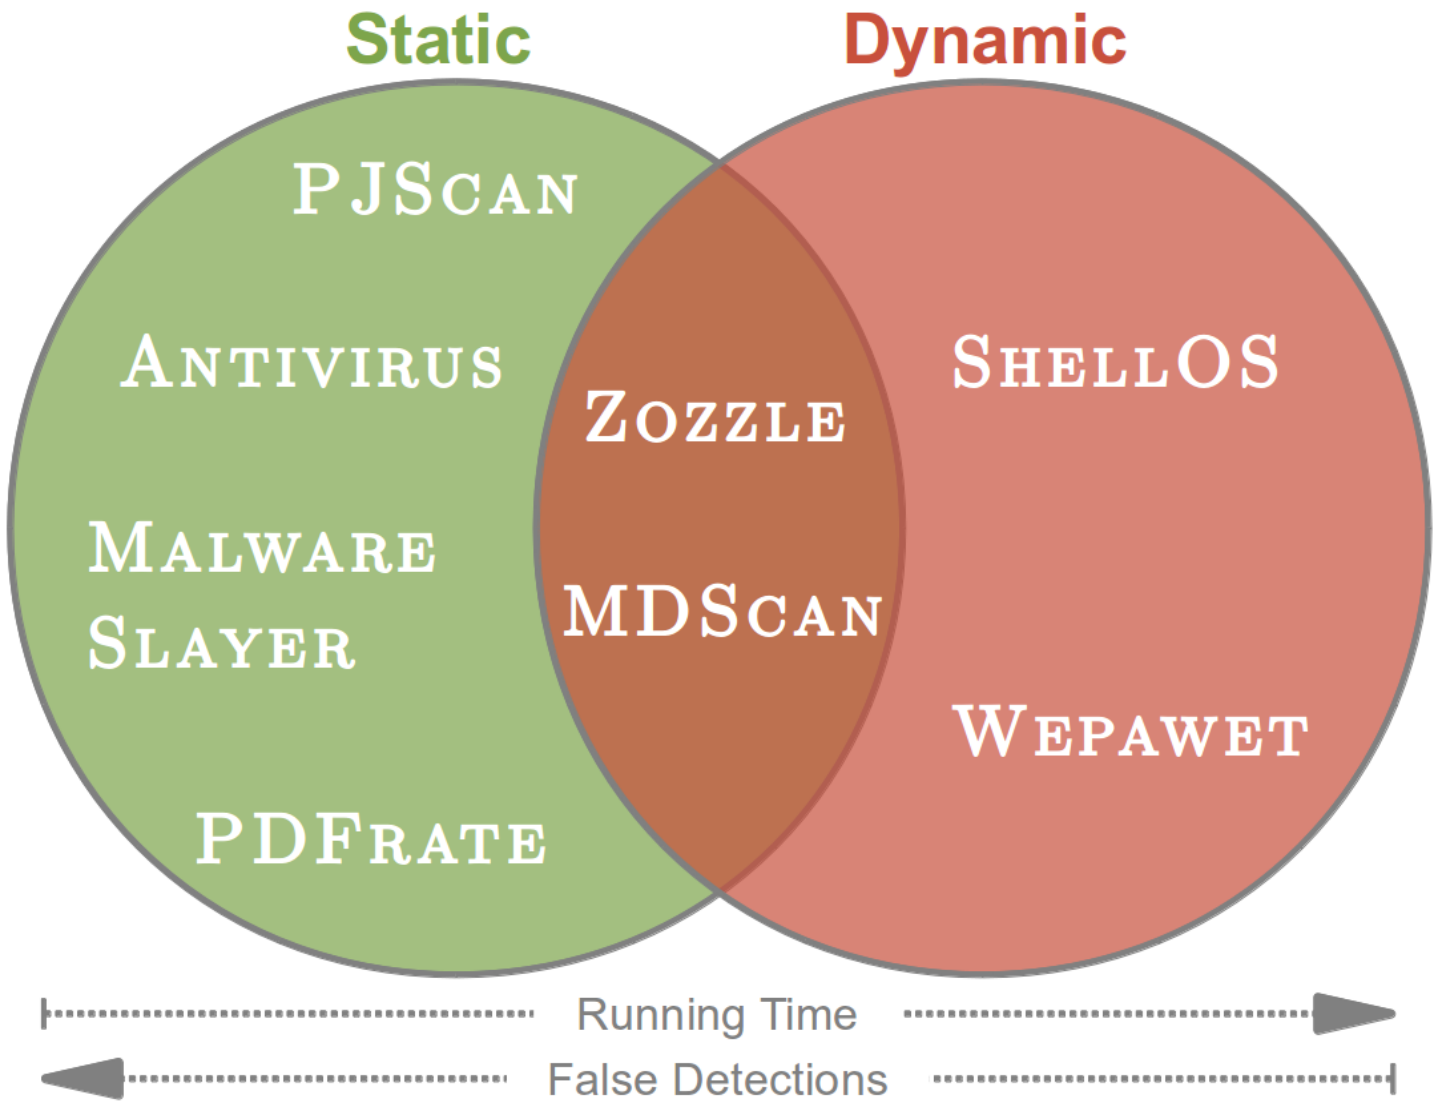

Goals

(1) Minimize false detections

(2) Minimize running time

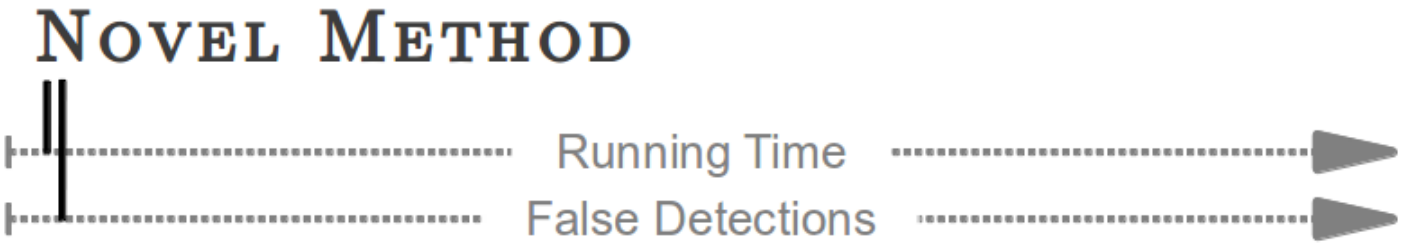

A Novel Approach

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- Is there a difference in how benign and malicious PDF files are organized internally?
  - working exploits require malicious content in specific context
- Examine structure, not content
  - static, fast and straightforward
  - no need for error-prone emulation
  - less prone to content obfuscation

PDF Document Structure

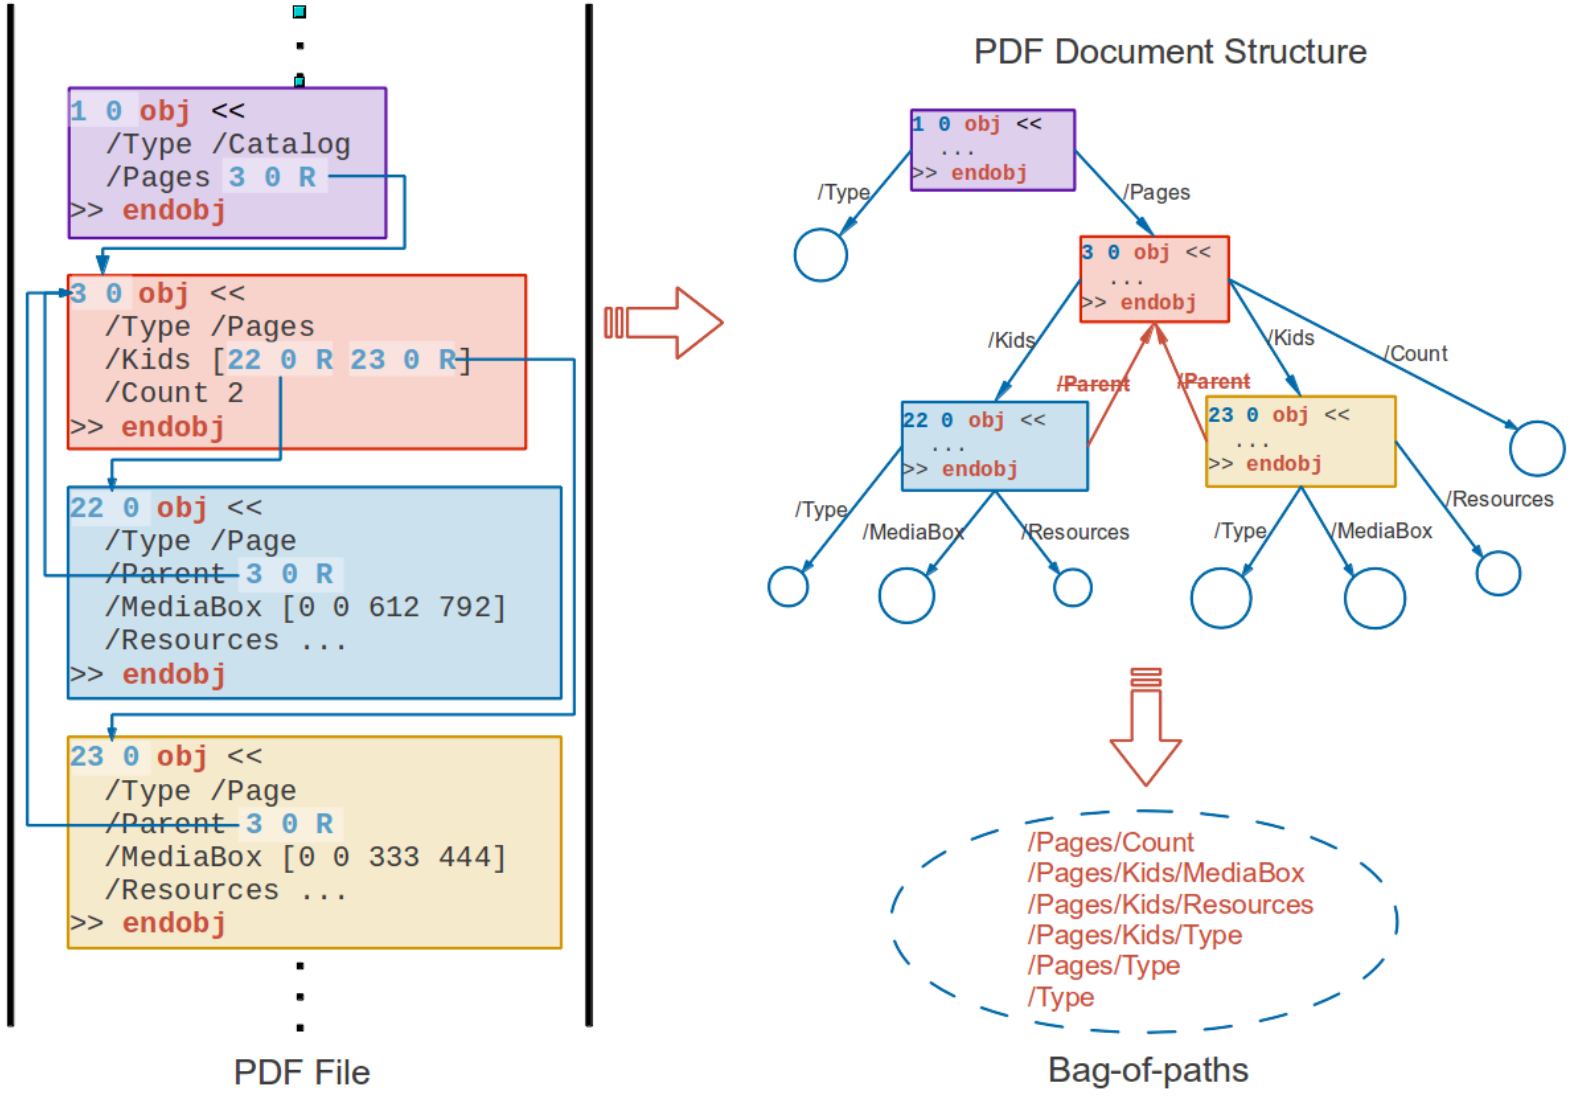

System Architecture

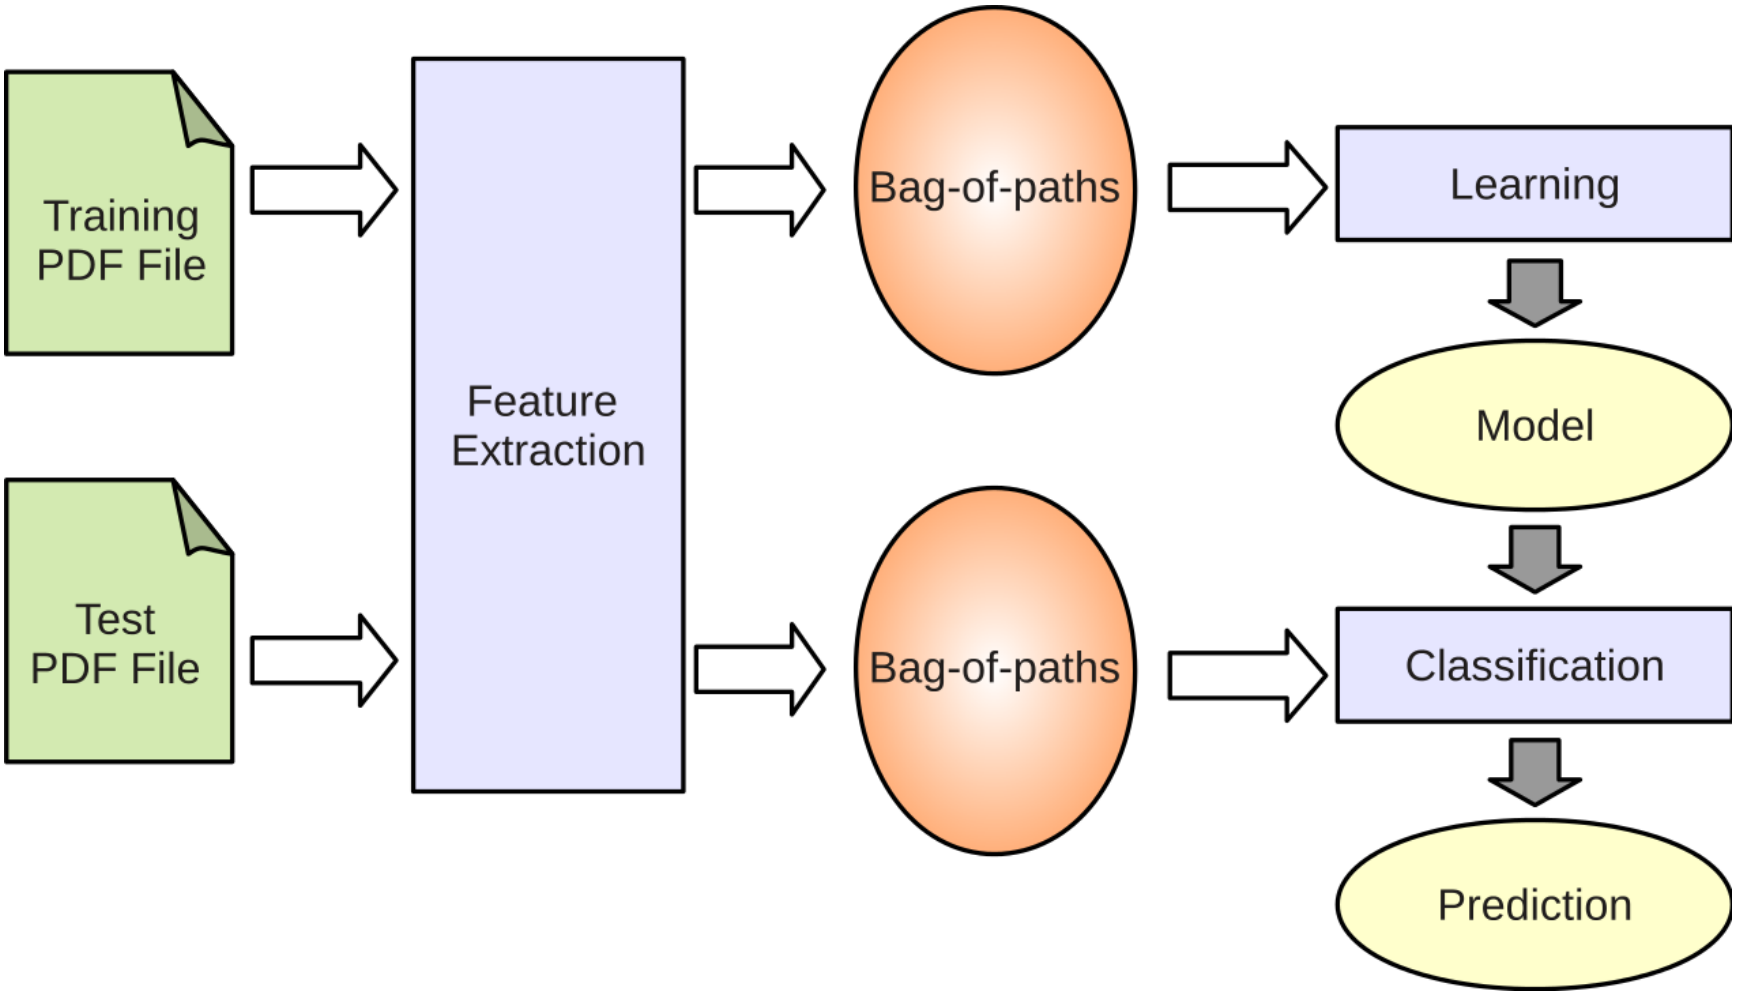

Feature Extraction

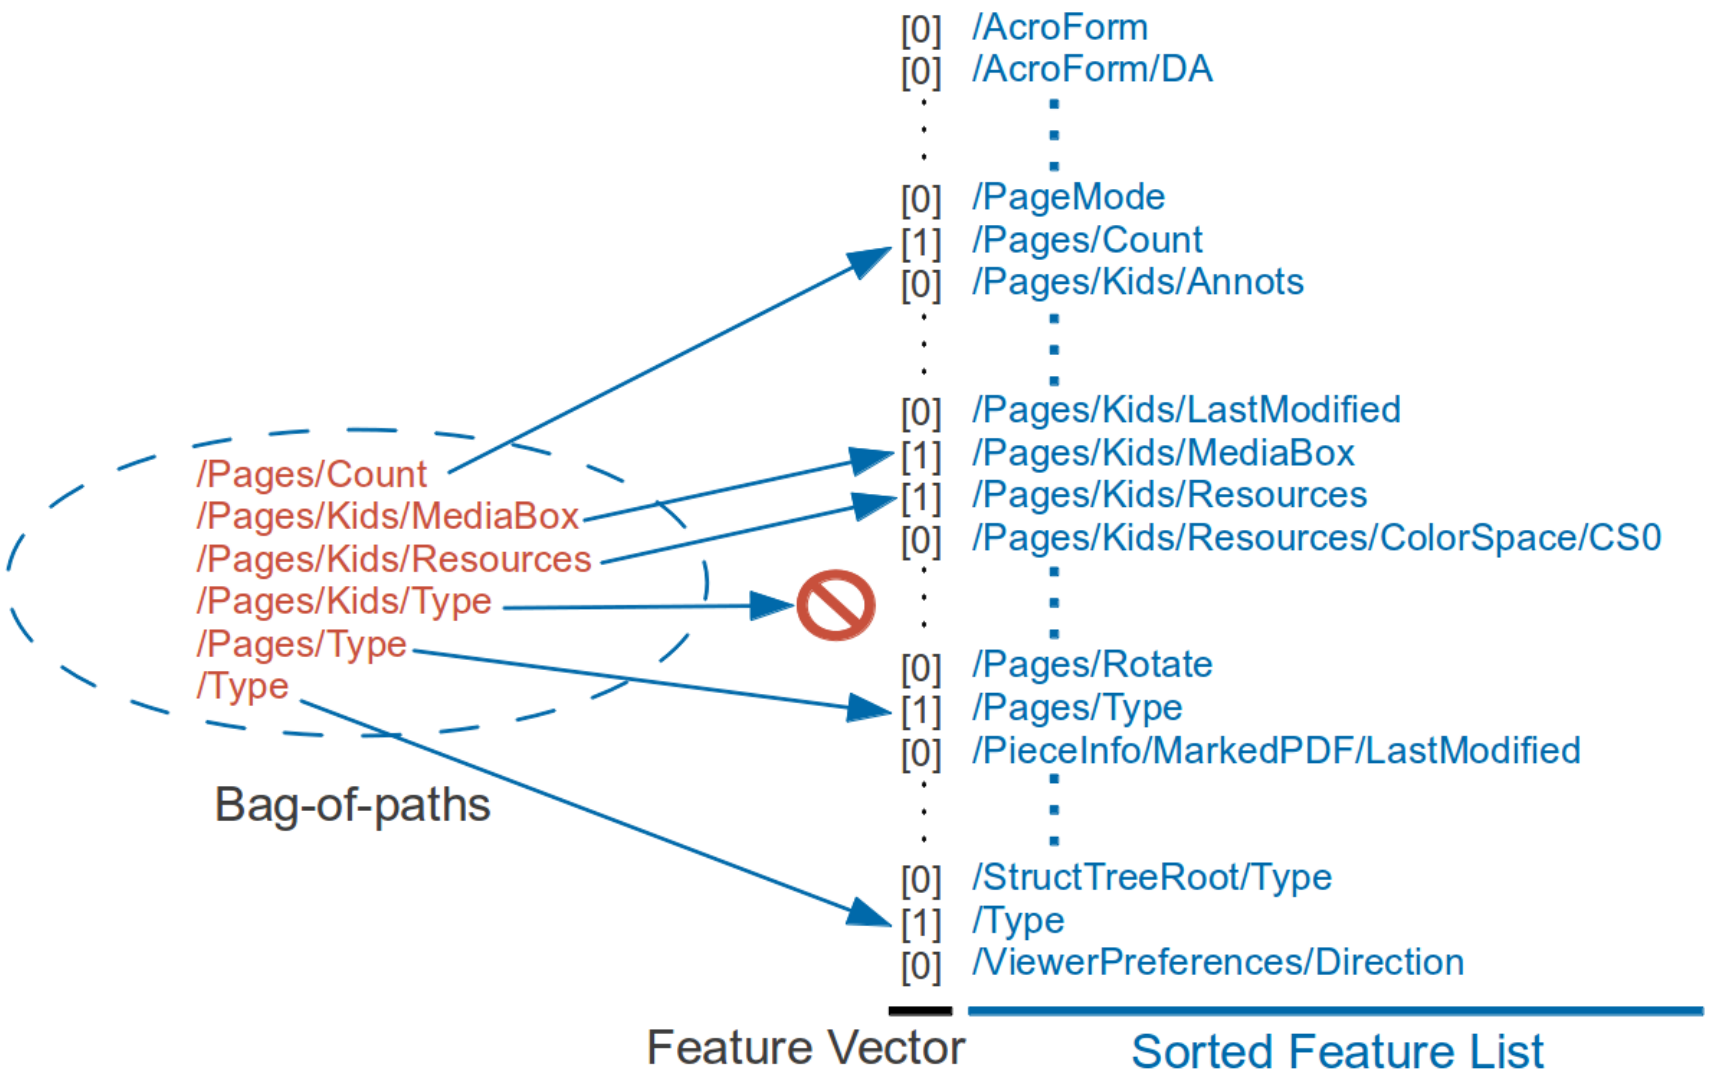

Learning Algorithms

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- Decision tree
  - implementation: C5.0
- Support Vector Machine
  - implementation: LIBSVM
  - RBF kernel, *C = 12*, *γ = 0.0025*

Experimental Datasets

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Dataset	Size
VirusTotal malicious* (old)	38,207 (1.4 GB)
VirusTotal malicious* (new)	11,409 (527 MB)
VirusTotal benign	79,200 (75 GB)
Google benign	90,384 (73 GB)
VirusTotal operational malicious*	35,526 (2.7 GB)
VirusTotal operational benign	407,037 (443 GB)
Total	**658,763 (595 GB)**

* Labeled as malicious by at least 5 out of 43 antiviruses.

Experiment 1: Standard

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- Objective: evaluate overall effectiveness
- Learning and evaluation (5-fold cross-validation):
  - malicious: VirusTotal old
  - benign: Google

	Decision tree	SVM
True Positives	38,102	38,163
False Positives	51	10
True Negatives	90,783	90,824
False Negatives	105	44
True Positive Rate	.99725	.99885
False Positive Rate	.00056	.00011
Detection Accuracy	.99879	.99958

- Goal (1) - minimize false detections - achieved!

Comparison to Antiviruses

- True positive count
- VirusTotal runs AVs through their command-line interface

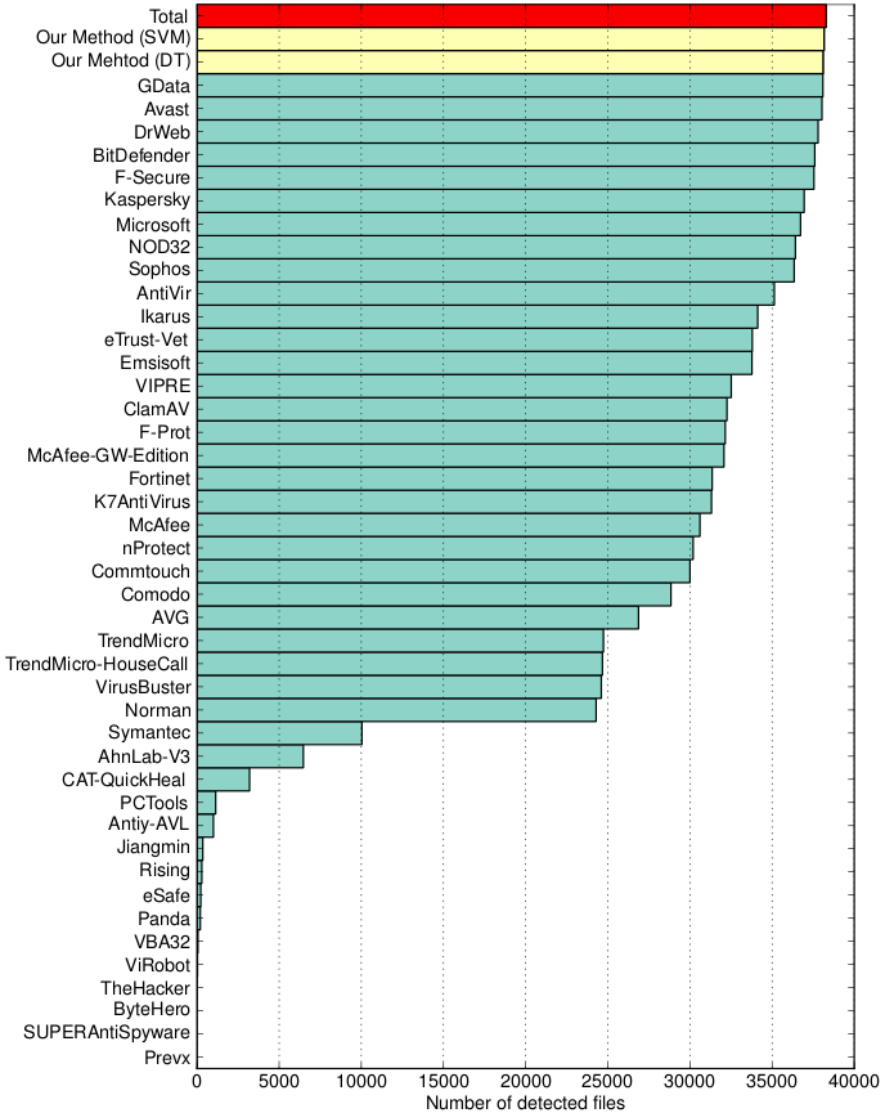

Experiment 5: 10 Weeks

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- Objective: evaluate real-world performance
- Learning and evaluation:
  - 14 weeks of VirusTotal operational ben. and mal.
  - weekly classify files gathered in weeks 5 to 14 using a model trained on the previous 4 weeks

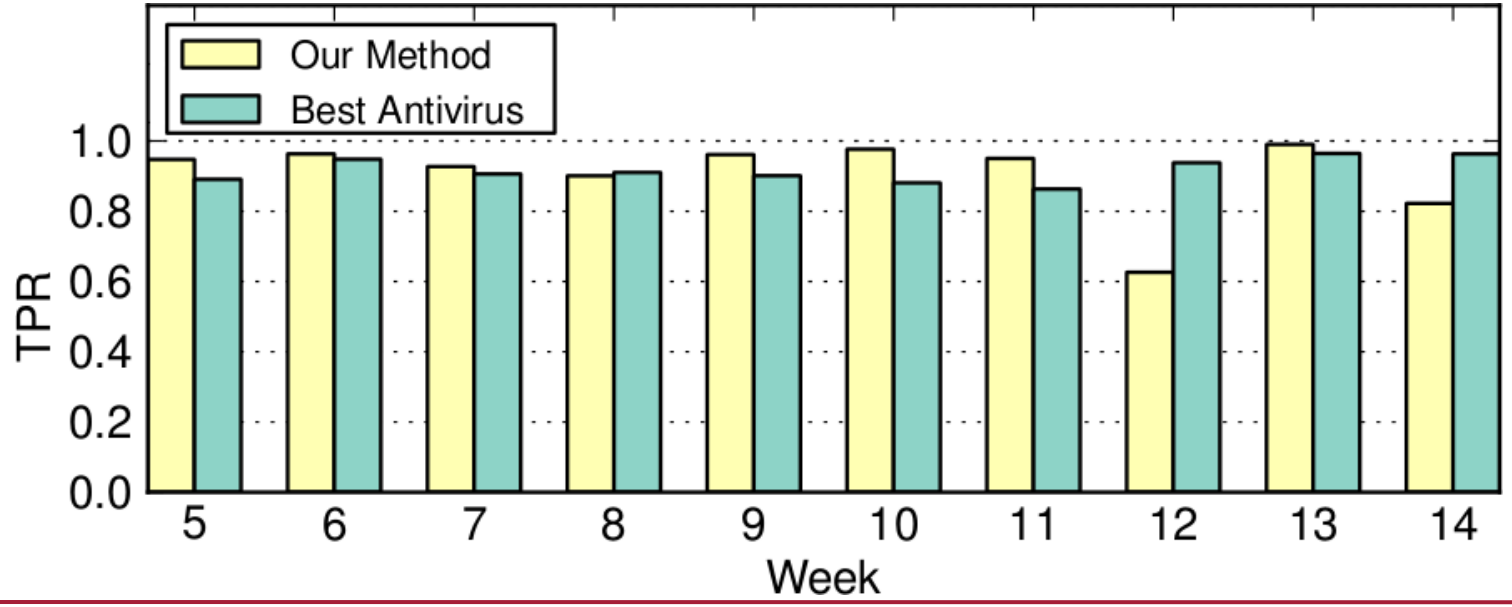

Running Time

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- Our method: 28 ms per file
- Others:
  - MDScan: 1,500 to 3,000 ms per file
  - ShellOS: 7,460 to 25,460 ms per file
  - PJScan: 23 ms per file (low detection performance)
- Goal (2) - minimize running time - achieved!

Evasion

- Goal: fool a classifier into producing a false negative by modifying a malicious sample
- Constraints: features can be added, cannot be removed
- Results:
  - decision tree - simple algorithm, successful
  - linear SVM - simple algorithm, successful
  - RBF SVM - mimicry attack, unsuccessful

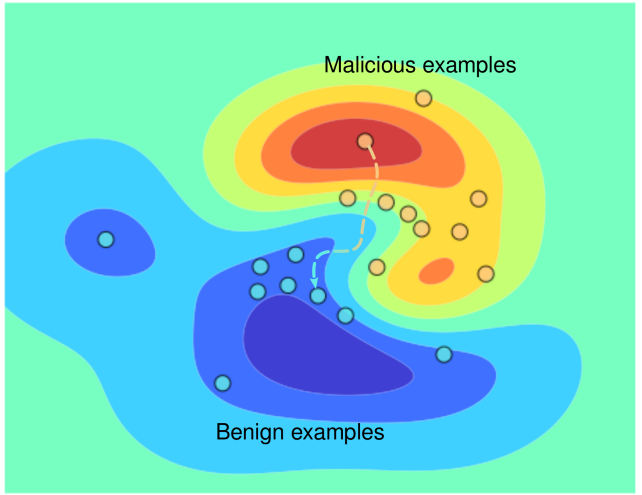

Conclusions

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

- PDF document structure **is highly discriminative** between benign and malicious PDF files
- Accurate and fast method
- Evasion is not trivial
- Explainability vs. robustness against evasion
  - do not disclose your decision trees in security applications