TokDoc – The Token Doctor

Tammo Krueger

26.03.2010
25th Symposium On Applied Computing

Outline

Web Application Firewall (WAF) – Overview

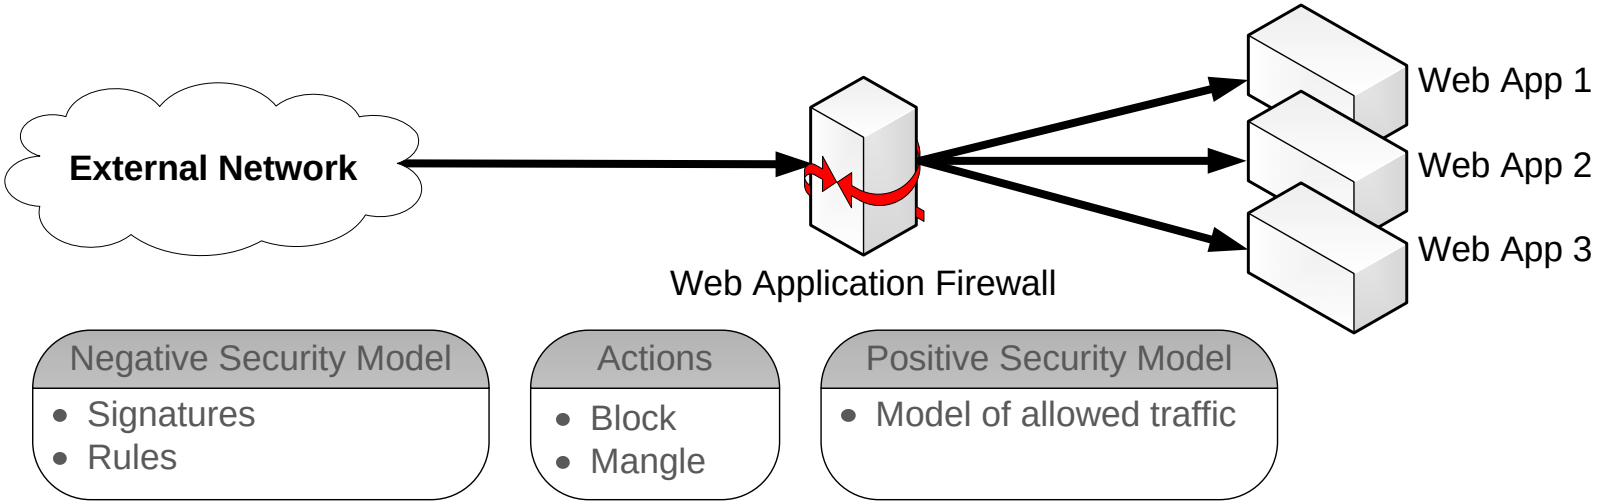

- We focus on reverse proxy setup
- Web Application Firewall monitors web applications
- Negative security model: set of fixed **signatures**/**rules**
- Positive security model: fixed allowed traffic
- re**acts** in real time to suspicious activities by
  - **blocking** the traffic
  - **replacing** parts of the request

TokDoc – Teach WAFs New Tricks

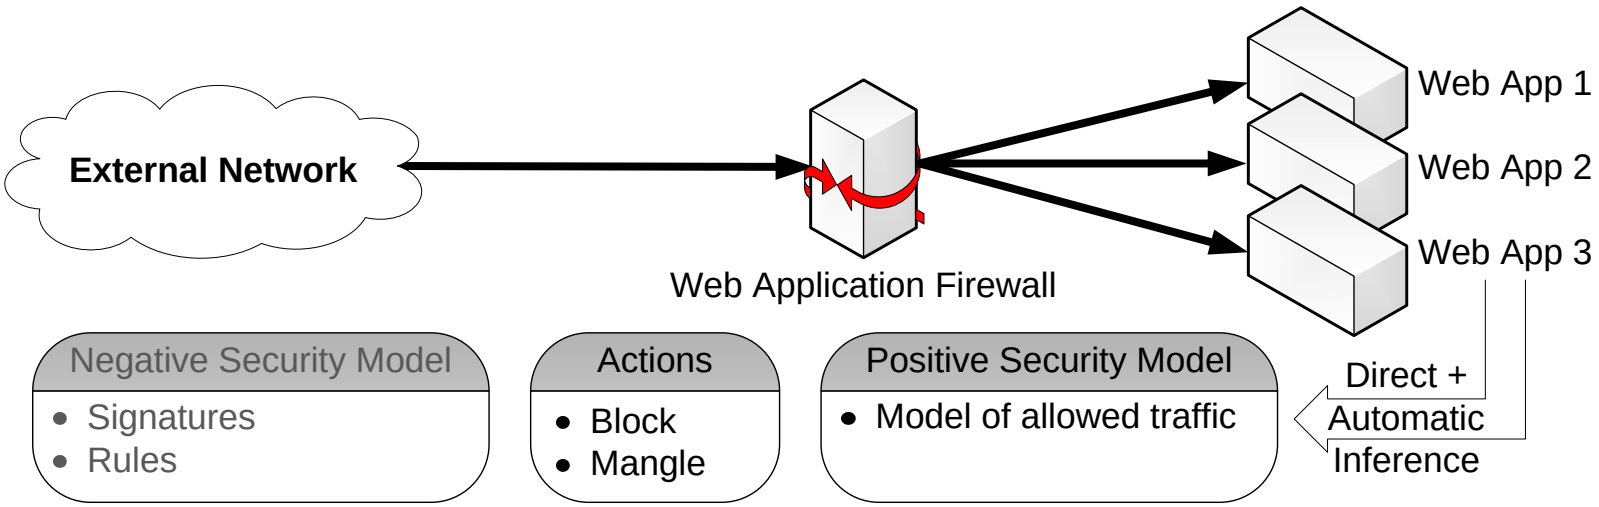

Tune up existing WAF techniques with machine learning:

- use syntactical structure for local per-token model building
- decide on this token-base,
  - which detector is suitable
  - which action is applicable
- learn these assignments **automatically** from collected data
- **no** additional attack dataset needed

Outline

1 Introduction

2 TokDoc

- Walkthrough
- Anomaly Detectors
- Healing Actions
- Setup Process

3 Evaluation

- Ensemble of Learners
- Comparison to Other Detectors
- Runtime

4 Conclusion and Further Work

TokDoc – Walkthrough 1/4

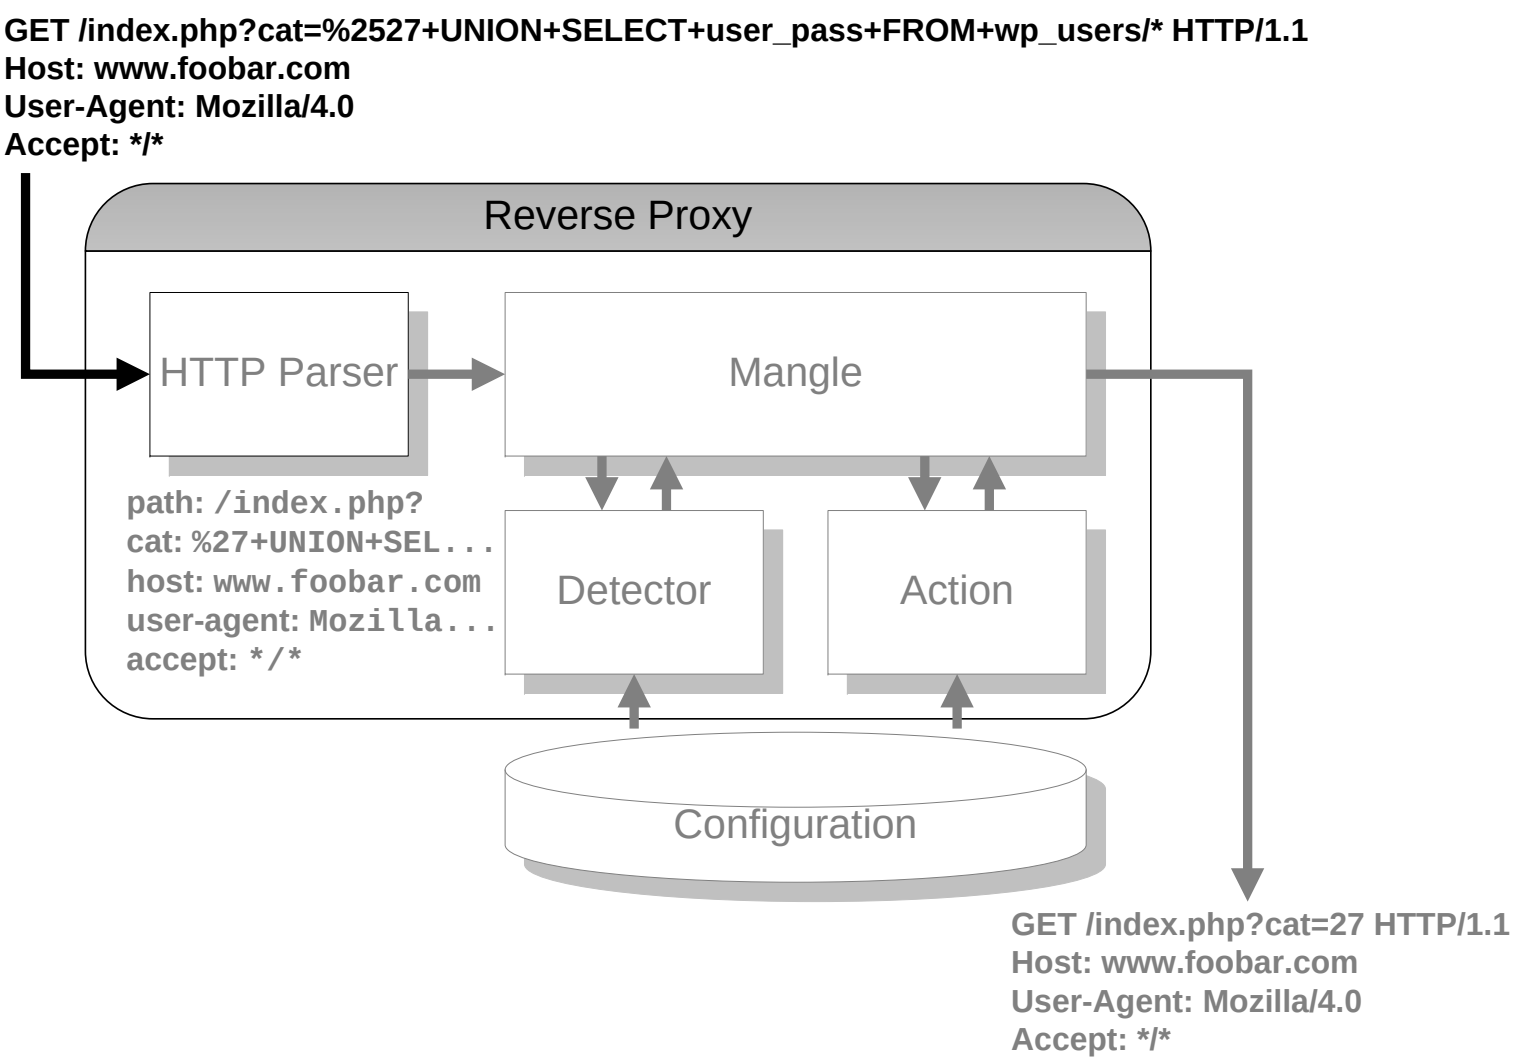

TokDoc – Walkthrough 2/4

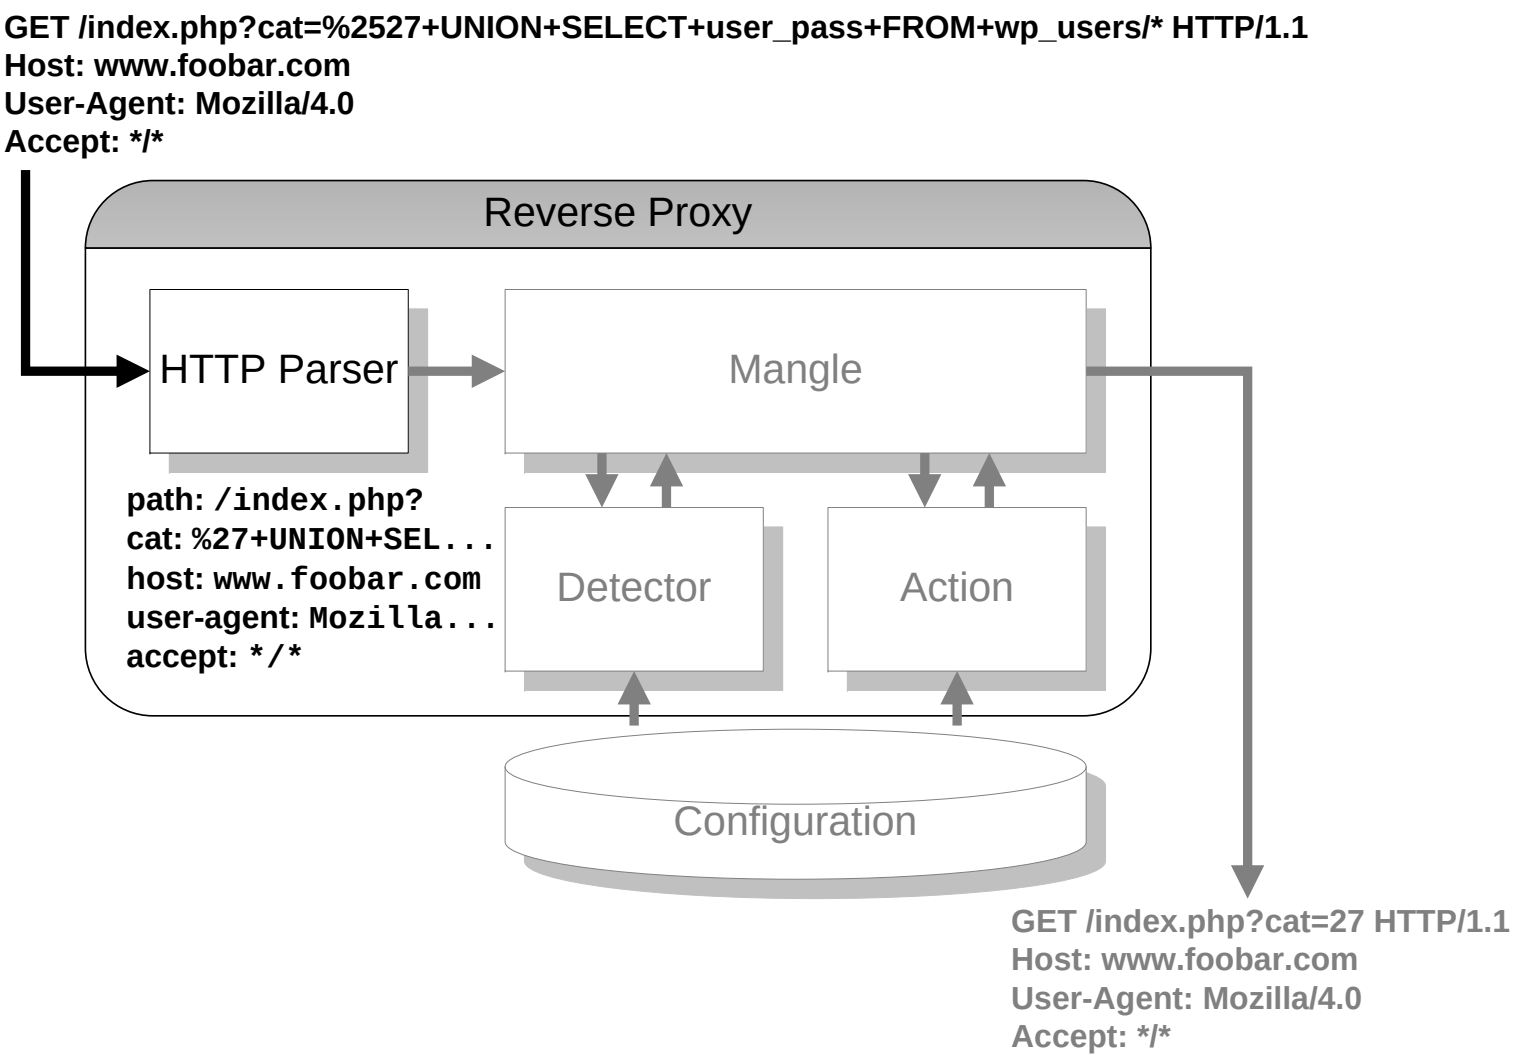

TokDoc – Walkthrough 3/4

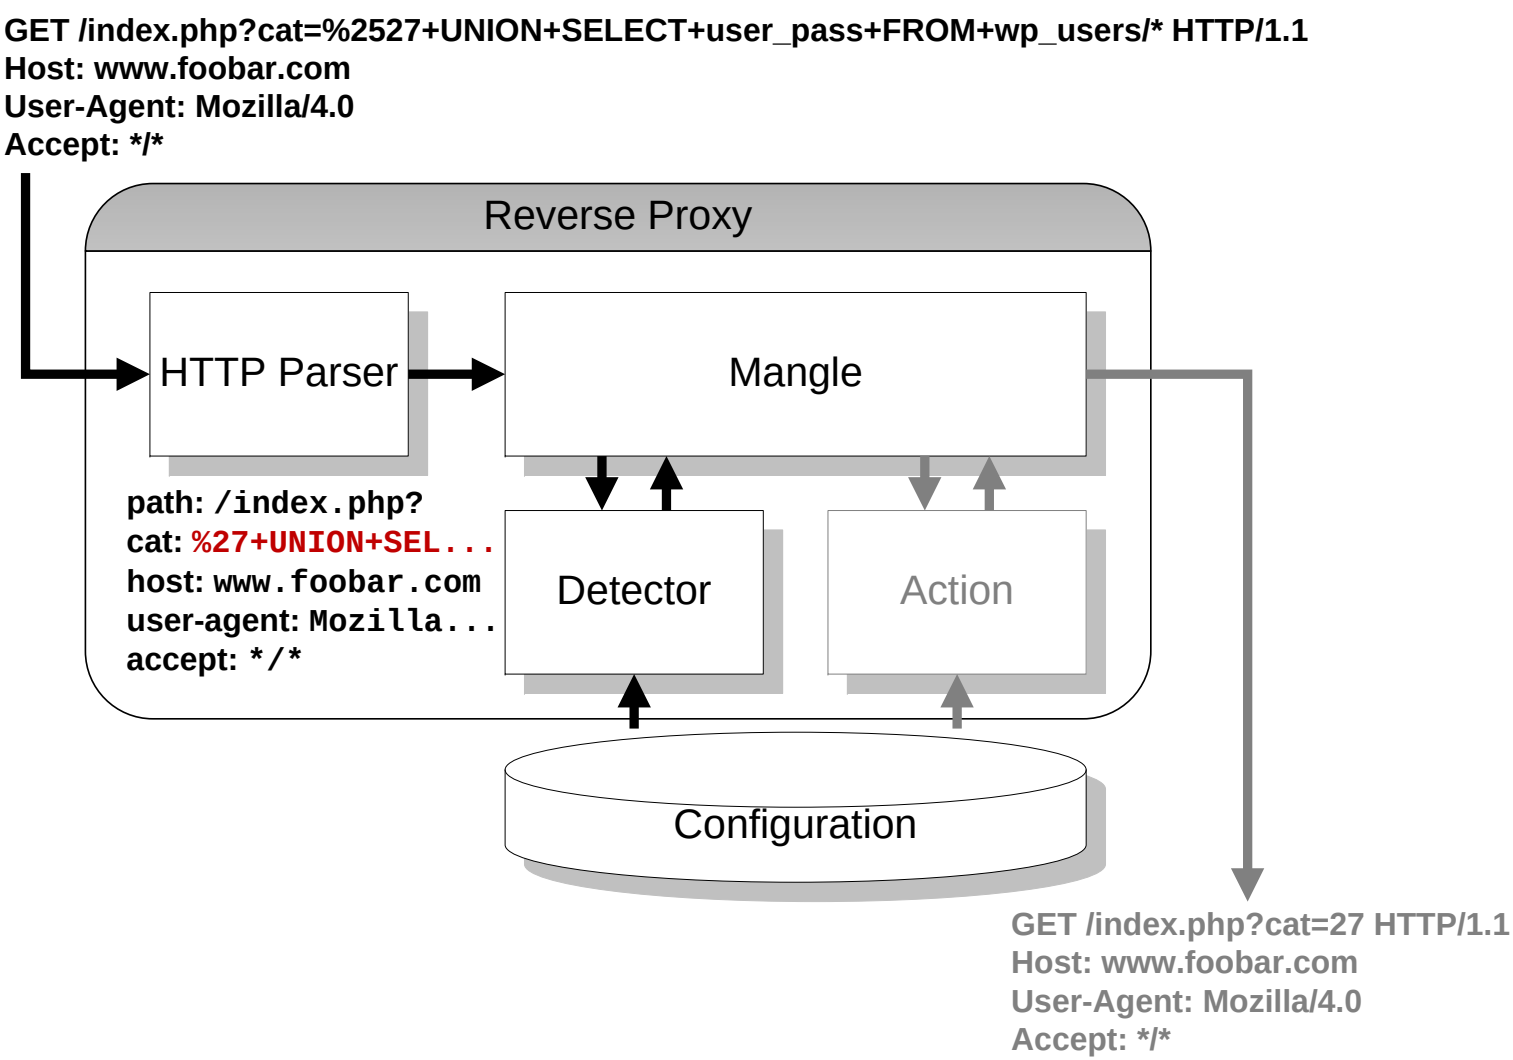

TokDoc – Walkthrough 4/4

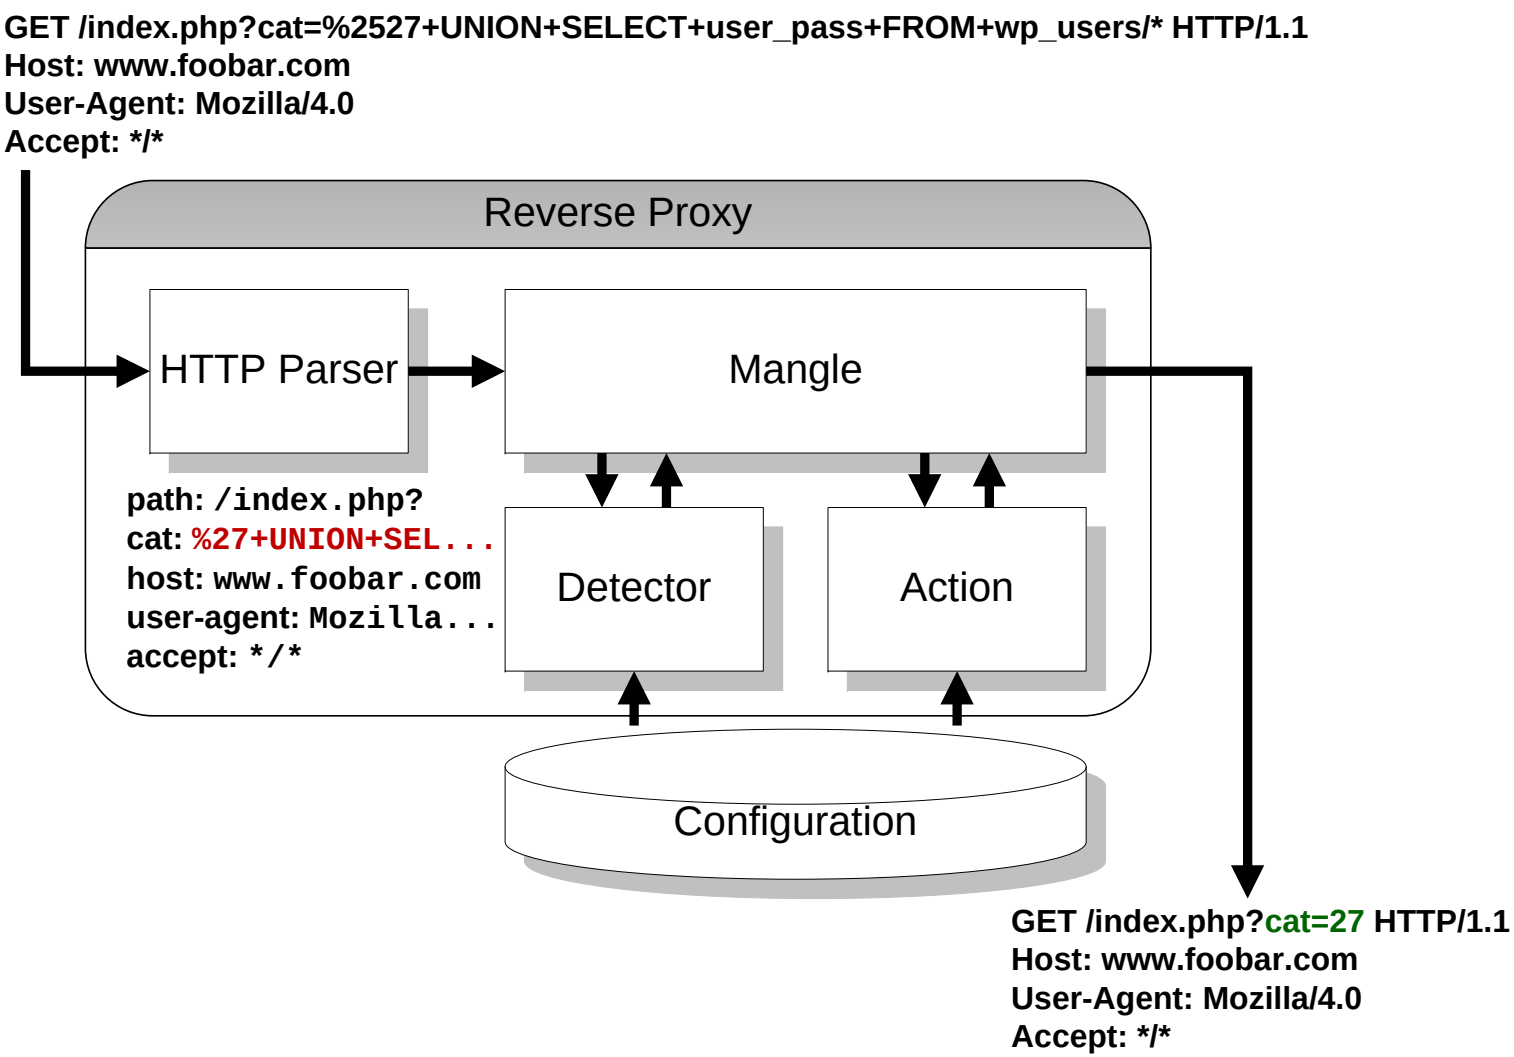

Outline

1 Introduction

2 TokDoc

- Walkthrough
- Anomaly Detectors
- Healing Actions
- Setup Process

3 Evaluation

- Ensemble of Learners
- Comparison to Other Detectors
- Runtime

4 Conclusion and Further Work

N-gram Centroid Anomaly Detector (NCAD)

Foundation n-gram vector space

Model Distance $d(\mu, x)$ of new x to mean μ of data

Decision Based on FP-tuned threshold t_a:

$$\text{score}_{\text{NCAD}}(x) = \begin{cases} \text{normal}, & \text{if } d(\mu, x) \leq t_a \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

PRO Good general model, capable of learning "normal" behavior

CON Fails for multimodal data distribution

TokDoc – Anomaly Detectors MCAD

Markov Chain Anomaly Detector (MCAD)

Foundation Markov Chain based on byte transitions (256 states with 256 transitions each)

Model Probablity $P(x \mid C)$ of new x in Markov Chain C

Decision Based on FP-tuned threshold p_a:

$$\text{score}_{\text{MCAD}}(x) = \begin{cases} \text{normal}, & \text{if } P(x \mid C) \geq p_a \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

PRO Can cope with multimodal data distribution

CON Not as tight as NCAD due to length dependence

TokDoc – Anomaly Detectors LAD

Length Anomaly Detector (LAD)

Foundation Robust estimation via bootstrap

Model Confidence interval around length quantile L_β of seen data

Decision Estimate variablity σ of length quantile and decide for new data point x:

$$\text{score}_{\text{LAD}}(x) = \begin{cases} \text{normal}, & \text{if } \text{len}(x) \leq L_\beta + c\sigma \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

PRO Works even for scarce data situations

CON No “deep” data inspection, just using length

TokDoc – Anomaly Detectors LIST

List Anomaly Detector (LIST)

Foundation Lists

Model Unique list L of seen values

Decision For new data point x:

$$\text{score}_{\text{LIST}}(x) = \begin{cases} \text{normal}, & \text{if } x \in L \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

PRO Fast and simple

CON We need to see all values beforehand (no "generalization" possible)

Outline

1 Introduction

2 TokDoc

- Walkthrough
- Anomaly Detectors
- Healing Actions
- Setup Process

3 Evaluation

- Ensemble of Learners
- Comparison to Other Detectors
- Runtime

4 Conclusion and Further Work

TokDoc – Healing Actions

1. **Dropping of tokens**:
   - remove anomalous token from request
   - default action for LAD detector.
2. **Preventive encoding**:
   - encode the anomalous value using HTML entities
   - manual assignment.
3. **Replacement with most frequent value**:
   - replace anomalous value with the most frequent normal value
   - Default action for LIST detector.
4. **Replacement with nearest value**.
   - replace anomalous value with its nearest-neighbor from the training set
   - Default action for both MCAD and NCAD

Outline

1 Introduction

2 TokDoc

- Walkthrough
- Anomaly Detectors
- Healing Actions
- Setup Process

3 Evaluation

- Ensemble of Learners
- Comparison to Other Detectors
- Runtime

4 Conclusion and Further Work

TokDoc – Setup Process

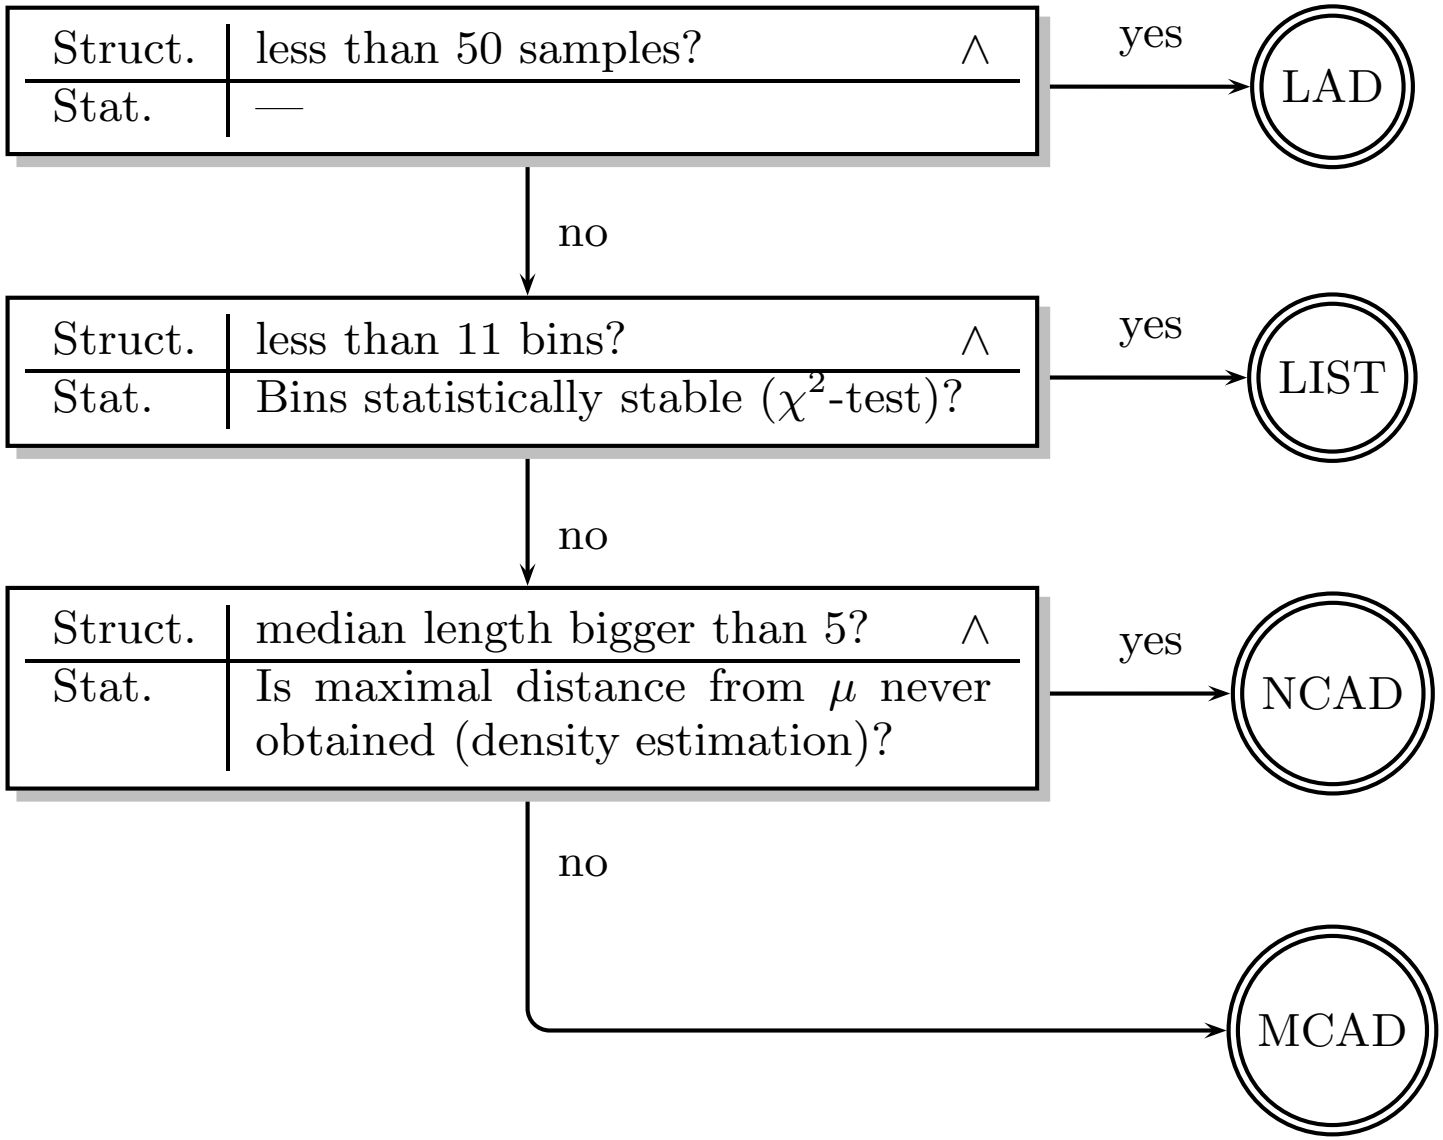

TokDoc – Setup Console

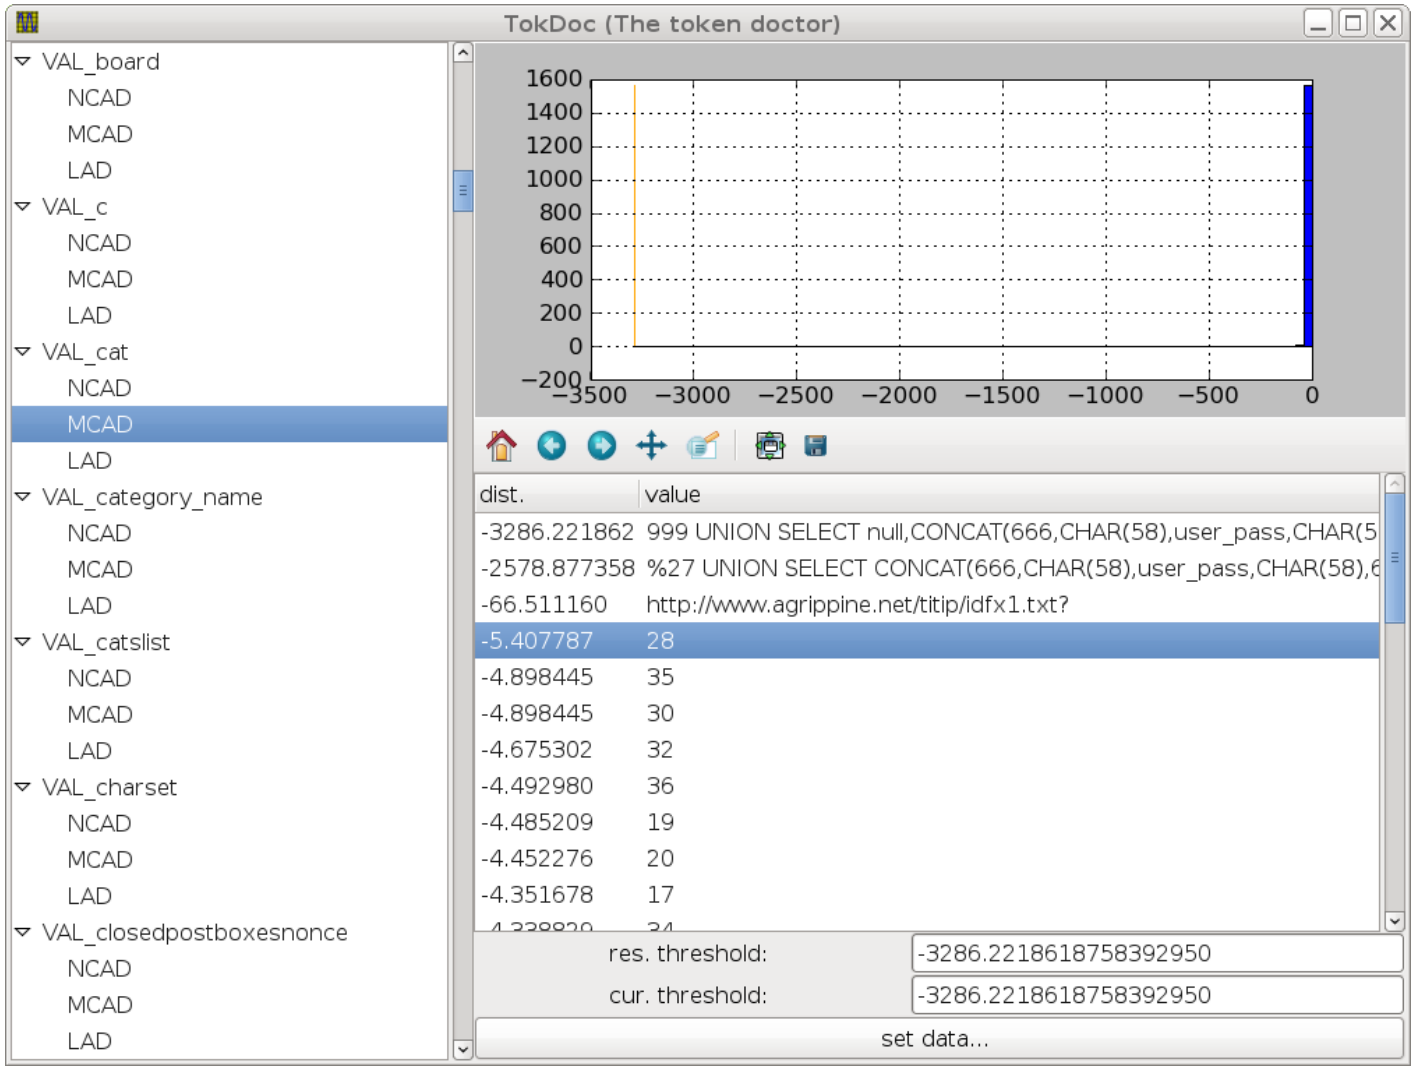

Outline

1 Introduction

2 TokDoc

- Walkthrough
- Anomaly Detectors
- Healing Actions
- Setup Process

3 Evaluation

- Ensemble of Learners
- Comparison to Other Detectors
- Runtime

4 Conclusion and Further Work

Evaluation – Ensemble of Learners

Question: Do we really need 4 different detectors?

Setup: Use just one type of detector and compare to TokDoc performance (FP = false-positive rate. TP = attacks found in normal traffic. FN = false-negative rate):

Dataset	Detector	FP	TP	FN
FIRST08	TokDoc	0.00002	0	0.00000
	TD_{LAD}	0.00000	0	0.02247
	TD_{MCAD}	0.00001	0	0.00000
	TD_{NCAD}	0.00002	0	0.22472
BLOG09	TokDoc	0.00003	212	0.04124
	TD_{LAD}	0.00001	68	0.15464
	TD_{MCAD}	0.00009	186	0.04124
	TD_{NCAD}	0.00003	0	0.22680

Answer: YES!

Outline

1 Introduction

2 TokDoc

- Walkthrough
- Anomaly Detectors
- Healing Actions
- Setup Process

3 Evaluation

- Ensemble of Learners
- Comparison to Other Detectors
- Runtime

4 Conclusion and Further Work

Evaluation – Comparison to Other Detectors

Question: Do we really need TokDoc at all?

Setup: Compare TokDoc performance to other anomaly-based detectors (FP_{TD} = false-positive rate of detector when calibrated to the true-positive rate of TokDoc. FN_{TD} = rate of missed regular attacks when detector is calibrated to the false-positive rate of TokDoc):

Dataset	**Detector**	**FP_{TD}**	**FN_{TD}**
FIRST08	TokDoc	0.00002	0.00000
	Markov Chain	0.02005	0.80899
	Anagram	0.00004	0.16854
BLOG09	TokDoc	0.00003	0.04124
	Markov Chain	0.16698	0.18557
	Anagram	1.00000	0.39175

Answer: YES!

Outline

Evaluation – Runtime

Question: Is TokDoc ready for deployment?
Setup: Measure median runtime in miliseconds per request and compare to other proxies:

	Proxy			
Dataset	Squid	ModSec.	twisted	TokDoc
FIRST08	1.387	1.536	2.552	2.768
BLOG09	1.500	1.694	2.430	2.902

Answer: YES!

Conclusion and Further Work

Protocol-aware reverse proxy TokDoc:

- fine-grained decisions at token level
- automatic setup procedure, which determines suitable model for each token
- intelligent mangling strategies for anomalous tokens

Future work:

- Integration of TokDoc into Squid or the ModSecurity platform
- Incorporation of a feedback loop
- Integration of session-awareness and “long term memory”
- Test in the wild...

Questions? Comments?
Thanks for your attention!

TokDoc – Anomaly Detectors NCAD

Given the set of all possible n-grams over byte sequences $S = \{0, \ldots, 255\}^n$, we define the embedding function ϕ for a token value x as follows:

$$\phi(x) = (\phi_s(x))_{s \in S} \in \mathbb{R}^{|S|} \quad \text{with} \quad \phi_s(x) = s \sqsubseteq x$$

$$\mu = \frac{1}{N} \sum_{i=1}^{N} \phi(x_i)$$

$$d(x, z) = \|\phi(x) - \phi(z)\|_2 = \sqrt{\sum_{s \in S} |\phi_s(x) - \phi_s(z)|^2}$$

$$\text{score}_{\text{NCAD}}(x) = \begin{cases} \text{normal}, & \text{if } d(\mu, x) \leq t_a \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

TokDoc – Anomaly Detectors MCAD

Having learned the transition probabilities, we can estimate the probability of a token value x of length n based on the learned Markov chain C:

$$P(x \mid C) = P(X_1 = x[1]) \prod_{i=1}^{n} P(X_{i+1} = x[i+1] \mid X_i = x[i])$$

$$\text{score}_{\text{MCAD}}(x) = \begin{cases} \text{normal}, & \text{if } P(x \mid C) \geq p_a \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

TokDoc – Anomaly Detectors LAD

Given a predefined significance level α_{LAD} we estimate the $1 - \alpha_{\text{LAD}}$ quantile of the length distribution of the train and validation data L, namely $\hat{L}_{1-\alpha_{\text{LAD}}}$. Now we construct a confidence interval for $L_{1-\alpha_{\text{LAD}}}$ by first calculating the bootstrap estimate of the standard error of $\hat{L}_{1-\alpha_{\text{LAD}}}$, namely $\hat{\sigma}$, and determining the parameter c, so that the following interval has probability coverage of $1 - \alpha_{\text{LAD}}$:

$$(\hat{L}_{1-\alpha_{\text{LAD}}} - c\hat{\sigma}, \hat{L}_{1-\alpha_{\text{LAD}}} + c\hat{\sigma})$$

$$\text{score}_{\text{LAD}}(x) = \begin{cases} \text{normal}, & \text{if } \text{len}(x) \leq \hat{L}_{1-\alpha_{\text{LAD}}} + c\hat{\sigma} \\ \text{anomaly}, & \text{otherwise.} \end{cases}$$

TokDoc – Setup Process

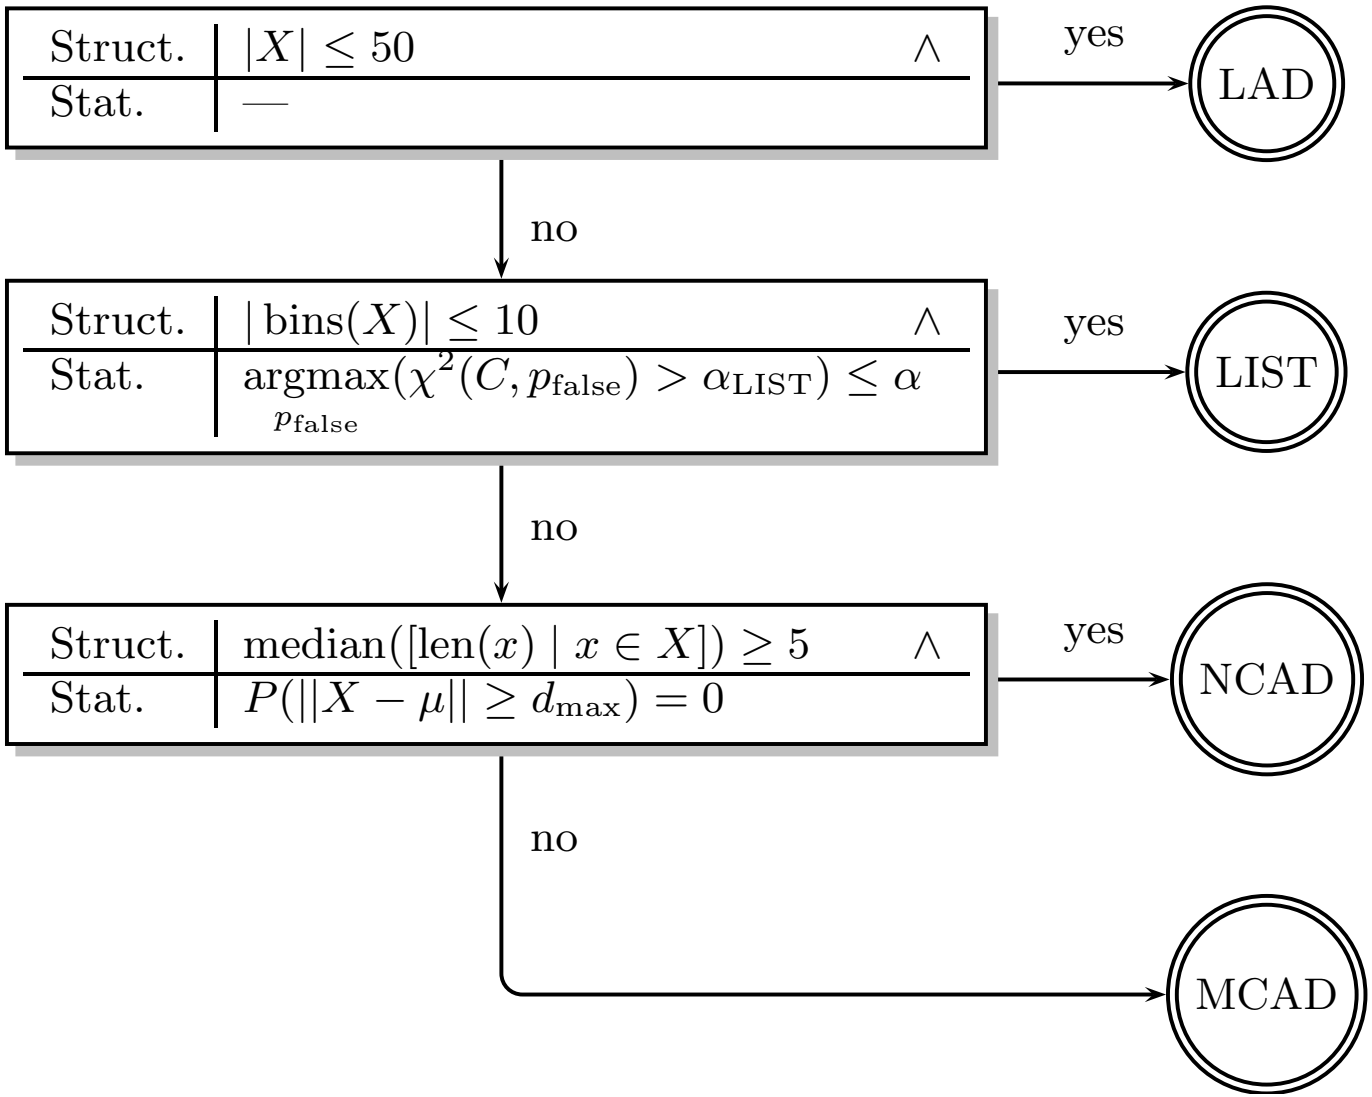

TokDoc – Setup Process

	Detectors **FIRST08**				
Category	LIST	LAD	MCAD	NCAD	Σ
Header	14	14	5	10	43
Parameter	9	3	4	—	16
Path	—	—	1	—	1
Σ	23	17	10	10	60

	Detectors **BLOG09**				
Category	LIST	LAD	MCAD	NCAD	Σ
Header	22	77	15	17	131
Parameter	14	166	28	7	215
Path	—	—	1	—	1
Σ	36	243	44	24	347