

Automated Detection and Fingerprinting of Censorship Block Pages

Ben Jones
Nick Feamster

Georgia Tech

Tzu Wen Lee
Phillipa Gill

Stony Brook University

Block Pages: a Common Form of Censorship



Access has been denied

The access to requested URL has been denied.

"www.cnn.com"

To have the rating of this web page re-evaluated please contact your [web moderator](#).


Toddlers' rape, killing shocks
South African community

Europe falls out of love with Obama


Block Page Detection is Important

- Censorship measurement is important
 - Censorship is pervasive
 - Transparency and accountability
- Block page detection is important
 - Common form of censorship
 - State of the art is manual analysis
 - We need automated, reliable, consistent methods




Detecting Block Pages



Challenges to Detecting Block Pages



Our Solution

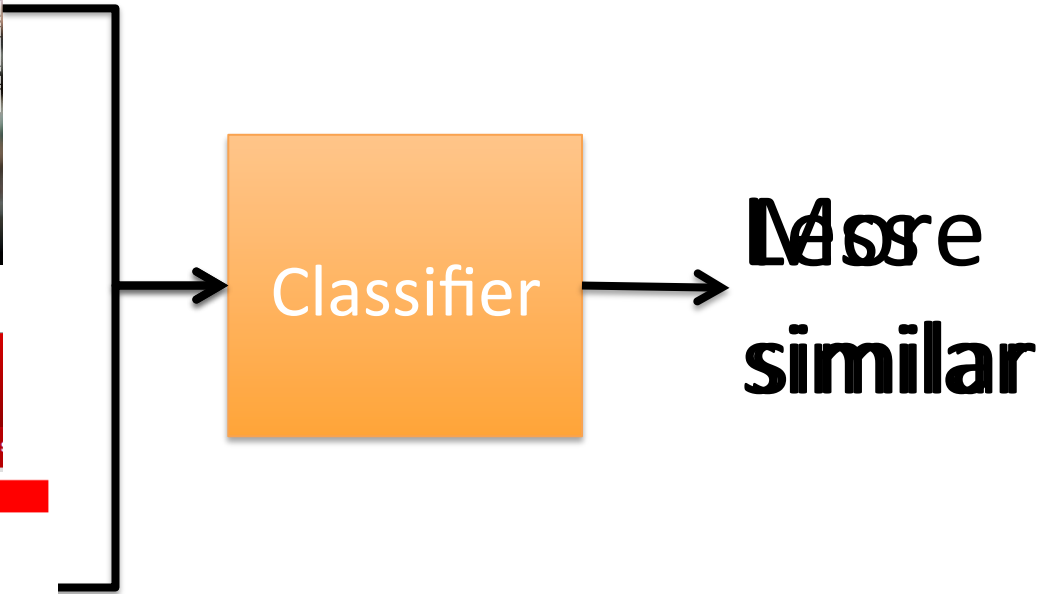


Access has been denied

The access to requested URL has been denied.

"www.cnn.com"

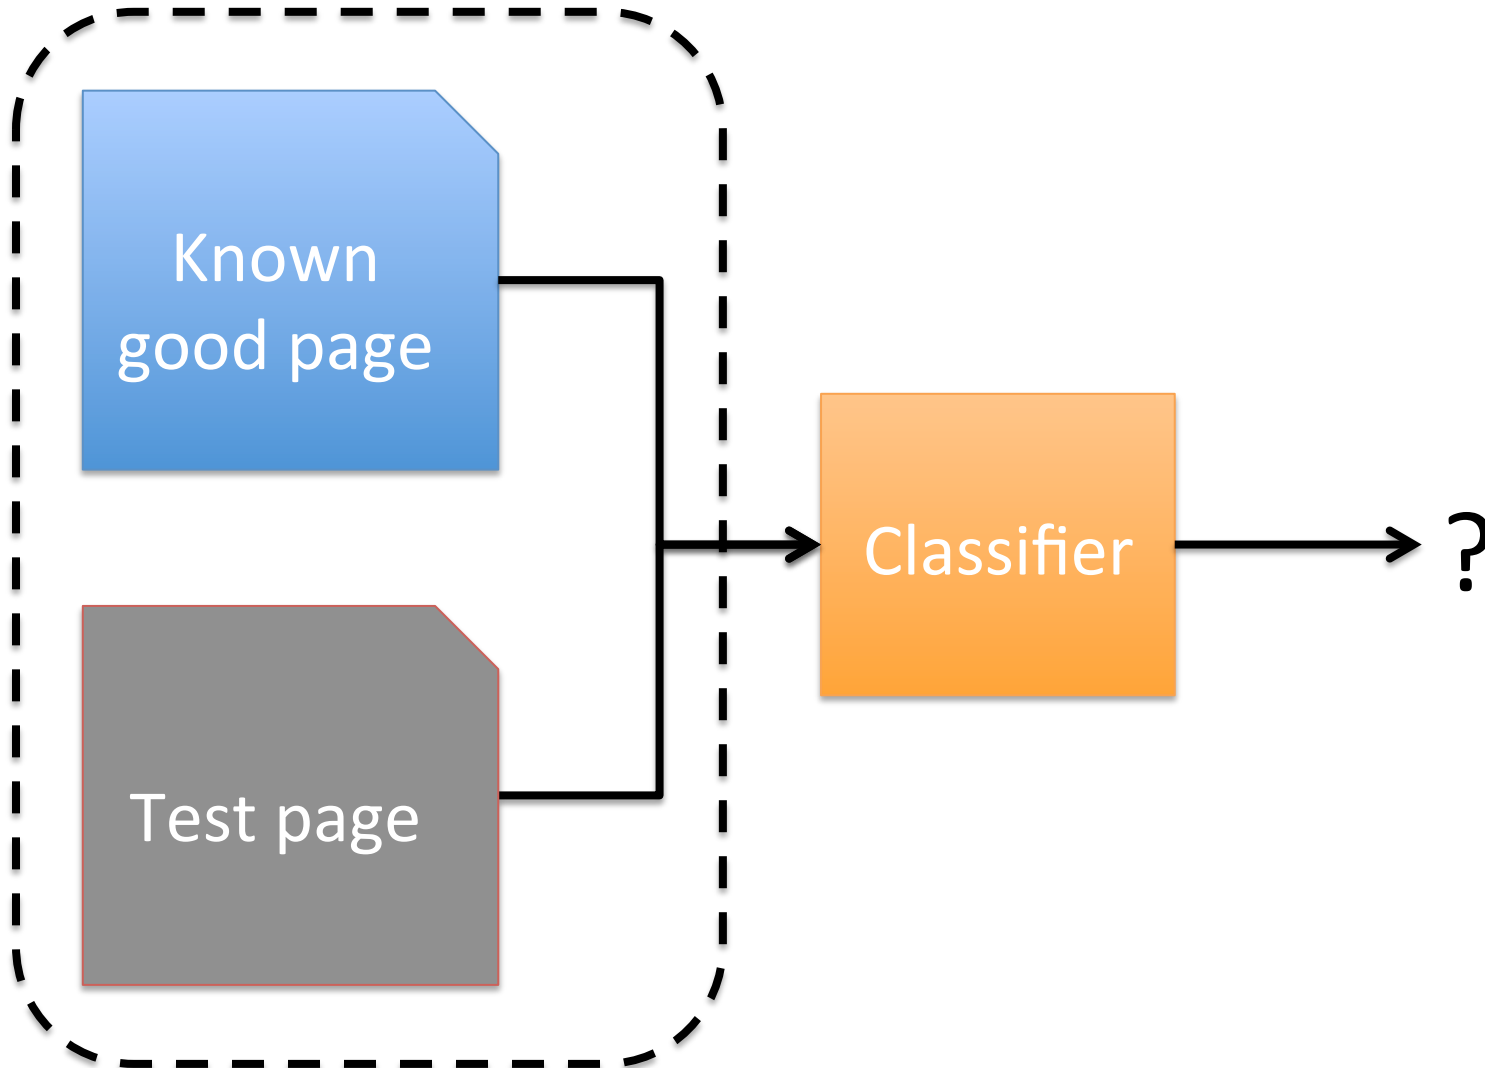
To have the rating of this web page re-evaluated please contact your [web moderator](#).



Problem Statements

- We need reliable, consistent, automated ways:
 - To detect censorship block pages
 - To detect the appearance of new censorship tools
- These techniques are necessary because:
 - The state of the art is not automated or consistent
 - Censorship measurement is important


Our Data



Our Data

- Gathered by the OpenNet Initiative (ONI)
 - Collected in 49 countries from 2007 to 2012
- Has pairs of test pages and known good pages
 - Each pair is labeled as censored or not
 - ~480,000 pairs where the test page is accessible
 - ~28,000 pairs where the test page is a block page


Block Page Detection



Detecting Block Pages

- Feature extraction
- Evaluate features with precision-recall curve
- Best feature: Page length ratio
 - Size of smaller page / size of larger page

Detecting Block Pages



Precision-Recall curve for varying thresholds

- Page length ratio was best classifier
- Best classifier achieved 80% precision, 95% true positive rate, 1.37% false positive

Finding New Censorship Tools

مشترک گرامی

دسترسی به این سایت امکان پذیر نمیباشد ^a

2. Different censorship tools use different templates

Access has been denied

The access to requested URL has been denied.

- Solution

"www.cnn.com"

- Cluster b

To have the rating of this web page re-evaluated please contact your [web moderator](#).

- Analyze cluster exemplars to determine what tool generated the template

Fingerprinting Evaluation

- Experimental Validation
 - Use F-1 measure to compare clusters against manually generated regular expressions
 - Best clustering F-1 = 0.98
 - Only 36 unique block page HTML structures
- Practical validation with retroactive analysis
 - Found a new filtering tool in Saudi Arabia: WireFilter

Real World Impact



- Real world deployment
 - ICLab collaboration between Citizen Lab, Stony Brook, and Georgia Tech
 - Deploying with censorship measurement platform

Conclusion

- We can accurately detect block pages and fingerprint the tools that create them
- Download our code from <https://github.com/iclab/blockpage-detection>
- Email me: bjones99@gatech.edu