

A User-Guided Approach to Program Analysis

Xin Zhang

Georgia Tech & University of Pennsylvania

Joint work with:

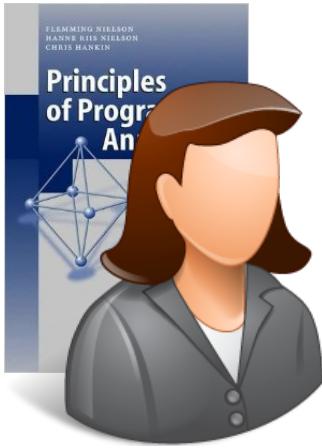
Ravi Mangal, Aditya Nori, and Mayur Naik

Motivation



Program
Analysis

Motivation

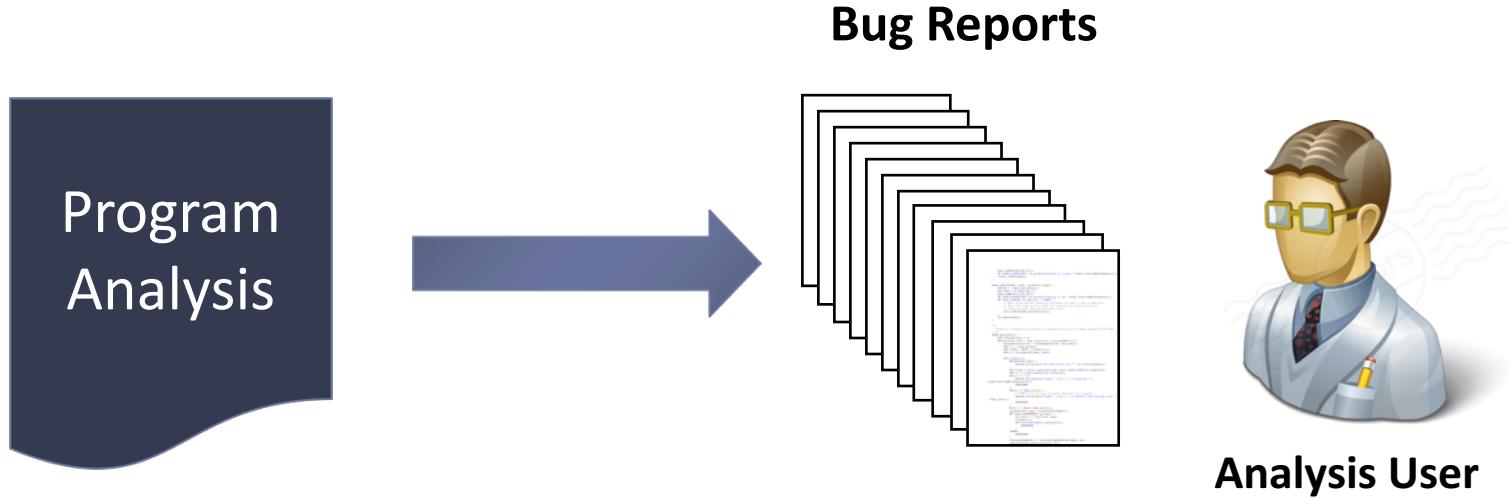


Analysis Writer



Computing exact solutions impossible
Imprecisely defined properties
Missing program parts
...

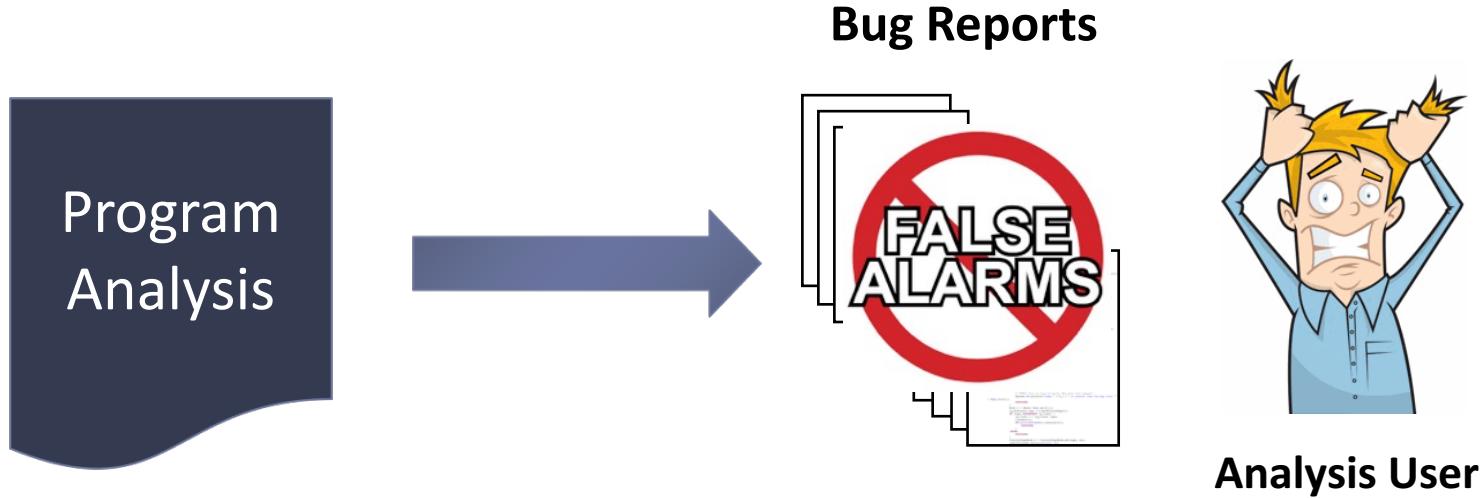
Motivation



Motivation



Motivation



“People ignore the tool if more than
30% false positives are reported ...”

[Coverity, CACM’10]

Our Key Idea

Shift decisions about **usefulness** of results
from **analysis writers** to **analysis users**



Example: Static Datarace Detection

Code snippet from Apache FTP Server

```
1 public class RequestHandler {  
2     FtpRequestImpl request;  
3     FtpWriter writer;  
4     BufferedReader reader;  
5     Socket controlSocket;  
6     boolean isConnectionClosed;  
7     ...  
8     public void getRequest() {  
9         return request; // x0  
10    }
```

```
11    public void close() {  
12        synchronized (this) {  
13            if (isConnectionClosed)  
14                return;  
15            isConnectionClosed = true;  
16        }  
17        request.clear(); // x1  
18        request = null; // x2  
19        writer.close(); // y1  
20        writer = null; // y2  
21        reader.close();  
22        reader = null;  
23        controlSocket.close();  
24        controlSocket = null;  
25    }
```

R1

R2

R3

R4

R5

Before User Feedback

Detected Races

R1: Race on field `org.apache.ftpserver.RequestHandler.m_request`



`org.apache.ftpserver.RequestHandler: 9`

`org.apache.ftpserver.RequestHandler: 18`

R2: Race on field `org.apache.ftpserver.RequestHandler.m_request`



`org.apache.ftpserver.RequestHandler: 17`

`org.apache.ftpserver.RequestHandler: 18`

R3: Race on field `org.apache.ftpserver.RequestHandler.m_writer`



`org.apache.ftpserver.RequestHandler: 19`

`org.apache.ftpserver.RequestHandler: 20`

R4: Race on field `org.apache.ftpserver.RequestHandler.m_reader`



`org.apache.ftpserver.RequestHandler: 21`

`org.apache.ftpserver.RequestHandler: 22`

R5: Race on field `org.apache.ftpserver.RequestHandler.m_controlSocket`



`org.apache.ftpserver.RequestHandler: 23`

`org.apache.ftpserver.RequestHandler: 24`

Eliminated Races

E1: Race on field `org.apache.ftpserver.RequestHandler. m_isConnectionClosed`

`org.apache.ftpserver.RequestHandler: 13`

`org.apache.ftpserver.RequestHandler: 15`



Before User Feedback

Detected Races

R1: Race on field org.apache.ftpserver.RequestHandler.m_request



org.apache.ftpserver.RequestHandler: 9

org.apache.ftpserver.RequestHandler: 18

R2: Race on field org.apache.ftpserver.RequestHandler.m_request



org.apache.ftpserver.RequestHandler: 17

org.apache.ftpserver.RequestHandler: 18

R3: Race on field org.apache.ftpserver.RequestHandler.m_writer



org.apache.ftpserver.RequestHandler: 19

org.apache.ftpserver.RequestHandler: 20

R4: Race on field org.apache.ftpserver.RequestHandler.m_reader



org.apache.ftpserver.RequestHandler: 21

org.apache.ftpserver.RequestHandler: 22

R5: Race on field org.apache.ftpserver.RequestHandler.m_controlSocket



org.apache.ftpserver.RequestHandler: 23

org.apache.ftpserver.RequestHandler: 24

Eliminated Races

E1: Race on field org.apache.ftpserver.RequestHandler. m_isConnectionClosed

org.apache.ftpserver.RequestHandler: 13

org.apache.ftpserver.RequestHandler: 15

After User Feedback

Detected Races

R1: Race on field `org.apache.ftpserver.RequestHandler.m_request`



`org.apache.ftpserver.RequestHandler: 9`

`org.apache.ftpserver.RequestHandler: 18`

Eliminated Races

E1: Race on field `org.apache.ftpserver.RequestHandler. m_isConnectionClosed`

`org.apache.ftpserver.RequestHandler: 13`

`org.apache.ftpserver.RequestHandler: 15`

E2: Race on field `org.apache.ftpserver.RequestHandler.m_request`

`org.apache.ftpserver.RequestHandler: 17`

`org.apache.ftpserver.RequestHandler: 18`

E3: Race on field `org.apache.ftpserver.RequestHandler.m_writer`

`org.apache.ftpserver.RequestHandler: 19`

`org.apache.ftpserver.RequestHandler: 20`

E4: Race on field `org.apache.ftpserver.RequestHandler.m_reader`

`org.apache.ftpserver.RequestHandler: 21`

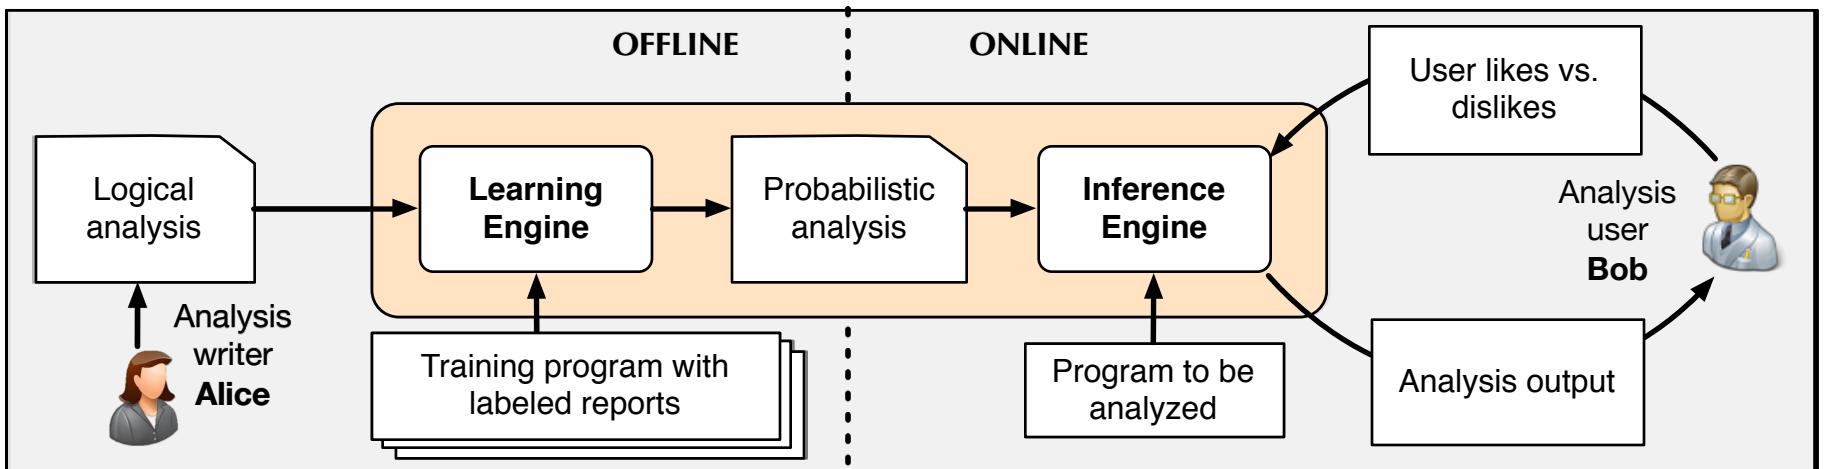
`org.apache.ftpserver.RequestHandler: 22`

E5: Race on field `org.apache.ftpserver.RequestHandler.m_controlSocket`

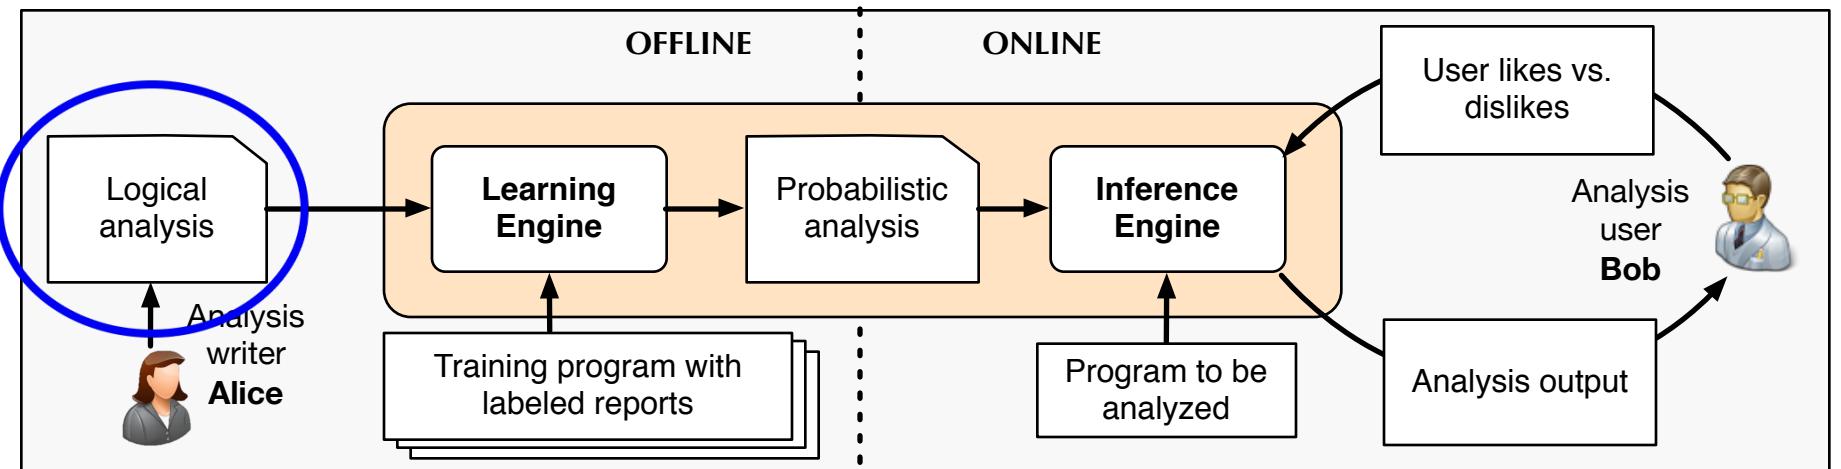
`org.apache.ftpserver.RequestHandler: 23`

`org.apache.ftpserver.RequestHandler: 24`

Our System For User-Guided Analysis



Logical Analysis



Logical Datarace Analysis Using Datalog

Input relations:

next(p1, p2), mayAlias(p1, p2), guarded(p1, p2)

Output relations:

parallel(p1, p2), race(p1, p2)

Rules:

parallel(p3, p2) :- parallel(p1, p2), next (p3, p1).

parallel(p1, p2) :- parallel(p2, p1).

race(p1, p2) :- parallel(p1, p2), mayAlias(p1, p2), \neg guarded(p1, p2).

Logical Datarace Analysis Using Datalog

Input relations:

next(p1, p2), mayAlias(p1, p2), guarded(p1, p2)

Output relations:

p1 is immediate
successor of p2.

p1 & p2 may
access the same
memory location.

p1 & p2 are
guarded by the
same lock.

Rules:

parallel(p3, p2) :- parallel(p1, p2), next (p3, p1).

parallel(p1, p2) :- parallel(p2, p1).

race(p1, p2) :- parallel(p1, p2), mayAlias(p1, p2), \neg guarded(p1, p2).

Logical Datarace Analysis Using Datalog

Input relations:

$\text{next}(p1, p2)$, $\text{mayAlias}(p1, p2)$, $\text{guarded}(p1, p2)$

Output relations:

$\text{parallel}(p1, p2)$, $\text{race}(p1, p2)$

Rules:

parallel($p1 \wedge p2$ may happen in parallel.)
parallel($p1 \wedge p2$ may have a datarace.)

$\text{race}(p1, p2) :- \text{parallel}(p1, p2), \text{mayAlias}(p1, p2), \neg \text{guarded}(p1, p2).$

Logical Datarace Analysis Using Datalog

Input relations:

next(p1, p2), mayAlias(p1, p2), guarded(p1, p2)

Output relations:

parallel(p1, p2), race(p1, p2)

If p1 & p2 may happen in parallel,
and they may access the same memory location,
and they are not guarded by the same lock,
then p1 & p2 may have a datarace.

Rules:

parallel(p3, p2) :-

parallel(p1, p2) :- parallel(p1, p3), parallel(p3, p2).

race(p1, p2) :- parallel(p1, p2), mayAlias(p1, p2), \neg guarded(p1, p2).

Why Datalog?

- ▶ Easier to specify

Analysis in Java

```
bdd1.andWith(bdd1.id());
if (TFAKE_RELATION(m)) out.println("Adding to formal: "+bdd1.toStringWithDomains());
formal.orWith(bdd1);
}

void addToIE(BDD bdd1, iq_Method target) {
    BDD bdd1 = M.lthVar(MC_1);
    bdd1.andWith(bdd1.id());
    if (TFAKE_RELATION(m)) out.println("Adding to IE: "+bdd1.toStringWithDomains());
    if (TFAKE_CONTEXT && IEfilter != null) {
        // When doing context-sensitive analysis, we need to add to IEs too.
        // This is done by adding all contexts for this invocation.
        // "and"-ing with IEfilter achieves this.
        IEor.orWith(bdd1, and(IEfilter));
    }
    IE.orWith(bdd1);
}

/** Finds all invocation sites with no targets and tries to create targets from them.
 * @param analyzeIE
 */
int noTargetsCalls = 0;
for (ProgramLocationIterator iter = ProgramLocation.mc.iterator(); iter.hasNext();) {
    ProgramLocation mc = (ProgramLocation) iter.next();
    int i_z = Imp.mc();
    BDD _bdd = BDD.bdd();
    BDD z_bdd = BDD.bdd();
    BDD r_bdd = BDD.bdd();
    Imp.mc();
    if (iter.isLast()) {
        System.out.println("No destination for " + mc.toStringLong());
    }
    BDD V_bdd = actual.relprod(r_bdd, z_bdd).restrictWith(z.lthVar(0));
    int V_z = V_bdd.scanVar(V2_l, intValue());
    if (V_z == -1) {
        System.out.println("Index " + V_z + " is scanning " +
            V_bdd.toStringWithDomains());
        continue;
    }
    if (V_z > Vmap.size()) {
        System.out.println("Index " + V_z + " is greater than the map size: " +
            Vmap.size());
        continue;
    }
    Node n = (Node) Vmap.get(V_z);
    iq_Type type = n.getType();
    if (type instanceof iq_Class) {
        iq_Class c = (iq_Class) type;
        c.prepare();
        if (!c.withInitFor.contains(o)) {
            continue;
        }
    } else {
        continue;
    }
    ConcreteTypeNode b = ConcreteTypeNode.get(type, mc);
    addToIE(b.replace(V2of1, h1));
}
```

VS.

Analysis in Datalog

```
### Context-sensitive inclusion-based pointer analysis using cloning
#
# Calculates the numbering based on the call graph relation.
#
# Author: John Whaley

include "fielddomains.pa"
bddnodes 10000000
bddcache 1000000
bddvarorder NO_F0_I0_M1_M0_V1_V0_VC1VC0_T0_Z0_T1_H0_H1

m10(m,i) :- ml(m,i,_).
IEnum(i,m,vc2,vc1) :- roots(m), m10(m,i), IE0(i,m). number

cvP(_ ,v,h) :- vp0(v,h).
ca(_ ,v1,_ ,v2) :- A(v1,v2).
IEcs(vc2,i,vc1,m) :- IE0(i,m), IEnum(i,m,vc2,vc1).
vpfilter(v,h) :- vt(v,tv), at(tv,th), ht(h,th).
ca(vc1,v1,vc2,v2) :- formal(m,z,v1), IEcs(vc2,i,vc1,m), actual(i,z,v2).
ca(vc2,v2,vc1,v1) :- Mret(m,v1), IEcs(vc2,i,vc1,m), Iret(i,v2).
ca(vc2,v2,vc1,v1) :- Mthr(m,v1), IEcs(vc2,i,vc1,m), Ithr(i,v2).

cvP(vc1,v1,h) :- ca(vc1,v1,vc2,v2), cvP(vc2,v2,h), vpfilter(v1,h).
hP(h1,f,h2) :- S(v1,f,v2), cvP(vc1,v1,h1), cvP(vc1,v2,h2).
cvP(vc1,v2,h2) :- L(v1,f,v2), cvP(vc1,v1,h1), hP(h1,f,h2), vpfilter(v2,h2). split

IE(i,m) :- IEcs(_,i,_,m).
```

Why Datalog?

- ▶ Easier to specify
- ▶ Leverage efficient solvers

Paddle

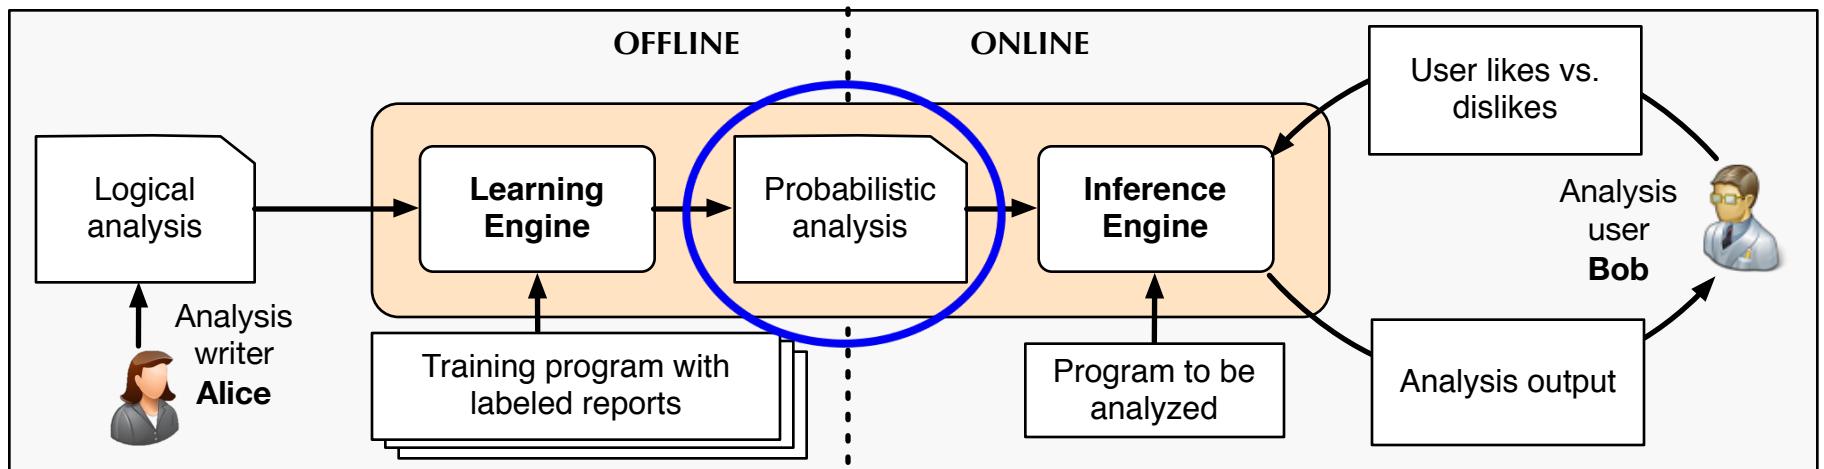


- ▶ Widely adaptable

Soot



Probabilistic Analysis



Datarace Analysis: Logical \rightarrow Probabilistic

Input relations:

$\text{next}(p1, p2)$, $\text{mayAlias}(p1, p2)$, $\text{guarded}(p1, p2)$

Output relations:

$\text{parallel}(p1, p2)$, $\text{race}(p1, p2)$

Rules:

$\text{parallel}(p3, p2) :- \text{parallel}(p1, p2), \text{next}(p3, p1).$ **weight 5**

$\text{parallel}(p1, p2) :- \text{parallel}(p2, p1).$

$\text{race}(p1, p2) :- \text{parallel}(p1, p2), \text{mayAlias}(p1, p2), \neg\text{guarded}(p1, p2).$

$\neg\text{race}(x2, x1).$ **weight 25**

Datarace Analysis: Logical \rightarrow Probabilistic

Input relations:

$\text{next}(p1, p2)$, $\text{mayAlias}(p1, p2)$, $\text{guarded}(p1, p2)$

Output relations:

$\text{parallel}(p1, p2)$, $\text{race}(p1, p2)$

Rules:

$\text{parallel}(p3, p2) :- \text{next}(p3, p1).$ “Hard” Rule weight 5

$\text{parallel}(p1, p2) :- \text{parallel}(p2, p1).$

$\text{race}(p1, p2) :- \text{parallel}(p1, p2), \text{mayAlias}(p1, p2), \neg\text{guarded}(p1, p2).$

$\neg\text{race}(x2, x1).$ weight 25

A Semantics for Probabilistic Analysis

- ▶ Probabilistic Analysis => Markov Logic Network (MLN)
[Richardson & Domingos, Machine Learning'06]
- ▶ MLN defines a probability distribution over all possible analysis outputs
- ▶ Probability of an output x :

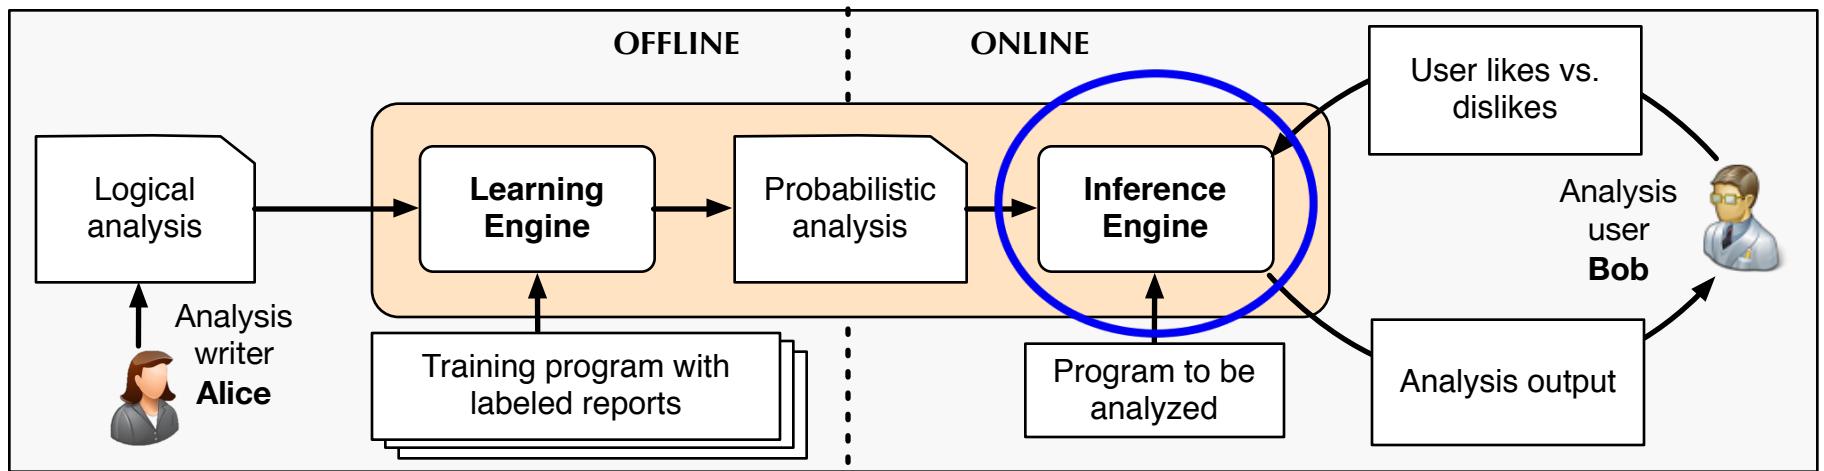
$$P(X=x) = \frac{1}{Z} \exp \left(\sum_i w_i n_i(x) \right)$$

Normalization factor

Weight of rule i

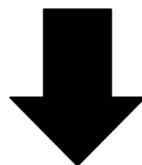
Number of true instances of rule i in x

Inference Engine



Probabilistic Inference

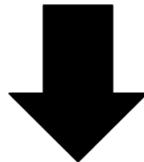
Find the most likely output given the input program



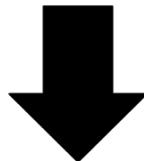
$$\begin{aligned}\arg \max_x P(x) &= \arg \max_x \frac{1}{Z} \exp \left(\sum_i w_i n_i(x) \right) \\ &= \arg \max_x \sum_i w_i n_i(x)\end{aligned}$$

What is MaxSAT?

```
¬ b1 ∨ ¬ b2 ∨ b3    weight 5   ∧  
b3 ∨     b4          weight 10  ∧  
¬ b4 ∨ ¬ b2          weight 7   ∧  
...
```



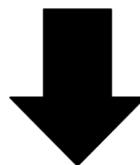
Find a boolean assignment such that the sum of the weights of the satisfied clauses is maximized



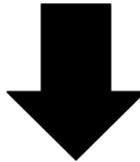
$$\arg \max_x \sum_i w_i n_i(x)$$

Probabilistic Inference → MaxSAT

Find the most likely output given the input program



$$\begin{aligned}\arg \max_x P(x) &= \arg \max_x \frac{1}{Z} \exp \left(\sum_i w_i n_i(x) \right) \\ &= \arg \max_x \sum_i w_i n_i(x)\end{aligned}$$



Solve the MaxSAT instance entailed by the MLN

Example: Static Datarace Detection

Code snippet from Apache FTP Server

```
1 public class RequestHandler {  
2     FtpRequestImpl request;  
3     FtpWriter writer;  
4     BufferedReader reader;  
5     Socket controlSocket;  
6     boolean isConnectionClosed;  
7     ...  
8     public void getRequest() {  
9         return request; // x0  
10    }
```

```
11    public void close() {  
12        synchronized (this) {  
13            if (isConnectionClosed)  
14                return;  
15            isConnectionClosed = true;  
16        }  
17        request.clear(); // x1  
18        request = null; // x2  
19        writer.close(); // y1  
20        writer = null; // y2  
21        reader.close();  
22        reader = null;  
23        controlSocket.close();  
24        controlSocket = null;  
25    }
```

R1

R2

R3

R4

R5

How Does Online Phase Work?

Input facts:

mayAlias(x2, x1), \neg guarded(x2, x1), next(x2, x1), parallel(x2, x0), race(x2, x0),
mayAlias(y2, y1), \neg guarded(y2, y1), next(y1, x2), parallel(x2, x1), **race(x2, x1)**,
next(y2, x1) parallel(y2, y1), **race(y2, y1)**

MaxSAT formula:

(parallel(x1, x1) \wedge next(x2, x1) \Rightarrow parallel(x2, x1)) **weight 5** \wedge
(parallel(x2, x1) \Rightarrow parallel(x1, x2)) \wedge
(parallel(x1, x2) \wedge next(x2, x1) \Rightarrow parallel(x2, x2)) **weight 5** \wedge
(parallel(x2, x2) \wedge next(y1, x2) \Rightarrow parallel(y1, x2)) **weight 5** \wedge
(parallel(y1, x2) \Rightarrow parallel(x2, y1)) \wedge
(parallel(x2, y1) \wedge next(y1, x2) \Rightarrow parallel(y1, y1)) **weight 5** \wedge
(parallel(y1, y1) \wedge next(y2, y1) \Rightarrow parallel(y2, y1)) **weight 5** \wedge
(parallel(y2, y1) \wedge mayAlias(y2, y1) \wedge \neg guarded(y2, y1) \Rightarrow race(y2, y1)) \wedge
(parallel(x2, x1) \wedge mayAlias(x2, x1) \wedge \neg guarded(x2, x1) \Rightarrow race(x2, x1)) \wedge
 \neg race(x2, x1) weight 25

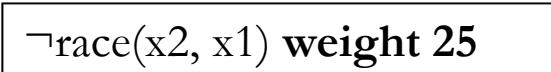
Output facts:

How Does Online Phase Work?

Input facts:

mayAlias(x2, x1), \neg guarded(x2, x1), next(x2, x1), parallel(x2, x0), race(x2, x0),
mayAlias(y2, y1), \neg guarded(y2, y1), next(y1, x2), parallel(x2, x1), **race(x2, x1)**,
next(y2, x1) parallel(y2, y1), **race(y2, y1)**

MaxSAT formula:

(parallel(x1, x1) \wedge next(x2, x1) \Rightarrow **parallel(x2, x1)**) **weight 5** \wedge 
(parallel(x2, x1) \Rightarrow parallel(x1, x2)) \wedge
(parallel(x1, x2) \wedge next(x2, x1) \Rightarrow parallel(x2, x2)) **weight 5** \wedge
(parallel(x2, x2) \wedge next(y1, x2) \Rightarrow parallel(y1, x2)) **weight 5** \wedge
(parallel(y1, x2) \Rightarrow parallel(x2, y1)) \wedge
(parallel(x2, y1) \wedge next(y1, x2) \Rightarrow parallel(y1, y1)) **weight 5** \wedge
(parallel(y1, y1) \wedge next(y2, y1) \Rightarrow parallel(y2, y1)) **weight 5** \wedge
(parallel(y2, y1) \wedge mayAlias(y2, y1) \wedge \neg guarded(y2, y1) \Rightarrow race(y2, y1)) \wedge
(parallel(x2, x1) \wedge mayAlias(x2, x1) \wedge \neg guarded(x2, x1) \Rightarrow race(x2, x1)) \wedge
 \neg race(x2, x1) **weight 25**

Output facts:

How Does Online Phase Work?

Input facts:

mayAlias(x2, x1), \neg guarded(x2, x1), next(x2, x1),

mayAlias(y2, y1), \neg guarded(y2, y1), next(y1, x2),

Output facts:

parallel(x2, x0), race(x2, x0),

~~parallel(x2, x1)~~, ~~race(x2, x1)~~,

next(y2, x1) parallel(y2, y1), race(y2, y1)

MaxSAT formula:

(parallel(x1, x1) \wedge next(x2, x1) \Rightarrow ~~parallel(x2, x1)~~) **weight 5** \wedge 

(parallel(x2, x1) \Rightarrow parallel(x1, x2)) \wedge

(parallel(x1, x2) \wedge next(x2, x1) \Rightarrow parallel(x2, x2)) **weight 5** \wedge

(parallel(x2, x2) \wedge next(y1, x2) \Rightarrow parallel(y1, x2)) **weight 5** \wedge

(parallel(y1, x2) \Rightarrow parallel(x2, y1)) \wedge

(parallel(x2, y1) \wedge next(y1, x2) \Rightarrow parallel(y1, y1)) **weight 5** \wedge

(parallel(y1, y1) \wedge next(y2, y1) \Rightarrow parallel(y2, y1)) **weight 5** \wedge

(parallel(y2, y1) \wedge mayAlias(y2, y1) \wedge \neg guarded(y2, y1) \Rightarrow race(y2, y1)) \wedge

~~parallel(x2, x1) \wedge mayAlias(x2, x1) \wedge \neg guarded(x2, x1) \Rightarrow race(x2, x1)~~) \wedge

race(x2, x1) weight 25

How Does Online Phase Work?

Input facts:

mayAlias(x2, x1), \neg guarded(x2, x1), next(x2, x1),

mayAlias(y2, y1), \neg guarded(y2, y1), next(y1, x2),

Output facts:

parallel(x2, x0), race(x2, x0),

~~parallel(x2, x1)~~, ~~race(x2, x1)~~,

~~parallel(y2, y1)~~, ~~race(y2, y1)~~

MaxSAT formula:

(parallel(x1, x1) \wedge next(x2, x1) \Rightarrow ~~parallel(x2, x1)~~) weight 5 \wedge 

~~(parallel(x2, x1) \Rightarrow parallel(x1, x2))~~ \wedge

~~(parallel(x1, x2) \wedge next(x2, x1) \Rightarrow parallel(x2, x2))~~ weight 5 \wedge

~~(parallel(x2, x2) \wedge next(y1, x2) \Rightarrow parallel(y1, x2))~~ weight 5 \wedge

~~(parallel(y1, x2) \Rightarrow parallel(x2, y1))~~ \wedge

~~(parallel(x2, y1) \wedge next(y1, x2) \Rightarrow parallel(y1, y1))~~ weight 5 \wedge

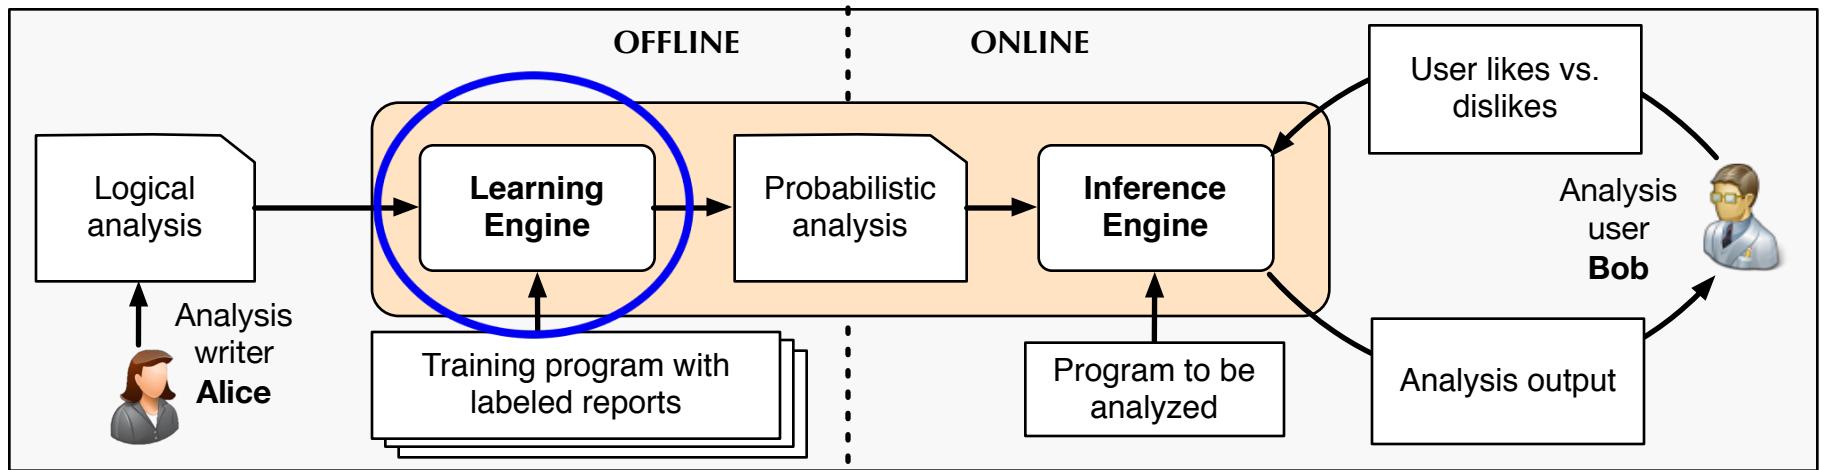
~~(parallel(y1, y1) \wedge next(y2, y1) \Rightarrow parallel(y2, y1))~~ weight 5 \wedge

~~(parallel(y2, y1) \wedge mayAlias(y2, y1) \wedge \neg guarded(y2, y1) \Rightarrow race(y2, y1))~~ \wedge

~~(parallel(x2, x1) \wedge mayAlias(x2, x1) \wedge \neg guarded(x2, x1) \Rightarrow race(x2, x1))~~ \wedge

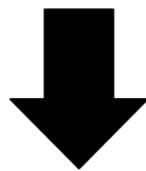
race(x2, x1) weight 25

Learning Engine



Weight Learning

Learn rule weights such that the probability
of the training data is maximized



Perform gradient descent
[Singla & Domingos, AAAI'05]

Empirical Evaluation Questions

- ▶ Does user feedback help in improving analysis precision?
- ▶ How much feedback is needed and does the amount of feedback affect the precision?
- ▶ How feasible is it for users to inspect analysis output and provide useful feedback?

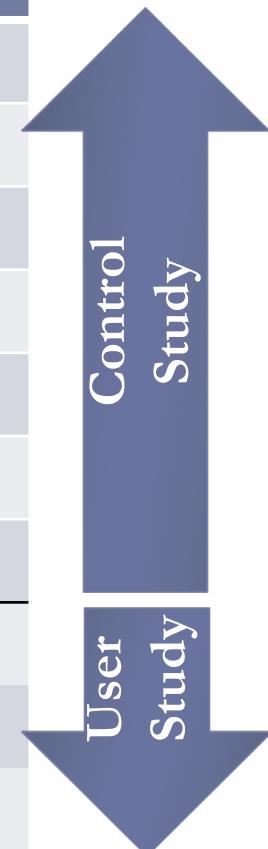
Empirical Evaluation Setup

- ▶ Control Study:
 - ▶ **Analyses:** (1) Pointer analysis, (2) Datarace analysis
 - ▶ **Benchmarks:** 7 Java programs (130-200 KLOC each)
 - ▶ **Feedback:** Automated [Zhang et.al, PLDI'14]

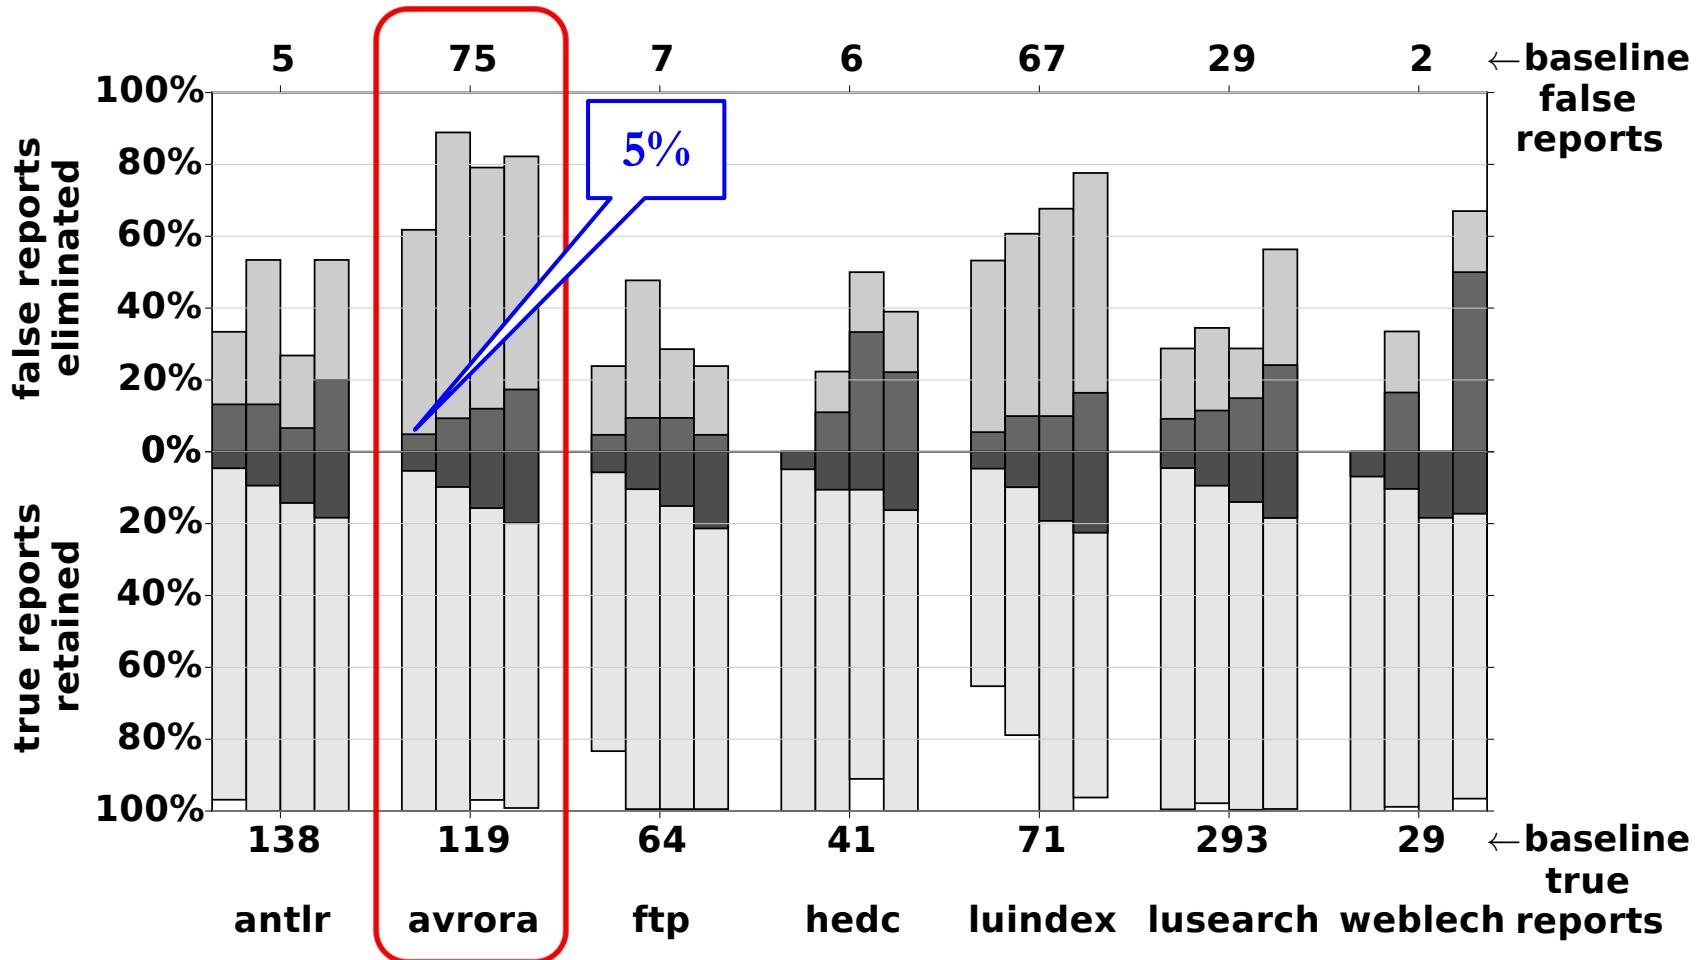
- ▶ User Study:
 - ▶ **Analyses:** Information flow analysis
 - ▶ **Benchmarks:** 3 security micro-benchmarks
 - ▶ **Feedback:** 9 users

Benchmarks Characteristics

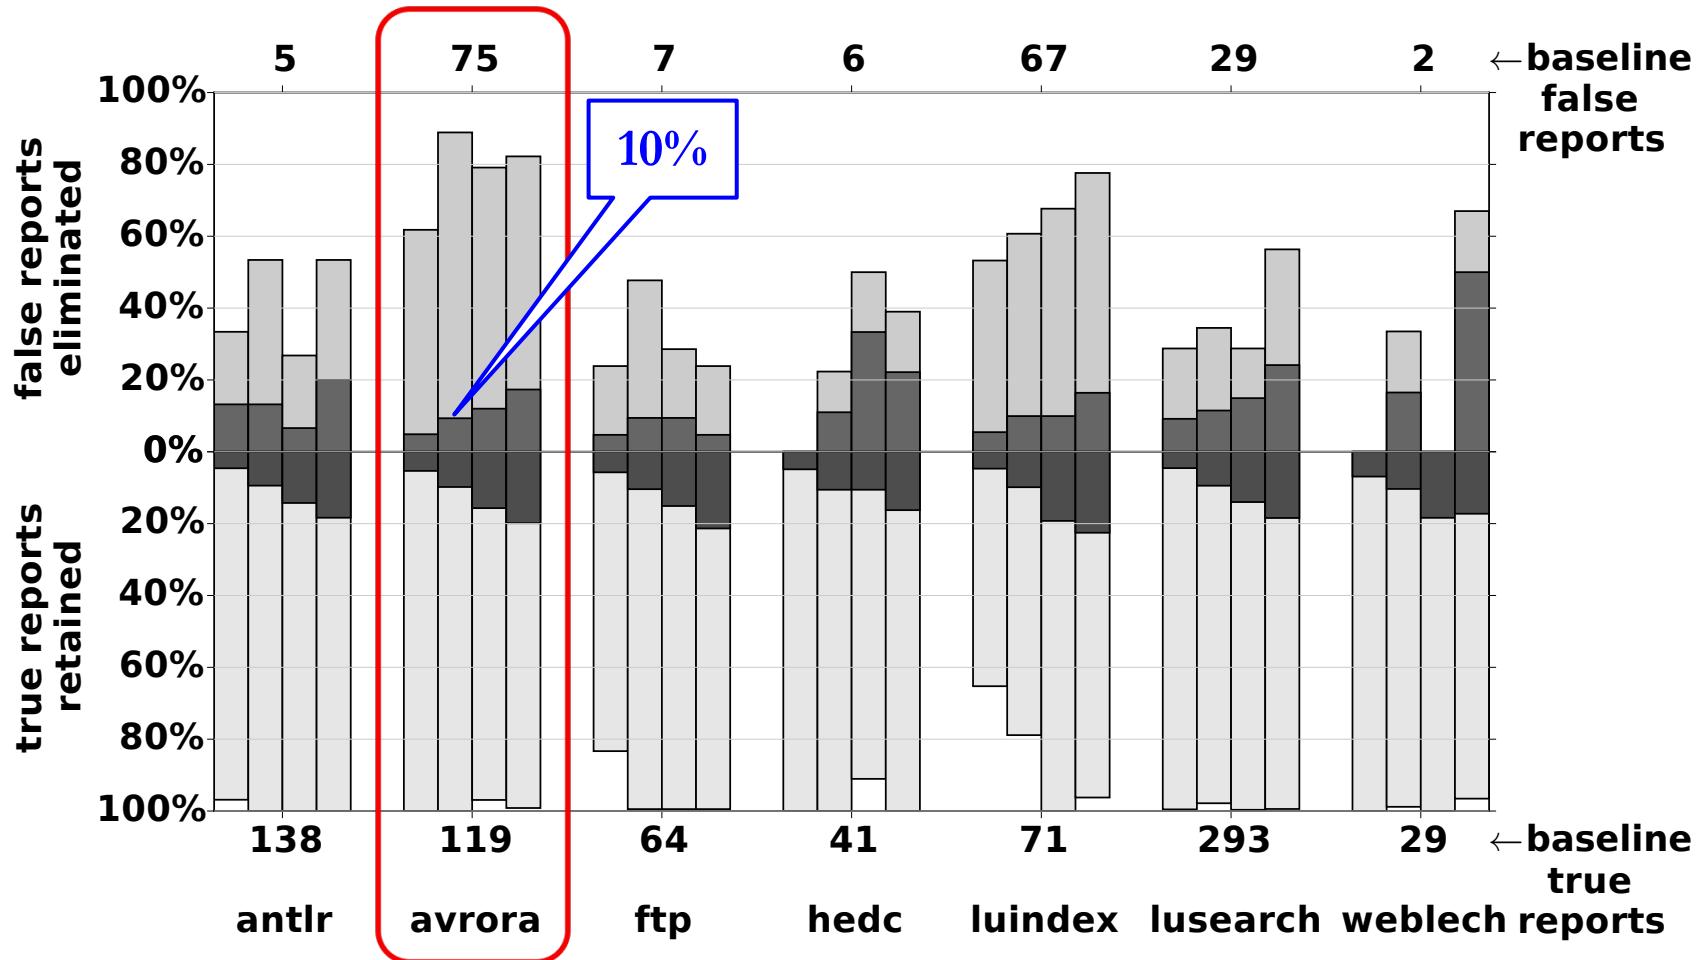
	classes	methods	bytecode(KB)	KLOC
antlr	350	2.3K	186	131
avrora	1,544	6.2K	325	193
ftp	414	2.2K	118	130
hedc	353	2.1K	140	153
luindex	619	3.7K	235	190
lusearch	640	3.9K	250	198
weblech	576	3.3K	208	194
secbench1	5	13	0.3	0.6
secbench2	4	12	0.2	0.6
secbench3	17	46	1.3	4.2



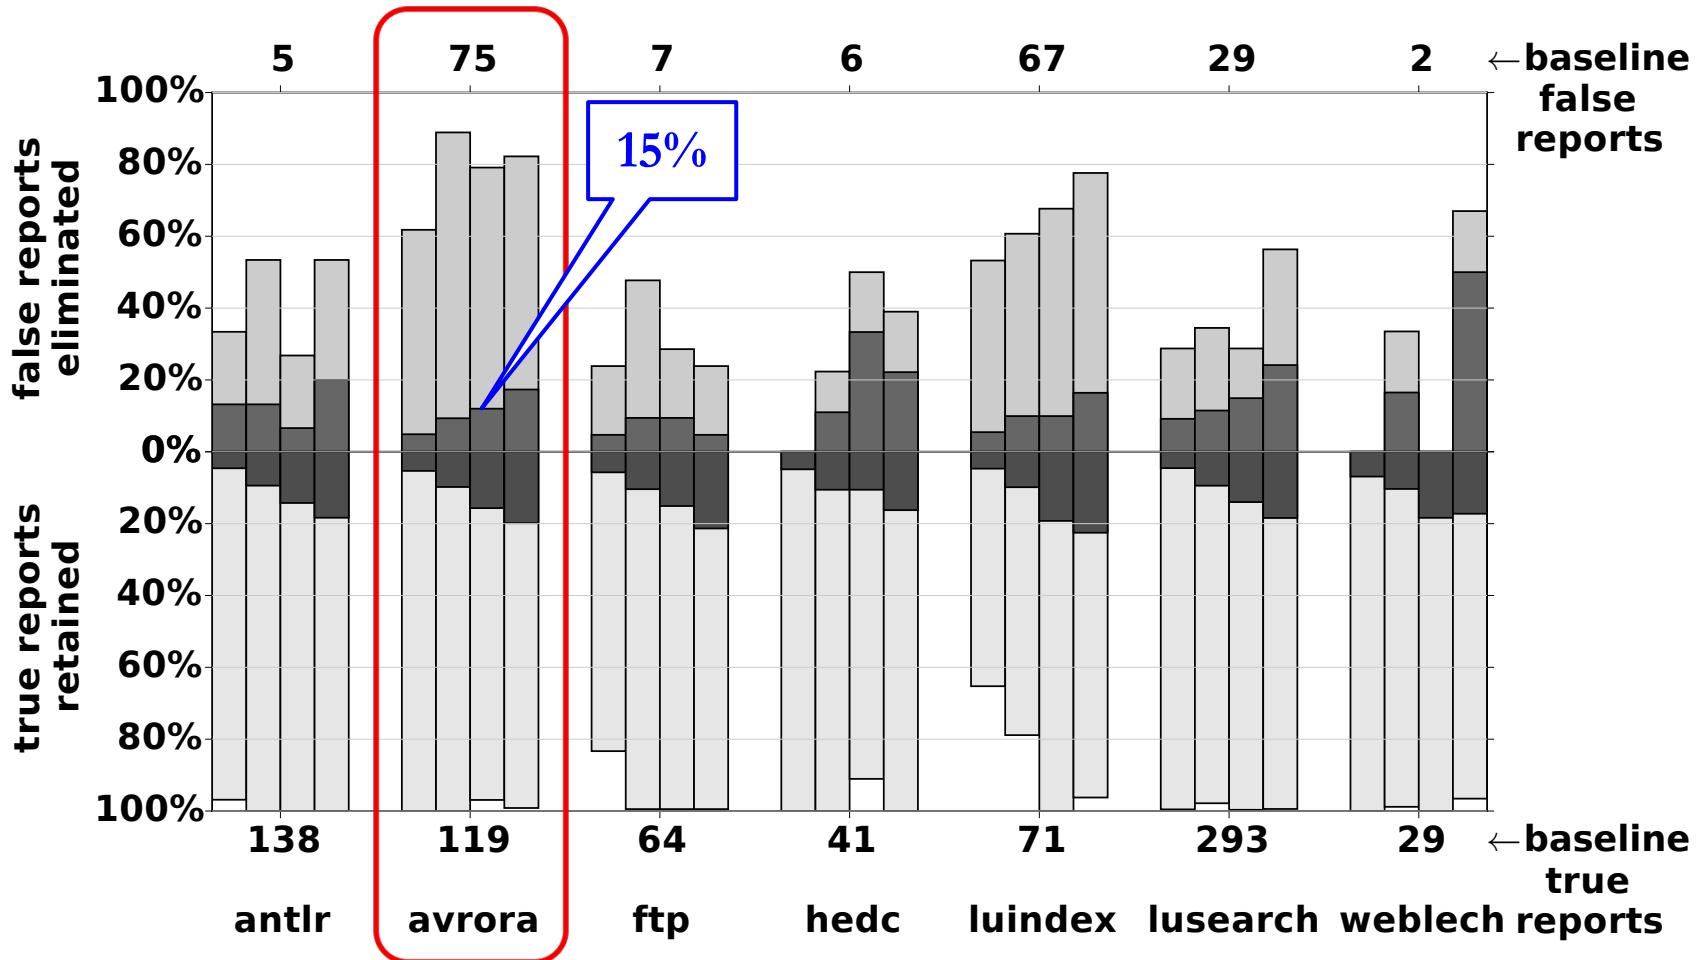
Precision Results: Pointer Analysis



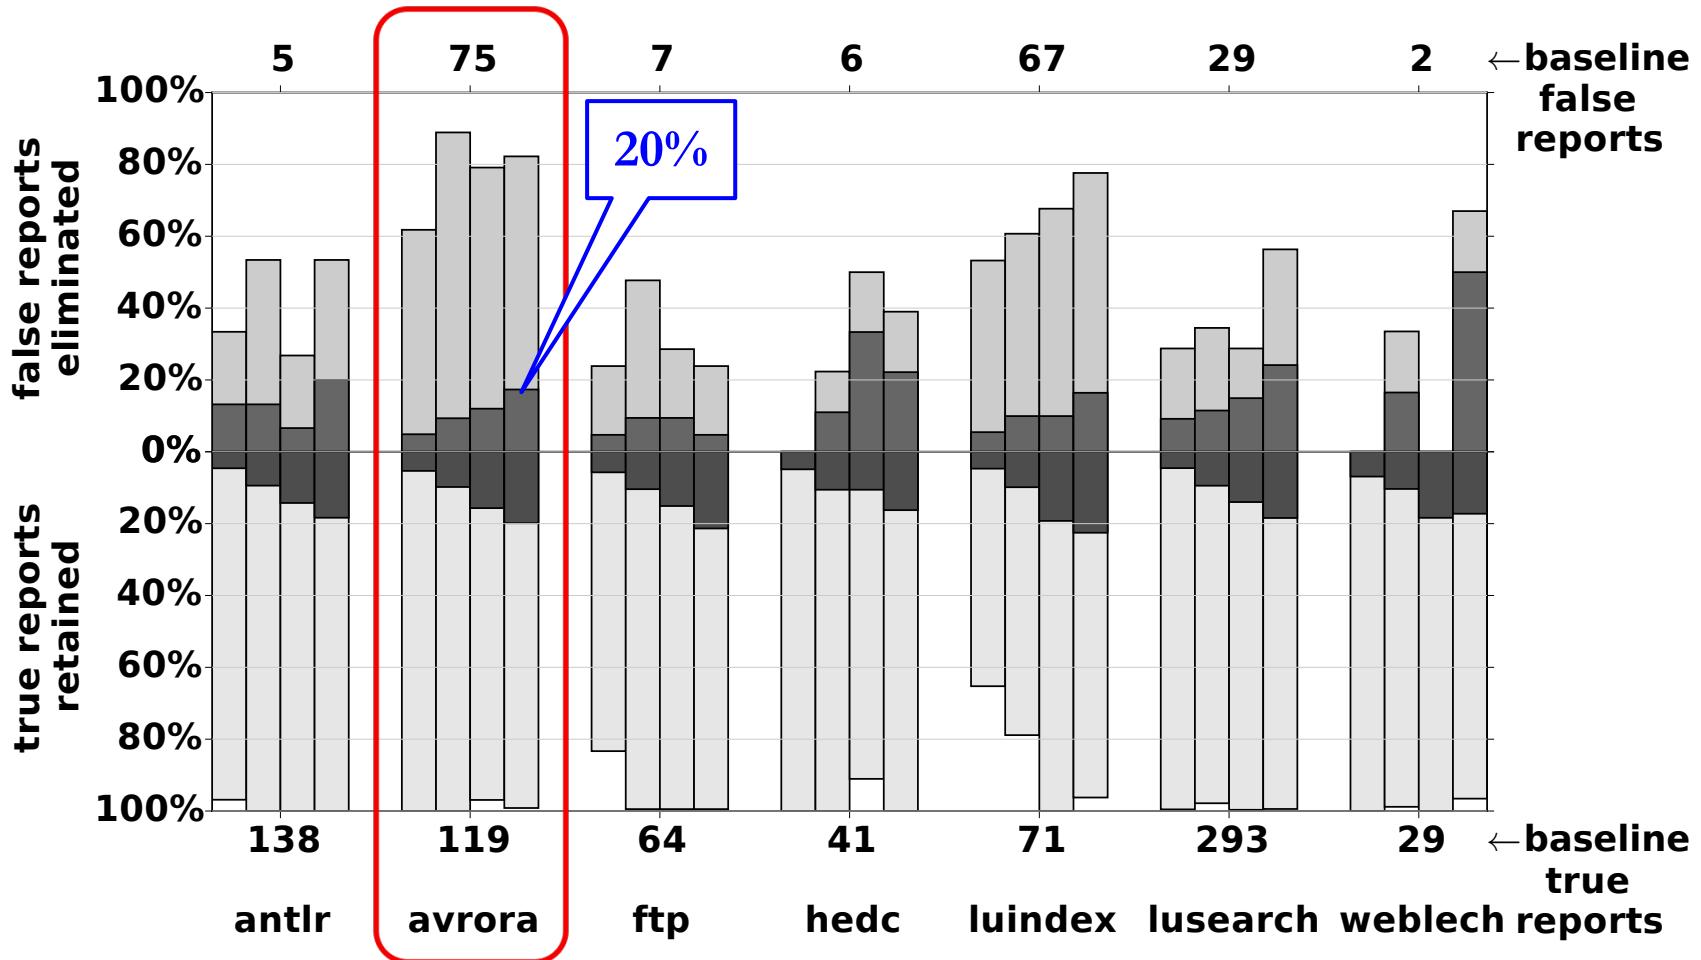
Precision Results: Pointer Analysis



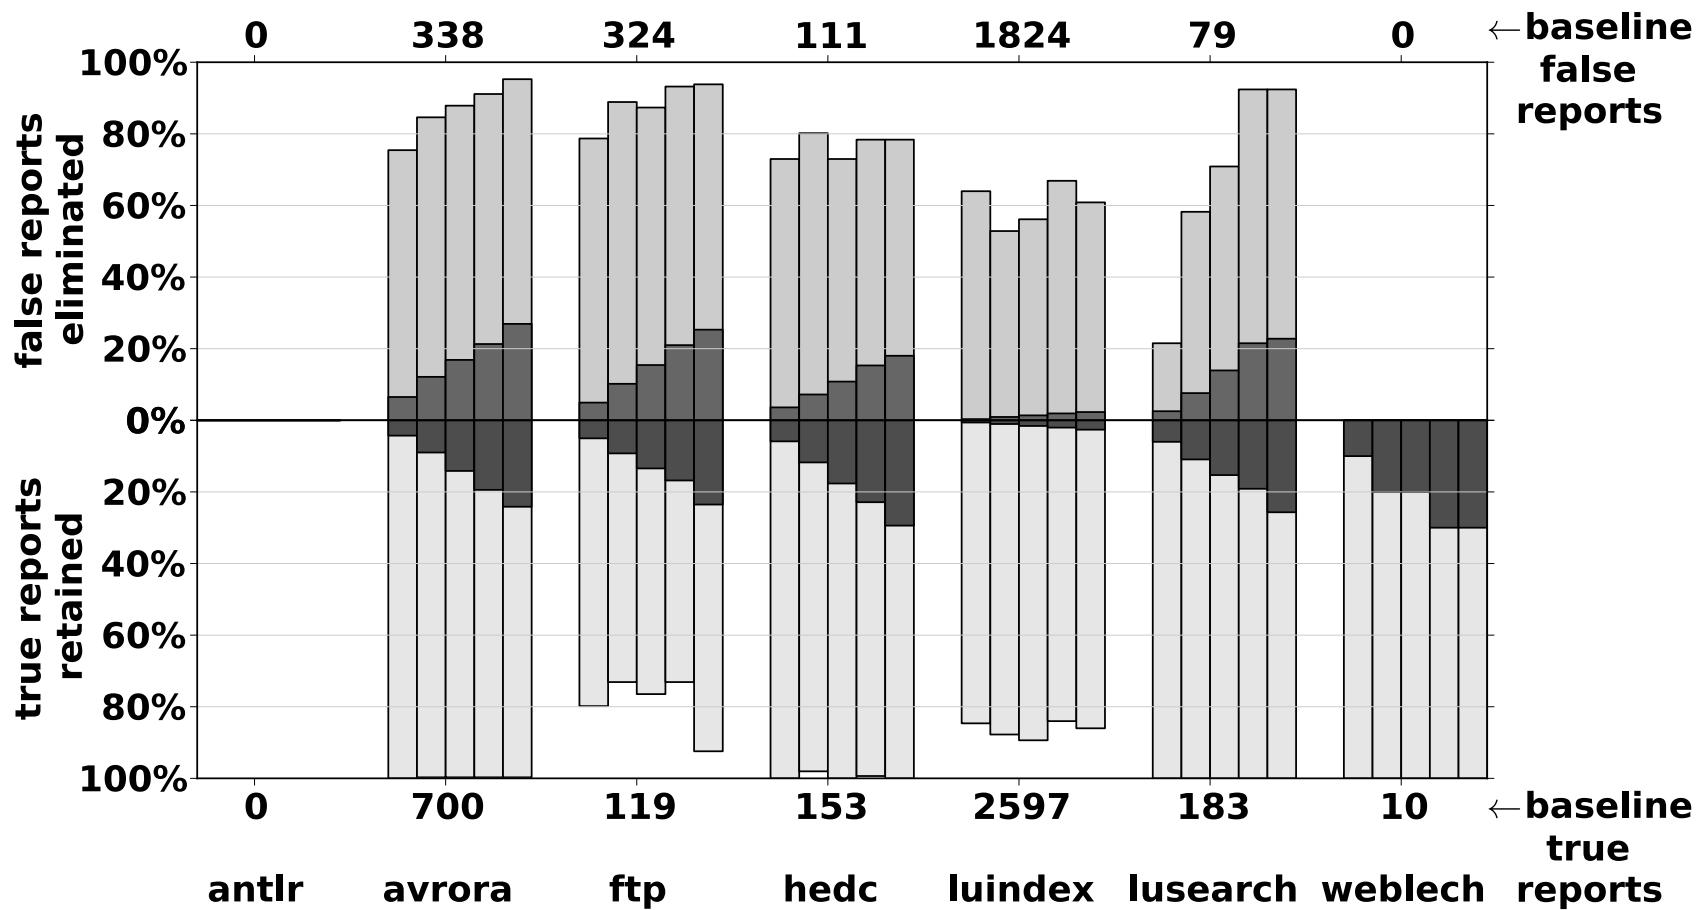
Precision Results: Pointer Analysis



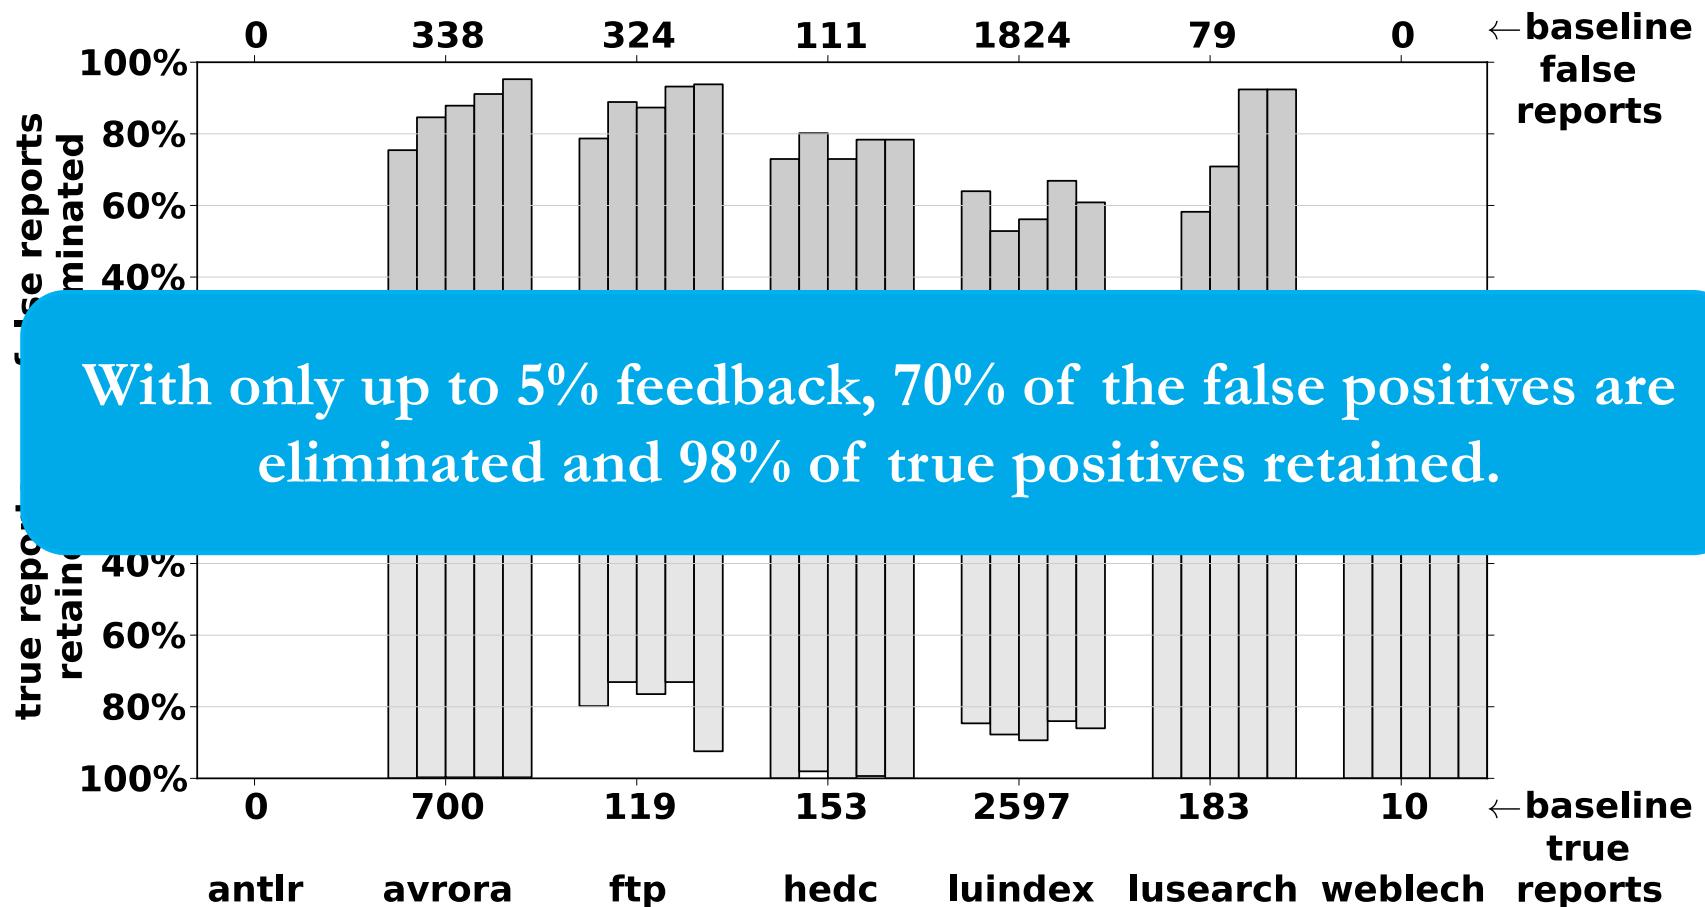
Precision Results: Pointer Analysis



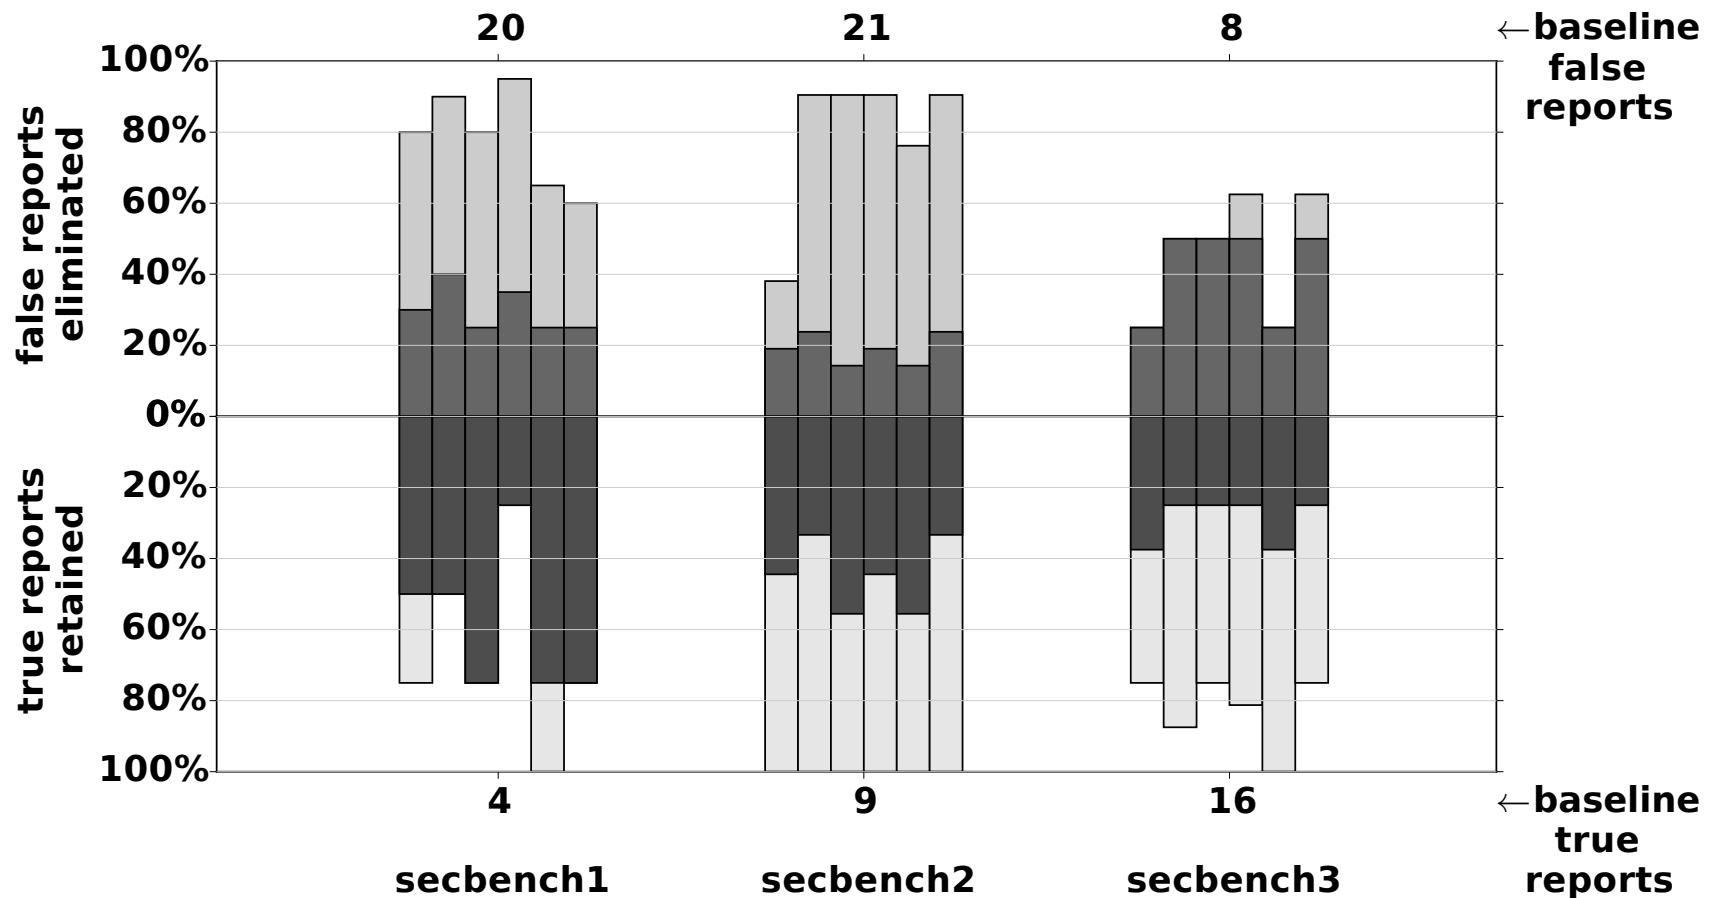
Precision Results: Datarace Analysis



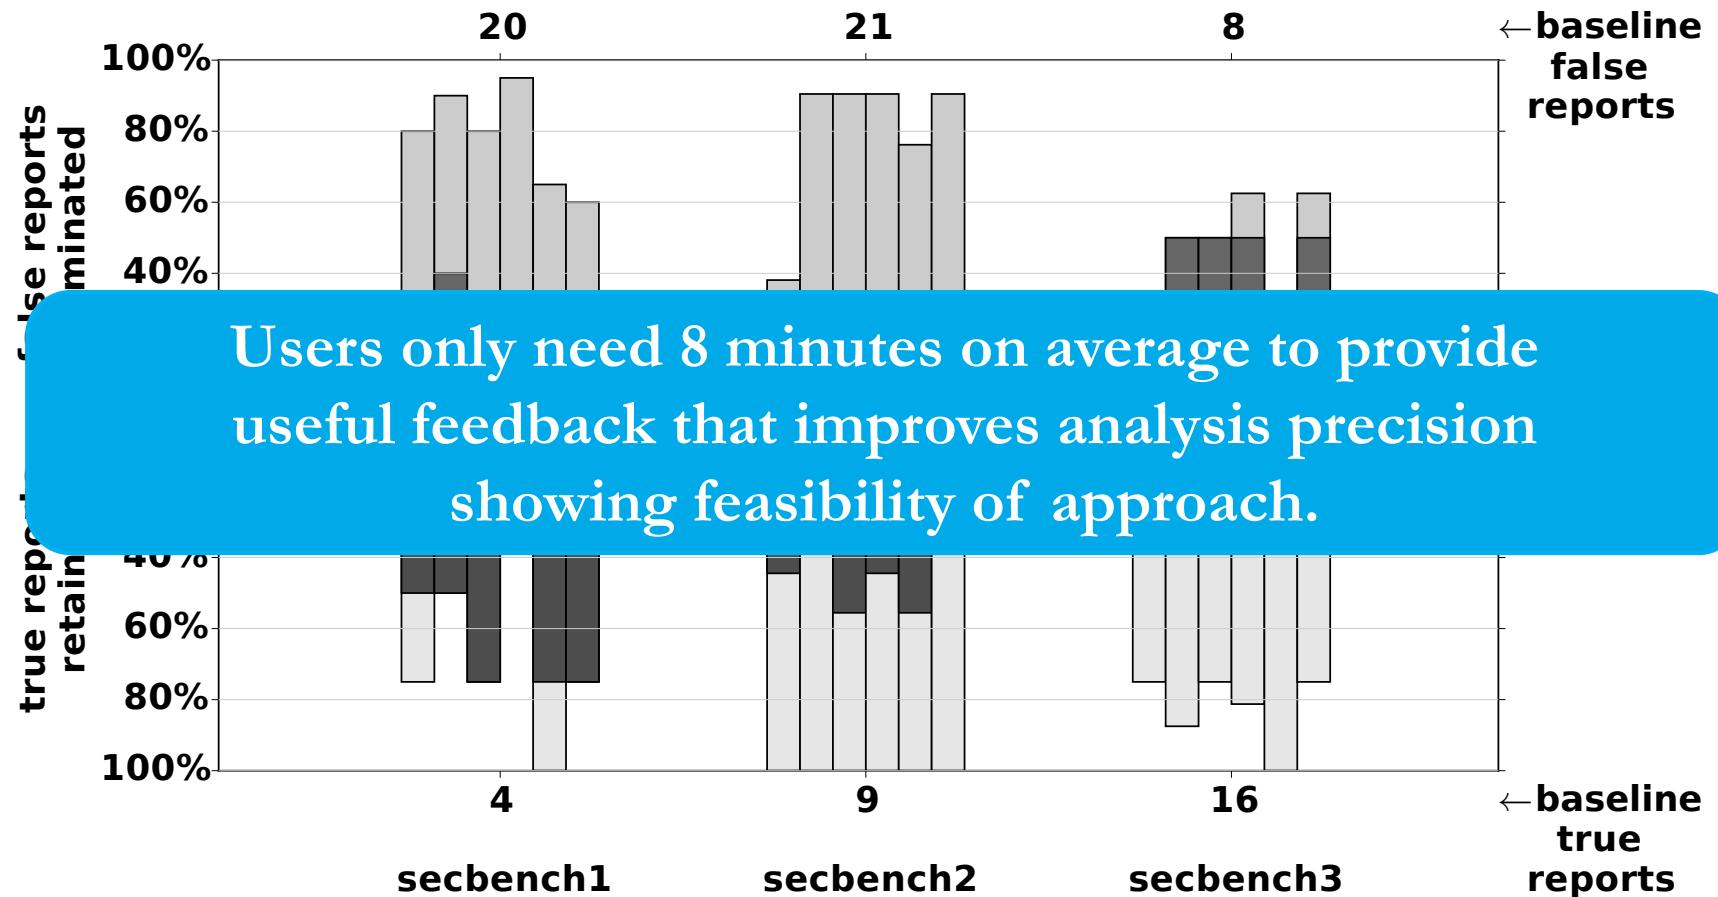
Precision Results: Datarace Analysis



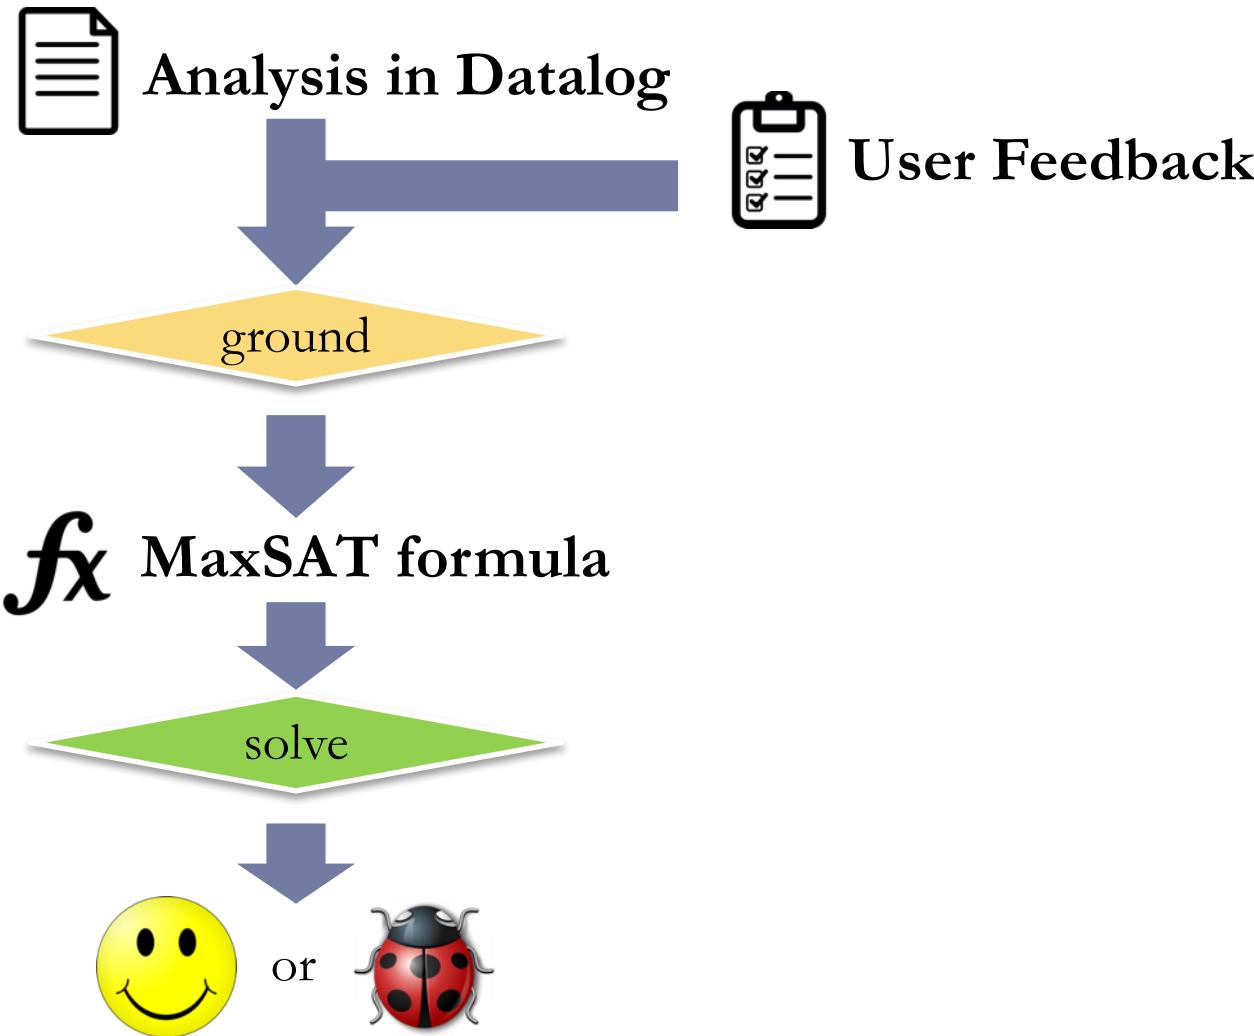
Precision Results: User Study



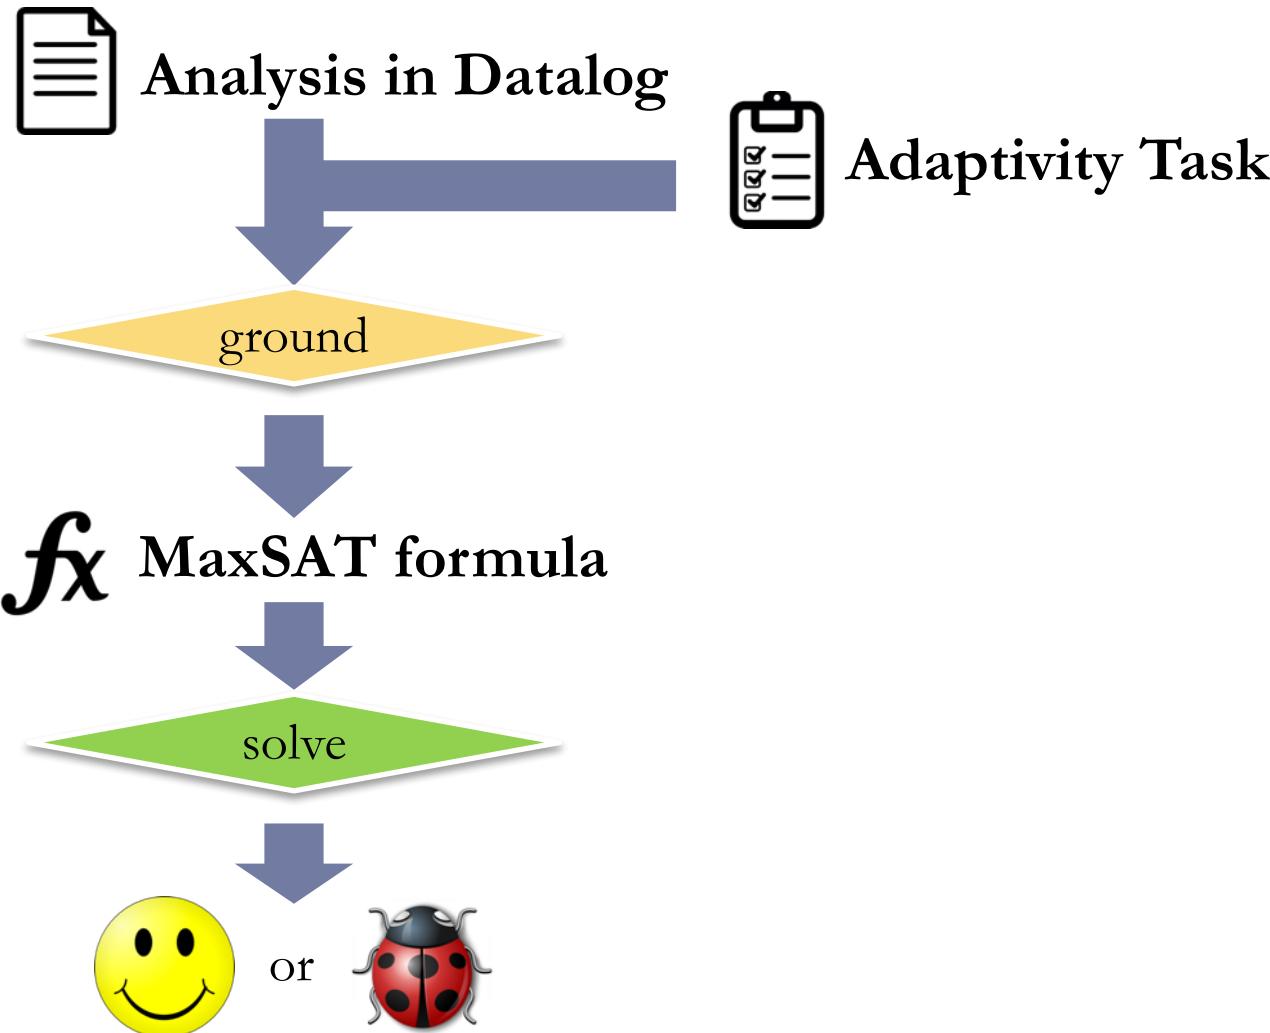
Precision Results: User Study



The Story So Far



A General Adaptivity Framework



A General Adaptivity Framework



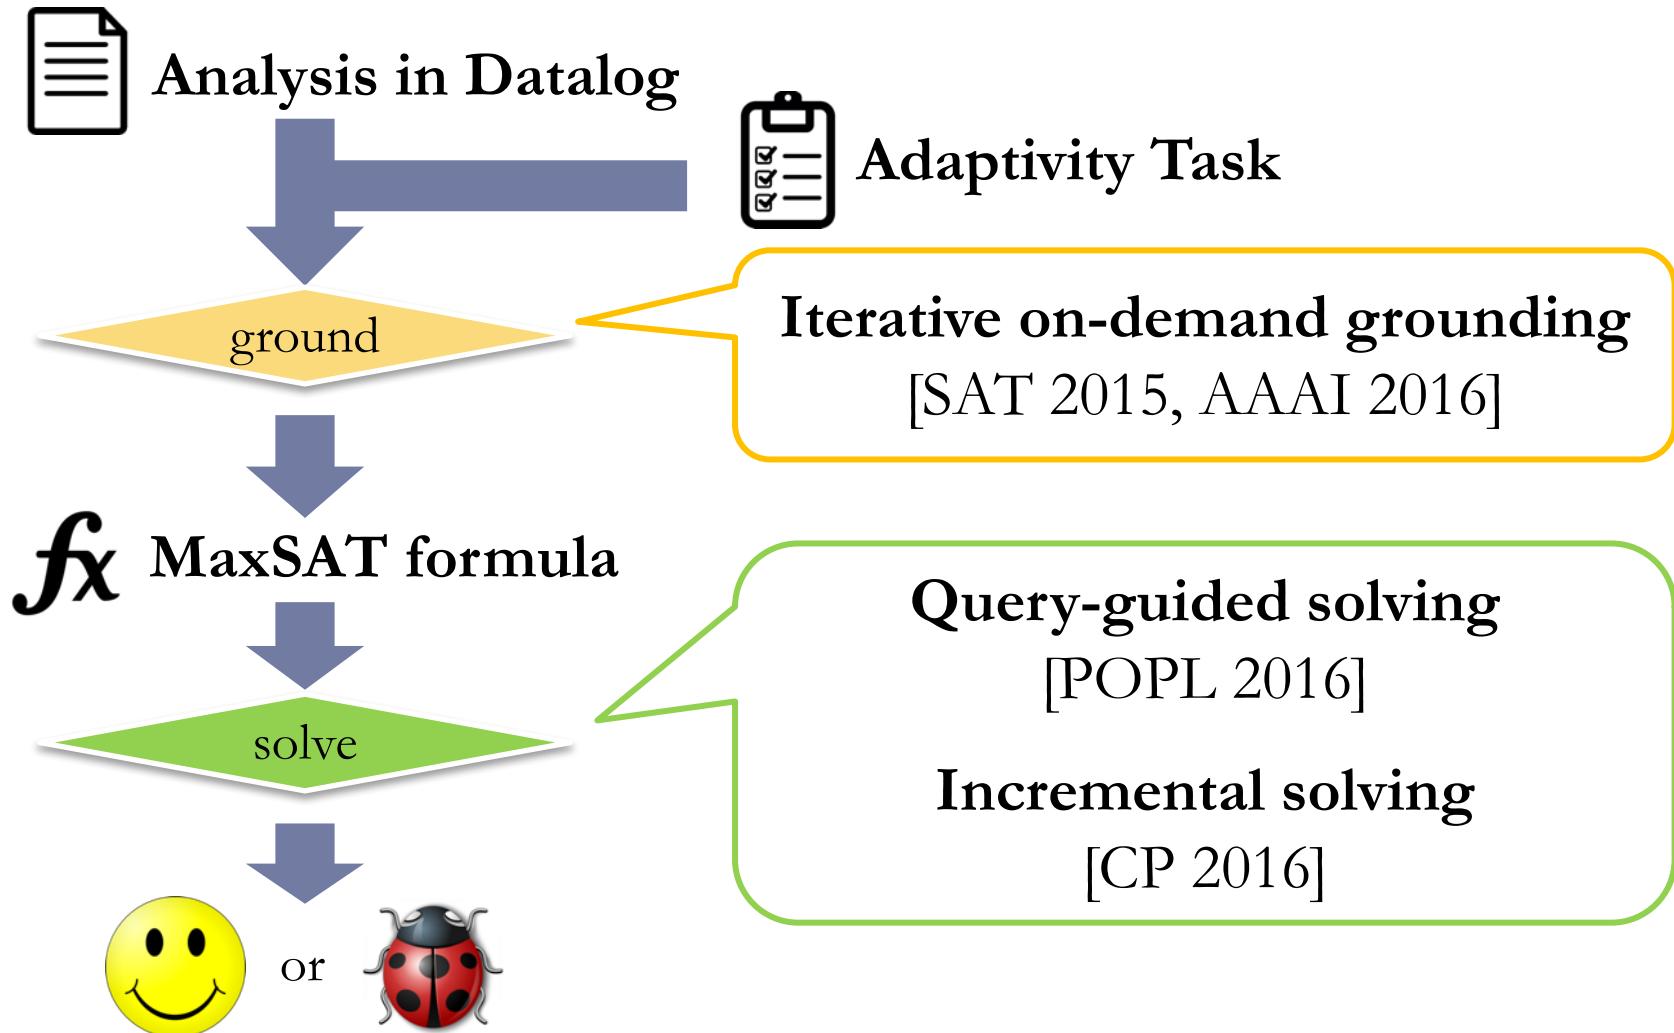
Analysis in Datalog



Adaptivity Task

What to adapt to	Technique	Focus	Tradeoff	Publications
User feedback	User-guided program analysis	Precision (bug-finding)	Soundness vs. completeness	FSE 2015
Assertions of interest	Abstraction refinement	Precision (verification)	Precision vs. scalability	PLDI 2013 PLDI 2014a
Procedure reuse within a program	Hybrid top-down, bottom-up analysis	Scalability (single-program)	Specialization vs. generalization	PLDI 2014b
Procedure reuse across programs	Transfer learning of analysis results	Scalability (across-program)		OOPSLA 2016

A General Adaptivity Framework



Conclusion

- ▶ A new paradigm that incorporates user feedback to guide program approximations
- ▶ Generalize user feedback on single bug report to other similar bug reports
- ▶ A unified framework for adapting program analyses in Datalog using MaxSAT