Towards automating reasoning over cybersecurity data

Jonathan M. Spring[1] and David Pym[1,2]

[1] University College London
[2] Alan Turing Institute

Data Science for Cybersecurity
September 26, 2017

Goal

Discuss results of Pym et al. [2017]

- Use of Separation Logic for finding memory management errors

Motivate application of similar approach to reasoning in cybersecurity

- Hope to avoid some shortcomings of machine learning
- Sketch application to incident response (IR) and investigation

This contributes to a broader discussion involving evidence evaluation [Spring and Illari, 2017] and philosophy of security [Spring et al., 2017, Hatleback and Spring, 2014].

J. Pearl

Hierarchy of causal importance [Pearl, 2016]:

- What is?
  - "can be computed efficiently using Bayesian Networks, or any of the graphical models that support deep-learning systems"
- What if?
  - Requires reasoning about interventions, see Woodward [2003], Halpern [2015]
- Why?
  - Requires manipulable models with structure
  - Pearl suggests Counterfactuals

Pearl's is not the only answer to "what if" & why

Separation Logic

But the answer is not traditional ML

SL introduces $*$ for "and, separately"

- Distinct from $\wedge$ for the usual sense of "and"

We take lessons from the history of SL implementation in Infer, file systems, OS scheduling, etc.
SL works

- Infer is open source, used by FB, others
- What we wanted to know was *why it works*

Enabling Innovations

- Overlap of two (types of) models:
  - Useful engineering model of RAM
  - Resource-based logical model and language
- Scalable proof theory
  - Compositional local reasoning via a modified Frame Rule
  - Automation of abductive inference
- Hard work adapting to actual software development practices (see O'Hearn [2015] for this)

Abduction

Third reasoning type, vice induction and deduction

Initially introduced by Peirce around 1900:

> *'Abduction is the process of forming an explanatory hypothesis. It is the only logical operation which introduces any new idea'*
>
> *[Bergman and Paavola, 2016, CP 5.171]*

Abduction

In Separation Logic, this looks like:

1. Attempt a proof of a code segment
   - 1.1 Fail
2. Using purpose-built proof rules check:
   - 2.1 If the heuristic matches the failed proof state
   - 2.2 What additional condition would be necessary for the proof to work
   - 2.3 Guess the program contains this condition somewhere
     - 2.3.1 And look for it

The proof-rules match engineer's heuristics for fixing bugs

Abduction

Instead of a software developer's bug-fixing heuristics, we want to re-make this for:

> *Heuristics to figure out how an attacker broke in during computer security incident response*

IR Primer

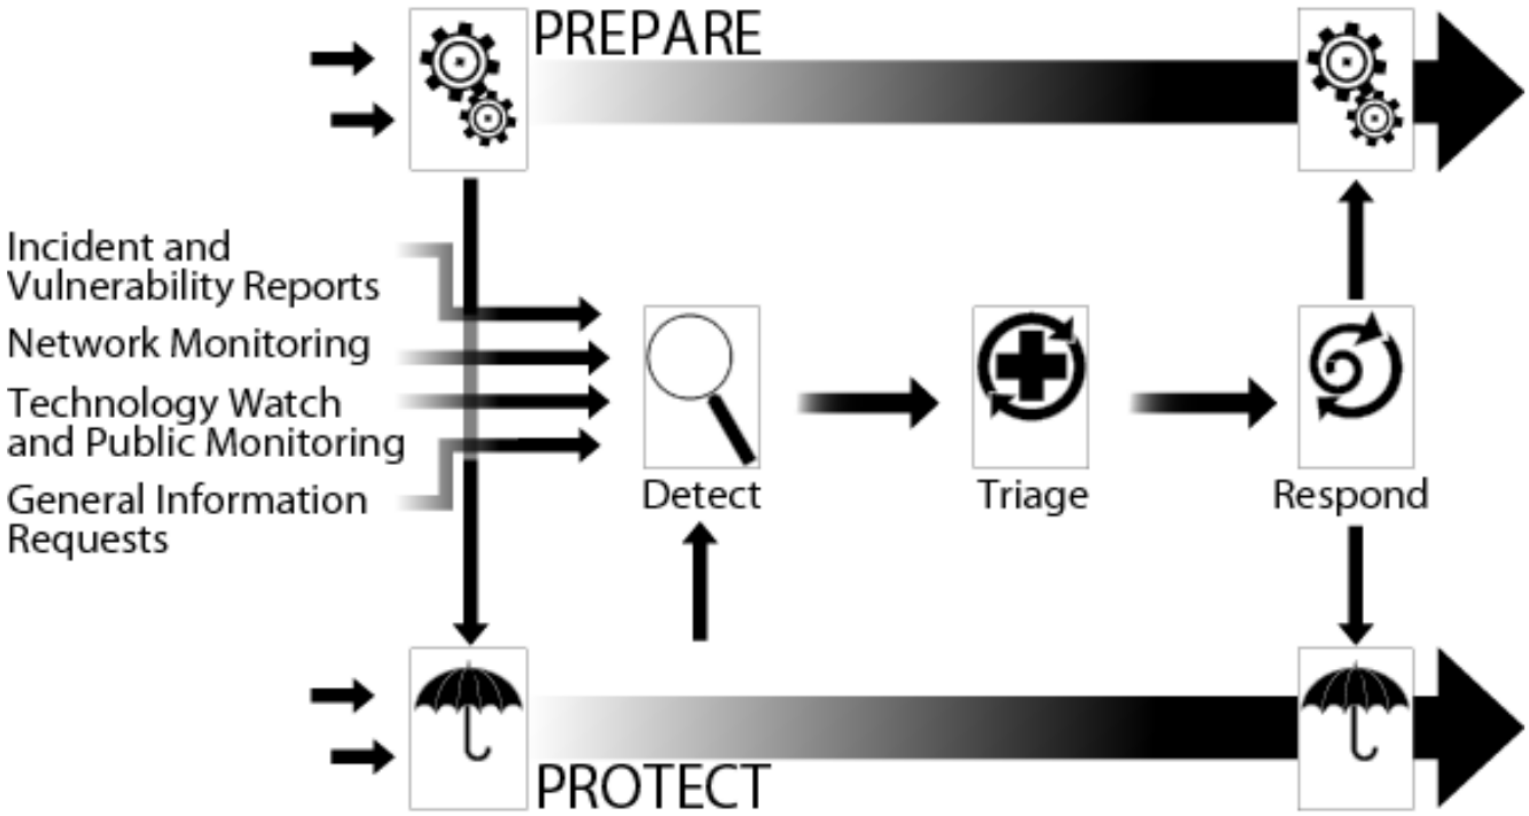

Incident management at a high level
(Figure 3 from Alberts et al. [2004])

IR Primer

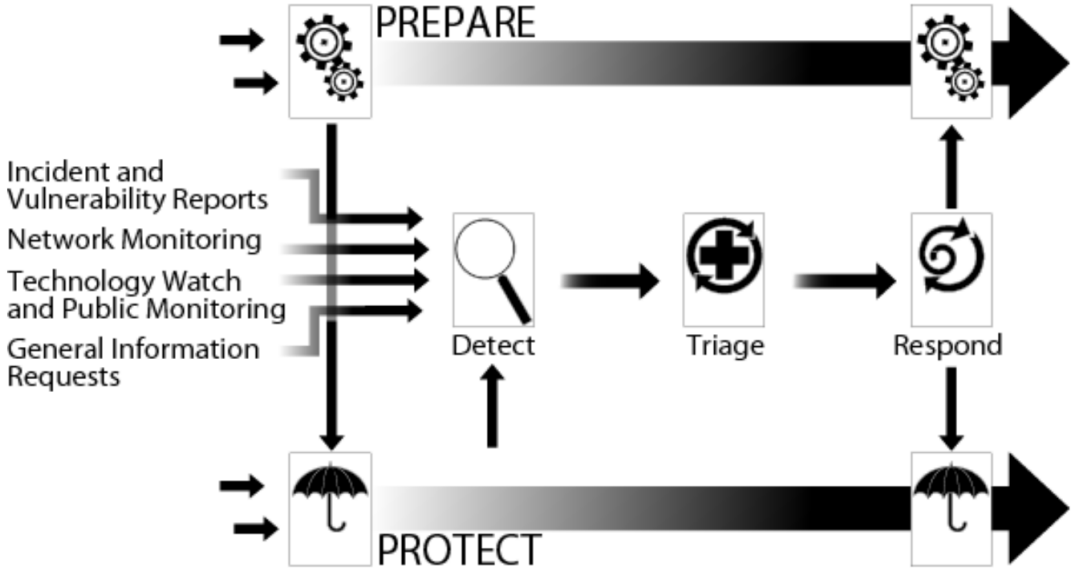

Intrusions

IR is a response to intrusions, several models:

- A "Common Language" [Howard and Longstaff, 1998]
- 5 steps [Bejtlich, 2004]
- 7-step "kill chain" [Hutchins et al., 2011]
- "Diamond model" [Caltagirone et al., 2013]

Investigators ask "what-if" questions based on models like these [Spring and Hatleback, 2017]

- Within "response" we have parts, too
- Investigation is: evidence collection, analysis, and reporting

Intrusions

IR is a response to intrusions, several models:

- A "Common Language" [Howard and Longstaff, 1998]
- 5 steps [Bejtlich, 2004]
- 7-step "kill chain" [Hutchins et al., 2011]
- "Diamond model" [Caltagirone et al., 2013]

Investigators ask "what-if" questions based on models like these [Spring and Hatleback, 2017]

- Within "response" we have parts, too
- Investigation is: evidence collection, analysis, and reporting

Investigation is:

A process carried out by an *agent* to build a*n* *explanation* of the *mechanism* by which security policy was violated.[1] Forensic investigation is historical.[2] The process includes *data* collection, in some *jargon*, via manipulation of available *resources*. The agent's *modelling decisions* are resource-constrained, and follow a *methodology*. The output of the process is that the agent *reports* results, namely, their new *beliefs* relative to *goals*.

[1] More generally, the phenomenon of interest

[2] As opposed to engineering or design which are future-oriented.

Investigation is:

A process carried out by an *agent* to build a*n* *explanation* of the *mechanism* by which security policy was violated.[1] Forensic investigation is historical.[2] The process includes *data* collection, in some *jargon*, via manipulation of available *resources*. The agent's *modelling decisions* are resource-constrained, and follow a *methodology*. The output of the process is that the agent *reports* results, namely, their new *beliefs* relative to *goals*.

[1] More generally, the phenomenon of interest

[2] As opposed to engineering or design which are future-oriented.

Investigation is:

A process carried out by an *agent* to build a*n* *explanation* of the *mechanism* by which security policy was violated.[1] Forensic investigation is historical.[2] The process includes *data* collection, in some *jargon,* via manipulation of available *resources*. The agent's *modelling decisions* are resource-constrained, and follow a *methodology*. The output of the process is that the agent *reports* results, namely, their new *beliefs* relative to *goals*.

[1] More generally, the phenomenon of interest

[2] As opposed to engineering or design which are future-oriented.

In a logic this would be...

We need a concept of time

- Temporal logic tends to be about the future
- Existing SL-TL mixtures keep a syntactic separation [Espinosa and Brotherston, 2017]

Plan to use histories of change as a resource

- Allows interleaving of different operators

Start with just one investigator for simplicity

- Other work talks about coordination between CSIRTs [Osorno et al., 2011]

In a logic this would be...

We need to understand what heuristics incident responders use

- Science can give generalization & discovery heuristics [Spring and Illari, 2017]
- Extract some from the process of mathematical modelling
- Draw on standards [Cichonski et al., 2012]
- Draw on case studies [Stoll, 1989, Mandiant, 2013]

The heuristics will not be perfect at first, but at least the language should allow for testing and evaluation

Summary

Reasoning about cybersecurity data cannot rely on black-box algorithms

- Need to know 'What if' and Why
- 'What is' questions have a place

There are languages that enable such questions
History of Infer shows Separation Logic is one such language

- Also teaches there are no shortcuts
- Heuristics and tools must be carefully tuned to the problem of interest

Questions?

"Why Separation Logic Works" draft available:
http://www0.cs.ucl.ac.uk/staff/D.Pym/recent.htm
Contact: jonathan.spring.15 (AT) ucl.ac.uk

References I

Chris Alberts, Audrey Dorofee, Georgia Killcrece, Robin Ruefle, and Mark Zajicek. Defining incident management processes for CSIRTS: A work in progress. Technical Report CMU/SEI-2004-TR-015, Software Engineering Institute, Carnegie Mellon University, 2004.

Richard Bejtlich. *The Tao of network security monitoring: beyond intrusion detection*. Pearson Education, 2004.

Mats Bergman and Sami Paavola. 'Abduction': term in The Commens Dictionary: Peirce's Terms in His Own Words. New Edition. http://www.commens.org/dictionary/term/abduction, Jul 14, 2016.

Sergio Caltagirone, Andrew Pendergast, and Christopher Betz. The diamond model of intrusion analysis. Technical report, Center for Cyber Intelligence Analysis and Threat Research, 2013. http://www.threatconnect.com/methodology/diamond_model_of_intrusion_analysis.

Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. Computer security incident handling guide. Technical Report SP 800-61r2, US Dept of Commerce, National Institute of Standards and Technology, Gaithersburg, MD, Aug 2012.

Gadi Tellez Espinosa and James Brotherston. Automatically verifying temporal properties of programs with cyclic proof. In *CADE-26*, volume 10395 of *LNAI*, pages 491–508. Springer, 2017.

Joseph Y. Halpern. A Modification of the Halpern-Pearl Definition of Causality. *ArXiv e-prints*, May 2015. URL http://arxiv.org/abs/1505.00162. extended from IJCAI 2015 version.

References II

Eric Hatleback and Jonathan M. Spring. Exploring a mechanistic approach to experimentation in computing. *Philosophy & Technology*, 27(3):441–459, 2014.

John D Howard and Thomas A Longstaff. A common language for computer security incidents. Technical Report SAND98-8667, Sandia National Laboratories, Oct 1998.

Eric M Hutchins, Michael J Cloppert, and Rohan M Amin. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research*, 1:80, 2011.

Mandiant. APT1: Exposing one of china's cyber espionage units. Technical report, 2013.

Peter W. O'Hearn. From Categorical Logic to Facebook Engineering. In *Logic in Computer Science (LICS)*, pages 17–20. IEEE, 2015.

Marcos Osorno, Thomas Millar, and Danielle Rager. Coordinated cybersecurity incident handling: Roles, processes, and coordination networks for crosscutting incidents. Technical report, Johns Hopkins Univ, Applied Physics Laboratory, Laurel, MD, Jun 2011.

Judea Pearl. Theoretical impediments to machine learning: A position paper. Nov 2016.

David Pym, Jonathan M. Spring, and Peter O'Hearn. Why separation logic works. *Under review – Philosophy & Technology*, 2017.

References III

Jonathan M. Spring and Eric Hatleback. Thinking about intrusion kill chains as mechanisms. *Journal of Cybersecurity*, 2(2), 2017.

Jonathan M. Spring and Phyllis Illari. Mechanisms and generality in information security. *Under review – Philosophy & Technology*, 2017.

Jonathan M. Spring, Tyler Moore, and David Pym. Practicing a science of security: A philosophy of science perspective. In *New Security Paradigms Workshop*, Islamorada, FL, Oct 2-4, 2017.

Clifford Stoll. *The cuckoo's egg: tracking a spy through the maze of computer espionage*. Pan Books, London, 1989.

James Woodward. *Making things happen: A theory of causal explanation*. Oxford University Press, Oxford, 2003.

Copyright

Acknowledgements

The Separation Logic paper, from which these slides are partly derived, is joint work with Peter O'Hearn.
Spring is supported by University College London's Overseas Research Scholarship and Graduate Research Scholarship.
orcid.org/0000-0001-9356-219X