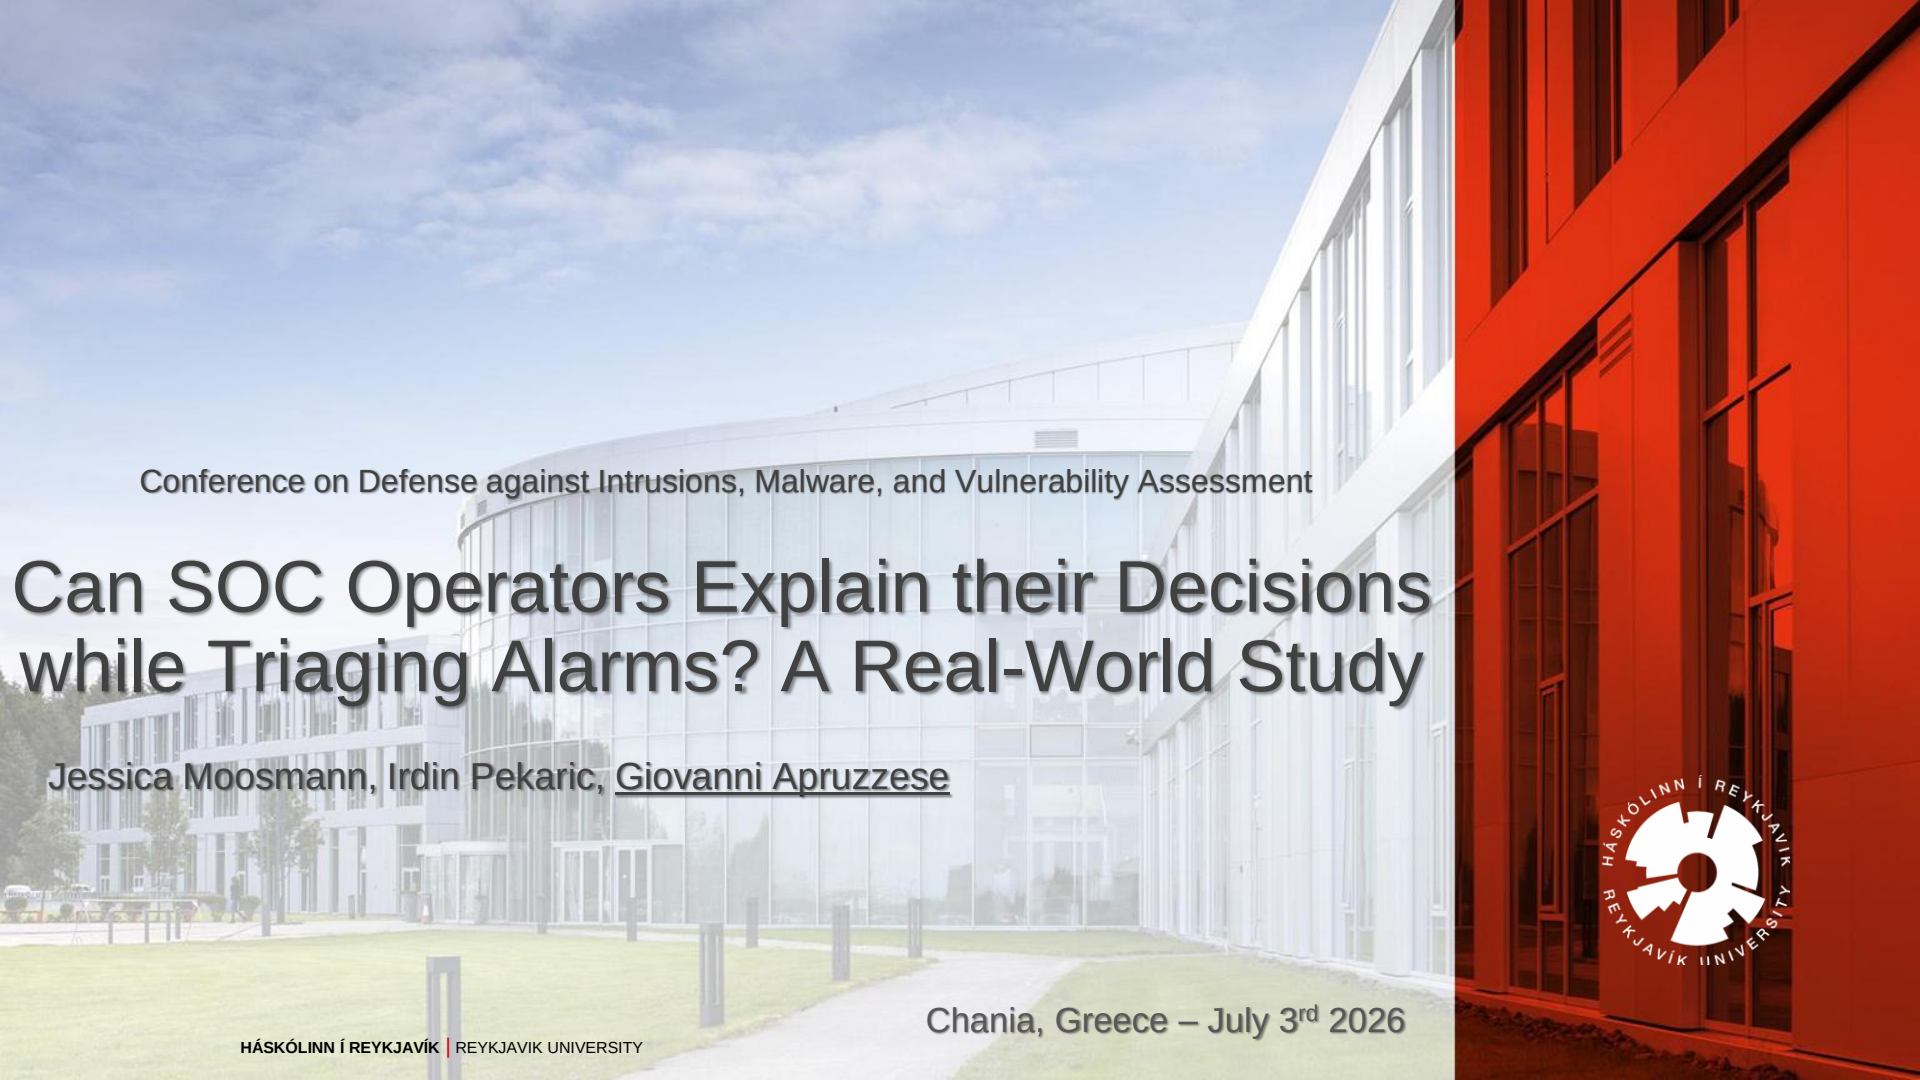

Conference on Defense against Intrusions, Malware, and Vulnerability Assessment

Can SOC Operators Explain their Decisions while Triaging Alarms? A Real-World Study

Jessica Moosmann, Irdin Pekaric, Giovanni Apruzzese

Chania, Greece – July 3rd 2026

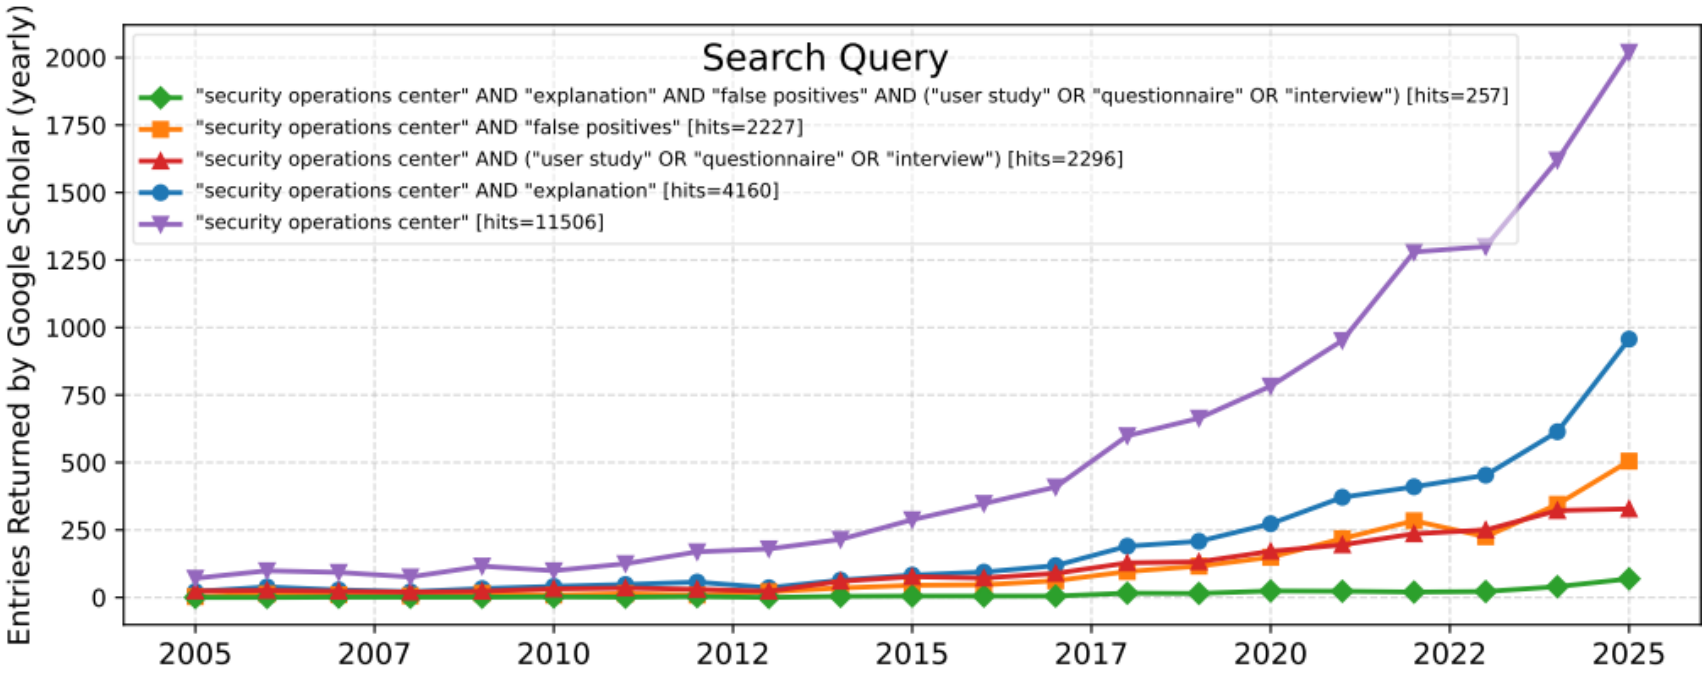

Giovanni Apruzzese
giovannia@ru.is

Systematic Literature Review (SLR)

RQ0: "*has prior work carried out user studies in which SOC practitioners were asked to triage alarms and explain their decisions?*"

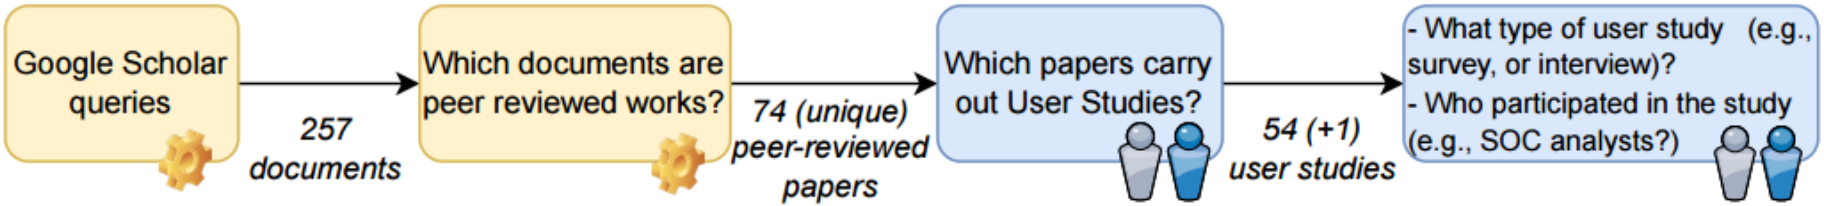

Fig. 3: Overview of the methodology of our SLR. The qualitative analysis was done by two authors.

Answer to RQ0: Only one work [29] could be said to have investigated whether real SOC practitioners can correctly explain the decisions they make.

[29] Kersten, L., Darr´e, S., Mulders, T., Zambon, E., Caselli, M., Snijders, C., Allodi, L.: A security alert investigation tool supporting tier 1 analysts in contextualizing and understanding network security events. In: ACSAC (2024)

Giovanni Apruzzese
giovannia@ru.is

Giovanni Apruzzese
giovannia@ru.is

Field Study in a Real SOC

- We partnered up with a SOC operating in the DACH area

Giovanni Apruzzese
giovannia@ru.is

Field Study in a Real SOC

- We partnered up with a SOC operating in the DACH area
- We recruited 12 employees of this SOC
- Their role: participate in a field study (~2 hours) in which they would be presented with a set of alarms shown in a security dashboard and:
    a) Guess whether they are ‘true’ or ‘false’ alarms
    b) Justify their decision (in writing)

Giovanni Apruzzese
giovannia@ru.is

Study Design and Challenges

- The alarms shall be drawn from the SOC---but shall not be drawn from those that participants may have seen in the past

Giovanni Apruzzese
giovannia@ru.is

Study Design and Challenges

- The alarms shall be drawn from the SOC---but shall not be drawn from those that participants may have seen in the past

- We aligned the setup with that of the SOC

- Time limited: 2 hours for the entire activity, of which:
  - 45 minutes allocated for debriefing and familiarization
  - 30 minutes for the main study

Giovanni Apruzzese
giovannia@ru.is

Cases

Table 1: **Summary of the Six Cases of Alerts included in our Study.** Brief description: *Case-1*: Scheduled vulnerability scan triggered alarms. *Case-2*: Firewall allows specific developer traffic. *Case-3*: DNS sinkhole detects C&C traffic; connection blocked. *Case-4*: HTTP redirect to `msftconnecttest` misclassified as suspicious. *Case-5*; WiFi client triggered alert due to network instability. *Case-6*: Endpoint protection quarantined infected file. (More details in our repository [1])

Case	Category	Alerts	Logged Events	Alert Type	Key Observables	Source	Action / Indicator
1: Network scan by internal system	False Positive - trivial (with context) / relevant (without)	15	2,034 over 4h	Port scanning / suspicious network activity	IP source, IP destination, hit count, firewall action	Internal IP (Nessus server)	Accept; system role identified in asset database
2: Developer activity triggering alerts	False Positive - trivial	4	1,340 over 48h	Outbound connections from internal IPs	IP source, IP destination, hit count, protocol, firewall action, category, properties	Internal developer-segment IPs	Accept; triggered by developer-specific firewall rule
3: Outbound C2 communication	True Positive	1	2 over 10h	Outbound connection to Command & Control server	IP source, IP destination, hit count, firewall action, malware action	Internal host	Prevent; DNS trap and malware action confirm C2 activity
4: Windows connectivity check	False Positive - trivial	2	31,518 in 24h	Outbound HTTP/S traffic to connectivity-check URL	IP source, IP destination, hit count, protocol, firewall action	Internal load balancer	Pass; traffic to `msftconnecttest` incl. HTTP redirect
5: Internal retransmission classified as malicious	False Positive - relevant	1	2 in 2h	Internal retransmission pattern flagged as malicious	IP source, IP destination, hit count, protocol number, detected host, domain, action, reason	Internal dynamically assigned WiFi IP	Block; triggered because IP source belongs to WiFi pool
6: Malicious PowerShell script	True Positive	1	2 in 20h	Endpoint protection alert: suspicious PowerShell script	Hit count, detected host, host, domain, user, label, reason	Endpoint (self-reported)	File blocked & quarantined; virus label, affected file, external-account username

Giovanni Apruzzese
giovannia@ru.is

Cases ☺

#	Type
1	False positive
2	False positive
3	True Positive
4	False Positive
5	False Positive
6	True Positive

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1												
P2												
P3												
P4												
P5												
P6												
P7												
P8												
P9												
P10												
P11												
P12												
Correct												

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓										
P2	✓	✗										
P3	✓	✓										
P4	✓	✗										
P5	✓	✓										
P6	✓	✓										
P7	✓	✓										
P8	✓	✓										
P9	✓	✓										
P10	✓	✓										
P11	✓	✗										
P12	✓	✗										
Correct	12	8										

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓	✓	☐								
P2	✓	✗	✓	✗								
P3	✓	✓	✓	✗								
P4	✓	✗	✓	✗								
P5	✓	✓	✓	☐								
P6	✓	✓	✓	✗								
P7	✓	✓	✓	✗								
P8	✓	✓	✓	✗								
P9	✓	✓	✓	✗								
P10	✓	✓	☐	☐								
P11	✓	✗	✓	✗								
P12	✓	✗	✓	✓								
Correct	12	8	11	1								

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓	✓	□	✓	✓						
P2	✓	✗	✓	✗	✓	✗						
P3	✓	✓	✓	✗	✗	✗						
P4	✓	✗	✓	✗	✓	✓						
P5	✓	✓	✓	□	✗	✗						
P6	✓	✓	✓	✗	✓	✗						
P7	✓	✓	✓	✗	✓	✓						
P8	✓	✓	✓	✗	✓	✗						
P9	✓	✓	✓	✗	✓	✓						
P10	✓	✓	□	□	✓	✗						
P11	✓	✗	✓	✗	✗	✗						
P12	✓	✗	✓	✓	✓	✓						
Correct	12	8	11	1	9	5						

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓	✓	□	✓	✓	✓	✗				
P2	✓	✗	✓	✗	✓	✗	✓	✗				
P3	✓	✓	✓	✗	✗	✗	✗	✗				
P4	✓	✗	✓	✗	✓	✓	✓	✗				
P5	✓	✓	✓	□	✗	✗	✓	✗				
P6	✓	✓	✓	✗	✓	✗	✓	✓				
P7	✓	✓	✓	✗	✓	✓	✗	✗				
P8	✓	✓	✓	✗	✓	✗	✗	✗				
P9	✓	✓	✓	✗	✓	✓	✓	✗				
P10	✓	✓	□	□	✓	✗	✓	✗				
P11	✓	✗	✓	✗	✗	✗	✓	✓				
P12	✓	✗	✓	✓	✓	✓	✓	✗				
Correct	12	8	11	1	9	5	9	2				

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓	✓	□	✓	✓	✓	✗	✓	✓		
P2	✓	✗	✓	✗	✓	✗	✓	✗	✓	✗		
P3	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗		
P4	✓	✗	✓	✗	✓	✓	✓	✗	✓	✗		
P5	✓	✓	✓	□	✗	✗	✓	✗	✓	✗		
P6	✓	✓	✓	✗	✓	✗	✓	✓	✓	✗		
P7	✓	✓	✓	✗	✓	✓	✗	✗	✓	✓		
P8	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗		
P9	✓	✓	✓	✗	✓	✓	✓	✗	✓	✗		
P10	✓	✓	□	□	✓	✗	✓	✗	□	□		
P11	✓	✗	✓	✗	✗	✗	✓	✓	✓	✓		
P12	✓	✗	✓	✓	✓	✓	✓	✗	✓	✗		
Correct	12	8	11	1	9	5	9	2	9	3		

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓	✓	□	✓	✓	✓	✗	✓	✓	✗	✗
P2	✓	✗	✓	✗	✓	✗	✓	✗	✓	✗	✓	✓
P3	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	□
P4	✓	✗	✓	✗	✓	✓	✓	✗	✓	✗	✓	✓
P5	✓	✓	✓	□	✗	✗	✓	✗	✓	✗	✓	✓
P6	✓	✓	✓	✗	✓	✗	✓	✓	✓	✗	✓	✗
P7	✓	✓	✓	✗	✓	✓	✗	✗	✓	✓	✓	✓
P8	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗	✓	✗
P9	✓	✓	✓	✗	✓	✓	✓	✗	✓	✗	□	□
P10	✓	✓	□	□	✓	✗	✓	✗	□	□	✓	✗
P11	✓	✗	✓	✗	✗	✗	✓	✓	✓	✓	✓	✗
P12	✓	✗	✓	✓	✓	✓	✓	✗	✓	✗	✓	✗
Correct	12	8	11	1	9	5	9	2	9	3	10	4

Giovanni Apruzzese
giovannia@ru.is

Results

Table 3: Per-Participant Classification and Explanation Results Across All Cases

	Case 1		Case 2		Case 3		Case 4		Case 5		Case 6	
Participant	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**	**C**	**E**
P1	✓	✓	✓	□	✓	✓	✓	✗	✓	✓	✗	✗
P2	✓	✗	✓	✗	✓	✗	✓	✗	✓	✗	✓	✓
P3	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	□
P4	✓	✗	✓	✗	✓	✓	✓	✗	✓	✗	✓	✓
P5	✓	✓	✓	□	✗	✗	✓	✗	✓	✗	✓	✓
P6	✓	✓	✓	✗	✓	✗	✓	✓	✓	✗	✓	✗
P7	✓	✓	✓	✗	✓	✓	✗	✗	✓	✓	✓	✓
P8	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗	✓	✗
P9	✓	✓	✓	✗	✓	✓	✓	✗	✓	✗	□	□
P10	✓	✓	□	□	✓	✗	✓	✗	□	□	✓	✗
P11	✓	✗	✓	✗	✗	✗	✓	✓	✓	✓	✓	✗
P12	✓	✗	✓	✓	✓	✓	✓	✗	✓	✗	✓	✗
Correct	12	8	11	1	9	5	9	2	9	3	10	4

It doesn't seem our SOC operators can properly explain their decisions ☹

Giovanni Apruzzese
giovannia@ru.is

Takeaway

- SOC Operators necessitate tools/support that allows them to better explain their decisions

Giovanni Apruzzese
giovannia@ru.is

Discussion

Giovanni Apruzzese
giovannia@ru.is

Discussion

- How was the correctness determined?
  → Three coders, which ultimately agreed

Giovanni Apruzzese
giovannia@ru.is

Discussion

- How was the correctness determined?
  → Three coders, which ultimately agreed

- In some instances, the explanation simply reported the name of the alert itself (i.e., “Suspected Botnet”).
- “the tool was recognized as malicious by the EDR” → Obviously true but not informative
- Valid: “it is a PowerShell script from Edge, so it is suspicious”

Giovanni Apruzzese
giovannia@ru.is

Discussion

- How was the correctness determined?
  → Three coders, which ultimately agreed

- In some instances, the explanation simply reported the name of the alert itself (i.e., “Suspected Botnet”).
- “the tool was recognized as malicious by the EDR” → Obviously true but not informative
- Valid: “it is a PowerShell script from Edge, so it is suspicious”

- Why are the results so bad?
  → Our participants were *not* used to this task

Giovanni Apruzzese
giovannia@ru.is

Takeaway(s)

- SOC Operators necessitate tools/support that allows them to better explain their decisions

- (for practitioners) SOC Operators need to practice the task of explaining their decisions

Giovanni Apruzzese
giovannia@ru.is

Takeaway(s)

- SOC Operators necessitate tools/support that allows them to better explain their decisions
- (for practitioners) SOC Operators need to practice the task of explaining their decisions
- (for researchers) We release our resources: future research should build on our work

Giovanni Apruzzese
giovannia@ru.is

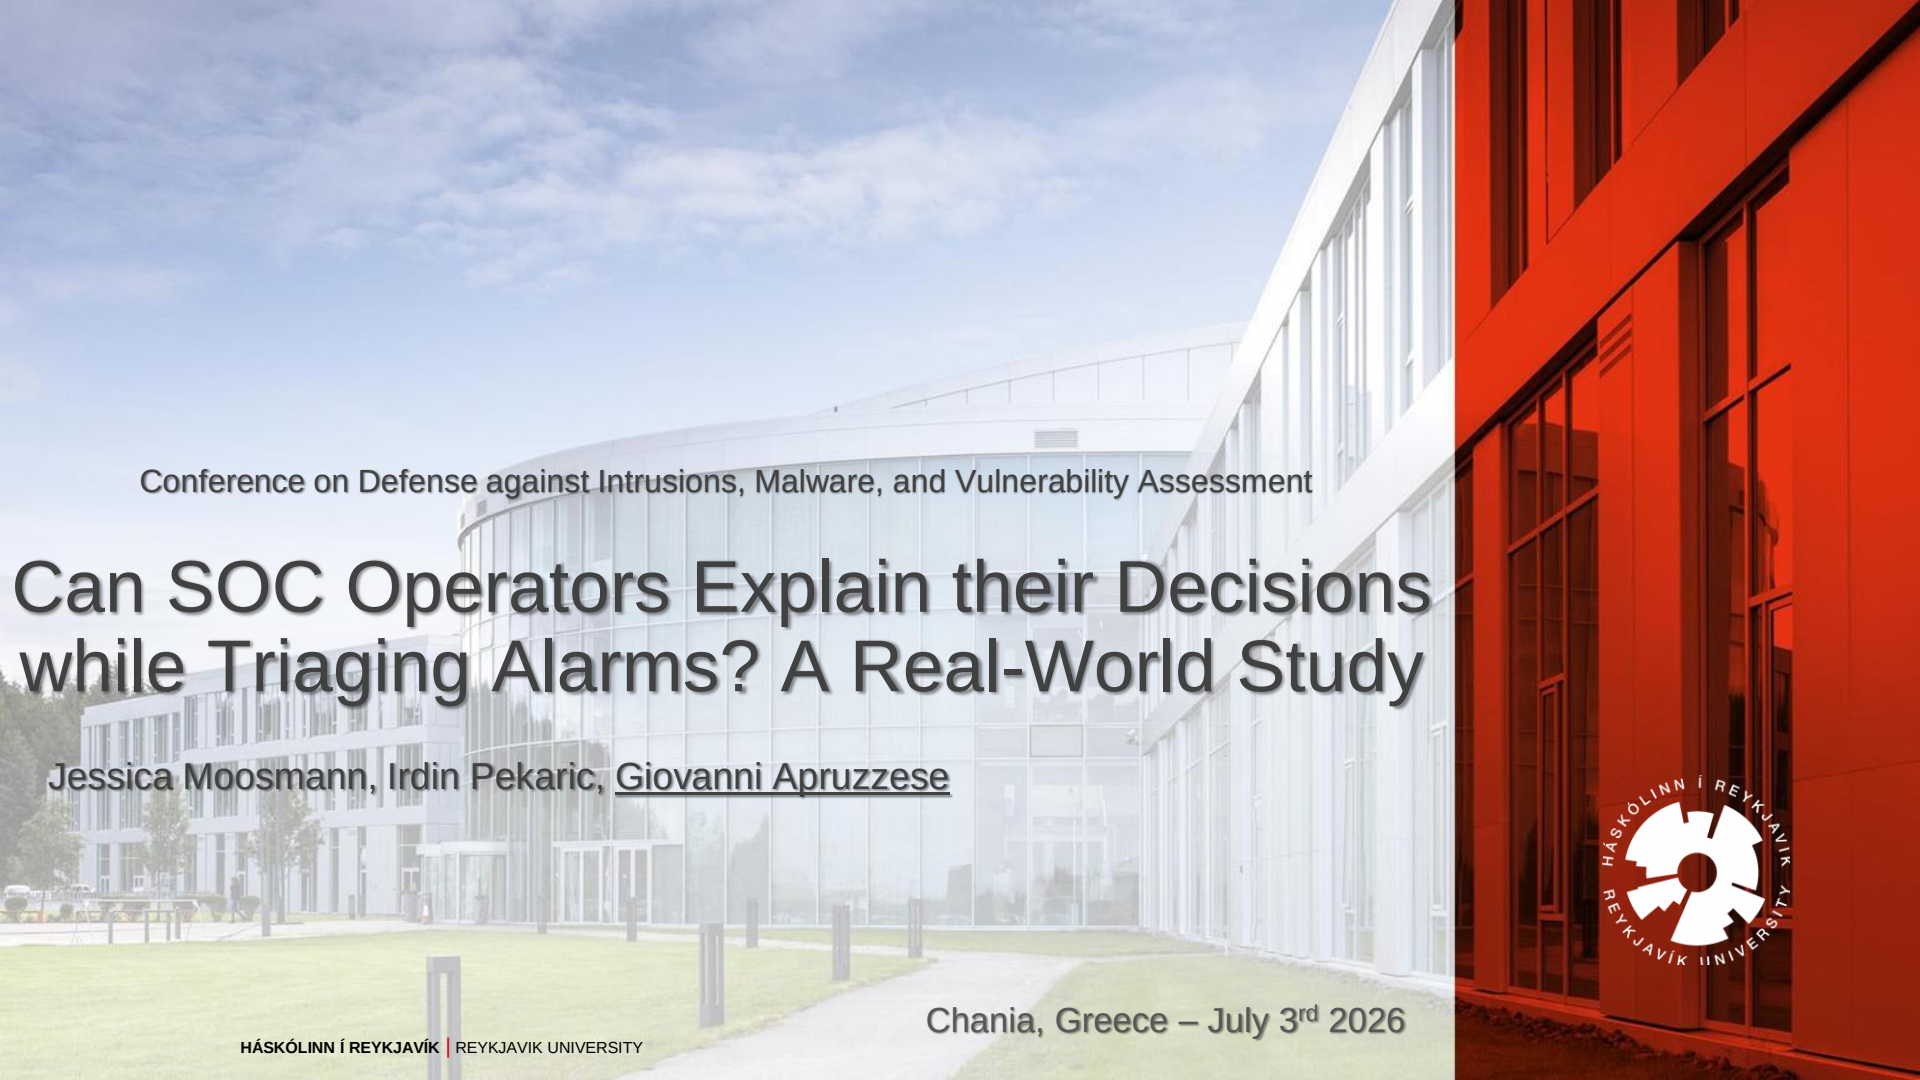

Conference on Defense against Intrusions, Malware, and Vulnerability Assessment

Can SOC Operators Explain their Decisions while Triaging Alarms? A Real-World Study

Jessica Moosmann, Irdin Pekaric, Giovanni Apruzzese

Chania, Greece – July 3rd 2026

Table 2: Distribution of classification-explanation result types across all participants.

Classification	Explanation	n (72)	%
Correct	Correct	22	31%
Correct	Incorrect / Vague	35	48%
Correct	Missing	3	4%
False / Missing	–	12	17%

Giovanni Apruzzese
giovannia@ru.is