SpacSec 2026
International Conference on Cyber Warfare, Security & Space Computing

Cyber threat intelligence data analysis system based on knowledge graphs

A.I. Lutskovich, A.M. Vulfin, A.D. Kirillova

Ufa University of Science and Technology

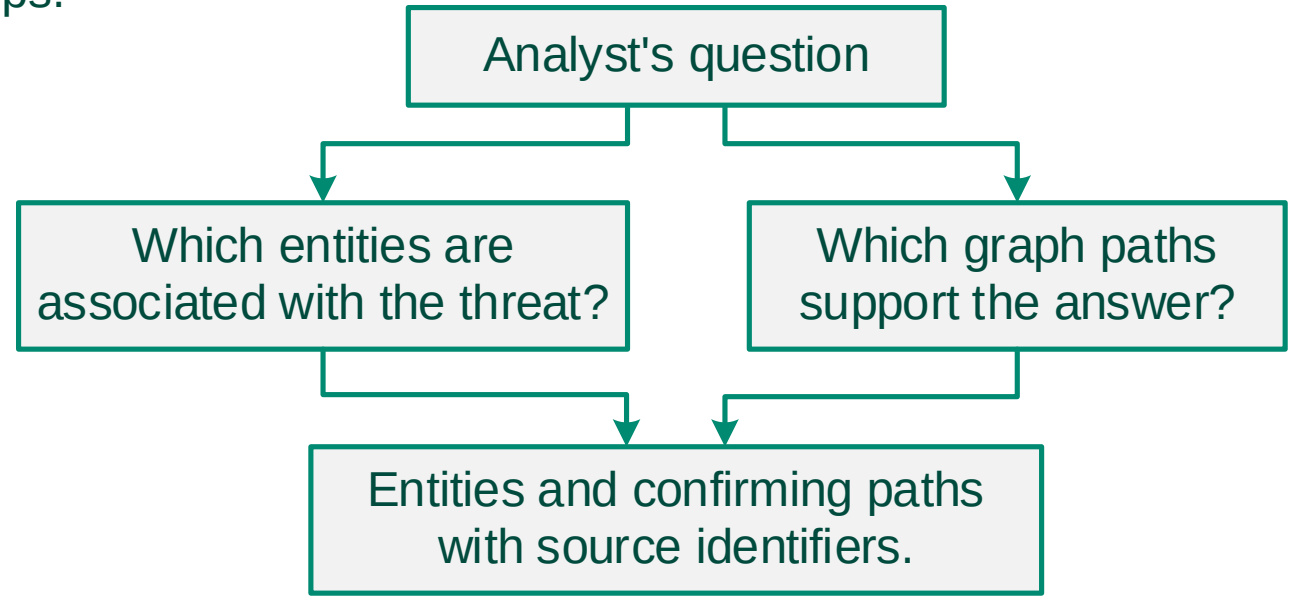
Speaker: A.I. Lutskovich

Saint Petersburg, 2026

The problem of cyber threat intelligence analysis

Descriptions of threats, intruders actions, and protective measures use different conceptual systems and identifiers.

Conventional intelligent retrieval systems have limited effectiveness in matching CTI reports, national threat taxonomies, and international knowledge bases across multiple steps.



Cyber threat intelligence: analyst tasks and the application of BDU-CTI-KG

CTI use cases



Challenges: data and analysis

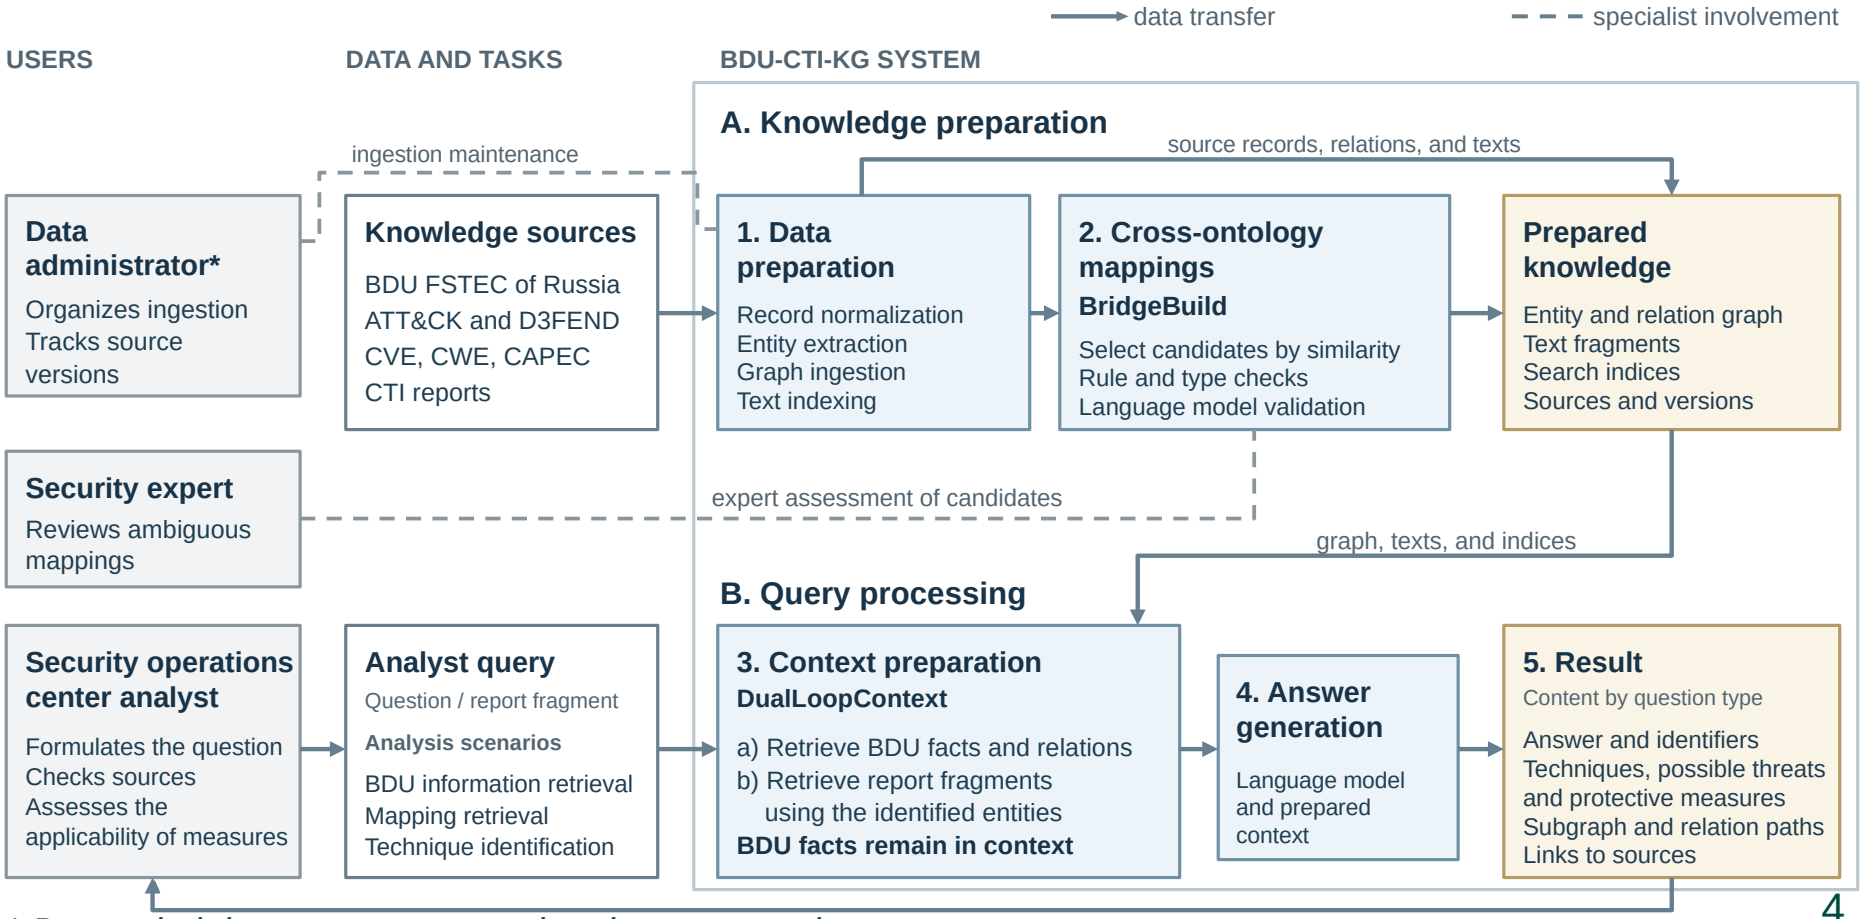


Using the system: from a threat report to a verifiable answer

Analyst's question	System processing	Verifiable result
Which techniques does the selected report describe?	Entity extraction; ATT&CK technique retrieval	Techniques and identifiers; supporting fragments
Which BDU threats are linked to the retrieved techniques?	Cross-ontology mapping retrieval	Possible BDU threats; relation paths and sources
Which measures are specified for these threats?	Retrieval of BDU measures and D3FEND defensive techniques	Protective measures and an explanation of their links to the retrieved threats

BDU – FSTEC of Russia Threat Data Bank; ATT&CK – database of attack techniques; D3FEND – database of defense techniques.

Use cases and data processing in BDU-CTI-KG



** Data administrator: a proposed maintenance role*

External data sources and knowledge bases

Source	Unit of representation	Role in the analysis
BDU FSTEC of Russia	Threat, method, protective measure	Normative threat model
MITRE ATT&CK	Tactic, technique, sub-technique	Intruder behavior
MITRE D3FEND	Defensive technique, digital artifact	Formal model of defensive actions
CVE / CWE / CAPEC	Vulnerability / weakness / attack pattern	Specified relations and realization conditions
CTI reports	Document, fragment, mention	Source texts and context

Ontology

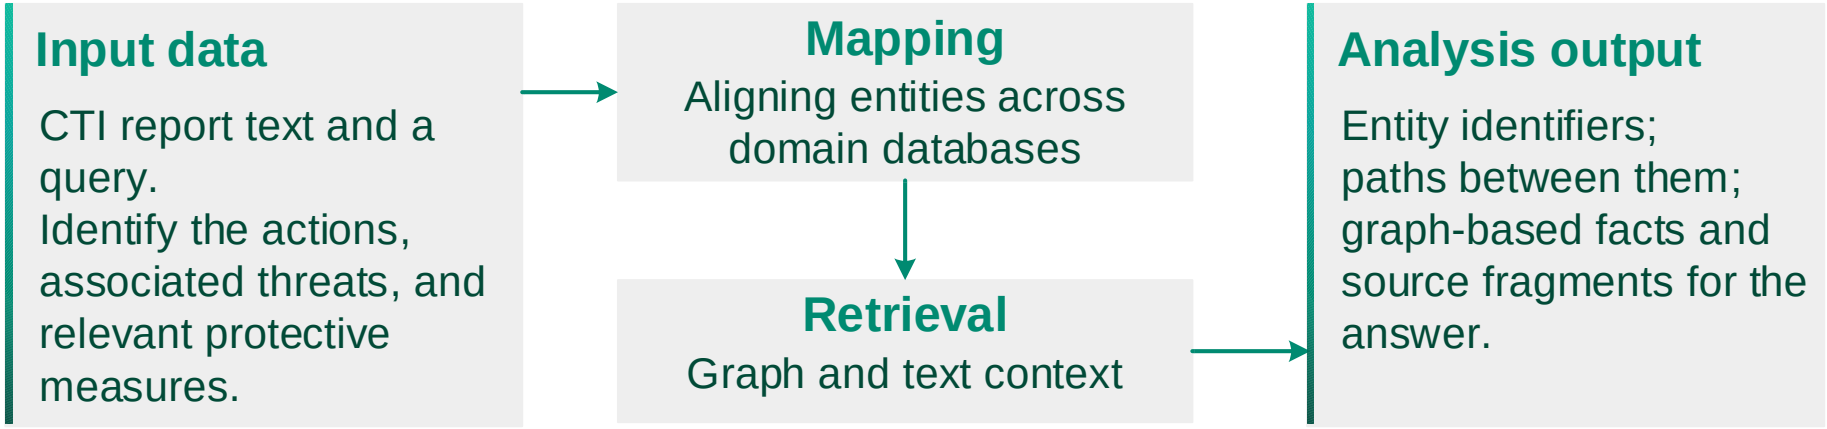
An explicit model of domain classes, relations, and constraints.

Knowledge graph

A graph representation of knowledge about entities, their types, relations, and properties.

RAG – Retrieval-Augmented Generation

Research objective and tasks

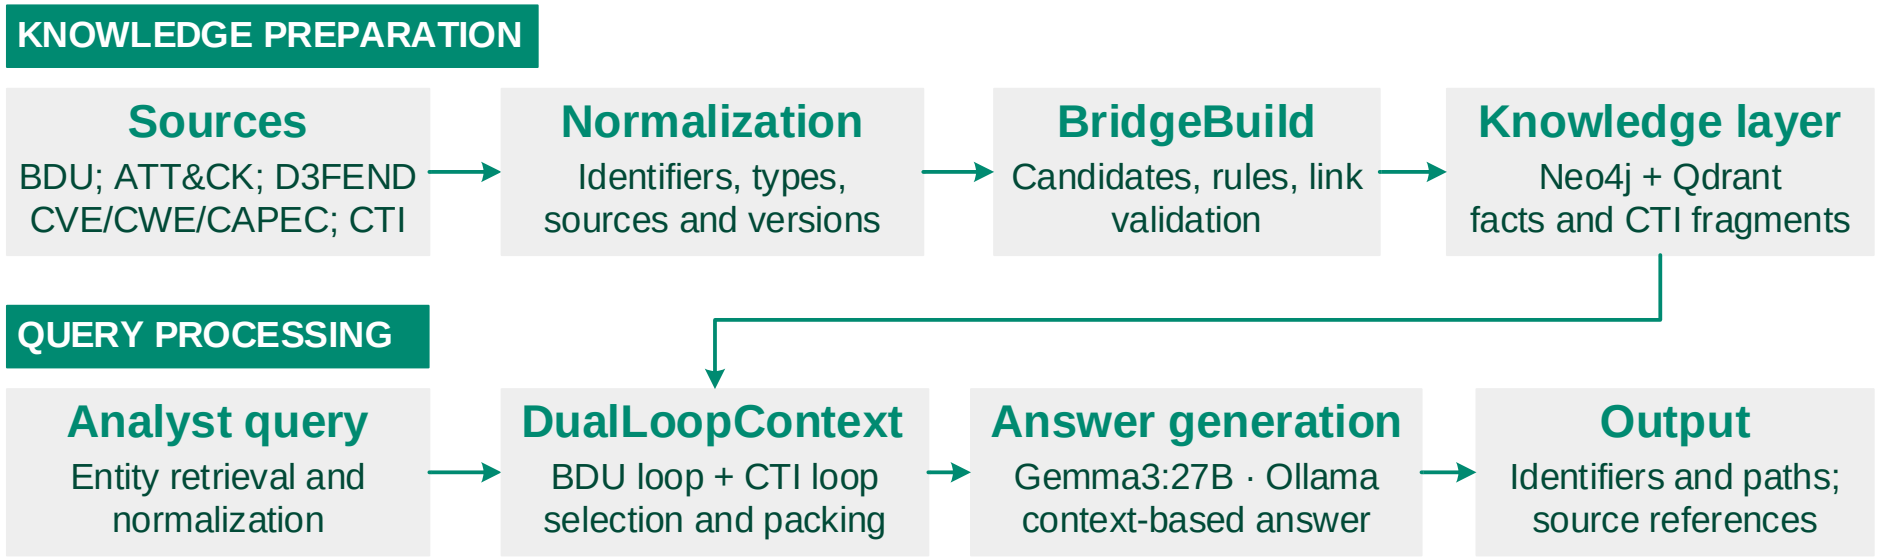


Research objective

To develop and experimentally evaluate a cyber threat intelligence analysis system based on an integrated knowledge graph, cross-ontology relations, and dual-loop context retrieval.

Unit of analysis: a connected set of entities, relations, and source text fragments.

Structural and functional organization of the system



Relation preparation is separate from query processing; sources are preserved at both stages.

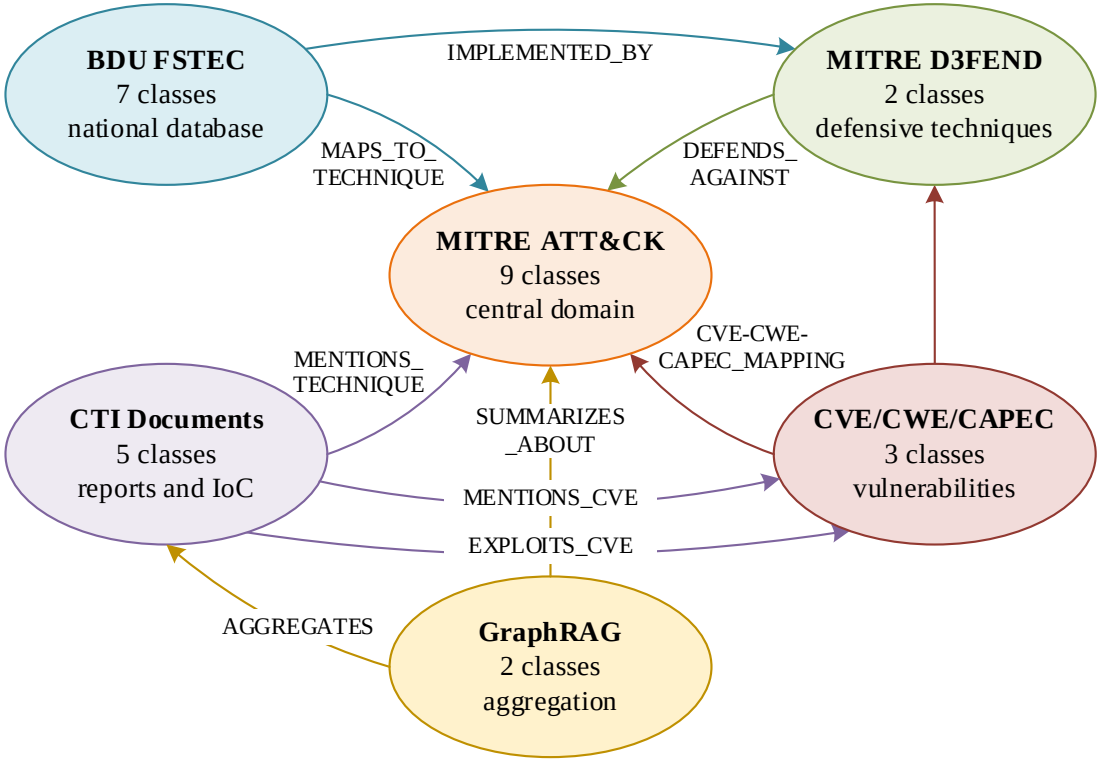
Conceptual model of the BDU-CTI-KG graph

$G = (V, E, \lambda_V, \lambda_E, A_V, A_E, \Pi)$

V, E – nodes and edges;
λ_V, λ_E – typing functions;
A_V, A_E – attributes;
Π – data provenance.

$b = (u, v, r, s, \varepsilon)$

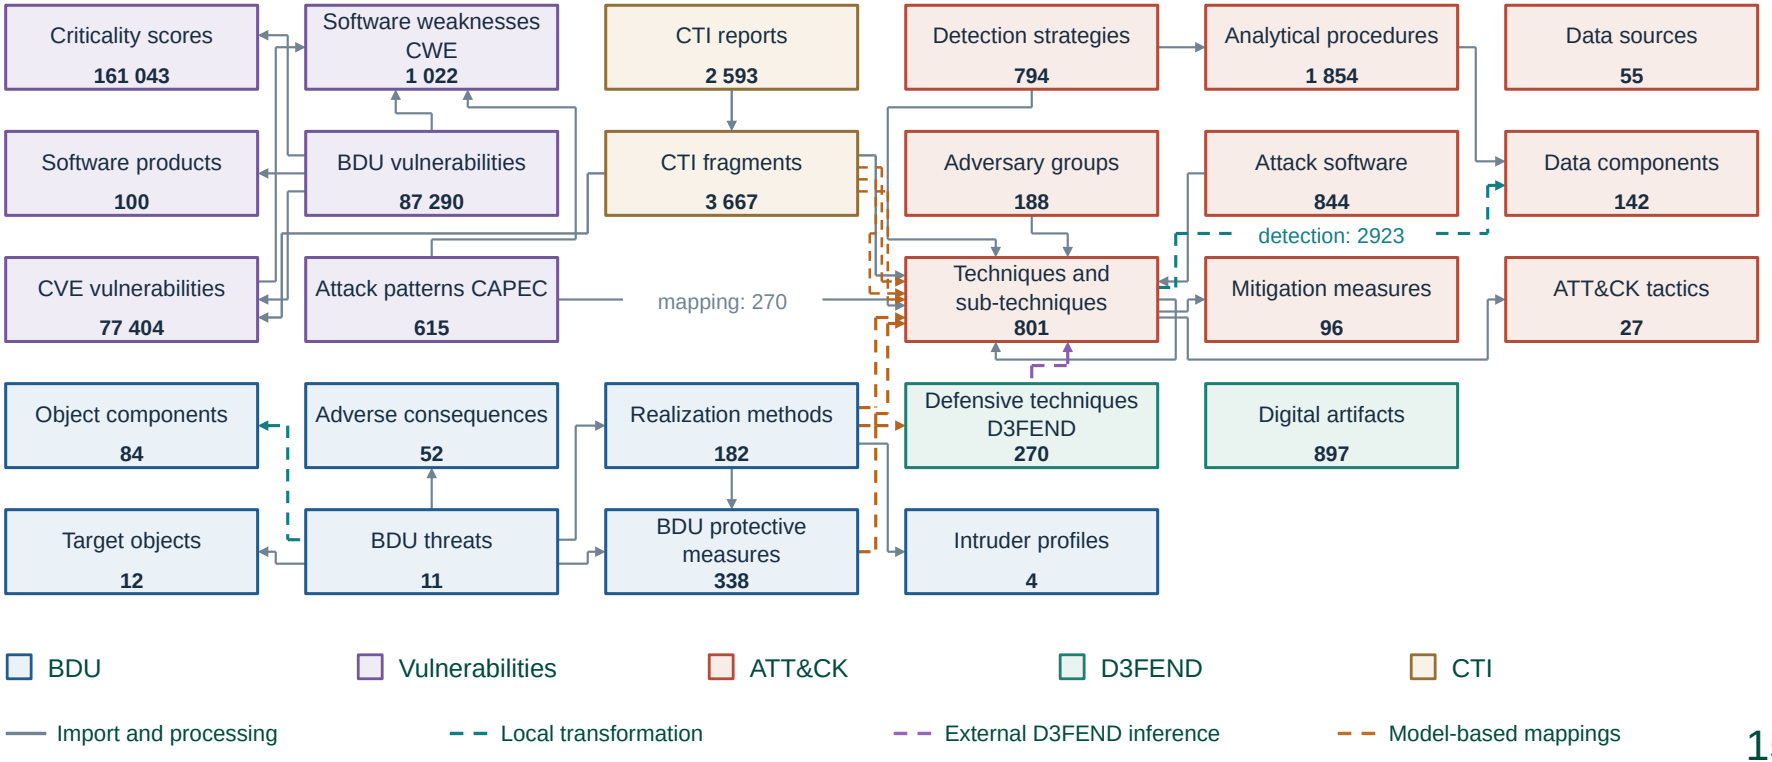
u, v – entities;
r – relation type;
s – mapping score;
ε – sources and checks.



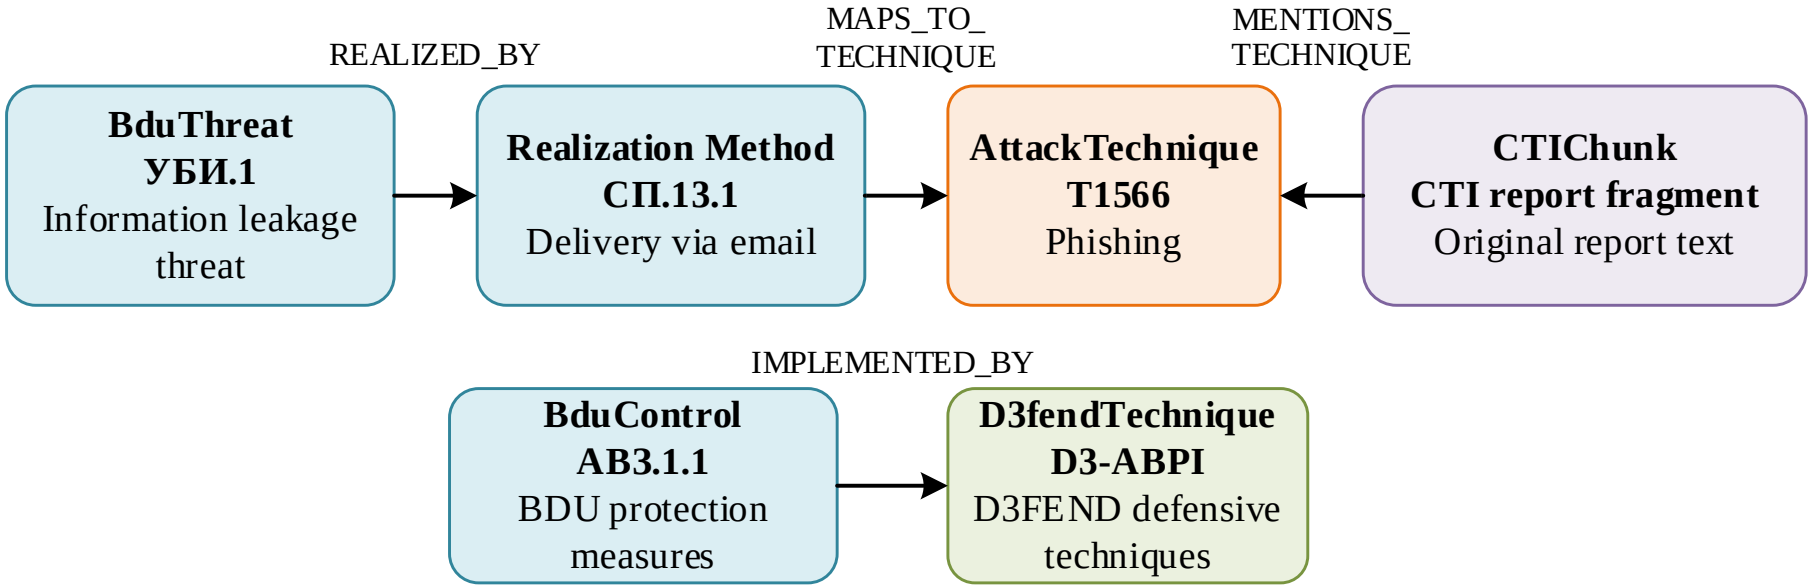
BDU-KG is the normative subgraph. The schema defines conceptual domains and relation types.

Metaschema of a stored knowledge graph

340 385 nodes and 516 424 relations; 26 groups by first label, 30 relation types, and 33 combinations.



Examples of cross-ontology relations



Two fragments: УБИ.1 – СП.13.1 – T1566 (V2 layer) and АВ3.1.1 – D3-ABPI.
The diagram shows logical relation types.

Roles of knowledge extraction and retrieval

Concept	Meaning
CTI / TTP	Cyber Threat Intelligence / Tactics, Techniques and Procedures
RAG / GraphRAG	RAG: answer generation with context retrieval; GraphRAG: context from graph entities and relations
Graph / BDU-CTI-KG	Labeled property graph; typed nodes and relations; attributes and data provenance information

Experiment 2 series	Primary context	Domain role
Experiment 2 – C0: Text-RAG	CTI corpus fragments	Relevant text retrieval
Experiment 2 – B1: graph	Neo4j nodes and relations	Structural queries
Experiment 2 – C1-C4	Facts; fragments; M1/M2 relations	Connected multi-hop analysis

GraphRAG denotes a class of approaches; specific results depend on the knowledge structure and retrieval method.

BridgeBuild: constructing cross-ontology mappings

1. Candidates
Vector similarity;
lexical overlap.
At most k candidates

→

2. Constraints
Types and direction;
structural constraints;
prohibitory rules.

→

3. LLM validation
Descriptions of the pair, facts. Mapping label; assessment and rationale.

→

4. Decision
Aggregation or classifier.
Threshold and preservation of ε.

$$C_k(u) = \mathrm{TopK}_{v \in V_M}[\mathrm{sim}(e(u), e(v))]$$

$e(\cdot)$ – description embedding;
$C_k(u)$ – candidate set.

$$N_{\mathrm{LLM}} \leq k|V_B|$$

Number of pairs evaluated

Input
Graphs GB and GM; domain rules R; candidate count k; threshold.

Output
Set of mappings B with relation type, score, and validation basis ε.

Source type	Computed relation	Target type
Threat realization method	MAPS_TO_TECHNIQUE	ATT&CK technique
BDU protection measure	IMPLEMENTED_BY	D3FEND defensive technique

A critical rule violation excludes a pair before a language model is invoked.

DualLoopContext: dual-loop context retrieval

1. NORMATIVE LOOP

Seed nodes S_0
Retrieving BDU entities for query q

→

Subgraph G_{norm}
Allowed relations and depth d_1

→

Context C_{norm}
Entity records and typed facts

2. CTI MATERIAL LOOP

Expanded query q_Q
Original question and subgraph entities

→

Text fragments T
CTI retrieval for the expanded query

→

Context C_{evid}
Selection, removal of duplicates, linking to entities

Selection rules

Entity coverage;
query relevance;
source diversity;
duplicate removal.

Parameters

k_1, k_2 – retrieval limits;
d_1 – expansion depth;
b_1, b_2 – layer budgets;
B_{tok} – total budget.

$$C(q) = \mathrm{Pack}(C_{norm}(q), C_{evid}(q), B_{tok}) \qquad \mathrm{tokens}\big(C(q)\big) \le B_{tok}$$

Selection: entity coverage, query similarity, source diversity, and redundancy.

Normative facts are included regardless of whether a supporting CTI fragment is available.

Three system evaluation scenarios

Scenario	Input and question	Expected result	Evaluation focus
BDU-QA	BDU threat: which protective measures are associated with it?	BDU protective measures and normative paths to them	Initial threat retrieval; type and completeness of related entities
Crosswalk-QA	BDU threat or ATT&CK technique: entities in the other database; paths to them	ATT&CK techniques or BDU threats; paths through realization methods	Cross-ontology relations; traversal direction; path validity
TTP-QA	CTI report fragment: which techniques are associated with it?	ATT&CK techniques and paths from the initial CTI fragment	Retrieval over a fixed graph; entity and path coverage

Experiment 1: BDU – 11; Crosswalk – 33 / 21; TTP – 216 / 208.
Experiment 2: a separate series of 250 questions per scenario (750 in total).

Experiment 1: normative, cross-ontology, and document retrieval.

Answer and retrieval evaluation metrics

Metric and evaluation target	Formula for question i
EM: exact answer match	$EM_i = \mathbf{1}[A_i = A_i^*]$
F_1: answer recall and precision	$F_{1,i} = \frac{2\lvert A_i \cap A_i^* \rvert}{\lvert A_i \rvert + \lvert A_i^* \rvert}$
ER@K: entity recall	$ER_i@K = \frac{\lvert E_i^{(K)} \cap E_i^* \rvert}{\lvert E_i^* \rvert}$
PR@K: path recall	$PR_i@K = \frac{\lvert \mathrm{Match}(P_i^{(K)}, P_i^*) \rvert}{\lvert P_i^* \rvert}$
Pset@8: window precision	$P_{\mathrm{set},i}@8 = \frac{\lvert U_{8,i} \cap A_i^* \rvert}{\lvert U_{8,i} \rvert}$

*A – answer identifiers; E – entities; P – paths; * – reference.*
U_8 – the canonical set of the first eight retrieved IDs.

Dataset results are averaged over questions.
Pset@8 – historical AP@8 field; Match – matched reference paths.

Scenario A: BDU-QA – questions about the BDU structure

Input and expected answer

Input: a threat identifier or description.
In this example, the answer consists of BDU protection measures.
Basis: normative relations.

Operations evaluated

1. Find the initial threat.
2. Select a permitted relation type.
3. Retrieve related entities and paths.
4. Generate an answer with valid Ids.

Example element	Archived record	Criterion
Question bdu_qa_0001	“How is information protected against unauthorized transmission outside the organization?”	Answer target: BduControl
Reference	20 measures; 20 paths from УБИ.1 through MITIGATED_BY_CONTROL	EM and F_1 for the answer; ER@50 and PR@10 for context

Scenario errors: incorrect threat, entities of another type, omitted measures, or excessive retrieval.

BDU-QA: reference and archived model answer

Question: “Which protection measures specified in the FSTEC of Russia Threat Data Bank correspond to the information security threat “Information leakage threat” and are intended to neutralize or mitigate it?”

Reference

20 protection measures .

Eight matching measures:
ИНЦ.4.3; ИПО.1.1; ОЦЛ.1.5; ИПО.3.1; ИНЦ.5.8; УПД.2.2; УПД.1.5; УПД.1.3.

The matching reference subset is shown; the full set contains 12 additional measures.

Gemma3:27B-based answer

The following protection measures are applied to neutralize or mitigate the “Information leakage threat” (BDU-UBI-УБИ.1)

[BDU-MEASURE-ИНЦ.4.3], [BDU-MEASURE-ИПО.1.1], [BDU-MEASURE-ОЦЛ.1.5], [BDU-MEASURE-ИПО.3.1], [BDU-MEASURE-ИНЦ.5.8], [BDU-MEASURE-УПД.2.2], [BDU-MEASURE-УПД.1.5], [BDU-MEASURE-УПД.1.3].

Answer citations : P = 8/8 = 1; R = 8/20 = 0.40; F_1 = 4/7 ≈ 0.571; EM = 0.

Scenario B: Crosswalk-QA – cross-domain analysis

Forward direction

Input: a BDU threat description or identifier.
Task: connect the domain databases and explain the transition.
Chain evaluated: threat → realization method → ATT&CK technique

Reverse direction

Input: an ATT&CK technique.
Task: identify related BDU threats.
Traversal proceeds from the technique to the method and threat; the directions of stored relations remain unchanged.

Evaluation	Reference and result	Typical error
Mapping	Entities from different databases; typed paths	Similar wording with different meanings
Path	Forward / reverse traversal; mapping layer version	Traversal without an intermediate realization method
No answer	Empty reference when no valid path exists	Textual similarity without a normative path

A correct cross-domain answer requires matching entities and valid connecting paths.

Crosswalk-QA: reference answer and generation evaluation

Question crosswalk_v2_0001: “Given the following description of an information security threat (BDU FSTEC): “The threat involves the possibility of unlawfully obtaining or transmitting information (confidential, configuration, authentication, etc.). [...]” Which techniques in the MITRE ATT&CK matrix correspond to this threat?”

Archived reference: 7 techniques

T1078.001; T1091; T1190;
T1204.002; T1566; T1566.001;
T1566.002.

8 paths; one example:
УБИ.1 → СП.13.1 → T1566.
REALIZED_BY;
MAPS_TO_TECHNIQUE_V2.

Historical generation evaluation

Evaluator model verdict: the graph-based answer includes all reference techniques and additional IDs.

faithfulness = 1,0;
citation_correctness = 1,0;
hallucinated_id_fraction = 0,75.

Reference agreement and contextual support are evaluated separately.

Crosswalk-QA: verbatim excerpt from a model answer

Question crosswalk_qa_0013: “Which MITRE ATT&CK techniques are associated with information security threat УБИ.11 “Threat of unauthorized mass collection of information”, considering its realization methods?”

Archived answer – excerpt

“The following MITRE ATT&CK techniques are associated with information security threat УБИ.11 :
[attack:enterprise:T1027], [attack:enterprise:T1553], [attack:enterprise:T1204], [attack:enterprise:T1112], [attack:enterprise:T1566], [...] [attack:enterprise:T1036].”

Result

One LLM call.
recall = 1.000
precision = 0.049
F_1 = 0.094
Latency : 8011.5 ms.

The context includes paths СП.22.24 → T1027 and СП.22.24 → T1553.

An answer can cover the reference while including a large number of additional identifiers.

Crosswalk-reverse: JavaScript and BDU threats

“Which information security threats (BDU FSTEC) can be realized using MITRE ATT&CK technique T1059.007 “JavaScript"? List the threats.”



Parameterized Cypher template

```
MATCH (t:AttackTechnique)
    <-[:MAPS_TO_TECHNIQUE_V2]-
    (m:RealizationMethod)
    <-[:REALIZED_BY]-(u:BduThreat)
WHERE t.id = $technique_id
RETURN DISTINCT u.id, m.id, t.id
ORDER BY u.id
```

technique_id = attack:enterprise:T1059.007

Reference: 8 threats and 8 paths.

Path union: 10 unique nodes and 9 relations.

Traversal direction and stored relation direction are specified separately.

Scenario C: TTP-QA – analysis of a CTI fragment

Input

Input: the text and exact identifier of a CTI report fragment.
Task: retrieve ATT&CK techniques associated with the fragment.
TTP – tactics, techniques, and procedures.

Reference and output

Reference: technique IDs and paths CTIChunk → AttackTechnique.
Evaluated result: ATT&CK techniques and paths from the fragment.

Evaluation	Practical meaning	Limitation
Semantics	The technique corresponds to the described action	A product or operating system name does not specify the action
Alignment	Paths and quotes refer to the same fragment_id	An adjacent fragment may describe a different action
Robustness	Removal of technique codes and names from the question text	The reference may remain dependent on the graph

Fixed graph: the evaluated V3 relations are available for retrieval.

TTP-QA: question, complete reference, and result

Question fragment – excerpt

“[...] Identify which MITRE ATT&CK techniques are demonstrated by the described behavior.
“Example 3
In the Indonesian incident, the attackers checked for the presence of a Kaspersky Endpoint Security for Windows agent and its version: [...]””

Full reference for this record

T1012; T1016; T1018;
T1069.002; T1082; T1482.

All 6 codes are in the full question.
Only answer metrics were retained.

From the archived result	Value	Interpretation
Local Search context	23 IDs; 4 matches	Context recall: 4/6 ≈ 0.667
Answer citations	5 matches	Citation recall: 5/6 ≈ 0.833

Context recall and citation recall characterize different stages of question processing.

TTP-QA: verbatim response and fragment alignment error

Original quotation: FIN13, chunk_0001

“FIN13 has remained focused on more traditional financially motivated cybercrime and has targeted both Linux and Windows systems throughout their operations.”

Supplied context: chunk_0000

A different fragment of the same report.
8 MENTIONS_TECHNIQUE_V3 paths;
8 selected IDs; 9 nodes.
Relation status: candidate_unaudited.

Gemma3:27B response, verbatim

The described activities use the following techniques : [attack:enterprise:T1047], [attack:enterprise:T1071.001], [attack:enterprise:T1190], [attack:enterprise:T1003.001], [attack:enterprise:T1053.005], [attack:enterprise:T1041], [attack:enterprise:T1059.001], [attack:enterprise:T1027].

Stored reference: 10 IDs

Only T1047 matched.
TP = 1; FP = 7; FN = 9.
P = 1/8 = 0.125; R = 1/10 = 0.10.
F_1 = 2/(2+7+9) ≈ 0.111.
EM = 0.

All 8 references occur in the supplied context, but their alignment with the original quotation is unconfirmed.

Experiment 2: relation and answer quality

EM: 250 questions per scenario

Configuration	BDU-QA	Crosswalk-QA	TTP-QA
C0 Vector Text-RAG	0.48	0.21	0.26
C1 BDU graph	**0.77**	**0.34**	**0.38**
C2 Expert mappings M1	**0.78**	**0.63**	**0.49**
C3 Automatic mappings M2	0.76	0.56	0.46
C4 C3 + community summaries	0.78	0.65	0.55

BDU-QA: C0 → C1

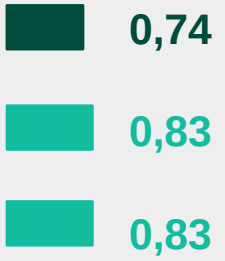
0.77 − 0.48 = +0.29

Crosswalk-QA: C1 → C2

0.63 − 0.34 = +0.29

Pair classification

Test F_1: 394 pairs



Mixed set of mapping types

The differences characterize the compared configurations, including the facts and relations available to them.

Experiment 1: retrieval on a fixed graph

Entity recall ER@50

Dataset	n	C0: BM25	C2: graph
TTP-QA V3	216	0.3528	**0.8753**
Masked technique codes and names	208	0.3273	**0.8740**

ER@50 difference and 95% CI



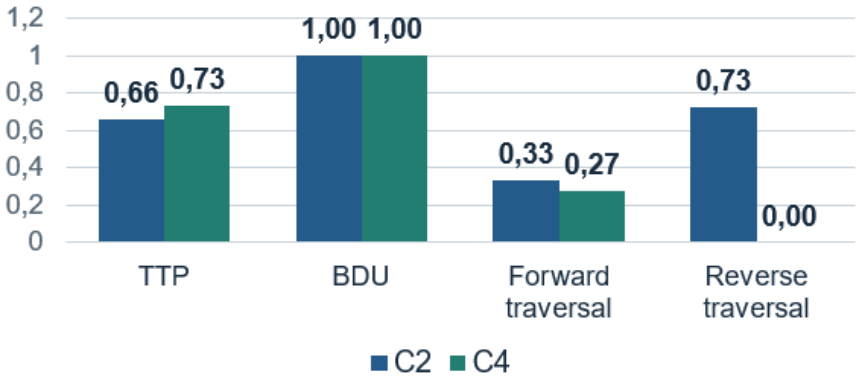
Conditions: retrieval without generation; vector search for seed nodes disabled.
The reference MENTIONS_TECHNIQUE_V3 relations are accessible to retrieval.
95% confidence interval: 10 000 resamples; the cluster is a report fragment.

C2 has higher ER@50 on both sets. Transfer to new reports requires separate evaluation.

Experiment 1: C2 and C4 across scenarios

C2: broad graph retrieval. C4: merge up to 50 IDs, rank, and select a window of up to 8.

Precision of the set of the first eight IDs



$$P_{set}@8 = \frac{|U_8 \cap A^*|}{|U_8|}$$

U_8 – canonicalized top 8 IDs; exact duplicates removed.
A – complete answer; A* – reference answer.

Scenario, n	Pset@8 C2 / C4	F_1 full list C2 / C4
TTP-QA, 216	0.6603 0.7280	0.1184 0.2635
BDU-QA, 11	1.0000 1.0000	0.1183 0.5714
БДУ to ATT&CK, 33	0.3295 0.2727	0.0960 0.1132
ATT&CK to БДУ, 21	0.7262 0.0000	0.7832 0.3130

$$F_1 = \frac{2|A \cap A^*|}{|A| + |A^*|}$$

Means across questions; full-list F_1: C3 = C4.
ATT&CK to BDU: C4 window has no reference IDs.

Main results and scope of applicability

Knowledge model

BDU normative records, ATT&CK techniques, D3FEND defensive techniques, and source texts are integrated. Relation types and data provenance are preserved.
Unified graph model.

Algorithms

BridgeBuild: construction of cross-ontology mappings.

DualLoopContext: retrieval of the normative core and textual context.

Experiment

Retrieval was compared on a fixed graph.
Retrieval recall and scenario-dependent window precision were evaluated.

Limitations and further evaluation

New CTI reports and analyst workflows require a separate experiment. Masking controls textual cues with the V3 layer available.

The system supports analysts in multistep alignment of cyber threat intelligence data.

Thank you for your attention!