The single-person (and several dozen AI agent) CTI team

Brendon Hawkins
IndependINT

What we'll be covering today

1. We'll examine what AI agents are and how they can be used to augment cyber threat intelligence capabilities.
2. I'll run you through some practical examples of using AI workflows from some of my own work.
3. We'll bed down some of the principles of what works when building AI tools for threat intelligence.
4. I'll discuss the potential applications of combining intelligence tradecraft with AI to build knowledge about the world.

I'm a senior intelligence professional with over 20 years of experience across Defence, NIC, policing, and corporate intelligence functions. My intelligence-related interests include intelligence training, prototyping tools, and experimenting with new processes.

The state of CTI in Australia

Most organisations in Australia don’t have a dedicated Cyber Threat Intelligence team. When they do, it’s often a single analyst, typically juggling multiple roles.

Some organisations outsource CTI, which can work, but might miss out on internal context.

FTE growth is a challenge, particularly when there are other security needs. Australia’s CTI workforce is also small and specialised, meaning hiring expert staff is difficult and expensive.

The question for us today: how can we use AI to augment CTI capabilities in Australian organisations.

I’ve been looking at this in my spare time for the past few years and have built some use cases which I’d like to share with you all.

There will be three main functions of a CTI analyst as AI matures:

Specialist Analysts

Intelligence Communicator

Intelligence Manager

How do we build the skills pipeline for junior analysts?

The tech team

I use a range of commercial and open-source tools when building my experiments. These include Telegram, Gemini, Chat GPT, Python, PostgreSQL, Scikit Learn, Claude, Cursor, and Spacy.

As for the human member:

- ✓ I am very experienced with intelligence process
- ✓ I've worked across the full intelligence cycle
- ✓ I've worked across a range of targets
- ✓ I can code (Python), build databases, use APIs
- ✓ I am very comfortable with data analysis
- ✓ I have someone to build infrastructure for me

- × I am not a software engineer
- × I'm an intelligence expert, not an AI expert
- × I am not a data scientist
- × Don't ask me to design a front end…

These illustrations are creative interpretations for educational purposes, intended to visualize technologies used in cyber threat intelligence workflows. They are not endorsed by, affiliated with, or created by the companies or products depicted.

What are AI Agents?

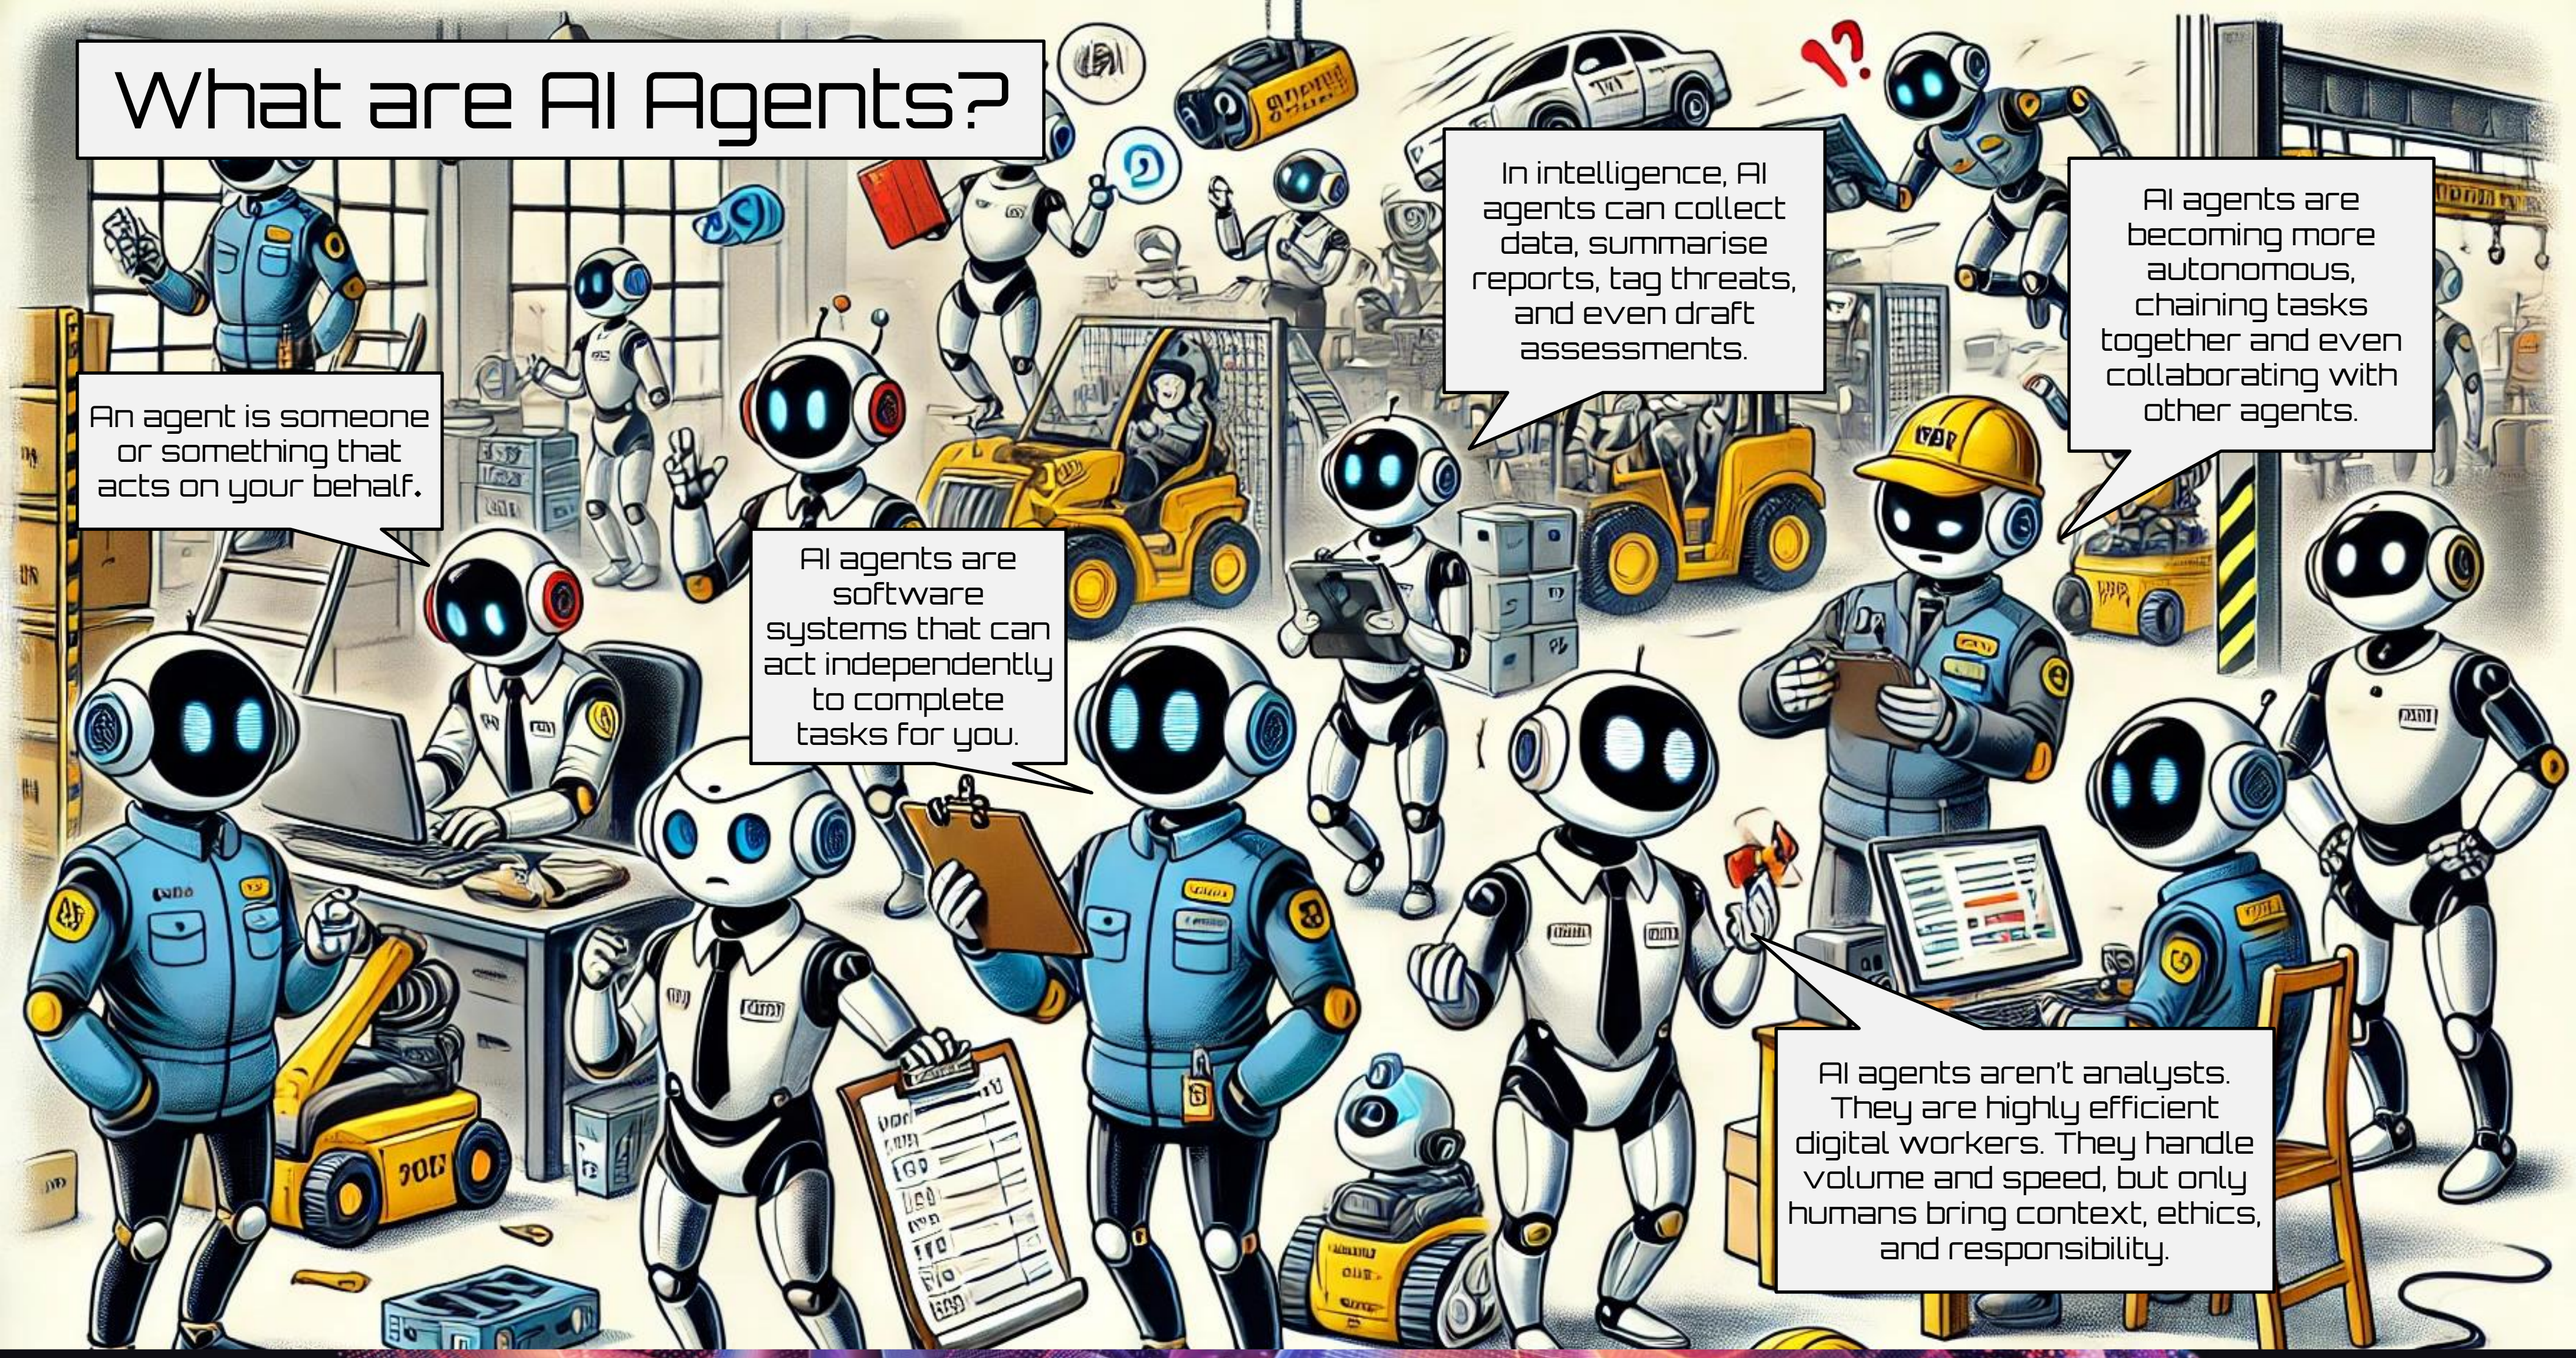

Working with the limitations

LLMs are fantastic for summarising, translating, triaging information, and speed.

LLMs are less effective for analysis*, long reports, referencing, and remembering.

Where I have had most success is in keeping AI focused on tasks by using robust **intelligence requirements.**

Then you need to understand your own **intelligence processes** and break them down into manageable chunks.

LLMs, like human analysts, make mistakes. But **good process** can minimise these.

AI is faster if you can tell it **what you need**.

The Intelligence Cycle

Intelligence is an ancient profession. But it was only systemised during the 20th century. In the western military context, it was structured using the intelligence cycle.

The intelligence cycle is a simplified framework for the activity of intelligence. Each part of the cycle traditionally uses specialised professionals to perform their part.

It's the same with AI agents in intelligence – they should be specialised to perform their role in the intelligence cycle.

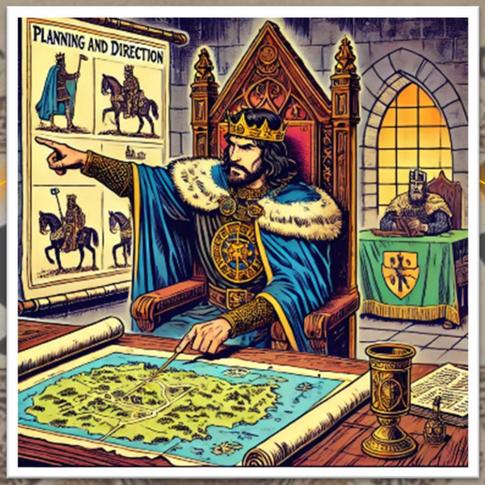

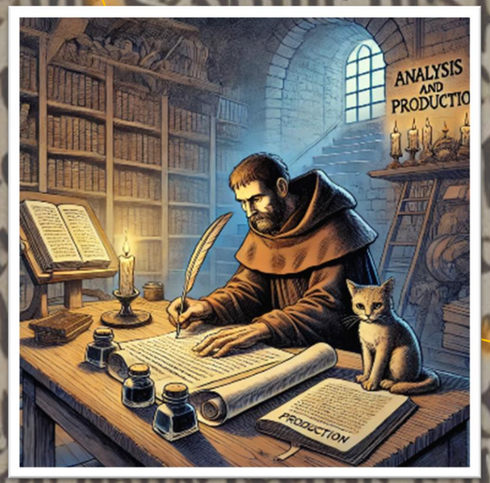

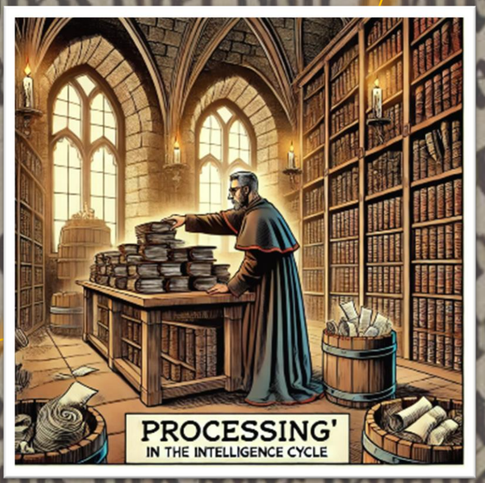

Refining intelligence requirements

Intelligence teams service a range of business areas.

They need to engage with stakeholders…

…and bring the results together…

…to generate perfect information needs!

A challenge for any intelligence function is that they are servicing a range of stakeholders. AI can be used to help refine requirements and make sure that the right intelligence is reaching the parts of the business that need it.

The QR code below links to a custom GPT which interviews a cyber security stakeholder from the company TelcoTechCom to determine how their needs align to intelligence requirements.

Scan it, open the web page, and have a go at using it after the presentation. It works best with voice mode.

Collection: Survey tool

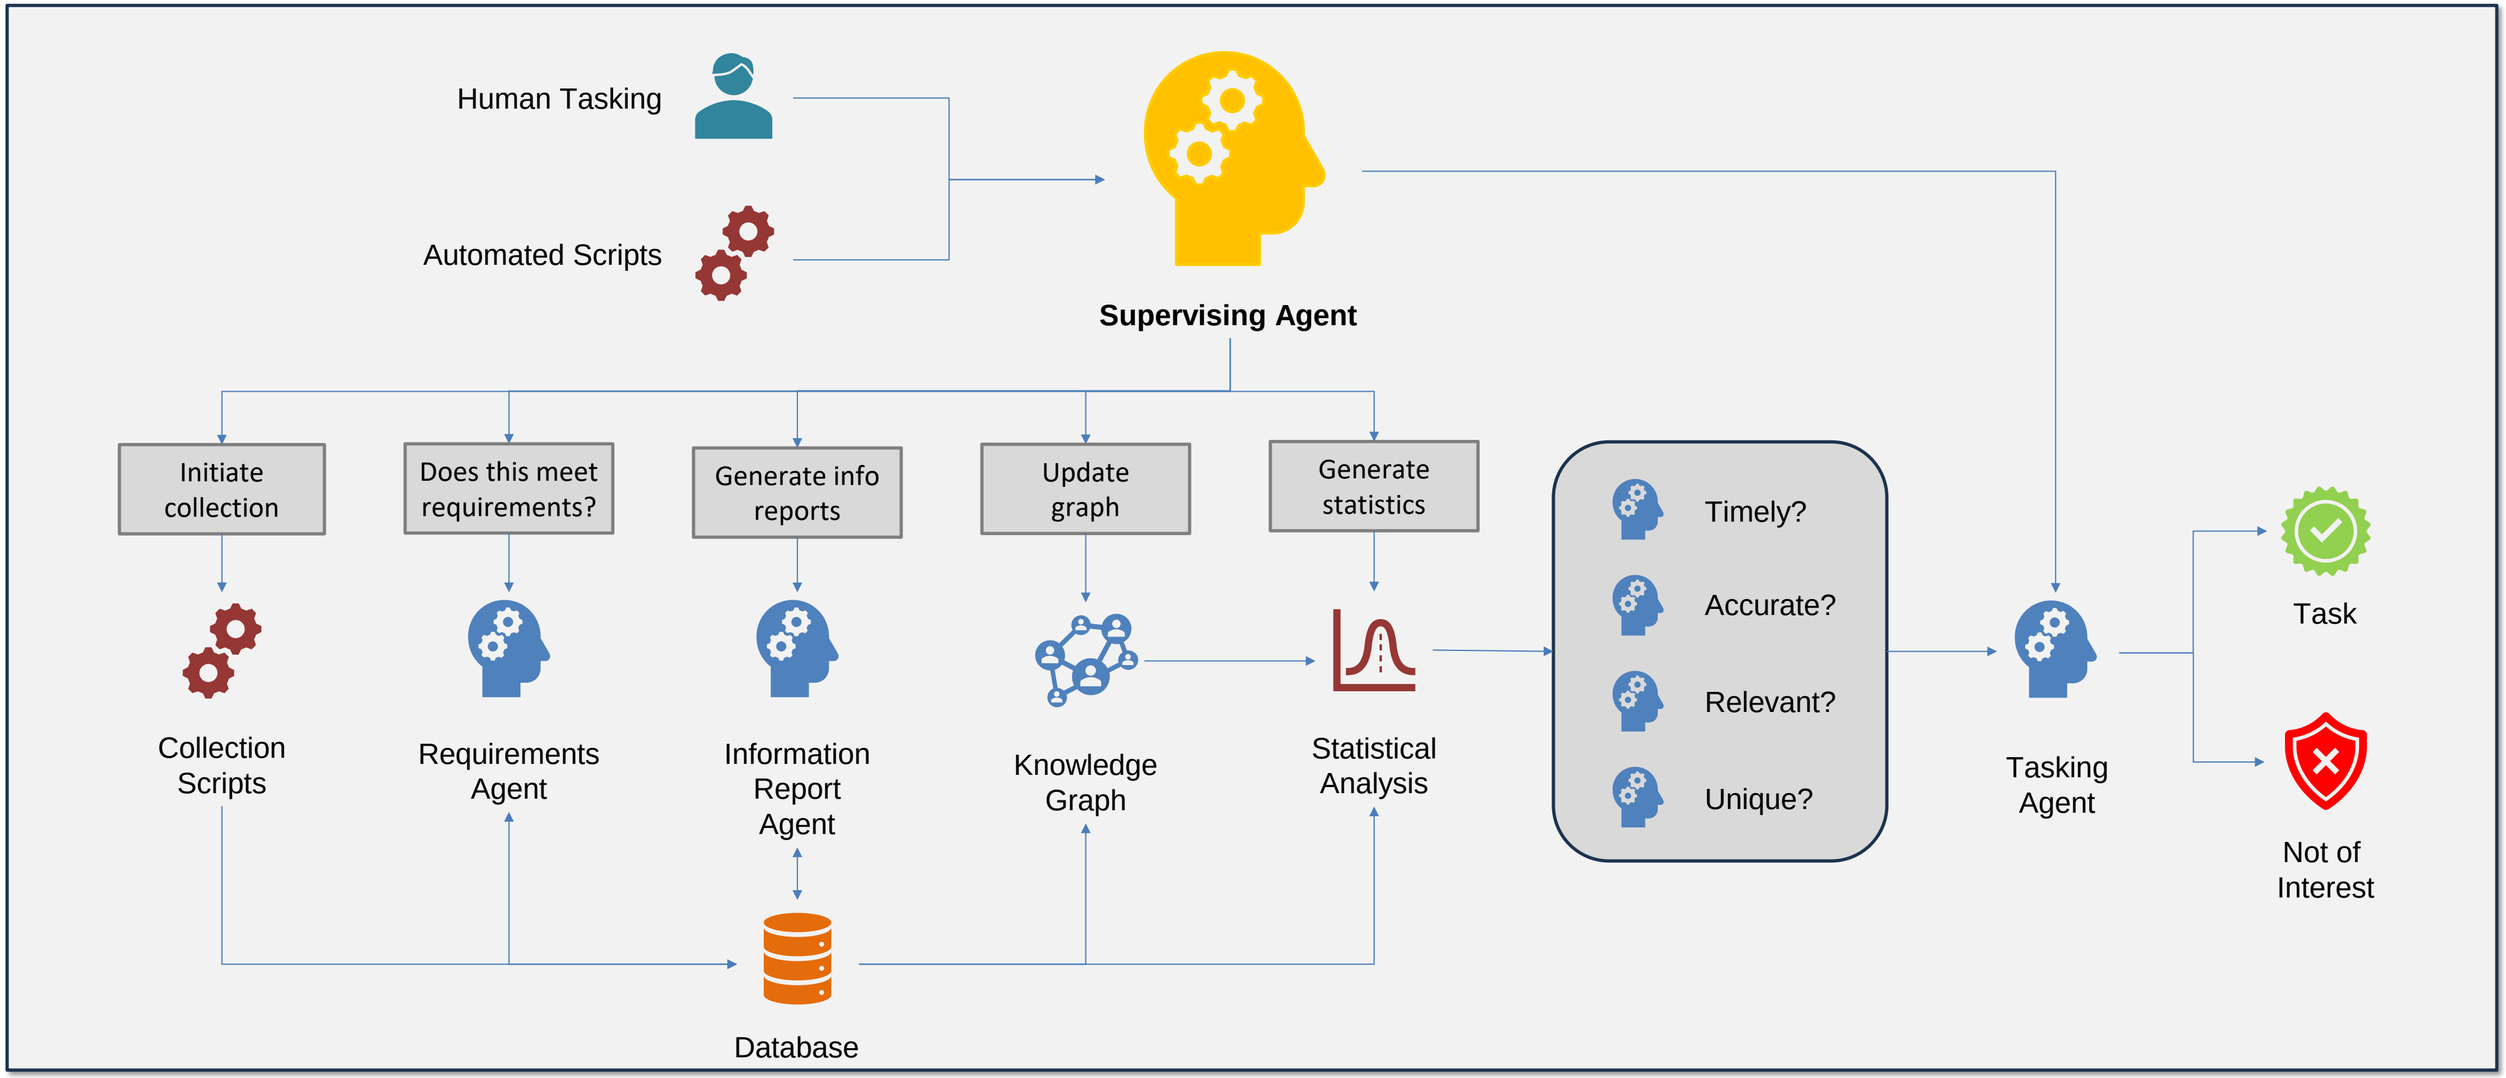

Processing

The real strengths of AI is in the processing phase of the intelligence cycle. These strengths include:

- Translating.
- Surfacing priority information.
- Formatting unstructured data.
- Working fast with good accuracy.

To get the most out of AI for processing, you need a well-managed intelligence function:

- A comprehensive set of intelligence requirements.
- Good collection management.
- A flexible data processing environment.
- A work environment that encourages the use of AI.

Most of these are simple LLM, ML, or statistical workflows

Writing information reports

The activity best suited to the capabilities of LLMs is generating information reports from unanalysed collected information. Asking an LLM to summarise a piece of information in a standard, repeatable way is well within its abilities, particularly when providing it with a good understanding of the context.

Take a social media post.

Check against intelligence requirements.

Pass to an LLM to summarise content and generate data.

Produce an information report and data.

Push to a database.

I've had a lot of success producing information reports from Telegram posts. The target sets that I've focused on are Hacktivists, the Russia-Ukraine War, and Right-Wing Extremism.

Hacktivism


```
{
  "cid": "CTI-TGM-1132964271-13567-20231204220140",
  "requirement_id": ["CTI-1.3.1", "CTI-1.3.3"],
  "information_report": "On 04 December 2023 at 22:01:40, the Telegram
        channel 'Black Security Team' posted a message by 'Tencher
        Scott' announcing a free cybersecurity course focused on SQL
        Injection vulnerabilities and countermeasures. The post
        explains that SQL Injection occurs when a backend developer
        fails to implement proper filtering while executing database
        queries. The message includes a link to the course hosted on
        'BlackSecurityTeam.com' and promotes it as a comprehensive web
        security training. An attached promotional image indicates that
        the course instructor is 'Mehdi Hassani' from the Black
        Security Team. The post also provides Telegram and website
        links for further engagement.",
  "analyst_comment": "This post promotes cybersecurity training with a
        focus on SQL Injection, a widely exploited web application
        vulnerability. While the course appears to be for educational
        and defensive purposes, similar materials can be leveraged for
        offensive security and penetration testing. The presence of a
        dedicated cybersecurity community and training website suggests
        an organised effort to spread cybersecurity knowledge,
        potentially attracting both security professionals and
        individuals with malicious intent.",
  "languages": ["Persian", "English"],
  "entities": [
        "04 December 2023 (DATE)",
        "22:01:40 (TIME)",
        "Black Security Team (ORG)",
        "Tencher Scott (PERSON)",
        "Mehdi Hassani (PERSON)",
        "SQL Injection (TECHNIQUE)",
        "BlackSecurityTeam.com (DOMAIN)",
        "T.me/Black_Security (ORG)"
        ]
}
```

Hacktivism

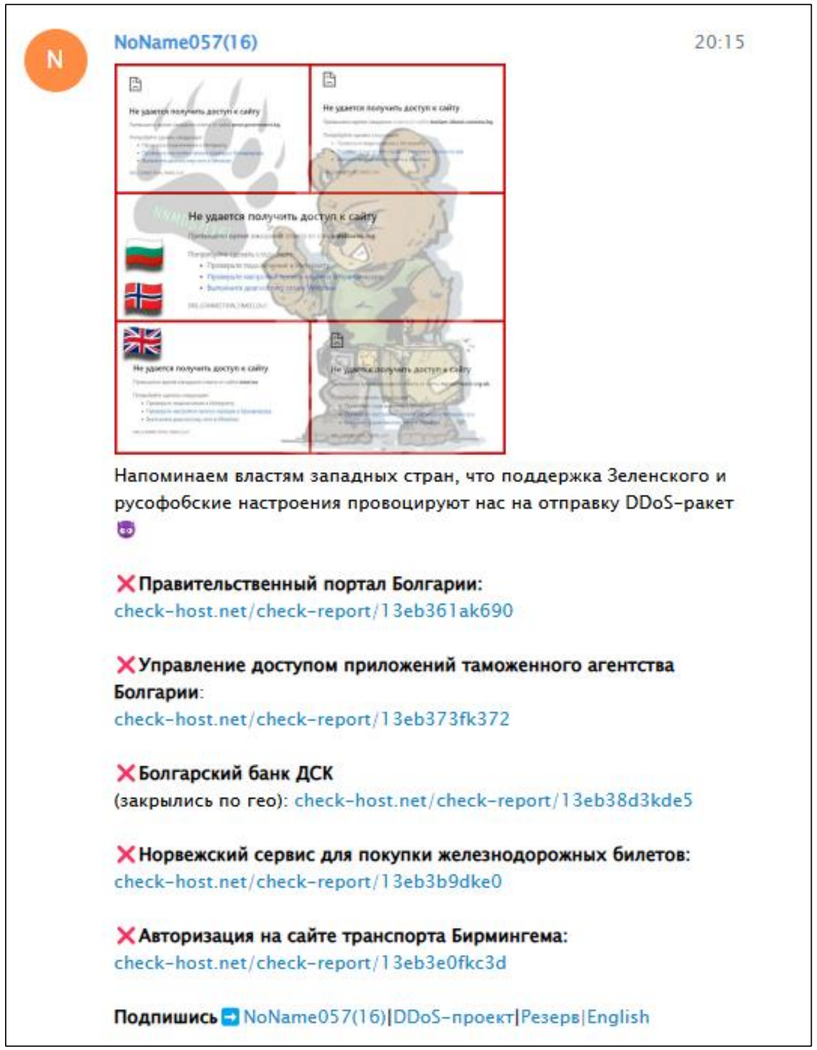


```
{
  "cid": "CTI-TGM-1732250465-5429-20231209030235",
  "requirement_id": ["CTI-1.1.1", "CTI-2.2.1", "CTI-2.2.2"],
  "information_report": "On 9 December 2023, NoName057(16) posted an image on their
        Telegram channel, showing website outage messages for multiple entities. The
        image indicated denial of service attacks targeting websites related to the
        government and financial sector in Bulgaria, as well as transportation
        services in Norway and the United Kingdom. The affected sites included: the
        Bulgarian government portal (government.bg), the Bulgarian Customs Agency
        application access portal (testiam-ids.ext.customs.bg), DSK Bank in Bulgaria
        (dskbank.bg), the Norwegian railway ticketing service (ruter.no), and the UK
        Swift transport card service (swiftcard.org.uk). The image displayed error
        messages in Russian stating 'Unable to access site' and 'Connection timed
        out'. This image was published alongside a text post discussing external
        media protection.",
  "analyst_comment": "This image is highly likely part of NoName057(16)'s ongoing
        pro-Russian politically motivated DDoS campaign. The selection of targets
        aligns with previous campaigns, focusing on entities in countries supporting
        Ukraine. The image serves as visual 'proof of success' for the group's
        attacks, aimed at bolstering credibility within their support base.",
  "languages": "Russian",
  "entities": [
        "9 December 2023 (DATE)",
        "NoName057(16) (ORG)",
        "Bulgaria (GPE)",
        "Norway (GPE)",
        "United Kingdom (GPE)",
        "Bulgarian Government Portal (ORG)",
        "Bulgarian Customs Agency (ORG)",
        "DSK Bank (ORG)",
        "Norwegian Railway Ticketing Service (ORG)",
        "Swift Transport Card Service (ORG)",
        "government.bg (URL)",
        "testiam-ids.ext.customs.bg (URL)",
        "dskbank.bg (URL)",
        "ruter.no (URL)",
        ]
}
```

Hacktivism

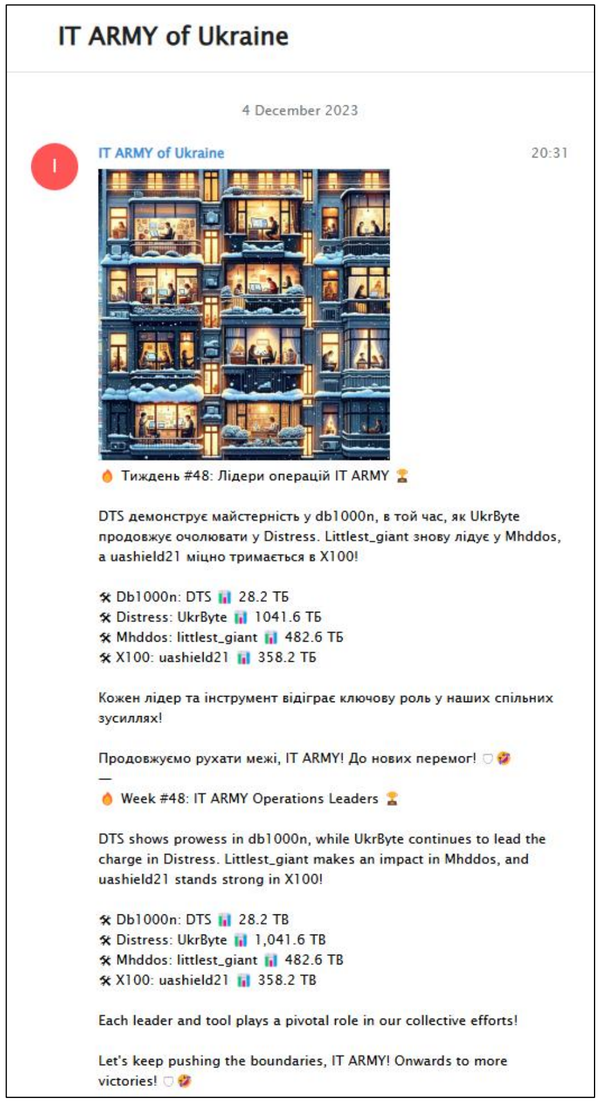

IT ARMY of Ukraine

4 December 2023

IT ARMY of Ukraine 20:31

🔥 Тиждень #48: Лідери операцій IT ARMY ⏳

DTS демонструє майстерність у db1000n, в той час, як UkrByte продовжує очолювати у Distress. Littlest_giant знову лідує у Mhddos, а uashield21 міцно тримається в X100!

🛠 Db1000n: DTS 📊 28.2 ТБ
🛠 Distress: UkrByte 📊 1041.6 ТБ
🛠 Mhddos: littlest_giant 📊 482.6 ТБ
🛠 X100: uashield21 📊 358.2 ТБ

Кожен лідер та інструмент відіграє ключову роль у наших спільних зусиллях!

Продовжуємо рухати межі, IT ARMY! До нових перемог! 🛡️❤️‍🔥
—
🔥 Week #48: IT ARMY Operations Leaders ⏳

DTS shows prowess in db1000n, while UkrByte continues to lead the charge in Distress. Littlest_giant makes an impact in Mhddos, and uashield21 stands strong in X100!

🛠 Db1000n: DTS 📊 28.2 TB
🛠 Distress: UkrByte 📊 1,041.6 TB
🛠 Mhddos: littlest_giant 📊 482.6 TB
🛠 X100: uashield21 📊 358.2 TB

Each leader and tool plays a pivotal role in our collective efforts!

Let's keep pushing the boundaries, IT ARMY! Onwards to more victories! 🛡️❤️‍🔥


```
{
  "cid": "CTI-TGM-1601423054-1828-20231204203150",
  "requirement_id": ["CTI-1.1.1", "CTI-2.1.3", "CTI-2.3.1", "CTI-2.3.2"],
  "information_report": "On 4 December 2023 at 20:31, the IT ARMY of Ukraine
        posted on their Telegram channel providing an update on operational
        leaders for week 48. The post identified four individuals or teams
        leading in the use of different cyber tools during the week. DTS led in
        the use of 'db1000n', generating 28.2 TB of traffic. UkrByte led
        operations using the 'Distress' tool, generating 1,041.6 TB of traffic.
        Littlest_giant led in the use of 'Mhddos', contributing 482.6 TB of
        traffic. Uashield21 led in 'X100' operations, producing 358.2 TB of
        traffic. The post highlighted that each of these leaders and tools
        played a key role in the group's collective efforts.",
  "analyst_comment": "This post is almost certainly related to ongoing
        distributed denial of service (DDoS) campaigns conducted by the IT ARMY
        of Ukraine against Russian or Russian-affiliated targets. The naming of
        specific tools (db1000n, Distress, Mhddos, X100) aligns with known
        tools used in crowdsourced DDoS attacks. The identification of
        operational leaders is likely intended to both motivate participants
        and publicly demonstrate the IT ARMY's continued activity and
        effectiveness. The use of both Ukrainian and English text indicates the
        message was intended for both domestic and international audiences.",
  "languages": ["Ukrainian", "English"],
  "entities": [
        "4 December 2023 (DATE)",
        "20:31 (TIME)",
        "IT ARMY of Ukraine (ORG)",
        "Telegram (ORG)",
        "DTS (PERSON)",
        "UkrByte (PERSON)",
        "Littlest_giant (PERSON)",
        "Uashield21 (PERSON) ",
        "db1000n (PRODUCT) ",
        "Distress (PRODUCT)",
        "Mhddos (PRODUCT)",
        "X100 (PRODUCT)"
        ]
}
```

Russia-Ukraine war

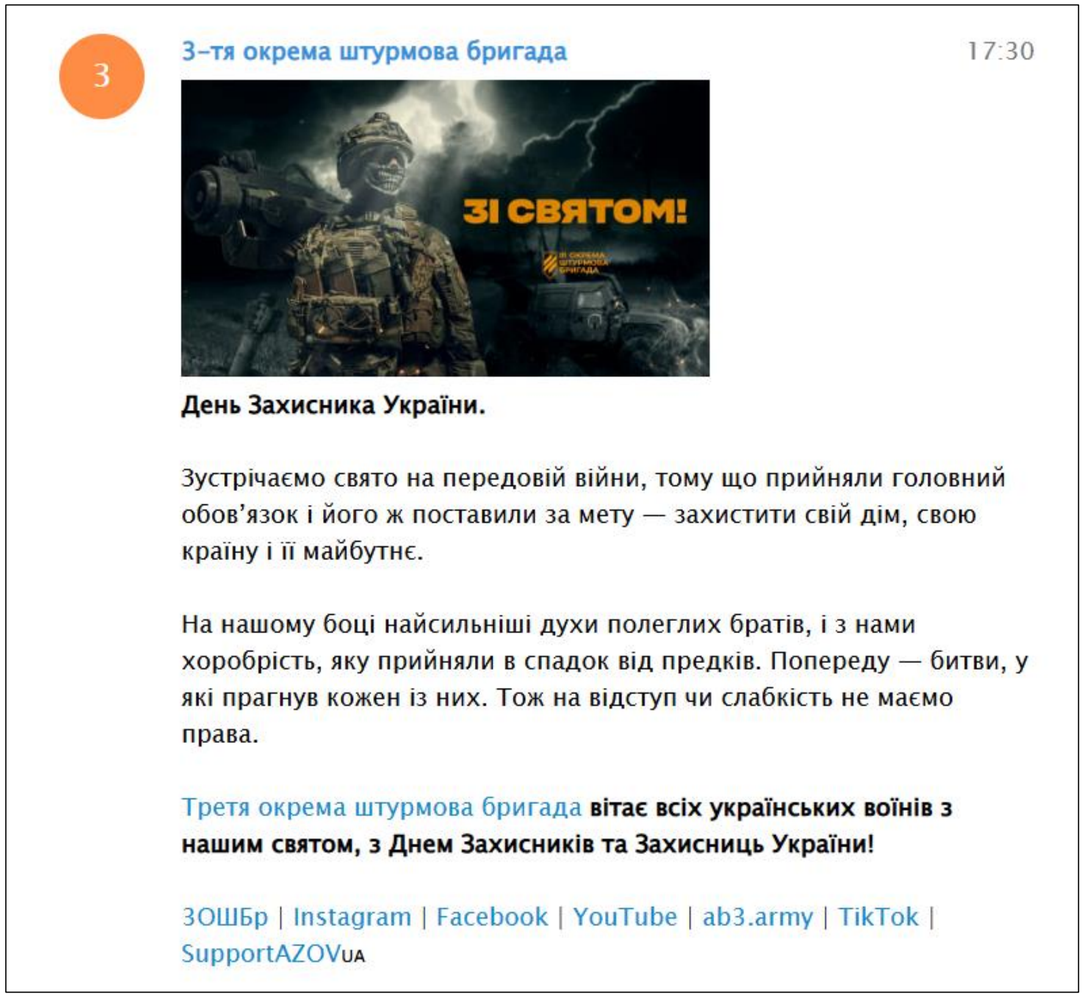


```
{
    "cid":, "RUK-TGM-1639691719-003203-20231001173007",
    "requirement_id": ["RUK-6.1.2"],
    "information_report": "On 01 October 2023 at 17:30 UTC, the 3rd Separate
        Assault Brigade (3 ОШБр) posted a message on their Telegram channel
        celebrating Defender of Ukraine Day. The post states that the brigade is
        marking the occasion while deployed on the frontlines, emphasizing their
        commitment to defending Ukraine, their homeland, and its future. It
        references fallen comrades and inherited bravery from ancestors, stating
        that retreat or weakness is not an option. The brigade extends greetings
        to all Ukrainian servicemen and women in honor of the national holiday.
        The message includes links to the brigade's social media and support
        channels, including Telegram, Instagram, Facebook, YouTube, and TikTok.",
    "analyst_comment": "This post follows a common Ukrainian military narrative,
        reinforcing themes of resilience, sacrifice, and national unity. The
        invocation of fallen comrades and ancestral bravery aims to boost morale
        and frame continued combat as an honorable duty. The inclusion of
        multiple social media links suggests an organized effort to increase
        public engagement and support. The mention of the national holiday ties
        the post to broader Ukrainian state messaging, which often emphasizes the
        military's role in national survival. The SupportAZOV link may indicate
        ties to the broader nationalist military movement, a common theme in some
        Ukrainian units' outreach efforts.",
    "languages": ["Ukrainian"],
    "entities": [
        "01 October 2023 (DATE)",
        "17:30 UTC (TIME)",
        "3rd Separate Assault Brigade (ORG)",
        "Ukraine (GPE)",
        "Defender of Ukraine Day (EVENT)",
        "Telegram (ORG)",
        "Instagram (ORG)",
        "Facebook (ORG)",
        "YouTube (ORG)",
        "TikTok (ORG)",
        "SupportAZOV (ORG)"
    ]
}
```


That was a good pickup this morning. Let's get a report together.
Tasking
Yep, they're bad guys...
Research
OK team, this is what I've got, what do you think?
Planning
Need at least 3,000 words...
Production
Good use of words of estimative probability...
Editing
As you can see, this is our guy, go get him team!
INTELLGENCE
Dissemination

Writing longer intelligence reports with AI

There are significant challenges in getting LLMs to write longer intelligence reports:

- Replicating the full process that experienced analysts use.
- Asking an LLM to extract the most important points from a corpus.
- Problems with context windows and hallucination (particularly at +80%).
- Capturing expert target knowledge.
- Referencing intelligence source information in a reliable way.
- Effectively assessing information.

The best solution at this stage is **AI-assisted intelligence production**.

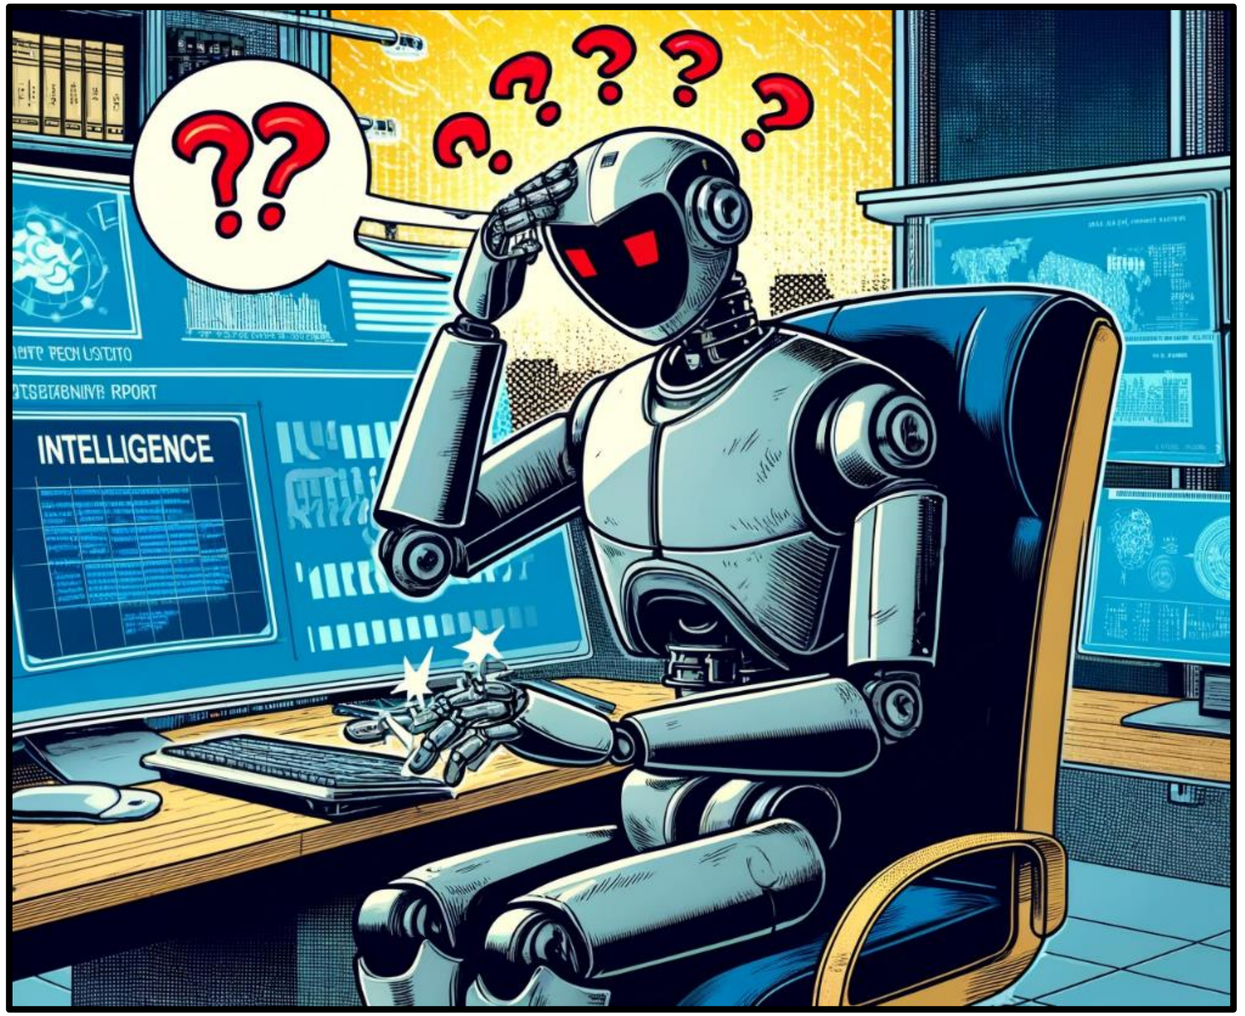

An example workflow for intelligence reports

Each AI icon is a multi-agent process.

AI does the heavy lifting, but humans own the judgment and the meaning.

Trying this workflow

I decided to have a go at building out this workflow using some old scripts and Cursor + Claude 3.7.

Writing* the code took about half an hour. The report used Gemini 1.5 Pro, took three minutes and cost $3.72.

Some issues with the report:

- Shorter than I would like.
- Paragraphs don’t go into enough detail.
- Some issues with the referencing.
- It’s clearly missed some attacks and countries.
- I didn’t sense check along the way.
- It didn’t have to generate the query.

Still, it produced something that is consistent with my understanding faster than any analyst could over that much data. I wouldn’t normally do this in one pass.

```
# Step 1: Generate key points from the reports
def generate_key_points(model, reports, requirements, num_versions=4):
    prompt = f"""
You are an expert intelligence analyst focusing on cyber threat intelligence.

SPECIFIC INSTRUCTIONS:
You are to produce a report on the Cyber Army of Russia Reborn (CARR) for the period of 2024.
Your focus is to be on the following points:
1. Their targeting prioirites and any rationale for their choices.
2. The methods they use to conduct their operations.
3. Any information about direction given by nation states.
4. Any other information that you believe would be useful to the intelligence community.

INTELLIGENCE REQUIREMENTS:
{json.dumps(requirements, indent=2)}

INTELLIGENCE REPORTS:
{json.dumps(reports, indent=2)}

Based on these reports, generate {num_versions} different versions of the top 5 key points that emerge from the data.
Each key point should:
1. Be directly tied to one or more specific intelligence requirements (reference the code)
2. Be supported by evidence from multiple reports where possible
3. Highlight significant patterns, trends, or notable events
4. Be concise but informative

For each version, take a slightly different analytical approach or emphasis.

Format your response as a JSON object with this structure:
{{
  "version_1": [
    {{
      "point": "Key point 1",
      "requirements": ["CTI-X.X.X"],
      "supporting_evidence": "Brief summary of supporting evidence"
    }},
```

```
File: data/2024_reports.json
Token count: 162,359
File size: 577.04 KB
Characters: 590,893
```

```
PS D:\(                                                                    .py
Step 1: Generating key points versions...
Step 2: Selecting best key points version...
Step 3: Generating detailed paragraphs...
Step 4: Editing paragraphs...
Step 5: Generating summary and assessment...
Step 6: Final editing...
Report generated and saved as intel_report_20250311_163348.md

Final Report Preview:
================================================================================
TITLE: CARR Cyber Group: Expanding Targeting and Evolving Capabilities

EXECUTIVE SUMMARY:
CARR, a suspected Russian state-influenced cyber group, has broadened its targeting from Ukrainian entities to include nations
n in geopolitical alignment.  The group primarily uses publicly available website vulnerability scanning tools to identify tar
 critical infrastructure, it currently lacks the demonstrated capability to pose a significant threat.  CARR's organizational
nternal conflict and marked by fluctuating collaborations with other pro-Russian cyber groups.


ASSESSMENT:
CARR’s expanded targeting almost certainly reflects increasing alignment with Russian strategic objectives and elevates the ri
onfidence that CARR will continue targeting nations supporting Ukraine, likely escalating to mor...
================================================================================
```

TITLE: CARR Cyber Group: Expanding Targeting and Evolving Capabilities

EXECUTIVE SUMMARY:

CARR, a suspected Russian state-influenced cyber group, has broadened its targeting from Ukrainian entities to include nations providing aid to Ukraine, indicating an escalation in geopolitical alignment. The group primarily uses publicly available website vulnerability scanning tools to identify targets. While CARR has expressed intent to target US critical infrastructure, it currently lacks the demonstrated capability to pose a significant threat. CARR's organizational structure appears fluid, potentially hampered by internal conflict and marked by fluctuating collaborations with other pro-Russian cyber groups.

ASSESSMENT:

CARR's expanded targeting almost certainly reflects increasing alignment with Russian strategic objectives and elevates the risk of wider cyber conflict. We assess with high confidence that CARR will continue targeting nations supporting Ukraine, likely escalating to more disruptive attacks, particularly against government and infrastructure targets. CARR may attempt to develop or acquire more sophisticated capabilities, though its current reliance on readily available tools suggests resource limitations. The fluidity of CARR's structure makes predicting future behavior challenging; however, continued monitoring of its tactics, techniques, and procedures (TTPs) and affiliations remains critical.

DETAILED ANALYSIS:

KEY POINT 1: Expanded Targeting Beyond Ukraine (CTI-3.1, CTI-2.2)

CARR has expanded its targeting to nations providing military and humanitarian aid to Ukraine, significantly broadening its operational scope. Initial attacks focused on Ukrainian government and infrastructure, including recruitment centers (CTI-TGM-1781536189-006102). Subsequent targeting of government and aid-related websites in Norway (CTI-TGM-1781536189-006125), Bulgaria (CTI-TGM-1781536189-006165), and Latvia (CTI-TGM-1781536189-006231) suggests an intent to pressure and disrupt Ukraine's international support network. This geographic expansion increases the risk of broader cyber conflict and necessitates a coordinated international response. This shift likely signifies a transition from patriotically motivated hacktivism to a geopolitically driven campaign aligned with Russian strategic objectives.

KEY POINT 2: Reliance on Publicly Available Tools (CTI-1.3, CTI-2.3)

CARR frequently uses publicly available website vulnerability scanning tools, such as check-host.net (CTI-TGM-1781536189-006091, CTI-TGM-1781536189-006137, CTI-TGM-1781536189-006170), indicating a pragmatic but potentially low-sophistication approach. This allows CARR to rapidly identify vulnerabilities without specialized expertise. While this reliance on public resources complicates attribution, it does not preclude the group from possessing or acquiring more advanced capabilities. Continued TTP monitoring is necessary to identify any evolution in sophistication.

KEY POINT 3: Suspected Russian State Influence (CTI-2.1.1, CTI-3.1.2)

Several indicators suggest a strong link between CARR and Russian intelligence services, though definitive attribution remains challenging. CARR's targeting aligns with Russian geopolitical interests, specifically pressuring nations aiding Ukraine (see Key Point 1). Its rhetoric often mirrors themes in Russian state-sponsored propaganda (CTI-TGM-1781536189-006158, CTI-TGM-1781536189-006220). An unverified report mentioning potential FSB contracts (CTI-TGM-1781536189-006094) further strengthens this assessment. While conclusive evidence of direct control is absent, these factors suggest CARR's operations are likely influenced, if not coordinated with, Russian intelligence objectives, raising concerns about potential escalation and the use of CARR as a proxy force.

KEY POINT 4: Threats Against Critical Infrastructure (CTI-1.1.4, CTI-2.2)

CARR has expressed intent to target US critical infrastructure, including water supply systems and energy companies (CTI-TGM-1781536189-006265, CTI-TGM-1781536189-006428). However, no confirmed successful attacks causing significant disruption or damage have been observed, suggesting limited capabilities or a prioritization of other targets. Despite this, CARR's stated intent necessitates vigilance and proactive defensive measures by potential target organizations.

KEY POINT 5: Fluid Organizational Structure (CTI-2.4)

CARR's organizational structure appears fluid and evolving, potentially marked by internal conflict, shifting allegiances, and varying levels of coordination with other pro-Russian cyber groups, such as 22C (CTI-TGM-1781536189-006256) and NoName057(16) (CTI-TGM-1781536189-006412). Reports indicate internal disputes and shifting allegiances within CARR (CTI-TGM-1781536189-006094, CTI-TGM-1781536189-006882). Understanding these internal dynamics is crucial for anticipating future actions, but this fluidity complicates predicting behavior and assessing overall capabilities. Continuous monitoring of CARR's internal and external relationships is necessary to accurately assess the group's evolving threat landscape.

What’s the trick to it all?

It comes down to knowing how intelligence works inside out.

There are still pieces to the puzzle that I haven’t quite figured out, particularly around the assessment of the reliability and accuracy of the information.

But if you can **critically assess** your own processes, **break them up** into meaningful chunks, and produce **clear instructions**, then you can build an army of AI assistants.

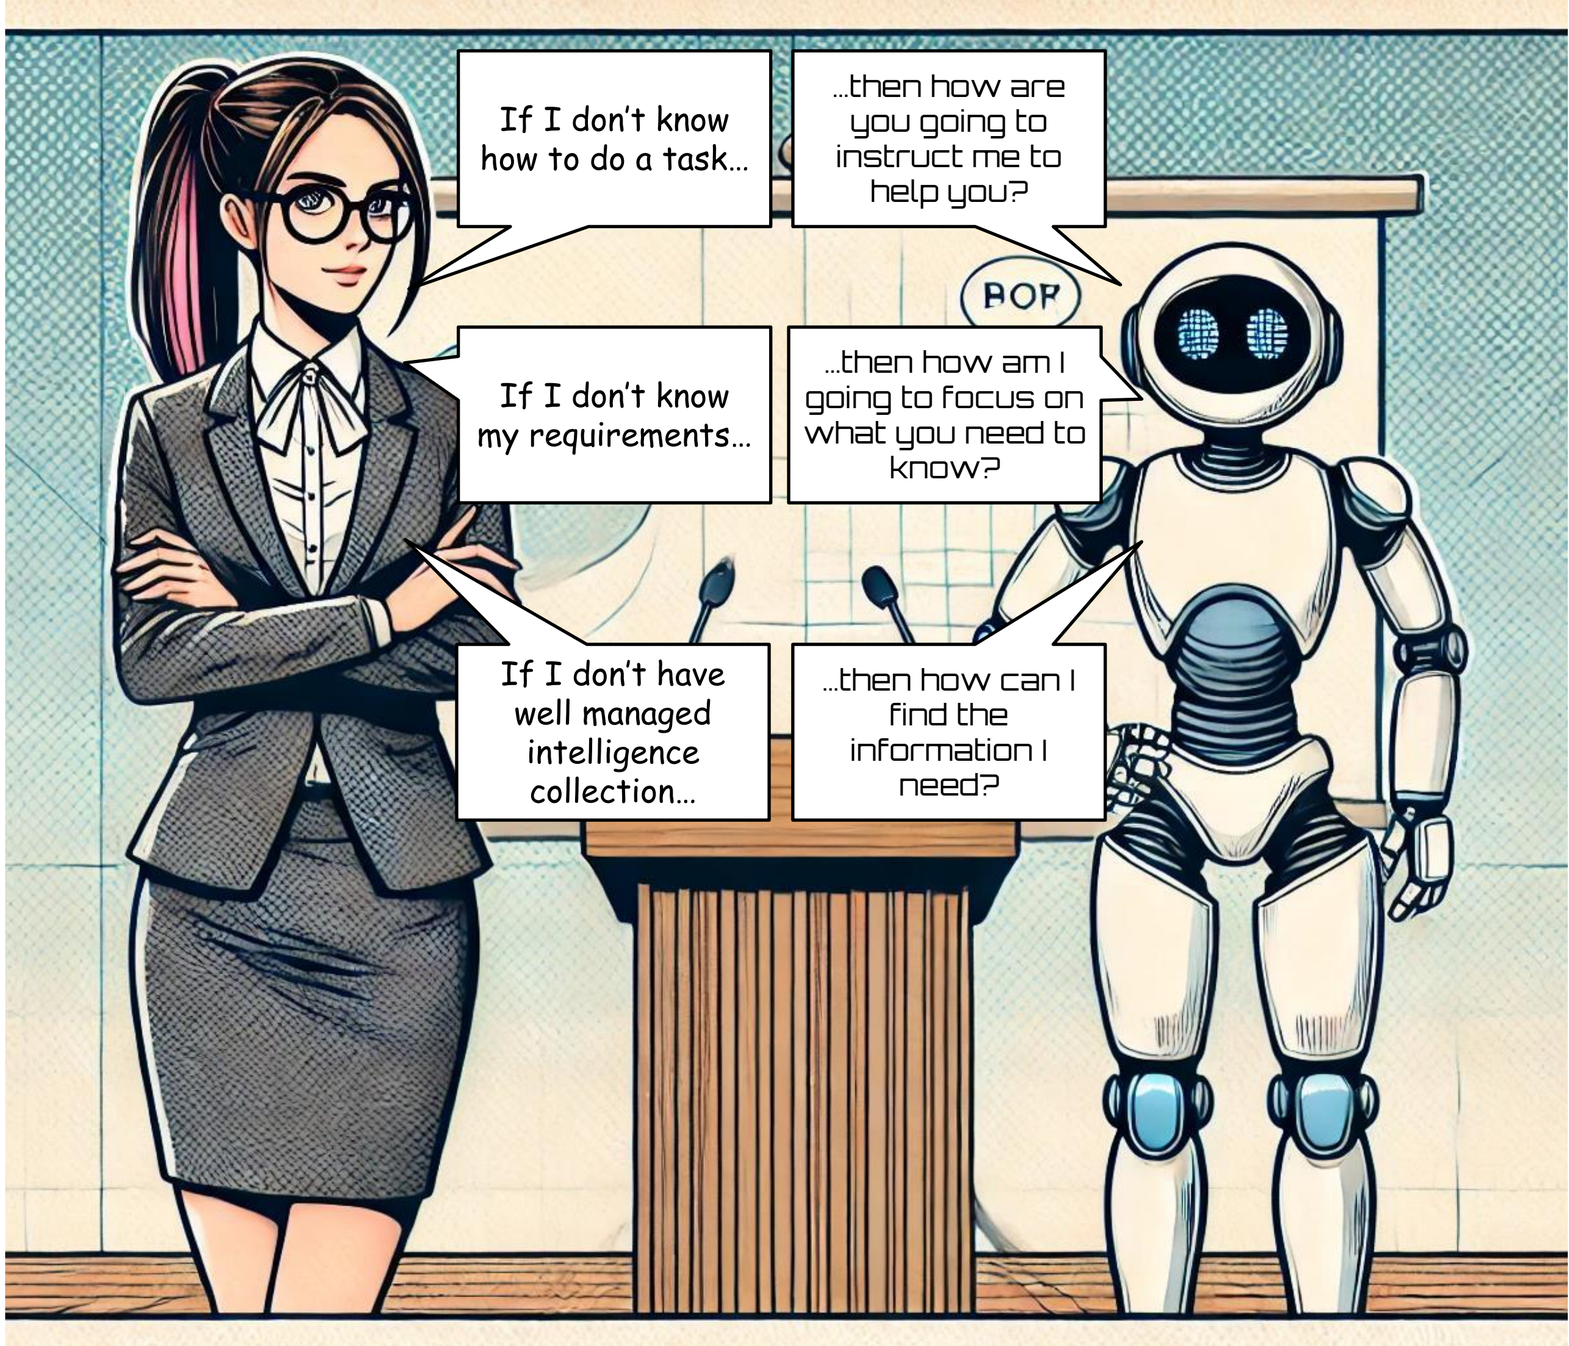

Intelligence process + AI to generate knowledge

We’ve only scratched the surface today, but there is more going on here than just threat intelligence. I’ve applied these principles and processes, in limited ways, to other domains of knowledge.

It’s more than a workflow; you can use these principles for trustworthy machine-assisted knowledge creation in any domain.

Start with the Requirement

Follow a Transparent Process

Preserve the Epistemic Trace

Structure the output

Keep human judgment in the loop

Contact

Brendon Hawkins
brendonhawkins@independint.com.au