
// Transforming High-Stakes Decision Making with Explainable and Reliable AI
// Do more, better & faster, with AI trained on your wisdom.
Bring human and machine intelligence together with Tulpa AI
tulpa®

Your Organisational Challenges

We understand that your organisation is grappling with unprecedented challenges—including rapid technological advancements, stringent regulatory pressures, sophisticated cyber threats, intense global competition, and a scarcity of highly skilled resources—that make critical decision-making more complex and demanding than ever before.

Navigating a Complex and Volatile Environment

Rapid Technological Advances

- Staying ahead in a world where technology evolves at breakneck speed, introducing new vulnerabilities and attack vectors.

Regulatory Pressures

- Adhering to increasingly stringent regulations demanding transparency, data protection, and accountability—such as GDPR, CCPA, HIPAA, and industry-specific compliance standards—is a significant challenge.
- Navigating emerging AI legislation like the EU AI Act, which imposes strict requirements on AI transparency, explainability, and risk management, along with new cyber-specific regulations such as the NIS Directive and its successor NIS2, adds complexity to compliance efforts.

Sophisticated Cyber Threats

- Defending against advanced persistent threats (APTs), zero-day exploits, ransomware, and other sophisticated cyber attacks that could compromise your operations and reputation.

Global Competition

- Maintaining a competitive edge amidst global market fluctuations, emerging competitors, and the need to protect intellectual property and sensitive data from cyber espionage.

Talent Shortages & Resource Constraints

Scarcity of Skilled Cybersecurity Professionals

- Securing highly skilled professionals with specialised expertise in cybersecurity is increasingly difficult due to a global talent shortage.

Overburdened Teams

- Limited availability of seasoned talent can hinder your organisation's ability to innovate and respond swiftly to emerging threats and opportunities, leading to overworked teams and potential burnout.

Risk of Losing Critical Knowledge

- The competitive market for skilled professionals leads to higher turnover, risking the loss of critical institutional knowledge and weakening your security posture.

Time-Consuming Talent Development

- Developing the necessary wisdom and expertise in new talent requires significant time and investment, often taking years to cultivate—time that organisations under constant threat may not have.

The Challenge with Current AI Models

While popular AI capabilities - such as large language models (LLMs), reinforcement learning (RL) and deep learning - excel in approximately 90% of applications, they often fall short in high-stakes tasks and environments where decisions carry significant consequences. We recognise these critical challenges and have developed tools to overcome the inherent limitations of these capabilities.

Limitations of Current AI Models in High-Stakes Environments

Opacity and Lack of Explainability Black Box Functionality

Current AI systems often operate as “black boxes,” providing outputs without insight into their reasoning process.
This opacity hinders trust and makes it difficult for security teams to rely on AI in critical decision-making.

Challenges in Incident Response

Without understanding the AI’s reasoning, analysts may hesitate to act on alerts, delaying response times.
Inability to explain decisions complicates communication with stakeholders and regulators during and after security incidents.

Inadequate Reasoning Ability

Lack of Causal Understanding

Current AI models often fail to understand causal relationships, leading to recommendations that lack contextual relevance.
This limitation hampers the detection of sophisticated threats that do not follow historical patterns or known signatures.

Inability to Adapt to Novel Threats

Without the ability to reason like a human expert, these models can’t adequately address complex, emerging cyber threats. Zero-day exploits and advanced persistent threats that require deeper contextual analysis may be missed.

Impact on Your Organisation

Risk Exposure

// Overlooked Threats
- Decisions based on unexplained AI outputs can lead to errors with significant consequences, such as failing to detect a critical breach.
- The inability to trust AI alerts may result in security gaps and increased vulnerability to attacks.

// Delayed Response Times
- Hesitation to act on AI recommendations can slow down incident response, allowing threats to escalate.

// Unnecessary business disruption
- An over-reliance on AI outputs could lead to unnecessary remedial action which costs time, money and reputational damage

Compliance Challenges

// Regulatory Scrutiny
- Regulatory environments increasingly demand transparency and accountability in AI-driven decisions, especially regarding data protection and privacy laws like GDPR and CCPA.
- Inability to explain AI actions can lead to non-compliance penalties and legal repercussions.

// Audit Difficulties
- Lack of explainability complicates the auditing process, making it challenging to justify actions, and to demonstrate due diligence and effective security measures to regulators and stakeholders.

Operational Inefficiencies

// Resource Drain
- Teams spend valuable time interpreting ambiguous AI outputs or investigating false positives, detracting from proactive security measures.
- A lack of trust in AI outputs could lead to counterproductive work such as an overreliance on manual analysis that increases workload and the potential for human error.

// Talent Utilisation
- Skilled professionals are bogged down with validating AI outputs rather than focusing on strategic initiatives and threat hunting.

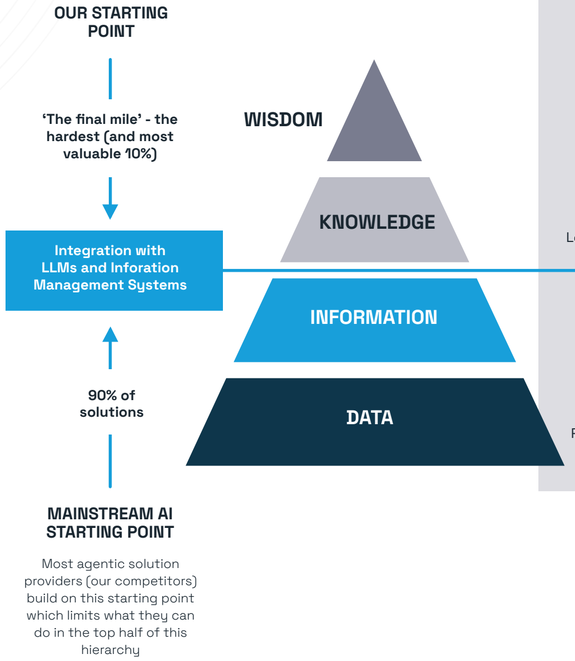

REASONING TYPES	VALUE PROPOSITION	BENEFIT
Counterfactual 'What if?' scenarios Action planning	**Understanding why** **Predicting outcomes** **Guiding principles**	**Better decisions through causal understanding**
Causal Cause-effect chains Intervention effects Learning new concepts	**Domain expertise** **Best practices** **Expert rules and heuristics**	**Preserved expertise through scalable knowledge solutions**
Associative Pattern matching Similarity finding	**Pattern recognition** **Data analytics** **Statistical relationships**	**Enhanced insights through integration**
Statistical Correlation analysis Probability estimation	**Raw data** **Data processing**	**Validated results through data**

tulpa®

Introducing Tulpa AI: A New Paradigm

A Solution Designed for Your Needs

At Tulpa AI, we recognise not only the technical challenges but also the human resource constraints your organisation faces. Our causal neurosymbolic AI platform is crafted to:

Enhance Your Decision-Making Processes

Contextually Relevant Insights
By providing AI that reasons like your experts, you gain insights that are contextually relevant and actionable against sophisticated threats.

Empower You with Explainability

Transparent AI Recommendations
Every AI-driven alert and recommendation comes with clear explanations, so you know exactly why a decision was suggested, building trust within your team and with stakeholders.

Reduce Dependence on Scarce Talent

Accessible to All Team Members
Our user-friendly interface allows your existing team members, regardless of their technical background, to customise and manage AI agents without the need for specialised AI expertise.

Facilitate Seamless Human-Machine Collaboration

Augmenting Human Expertise
Our platform acts as a "second brain," augmenting your team's capabilities without replacing the invaluable human touch, enhancing efficiency and effectiveness in threat detection and response.

Customise AI Agents to Your Specific Needs

Leverage Your Existing Expertise

Scaling Internal Knowledge
We encode the wisdom of your experts into a 'wisdom model' which makes their expertise and know-how available to colleagues and AI around the clock

User-Friendly Interface

Easy Modification and Adaptation
Modify and adapt AI agents quickly without needing extensive technical knowledge, allowing you to respond to new threats and regulatory changes efficiently.

Focus on High-Stakes "10% Problems"

Strategic Advantage in Critical Areas
Target the critical security issues often overlooked by or unfit for mainstream solutions, giving you a strategic advantage despite resource constraints.

How Tulpa AI Addresses Your Talent Challenges

Scaling Your Cybersecurity Expertise

Amplifying Your Team's Capabilities

- **Enhanced Productivity Without Additional Hiring**
  // Tulpa AI enables your current cybersecurity workforce to achieve more without the immediate need to hire additional specialised talent.
  // By automating routine tasks and providing advanced insights, your team can focus on strategic initiatives and critical decision-making.

Retaining and Utilising Institutional Knowledge

- **Capturing Expert Insights**
  // By embedding your experts' insights into our Tulpa AI agents, you ensure that critical knowledge is captured, retained, and accessible across your organisation, mitigating the impact of staff turnover.
  // This continuity strengthens your security posture by preserving institutional memory and best practices.

Empowering Non-Technical Staff

- **Accessible Advanced AI Capabilities**
  // With our intuitive platform, team members from various backgrounds can harness advanced AI capabilities, fostering a culture of security awareness and reducing bottlenecks caused by talent shortages.
  // This democratisation of AI tools allows for broader participation in a wider spectrum of cybersecurity efforts, enhancing overall resilience.

Reducing Dependence on Scarce Talent

Simplified AI Management

- **User-Friendly Interface**
  // Our platform eliminates the need for a large team of data scientists and AI specialists, making advanced AI accessible to organisations with limited specialised resources.
  // Security professionals can easily customise and manage AI agents without deep technical expertise.

Cost Efficiency

- **Optimising Resources**
  // By optimising your existing team's productivity, you can achieve more with less, allocating resources more effectively across your organisation.
  // Reduced reliance on external consultants or additional hires leads to significant cost savings.

Improved Productivity

- **Accelerated Response Times**
  // Tulpa AI accelerates offensive security practices, enabling your team to act swiftly and effectively.
  // Streamlined workflows and AI-assisted analysis and automated actions reduce manual workload and improve efficiency.

The Tulpa Advantage for Your Organisation

Tulpa AI eliminates the "black-box" problem by providing causal explanations, not just correlations, and enhances human decision-making, rather than replacing it.

Transparent and Explainable AI

Build Trust Through Clarity

- Gain confidence with AI that shows you the 'why' behind every recommendation, enhancing trust and facilitating quicker, more informed decisions.
- Simplify compliance efforts by having clear explanations readily available for audits and regulatory reviews.

Customised Solutions

Tailored to Your Security Challenges

- Easily customise AI agents to address your organisation's specific cybersecurity threats without the need for deep technical expertise.
- Adapt quickly to emerging threats by modifying AI agents as needed.

Enhanced Human-Machine Teaming

Amplify Your Team's Capabilities

- Empower your security professionals by augmenting their expertise with advanced AI tools, allowing for more effective and informed decision-making.
- Free up valuable time for your team to focus on strategic initiatives rather than routine tasks.

Focus on Critical Challenges

Address High-Stakes Security Issues

- Tackle the sophisticated cyber threats unique to your industry that mainstream AI solutions often overlook.
- Prioritise threats based on potential impact, ensuring critical issues receive immediate attention.

Optimising Human Resources

Maximise Existing Talent

- Reduce reliance on scarce cybersecurity talent by empowering more of your cyber and IT workforce with user-friendly, advanced AI tools.
- Improve productivity and efficiency without the immediate need to expand your team.

Use Cases and Applications

Tulpa AI's advanced cybersecurity agent offers a range of applications that enhance your organisation's security operations, efficiency, and overall defence posture. Below are key use cases demonstrating how our agent can deliver significant benefits:

1. Assistant to Experts / Automation

Benefits:

// **Time Savings:** Save 20-30% of penetration testers' time by automating low-value, repetitive tasks such as running and waiting for scans and collating reports.
// **Increased Team Capacity:** Enable your team to handle more projects without increasing headcount.
// **Reduced Human Error:** With state tracking and intuitive visualisation, action options are presented at strategic and tactical levels, minimising mistakes and improving decision-making.

2. Empowering Junior Penetration Testers

Benefits:

// **Skill Enhancement:** Junior pen testers improve their skills and abilities while reducing the overhead on senior pen testers for training.
// **Improved Accuracy:** Increase the accuracy of junior pen testers' decisions by 40% through direct learning from expert knowledge encoded in the agent.
// **Increased Efficiency:** Boost the speed of pen tests performed by junior pen testers by 80%.

3. Enhancing SOC Team Capabilities

Benefits:

// **Hybrid Role Development:** Security Operations Centre (SOC) teams develop hybrid roles where staff perform both defensive and offensive security tasks as part of their normal workflow.
// **Cost Savings:** Save £40-50k per person by recruiting staff earlier in their careers and upskilling them using our agent.
// **Skill Diversification:** Enable team members to perform offensive security aspects without needing to be experts.

4. Breach Attack Simulations and Cyber Exercises

Benefits:

// **Enhanced Training:** Deploy our agent as part of breach attack simulation exercises or other cyber exercises to assist traditional red teamers or act as a machine-speed adversary.
// **Improved Defence Responses:** Expand the quality of cyber defence responses by 10-20%.
// **Proactive Vulnerability Management:** Identify and remediate vulnerabilities against machine-speed cyber adversaries on the horizon.

5. Training Aid for Defensive AI Capabilities

Benefits:

// **Cost Efficiency:** Save £700-£1800 per day compared to hiring expert human red teamers generating static datasets.
// **Adaptive Training:** Deploy an adaptive agent on-demand, complete with the latest exploits, fitting into defensive AI training pipelines.
// **Enhanced Learning:** Provide a dynamic training environment for defensive teams to improve their AI capabilities.

6. Post-Penetration Test Validation

Benefits:

// **Cost Savings:** Save £20-100k by avoiding the need to re-run pen tests manually after patches are applied.
// **Improved Security Posture:** Confirm that vulnerabilities identified in previous pen tests have been effectively addressed.
// **Executive Assurance:** Provide significant value to CISOs who can report with more confidence to the Board on the organisation's security status.

7. Continuous Testing of Defensive Tools

Benefits:

// **Maximised ROI:** Increase the return on investment on defensive tools and systems by ensuring they perform optimally.
// **Ongoing Assurance:** Improve assurance of your organisation's defensive posture by continuously testing tools like intrusion detection systems against sophisticated threats.
// **Threat Preparedness:** Ensure defensive tools are calibrated to detect 'low and slow' adversaries and other advanced attack techniques.

8. Verification of Third-Party Penetration Tests

Benefits:

// **Time Efficiency:** Reduce the time required to validate or verify third-party pen tests by 90% through automation.
// **Enhanced Verification Quality:** Perform a much larger sample size than economically feasible by humans, improving the quality of the verification process.
// **Trust but Verify:** Ensure that outsourced pen testing services meet your organisation's security standards.

9. Automated Penetration Test Reporting

Benefits:

// **Administrative Efficiency:** Significantly reduce the administrative overhead of documenting a pen test.
// **Improved Reporting Accuracy:** Enhance the quality and accuracy of pen test reports, leading to better remediation planning.
// **Workflow Optimization:** Help pen testers keep track of their activities, reducing the overall time taken to perform tests and improving project management.

By leveraging Tulpa AI's cybersecurity agent across these use cases, you can enhance your organisation's security operations, optimise resources, and strengthen your defences against evolving cyber threats.

Technical Insights: How Tulpa AI's Causal Neurosymbolic AI Works

Understanding the technology behind Tulpa AI is essential to appreciating how it revolutionises decision making in high-stakes environments. Our platform is built upon a unique combination of causal reasoning, neurosymbolic integration, and continuous machine learning. Below, we delve into each component to explain how they work together to provide transparent, explainable, and highly effective AI solutions.

Causal Reasoning Engine

At the core of Tulpa AI is the Causal Reasoning Engine, which fundamentally changes how AI interprets data:

- **Understanding Cause-Effect Relationships:**
  // Unlike popular AI models that rely heavily on pattern recognition and correlations, our causal reasoning engine focuses on identifying and understanding the underlying cause-effect relationships within data.
  // This approach enables our AI to make informed decisions based on how different variables interact with and influence one another, much like human reasoning.
- **Contextual Relevance:**
  // By grasping causality, our AI provides insights that are contextually relevant and actionable.
  // This ensures that recommendations are not just statistically significant but also make practical sense within the specific context of your organisation's operations.

Benefits for Your Organisation:

- **Improved Decision Accuracy:**
  // Decisions are based on a deep understanding of causal mechanisms, reducing the likelihood of errors.
- **Enhanced Trust and Transparency:**
  // Clear explanations of why and how decisions are made build confidence among users and stakeholders.

Neurosymbolic Integration

Tulpa AI's innovation lies in its ability to merge two traditionally separate AI paradigms through Neurosymbolic Integration:

- **Merging Symbolic Logic with Neural Networks:**
  // Symbolic AI involves using explicit, human-readable symbols and rules to represent knowledge and reasoning processes.
  // Neural Networks excel at pattern recognition and handling unstructured data but often lack explainability.
  // Tulpa AI integrates these two approaches, allowing the system to leverage the strengths of both.
- **Advanced Reasoning Capabilities:**
  // The symbolic component handles high-level reasoning tasks, such as planning, logic, and understanding complex relationships.
  // The neural network component processes raw data inputs, such as text, images, or sensor data, extracting meaningful patterns.
- **Synergistic Functionality:**
  // By combining these elements, Tulpa AI can interpret data inputs, reason about them symbolically, and produce outputs that are both accurate and explainable.
  // This synergy enables our AI to handle complex tasks that require both data-driven insights and logical reasoning.

Benefits for Your Organisation:

- **Comprehensive Analysis:**
  // Ability to process and understand complex data sets that require nuanced reasoning.
- **Explainable Outputs:**
  // Results are presented with clear logic and reasoning paths, making it easier for users to understand and trust our AI's conclusions.

Machine Learning Component

The Machine Learning Component ensures that Tulpa AI remains dynamic and continuously improves over time:

- **Continuous Learning from New Data Inputs:**
  // Our AI system is designed to learn and adapt from every interaction and data input it receives.
  // This ongoing learning process allows the AI to stay up-to-date with the latest information, trends, and patterns relevant to your organisation.

- **Adaptability to Changing Environments:**
  // As your organisation's needs evolve, Tulpa AI evolves with you, ensuring that the insights and recommendations remain relevant and effective.

- **Data-Driven Enhancements:**
  // Machine learning algorithms analyse outcomes and user feedback to refine decision-making processes.
  // This iterative improvement cycle enhances the accuracy and reliability of the AI over time.

Benefits for Your Organisation:

- **Future-Proof Solutions:**
  // The Tulpa agent's ability to learn ensures long-term value and relevance, even as industry landscapes change.
- **Personalisation:**
  // The system becomes increasingly tailored to your organisation's specific patterns, preferences, and requirements.
- **Catalyse the Human-AI flywheel:**
  // A self-reinforcing cycle where increased usage leads to enhanced AI capabilities, which in turn empowers users to be more effective, creating a loop of mutual growth and knowledge advancement.

Outcome:

- **Informed Decision-Making:**
  // Your team can make a well-informed choice, understanding not just what the AI recommends but why.
- **Regulatory Compliance:**
  // The explainable nature of the AI's reasoning supports compliance with transparency requirements.
- **Increased Confidence:**
  // The clarity and depth of insights build trust in the AI's capabilities, enhancing collaboration between human and machine.
- **Improved Productivity**
  // By streamlining processes and reducing the time spent on data interpretation, your team can focus on higher-value tasks, leading to increased efficiency and productivity.

Summary of Benefits:

- **Transparency and Explainability:**
  // Every decision is backed by clear reasoning, enhancing trust and compliance.
- **Enhanced Human-Machine Collaboration:**
  // The AI serves as a "second brain," augmenting human expertise without replacing it.
- **Scalability and Adaptability:**
  // The system grows and evolves with your organisation, ensuring long-term value.
- **Addressing High-Stakes Challenges:**
  // Tulpa AI's advanced reasoning capabilities make it uniquely suited for environments where decisions have significant consequences.

Tulpa AI's causal neurosymbolic technology brings transformative potential to your organisation's decision-making processes.

Our Tulpa AI agent doesn't just provide answers—it offers understanding, empowering you to navigate complex challenges with confidence and clarity.

Tulpa AI Results

Through our research in collaboration with world-class national security agencies and cybersecurity teams, Tulpa AI has demonstrated significant improvements in productivity and efficiency through Human-Machine Teaming (HMT). By combining our causal neurosymbolic AI platform with human expertise, teams achieved faster results and generated insights that surpassed what experts could accomplish alone. This synergy between human intuition and AI reasoning not only accelerated analysis but also enhanced decision-making processes, showcasing the transformative impact of HMT in high-stakes environments.

Human Machine Testing - A/B trials

In an A/B test simulating penetration testing tasks, participants were divided into two groups: one collaborated with the Tulpa AI agent (HMT group), and the other worked without it (User Only group). The objective was to discover services, assets, and subnets, ultimately gaining access to a target asset called *BoardChief*.

A/B HMT results

// HMT graphs with “Agent Followed” and “Agent Not Followed”. “Agent Followed” was verified by users choosing the “Execute” agent recommendation button in the tool. Users could still choose to follow the agent, manually executing the same action, but this was not recorded as “Agent Followed”.

The **first set of results** compares the HMT group and Users Only group for certain discovery metrics (Figure 1) including the number of services (ssh, ftp, smb etc), assets and subnets found, and against exploit metrics (Figure 2), including for gaining access or privileges of an asset. In both cases, these were tracked as a turn based timeline, where a turn is defined as the completion of an action, for example Portscan against 0.0.0.0 or an action failure.

In both Figure 1 and 2, the HMT users, on average, discovered more services, assets and subnets and gained access and privileges on a greater number of assets.

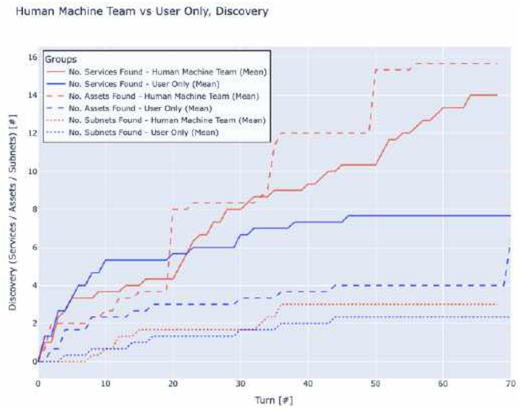

Figure 1: Number of discoveries achieved for HMTs and Users Only.

Figure 2: Number of exploits achieved for HMTs and Users Only.

Through our research in collaboration with world-class national security agencies and cybersecurity teams, Tulpa AI has demonstrated significant improvements in productivity and efficiency through Human-Machine Teaming (HMT). By combining our causal neurosymbolic AI platform with human expertise, teams achieved faster results and generated insights that surpassed what experts could accomplish alone. This synergy between human intuition and AI reasoning not only accelerated analysis but also enhanced decision-making processes, showcasing the transformative impact of HMT in high-stakes environments.

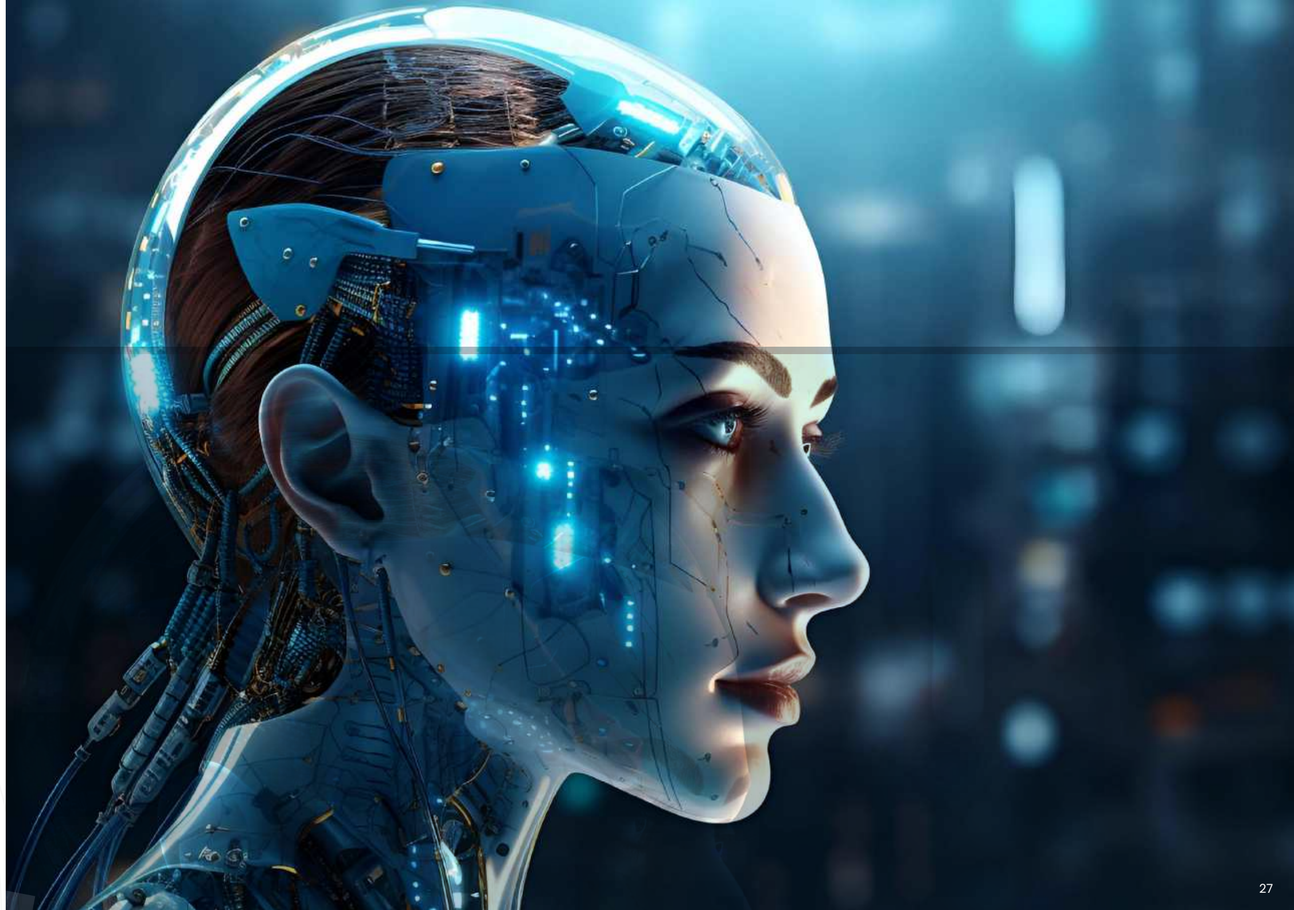

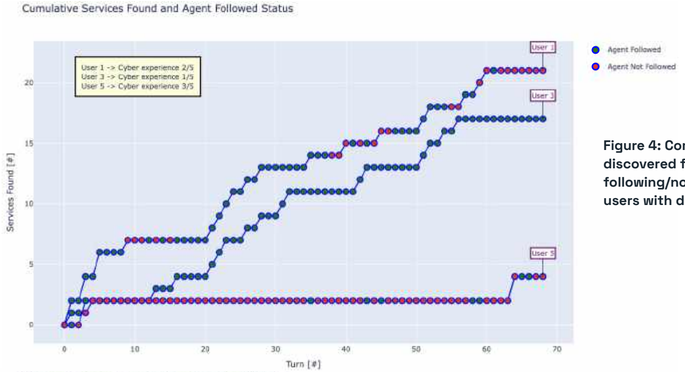

Figure 4: Comparison of the number of services discovered for instances of the human in the HMTs following/not following the agent recommendations for users with different cyber experience.

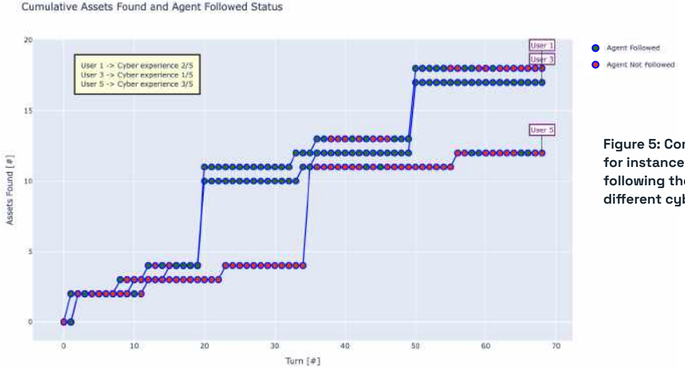

Figure 5: Comparison of the number of assets discovered for instances of the human in the HMTs following/not following the agent recommendations for users with different cyber experience.

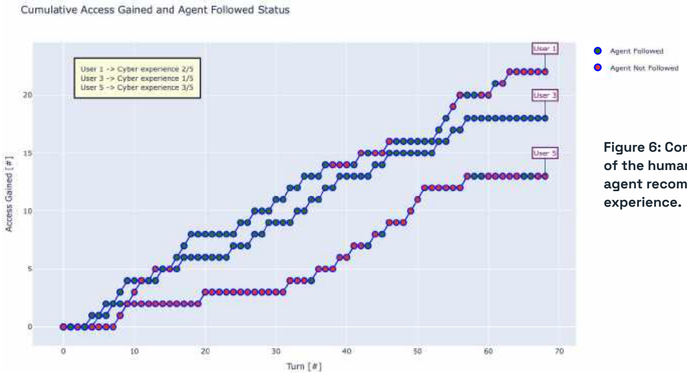

Figure 6: Comparison of the access gained for instances of the human in the HMTs following/not following the agent recommendations for users with different cyber experience.

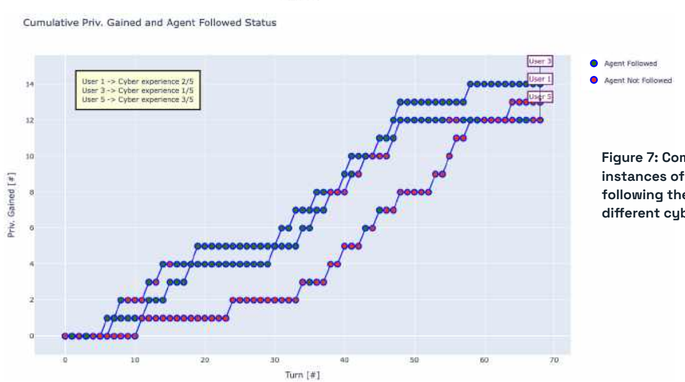

Figure 7: Comparison of the privilege gained for instances of the human in the HMTs following/not following the agent recommendations for users with different cyber experience.

Human-Machine Teaming: The Future of Enhanced Decision-Making

Challenges are becoming increasingly complex and high-stakes decisions are becoming the norm. Solely relying on either human intuition or artificial intelligence is no longer sufficient. Human-machine teaming represents a paradigm shift where the strengths of both humans and AI are combined to create a more powerful decision-making entity than either could achieve alone.

What is Human-Machine Teaming?

Human-machine teaming is the collaborative interaction between humans and artificial intelligence systems, working together to achieve shared goals. This partnership leverages:

Human Expertise and Intuition

Contextual Understanding: Humans possess deep contextual awareness and can interpret nuances that machines might overlook.

Ethical Judgement: The ability to apply ethical considerations and understand societal impacts.

Creative Problem-Solving: Humans excel at thinking outside the box and generating innovative solutions.

Machine Strengths

Data Processing Power: AI can analyse vast amounts of data at speeds unattainable by humans.

Pattern Recognition: Machines can identify complex patterns and correlations within data sets.

Consistency and Reliability: AI systems provide consistent outputs without fatigue or emotional bias.

Why Human-Machine Teaming is the Way Forward

Adopting human-machine teaming is not just a technological upgrade; it's a strategic move toward future-proofing your organisation. As industries evolve and challenges become more intricate, organisations that harness the synergy between human intellect and machine efficiency will lead the way.

Enhanced Decision Accuracy and Efficiency

Complementary Strengths: By combining human judgement with AI's analytical capabilities, organisations can make more accurate and timely decisions.

Reduced Errors: AI can handle repetitive tasks and complex calculations, minimising human errors, while humans oversee and interpret AI outputs for contextual relevance.

Improved Explainability and Trust

Transparent Collaboration: When AI systems provide explainable insights, humans can understand and trust the recommendations, leading to better adoption and more confident decision-making.

Accountability: Human oversight ensures that AI decisions align with ethical standards and organisational values.

Addressing Talent Shortages and Resource Constraints

Scaling Expertise: Human-machine teaming allows organisations to extend the reach of their existing talent, making specialised knowledge more widely accessible across the organisation.

Efficient Training: AI systems can encapsulate expert knowledge, reducing the time and expense required to train new staff to the same level of expertise.

Adaptive and Resilient Operations

Continuous Learning: AI systems can learn from human feedback, improving over time and adapting to new challenges.

Resilience to Change: The combination of human adaptability and AI's data-driven insights allows organisations to respond swiftly to evolving circumstances.

Strategic Allocation of Resources: With AI handling routine tasks, human experts can focus on complex, high-impact issues that require critical thinking and strategic insight.

Tulpa's Story

Our Mission

To empower organisations with trustworthy AI that amplifies human potential and organisational excellence

Our Vision

A future where human and AI collaboration drives innovation and solves the most critical challenges.

Our Values

Integrity: Commitment to ethical AI practices.

Innovation: Constantly pushing the boundaries of technology.

Collaboration: Working closely with clients and partners

The Meaning of "Tulpa"

"Tulpa" is a term originating from Tibetan Buddhism, referring to a being or object created through spiritual or mental powers—a manifestation brought into existence by the focused intention and concentration of the mind. Essentially, a tulpa is a thought-form that becomes a quasi-physical entity through the power of the mind.

At Tulpa, we embrace this concept by enabling organisations to create customised, transparent, and explainable AI agents that act as extensions of human expertise. Just as a tulpa is shaped by its creator's thoughts and intentions, our AI agents are formed and guided by your organisation's knowledge and needs. This synergy enhances human-machine collaboration, allowing you to scale intellectual capital, address complex challenges, and foster innovation through AI that is both ethical and trustworthy.

Our Team

We are a multidisciplinary team of mathematicians, scientists, engineers, and psychologists drawn from world-class institutions.

Rich Carter
PhD FBCS FRSA
CEO

International profile in national security and artificial intelligence.

Senior Visiting Fellow at
The Alan Turing Institute specialising in human-machine teaming.

Joe Ham
BA MSc
CPO

Leader of high-performing AI teams. Expertise in the design + implementation of data solutions.

Substantive knowledge of national security, telecoms and engineering.

Richard Porter
PhD
COO

Successful track record in large programme delivery and innovation in defence and government.

Cybersecurity Research Fellow at the University of Warwick.