

AI-Driven Cognitive Modelling applied to Agent-vs-Agent Simulation for Decision Support in Defence Wargaming

*Diego Tornese, Annalisa Letizia, Gianluca
Cerio, Gianluca Toscano*



Agenda

Scope and Structure of the Presentation

 01 Context and operational need

 02 Proposed architecture, at high level

 03 Stage 1 and Stage 2 in detail

 04 Scientific contribution and applicability

 05 Limitations and future work



Context and Need

Limited Behavioural Differentiation in Synthetic Actors

Computer Assisted Exercises and Agent-vs-Agent environments depend on synthetic actors whose behaviour approximates the human roles they represent. Their behavioural parameters are defined manually by subject matter experts, scenario by scenario. The practice does not scale across exercises, is difficult to keep internally consistent, and offers little basis for justifying why a given actor behaves as it does.

Current practice



One hand-authored parameter set is reused across roles, so distinct actors behave alike.

Requirement



Behavioural models derived from data, distinguishable from one another and individually justifiable.



Context and Need

COA Robustness Requires Multiple Distinguishable Behavioural Models

Modelling and Simulation activities within Multi-Domain Operations require decision-making agents that reflect a plausible diversity of human behaviour rather than a single averaged response pattern. Assessing a plan against different adversarial or partner behaviours presupposes multiple distinguishable behavioural models. Introducing a language model into the loop does not by itself produce that differentiation.



Risk-averse model

Low tolerance for risk,
with preference for
preserving resources over
positional gain.



Cautious model

Slower preferred pacing,
committing only once the
interaction state is
confirmed.



Aggressive model

Higher risk tolerance,
accepting exposure in
exchange for tempo and
initiative.



The same course of action, tested against each of them





Selection of the Agentic Architectural Paradigm

Narrowly scoped responsibility as a condition of inspectability

The paradigm is selected against two requirements: behavioural variety and inspectability. The adopted solution is a pipeline in which each LLM-based agent is responsible for a single well-defined transformation, rather than one general-purpose agent performing all functions.



Single LLM call

No separation of responsibility: the resulting behaviour cannot be inspected or corrected at a specific point of failure.



Rule-based system

Rigid: it captures anticipated cases well but does not model the behavioural variety in question.



Multi-agent pipeline

Narrowly scoped, individually inspectable stages, adopted for explainability and modularity rather than for novelty.



Architectural Key Features

Three Classes of Components in the Demonstrator

Three classes of components make the architecture more than a sequence of prompts: scoped AI entities, explicit contracts between stages, and a point of truth external to natural language.



Narrowly scoped LLM agents

One well-defined transformation each: validation, enrichment, profiling, strategic planning, tactical evaluation.



Structured data contracts

A predefined, machine-checkable format at every stage boundary, so independently developed components compose without ambiguity.



Lightweight deterministic rule engine

Anchors the numerical evaluation of candidate actions to a verifiable computation rather than to the numerical judgment of a model.



Sandbox environment selection

Structural portability distinguished from empirical validity

Why not chess or Go

Decades of accumulated theory, opening repertoires and engine-derived heuristics mean that a strong policy is only loosely coupled to the profile that generated it.

An abstract game with a limited body of prior study still demands genuine strategic and tactical reasoning, but offers fewer pre-existing shortcuts.

Scope of the current work

A methodological demonstrator: it establishes the feasibility of the cascade end to end and exposes the assumptions, interfaces and failure modes that extension toward C2-oriented environments would need to address.

It is not yet a validated decision-support product. Structural portability of the pattern is not empirical validity of the behaviour it produces.



Stage 1: Human Cognitive-Behavioural Modelling

From information about an individual to a natural-language behavioural style

The first stage operates on unstructured or semi-structured information about an individual and produces four outputs of increasing specificity. The profile is the domain-agnostic core artefact of the stage.



Intake and validation

A data-quality gate on role, seniority, sector, organisational context and stated objectives.



Enrichment

A narrative in fixed categories, with identifying elements replaced by neutral placeholders.



Profiling

A structured cognitive-behavioural profile, including decision-making style and identifiable biases.



Game-style description

Free text on risk tolerance, preferred pacing and patterns to favour or avoid.



Transparency Mechanisms Embedded in Stage 1

Data-quality gate, anonymisation by design, rationale log with confidence

The explainability layer responds to the intended audience: reviewers and end users are experts in the operational domain rather than in AI methods and therefore require an inspectable justification for each derived trait.



Data-quality gate

Underspecified input is met with a structured request identifying what is missing, so the pipeline does not proceed on weak foundations.



Anonymisation by design

Identifying elements are replaced from the enrichment step onward, so all downstream processing operates on an anonymised representation.



Rationale log with confidence

Each stated preference is attributed to the profile trait it was derived from, with a qualitative confidence indicator for that attribution.



Stage 2: An Agent-vs-Agent Four-Step Cycle

From a natural-language description to a bounded, repeatable confrontation

The second stage begins from a natural-language game-style description and is functionally independent of how that description was produced. The cycle executes at every turn for each of two agents instantiated from distinct profiles, interacting in the same rule-bound environment under identical conditions.






Two Independent Axes in the Tactical Evaluation

Verifiable rule-based assessment and qualitative alignment, examined separately

Combining a verifiable rule-based assessment with a qualitative profile-alignment assessment allows the two sources of uncertainty, those inherent to rule-based domain modelling and those inherent to model inference, to be examined separately rather than conflated into a single opaque score.






Scientific Contribution

Structural Decoupling of Domain-Agnostic Inference from Domain-Specific Execution

The two stages are connected through a natural-language interface rather than a shared data structure. The cognitive modelling methodology is structurally independent of any particular downstream simulation, and the simulation stage constrains neither the provenance nor the internal representation of the description it receives.



The interpretation-strategy-tactics-selection cycle is structurally applicable to any rule-bound environment exposing available actions and an evolving state. Whether behaviour produced through the pattern remains faithful in a given new domain is a separate question.



Scientific Contribution

Explainability as First-Class Output and External Verification

Opacity of AI decision processes is a recognised impediment to trust in AI-supported wargaming. Explicit attribution makes the approach usable in human factors analysis, where the traits driving an agent's behaviour can be examined directly rather than inferred from observed play.



Explainability as first-class output

Each behavioural trait carries an explicit attribution to the profile element it was derived from and a confidence indicator, rendering the full derivation chain inspectable by a domain expert.



External deterministic anchor



An external component anchors the quantitative half of the tactical evaluation, so the two sources of uncertainty are examined separately rather than conflated.



Applications

COA Analysis Across a Range of Plausible Counterpart Behaviours

The architecture supports the generation of multiple synthetic actors with distinguishable, traceable behavioural profiles, allowing the robustness of a plan to be examined across a range of plausible counterpart behaviours rather than against a single assumed response.





Applications

From Bespoke Authoring to a Repeatable, Data-Grounded Process

For training and exercise preparation, the production of behavioural variants shifts from bespoke authoring toward a repeatable process grounded in available data about the roles represented, reducing the dependence on subject matter expert availability identified as a constraint in prior wargaming work.

	Bespoke authoring	Pipeline-driven derivation
 Production of variants	Authoring is repeated by hand for each scenario, constraining the number of variants realistically produced.	A repeatable process is executed on data already available about the roles represented.
 Grounding and consistency	Parameters rest on individual judgment, with limited traceability to underlying data.	Each trait is derived from source material and carries its own attribution and confidence.
 Role of the expert	Expert availability during development and execution is a recurring constraint.	The expert moves from authoring toward review and high-level supervision.



Limitations Internal to the Pipeline

Unverified rationale, cross-stage fidelity, unequal grounding of the two axes

Section 5 claims structural portability of the architecture; it does not claim that the behaviour produced through it generalises. Three limitations are internal to the pipeline itself.



Plausible account, not verified explanation

The rationale and confidence indicators are themselves model-generated: a structured aid to expert review, not a substitute for it.



Cross-stage information fidelity

Two independent transformations in sequence: divergence between intended and instantiated profile cannot be excluded without consistency checks, not yet implemented.



Unequal grounding of the two axes

The rule-based assessment is externally verifiable; alignment and strategic planning depend on model judgment, under a fixed rather than calibrated weighting.



Data and Governance: Two Orthogonal Categories of Risk

Profile representativeness and the obligations that hold independently of it

The two risks are deliberately kept apart: neither one causes the other, and neither is mitigated by addressing the other.



Profile representativeness

Epistemic risk

A profile derived from partial or openly available information reflects the biases and gaps of those sources. A confidently stated but inaccurate profile is more problematic than a visibly incomplete one, since it invites unwarranted trust in the resulting agent.



Obligations on identifiable individuals

Ethical-legal obligation


Anonymisation, data minimisation and clarity on consent and intended use are design constraints on the architecture itself, not downstream compliance concerns. The dual-use sensitivity of profiling in a defence context warrants explicit governance.



From Sandbox to Operational Deployment

A Precondition, Not Evidence

Behaviour observed within the sandbox does not transfer automatically to higher-fidelity C2 environments, where the action space and its consequences are substantially larger and less well defined. Producing distinguishable, internally consistent agents in a constrained setting is a necessary precondition for such extension, not evidence of it.



The architecture targets exercise and analysis environments rather than live decision advice. Any decision-support use should retain human-in-the-loop review of both the generated profiles and the resulting behavioural models.



Future Improvements

Development Directions in Order of Logical Priority

Each direction addresses one of the limitations stated before. These are priorities in logical order rather than a roadmap.



Cross-stage consistency checks

Checks between the two sequential transformations, to reduce the risk of divergence between intended and instantiated profile.



Calibration of tactical weights

Calibration of the tactical evaluation weights against domain expert judgment, in place of the fixed weighting of the current implementation.



Extension of validation

Extension beyond the current sandbox toward environments closer to a realistic C2 context, with the same methodological caution.



Conclusions

Structural Portability and Empirical Validity: Distinct Claims

The methodological contribution rests on the separation between the two stages, on the treatment of explainability as an inspectable first-class output, and on the use of an externally verifiable rule-based component to anchor part of the tactical evaluation.



Structural portability of the pattern

Follows from the design: the natural-language interface preserves modularity between a domain-agnostic inference stage and a domain-specific execution stage.



Empirical validity of the behaviour

Remains open, bounded by profile representativeness, cross-stage information fidelity, the grounding of the strategic and alignment components, and the gap toward higher-fidelity decision environments.

Thank you for attention



Disclaimer

This report is prepared by Teoresi S.p.A. ("Teoresi" or the "Company") for the exclusive use and interest of the recipient.

The contents of this report may be subject to changes and even substantial updates.

This report as a whole and each of its contents must be considered as strictly confidential and any disclosure and / or communication to third parties other than the recipient is prohibited. This report may in no way be used, taken into consideration or mentioned by third parties other than the recipient.

No binding assessments of adequacy or feasibility are expressed in this report. No commitment, guarantee or other form of expressed or implied assurance is given by Teoresi regarding the accuracy or completeness of the information and / or opinions contained in this report.

To the maximum extent permitted by law, any liability of Teoresi, its shareholders, directors, managers, employees and consultants relating to or arising from such information and / or opinions is excluded. This content may contain legally protected information and is intended solely for the specified use.



www.teoresigroup.com

