

## FROM OVER-RELIANCE TO SMART INTEGRATION: USING LARGE-LANGUAGE MODELS AS TRANSLATORS BETWEEN SPECIALIZED MODELING AND SIMULATION TOOLS

Philippe J. Giabbanelli<sup>1</sup>, John Beverley<sup>2</sup>, Istvan David<sup>3</sup>, and Andreas Tolk<sup>4</sup>

<sup>1</sup>Virginia Modeling, Analysis, and Simulation Center (VMASC), Old Dominion University, VA, USA

<sup>2</sup>Dept. of Philosophy, University at Buffalo, Buffalo, NY, USA

<sup>3</sup>McMaster Centre for Software Certification (McSCert), McMaster University, Hamilton, CANADA

<sup>4</sup>The MITRE Corporation, Charlottesville, VA, USA

### ABSTRACT

Large Language Models (LLMs) offer transformative potential for Modeling & Simulation (M&S) through natural language interfaces that simplify workflows. However, over-reliance risks compromising quality due to ambiguities, logical shortcuts, and hallucinations. This paper advocates integrating LLMs as middleware or translators between specialized tools to mitigate complexity in M&S tasks. Acting as translators, LLMs can enhance interoperability across multi-formalism, multi-semantics, and multi-paradigm systems. We address two key challenges: identifying appropriate languages and tools for modeling and simulation tasks, and developing efficient software architectures that integrate LLMs without performance bottlenecks. To this end, the paper explores LLM-mediated workflows, emphasizes structured tool integration, and recommends Low-Rank Adaptation-based architectures for efficient task-specific adaptations. This approach ensures LLMs complement rather than replace specialized tools, fostering high-quality, reliable M&S processes.

### 1 INTRODUCTION

Large Language Models (LLMs) offer the temptation of a one-stop-shop by handling seemingly any Modeling & Simulation (M&S) task, at the easily satisfiable condition of expressing the task through natural language and/or through images. This vastly simplifies the workflow of modelers, by forging the illusion that it is no longer necessary to navigate different tools and data representations. As a result, LLMs often are adopted due to the perceived *convenience* of bypassing the strenuous steps of de- and re-composing complex problems for modeling and analysis purposes. Such convenient shortcuts conceal accidental and essential complexity of the problem at hand and therefore, excessive reliance on LLMs may have detrimental effects for M&S research.

Research quality is negatively impacted when the LLM output is *suboptimal* relative to specialized tools or expertise. Modelers may deem the LLM outputs sufficiently accurate if they are not aware that a better, specialized solution exists. For example, the LLM may be asked to think about whether changes in the input of a simulation model would create a desired effect. Directly using a LLM may result in logical shortcuts or hallucinations. However, LLMs are proficient at translating such queries into symbolic representations (e.g., first order logic) and they can call a symbolic solver to guarantee logical coherence during reasoning (Lam et al. 2024).

In this paper, we focus on leveraging the convenience and simplicity of using LLMs without sacrificing quality or losing the benefits of existing investments in specialized tools. **We posit that LLMs should be used to mitigate accidental complexity, so that the human can focus on the essential complexity of the problem at hand.** Since accidental complexity in M&S research typically stems from multi-formalism, multi-semantics, multi-paradigm modeling, and the supporting heterogeneous toolchains, we propose that modelers should leverage LLMs as a “glue” or middleware “translator” between specialized tools and

expertise. That is, instead of aiming to perform any task solely *within* a LLM, the emphasis should be on placing LLMs *between* modeling and simulation systems to make them more easily interoperable.

The realization that LLMs need to be integrated with specialized tools is gaining traction in M&S research (Zhang et al. 2024) and echoes a broader interest in positioning LLMs as middleware instead of the sole solution (Lehmann 2024; Calvet et al. 2024). In this paper, our main contribution is to address the two core questions raised by the opportunities and challenges of integrating LLMs with specialized tools to accomplish modeling and simulation tasks:

- Q1** If the LLM is to act as a translator, then which *language* should it translate into and which *tool* should it call depending on the M&S task?
- Q2** Which software *architecture* could efficiently integrate LLM as a middleware to support the interoperability of specialized M&S tools without creating performance bottlenecks?

The remainder of this paper is organized as follows. Section 2 succinctly demonstrates across fields that LLMs have been used for their convenience, occasionally at the expense of quality. This evidence base, across applications, motivates our paper by showing that the understandable need for convenience from M&S practitioners should not result in an exclusive reliance on LLMs, but rather by employing LLMs in well-defined roles such as a middleware between specialized tools. In Section 3, we identify languages and tools to perform modeling tasks (**Q1**), ranging from building models (e.g., by merging or from a corpus) to validating and explaining their structure. Similarly, Section 4 covers the languages and tools for simulation tasks (**Q1**). In Section 5, we review the architectures available (**Q2**) and propose one solution that has yet to be employed for M&S but has a strong potential for high performances in both training and execution. Finally, Section 6 examines key issues in using LLMs as middleware, including *trust*, the *reciprocal benefits* between LLMs and specialized tools, and the need for extensive *benchmarks*.

## 2 BACKGROUND

### 2.1 Evidence for the Convenient Use of LLMs at the Detriment of Quality Across Applications

An already infamous Forbes report from 2023 illustrates that LLMs are used as a ‘Swiss-Army Knife’ for diverse tasks because they are easy to use, despite the existence of tools that might produce higher-quality outputs. This report describes a lawyer who used ChatGPT to draft and submit a defense brief that contained fictitious case precedents (Novak 2023). While the efficacy and accuracy of LLMs used in the legal domain has significantly increased since 2023, the precision demanded in legal discourse within and across languages raises barriers just as significant (Padiu et al. 2024). As a result, the ease of using LLMs in such cases is outweighed by the risk of legal missteps that are more likely to be avoided by legal experts.

In the clinical domain, Brown et al. (2024) compared traditional machine learning (ML) methods with LLMs (GPT-3.5 and GPT-4) for prediction tasks using electronic health records. Their study found that while LLMs are attractive for their simplicity and ease of deployment, traditional ML approaches—such as gradient-boosted trees—achieved significantly higher predictive performance and calibration. From another direction, Van Veen et al. (2024) reported that although LLMs can produce clinical summaries that, in many cases, are judged comparable to those generated by experts, the models still exhibit safety concerns such as hallucinatory content, which specialized summarization systems or curation methods employed by experts are better at avoiding. This is not to say continued research will not ultimately address LLM limitations in this domain or others. Rather, it is to highlight that despite the impressive applications of LLMs in this area (Yu et al. 2024; Zheng et al. 2024), general applicability should be viewed with caution.

In the area of code generation, Tang (2024) and Basic and Giaretta (2024) highlight that LLM-generated code, while easy to obtain, often contains bugs and security vulnerabilities that require manual, and often time-consuming, efforts to untangle. In a reverse engineering context, Pordanesh and Tan (2024) demonstrated that GPT-4 can offer broad insights when analyzing decompiled code, but it frequently struggles with the intricacies of complex code structures—leading to superficial analyses compared to

established reverse engineering tools, such as the disassembler and decompiled Binary Ninja and the Ghidra suite.

## **2.2 Can LLMs address the Elusiveness of Simulation Interoperability**

Simulation interoperability has been pursued through various research and engineering efforts for over three decades (Tolk 2024). Many approaches have fallen short because they treated simulation systems as mere software solutions. In reality, simulations implement models that are developed for specific purposes: to address particular research questions or support distinct tasks. This means that representations of simulated entities are rarely identical across different systems – their scope and structure naturally differ. This underlying model diversity is the fundamental reason why simulation interoperability remains elusive: technical standards that interconnect simulations cannot resolve conceptual differences between models.

LLMs offer promising potential to address this challenge. They can help bridge both referential differences (what information a simulation provides) and methodological differences (how this information is structured). Frydenlund et al. (2024) demonstrated using the same model to create simulations employing different paradigms, achieving varying levels of success. Jackson and Rolf (2023) utilized LLMs to generate natural language descriptions from simulation data.

Evidence suggests LLMs have strong translation capabilities. OpenAI’s GPT-4 has demonstrated ability to translate between numerous languages according to their technical documentation and public demonstrations, while Google has progressively integrated LLM technology into their translation services, beginning with Neural Machine Translation systems and evolving to include models like T5 and PaLM. Although neither company fully discloses their production architecture details, their capabilities suggest that if LLMs can effectively translate between different linguistic structures and grammars, they should increasingly contribute to addressing the resolution, paradigm, and scope differences that have historically challenged simulation interoperability.

A new research domain are execution-aware LLMs (Di Menna et al. 2025). While the current articles focusing predominantly on source code generation that does not only consider the static source code but also considers how this source code is executed and what effect such executions have, these ideas clearly point beyond the usual semantic understanding of an artifact towards a pragmatic understanding of how this artifact is used. This new domain has the potential to go beyond source code engineering and address the higher levels of interoperability as requested in (Tolk 2024). However, the research is still in its infancy and the authors are not aware of any simulation-specific applications thereof.

## **3 LLMS FOR MODELING**

The flexibility of LLMs often comes at the cost of output precision, formal reasoning capabilities, and domain-specific accuracy. Modeling is a task driven activity to select the scope of the model, as well as the appropriate abstraction level that results in a specification that can be implemented. As such, it involves tasks that range from constructing model specifications (e.g., by merging existing models or extracting models from data) to verifying and explaining model structures. Many such tasks require structured data representations, formal logic, and domain-aware processing, which specialized tools handle more effectively than LLMs. Table 1 summarizes key modeling tasks and examples of languages and tools that directly support each task out-of-the-box, without requiring custom integration code.

Table 2 provides an overview of contemporary tools that directly support major modeling tasks without requiring custom integration. These tools are designed to reduce manual effort and improve modeling reliability across M&S domains. For example, OntoAligner and CoMerger support semi-automated merging of OWL ontologies with built-in alignment and consistency checks, ideal for collaborative modeling or integrating heterogeneous sources in domains like defense and logistics. Legacy tools like the PROMPT plugin for Protégé remain useful for smaller or more controlled merge tasks. AutoMap enables rapid model generation from unstructured text by extracting and visualizing concept networks. This is particularly

Table 1: Modeling tasks and examples of languages/tools with native out-of-the-box support.

Modeling Task	Example Language/Notation	Specialized Tool (Direct Support)
Merging Models	OWL ontologies	<i>OntoAligner</i> , <i>CoMerger</i> , <i>Protégé</i> with PROMPT plugin (Horridge and Bechhofer 2004)
Building Model from Corpus	Network text analysis	<i>AutoMap</i> – extracts and visualizes concept networks from text (Carley 2013)
Validating Model Structure	First-order logic (Alloy), OWL, UML	<i>Alloy Analyzer</i> (Jackson 2006), <i>Hermit</i> reasoner in <i>Protégé</i> (Glimm et al. 2014), <i>UModel</i>
Explaining/Documenting Models	Structured model descriptions	<i>Enterprise Architect</i> from Sparx, <i>Protégé</i> OWL2NL plugin
Data/Schema Modeling	ER diagrams, UML	<i>ER/Studio</i> , <i>Erwin Data Modeler</i> (Quest Software, Inc. 2023)
System Modeling and Simulation	Modelica	<i>Wolfram SystemModeler</i>

Note: OWL is the Web Ontology Language, UML is the Unified Modeling Language.

useful in healthcare or intelligence analysis, where narrative data, such as that reflecting patient reports or incident logs, can be converted directly into influence diagrams or causal maps.

Table 2: LLM-mediated modeling workflows: tools, roles, and integration patterns.

Modeling Subtask	Recommended Tool(s)	LLM Role	Integration Notes
Drafting initial model structure from natural language	GPT-4 + <i>UModel</i> , <i>ER/Studio</i> , <i>Protégé</i>	Translates textual input to structured representation (UML, OWL, ER)	Output should target tool-compatible syntax (e.g., XMI, RDF/XML); tool validates structure
Ontology/model alignment	<i>OntoAligner</i> , <i>CoMerger</i> , <i>Protégé</i> PROMPT	Suggests initial mappings between terms; explains inconsistencies	Specialized tool performs alignment, consistency check; LLM mediates error resolution
Data schema mapping and reconciliation	openrefine.org, <i>Neo4j</i> , <i>spaCy</i>	Translates error logs; proposes entity matches	LLM can explain match confidence, suggest reconciliation rules
Model verification and logic enforcement	<i>Alloy Analyzer</i> (Jackson 2006), <i>Hermit</i> Reasoner (Glimm et al. 2014), SHACL validators	Converts informal constraints to formal rules; explains violations	LLM provides human-readable feedback on validation results; tools execute logic checks
Process modeling and simulation specification	<i>Wolfram SystemModeler</i> , MATLAB/Simulink, Cameo	Converts scenarios into system dynamics / process models	LLM prompts must be aligned with domain-specific syntax; tool ensures compile-ready output
Business rule formalization	<i>Erwin</i> (Quest Software, Inc. 2023), <i>Enterprise Architect</i> from Sparx	Converts high-level requirements to constraint logic	Tool checks completeness, syntax; LLM supports traceability explanation
Documentation and model explanation	OWL2NL, GPT-4 + LangChain (LangChain Contributors and Harrison Chase 2023)	Verbalizes formal models; generates stakeholder reports	LLM expands sparse documentation; tool outputs serve as structured input
Workflow orchestration (multi-tool)	GPT-4 + Toolchain (via LangChain and API)	Dispatches subtasks to appropriate tools; sequences steps	Requires modular design; tools expose APIs or plugin interfaces

Note: RDF is the Resource Description Framework, SHACL is the Shapes Constraint Language for describing RDF graphs.

Model validation is natively supported by tools such as the Alloy Analyzer (Jackson 2006), which provides bounded model checking for first-order logic specifications, and Hermit (Glimm et al. 2014), a

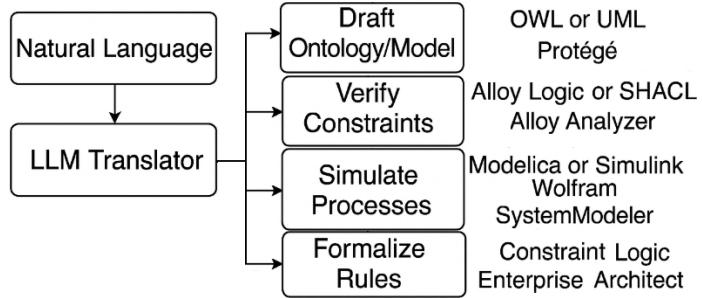


Figure 1: LLMs as Middleware for Modeling Tasks

widely-used reasoner for OWL ontologies. These tools help verify that structural constraints and logical rules are satisfied, reducing errors prior to simulation or implementation into broader workflows. Tools such as Sparx’s Enterprise Architect offer automatic report generation from structured models, while OWL2NL translates ontology axioms into readable natural language. These features support stakeholder communication, requirements documentation, and training. Additionally, tasks such as data and schema modeling are directly supported out-of-the-box by tools such as ER/Studio and Erwin Data Modeler (Quest Software, Inc. 2023), both of which allow analysts to design ER or UML diagrams, enforce schema constraints, and integrate with enterprise databases. Lastly, system modeling and simulation is supported natively by Wolfram System Modeler, which uses the Modelica language to represent complex multi-domain systems and provides built-in simulation capabilities—particularly useful in biomedical, engineering, and manufacturing applications.

In some cases, general-purpose modeling tools or libraries can accomplish identified modeling for M&S tasks with user-driven integration. Table 2 enumerates such scenarios. No single tool from the table can, in isolation, produce a validated, domain-specific model; modelers must orchestrate multiple tools in sequence to that end. Prior to LLMs, such orchestration was already common: one might apply an NLP pipeline—e.g., using a library like spaCy—to extract candidate entities or relationships from text, then employ a schema matching system or data cleaning tool, such as OpenRefine (openrefine.org), to reconcile terminology, and finally use an ontology editor or modeling environment to formalize the concepts. LLM-integrated workflows build upon the same division of labor, where an LLM takes on a ‘glue’ role: mediating between human language and the formalisms required by each specialized tool.

While many core modeling tasks are supported by specialized tools out-of-the-box, numerous other tasks remain difficult to execute without integrating multiple tools or bridging semantic gaps between representations. These include cases where natural language must be translated into formal structure, where informal business rules must be formalized, or where model components from heterogeneous systems must be aligned. Here, LLMs can serve as a middleware, translating, sequencing, and coordinating modeling subtasks (Figure 1). Table 2 presents LLM-mediated modeling workflows from initial model drafting to rule formalization and documentation, highlighting common subtasks, the target representation language(s) that the LLM must generate (e.g., OWL, UML), the specialized tool(s) capable of ingesting and executing that representation (e.g., Protégé), and the role of the LLM as a mediator between natural language and formal artifacts (e.g., translation, summarization).

The LLM serves not as a substitute for the formal tool, but as an intelligent bridge—converting human intent into machine-verifiable representations (Figure 1). The table operationalizes this division of labor, offering practical guidance on how to scope LLM-mediated translation workflows depending on the modeling need. This directly addresses **Q1** by establishing clear, task-specific routes from language to tool.

For example, when drafting a model from natural-language requirements, an LLM such as GPT-4 can be used to generate structured output in OWL, UML, or ER formats, which can then be validated and refined in tools like Protégé, UModel, or ER/Studio. Here, the LLM is a translator that outputs standards-compliant syntax – such as RDF/XML – enabling handoff to verification-capable environments. In more complex

tasks like ontology alignment or schema reconciliation, LLMs may again play a mediating role, revealing candidate terms for mappings, resolving ambiguity, or interpreting tool-generated error logs. Tools such as OntoAligner and OpenRefine remain central for alignment and cleanup, but LLM can offer contextual explanations and generate documentation. Related, when enforcing model constraints or validating logical coherence, LLMs may assist by transforming high-level rules into formal expressions suitable for execution in tools like Alloy Analyzer or HermiT. Here again, LLMs also help translate validation failures into actionable, human-readable documentation that reduces debugging time.

In process modeling and simulation workflows, LLMs may convert textual descriptions into system dynamics models in Modelica, block diagrams in Simulink, or behavior trees in Cameo. These conversions require structured prompting and domain-specific syntax, which the modeling tool ultimately compiles and validates. Overall, LLMs can orchestrate multi-step modeling workflows by dispatching tasks to different tools, coordinating formats, and sequencing outputs. This pattern is increasingly supported by frameworks like LangChain and toolchains that expose plugin APIs, allowing for modular LLM-driven pipelines.

The utility of LLMs as middleware is made more apparent when combining tasks. For example, if an LLM proposes a set of UML class definitions, a UML parser or modeling environment can be invoked to parse that proposal. If the parser returns errors (say, an undefined class reference or a syntax mistake in the UML), the error messages can be fed back to the LLM with a prompt to fix the issues. The LLM, now aware of where it went wrong, can adjust the class definitions accordingly. This iterative loop continues until the specialized tool confirms that the model is syntactically and structurally sound. Such a strategy leverages the LLM's strength in rapidly generating and regenerating content, while the tool's strictness guarantees eventual correctness. It is a form of algorithmic scaffolding, where the LLM is never left unguided for too long – the tool provides guardrails at each step. The distinction between Table 1 and Table 2 underlines a recommendation: whenever a task can be done with a purpose-built tool, prefer to integrate the LLM with that tool, rather than relying on the LLM alone or reinventing logic via scripting.

## 4 LLMS FOR SIMULATION

After we discussed in the last section how LLMs can support the modeling process better, we focus in this section on the simulation development process.

### 4.1 Requirements Analysis

In the earliest phase of the M&S lifecycle, artifacts are typically encoded in natural language (e.g., requirements and specifications), which opens some opportunities for LLMs. LLMs can be used, e.g., to *analyze project requirements* at a high level before commencing more resource-intensive formal specification and problem decomposition tasks. Such analyses can be used, e.g., to *recommend appropriate simulation paradigms, formalisms, and tools* for the project. LLMs can rely on very diverse contextual data to drive such recommendations, including the comprehensive analysis of a corpus describing a problem, available data, and the skill set at hand. Such use cases raise the question whether LLMs can be used to *bypass modeling* in the traditional requirements – modeling – simulation process. Finally, LLMs can be used for validating requirements, e.g., by *producing pseudocode* (Xu et al. 2024).

### 4.2 Source Code Engineering

The strong link between LLMs and software code is rooted in the well-recognized naturalness hypothesis of software (Hindle et al. 2016), which states that program code is a form of human communication bearing similar statistical properties to those of natural languages. *Source code generation* is among the key techniques in this lifecycle phase, and is has been a subject to thorough investigation. Gerstmayr et al. (2024) use LLMs to generate multi-body system dynamics models from natural language specifications. Miceli-Barone et al. (2023) translate user instructions to code for self-driving car simulations. Liu and Li (2024) generate executable simulation code from human language descriptions. (Jackson et al. 2024)

generate executable simulation for a logistic systems for inventory and process control. Wang et al. (2024) generate code for rich simulation environments and expert demonstrations for robot training. *Code translation and adaptation* between tools in simulation tool chains is another important case of source code engineering by LLMs. This case is gaining attention in research. For example, Shrestha and Csallner (2021) use a GPT model to generate graphical block-diagram models, such as Simulink models, and pre-process the generated models for rapid search. However, as shown by Pan et al. (2023) through their experiments with 1 700 code samples in C, C++, Go, Java, and Python, LLMs are not yet reliable enough for code translation, with the reported proportion of correct translations ranging between 2.1% and 47.3%.

### 4.3 Testing and V&V

In the testing and V&V lifecycle phase, LLMs are typically used for *generating test cases and test artifacts* from simulation models or code, and usage examples. Such avenues have been explored by Schäfer et al. (2024) who generate unit tests with a median statement coverage of above 70% and branch coverage above 52%. LLMs can also generate various tests artifacts for simulation testing. For example, Duvvuru et al. (2025) generate environmental configurations, test cases, and test properties (specific metrics, e.g., target detection and identification accuracy) for testing uncrewed aerial systems. Another application of LLMs is *verification*, when more rigorous evidence, e.g., formal proof is required to guarantee key system properties. For example, Hassan et al. (2024) use LLMs to generate formal specifications from natural language specifications, and to generate system invariants for a high-performance theorem prover (Z3) and summarize them in natural language.

### 4.4 Documentation

Documentation is a natural use case of LLMs, for the same reason code engineering is: due to the naturalness hypothesis of software (Hindle et al. 2016). Since source code exhibits the same statistical traits as human language, it is rather straightforward to employ LLMs for *generating human-readable explanations* of source code, simulation models, and simulation results at various levels of technical depth. An example is given by Jackson and Rolf (2023). Similar avenues have been explored, e.g., by Nam et al. (2024) to generate comprehensive documentation for complex simulations. A special form of documentation is *summarization* of source code and simulation models that aims to produce brief and target human-readable descriptions for experts. For example, Ahmed and Devanbu (2022) use GPT models with few-shot learning to summarize source code, and their evidence suggests that LLMs significantly surpass other state-of-the-art models for code summarization. Similarly, Khan and Uddin (2023) generate documentation using a GPT-3 based model, Codex, pre-trained on both natural and programming languages, and achieve a BiLingual Evaluation Understudy (BLEU) score over 0.2, a 11% improvement over state-of-the-art techniques.

## 5 ARCHITECTURES TO EFFICIENTLY INTEGRATE LLMS AS M&S MIDDLEWARE

A simple use of LLMs is to prompt them with a task and directly use their answer (Figure 2-A), possibly through a follow-up prompt. However, as the use of LLMs within M&S matured, we observe a shift towards training and refining LLMs for specific tasks and/or integrating LLMs with other tools through a workflow. As a case in point, researchers first used GPT through prompts to turn a prose-based narrative into simulation code (Frydenlund et al. 2024), and later Martínez et al. (2024) fine-tuned the system (few-shot prompting) and connected it to a database (retrieval-augmented generation). Newer architectures reflect these changes in practices (Figure 2-B), for instance when a modeling tool such as a class diagram editor internally prompts the LLM to recommend classes or attributes (Chaabén et al. 2024). This section thus examines approaches that employ LLMs as a middleware, emphasizing support for M&S tasks.

There is also the potential to use LLMs as translators between two different representations of knowledge. While several studies bridged natural language and SQL commands for a database (Zhou et al. 2024; Li et al. 2024), mediating between different schema has received less attention. Developing these methods

further can lead to the automatic configuration of data mediation services between different structures of knowledge representation while ensuring their conceptual alignment (Tolk 2024). Calvet et al. (2024) compared two architectures to use LLMs as data conversion tools, possibly within a single system. They noted that prior works performed a *direct* conversion: users provide one data format to the LLM, which must generate data in the target format. The authors introduced an *indirect* step by generating Python code and calling it to execute the conversion (Figure 2-C). The advantage is subtle yet important: a direct conversion may work once but there is no guarantee that it will work next time, due to stochasticity in the LLMs. In contrast, if a Python code is correctly generated once, then it can be used to correctly perform all data conversions. The authors also showed that LLM-based translations can be challenging, as the relatively simple data schema used in the study could not always be fully converted and some LLMs were never able to generate Python code to perform the conversion. The takeways are twofold: (i) transforming a problem instance into a target domain language and executing it can improve system reliability compared to performing all tasks via the LLM, but (ii) some transformations can require significant engineering efforts.

Lehmann (2024) proposed an architecture consisting of two LLMs, one for the ‘consumer’ and another for the ‘provider’ (Figure 2-D). His vision was unique in using natural text as intermediate between LLMs: the consumer LLM does not *directly* pass calls onto the target system or the provider LLM; rather, it translates the call into text, which is translated back into a call by the provider LLM. As a result, this architecture involves four translations. The author noted the cumulative risks posed by these multiple translations, “adding multiple components into the communication between consumer and provider also introduces potential points of failure”, thus we do not retain this architecture for M&S applications.

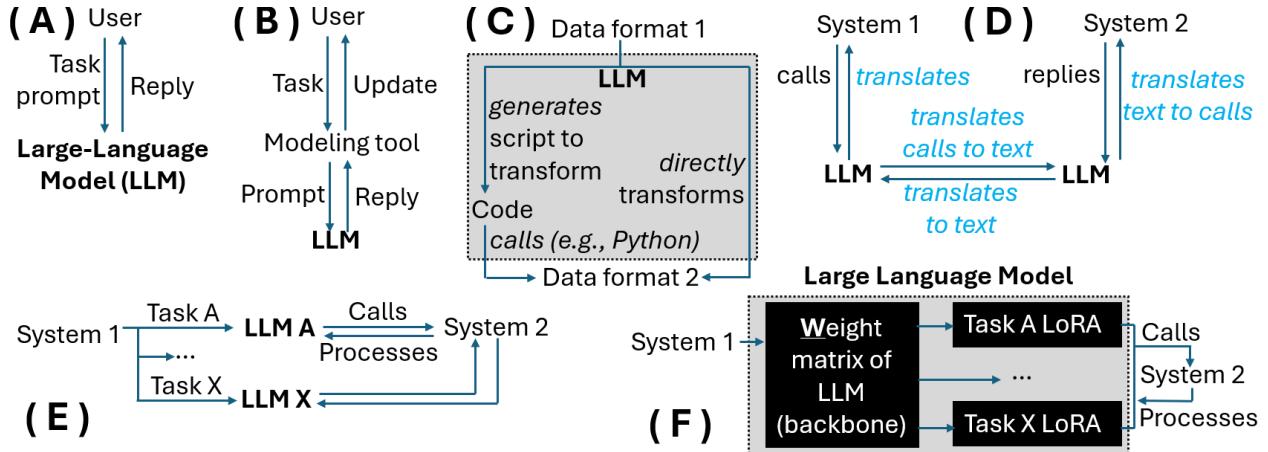


Figure 2: The simplest use of LLMs is to set a task and use the reply either directly (A) or via a software (B), but a LLM may generate a script instead of directly solving a task itself (C) or a set of LLMs may support system interoperability (D–E). In contrast with these prior frameworks, we recommend handling for M&S tasks with one LLM to improve computational efficiency and maintainability (F).

The ‘pretrain-then-finetune’ paradigm results in a multiplicity of specialized LLMs, each with (sometimes minor) differences in the weights of their neural network. As a result, if a modeler switches from task A to task B, they need to unload the LLM for task A and load the LLM for task B (Figure 2-E). In application settings where data is sensitive and computations must be performed locally (e.g., enterprise modeling, healthcare, defense), repeatedly unloading and loading a locally deployed LLM becomes a performance bottleneck as it is taxing on memory and takes time. In contrast, a parameter-efficient fine-tuning method such as Low-Rank Adaptation (LoRA) introduces small (low-rank) matrices into the weight updates of a pre-trained model, rather than fine-tuning all parameters (Yang et al. 2024). As a result, the bulky LLM is loaded only once, and light-weight matrices build on this shared backbone to support specific tasks (Figure 2-F). This avoids the taxing loading-and-unloading steps for memory management (Chen 2024),

and it also improves the training process since training a LoRA is significantly faster and more efficient than fully fine-tuning a large neural network – computations can be further improved through approximated LoRA (Chen et al. 2025). Several systems can manage LoRA adapters on a shared backbone and route user requests along different task-specific LoRAs, from the early *Punica* (Chen et al. 2024) to the recent *dLoRA* (Dou et al. 2024).

We note that this research area is still in its infancy, thus there is ongoing work on integrating adapters for optimal performance. In particular, Tian et al. (2025) proposed creating a ‘middleman adapter’ (or smart routing system) that identifies a task from the user input and subsequently matches the task with the right LoRA. It is thus *essential for the M&S field to develop high-quality specialized LoRAs and automatically map them onto tasks*. The literature on M&S building blocks suggests that it is preferable to identify the tasks faced by modelers and then create the corresponding LoRAs (Cheng et al. 2023; Schroeder et al. 2022), rather than creating specialized systems and then trying to map them to user needs. We thus recommend a comprehensive analysis of modeling tasks (i.e., user needs gathering), culminating in the generation of LoRAs covering several orthogonal needs. We posit that such a bottom-up process would also facilitate the alignment of user prompts and their correct routing through a LoRA.

## 6 DISCUSSION

Trusting LLMs as translators between systems depends on the definition of ‘trust’ and the specific context in which they are used. If trust is defined as *absolute reliability* with zero distortion or hallucination, then LLMs cannot be fully trusted, as they can introduce subtle errors or fabrications. However, trust can be framed in a *comparative sense* by evaluating the distortion caused by LLMs against existing solutions. For instance, Yoon et al. (2024) converted data back-and-forth between two formats, with an LLM and with traditional rule-based methods. In their context of health data, the LLM distorted and lost data less than the rule-based methods currently in use. In addition to measure the amount of errors, it is useful to understand the nature of these errors. Errors do not all have the same consequences and translators do not all have the same predictability: if a rule-based system tends to misclassify certain types of entries then we *know* the errors, whereas hallucinations from LLMs may not be as predictable. Ultimately, trust should not be considered as binary but as a balance of accuracy, predictability, and the ability to mitigate errors through oversight and verification.

This paper is an important step in identifying potential tools and languages for each step encountered in the M&S process, but these choices have an impact. Lam et al. (2024) reported a near 50% performance variation when the LLM translated between different representations and employed various tools to determine whether a conclusion followed from a set of premises. It is thus necessary to extend emerging M&S benchmarks (Giabbanelli et al. 2025) to account for the choice of tools and representations, such that modelers identify the best option available in their context. Designing benchmarks is challenging, as even small elements such as fictional character names may impact the LLM’s performance (Saparov et al. 2023).

In addition, training LLMs to support multiple translations (e.g., via plethora of LoRAs) may become cumbersome, particularly if we seek to certify these translations or maintain their accuracy in the face of evolving language specifications. Rather than envisioning that all translations *directly* produce the language of interest, it may be more pragmatic to consider *indirect* or multi-hops translations. That is, the M&S community may maintain a core set of translations (e.g., if the user’s query calls for first-order logic then use *Prover9* syntax) and leave it to users if they need minor adjustments afterward (e.g., turn *Prover9* into the related *Pyke* syntax to use a different solver). As previously noted, supporting interoperability by including indirect translations may lead to different sequences (e.g., turn task description into *Z* then *C* or *A, B* then *C*), which may have different computational costs or accuracies (Schuerkamp et al. 2023). Such problems have been studied in multi-paradigm modeling (MPM) (Vangheluwe et al. 2002) using *formalism-transformation graphs* (FTGs) that represent formalisms as nodes of a graph and transformations as edges (Challenger et al. 2020). A process model that instantiates formalisms as models, and transformations as activities, allows for the thorough investigation of complex transformation chains (Lúcio et al. 2013).

An approach worth following is the combination of ontological methods and LLMs, such as exemplified in (Allemang and Sequeda 2024). The underlying idea is to use ontologies as the formal specification of the conceptualization of the information to be shared as well as its structure. LLM expression must fall into this specification to be producible for the specified sending systems and vice versa to be understandable by the receiving system. Using the ontological specification of sender and receiver avoids the acceptance of hallucinations outside of the valid range of scope and structure while benefiting from the convenience of using the LLM. As demonstrated by Jackson and Rolf (2023), LLMs can describe even simulation systems in natural language, so that the description of interfaces and their constraints should be supported similarly. However, this approach is still in its early stages.

While our paper focused on the use of LLMs as middleware, it does not mean that LLMs merely act as messengers that translate requests between users and specialized tools. On the one hand, the tools can improve the LLM by amplifying its capabilities (e.g., a solver ensures that reasoning is coherent) or giving it feedback (e.g., errors reported by tools can be integrated in error-solving mechanisms as in Pan et al. 2023). On the other hand, *the LLM may help the tools*. For instance, the user may have a query in which some conditions are so obvious that they are not explicitly stated. A symbolic solver by itself may struggle when it expects a full chain of logic, but the LLM may provide the missing information. The architectures envisioned here are thus a simplification of the exciting potential mutual benefits that remain to be explored between LLMs and tools.

## REFERENCES

Ahmed, T. and P. T. Devanbu. 2022, October 10–14. “Few-shot Training LLMs for Project-Specific Code-Summarization”. In *Proc. 37th IEEE/ACM Int. Conf. Automated Software Engineering*, 177:1–177:5. Rochester, MI, USA.

Allemang, D. and J. Sequeda. 2024. “Increasing the Accuracy of LLM Question-Answering Systems with Ontologies”. In *The Semantic Web – ISWC 2024: 23rd International Semantic Web Conference, Baltimore, MD, USA, November 11–15, 2024, Proceedings, Part III*, 324–339: Springer-Verlag.

Basic, E. and A. Giaretta. 2024. “Large Language Models and Code Security: A Systematic Literature Review”. arXiv preprint arXiv:2412.15004.

Brown, K. et al. 2024. “Not the Models You Are Looking For: Traditional ML Outperforms LLMs in Clinical Prediction Tasks”. *medRxiv*:2024–12.

Calvet, E., R. Falcão, and J. Ginkel. 2024. “Can LLMs Make Systems Interoperate at Runtime? An Exploratory Study”. *SSRN*.

Carley, K. M. 2013. *AutoMap User’s Guide: Dynamic Network Analysis Tool*. Concept network extraction from text.

Chaabén, M. B., L. Burgueño, I. David, and H. Sahraoui. 2024. “On the Utility of Domain Modeling Assistance with Large Language Models”. arXiv preprint arXiv:2410.12577.

Challenger, M., K. Vanherpen, J. Denil, and H. Vangheluwe. 2020. *FTG+PM: Describing Engineering Processes in Multi-Paradigm Modelling*, 259–271. Springer.

Chen, G., Y. He, Y. Hu, K. Yuan and B. Yuan. 2025. “CE-LoRA: Computation-Efficient LoRA Fine-Tuning for Language Models”. arXiv preprint arXiv:2502.01378.

Chen, J. 2024. “Comparative Analysis and Optimization of LoRA Adapter Co-serving for Large Language Models”. In *Proceedings of the 25th International Middleware Conference: Demos, Posters and Doctoral Symposium*, 27–28.

Chen, L., Z. Ye, Y. Wu, D. Zhuo, L. Ceze and A. Krishnamurthy. 2024. “Punica: Multi-tenant LoRA Serving”. *Proceedings of Machine Learning and Systems* 6:1–13.

Cheng, S., P. J. Giabbanelli, and Z. Kuang. 2023. “Identifying the Building Blocks of Social Simulation Models: A Qualitative Analysis using Open-Source Codes in NetLogo”. In *Annual Modeling and Simulation Conference*, 306–317. IEEE.

LangChain Contributors and Harrison Chase 2023. “LangChain: A Framework for Developing Applications Powered by Language Models”. <https://github.com/langchain-ai/langchain>. GitHub repository. Accessed: 2025-06-03.

Di Menna, F., L. Traini, G. Bavota, and V. Cortellessa. 2025. “Investigating Execution-Aware Language Models for Code Optimization”. arXiv preprint arXiv:2503.08228.

Dou, S., E. Zhou, Y. Liu, et al. 2024. “LoRAMoE: Alleviating World Knowledge Forgetting in Large Language Models via MoE-style Plugin”. In *Proc. 62nd Annual Meeting of the Association for Computational Linguistics*, 1932–1945.

Duvvuru, V. S. A., B. Zhang, M. Vierhauser, and A. Agrawal. 2025. “LLM-Agents Driven Automated Simulation Testing and Analysis of small Uncrewed Aerial Systems”. arXiv preprint arXiv:2501.11864.

Frydenlund, E., J. Martínez, J. J. Padilla, K. Palacio and D. Shuttleworth. 2024. “Modeler in a Box: How Can Large Language Models Aid in the Simulation Modeling Process?”. *Simulation*:00375497241239360.

Gerstmayr, J., P. Manzl, and M. Pieber. 2024. “Multibody Models Generated from Natural Language”. *Multibody System Dynamics* 62(2):249–271.

Giabbanelli, P. J., T. J. Gandee, A. Agrawal, and N. Hosseinichimeh. 2025. “Benchmarking and Assessing Transformations Between Text and Causal Maps via Large Language Models”. *Applied Ontology*:15705838241304102.

Glimm, B., I. Horrocks, B. Motik, G. Stoilos and Z. Wang. 2014. “HermiT: An OWL 2 Reasoner”. *Journal of Automated Reasoning* 53(3):245–269.

Hassan, M., S. Ahmadi-Pour, K. Qayyum, C. K. Jha and R. Drechsler. 2024. “LLM-Guided Formal Verification Coupled with Mutation Testing”. In *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 1–2.

Hindle, A., E. T. Barr, M. Gabel, Z. Su and P. Devanbu. 2016. “On the Naturalness of Software”. *Commun. ACM* 59(5):122–131.

Horridge, M. and S. Bechhofer. 2004. “The Protégé OWL Plugin: An Open Development Environment for Semantic Web Applications”. In *Proc. 3rd Int. Semantic Web Conference (ISWC)*, 303–317: Springer. Protégé software and plugins.

Jackson, D. 2006. *Software Abstractions: Logic, Language, and Analysis*. MIT Press. Alloy Analyzer documentation.

Jackson, I., M. Jesus Saenz, and D. Ivanov. 2024. “From Natural Language to Simulations: Applying AI to Automate Simulation Modelling of Logistics Systems”. *International Journal of Production Research* 62(4):1434–1457.

Jackson, I. and B. Rolf. 2023. “Do Natural Language Processing Models Understand Simulations? Application of GPT-3 to Translate Simulation Source Code to English”. *IFAC-PapersOnLine* 56(2):221–226.

Khan, J. Y. and G. Uddin. 2023. “Automatic Code Documentation Generation Using GPT-3”. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering, ASE ’22*: ACM.

Lam, L. H. M., R. K. Thatikonda, and E. Shareghi. 2024. “A Closer Look at Logical Reasoning with LLMs: The Choice of Tool Matters”. arXiv preprint arXiv:2406.00284.

Lehmann, R. 2024. “Towards Interoperability of APIs—an LLM-based approach”. In *Proceedings of the 25th International Middleware Conference: Demos, Posters and Doctoral Symposium*, 29–30.

Li, J., B. Hui, G. Qu, B. Li, J. Yang, B. Li *et al.* 2024. “Can LLM Already Serve as a Database Interface? A Big Bench for Large-scale Database Grounded Text-to-SQLs”. *Advances in Neural Information Processing Systems* 36.

Liu, H. and L. Li. 2024. “On Languaging a Simulation Engine”. arXiv preprint arXiv:2402.16482.

Lúcio, L., S. Mustafiz, J. Denil, H. Vangheluwe and M. Jukss. 2013. “FTG+PM: An Integrated Framework for Investigating Model Transformation Chains”. In *SDL 2013: Model-Driven Dependability Engineering*, 182–202: Springer.

Martínez, J., B. Llinas, J. G. Botello, J. J. Padilla and E. Frydenlund. 2024. “Enhancing GPT-3.5’s Proficiency in NetLogo Through Few-Shot Prompting and Retrieval-Augmented Generation”. In *2024 Winter Simulation Conference (WSC)*, 666–677. IEEE <https://doi.org/10.1109/WSC63780.2024.10838967>.

Miceli-Barone, A. V., A. Lascarides, and C. Innes. 2023. “Dialogue-Based Generation of Self-Driving Simulation Scenarios Using Large Language Models”. arXiv preprint arXiv:2310.17372.

Nam, D., A. Macvean, V. Hellendoorn, B. Vasilescu and B. Myers. 2024. “Using an LLM to Help with Code Understanding”. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 1–13.

Novak, M. 2023. “Lawyer Uses ChatGPT in Federal Court and it Goes Horribly Wrong”. Accessed 28th July 2025.

Padiu, B., R. Iacob, T. Rebedea, and M. Dascalu. 2024. “To What Extent Have LLMs Reshaped the Legal Domain so Far? A Scoping Literature Review”. *Information* 15(11):662.

Pan, L., A. Albalak, X. Wang, and W. Y. Wang. 2023. “Logic-LM: Empowering Large Language Models with Symbolic Solvers for Faithful Logical Reasoning”. arXiv preprint arXiv:2305.12295.

Pan, R., A. R. Ibrahimzada, R. Krishna, D. Sankar, L. P. Wassi, M. Merler *et al.* 2023. “Understanding the Effectiveness of Large Language Models in Code Translation”. *CoRR* abs/2308.03109.

Pordanesh, S. and B. Tan. 2024. “Exploring the Efficacy of Large Language Models (GPT-4) in Binary Reverse Engineering”. arXiv preprint arXiv:2406.06637.

Quest Software, Inc. 2023. “Erwin Data Modeler”. <https://www.erwin.com/products/erwin-data-modeler>. Accessed 28th July 2025.

Saparov, A., R. Y. Pang, V. Padmakumar, N. Joshi, M. Kazemi, N. Kim *et al.* 2023. “Testing the General Deductive Reasoning Capacity of Large Language Models Using OOD Examples”. *Advances in Neural Information Processing Systems* 36:3083–3105.

Schäfer, M., S. Nadi, A. Eghbali, and F. Tip. 2024. “An Empirical Evaluation of Using Large Language Models for Automated Unit Test Generation”. *IEEE Transactions on Software Engineering* 50(1):85–105.

Schroeder, S. A., C. Vendome, P. J. Giabbanelli, and A. M. Montfort. 2022. “Towards Reusable Building Blocks to Develop COVID-19 Simulation Models”. In *2022 Winter Simulation Conference (WSC)*, 569–580 <https://doi.org/10.1109/WSC57314.2022.10015340>.

Schuerkamp, R., P. J. Giabbanelli, and N. Daclin. 2023. “Facilitating the Interoperability and Reuse of Extensions of Fuzzy Cognitive Maps”. In *Annual Modeling and Simulation Conference*, 708–719. IEEE.

Shrestha, S. L. and C. Csallner. 2021. "SLGPT: Using Transfer Learning to Directly Generate Simulink Model Files and Find Bugs in the Simulink Toolchain". In *Proc. 25th Int. Conf. Evaluation and Assessment in Software Engineering, EASE '21*, 260–265: ACM.

Tang, N. 2024. "Towards Effective Validation and Integration of LLM-Generated Code". In *2024 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, 369–370. IEEE.

Tian, Y., B. Zhang, Z. Tu, and D. Chu. 2025. "Adapters Selector: Cross-Domains and Multi-Tasks LoRA Modules Integration Usage Method". In *Proceedings of the 31st International Conference on Computational Linguistics*, 593–605.

Tolk, A. 2024. "Conceptual Alignment for Simulation Interoperability: Lessons Learned From 30 Years of Interoperability Research". *Simulation* 100(7):709–726.

Van Veen, D., C. Van Uden, L. Blankemeier, J.-B. Delbrouck, A. Aali, C. Bluethgen *et al.* 2024. "Adapted Large Language Models can Outperform Medical Experts in Clinical Text Summarization". *Nature medicine* 30(4):1134–1142.

Vangheluwe, H., J. de Lara, and P. J. Mosterman. 2002. "An Introduction to Multi-Paradigm Modelling and Simulation". In *Proc. AIS2002 Conference*, 9–20.

Wang, L., Y. Ling, Z. Yuan, M. Shridhar, C. Bao, Y. Qin *et al.* 2024. "GenSim: Generating Robotic Simulation Tasks via Large Language Models". arXiv preprint arXiv:2310.01361.

Xu, S., Z. Li, K. Mei, and Y. Zhang. 2024. "AIOS Compiler: LLM as Interpreter for Natural Language Programming and Flow Programming of AI Agents". arXiv preprint arXiv:2405.06907.

Yang, M., J. Chen, Y. Zhang, J. Liu, J. Zhang, Q. Ma *et al.* 2024. "Low-Rank Adaptation for Foundation Models: A Comprehensive Review". arXiv preprint arXiv:2501.00365.

Yoon, D., C. Han, D. W. Kim, S. Kim, S. Bae, J. A. Ryu *et al.* 2024. "Redefining Health Care Data Interoperability: Empirical Exploration of Large Language Models in Information Exchange". *Journal of Medical Internet Research* 26:e56614.

Yu, H., L. Fan, L. Li, J. Zhou, Z. Ma, L. Xian, *et al.* 2024. "Large Language Models in Biomedical and Health Informatics: A Review with Bibliometric Analysis". *Journal of Healthcare Informatics Research* 8(4):658–711.

Zhang, L., Z. Chen, and V. Ford. 2024. "Advancing Building Energy Modeling with Large Language Models: Exploration and Case Studies". *Energy and Buildings* 323:114788.

Zheng, Y., W. Gan, Z. Chen, Z. Qi, Q. Liang and P. S. Yu. 2024. "Large Language Models for Medicine: A Survey". *International Journal of Machine Learning and Cybernetics*:1–26.

Zhou, X., Z. Sun, and G. Li. 2024. "DB-GPT: Large Language Model Meets Database". *Data Sci. Eng.* 9(1):102–111.

## AUTHOR BIOGRAPHIES

**PHILIPPE J. GIABBANELLI** is a Research Professor at the Virginia Modeling, Analysis and Simulation Center (VMASC) at Old Dominion University. He has authored over 140 articles (of which 15 have appeared at the Winter Simulation conference) and several books in Modeling & Simulation. His technical expertise includes combining M&S with Machine Learning methods such as Generative AI, including Large Language Models. His email address is [pgiabban@odu.edu](mailto:pgiabban@odu.edu).

**JOHN BEVERLEY** is an Assistant Professor at the University of Buffalo, Co-Director of the National Center for Ontological Research, co-lead developer of BFO, and Associate Director of the Institute of Artificial Intelligence and Data Science. He holds a PhD in Philosophy from Northwestern University. His research is at the intersection of ontology engineering, formal logic, and machine learning. His email is [johnbeve@buffalo.edu](mailto:johnbeve@buffalo.edu).

**ISTVAN DAVID** is an Assistant Professor at McMaster University in Canada, where he leads the Sustainable Systems and Methods Lab, and is a faculty member of the McMaster Centre for Software Certification (McSCert). His research focuses on engineering complex heterogeneous systems (e.g., digital twins and cyber-physical systems) through model-driven engineering, simulation, and collaborative modeling. His email is [istvan.david@mcmaster.ca](mailto:istvan.david@mcmaster.ca).

**ANDREAS TOLK** is Chief Scientist for Complex Systems Modeling at The MITRE Corporation in Charlottesville, Virginia. He holds a PhD and M.Sc. in Computer Science from the University of the Federal Armed Forces, Germany. His research focuses on interoperability, data science, and the integration of simulation with real-world systems, particularly in defense and engineering applications. He delivered the military keynote at WinterSim'24. His email address is [atolk@mitre.org](mailto:atolk@mitre.org).