AI Beyond Rules, Heuristics, and Dreams: Ergative-Absolutive Agents for Participatory Simulations

Denisa Reshef Kera[1,*], Avital Dotam[2]

[1]Design & Policy Lab, STS, Bar Ilan University, 5290002 Ramat Gan, Israel
[2]Avital Dotan, Science, Technology and Society studies, 5290002 Ramat Gan, Israel

Abstract

This paper opens with a critique of the familiar division between rule-based Symbolic AI and data-driven Subsymbolic methods, suggesting that even their neurosymbolic convergence remains shaped by an inherited model of agency in which discrete subjects act upon passive objects, a structure naturalized by Indo-European subject-predicate-object grammars. We propose a linguistic turn in designing AI agents, using typologically diverse alignment systems, particularly ergative-absolutive languages such as Basque, as tools to rethink how agency is modeled and enacted in language-trained systems. We argue that Large Language Models (LLMs), far from being mere predictive tools, function as performative stages where grammars of agency are enacted rather than encoded. This reframing invites a shift: from optimizing systems to express predefined meanings, to interpreting the emergent structures that unfold through interaction. Drawing on the metaphor of the dreaming machine, we treat unpredictability and improvisation not merely as limitations of reasoning, but as openings for enacting alternative ontologies of action. To explore this, we propose a two-step framework. First, we examine how alignment patterns surface in LLM-generated interaction, not as imposed rules, but as constraints enacted by the grammar in context. Second, we stage participatory simulations in which stakeholders co-design agents with contrasting grammatical alignments, testing how such reconfigurations may support more adaptive, negotiated, and accountable forms of agency.

Keywords

ergative AI, ergative languages, language alignment, performative simulations, AI agency, deliberative AI, AI-human collaboration

1. Beyond Rules and Dreams: The Linguistic Reconfiguration of AI Agency

1.1. Symbolic, Subsymbolic, and Neurosymbolic AI

Designing Artificial Intelligence (AI) agents involves navigating a fundamental trade-off between Symbolic and Subsymbolic (connectionist) approaches [1, 2, 3, 4]. Symbolic AI, rooted in the principles of formal logic and explicit knowledge representation, constructs agents that operate through a system of predefined rules and symbolic manipulation [5]. These agents rely on explicit representations of knowledge, such as facts, concepts, and relationships, that make their actions determined by rule-based inference mechanisms and enables them to make decisions based on logical deductions.

Symbolic AI's strength is its transparency and controllability. This makes reasoning traceable and behavior predictable by design, enacting a view of agency rooted in clear intentionality and structure. The weakness, however, lies in handling the inherent uncertainty and complexity of the real world [6, 7] as environments where knowledge is rarely complete or perfectly definable. Furthermore, the processes of encoding all the necessary knowledge and rules for even relatively simple tasks can be extremely time-consuming and difficult.

In contrast, Subsymbolic AI, exemplified by today's dominant neural net-works, excels where symbolic approaches fail, often embodying the "heuristic" or even "dream-like" aspects of intelligence. Rather than relying on explicit rules and symbolic representations, AI "learns" from the data through statistical

ESWC 2025 Workshops and Tutorials Joint Proceedings
*Corresponding author.
denisa.reshef@biu.ac.il (D. Reshef Kera)
0000-0001-7242-9605 (D. Reshef Kera)

adaptation and pattern recognition over interconnected, adaptable nodes (neurons). The interconnected neurons ad-just their weights based on the input data they receive, which enables the system to react to uncertainty and change, generalize from incomplete data, and handle the complexity of real-world contexts. Such adaptability, however, typically sacrifices interpretability and direct control. The internal workings of the models are often opaque, creating "black box" systems that resemble associative "dreams" more than reasoned deliberation. The lack of interpretability and explicit control creates a persistent challenge for trust, safety, and accountability.

In response to the tension between logic-driven transparency and data-driven adaptability, many proposed various neurosymbolic approaches [8], aiming to layer formal constraints onto neural architectures. While these models offer valuable gains in interpretability, most remain committed to an architectural logic that treats symbolic and neural components as discrete, interfaced modules. This modularity reflects a representationalist paradigm: symbols encode fixed structures, while neural systems approximate statistical regularities, yet neither accommodates forms of agency that emerge relationally through interaction, context, and negotiated roles. As a result, neurosymbolic systems often struggle to support dynamic, co-constructed meaning, reverting instead to predefined schemas and inference paths.

What is missing is a framework in which structure and adaptability are not externally engineered and pre-coded but co-emerge from within a relational system. Rather than fusing rule-based and connectionist systems, our approach examines how natural language itself encodes neuro-symbolic orderings, particularly through alignment systems like ergativity, which bind agency, causation, and hierarchy into flexible, interpretable patterns. By treating LLMs as sites where these deep linguistic structures are instantiated and can be experimentally reconfigured, we move beyond the logic of integration toward a situated grammar of action and participation.

1.2. Rethinking AI Agents through Linguistic Alignment

Instead of treating intelligence as either structured reasoning or emergent adaptation, we approach it as fundamentally embedded in language and culture. What makes Large Language Models (LLMs) unique is that they are not only capable of capturing this embeddedness but also per-forming it, enacting latent structures of meaning, role, and inference that have evolved across linguistic systems. This perspective motivates our turn to ergative-absolutive alignment not merely as a linguistic curiosity, but as a conceptual resource for rethinking AI agency.

Linguistic systems like ergativity encode deep, often overlooked constraints on causation, agency, and hierarchy. Unlike subject-object-predicate schemas, which presuppose a stable and centralized agent, ergative constructions foreground the relational and distributed nature of action. By treating LLMs as laboratories for surfacing these "neurosym-bolic" grammars, we move beyond the opposition between symbolic rigidity and subsymbolic emergence. More broadly, we aim to explore a range of typologically diverse linguistic alignments, including systems that vary in marking (e.g., case marking, agreement marking) and in whether they are syntactically or semantically based [9]. While many languages rely on syntactic structure to assign argument roles (English), others, like Chukchee or Folopa, use semantic cues such as volitionality to determine grammatical marking, leading to dynamic realignments of agency [9]. By exploring such diversity, we can challenge the implicit biases embedded in conventional AI design.

While neurosymbolic architectures attempt to integrate symbolic transparency and neural adaptability, they often remain committed to modular design, unable to support relational or performative forms of agency. What they fail to capture is a form of agency that is simultaneously structured and adaptive, relational and accountable. Our approach differs by starting from the premise that language already integrates symbolic and neural dynamics. We are not proposing to build neurosymbolic systems, but to interpret LLMs as already bearing traces of these dynamics, particularly where typological variation, such as ergativity, encodes different cultural situated assumptions about action.

This reorientation reframes AI agency not as the execution of pre-programmed actions upon a predefined environment, but as participation in the situated co-construction of meaning. Rather than

optimizing for prediction or goal-completion, agents become responsive to the evolving grammar of interaction, where roles are fluid, hierarchies are negotiated, and agency is not assigned in advance but emerges relationally. In this context, ergativity is not a syntactic feature to be grafted onto existing models, but a structural affordance that challenges dominant agent-patient schemas encoded in nominative-accusative languages.

While some [10] treated ergative alignment as a surface phenomenon with no real impact on semantic role structure, subsequent work in typology and syntax [9, 11] has shown that ergativity can reshape core grammatical operations: relativization, coordination, argument omission, and clause chaining all operate differently in syntactically ergative systems. These differences are not reducible to static semantic roles but reflect alternative grammars of salience, control, and responsibility. When such grammars are internalized by LLMs, through training on ergative-aligned corpora, they do not function as symbolic parameters to be toggled, but as deep structural tendencies that shape interactional behavior.

Our proposal is to activate these grammars not through rule imposition, but through performative simulation: staged interactions in which agents operating under different alignment systems negotiate meaning, distribute roles, and respond to shifting contextual cues. Ergativity, in this frame, is not a syntactic feature to emulate, but a structural affordance to be enacted and tested in participatory settings. It offers a pathway for reimagining agency not as a predefined capacity, but as a situated, negotiated, and collectively distributed function. As Polinsky [11] demonstrates, syntactic ergativity reshapes not only morphosyntactic relations but also the allocation of grammatical pivots and syntactic operations, revealing that alignment systems encode distinct architectures of control and agency. Recognizing this invites a shift in AI design: from engineering hybrid symbolic-subsymbolic architectures toward interrogating the epistemologies these systems perform. This reorientation echoes a deeper philosophical divide between rationalist ideals of rule-based reasoning and the improvisational, context-sensitive dynamics of social interaction.

The tension between symbolic clarity and emergent adaptability has long shaped the conceptual field of AI. But it also mirrors a longstanding divide in the philosophy of language and mind: on one side, the rationalist tradition that treats language as an internal, rule-governed system for representing thought [12, 13], and on the other, a pragmatic and embodied view that understands language as action situated in social and material contexts [14, 15, 16]. In the next section, we explore how this tension animates the figure of the "dreaming machine," a metaphor that captures both the generative potential and epistemic instability of subsymbolic systems. Understanding this metaphor allows us to situate our proposal for agents who do not merely generate language, but inhabit its grammars, perform its roles, and co-construct meaning through interaction.

2. Dreaming Machine's Paradox: From Subsymbolic Dreams to Relational Agency

2.1. The Metaphor of the "Dreaming Machine"

Having argued for a shift in AI agent design, from engineered systems balancing rules and neural networks to a focus on the latent grammars of agency enacted through language, we propose a different point of departure: not the architecture of intelligence, but the performance of language. We view AI agency as relational and situated, not engineered in advance but discovered through linguistic interaction. This requires a shift from designing systems to interpreting their enactments. But what does it mean to say that LLMs enact language and, with it, culture and agency? How do we make sense of its outputs not as answers, but as performances within latent grammatical structures?

To explore these questions, we turn to a metaphor that has long shaped how subsymbolic intelligence is imagined: the dreaming machine. The figure of the dream reveals the epistemological rupture between rule-based and generative systems and opens space for rethinking what it means for AI to act, interpret, and participate. It is a metaphor that captures both the allure and anxiety surrounding deep learning systems, shaping public imagination and theoretical discourse alike. Subsymbolic models generate

associative patterns from vast textual and sensory data mirroring the process of human dreams: an improvisational entangling of memories, stimuli, and abstraction.

The resemblance is not merely aesthetic. Deep learning models, such as DeepDream and Generative Adversarial Networks (GANs), produce out-puts that often defy logical structure, resembling the distortions and juxtapositions characteristic of dreams [10 - 12]. The unpredictability of their results, such as hallucinatory visuals, latent-space interpolations, and emergent stylistic transformations, suggests a fundamental departure from the epistemological foundations of Symbolic AI. This raises a critical question for AI agent design: if heuristic and generative models resemble dreamers more than rule-bound deliberators, should they still be evaluated by criteria rooted in logical precision and representational fidelity? If AI is fundamentally improvisational, producing meaning through enactment of patterns, context, and ambiguity, then its outputs call for interpretive frameworks that accommodate indeterminacy rather than constrain it within inherited structures of correctness.

Rather than resolving the tension between symbolic and subsymbolic models by layering symbolic scaffolding onto neural architectures, we argue that the tension itself calls for rethinking the very notion of agency. We propose a turn to linguistic alignment systems, specifically, ergative-absolutive structures [9, 17], not as models to emulate, but as grammatical resources for reimagining interaction. Ergativity decentralizes the figure of the agent, foregrounding context, dependency, and relational roles. Where subject-object-predicate constructions presume an initiating actor, ergative alignments allow roles to shift based on tense, volition, or discourse frame. These alignmnets are not rules to be coded but grammars to be enacted. In this light, the dreaming machine is not a failure of symbolic order, but a space for activating alternate grammars of sense and responsibility. It becomes a stage where language, culture, and agency unfold, not through rules or randomness, but through performative structure.

2.2. Staging Grammar: From Dreaming Metaphors to Simulated Interaction

If the dreaming machine marks a shift from representational models to performative frameworks, then it invites a corresponding shift in method: from designing fixed behaviors to staging interactions that reveal how agency is enacted. Rather than resolving the symbolic–subsymbolic tension through architectural integration, we propose a methodological re-framing: using AI agent simulations as a space to activate and observe how grammatical structures, particularly alignment systems, shape inter-action.

AI agent simulations are typically used to evaluate performance against predefined goals, such as efficiency, accuracy, coherence, or task completion. In contrast, we propose simulations as sites for exposing and testing the grammatical and interactional assumptions embedded in the agents' behavior. Rather than assess how well agents meet external criteria, we are proposing simulations that explore how agents enact roles, negotiate meaning, and redistribute agency under different linguistic alignments. By varying the conditions of alignment (nominative-accusative versus ergative-absolutive), we observe how the internalized grammars of LLMs affect relational dynamics and responsibility framing in interaction[1].

Our focus is not on engineering these grammars into the models, but on staging them as conditions for interaction. Alignment is treated not as a rule to be imposed but as a latent affordance to be activated through context-sensitive prompts, discourse framing, and participatory role design. Participants in such simulations co-create agent roles and interaction scenarios, experimenting with different linguistic constraints to ob-serve how agents negotiate, defer, or assume roles under varying conditions.

This approach supports a performative understanding of AI agent simulations. If alignment systems carry with them epistemologies of causation and control, then to simulate alignment is to simulate worldviews. Simulation, in this sense, becomes a method of epistemic inquiry rather than verification, one that reveals how language-trained models inhabit, rather than merely generate, linguistic structure. In what follows, we ex-amine how ergative-absolutive alignment functions not only as a grammatical system, but as a situated constraint that reconfigures how agency is enacted and interpreted. This forms the linguistic core of our proposal.

[1]We are in an early stage of testing the research design of that hypothesis: https://github.com/anonette/ai-debate-simulator

3. The Linguistic Turn: Rethinking AI Agency through Ergative-Absolutive Alignment

3.1. Reconfiguring Agency: Linguistic Alignment as Design Constraint

Building on the framing established in Section 1.1, we shift from architectural critiques of symbolic/-subsymbolic models to a deeper examination of the grammatical assumptions shaping agency in LLMs Beyond the technical divide between rule-based control and statistical emergence lies a deeper constraint: the grammatical bias encoded in subject-object-predicate structures dominant in Indo-European languages. These alignments, which shape much of the data used to train LLMs, subtly reinforce hierarchical models of action, where agents are cast either as initiators acting on objects or as reactive systems governed by input. This grammar of agency limits the space for negotiated roles, distributed responsibility, and adaptive, context-sensitive interaction.

A central objection, grounded in linguistic theory, holds that ergative-absolutive alignment does not fundamentally alter semantic roles such as agent and patient. Edward Keenan, a foundational figure in the typology of grammatical relations, argued [10] that while case-marking patterns vary across languages, the underlying semantic structure of verb arguments (who does what to whom) remains stable. On this view, alignment systems reflect surface-level syntactic variation rather than deep differences in how agency or causation are conceptually organized. The implication is that ergativity, while typologically interesting, has limited relevance for modeling agency in AI systems trained on natural language.

Our proposal does not dispute the existence of cross-linguistic semantic regularities, nor does it claim that alignment systems redefine the ontology of events. Rather, we challenge the assumption that these regularities render alignment typologically irrelevant for modeling interactional agency. Drawing on more recent linguistic research [9, 11], we build on work that complicates the universalist position associated with Keenan [10]. Empirical studies by Dryer [18], Tsunoda [19], and Moravcsik [20] document languages with mixed or split alignment systems, in which morphological and syntactic patterns diverge in systematic ways. These data reveal that alignment does not operate as a uniform grammatical layer but interacts variably with tense, aspect, person, and clause structure. Furthermore, Gildea [21] and McGregor [22] argue that alignment is often sensitive to discourse-pragmatic factors such as perspective, topicality, and volitionality, indicating that alignment systems reflect interactional and cultural logics rather than fixed cognitive universals. These findings support our claim that alignment functions not as a static syntactic parameter, but as a context-sensitive constraint that shapes how agency is foregrounded and interpreted in interaction.

In LLMs, alignment is not solely a product of fine-tuning or reinforcement learning, but emerges through the enactment of patterns already embedded in language itself, patterns that models internalize through statistical exposure and reproduce in interaction. In light of this, we propose that ergativity can serve as a design-relevant affordance when modeling interactional grammars in AI. It does not eliminate core semantic roles, but it reorganizes how those roles are enacted, how initiative is attributed, and how responsibility circulates in discourse. This becomes particularly significant in the context of LLMs. When trained on corpora that encode ergative alignment, explicitly or implicitly, LLMs may internalize not rules, but weighting biases that shape their outputs in pragmatically meaningful ways. Rather than asserting linguistic determinism, we suggest that alignment offers an underutilized resource for reconfiguring the relational dynamics of agency in AI interaction, particularly when enacted through participatory simulations where these grammars become active constraints rather than abstract representations.

3.2. Performative Agency: Designing AI Through Alignment and Interaction

Having proposed a reorientation of agency through grammatical alignment, we now turn to its implications for interaction design. This section transitions from conceptual critique to design practice, examining how ergative structures can inform participatory simulations. Rather than re-iterating the

improvisational potential of LLMs already discussed in earlier sections, we focus on how alignment systems serve as dynamic constraints that shape how agents assume, negotiate, and redistribute roles in context.

Preliminary tests[1] suggest that agents influenced by ergative-absolutive grammars exhibit greater sensitivity to interactional con-text. Their behavior reflects relational adaptivity, distributed role-taking, and flexible decision-making. Unlike conventional agents that follow hierarchical role logic, these agents engage in negotiated sequences of action where agency is co-constructed through interaction rather than pre-assigned by architecture or goal states. This design logic supports a performative model of agency: roles are enacted rather than hard-coded, and responsibility emerges in response to evolving scenarios. In participatory simulations, such agents do not execute tasks in isolation but participate in shaping meaning and accountability through discourse. Their actions are not determined by static inference rules but shaped by linguistic affordances that respond to grammar, framing, and participant cues.

This framing lays the groundwork for our simulation architecture, in which alignment is not a symbolic rule to be encoded, but a constraint on interpretation enacted through interaction. Agency becomes a product of engagement rather than an internal property of the model, enabling AI systems to operate in deliberative, context-sensitive environments that demand responsiveness, adaptability, and negotiation.

4. Operationalizing Linguistic Agency in AI Systems

4.1. From Linguistic Bias to Design Hypothesis

We now translate the preceding conceptual framework into a testable hypothesis: that linguistic alignment, particularly the dominance of subject-object-predicate structures in Indo-European languages, shapes how AI agents negotiate roles and distribute responsibility. These structures tend to reinforce agent-object hierarchies and fixed task sequences. In contrast, ergative-absolutive alignments offer an alternative interactional logic, in which agency is more distributed and contingent on context.

To move beyond critique, we propose an implementable design approach that tests whether alignment systems condition patterns of role assignment, decision-making, and collaborative reasoning. We hypothesize that AI agents exposed to ergative structures will display more adaptive and context-sensitive forms of interaction. This shifts the framing of simulations: instead of treating agents as logic-driven or optimization-bound units, we model them as participants in a performative process of deliberation. Building on performative and role-theoretic approaches to design [23, 24, 25], our agents enact agency through co-constructed meaning rather than predefined scripts.

Preliminary results from a pilot conducted in March 2025[1] suggest that ergative-trained agents engaged in more distributed turn-taking and deferral behavior, while Indo-European-trained agents favored directive action and centralized decision proposals. These early findings indicate the potential of alignment-informed agent design to influence discourse dynamics in simulated deliberation.

Grounding AI design in ergative alignment enables the development of agents that are relational, flexible, and responsive to shifting interactional cues. These agents define their roles through ongoing engagement with human and artificial interlocutors, rather than executing fixed behaviors. This participatory grammar of agency opens a new design space for deliberative and accountable AI.

4.2. Experimental Setup: Comparing Role Distribution Across Alignments

To validate this framework, we propose a comparative experimental design using two primary groups of AI agents, each trained using distinct linguistic models. While the first group will use model trained on Indo-European languages, the second group, will use models trained in ergative-absolutive languages such as Basque that decentralize agency by, for example, aligning the subject of intransitive verbs with the object of transitive verbs.

Both groups will engage in structured problem-solving tasks and emergent discussion scenarios, designed to evaluate agency, decision-making, and communicative dynamics. These tasks will include negotiation, explanation, fairness definitions, and ethical debate, allowing for an assessment of how agents interact, adapt, and engage in co-constructed reasoning. Our primary aim is to determine whether ergative-trained agents exhibit more fluid role assignments, adaptive decision-making, and decentralized negotiation structures compared to their Indo-European-trained counterparts.

The study will compare two sets of AI agents:

- Set A: LLM-based agents trained on Indo-European languages.
- Set B: LLM-based agents trained on ergative-absolutive languages such as Basque or Inuktitut.

Each group will engage in a series of interactive tasks, including negotiation, collaborative decision-making, and ethical deliberation. Scenarios will be designed to test:

- Role assignment and fluidity
- Responsiveness to context and shifting participant structures
- Patterns of initiative, turn-taking, and deference

The aim is to test whether alignment bias translates into observable differences in how agents engage with uncertainty, interpret roles, and structure interaction in multi-agent settings.

4.3. Hypotheses and Towards Participatory AI Simulations

We propose several key hypotheses. First, we anticipate that Indo-European-prompted AI agents will favor hierarchical, role-fixed interaction patterns with more centralized decision-making. Second, ergativity-informed agents will demonstrate more fluid role allocation and adaptive, relational reasoning across discursive contexts. Third, we expect structural misalignment when agents trained in different linguistic models interact, leading to divergent negotiation and reasoning patterns.

These outcomes support the broader claim that grammatical structures, embedded in training data and enacted in interaction, shape not only linguistic outputs but the distribution of agency in deliberative contexts. The implications extend to explainability, accountability, and the epistemic assumptions underpinning AI model design.

Beyond evaluation, we propose a participatory simulation framework in which human stakeholders co-design agents using ergative principles. Participants define agent personas, rhetorical strategies, and interaction protocols, enabling deliberative engagement rather than passive observation. In this setup, AI agents function as rhetorical participants rather than black-box tools. Prompt design, alignment framing, and discourse style become levers for shaping how agency is enacted and negotiated in hybrid public settings. The simulation environment serves both as a research testbed and a civic theatre, where the grammar of action is co-produced. This enables us to explore how alternative linguistic alignments may foster more accountable, situated, and collaborative models of AI-driven governance.

5. Conclusion

This paper reframes the challenge of AI agency by moving beyond the engineering of neurosymbolic hybrids, where symbolic rules and neural networks remain modular and externally integrated. We propose instead a linguistic turn: treating language as the medium in which symbolic and subsymbolic dynamics are already entangled and enacted. Alignment systems such as ergativity are not merely classificatory devices; they organize how agency, causation, and responsibility are distributed and negotiated within discourse.

Embedding AI in participatory and performative contexts shifts the design focus from deterministic optimization to deliberative engagement. Ergativity-informed models enable AI agents to act not as fixed executors of pre-assigned tasks, but as relational actors shaped by interaction. This reconceptualization opens space for fluid decision-making, dynamic forms of governance, and stakeholder-defined modalities

of agency. Through ergative-aligned agent design, participatory simulation, and co-configured rhetorical protocols, we explore how AI may function less as a predictive tool and more as a co-actor in a hybrid public sphere of collaborative meaning-making.

Declaration on Generative AI Use

During the preparation of this work, the authors used GPT-4o for the fol-lowing roles, as defined in the GenAI Usage Taxonomy: Grammar and spelling check, paraphrase and re-word selected sentences for improved clarity and conciseness. We also used it for peer review simulation and content enhancement to suggest additional counterpoints and test the conceptual robustness of the argumentation. The authors critically reviewed all suggestions and generative outputs, made substantial edits where appropriate, and take full responsibility for the content and its scholarly integrity. GPT-4o was used strictly as a tool to support and refine the author's independent research and writing.

References

[1] E. Ilkou, M. Koutraki, Symbolic Vs Sub-symbolic AI Methods: Friends or Enemies?, in: Proceedings of the CIKM 2020 Workshop, volume 2699, Galway, Ireland, 2020.

[2] C. Núñez Molina, P. Mesejo, J. Fernández-Olivares, A Review of Symbolic, Subsymbolic and Hybrid Methods for Sequential Decision Making, ACM Comput. Surv. 56 (2024). doi:10.1145/3663366.

[3] A. Platzer, Intersymbolic ai, in: Leveraging Applications of Formal Methods, Verification and Validation. Software Engineering Methodologies, Springer Nature Switzerland, Cham, 2025, pp. 162–180. doi:10.1007/978-3-031-75387-9_11.

[4] H. Xiong, Z. Wang, X. Li, J. Bian, Z. Xie, S. Mumtaz, A. Al-Dulaimi, L. E. Barnes, Converging paradigms: The synergy of symbolic and connectionist ai in llm-empowered autonomous agents, 2024. doi:10.48550/arXiv.2407.08516. arXiv:2407.08516.

[5] S. Russell, P. Norvig, Artificial Intelligence: A Modern Approach. 4th ed., Pearson, 2021.

[6] Y. LeCun, Y. Bengio, G. Hinton, Deep Learning, Nature 521 (2015) 436–444. doi:10.1038/nature14539.

[7] G. Marcus, Deep Learning: A Critical Appraisal (2018). doi:10.48550/arXiv.1801.00631. arXiv:1801.00631.

[8] A. d'Avila Garcez, M. Gori, L. Lamb, L. Serafini, M. Spranger, S. Tran, Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning (2019). doi:10.48550/arXiv.1905.06088. arXiv:1905.06088.

[9] R. M. W. Dixon, Ergativity, Cambridge Studies in Linguistics, Cambridge University Press, 1994. doi:10.1017/CBO9780511611896.

[10] E. Keenan, Semantic correlates of the ergative/absolutive distinction, Linguistics 22 (1984) 197–224. doi:10.1515/ling.1984.22.2.197.

[11] M. Polinsky, Deconstructing Ergativity: Two Types of Ergative Languages and Their Features, Oxford University Press, 2016. doi:10.1093/acprof:oso/9780190256586.001.0001.

[12] N. Chomsky, Syntactic Structures, The Hague: Mouton, 1957.

[13] J. Fodor, The Language of Thought, volume 5, Harvard University press, Cambridge, MA, USA, 1975.

[14] J. L. Austin, How to Do Things with Words, Oxford: Clarendon Press, 1962.

[15] R. J. Bernstein, What is the Difference That Makes a Difference? Gadamer, Habermas, and Rorty, PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 1982 (1982) 331–359. URL: http://www.jstor.org/stable/192429.

[16] A. Clark, Being There: Putting Brain, Body, and World Together Again, The MIT Press, Cambridge, MA, 1996. doi:10.7551/mitpress/1552.001.0001.

[17] A. Johns, D. Massam, J. Ndayiragije, Ergativity: Emerging Issues, volume 65 of *Studies in Natural Language and Linguistic Theory*, Springer, 2006. doi:10.1007/1-4020-4188-8.

[18] M. S. Dryer, Rewiew of: "Ergativity: Towards a Theory of Grammatical Relations", Canadian Journal of Linguistics 30 (1985) 207–212. doi:10.1017/S0008413100010938.
[19] T. TASAKU, Split case-marking patterns in verb-types and tense/aspect/mood, Linguistics 19 (1981) 389–438. doi:10.1515/ling.1981.19.5-6.389.
[20] E. A. Moravcsik, On the distribution of ergative and accusative patterns, Lingua 45 (1978) 233–279. doi:10.1016/0024-3841(78)90026-8.
[21] S. Gildea, Are there universal cognitive motivations for ergativity, Studies in Language 30 (2006) 301–357.
[22] W. B. McGregor, Typology of ergativity, Language and Linguistics Compass 3 (2009) 480–508.
[23] R. Schechner, Performance Studies: An Introduction. 2nd ed., Routledge, 2002.
[24] V. Turner, From Ritual to Theatre: The Human Seriousness of Play, Performing arts journal publ., 1998.
[25] I. Bogost, Unit Operations: An Approach to Videogame Criticism, MIT Press, 2006. doi:10.7551/mitpress/6997.001.0001.