

ARS as a Meta-Methodology

The Conditions for Explanatory Models

in the Age of Generative AI

Paul Koop

1994–2026

Abstract

This paper examines the Algorithmic Recursive Sequence Analysis (ARS) not merely as a method but as a *meta-methodology*—a framework that specifies the conditions under which a model counts as *explanatory* rather than merely *descriptive* or *simulative*. Drawing on the ARS’s core principles—the primacy of interpretation, the separation of structure and statistics, controlled falsification, and XAI validation—I argue that these principles constitute necessary criteria for explanatory models in any discipline that deals with sequential social processes. The paper systematically relates ARS to five contemporary research programs: (1) Formal Verification and Model Checking, (2) Interpretable Machine Learning and Rule Extraction, (3) Grounded Theory, (4) Causal Inference, and (5) Process Mining. For each, I demonstrate what ARS can learn from these approaches and, crucially, what these approaches can learn from ARS. The paper concludes that ARS provides a transdisciplinary benchmark for distinguishing genuine explanation from statistical or structural description.

Contents

1	Introduction: From Method to Meta-Methodology	3
2	ARS as a Meta-Methodology: The Core Principles	4
2.1	Principle 1: Interpretive Grounding	4
2.2	Principle 2: Separation of Structure and Statistics	4
2.3	Principle 3: Generative Transparency	4
2.4	Principle 4: Controlled Falsification	4
2.5	Principle 5: XAI Validation	5
3	Five Approaches in Dialogue with ARS	5
3.1	Formal Verification and Model Checking	5
3.1.1	What Formal Verification Is	5
3.1.2	What ARS Can Learn from Formal Verification	5
3.1.3	What Formal Verification Can Learn from ARS	6
3.2	Interpretable Machine Learning and Rule Extraction	6
3.2.1	What IML and Rule Extraction Are	6
3.2.2	What ARS Can Learn from IML and Rule Extraction	6
3.2.3	What IML and Rule Extraction Can Learn from ARS	7
3.3	Grounded Theory	7
3.3.1	What Grounded Theory Is	7
3.3.2	What ARS Can Learn from Grounded Theory	7
3.3.3	What Grounded Theory Can Learn from ARS	8
3.4	Causal Inference and Causal Graphical Models	8
3.4.1	What Causal Inference Is	8
3.4.2	What ARS Can Learn from Causal Inference	8
3.4.3	What Causal Inference Can Learn from ARS	9
3.5	Process Mining	9
3.5.1	What Process Mining Is	9
3.5.2	What ARS Can Learn from Process Mining	9
3.5.3	What Process Mining Can Learn from ARS	10
4	Toward a Transdisciplinary Benchmark for Explanation	10
4.1	The Explanation vs. Simulation Distinction Revisited	11
4.2	The Role of the Human Interpreter	12
5	Conclusion and Outlook	12

1 Introduction: From Method to Meta-Methodology

The Algorithmic Recursive Sequence Analysis (ARS), in its various versions (2.0–4.0), has been presented primarily as a method for the formal analysis of sequential interactions. It transforms interpretively obtained categories into formal models: probabilistic context-free grammars (PCFG), Petri nets, Bayesian networks, and deterministic finite automata (DFA). This paper takes a different perspective. It argues that ARS is not just a method but a *meta-methodology*—a framework that specifies the *conditions* under which a model can be considered explanatory.

This shift in perspective is motivated by a recurring confusion in contemporary AI and data science: the conflation of **simulation** with **explanation**. A large language model (LLM) trained on sales conversations can simulate plausible dialogues with high fidelity. Yet, as the ARS notebooks have demonstrated, it cannot explain *why* a particular sequence of speech acts is well-formed or what rules constitute its generation. The LLM provides a statistical shadow of the process; the ARS provides its structural skeleton.

The meta-methodological claim of this paper is that the ARS principles constitute **necessary conditions** for explanatory models in any discipline dealing with sequential social processes. These conditions are:

1. **Interpretive Grounding:** Symbols must be tied to documented qualitative interpretations, not merely to statistical correlations.
2. **Structural Decidability:** The well-formedness of sequences must be formally decidable (e.g., by a DFA), independent of empirical frequencies.
3. **Generative Transparency:** The model must be able to generate sequences in a way that every step is traceable to explicit rules.
4. **Falsifiability:** Interpretations must be subject to controlled falsification; counterexamples must be able to refute rules.
5. **XAI Validation:** The model must satisfy the NIST XAI criteria—meaningfulness, accuracy, and knowledge limits.

To substantiate this claim, I systematically relate ARS to five contemporary research programs that share overlapping concerns: formal verification, interpretable machine learning, grounded theory, causal inference, and process mining. For each, I ask two reciprocal questions:

1. What can ARS learn from this approach? (Technical or conceptual enhancements)
2. What can this approach learn from ARS? (Methodological safeguards, criteria for explanation)

2 ARS as a Meta-Methodology: The Core Principles

Before examining the five approaches, it is necessary to state the ARS principles that serve as the meta-methodological benchmark.

2.1 Principle 1: Interpretive Grounding

In ARS, every terminal symbol (KBG, VBBd, KAA, etc.) is the product of a documented qualitative interpretation. The coding process is not a black box; it is recorded, justified, and subject to intersubjective validation. This principle excludes purely data-driven category formation (e.g., clustering embeddings) from counting as explanation.

2.2 Principle 2: Separation of Structure and Statistics

As developed in `ARS_XAI_Aut2_Ger.tex`, ARS maintains a strict separation between structural rules (which are deterministic and context-free) and statistical regularities (which are empirical and contingent). A structural rule holds or does not hold—independent of how often it is violated. This separation is absent in purely statistical models.

2.3 Principle 3: Generative Transparency

The induced grammar must be able to generate sequences in a traceable manner. The transducer in Lisp, the parser in Pascal, and the inductor in Scheme each provide a different window into this transparency. A model that cannot generate exemplars from its own rules is not explanatory.

2.4 Principle 4: Controlled Falsification

Interpretations are not asserted once and for all. They are produced as readings and then falsified by subsequent sequence positions (following Oevermann's sequential

analysis). This creates a spiral of interpretation and refutation that mirrors Popperian falsificationism adapted to qualitative material.

2.5 Principle 5: XAI Validation

The three NIST XAI criteria—meaningfulness, accuracy, knowledge limits—are not optional add-ons but constitutive elements of explanatory models. Meaningfulness requires semantic interpretability; accuracy requires correspondence with the material; knowledge limits require explicit documentation of the model’s boundaries.

3 Five Approaches in Dialogue with ARS

3.1 Formal Verification and Model Checking

3.1.1 What Formal Verification Is

Formal verification, particularly model checking, is a method from theoretical computer science that systematically checks whether a formal model (e.g., a finite automaton, a Petri net, or a Bayesian network) satisfies specified properties. These properties include *safety* ("something bad never happens") and *liveness* ("something good eventually happens").

Model checking is exhaustive: it explores the entire state space of the model. Unlike statistical testing, which provides probabilistic guarantees, model checking provides *deterministic* guarantees about the model’s behavior.

3.1.2 What ARS Can Learn from Formal Verification

- **Property Specification Languages:** ARS could adopt temporal logics (LTL, CTL) to specify what properties the induced grammar should satisfy. For example: $\Box(KBG \rightarrow \Diamond VBG)$ ("Always, if a customer greets, eventually the seller greets back").
- **Counterexample Generation:** When a property fails, model checkers produce a counterexample trace. This could serve as a powerful falsification tool for ARS interpretations.
- **State Space Explosion Awareness:** ARS modelers should be aware that hierarchical grammars can lead to large state spaces. Formal verification offers abstraction techniques to manage this.

3.1.3 What Formal Verification Can Learn from ARS

- **Interpretive Grounding of States:** In standard model checking, states are abstract symbols. ARS insists that each state must be interpretively grounded. This could lead to a new subfield: *interpretive model checking*, where properties are checked *and* the meaning of states is documented.
- **The Separation of Structure and Statistics:** Model checking typically assumes a deterministic model. ARS shows how to separate structural rules (which can be verified) from statistical variations (which cannot). This could enrich probabilistic model checking.
- **XAI Criteria for Verification:** Model checking results are often opaque to non-experts. ARS’s XAI criteria could guide the development of more understandable verification outputs.

3.2 Interpretable Machine Learning and Rule Extraction

3.2.1 What IML and Rule Extraction Are

Interpretable Machine Learning (IML) aims to make black-box models (neural networks, gradient boosting, random forests) understandable to humans. *Rule extraction* is a specific IML technique that attempts to describe the learned function approximately through a set of if-then rules (e.g., using RIPPER, CART, or analyzing activation patterns).

Unlike ARS, which induces rules directly from data, rule extraction typically works *post-hoc*: the model is already trained, and rules are extracted as an explanation.

3.2.2 What ARS Can Learn from IML and Rule Extraction

- **Scalable Rule Induction:** ARS currently induces rules from small corpora ($n = 8$). IML offers techniques for extracting rules from large datasets, albeit with less interpretive control.
- **Quantitative Rule Evaluation:** IML provides metrics for evaluating extracted rules (coverage, fidelity, stability). ARS could adopt these to assess how well a grammar generalizes.
- **Hybrid Rule Sets:** Some IML methods combine global and local rules. ARS could explore hybrid grammars that have both a core structural grammar and local statistical variations.

3.2.3 What IML and Rule Extraction Can Learn from ARS

- **Explication vs. Approximation:** IML's rule extraction is almost always approximate. ARS insists on *exact* rules for the given corpus. This raises a fundamental question: Is approximation ever acceptable for explanation? ARS suggests a clear answer: approximation is acceptable only if the approximation error is documented and the structural core is exact.
- **Interpretive Validation of Rules:** IML extracts rules that statistically fit the data. ARS adds a layer of *interpretive validation*: rules must also make sense to human interpreters. This could prevent the extraction of statistically correct but semantically meaningless rules.
- **Falsifiability as a Criterion:** IML rarely discusses how extracted rules could be falsified. ARS makes falsifiability a central criterion. IML could adopt this: a rule set is not explanatory if no conceivable counterexample could refute it.

3.3 Grounded Theory

3.3.1 What Grounded Theory Is

Grounded Theory (GT), developed by Glaser and Strauss, is a classic methodology in qualitative social research for developing theories from data. It involves procedures such as open coding, axial coding, and selective coding. The goal is to generate middle-range theories that are "grounded" in empirical material.

In recent years, there have been attempts to formalize parts of GT or to support it computationally (e.g., through natural language processing or topic modeling). However, GT remains largely informal in its final output.

3.3.2 What ARS Can Learn from Grounded Theory

- **Systematic Coding Procedures:** GT offers a rich vocabulary and set of procedures for coding that could enrich ARS's interpretive phase. Concepts like "axial coding" (relating categories to subcategories) are similar to ARS's hierarchical compression but more fine-grained.
- **Theoretical Sampling:** GT's principle of theoretical sampling—selecting new cases based on emerging theoretical insights—could guide ARS's case selection in larger studies.
- **Constant Comparison:** GT's method of constant comparison (comparing

each new case with already developed categories) is already implicit in ARS's systematic case comparison (Phase 4) but could be made more explicit.

3.3.3 What Grounded Theory Can Learn from ARS

- **From Theory to Generative Model:** GT typically stops at the level of narrative theory or category systems. ARS goes further: it transforms the theory into a *generative grammar* that can produce new sequences. GT could adopt this to make its theories testable and executable.
- **Formal Falsifiability:** GT's validation procedures are primarily qualitative (e.g., member checking, peer debriefing). ARS adds formal falsifiability: the grammar can be wrong in a way that can be demonstrated mechanically (e.g., by a parser rejecting a sequence).
- **XAI Criteria for Grounded Theory:** ARS's XAI criteria (meaningfulness, accuracy, knowledge limits) provide a checklist for evaluating grounded theories. A GT theory that cannot specify its knowledge limits is incomplete.

3.4 Causal Inference and Causal Graphical Models

3.4.1 What Causal Inference Is

Causal inference, particularly with causal graphical models (e.g., DAGs, DoWhy, CausalNex), goes beyond mere correlation and attempts to identify and quantify causal relationships between variables. It uses techniques such as the do-calculus, instrumental variables, and counterfactual reasoning.

A key insight of causal inference is that correlation is not causation. Directed acyclic graphs (DAGs) are used to represent assumptions about causal structure.

3.4.2 What ARS Can Learn from Causal Inference

- **Causal Interpretation of Grammars:** ARS grammars describe sequential dependencies. Causal inference could help distinguish whether these dependencies are merely sequential or genuinely causal. For example: Does the question "Anything else?" *cause* an additional purchase, or is it merely correlated with it?
- **Counterfactual Reasoning:** Causal inference excels at answering counterfactual questions ("What would have happened if the seller had not asked?"). ARS

could adopt counterfactual simulation (already present in Phase 3 of CGTI) as a standard validation tool.

- **Instrumental Variables for Sequential Data:** ARS deals with sequential data where confounding is common. Instrumental variable techniques could help identify causal effects even in observational sequential data.

3.4.3 What Causal Inference Can Learn from ARS

- **Interpretive Grounding of Causal Graphs:** In standard causal inference, the DAG is often assumed or learned from data without interpretive documentation. ARS insists that every node and edge must be interpretively grounded. This could lead to *interpretive causal inference* as a new subfield.
- **Sequential Grammars as Causal Structures:** ARS grammars are a form of causal structure over sequences. Causal inference typically deals with static or time-series data, not with grammatical sequences. ARS could inspire a new class of *grammatical causal models*.
- **XAI for Causal Models:** Causal models are often presented as DAGs with probabilities, which are not self-explanatory. ARS's XAI criteria could guide the documentation of causal models, making them more accessible to domain experts.

3.5 Process Mining

3.5.1 What Process Mining Is

Process mining is a research field at the intersection of data mining, machine learning, and process modeling. It aims to discover, conform, and enhance process models (often in the form of Petri nets, BPMN diagrams, or directly follows graphs) from event logs—sequential recordings of process steps, e.g., in workflow management systems or ERP systems.

Process mining typically works with large, anonymized, and weakly annotated logs. It discovers the "average process" or the "most common paths." It does not aim for a complete reconstruction of a single case's constitutive rules.

3.5.2 What ARS Can Learn from Process Mining

- **Scalable Discovery Algorithms:** Process mining offers sophisticated algorithms (e.g., Alpha miner, Heuristics miner, Inductive miner) for discovering

Petri nets from large logs. ARS could adopt or adapt these for larger corpora while preserving interpretive control.

- **Conformance Checking:** Process mining includes techniques for checking whether an event log conforms to a given model. This could be used to validate ARS grammars against new data.
- **Performance Analysis:** Process mining adds performance dimensions (time, cost, frequency). ARS could be extended to incorporate temporal and resource dimensions more systematically.

3.5.3 What Process Mining Can Learn from ARS

- **Interpretive Discovery for Small Logs:** Process mining typically requires large logs to produce reliable models. ARS shows how to discover models from a single case ($n = 1$) through interpretive depth. This could be valuable for process mining in domains where data is scarce (e.g., medical procedures, legal cases).
- **Documentation of Discovery Decisions:** Process mining algorithms make many decisions (e.g., which paths to include, how to handle noise). These decisions are rarely documented in an interpretively accessible way. ARS's reflexive documentation could serve as a model.
- **Separation of Structure and Statistics:** Process mining often produces models that mix structural rules with statistical noise. ARS's strict separation could improve the quality of discovered models by clearly distinguishing what is structurally necessary from what is merely empirically frequent.
- **XAI for Process Mining:** The models produced by process mining (e.g., spaghetti-like Petri nets) are often hard to understand. ARS's XAI criteria could guide the development of more explainable process mining outputs.

4 Toward a Transdisciplinary Benchmark for Explanation

The five comparisons above reveal a common pattern. Each contemporary approach has technical strengths that could enhance ARS: scalability (IML, Process Mining), formal rigor (Verification, Causal Inference), and systematic coding procedures (Grounded Theory). Conversely, each approach lacks some of the meta-methodological

safeguards that ARS provides: interpretive grounding, structural decidability, generative transparency, controlled falsification, and XAI validation.

This symmetry suggests that ARS is not merely one method among others but a **transdisciplinary benchmark** for what counts as an explanatory model.

Table 1: ARS as a Transdisciplinary Benchmark

Criterion	Question	Absence indicates...
Interpretive Grounding	Are the symbols meaningfully documented?	Statistical correlation without understanding
Structural Decidability	Is well-formedness formally decidable?	Probabilistic guessing rather than rule-following
Generative Transparency	Can the model generate exemplars traceably?	Simulation without explanation
Controlled Falsification	Can counterexamples refute rules?	Unfalsifiable post-hoc storytelling
XAI Validation	Are meaningfulness, accuracy, and knowledge limits documented?	Technical sophistication without epistemic accountability

4.1 The Explanation vs. Simulation Distinction Revisited

The core distinction that emerges from this analysis is between **explanation** and **simulation**. A model *simulates* if it reproduces the statistical properties of the data. A model *explains* if it specifies the constitutive rules that generate the phenomenon.

- **Simulation** is sufficient for prediction. An LLM that accurately predicts the next token is a good simulator.
- **Explanation** is necessary for understanding, intervention, and normative evaluation. An ARS grammar that specifies the rules of a sales conversation is an explanation.

The five criteria above are the conditions under which a model qualifies as an explanation rather than merely a simulation.

4.2 The Role of the Human Interpreter

A recurring theme across all five comparisons is the role of the human interpreter. In ARS, the human is constitutive: interpretation is a human act that cannot be fully automated. In the other approaches, the human is often external—designing algorithms, providing training data, evaluating outputs.

This is not a weakness of ARS but a strength. ARS makes explicit what is often implicit: that explanation is a human practice, not a property of a model considered in isolation. A model is explanatory *for someone* who can understand it, use it, and be held accountable for it.

5 Conclusion and Outlook

This paper has argued that the Algorithmic Recursive Sequence Analysis (ARS) is not merely a method but a meta-methodology—a framework that specifies the conditions under which a model counts as explanatory. Five core principles were identified: interpretive grounding, structural decidability, generative transparency, controlled falsification, and XAI validation.

The paper then systematically related ARS to five contemporary research programs: formal verification, interpretable machine learning, grounded theory, causal inference, and process mining. For each, I demonstrated what ARS can learn from these approaches (technical enhancements) and, crucially, what these approaches can learn from ARS (methodological safeguards, criteria for explanation).

The meta-methodological claim is that the ARS principles constitute necessary conditions for explanatory models in any discipline dealing with sequential social processes. They provide a transdisciplinary benchmark for distinguishing genuine explanation from statistical or structural description.

For future research, three directions are particularly promising:

1. **Implementation of hybrid systems:** Integrate ARS grammars with model checkers, rule extractors, or process mining algorithms while preserving the meta-methodological safeguards.
2. **Empirical testing of the benchmark:** Apply the five criteria to existing models in different disciplines and assess whether they predict perceived explanatory quality.

3. **Extension to non-sequential domains:** While ARS was developed for sequential data, the meta-methodological principles may generalize to other types of models (e.g., classification, clustering, regression).

In conclusion: The question is not whether a model fits the data. Statistical fit is necessary but not sufficient. The question is whether the model meets the meta-methodological criteria that make explanation possible. ARS provides a concrete, operationalized answer.

References

- Baier, C., & Katoen, J.-P. (2008). *Principles of Model Checking*. MIT Press.
- Glaser, B. G., & Strauss, A. L. (1967). *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Aldine.
- Koop, P. (1992). *Demo-Parser Chart-Parser Version 1.0*. Pascal source code.
- Koop, P. (1994). *Grammatikinduktion empirisch gesicherter Verkaufsgespräche*. Scheme source code.
- Koop, P. (1994). *Sequenzanalyse empirisch gesicherter Verkaufsgespräche*. Lisp source code.
- Koop, P. (2023). *Qualitative Sozialforschung und Große Sprachmodelle*. Jupyter Notebook.
- Koop, P. (2024/2026). *Between Interpretation and Computation: Algorithmic Recursive Sequence Analysis as a Bridge between Qualitative Hermeneutics and Formal Modeling*. the-last-freedom.org.
- Koop, P. (2026). *Social Structures and Processes: Causal Inference with Probabilistic Context-Free Grammars and Bayesian Networks*. the-last-freedom.org.
- Molnar, C. (2022). *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. leanpub.com.
- Oevermann, U., Allert, T., Konau, E., & Krambeck, J. (1979). The methodology of objective hermeneutics. In H.-G. Soeffner (Ed.), *Interpretative Procedures in the Social and Text Sciences* (pp. 352-434). Metzler.
- Ortigossa, E. S., Gonçalves, T., & Nonato, L. G. (2024). Explainable Artificial Intelligence (XAI)—From Theory to Methods and Applications. *IEEE Access*, 12, 80799-80846.
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.
- van der Aalst, W. M. P. (2016). *Process Mining: Data Science in Action* (2nd ed.). Springer.