

Generating Literature-Driven Scientific Theories at Scale

Peter Jansen^{1,2}, Peter Clark¹, Doug Downey¹, Daniel S. Weld^{1,3}

¹Allen Institute for Artificial Intelligence ²University of Arizona

³University of Washington

peterj@allenai.org

Abstract

Contemporary automated scientific discovery has focused on agents for generating scientific experiments, while systems that perform higher-level scientific activities such as theory building remain underexplored. In this work, we formulate the problem of synthesizing theories consisting of qualitative and quantitative laws from large corpora of scientific literature. We study theory generation at scale, using 13.7k source papers to synthesize 2.9k theories, examining how generation using literature-grounding versus parametric knowledge, and accuracy-focused versus novelty-focused generation objectives change theory properties. Our experiments show that, compared to using parametric LLM memory for generation, our literature-supported method creates theories that are significantly better at both matching existing evidence and at predicting future results from 4.6k subsequently-written papers.¹

1 Introduction

Automated scientific discovery systems aim to generate novel scientific insights by leveraging data, experiments, and prior literature (e.g. Thilakarathne et al., 2019; Majumder et al., 2024; Tobias and Wahab, 2025). Recent advances in language model agents have enabled the integration of these paradigms, leading to a proliferation of experiment-driven discovery systems across domains including computer science, biology, and astronomy (e.g. Lu et al., 2024; Jansen et al., 2025; Schmidgall et al., 2025; Ghafarollahi and Buehler, 2024; Moss, 2025).

While experiment-driven discovery is central to scientific progress, individual experiments are typically conducted in service of higher-level scientific activities such as theory formation and synthesis (Kuhn, 1962). Theories consolidate large bodies of

empirical results into compact, generalizable statements that support explanation and prediction. For example, Kepler’s laws compressed centuries of astronomical observations into a set of laws that describe planetary motion. This raises a fundamental question: *can AI systems, built on language models, synthesize scientific theories by reading the research literature?*

In this work, we first define the problem of literature-based theory synthesis — generating sets of qualitative and quantitative laws that generalize scientific knowledge in a subect area. We then specify five desirable criteria for these generated theories: specificity, empirical support, predictive accuracy, novelty, and plausibility. We introduce a novel system, THEORIZER, that reads tens of thousands of papers spanning diverse topics in artificial intelligence and natural language processing to generate numerous candidate theories. We explore several variants of THEORIZER— a RAG-style, literature-supported method and a simpler LLM baseline. For each of these methods, we consider two variants — one that optimizes accuracy and another that aims to increase novelty.

Evaluating thousands of theories via new experiments is infeasible at scale. To address this, we develop several evaluation methods, including a backtesting paradigm (Yetisgen-Yildiz and Pratt, 2009) where theories are generated using a fixed knowledge cutoff, and their predictions are evaluated against experimental results reported in subsequently published literature. Experiments show that, while our literature-supported method is almost 7x more expensive than a simple LLM approach, it yields significantly more accurate and predictive theories - as measured both by using an LLM judge on theory-law statements, and by comparing the generated theories to 4,554 subsequently-written papers in our backtesting paradigm. We further analyze how theory quality varies across the generation conditions, comparing generation

¹<https://github.com/allenai/asta-theorizer>

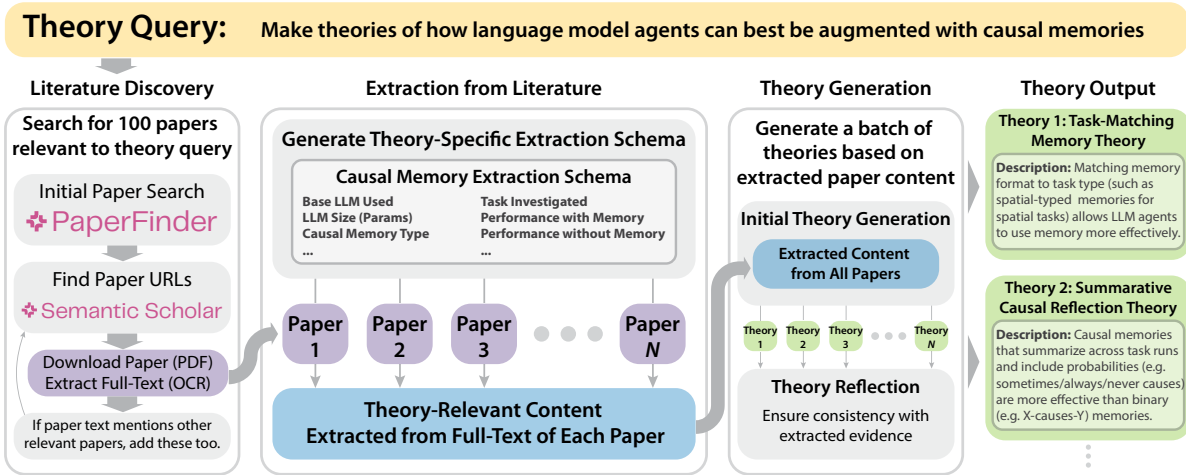


Figure 1: An overview of synthesizing theories from scientific literature with THEORIZER. A user-provided theory query guides a search for scientific papers, then theory-relevant knowledge is extracted from each paper. That knowledge is provided to a language model which generates and refines a set of theories. Full example theories are large and provided in the APPENDIX.

using literature-supported versus purely parametric knowledge, and accuracy-focused versus novelty-focused objectives. In summary, we make the following contributions:

1. We formulate the problem of literature-based scientific theory synthesis.
2. We present the THEORIZER system, which has created 2,856 theories by synthesizing evidence from 13,744 scientific papers.
3. Using THEORIZER, we systematically compare theory generation using literature-supported versus parametric knowledge, and accuracy-focused versus novelty-focused objectives. We analyze how these conditions affect theories’ specificity, empirical support, predictive accuracy, novelty, and plausibility. To measure predictive accuracy, we use a backtesting paradigm, evaluating how well each generated law matches results from 4,554 papers that were published subsequent to THEORIZER’s knowledge-cutoff.
4. Our experiments show that our literature-supported theories are significantly more accurate at matching both *existing* evidence, and also at predicting *future* results.

2 Related Work

Scientific theories are typically described as sets of qualitative and quantitative laws that predict empirical observations given initial experimental conditions, while offering some explanatory mechanism

for these observations (Džeroski et al., 2007). Prior work in automated scientific discovery has largely focused on inducing either quantitative or qualitative laws, typically from structured data rather than from the research literature. Our work addresses recent calls for integrated discovery systems that synthesize multiple forms of scientific knowledge into higher-level theories (Langley, 2024).

Quantitative Law Induction. Quantitative laws make specific numerical commitments, such as “*force is equal to mass times acceleration*”. A substantial body of work studies the discovery of quantitative laws, often framed as symbolic search or equation discovery problems predating language models (e.g. Langley et al., 1987). More recent systems combine neural and symbolic methods to rediscover known physical laws, such as Kepler’s third law, from data (Cornelio et al., 2023). Language models have also been incorporated into search-based frameworks for quantitative discovery, including evolutionary approaches such as FunSearch (Romera-Paredes et al., 2023) and LLM-SR (Shojaee et al., 2024), as well as probabilistic modeling systems like BoxLM (Li et al., 2024). Several benchmarks evaluate data-driven scientific discovery (Majumder et al., 2025; Cranmer, 2023; Gu et al., 2024). In contrast to this line of work, we focus on synthesizing theories from experimental evidence reported across up to one hundred papers in the scientific literature per theory, rather than inducing equations directly from raw data.

Qualitative Law Induction. Qualitative laws capture relational or directional regularities without

precise numerical commitments, such as “*acids and alkalis react to produce salts*”, and play a central role in scientific reasoning (Forbus, 1984, 2019). Prior work has extracted qualitative relationships from text to support tasks such as information extraction, causal modeling, and rule induction (e.g. Dagdelen et al., 2024; Friedman et al., 2022; Yang et al., 2024). Historical systems for qualitative scientific discovery typically relied on logical or symbolic induction (Langley et al., 1985; King et al., 1996). More recently, language models have been used to induce qualitative relations from parametric knowledge alone (Bosselut et al., 2019; Bouraoui et al., 2019). Despite this progress, inducing qualitative scientific laws from empirical evidence remains relatively under-explored, and in this work we synthesize empirically-grounded qualitative laws by explicitly aggregating experimental evidence across collections of scientific papers.

Hypothesis Generation. Scientific theories can also be used to generate testable predictions for future experiments in the form of hypotheses. A substantial body of prior work studies hypothesis generation from scientific literature and related evidence sources (see Zhou et al. (2025) for a recent review). SCIMON (Wang et al., 2024) shows that literature-inspired hypothesis generation can increase novelty and technical depth relative to a parametric baseline. Baek et al. (2025) use literature-grounded agents to propose and iteratively refine research ideas. Vasu et al. (2025a) show that explicitly modeling reasoning paths can improve the feasibility and empirical grounding of generated hypotheses. Liu et al. (2025b) generate hypotheses from combinations of literature and data, improving user ratings of clarity, novelty, and plausibility. In this work, we use induced theories to generate hypotheses, and evaluate theory quality in part by testing whether these hypotheses are supported by held-out literature published after the theory generation model’s knowledge cutoff.

3 Problem Formulation

We formulate the task of *theory synthesis from scientific literature* as follows. The synthesis process takes as input a user-specified *theory query* (Q) and a corpus of scientific papers (C), and produces a set of theories (T): $(Q, C) \rightarrow T$. Each theory $t_i \in T$ consists of a set of j structured <LAW, SCOPE, EVIDENCE> tuples, $\langle l_j, s_j, e_j \rangle$. A law (l_j) expresses a qualitative or quantitative relationship,

represented either as a mathematical expression or a natural-language statement (e.g., X causes Y , A increases B , $X = Y/Z$). The *scope* (s_j) specifies the conditions under which the law is expected to hold, including domain constraints or known exceptions (e.g., *applies only for small R , does not apply when P is present*). The associated *evidence* (e_j) lists empirical or observational support for the law, typically drawn from the input corpus.

In addition to its constituent laws, each theory t_i has a natural-language *theory name* (n_i) and a high-level *theory description* (d_i), which together provide contextual grounding and situate the theory within the broader scientific literature.

3.1 Theory Desiderata

A high-quality scientific theory and its component laws should include these desirable characteristics:

1. *Specificity*: Laws should make specific, testable claims and predictions.
2. *Empirical Support*: These claims should be consistent with existing empirical evidence.
3. *Predictive Accuracy*: These claims should be supported when evaluated against future or held-out experimental results.
4. *Novelty*: Laws should introduce insights not explicitly stated in prior work.
5. *Plausibility*: Laws should have a plausible scientific rationale or mechanistic explanation.

We operationalize approximate, automated measures of these desiderata in Section 5.

4 Theory Generation

Figure 1 provides an overview of THEORIZER, a system for inducing scientific theories from the research literature given a user-specified *theory query*. The system discovers relevant papers, extracts structured evidence, and uses this evidence to synthesize structured theories. A system description is below, with additional system details and hyperparameters provided in the APPENDIX, with source code provided in the source repository.

4.1 Literature Discovery

Given a theory query Q , THEORIZER builds corpus C by retrieving up to K relevant research papers to serve as evidence. The theory query is first reformulated into a literature search query and submitted to AI2 PAPERFINDER.² For each candidate paper, SEMANTIC SCHOLAR (Ammar et al., 2018;

²<https://github.com/allenai/asta-paper-finder>

Dimension	Accuracy-Focused Generation			Novelty-Focused Generation		
	Param.	Lit.	Δ	Param.	Lit.	Δ
LLM-AS-A-JUDGE						
SPECIFICITY	5.3	6.5	23% [‡]	7.9	8.0	1% [†]
EMPIRICAL SUPPORT	3.9	5.8	49% [‡]	2.0	3.5	75% [‡]
PREDICTIVE ACCURACY	Evaluated in Section 5.2					
NOVELTY	5.8	5.8	0%	6.3	6.1	-3%
PLAUSIBILITY	7.1	7.9	11% [‡]	5.0	6.3	26% [‡]
NUMBER OF SAMPLES (LAWS)	1187	927	–	799	790	–

Table 1: Average LLM-AS-A-JUDGE scores comparing parametric-generated theories (*Param.*) with literature-supported theories (*Lit.*) in two conditions: *accuracy-focused* and *novelty-focused* theory generation objectives. LLM-AS-A-JUDGE scores rate common dimensions of ideation, on a 1-10 Likert scale (1 signifies low, 10 signifies high). Base judge model is CLAUDE SONNET 4.5. Δ represents the relative difference between the parametric-only and literature-supported score, with positive signifying the literature-grounded score is higher. Inferential statistical testing using one-sided non-parametric bootstrap resampling (N=10,000 resamples); significant at $p < 0.01$ ([‡]), $p < 0.05$ ([†]).

Kinney et al., 2023) is used to identify links to open-access PDFs, which are downloaded and converted to full text using an OCR-based extraction pipeline. If fewer than K papers are obtained in the initial search, the system expands the candidate set by searching the full text of retrieved papers for additional references (i.e. crawling the citation network starting from the seed papers), ranking them by relevance, and backfilling the corpus with the highest-scoring papers. For the experiments reported here, K is set at 100 papers to balance context window token limits and extraction time.

4.2 Evidence Extraction

To gather evidence to build each theory, the system generates an extraction schema tailored to the theory query, and uses this to extract theory-relevant information from each of the K papers. The schema specifies the entities, variables, and empirical results required to support theory induction. For example, for queries concerning memory-augmented language models, the schema might include the task type examined (e.g. question answering), specific memory mechanisms used in experiments (e.g. causal memory, spatial memory, etc.), evaluation setting, and performance comparisons (such as task performance with and without using the memory). An inexpensive extraction model is used to populate the schema with document-specific values from each paper. The output of this stage is a collection of structured, JSON-formatted extraction records, which serve as input for theory synthesis. Depending on how much paper content matches the extraction schema, each paper may generate zero, one, or more than one extraction record. An example schema and

extraction record are provided in the APPENDIX.

4.3 Theory Synthesis

The theory synthesis stage aggregates extracted evidence across papers and induces candidate theories grounded in that evidence using a language model prompt. After generation, the initial batch of theories is refined using a self-reflection step (Madaan et al., 2023), which is applied independently to each theory to improve internal consistency, evidence attribution, and specificity.

In addition to the theories themselves, our system generates metadata to increase user utility. This includes identifying not only supporting evidence, but also potentially conflicting evidence for each theory. It also produces two sets of predictions in the form of hypotheses: those likely to be true given the available evidence, and higher-entropy predictions for which the generating model cannot identify strong support or contradiction in advance, and which may therefore be especially informative to test.

5 Experiments

We investigate two research questions in literature-grounded theory synthesis. First, how does having access to *literature* during the theory generation process affect theory quality over using *parametric knowledge* alone? Second, how does steering theory generation toward either *accuracy-focused* or *novelty-focused* generation objectives change the characteristics of model-generated theories.

Knowledge Window: To support the backtesting evaluation paradigm, we use a generation model with a reported knowledge cutoff of June 2024 (GPT 4.1), which provides an 18 month window

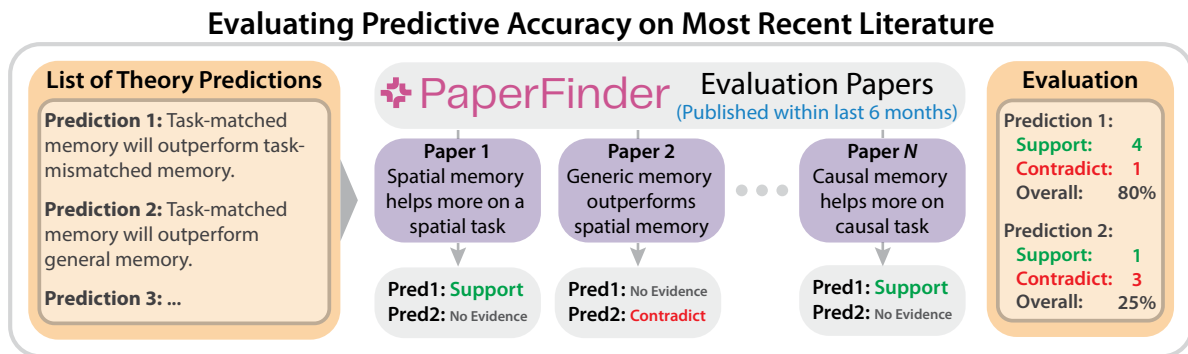


Figure 2: An overview of the *predictive accuracy* evaluation procedure. For each generated theory law, a language model is used to generate a detailed list of predictions. PAPERFINDER is used to find papers that may speak to those predictions, and each paper is rated as *supporting*, *contradicting*, or having *no evidence* for each prediction. This evidence is tallied across papers to arrive at final estimates of predictive precision and recall for a given law.

to the present date at submission time. The first 12 months of this window are used to supplement the model’s own parametric knowledge with recent literature during theory generation, while papers published in the most recent 6 months are held-out from theory generation, and used for evaluation.

Theory Queries from Literature Cross-section:

To obtain a broad and representative set of theory queries, we randomly sampled 50 papers from recent NLP and AI venues published within several months of the model’s knowledge cutoff, including ACL 2023, EMNLP 2023, AAAI 2023, and NeurIPS 2024. For each paper, we automatically generated two theory queries, one general and one specific, targeting the paper’s central research theme, yielding a total of 100 theory queries.

Theory Generation: Using these 100 theory queries, we generated four sets of theories corresponding to each of the two generation conditions: *accuracy-focused* versus *novelty-focused*, and *parametric* versus *literature-supported*. All subsequent analyses report aggregate properties within these four groups. While the 2,856 total generated theories are not restricted in the number of laws they can contain, each typically contained one or two laws. To control for variation in theory size, we perform all evaluations at the level of individual laws rather than entire theories. This resulted in between 790 and 1187 laws per condition.

5.1 Hypothesis Generation Metrics

To assess how access to literature and generation focus affect theory quality, we first use an LLM-AS-A-JUDGE paradigm (Zheng et al., 2023) to assess theory characteristics relative to the desiderata in Section 3.1, before performing more detailed

literature-based analyses for two core metrics, *predictive accuracy* and *novelty*, in the subsequent section. These LLM-AS-A-JUDGE metrics are similar to those commonly employed in the ideation literature to assess model-generated research hypotheses (Wang et al., 2023; Radensky et al., 2024; Vasu et al., 2025a; Liu et al., 2025a; Vasu et al., 2025b). All ratings are performed on a 10-point Likert scale (Likert, 1932), with a static rubric, a held-constant judge model (CLAUDE SONNET 4.5), and evaluated in isolation to avoid cross-contamination. Additional details and prompts are provided in the APPENDIX and GITHUB repository.

Table 1 summarizes results across all conditions. In the accuracy-focused setting, literature-supported theories outperform parametric theories across nearly all metrics, with particular gains in empirical support and specificity, likely owing to having access to specific details from the literature-derived evidence during theory generation. In the novelty-focused setting, literature-supported theories are rated as more empirically supported and plausible, though show no gains in specificity, likely owing to the novelty-focused prompt emphasizing making particularly specific predictions to attempt to separate these theories from the high-prior “safe-zone” accessed in the accuracy-focused condition. Across all conditions, model-assessed novelty remains approximately constant.

5.2 Predictive Accuracy

One of the core criteria for evaluating a scientific theory is whether its predictions agree with future empirical results (Popper, 1959). We operationalize this criterion using a backtesting paradigm: theories are generated using knowledge available up to June 2025, and are then evaluated against empirical

Dimension	Accuracy-Focused			Novelty-Focused		
	Param.	Lit.	Δ	Param.	Lit.	Δ
PREDICTIVE ACCURACY (GROUNDED IN LITERATURE)						
PRECISION: SUPPORTING EVIDENCE	0.88	0.90	2% [†]	0.34	0.61	76% [‡]
RECALL: LAWS WITH SOME EVIDENCE	0.45	0.51	13% [‡]	0.04	0.16	300% [‡]
RECALL: PREDICTIONS WITH SOME EVIDENCE	0.20	0.24	20% [‡]	0.01	0.06	500% [‡]
LAW-PAPER EVALUATIONS	5739	4692	–	3199	3083	–
PAPERS WITH RELEVANT EXPERIMENTS	1493	1649	–	47	221	–
LAWS WITH AT LEAST 1 RELEVANT PAPER	515	467	–	34	121	–
AVG. PAPERS PER LAW WITH EVIDENCE	2.9	3.5	–	1.4	1.8	–
SELF-ASSESSED BELIEF (LLM-AS-A-JUDGE)						
LIKELIHOOD OF LAW BEING TRUE	0.79	0.83	5% [‡]	0.67	0.76	13% [‡]

Table 2: Average *predictive accuracy* of laws as measured by evaluating each law’s predictions on recent literature published after the knowledge cutoff used for theory generation. (*top*) Measurements of *predictive precision* and *predictive recall*, evaluated using recent papers. A total of 16,713 evaluations were conducted against 4,554 unique papers (as the same paper can be retrieved when evaluating different but related laws). Definitions of predictive precision and recall are found in Section 5.2. Δ represents the relative difference between the parametric-only and literature-supported score, with positive signifying the literature-grounded score is higher. (*bottom*) Average self-assessed likelihood of law plausibility as assessed by the theory generating model. Inferential statistical testing using one-sided non-parametric bootstrap resampling (N=1,000 resamples); significant at $p < 0.01$ ([‡]), $p < 0.05$ ([†]).

results reported in papers published in the subsequent 6 months (July through December 2025).

Evaluation Procedure: Figure 2 summarizes our approach for evaluating predictive accuracy. At a high level, for each law, a language model first generates a detailed set of empirical predictions that would be expected to hold if the law were true. We then use PAPERFINDER to search for recently published papers whose experimental results may speak to these predictions. Each paper is evaluated independently to determine whether it provides evidence that supports, contradicts, or does not inform each prediction.

For a given prediction, the proportion of supporting versus contradicting evidence serves as a proxy for predictive *precision*. Similarly, the proportion of laws (or predictions) for which at least one relevant paper can be identified serves as a proxy for predictive *recall*. Together, these measures capture both how often predictions are empirically supported when tested, and how frequently predictions can be tested using recently published work.

The set of predictions is used as a grading rubric. Each prediction is represented as a schema that includes (i) the *specific prediction*, and (ii) a set of *operational signals* commonly used to measure that prediction in practice. For example, if a law predicts that a particular intervention improves machine translation performance, operational signals may include improvements in BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), or related metrics. To avoid overly permissive matching, each

prediction schema also specifies (iii) a *strong test requirement*, describing the concrete experimental comparison a paper must report to be considered a valid test (e.g., reporting machine translation performance both with and without the intervention). Finally, the schema defines (iv) explicit criteria for what constitutes *supporting* versus *contradicting* evidence. Taken together, these constraints help impose sharp boundaries on evaluation and work to prevent marginally related results from being misinterpreted as direct evidence.

Precision and Recall: For a given prediction, its precision (P_{pred}) is defined as the number of papers from corpus C providing supporting evidence (C_{supp}) divided by the total number of papers from C providing either supporting or contradicting (C_{supp}) evidence, i.e., $P_{\text{pred}} = C_{\text{supp}} / (C_{\text{supp}} + C_{\text{cont}})$. The precision of a law is computed as the average precision of its predictions, excluding predictions for which no evidence is found. Recall is defined as the proportion of laws for which at least one paper provides direct evidence evaluating one or more of its predictions. We additionally report recall at the level of individual predictions.

Self-Assessed Belief: As a complementary signal, we also estimate each law’s plausibility using the generating model’s own self-assessed belief regarding whether its predictions are likely to hold, based solely on parametric knowledge (Kadavath et al., 2022). This self-belief measure provides a model-internal estimate of likelihood, independent of direct literature grounding. We use the repeated sam-

Novelty Type	Accuracy-Focused		Novelty-Focused	
	Param.	Lit.	Param.	Lit.
QUALIFIED NOVELTY (RELATIVE TO PAPERS USED FOR THEORY GENERATION)				
PHENOMENON/EFFECT	0.69	0.36	0.99	0.90
EXPLANATION	0.79	0.49	0.96	0.90
UNIFICATION	0.80	0.68	0.97	0.92
GENERALIZATION/SCOPE EXPANSION	0.87	0.82	0.99	0.96
LIMITATION/SCOPE CONSTRAINT	0.42	0.36	1.0	0.96
CONCEPTUAL REFRAMING/ABSTRACTION	0.76	0.50	0.99	0.77
META-ANALYSIS/EMPIRICAL SYNTHESIS	0.82	0.70	1.0	0.98
NUMBER OF LAWS EVALUATED	100	100	100	100
AVERAGE PAPERS PER EVALUATION	66	66	71	71
TOTAL PAPERS EVALUATED	6632	6632	7112	7112

Table 3: *Qualified Novelty Analysis*, with novelty evaluated with respect to the input papers found by PAPERFINDER when generating a given theory. Due to pragmatic limitations of cost, each of the four generation conditions is evaluated on 100 randomly selected laws, using a total of 13,744 papers. Values represent average proportions of laws rated as having a given type of novelty, with higher proportions reflecting more novelty relative to the reference corpus C .

pling procedure of Agarwal et al. (2025) to estimate the probability that a law’s predictions are true. We report this measure alongside literature-grounded predictive accuracy to contrast the model’s internal beliefs with empirical support observed in the literature.

Results: Predictive accuracy results are shown in Table 2. CLAUDE SONNET 4.5 was used as the evaluation model. On average, 5 predictions were generated per law. Across all conditions, 2,983 laws were evaluated against the full text of 4,554 unique papers published in the most recent 6 months. As the same papers are frequently retrieved when evaluating similar laws, a total of 16,713 law–paper evaluations (5.6 papers per law) were conducted.

In the accuracy-focused condition, laws exhibit similarly high precision across both parametric and literature-supported generation (0.88 and 0.90), with recall ranging from 0.45 to 0.51. That is, for roughly half of the generated laws, at least one recent paper directly tests one or more of their predictions, and when such tests exist, the evidence overwhelmingly supports the predictions. This suggests that accuracy-focused theories can achieve broad predictive validity regardless of whether they are generated with explicit literature support, or solely using parametric knowledge.

In contrast, laws generated in the novelty-focused condition exhibit substantially lower predictive accuracy and recall. Among laws with literature-evaluable predictions, 61% of predictions from literature-supported theories are supported by recent evidence, compared to only 34% for parametric theories. Similarly, recall is substan-

tially lower in this setting: only 16% of literature-supported laws and 4% of parametric laws have any directly relevant recent papers. This pattern is consistent with the novelty-focused prompt producing more uncommon or speculative predictions that are less frequently tested in the current literature. Taken together, these results indicate that while novelty-focused theories are riskier, access to literature during the theory generation process substantially improves their predictive accuracy when evaluated using literature-grounded backtesting.

Finally, model self-assessed belief aligns well with literature-grounded predictive accuracy in the accuracy-focused condition, but not the novelty-focused condition: the model estimates that 79–83% of accuracy-focused predictions will hold, which approximately matches the observed empirical precision of 88–90%. However, in the novelty-focused condition, the generating model substantially overestimates predictive accuracy (67–76%) relative to observed support in the literature (34–61%). This discrepancy may reflect limitations of backtesting, in that novel predictions may not have been empirically tested in the literature. Similarly, it may also reflect that novelty-focused theories are inherently higher-risk by emphasizing more novel, specific, and higher-entropy laws.

5.3 Novelty Analysis

A core desideratum for theories in Section 3.1 is the novelty of their contributions relative to existing work. Exhaustively evaluating novelty against the full scientific literature is impractical at our scale. We therefore evaluate *qualified novelty*: novelty assessed *with respect to the set of retrieved papers*

used during theory generation for a given query.

Reference Corpus: Ideally, each law would be evaluated for novelty against the full literature of the field, though this is currently intractable. Accordingly, we evaluate novelty relative to the scoped reference corpus C consisting of the ≈ 66 – 71 papers (on average) retrieved by PAPERFINDER for each law’s theory query. For literature-supported generation, this corpus corresponds directly to the evidence set used during theory synthesis. For parametric generation, we evaluate novelty against the same set of retrieved papers to enable a matched comparison. We therefore interpret the resulting scores as *qualified novelty estimates*: they indicate whether a law is novel with respect to the literature most directly relevant to its generation, while acknowledging that overlap with work outside this retrieved set may exist.

Novelty Dimensions: We evaluate novelty on seven common dimensions on which scientific ideas are frequently considered as novel. PHENOMENON/EFFECT captures the discovery of a new effect or empirical regularity. EXPLANATION captures proposing a new mechanism for an existing effect. UNIFICATION captures linking previously separate concepts under a shared framework. GENERALIZATION/SCOPE EXPANSION and LIMITATION/SCOPE CONSTRAINT capture changes to the conditions under which an existing relationship is claimed to hold. CONCEPTUAL REFRAMING/ABSTRACTION captures reparameterizing or reframing a phenomenon at a different level of description. Finally, META-ANALYSIS/EMPIRICAL SYNTHESIS captures new insights obtained by synthesizing results across prior empirical studies.

Evaluation Procedure: Evaluation proceeded similarly to predictive accuracy. For each law, we evaluate novelty relative to each retrieved reference paper in C using a per-paper judge model. The resulting per-paper novelty assessments are then provided to a consolidating judge model, which produces a final decision for each novelty dimension, as well as a human-readable summary of what is known, and what appears novel. We evaluate each dimension independently (i.e. 7 consolidation calls per law) and sort per-paper results in order of descending novelty to conservatively present evidence against novelty in the portion of the context window where the model is most sensitive (Liu et al., 2024). GPT-5-MINI is used as the per-paper judge model and CLAUDE SONNET 4.5 as the con-

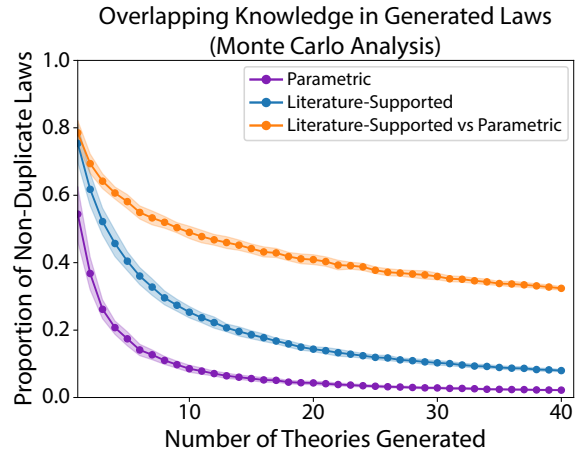


Figure 3: Monte Carlo analysis of theory law overlap when repeatedly generating theories using the same theory query. Parametric and literature-supported series measure duplicates *within* group (i.e. randomly select a *parametric* theory, then check whether it is duplicated in a random sample of N *parametric* theories). The literature-supported vs parametric series measures duplicates *across* groups (i.e. randomly select a *parametric* theory, then check whether it is duplicated in a random sample of N *literature-supported* theories), and is designed to measure whether literature-supported vs parametric theories are generating similar content. Duplication is measured using a pairwise LLM-AS-A-JUDGE using two input laws. Each point reflects the average of 50 samples. Shaded areas represent standard deviation across the 50 samples. Additional details in the APPENDIX.

solidating judge model. Due to evaluation cost (\approx \$3 per law, per condition), we evaluate a randomly selected subset of 100 laws per condition.

Results: Qualified novelty results are shown in Table 3. Across conditions, the system frequently generates laws that are judged novel relative to the reference papers used during theory generation, suggesting a nontrivial capacity to synthesize ideas not explicitly present in the retrieved literature. In the accuracy-focused setting, novelty most often arises through GENERALIZATION/SCOPE EXPANSION and least often through discovery of new PHENOMENA/EFFECTS. Literature-supported laws are generally rated as less novel than parametric laws, consistent with anchoring effects from explicit access to retrieved evidence during generation, or alternatively, a genuine capacity for pretrained models to synthesize novel contributions from parametric knowledge alone. In the novelty-focused setting, laws are rated as highly novel across nearly all dimensions, reflecting the prompt’s emphasis on novel, specific, high-entropy theory generation.

6 Discussion

Novelty versus Predictive Accuracy. Across our experiments, novelty and predictive accuracy ex-

hibit a clear tradeoff. The qualified novelty analysis shows that the *novelty-focused* condition produces highly novel laws relative to the reference literature. However, novelty in this setting reflects absence from the retrieved papers, not empirical validity. Consistent with this distinction, the predictive accuracy results indicate that novelty-focused laws are substantially riskier when evaluated using literature-based backtesting. While a subset of these laws may represent genuinely new and correct hypotheses, others are likely novel because they are speculative or inconsistent with empirical results reported in recent work. Taken together, these findings suggest that novelty-focused generation expands the space of candidate theories, but at the cost of reduced reliability under a literature-grounded backtesting-style evaluation.

Novelty and Literature Grounding. We present two methods of measuring novelty: one self-assessed by the generating model (Table 1), and the other grounded in the source literature used to generate a given theory (Table 3). That self-assessed novelty remains similar across conditions, while corpus-relative novelty differs, suggests that literature-supported theories are judged comparably creative in isolation but are more anchored to the retrieved evidence when evaluated against those retrieved papers.

Parametric Saturation and Literature-Guided Exploration. A desirable property of a theory generation system is for it to be able to generate a diverse set of theories for a given theory query. To investigate how different knowledge sources affect the diversity of generated theories, we conducted a monte-carlo analysis that shows the proportion of duplicate theories that are generated when the same theory query is repeatedly submitted. This analysis, shown in Figure 3, shows that theories generated with parametric knowledge alone quickly saturate to generating only duplicates, mirroring observations in hypothesis generation (Si et al., 2025), while also showing theories generated with knowledge from literature have a lower proportion of duplicates. Critically, when comparing theories generated *across* conditions rather than *within* conditions — that is, comparing literature-supported theories against theories generated with parametric knowledge — the proportion of non-duplicate statements after generating 40 theories is 32%, suggesting that literature-supported and parametric theories are generating laws with different content in a

large proportion of cases. This suggests that rather than one generation method being obviously better than another, parametric and literature-supported generation may serve complementary roles.

7 Conclusion

We formulate the task of synthesizing theories from corpora of research papers, demonstrate a system that generates thousands of candidate theories from tens of thousands of input papers, and analyze how theory characteristics vary under accuracy-focused or novelty-focused objectives, with either literature support or parametric knowledge alone. Our literature-supported generation method, though more expensive, creates theories with significantly greater specificity, greater empirical support, and higher predictive accuracy when compared with subsequently reported results. In future work we hope to develop further improved techniques for theory synthesis, including methods that yield higher *novelty*.

8 Limitations

Backtesting vs Experimentation: We evaluate predictive accuracy using the backtesting paradigm (Yetisgen-Yildiz and Pratt, 2009). Nominally, when a theory makes experimental predictions, these are typically evaluated by conducting a series of new experiments to find evidence that supports or contradicts those predictions. Experimental science is expensive in both time and cost, and single experiments can take months, and cost thousands of dollars and other resources. Backtesting avoids much of this expense by holding out the most recent literature during training, and using this as an evaluation set. The central trade-off is in recall: the number of papers that happen to have been published that test the specific predictions a given theory makes will be small. In this work we observe that approximately 51% of *accuracy-focused* theories have at least one paper that tests its laws predictions, though this reduces substantially to only 4% of theories in the *novelty-focused* condition. As a result, many of the predictions that the model-generated theories make are not fully tested in the literature. Further, literature has a positive-result bias: negative results often go unreported, which may make finding contradictory evidence more challenging through literature search. An individual researcher using model-generated theories to help generate novel experimental predictions to

test may be able to reasonably test an individual theory with a normative number of experiments, but theory generation at scale pragmatically requires more scalable methods of evaluation, such as backtesting. The automation of some kinds of experimental science may allow some theories to be experimentally evaluated partially automatically in the near future. However, current experimental model performance can be low – for example, CODESCIENTIST (Jansen et al., 2025) reports that when its discoveries were examined by a human domain expert, only 30% of automatic discoveries passed a code review and replication.

Long-term Theoretical Developments: Our evaluation considers only short-horizon theoretical developments. Specifically, we use an 18-month window after the generating model’s knowledge cutoff date, using the first 12 months of literature to generate theories and the subsequent 6 months to test their predictions. This is not a design limitation in our methodology, but a pragmatic consequence of contemporary frontier language models having recent knowledge cutoff dates. As such, the experiments reported here primarily captures near-term developments rather than longer-range theoretical trajectories. An important direction for future work, though currently prohibitively expensive, would be to train strong models with substantially earlier knowledge cutoffs and examine whether their theoretical predictions are borne out over periods of years or decades.

Cost: Theory generation and (particularly) evaluation is currently somewhat expensive, limiting the capacity to scale. An estimate of LLM-associated API costs is provided in the APPENDIX (See Table 4). Similarly, the theory generation pipeline takes approximately 15-30 minutes per theory query, though this can be parallelized up to API rate limits for LLM calls and paper retrieval limits.

Open Access Literature: Currently our theory generation scope is limited to fields whose scholarly literature is primarily open access and available for automated download, such as artificial intelligence and natural language processing. Similarly, searching for and downloading these source papers for generation and evaluation is the primary time component of our system, largely due to rate limits in searching for and downloading literature.

Acknowledgements

We thank the members of the Asta, Aristo, and Semantic Scholar teams at the Allen Institute for Artificial Intelligence for their discussions and support through the development of this work, as well as the three anonymous reviewers for their comments.

References

- Dhruv Agarwal, Bodhisattwa Prasad Majumder, Reece Adamson, Megha Chakravorty, Satvika Reddy Gavireddy, Aditya Parashar, Harshit Surana, Bhavana Dalvi Mishra, Andrew McCallum, Ashish Sabharwal, and Peter Clark. 2025. [Autodiscovery: Open-ended scientific discovery via bayesian surprise](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, Rodney Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler Murray, Hsu-Han Ooi, Matthew Peters, Joanna Power, Sam Skjonsberg, Lucy Lu Wang, and 4 others. 2018. [Construction of the literature graph in semantic scholar](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 84–91, New Orleans - Louisiana. Association for Computational Linguistics.
- Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2025. [ResearchAgent: Iterative research idea generation over scientific literature with large language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6709–6738, Albuquerque, New Mexico. Association for Computational Linguistics.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. [COMET: Commonsense transformers for automatic knowledge graph construction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- Zied Bouraoui, José Camacho-Collados, and S. Schockaert. 2019. [Inducing relational knowledge from bert](#). In *AAAI Conference on Artificial Intelligence*.
- Cristina Cornelio, Sanjeeb Dash, Vernon Austel, Tyler R. Josephson, Joao Goncalves, Kenneth L. Clarkson, Nimrod Megiddo, Bachir El Khadir, and L. Horesh. 2023. [Combining data and theory for derivable scientific discovery with ai-descartes](#). *Nature Communications*, 14.

- Miles Cranmer. 2023. [Interpretable machine learning for science with pysr and symbolicregression.jl](#). *ArXiv*, abs/2305.01582.
- John Dagdelen, Alex Dunn, Sanghoon Lee, Nicholas Walker, Andrew S. Rosen, G. Ceder, Kristin A. Persson, and Anubhav Jain. 2024. [Structured information extraction from scientific text with large language models](#). *Nature Communications*, 15.
- Sašo Džeroski, Pat Langley, and Ljupčo Todorovski. 2007. Computational discovery of scientific knowledge. In *Computational Discovery of Scientific Knowledge: Introduction, Techniques, and Applications in Environmental and Life Sciences*, pages 1–14.
- Kenneth D. Forbus. 1984. [Qualitative process theory](#). *Artif. Intell.*, 24:85–168.
- Kenneth D. Forbus. 2019. *Qualitative Representations: How People Reason and Learn about the Continuous World*. The MIT Press, Cambridge, MA.
- Scott Friedman, Ian Magnusson, Vasanth Sarathy, and Sonja Schmer-Galunder. 2022. [From unstructured text to causal knowledge graphs: A transformer-based approach](#). *Preprint*, arXiv:2202.11768.
- Alireza Ghafarollahi and Markus J. Buehler. 2024. [Sci-agents: Automating scientific discovery through bio-inspired multi-agent intelligent graph reasoning](#). *Advanced Materials (Deerfield Beach, Fla.)*, 37.
- Ali Essam Ghareeb, Benjamin Chang, Ludovico Mitchener, Angela Yiu, Caralyn J. Szostkiewicz, Jon M. Laurent, Muhammed T. Razzak, Andrew D. White, Michaela M. Hinks, and Samuel G. Rodrigues. 2025. [Robin: A multi-agent system for automating scientific discovery](#). *Preprint*, arXiv:2505.13400.
- Ken Gu, Ruoxi Shang, Ruijen Jiang, Keying Kuang, Richard-John Lin, Donghe Lyu, Yue Mao, Youran Pan, Teng Wu, Jiaqian Yu, Yikun Zhang, Tianmai M. Zhang, Lanyi Zhu, Mike A Merrill, Jeffrey Heer, and Tim Althoff. 2024. [BLADE: Benchmarking language model agents for data-driven science](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13936–13971, Miami, Florida, USA. Association for Computational Linguistics.
- Peter Jansen, Oyvind Tafjord, Marissa Radensky, Pao Siangliulue, Tom Hope, Bhavana Dalvi Mishra, Bodhisattwa Prasad Majumder, Daniel S Weld, and Peter Clark. 2025. [CodeScientist: End-to-end semi-automated scientific discovery with code-based experimentation](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 13370–13467, Vienna, Austria. Association for Computational Linguistics.
- John M. Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A A Kohl, Andy Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, and 15 others. 2021. [Highly accurate protein structure prediction with alphafold](#). *Nature*, 596:583 – 589.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, and 17 others. 2022. [Language models \(mostly\) know what they know](#). *Preprint*, arXiv:2207.05221.
- Ross D. King, Stephen H. Muggleton, Ashwin Srinivasan, and Michael J. E. Sternberg. 1996. [Structure-activity relationships derived by machine learning: the use of atoms and their bond connectivities to predict mutagenicity by inductive logic programming](#). *Proceedings of the National Academy of Sciences of the United States of America*, 93 1:438–42.
- Rodney Michael Kinney, Chloe Anastasiades, Russell Authur, Iz Beltagy, Jonathan Bragg, Alexandra Buraczynski, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Arman Cohan, Miles Crawford, Doug Downey, Jason Dunkelberger, Oren Etzioni, Rob Evans, Sergey Feldman, Joseph Gorney, David W. Graham, F.Q. Hu, and 29 others. 2023. [The semantic scholar open data platform](#). *ArXiv*, abs/2301.10140.
- Thomas S. Kuhn. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press.
- Pat Langley. 2024. [Integrated systems for computational scientific discovery](#). In *AAAI Conference on Artificial Intelligence*.
- Pat Langley, Herbert A. Simon, Gary L. Bradshaw, and Jan M. Zytkow. 1987. *Scientific Discovery: Computational Explorations of the Creative Process*. MIT Press.
- Pat Langley, Herbert A. Simon, Jan M. Zytkow, and Douglas H. Fisher. 1985. [Discovering qualitative empirical laws](#).
- Michael Y. Li, Emily Fox, and Noah Goodman. 2024. [Automated statistical model discovery with language models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 27791–27807. PMLR.
- Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of Psychology*, (140).
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Haokun Liu, Sicong Huang, Jingyu Hu, Yangqiaoyu Zhou, and Chenhao Tan. 2025a. [Hypobench: Towards systematic and principled benchmarking for hypothesis generation](#). *Preprint*, arXiv:2504.11524.

- Haokun Liu, Yangqiaoyu Zhou, Mingxuan Li, Chenfei Yuan, and Chenhao Tan. 2025b. [Literature meets data: A synergistic approach to hypothesis generation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 245–281, Vienna, Austria. Association for Computational Linguistics.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob N. Foerster, Jeff Clune, and David Ha. 2024. [The ai scientist: Towards fully automated open-ended scientific discovery](#). *ArXiv*, abs/2408.06292.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Sanchaita Hazra, Ashish Sabharwal, and Peter Clark. 2024. [Position: Data-driven discovery with large generative models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 34350–34382. PMLR.
- Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Bhavana Dalvi Mishra, Abhijeetsingh Meena, Aryan Prakhar, Tirth Vora, Tushar Khot, Ashish Sabharwal, and Peter Clark. 2025. [Discoverybench: Towards data-driven discovery with large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- Amil Merchant, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk. 2023. [Scaling deep learning for materials discovery](#). *Nature*, 624:80 – 85.
- Adam Moss. 2025. [The ai cosmologist i: An agentic system for automated data analysis](#). *ArXiv*, abs/2504.03424.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Karl R. Popper. 1959. *The Logic of Scientific Discovery*. Harper & Row, New York.
- Marissa Radensky, Simra Shahid, Raymond Fok, Pao Siangliulue, Tom Hope, and Daniel S. Weld. 2024. [Scideator: Human-llm scientific idea generation grounded in research-paper facet recombination](#). *ArXiv*, abs/2409.14634.
- Bernardino Romera-Paredes, M. Barekatin, Alexander Novikov, Matej Balog, M. P. Kumar, Emilien Dupont, Francisco J. R. Ruiz, J. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, Alhussein Fawzi, Josh Grochow, Andrea Lodi, Jean-Baptiste Mouret, Talia Ringer, and Tao Yu. 2023. [Mathematical discoveries from program search with large language models](#). *Nature*, 625:468 – 475.
- Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Michael Moor, Zicheng Liu, and Emad Barsoum. 2025. [Agent laboratory: Using LLM agents as research assistants](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5977–6043, Suzhou, China. Association for Computational Linguistics.
- Parshin Shojaee, Kazem Meidani, Shashank Gupta, A. Farimani, and Chandan K. Reddy. 2024. [Llm-sr: Scientific equation discovery via programming with large language models](#). *ArXiv*, abs/2404.18400.
- Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2025. [Can LLMs generate novel research ideas? a large-scale human study with 100+ NLP researchers](#). In *The Thirteenth International Conference on Learning Representations*.
- M. Thilakarathne, K. Falkner, and Thushari Atapattu. 2019. [A systematic review on literature-based discovery workflow](#). *PeerJ Computer Science*, 5.
- Alexander V. Tobias and Adam Wahab. 2025. [Autonomous ‘self-driving’ laboratories: a review of technology and policy implications](#). *Royal Society Open Science*, 12.
- Rosni Vasu, Chandrayee Basu, Bhavana Dalvi Mishra, Cristina Sarasua, Peter Clark, and Abraham Bernstein. 2025a. [HypER: Literature-grounded hypothesis generation and distillation with provenance](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25424–25449, Suzhou, China. Association for Computational Linguistics.
- Rosni Vasu, Peter Jansen, Pao Siangliulue, Cristina Sarasua, Abraham Bernstein, Peter Clark, and Bhavana Dalvi Mishra. 2025b. [Harpa: A testability-driven, literature-grounded framework for research ideation](#). *Preprint*, arXiv:2510.00620.
- Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. 2023. [Scimon: Scientific inspiration machines optimized for novelty](#). In *Annual Meeting of the Association for Computational Linguistics*.

Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. 2024. [SciMON: Scientific inspiration machines optimized for novelty](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 279–299, Bangkok, Thailand. Association for Computational Linguistics.

Zonglin Yang, Li Dong, Xinya Du, Hao Cheng, Erik Cambria, Xiaodong Liu, Jianfeng Gao, and Furu Wei. 2024. [Language models as inductive reasoners](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 209–225, St. Julian’s, Malta. Association for Computational Linguistics.

Meliha Yetisgen-Yildiz and Wanda Pratt. 2009. [A new evaluation methodology for literature-based discovery systems](#). *Journal of Biomedical Informatics*, 42(4):633–643.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.

Zekun Zhou, Xiaocheng Feng, Lei Huang, Xiachong Feng, Ziyun Song, Ruihan Chen, Liang Zhao, Weitao Ma, Yuxuan Gu, Baoxin Wang, Dayong Wu, Guoping Hu, Ting Liu, and Bing Qin. 2025. [From hypothesis to publication: A comprehensive survey of AI-driven research support systems](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 11773–11803, Suzhou, China. Association for Computational Linguistics.

A Supplementary Material

The repository and supplementary material³ includes:

Theories: An example set of theories, extraction schema, extraction results, predictive accuracy evaluations, and qualified novelty evaluations.

Code: The GITHUB repository includes full code, getting-started instructions, a user interface, and API. The documentation also points to all major prompts used for generation and evaluation.

B Follow-Through Example

While the supplementary material includes hundreds of model-generated theories, schemas, and evaluations in easily-viewable HTML format, the following follow-through example on Intelligent Tutoring Systems is provided in the APPENDIX:

³<https://github.com/allenai/asta-theorizer>

1. Generated Theory: Table 5
2. Predictive Accuracy Evaluation: Table 6
3. Qualified Novelty Evaluation: Tables 7 and 8
4. Extraction Schema: Table 9
5. Extraction from Paper: Table 10

Note that these supplementary materials have been reduced for space, and we encourage viewing the original source materials provided in the GITHUB repository and data release.

C Theory Generation Hyperparameters

The theory generation procedure is described in Section 4, with full details including prompts in the code release.

Models: The generation models are user-selectable in the THEORIZER user interface and API. In the experiments reported here, the generation model was GPT-4.1, which is used for all aspects of theory generation (i.e. extraction schema generation, theory generation, theory reflection) except for extracting evidence from individual papers, which is performed by an inexpensive model (GPT-5-MINI) due to the scale required. Generation temperature was set to 0 in the high-accuracy condition, and 1 in the novelty-focused condition.

Self-Evaluated Novelty and Filtering: Independent of the qualified novelty evaluation in Section 5.3, the theory generation process itself requires generated theories to include self-assessments of novelty provided by the generating model, to encourage the model to generate laws that are not repetitions of existing/well-known laws. This assessment includes a chain-of-thought style assessment of the novelty (*what already exists, what is novel, provide an explanation/reasoning as to how the novelty of this law could be classified*), followed by a categorical assessment of the law’s novelty using one of 4 classification labels. Those labels, in order of decreasing novelty, are: 1. NEW, 2. SOMEWHAT-RELATED-TO-EXISTING, 3. CLOSELY-RELATED-TO-EXISTING, 4. EXISTING. To focus our evaluation primarily on novel theories/laws, we filtered any laws rated as CLOSELY-RELATED-TO-EXISTING or EXISTING from further analysis. The proportion of these laws filtered varied between 1% and 18%, depending on generation condition. The total number of laws reported in Table 1 is the total after this filtering.

Subsampling Evidence: Depending on the number of papers retrieved and the amount of evidence

in each paper, the total size (in tokens) of all retrieved evidence can extend beyond the generating model’s input context token limit. When this happens, the evidence is randomly subsampled.

Evaluation Models: Unless otherwise stated, the evaluation model for all evaluations was CLAUDE SONNET 4.5. In cases where a large number of individual evaluations was required (e.g. evaluations against a corpus of papers), this model is sometimes augmented with a more inexpensive model, GPT-5-MINI, and these cases are described in text.

D Generation Costs

Estimated LLM API costs are included in Table 4.

Activity	Cost Estimate
Theory Generation	
<i>Literature-Supported</i>	\$0.26 per theory
<i>Parametric Only</i>	\$0.04 per theory
LLM-as-a-Judge Metrics	\$0.03 per law
Predictive Accuracy	\$0.39 per law (\approx 5 papers)
Surprisal Analysis	\$0.03 per law
Novelty Evaluation	\$2.46 per law (\approx 70 papers)
Overlap Analysis	\$10 per law

Table 4: Estimated LLM API costs for the generation and evaluation of theories in this work.

E Theory Query Generation

Theory queries were abstracted from randomly selected scientific articles authored within 1 year prior to the base generation model’s (GPT-4.1) advertised knowledge cutoff of June 2024. The source venues included ACL 2023, EMNLP 2023, AAAI 2023, and NeurIPS 2024. Example theory queries are shown in Table 11.

F Knowledge Overlap Analysis

An example of the duplicate law evaluation schema is shown in Table 12, for two cases: two laws judged as duplicates, and two laws judged as not duplicates. The judge model was GPT-5-MINI. Due to cost, only 9 theory queries were included in the monte-carlo analysis, with a total of over 43k individual law-law comparisons.

G Broader Impacts and Ethical Considerations

Automated scientific discovery offers the promise of accelerating scientific developments, with high-impact recent successes in protein folding (Jumper et al., 2021) and materials discovery (Merchant

et al., 2023). Augmenting problem-specific discovery methods with information grounded in the scientific literature enables more capable, integrated end-to-end discovery systems (e.g. Ghareeb et al., 2025). In this work, we propose shifting some focus from automated experiment-driven discovery toward automated theory building through literature-based discovery. While theories can provide high-value guidance for future experiments, more fundamentally they aim to compress the knowledge within a scientific domain into a set of governing laws that accurately predict the outcomes of future experiments. Having accurate theories and laws for a given subdomain allows it to transition into a principled engineering discipline, and to more systematically translate empirically observed regularities into useful and impactful technologies. That being said, automated theory generation systems like the one proposed in this work have risks, such as the potential to generate partially-accurate, inaccurate, misleading, or hallucinated theories, and generated output such as candidate theories should be confirmed using standard scientific protocols before downstream use in any application.

Example Theory: Emergent Multi-Agent Social Presence Theory in LLM-ITS (Intelligent Tutoring Systems)	
Theory Query	
Theory Query	Build a theory of how the explicit integration of step-by-step scaffolding and dynamic conversational strategies, as operationalized in the CLASS framework, influences student cognitive gains, engagement, and motivation in LLM-powered intelligent tutoring systems across diverse subject domains.
Theory (General)	
Generation Condition	<i>literature-supported, accuracy-focused</i>
Theory ID	theory-215
Theory Name	Emergent Multi-Agent Social Presence Theory in LLM-ITS (Intelligent Tutoring Systems)
Theory Description	This theory asserts that the explicit integration of multi-agent conversational strategies (e.g., simulated classmates, assistants, note-takers) in LLM-powered ITS can create emergent social presence and collaborative learning dynamics, which in turn enhance engagement, motivation, and, under certain conditions, cognitive gains. The effect is mediated by the diversity of agent roles, the orchestration of turn-taking, and the degree of user participation. However, the benefits are domain- and context-dependent, and excessive agent complexity or poorly managed interaction can reduce learning efficiency.
One Law/Statement	
Law Name	Multi-Agent Social Presence Law
Law Statement	In LLM-powered ITS, the presence of multiple simulated agent roles (e.g., teacher, assistant, classmates) and dynamic conversational orchestration increases perceived social presence, engagement, and motivation, as measured by interaction metrics and self-report, compared to single-agent or non-interactive systems.
Law Type	<i>qualitative</i>
Scope/Domain	LLM-powered ITS and classroom simulations with multi-agent conversational architectures, across higher education and K-12 domains, especially in settings where social/collaborative learning is valued.
Special Cases	<ol style="list-style-type: none"> 1. In highly individual or assessment-driven tasks, social presence may have less impact on cognitive gains. 2. For learners with social anxiety, increased agent interaction may not increase motivation.
Supporting Evidence	<ol style="list-style-type: none"> 1. SimClass multi-agent classroom ablation showed that removing classmate agents reduced user speech length by 26.5% (TAGI) and 45.2% (HSU), and reduced Community of Inquiry (CoI) social and cognitive presence scores; full multi-agent systems had higher engagement and learning experience. (UUIDs: e2892.0) 2. SimClass FIAS coding showed high Student Initiation Ratios (SIR 0.9), indicating active participation in multi-agent settings. (UUIDs: e2892.0) 3. SRLAgent's gamified, multi-agent orchestration (Planning Agent, SubTask Tutor, Reflection Agent) increased engagement and SRL skills compared to baseline multimedia learning. (UUIDs: e2717.0) 4. EnglishBot's open conversational practice (simulated dialogue) led to greater engagement and learning gains than a listen-and-repeat interface. (UUIDs: e2760.3) 5. DBTS (Discussion-Based Teaching Systems) report 72% increase in engagement and 74% improvement in learning outcomes, attributed to dialogic, multi-agent interaction. (UUIDs: e2771.2)
Self-Assessed Law Novelty (<i>produced as part of theory generation; independent of later novelty evaluation</i>)	
What Already Exists	Social presence and collaborative learning are established in educational theory, and multi-agent systems have been explored in ITS, but the emergent effects of LLM-powered multi-agent orchestration on engagement and learning are not formalized.
What Is Novel	The explicit law linking multi-agent LLM orchestration to emergent social presence and its downstream effects in ITS is novel.
Classification Expl.	Existing work discusses social presence and collaborative learning, but not the emergent effects of LLM-driven multi-agent orchestration in ITS.
Classification	<i>new</i>
LLM-Generated Refs.	<ol style="list-style-type: none"> 1. Garrison et al. (2000) Critical Inquiry in a Text-Based Environment: Computer Conferencing in Higher Education [Community of Inquiry framework] 2. Dillenbourg (1999) Collaborative learning: Cognitive and computational approaches [collaborative learning, not LLM multi-agent orchestration]

Table 5: An example theory generated in this work, including the theory name, description, a single law, and a self-assessment of novelty made from the generation model. An example predictive accuracy evaluation for this theory is provided in Table 6, and the qualified novelty evaluation relative to the source papers used to generate this theory is in Table 7. The extraction schema used to generate this theory is in Table 9, and example extracted evidence shown in Table 10. Hundreds of additional examples are available in HTML format in the Supplementary Material.

Predictive Accuracy Example: Emergent Multi-Agent Social Presence Theory in LLM-ITS (Intelligent Tutoring Systems)	
Evaluation (General)	
Theory ID	theory-215
Theory Name	Emergent Multi-Agent Social Presence Theory in LLM-ITS
Theory Statement	In LLM-powered ITS, the presence of multiple simulated agent roles (e.g., teacher, assistant, classmates) and dynamic conversational orchestration increases perceived social presence, engagement, and motivation, as measured by interaction metrics and self-report, compared to single-agent or non-interactive systems.
Papers Examined	14 papers published between July 2025-December 2025 were examined.
Average Precision	1.0 (of the 2 rubric items with some evidence found)
Rubric Evaluation (Per Prediction) (Note: Only 3 of 5 rubric items shown for space)	
Prediction Short Name	agent_role_diversity_effect
Specific Prediction	Within multi-agent LLM-ITS, increasing the diversity of agent roles (e.g., adding classmate agents to teacher+assistant) increases engagement or social presence compared to reduced-role configurations.
Operational Signals	<ol style="list-style-type: none"> 1. Ablation studies varying number or types of agent roles 2. Engagement or social presence metrics across different role configurations 3. Explicit manipulation of agent diversity (e.g., teacher-only vs teacher+classmates)
Strong Test Requirement	A paper must report an ablation or comparison where agent role diversity is systematically varied, with engagement or social presence outcomes measured for each configuration.
What Does Support Look Like	Adding diverse agent roles (e.g., classmates) significantly increases engagement metrics or social presence scores compared to fewer-role configurations.
What Does Contradiction Look Like	Removing or adding agent roles has no significant effect on engagement or social presence, or reduced-role systems perform equivalently or better.
Counts	support: 1 paper, contradict: 0 papers, no_evidence: 13 papers
Proportions	support: 1.0, contradict: 0.0
Evidence	<p>1. <i>Source Paper</i>: OnlineMate: An LLM-Based Multi-Agent Companion System for Cognitive Support in Online Learning (Xian Gao, Zongyun Zhang, Ting Liu, Yuzhuo Fu; 2025/9; ArXiv: 2509.14803)</p> <p><i>Evidence Quote or Locator</i>: Figure 5 and associated text: 'Figure 5 illustrates the impact of varying the number of OnlineMate Agents on the highest cognitive level attained by students. As the teaching assistant role transitions from a single agent to multiple agents engaged in classroom discussions, a noticeable improvement in students' highest cognitive levels is observed... However, when the number of agents exceeds four, the increase in cognitive level begins to plateau, and a decline is even noted when the number of agents surpasses six.'</p> <p><i>Evidence in Support</i>: The paper reports an ablation study systematically varying the number of agents (from 1 to 6+) and measuring cognitive level outcomes. Results show that increasing from single to multiple agents (up to 4) produces noticeable improvements in cognitive levels, directly supporting the prediction that agent role diversity increases engagement/outcomes.</p>
Prediction Short Name	assessment_task_boundary
Specific Prediction	In highly individual or assessment-driven tasks (e.g., standardized testing, individual skill drills), multi-agent social presence has reduced or no significant impact on cognitive learning gains (e.g., test scores, knowledge assessments) compared to collaborative or exploratory learning contexts.
Operational Signals	<ol style="list-style-type: none"> 1. Comparison of multi-agent effects across task types: assessment-driven vs collaborative/exploratory 2. Learning outcome measures (test scores, knowledge gains) in assessment contexts 3. Explicit categorization of tasks as individual/assessment vs collaborative
Strong Test Requirement	A paper must compare multi-agent vs single-agent LLM-ITS effects on cognitive outcomes in an explicitly assessment-driven or individual task context, with clear outcome measures.
What Does Support Look Like	Multi-agent presence shows minimal or no significant advantage on cognitive gains in assessment-driven tasks, even if engagement increases.
What Does Contradiction Look Like	Multi-agent systems produce significant cognitive gains in highly individual or assessment-driven tasks comparable to those in collaborative contexts.
Counts	support: 0 papers, contradict: 0 papers, no_evidence: 14 papers
Proportions	support: 0, contradict: 0
Evidence	<i>No evidence found in recent papers (last 6 months) that test this prediction.</i>
Prediction Short Name	excessive_complexity_inefficiency
Specific Prediction	LLM-powered ITS with excessive agent complexity or poorly managed interaction (e.g., unclear turn-taking, overwhelming number of agents, confusing orchestration) reduces learning efficiency, as measured by time-on-task, task completion rates, or learning gains per unit time, compared to well-orchestrated multi-agent systems.
Operational Signals	<ol style="list-style-type: none"> 1. Comparison of different orchestration strategies or agent complexity levels 2. Learning efficiency metrics: time to mastery, learning gains per minute, task completion rates 3. User confusion, cognitive load, or negative feedback related to agent interaction complexity
Strong Test Requirement	A paper must compare different levels of agent complexity or orchestration quality in multi-agent LLM-ITS, measuring learning efficiency or related outcomes.
What Does Support Look Like	Overly complex or poorly orchestrated multi-agent systems show reduced efficiency metrics (longer time, lower completion rates, or reduced gains per unit time) compared to simpler or well-managed configurations.
What Does Contradiction Look Like	Increased agent complexity consistently improves or has no negative effect on learning efficiency, regardless of orchestration quality.
Counts	support: 1 paper, contradict: 0 papers, no_evidence: 13 papers
Proportions	support: 1.0, contradict: 0.0
Evidence	<p>1. <i>Source Paper</i>: OnlineMate: An LLM-Based Multi-Agent Companion System for Cognitive Support in Online Learning (Xian Gao, Zongyun Zhang, Ting Liu, Yuzhuo Fu; 2025/9; ArXiv: 2509.14803)</p> <p><i>Evidence Quote or Locator</i>: Figure 5 and associated text: "However, when the number of agents exceeds four, the increase in cognitive level begins to plateau, and a decline is even noted when the number of agents surpasses six. This suggests that, in online learning contexts, an excessive number of agent companions may not necessarily be beneficial. An overabundance of agents providing excessive information may lead to cognitive overload, hindering the students"</p> <p><i>Evidence in Support</i>: The paper directly tests varying levels of agent complexity (number of agents) and finds that excessive agents (>6) lead to declining cognitive outcomes, explicitly attributing this to cognitive overload from too much information. This directly supports the prediction that excessive complexity reduces learning efficiency.</p>

Table 6: Example Predictive Accuracy evaluation, including the rubric, and supporting evidence for two rubric elements.

Qualified Novelty Evaluation: Emergent Multi-Agent Social Presence Theory in LLM-ITS (Intelligent Tutoring Systems)	
Evaluation (General)	
Theory ID	theory-215
Theory Name	Emergent Multi-Agent Social Presence Theory in LLM-ITS
Theory Statement	In LLM-powered ITS, the presence of multiple simulated agent roles (e.g., teacher, assistant, classmates) and dynamic conversational orchestration increases perceived social presence, engagement, and motivation, as measured by interaction metrics and self-report, compared to single-agent or non-interactive systems.
Papers Evaluated	54 papers (<i>that were used as literature-support in the generation of this theory</i>)
Final Aggregated Qualified Novelty Evaluation (Per Dimension)	
Dimension	phenomenon_effect
What Is Known	Three to four recent papers (Ruffle&Riley studies and SimClass, 2023-2024) provide direct experimental evidence that multi-agent LLM conversational architectures with distinct simulated roles (teacher, assistant, classmates) produce statistically significant increases in self-reported engagement, perceived social presence (CoI framework), and interaction metrics compared to single-agent or non-interactive baselines. These studies report quantitative effect sizes across multiple validated instruments (CoI subscales, FIAS coding, message counts, satisfaction scores) in controlled experiments and semester-long deployments, establishing the phenomenon empirically within their specific implementations.
What Introduced	The Multi-Agent Social Presence Law articulates this as a generalizable empirical regularity: that explicit integration of multiple simulated agent roles with dynamic conversational orchestration systematically produces emergent social presence and collaborative learning dynamics that increase perceived engagement and motivation across LLM-powered ITS contexts. The theory names and frames this as a distinct, measurable phenomenon mediated by role diversity, turn-taking orchestration, and user participation, extending beyond isolated implementations to propose a general principle.
What Novel	While the core phenomenon has been explicitly demonstrated with rigorous experimental evidence in 3-4 papers, the overwhelming majority of the LLM-ITS literature (43/59 papers assessed) does not study, measure, implement, or discuss multi-agent social presence as a distinct effect. Most work focuses on single-agent personalization, hint generation, or pedagogical strategies without investigating role-diverse architectures or emergent social-presence dynamics. The phenomenon exists as an explicitly reported finding in isolated recent studies but has not achieved field-wide recognition, systematic replication across contexts, or treatment as established knowledge; experts outside these specific research groups would likely not cite it as a known regularity.
Degree of Novelty	explicit_peripheral
Dimension	explanatory_mechanistic
What Is Known	The literature documents various mechanisms in isolation or within single-agent paradigms: LLM tutors employ prompt engineering, scaffolding, and adaptive feedback; traditional ITS use learner modeling and personalization; some systems implement multi-agent architectures with role-specific prompts and orchestration (2 papers peripherally mention related components); social presence and engagement are recognized constructs. However, these elements are analyzed separately, and no paper articulates a unified causal mediation model linking multi-agent architectural features to social-psychological mediators to learning outcomes.
What Introduced	The theory proposes a novel causal mediation framework specific to multi-agent LLM-powered ITS: diversity of agent roles (teacher, peers, assistants), orchestration of conversational turn-taking, and degree of user participation act as joint architectural mediators that produce emergent social presence as a psychological mechanism, which in turn causally increases engagement, motivation, and (under boundary conditions) cognitive gains. This articulates an integrated causal chain: multi-agent architecture → emergent social presence → affective/motivational outcomes → conditional cognitive effects.
What Novel	Across 54 papers, the vast majority (44 papers, 81%) propose entirely different mechanisms (single-agent scaffolding, prompt engineering, personalization algorithms) or do not address multi-agent social dynamics at all. Two papers mention related mechanistic components peripherally but do not formalize them as causal models; eight papers acknowledge the components are derivable but unstated. Critically, no paper in the corpus articulates or tests the specific causal mediation structure proposed by the theory—that multi-agent role diversity and orchestration jointly produce emergent social presence as the proximate cause of engagement and learning gains. This represents a genuinely new explanatory framework that synthesizes previously disconnected architectural and psychological constructs into a testable causal theory.
Degree of Novelty	genuinely_new
Dimension	unification
What Is Known	The literature contains separate strands discussing ITS architectures, virtual agents, social presence theory in online learning, and collaborative learning dynamics. Multiple systematic reviews and surveys synthesize these areas and provide taxonomies (e.g., three-class tutoring-system taxonomy, pedagogical frameworks, LLM-education taxonomies), and several papers unify technical components within single-agent systems (e.g., LLM+ITS integration, student modeling, prompt engineering). However, no reviewed paper explicitly elevates multi-agent conversational architectures as a unifying design principle that subsumes engagement and motivation outcomes under an emergent social-presence framework.
What Introduced	The theory explicitly unifies multi-agent conversational strategies (simulated teacher, assistant, classmates, note-takers) and their orchestration as instances of a single underlying principle: that role diversity and turn-taking dynamics in LLM-powered ITS produce emergent social presence, which causally drives measurable improvements in engagement, motivation, and collaborative learning dynamics across educational contexts. This unification treats previously separate phenomena (agent-based tutoring, social presence effects, collaborative learning gains) as manifestations of one coherent mechanism.
What Novel	While component ideas exist separately across the literature—multi-agent systems in education, social presence theory, ITS engagement measures, collaborative learning benefits—the explicit unifying principle that diverse multi-agent conversational role architectures are manifestations of a single emergent social-presence mechanism is not stated in any reviewed paper. This specific conceptual unification could be derived by synthesizing insights across multiple papers (especially reviews discussing agents, ITS design, engagement theory, and collaborative learning), but it has not been explicitly formulated, validated, or presented as a unified theoretical principle in prior work. The shift from separate strands to an integrated explanatory framework represents a novel synthesis.
Degree of Novelty	derivable_unstated
Continued in Table 8	

Table 7: Example Qualified Novelty aggregated evaluation on 7 novelty dimensions. Due to space limitations, individual evaluations of the 54 source papers examined to create this evaluation are not shown, and evaluation for the remaining 4 novelty dimensions continues on Table 8.

Qualified Novelty Evaluation (Continued): Emergent Multi-Agent Social Presence Theory in LLM-ITS (Intelligent Tutoring Systems)	
Dimension	generalization_scope_expansion
What Is Known	The surveyed literature (54 papers) demonstrates LLM-based educational effectiveness in narrow, specific contexts: predominantly single-agent architectures evaluated within particular domains (programming, mathematics, language learning), specific educational levels (individual courses or grade levels), and limited populations. Multiple review papers discuss broad AI/ITS applicability across educational settings but do not establish that multi-agent conversational architectures produce consistent effects. Only one paper (Simulating Classroom Education) implements multi-agent LLM classroom simulation, but strictly in two university courses with explicit limitations on generalizability; 53 other papers use single-agent designs or do not address multi-agent architectures at all.
What Introduced	The theory simultaneously generalizes across three independent axes: (1) architectural: from single-agent tutors to multi-agent conversational systems with diverse simulated roles (teacher, assistant, classmates, note-takers) and dynamic turn-taking orchestration; (2) educational scope: across K-12 and higher education levels; (3) domain scope: across contexts where social/collaborative learning is valued. It asserts these multi-agent architectures produce emergent social presence, engagement, and motivation benefits compared to single-agent or non-interactive systems, with specified boundary conditions (e.g., assessment-driven tasks may show reduced benefits, learners with social anxiety may not benefit from increased interaction).
What Novel	The aggregated evidence shows 31/54 papers (57%) rate this scope expansion as genuinely new relative to their work, 21/54 (39%) as derivable but unstated, and only 2/54 (4%) as explicit but peripheral. The vast majority explicitly limit claims to their specific domains, populations, or single-agent architectures and provide no empirical or analytical evidence for the theory's three-fold generalization. The claim that multi-agent conversational architectures in LLM-powered ITS produce consistent cross-domain social-presence and engagement effects across educational levels is not established, demonstrated, or derived in the surveyed literature; while plausible extensions of component findings exist, their unification into a broad multi-agent, cross-context law represents a genuine scope expansion.
Degree of Novelty	genuinely_new
Dimension	constraint_limitation
What Is Known	Three Ruffle&Riley papers explicitly establish that multi-agent conversational ITS can increase engagement and perceived social presence without corresponding short-term learning gains, while incurring time costs and interaction quality problems (lenient feedback, repeated prompts). Broader literature documents general LLM/ITS limitations (hallucinations, bias, over-reliance risks, need for human oversight) but primarily for single-agent systems. Context-dependence of educational technology effectiveness across learner populations and task types is widely acknowledged as a general principle.
What Introduced	The theory explicitly enumerates a systematic set of multi-agent-specific constraints: (1) benefits are domain- and context-dependent, with particular reductions in highly individual or assessment-driven tasks; (2) excessive agent complexity or poorly managed multi-agent orchestration constitutes a distinct failure mode that reduces learning efficiency; (3) learner traits, specifically social anxiety, moderate or reverse motivational benefits of increased agent interaction. These are framed as necessary boundary conditions for multi-agent conversational architectures rather than general LLM caveats.
What Novel	The basic dissociation between engagement and learning gains is empirically established by Ruffle&Riley work. However, the identification of *excessive multi-agent complexity* and *orchestration failures* as specific, nameable failure modes is not explicitly recognized in the literature—it is at best derivable from scattered observations about interaction quality and complexity costs across 25 papers. The social-anxiety moderator and the explicit task-type taxonomy (assessment-driven vs. collaborative) as systematic boundary conditions are not stated in 48 of 63 papers (76%). The theory performs a novel synthesis: elevating derivable-but-unstated constraints into an explicit, testable framework of multi-agent-specific failure regimes.
Degree of Novelty	derivable_unstated
Dimension	conceptual_reframing_abstraction
What Is Known	The literature contains implementations of multi-agent and role-based LLM architectures (3 papers show explicit but peripheral mentions), established research on social presence and collaborative learning as pedagogical constructs, and various single-agent abstractions for ITS design (scaffolding taxonomies, personalization pipelines, prompt engineering frameworks). However, these elements exist as separate threads: multi-agent designs are typically framed as engineering decompositions for task distribution rather than as social-presence generators, and social presence research has not been systematically connected to multi-agent architectural choices as a primary ITS design principle across the surveyed literature.
What Introduced	The theory elevates 'emergent multi-agent social presence' from scattered implementations to a formalized, named organizing principle (the 'Multi-Agent Social Presence Law') that should serve as the central abstraction for LLM-ITS design. It proposes that intentional engineering of diverse simulated agent roles combined with dynamic conversational orchestration produces social presence as a predictable emergent property, and positions role diversity, turn-taking orchestration, and participation management as first-class design variables that mediate engagement, motivation, and learning outcomes. This reframes prior work by treating what were previously implementation details or secondary outcomes as the primary explanatory framework linking architecture to educational impact.
What Novel	The conceptual reframing is genuinely novel: 75% of papers (42/56) judge this organizing abstraction as absent from their frameworks and not derivable as a straightforward consequence of their work. While 11 papers (20%) find it derivable but unstated and 3 papers (5%) show explicit peripheral mentions of related concepts, no paper in the surveyed literature formalizes multi-agent role orchestration and emergent social presence as the primary organizing principle for ITS design. The novelty lies in proposing a new conceptual lens that reorganizes how diverse findings should be interpreted, designed, and evaluated—shifting from agent-as-tool or agent-as-implementation-detail to multi-agent-social-ecosystem-as-central-construct, with testable predictions about how architectural choices (role diversity, orchestration patterns) map to social-presence and motivational outcomes.
Degree of Novelty	genuinely_new
Dimension	empirical_synthesis_meta_regularity
What Is Known	General ITS effectiveness is well-established through multiple meta-analyses showing cognitive and engagement benefits. Single-agent LLM tutors have been shown in individual studies to improve various outcomes (engagement, satisfaction, learning gains). Three papers (7%) provide within-study evidence that multi-agent or persona-based approaches can improve user experience or engagement in specific contexts (DHH persona preferences, GenMentor deployment satisfaction, Ruffle&Riley UX ratings), but these are single-system results without cross-study generalization. Eight review papers (19%) synthesize LLM/ITS literature broadly but perform no meta-analyses comparing multi-agent versus single-agent architectures on social presence or engagement outcomes.
What Introduced	The theory asserts a cross-deployment empirical meta-regularity: that multi-agent simulated roles with dynamic conversational orchestration systematically and reliably increase perceived social presence, engagement, and motivation compared to single-agent or non-interactive systems across LLM-powered ITS implementations, measurable via interaction metrics and self-reports, with effects mediated by role diversity and orchestration quality. This claim requires aggregation of comparative evidence across multiple independent studies or deployments to establish the regularity.
What Novel	The claimed meta-regularity is not established in the surveyed literature. No paper synthesizes cross-study empirical evidence demonstrating that multi-agent architectures consistently outperform single-agent systems on social presence and engagement metrics across contexts. The vast majority of papers (32/43, 74%) find this claim entirely novel to their work, indicating they provide no relevant evidence. The three papers with supportive within-study patterns cannot establish a cross-study regularity, and the eight reviews do not perform the necessary comparative meta-analysis. Establishing this meta-regularity would require novel systematic comparative experiments across implementations or formal meta-analytic synthesis of multi-agent versus single-agent effects, neither of which exists in the current literature base.
Degree of Novelty	genuinely_new

Table 8: Example Qualified Novelty aggregated evaluation on 7 novelty dimensions. Continued from Table 7.

Example Evidence Extraction Schema (Intelligent Tutoring Systems)	
Theory Query	
Theory Query	Build a theory of how the explicit integration of step-by-step scaffolding and dynamic conversational strategies, as operationalized in the CLASS framework, influences student cognitive gains, engagement, and motivation in LLM-powered intelligent tutoring systems across diverse subject domains.
Extraction Schema Definition	
Extraction Query	Extract any mentions of LLM-powered intelligent tutoring systems (ITS) that use step-by-step scaffolding and/or dynamic conversational strategies (especially as operationalized in the CLASS framework), and report effects on student cognitive gains, engagement, or motivation across different subject domains.
Generation Model	openai/gpt-4.1-2025-04-14
Slot Name	Slot Description
its_name	The name of the LLM-powered intelligent tutoring system (ITS) being studied.
its_description	A brief description of the ITS, including the role of LLMs and any unique features.
uses_class_framework	Does the ITS explicitly use the CLASS framework (Conversational Learning and Scaffolding Strategies)? (true, false, or null if not specified)
scaffolding_strategies	Describe the step-by-step scaffolding strategies used in the ITS (e.g., breaking down problems, hinting, graduated guidance).
conversational_strategies	Describe the dynamic conversational strategies used (e.g., adaptive dialogue, real-time feedback, question-asking, personalization).
subject_domain	The subject domain(s) in which the ITS was applied (e.g., mathematics, science, language learning, etc.).
student_population	A brief description of the student population (e.g., age, grade, prior knowledge, number of participants).
outcomes_measured	Which outcomes were measured? (e.g., cognitive gains, engagement, motivation; specify how each was measured if possible).
results_summary	A concise summary of the results regarding the impact of scaffolding and conversational strategies on cognitive gains, engagement, and motivation (include quantitative results if available).
comparison_conditions	Describe any comparison or control conditions (e.g., ITS with vs. without scaffolding, with vs. without dynamic conversation, with vs. without CLASS framework).
study_design	Briefly describe the study design (e.g., randomized controlled trial, pre-post, observational, qualitative, etc.).
limitations_or_counter_evidence	Any reported limitations, null results, or counter-evidence regarding the effectiveness of scaffolding or conversational strategies in the ITS.

Table 9: An example model-generated extraction schema in response to a theory query about intelligent tutoring systems (ITS). An example of applying this extraction schema to extract information from a single paper is provided in Table 10.

Example Evidence Extraction from Paper (Intelligent Tutoring Systems)

Extraction Schema Definition	
Extraction Schema ID	extraction-schema-36 (See Table 9)
Single Extraction Example (From One Paper)	
Extraction Result ID	extraction-result-2696
uuid	e2696.0
source_info	Lyu et al. (2024). Evaluating the Effectiveness of LLMs in Introductory Computer Science Education: A Semester-Long Field Study; (Publication Date: 2024-04)
name_short	CodeTutor
name_full	CodeTutor (LLM-powered virtual teaching assistant)
brief_description	A browser-based LLM-powered conversational assistant (implemented with OpenAI GPT-3.5 API) used as an out-of-class virtual TA for an introductory Python programming course; supports multi-turn dialogue, message- and conversation-level feedback, explanations, corrections, and debugging help.
citation_title	Evaluating the Effectiveness of LLMs in Introductory Computer Science Education: A Semester-Long Field Study
mention_or_use	use
its_name	CodeTutor
its_description	CodeTutor is a web application integrating GPT-3.5 via OpenAI API, presented as a virtual teaching assistant for an entry-level Python course. It provides multi-turn conversational help (explanations, corrections, code snippets, debugging support), stores conversation history, collects message- and conversation-level feedback, and uses a system-role prompt to behave as a course TA.
uses_class_framework	false
scaffolding_strategies	No explicit formal step-by-step scaffolding protocol (e.g., CLASS) is reported. Operational scaffolding observed includes multi-turn decomposition via user-driven follow-up (users request further information, modifications, or debugging), corrective feedback (CodeTutor issues Corrections), explanatory responses (Explanations), and iterative refinement driven by user-provided additional context. The system relies on conversational iteration rather than a documented graduated-hint or staged-problem decomposition schema.
conversational_strategies	Dynamic multi-turn conversational interaction: real-time replies to user prompts, corrections and explanations, support for modification requests and debug requests, message-level upvote/downvote feedback, conversation-level ratings on comprehension/critical thinking/syntax/independence/TA-replacement; iterative clarification when users supply further information. Conversation behavior is user-adaptive in practice (responds to follow-up queries), but no explicit adaptive learner model or CLASS-framework dialogue policy was implemented.
subject_domain	Introductory computer science — Python programming (entry-level programming)
student_population	50 undergraduate students enrolled in an entry-level programming course at a U.S. 4-year university (semester-long study); participants were split into experimental (CodeTutor) and control groups, demographics reported (majors including Computer Science, Data Science, Mathematics, Biology, etc.).
outcomes_measured	Cognitive gains: pre-test and final exam/final grade comparisons (overall score, CodeTutor-Allowed components vs CodeTutor-Not-Allowed components); Engagement: system usage logs (82 conversation sessions, 2,567 messages, 1,190 prompts analyzed, counts of message categories); Attitudes/motivation proxies: survey items on Comprehension, Critical Thinking, Syntax Mastery, Independent Learning, and TA Replacement collected over time (repeated measures); Prompt quality and response effectiveness (chi-square test correlating prompt quality with 'working' responses).
results_summary	Use of CodeTutor (an LLM conversational assistant) was associated with statistically significant cognitive gains in an introductory programming course: experimental group average improvement (final minus pre-test) = +12.50 versus control = -3.17; paired t-test for experimental group improvement $p = 0.009$. Improvement persisted on CodeTutor-Not-Allowed components (final exam / pre-test) with experimental group mean increase = +7.33 (paired t-test $p = 0.026$), suggesting transfer beyond tool use. Students who had never used LLM tools previously showed larger gains (regression estimate = +18.877, $p = 0.032$). Engagement data: 82 sessions, 2,567 messages (1,288 user messages, 1,279 CodeTutor messages); 37% of prompts judged 'good quality' and prompt quality correlated strongly with effective responses (Chi-square $\chi^2 = 144.84$, $p < 0.001$). Attitudes: majority endorsed CodeTutor's comprehension ability (67% agree/strongly), and syntax help (58% combined agree/strongly), but few endorsed its help for critical thinking (19% agree/strongly); over time comprehension agreement declined (LME time $\beta = -0.114$, $p < 0.01$) and preference for TA replacement declined weakly ($\beta = -0.099$, $p < 0.1$). The paper reports no evidence that the system implemented formal stepwise scaffolding frameworks like CLASS, and notes concerns about supporting higher-order critical thinking despite measurable gains in course scores.
comparison_conditions	Between-subjects comparison: experimental group with CodeTutor (allowed for labs/assignments) versus control group using traditional learning resources and human TAs (no CodeTutor). Exams/midterms/final were CodeTutor-Not-Allowed, enabling comparison of transfer. No within-ITS manipulations (e.g., ITS with vs. without scaffolding or CLASS) were reported.
study_design	Semester-long, between-subjects field study (pre-test/post-test), combining quantitative analyses (paired t-tests, multiple regression, ANOVA, chi-square, linear mixed-effects models) with qualitative conversational coding (inductive thematic analysis) and system usage logging.
limitations_or_counter_evidence	Limitations noted: modest sample size ($n=50$) and single-institution entry-level CS context (limits generalizability); use of GPT-3.5 (model limitations); groups balanced by pre-test and prior LLM experience but not described as fully randomized; no explicit operationalization or implementation of CLASS or other formal scaffolding frameworks, so effects of formal stepwise scaffolding were not tested; attitudinal declines over time (comprehension and TA-replacement) indicate waning trust/engagement with prolonged use; students reported limited impact on critical thinking despite score gains, and heavy over-reliance behaviors (many prompts were direct copy-pastes of assignments), which may reflect superficial engagement rather than deep learning; motivation measures were indirect (attitudes and usage) rather than direct validated motivation scales.

Table 10: An example of using the extraction schema in Table 9 on a single paper (Lyu et al., 2024), using GPT-5-MINI as an extraction model. Each theory is generated by supplying extraction results (such as those shown here) from up to 100 papers together in a single theory generation prompt.

Example Literature-Derived Theory Queries

Build a theory of how multi-agent interactions between language models can be leveraged for automated factuality assessment in natural language generation.
Build a theory of how structured cross-examination dialogues between an examiner and examinee language model can systematically expose factual errors in generated claims, including the mechanisms by which inconsistencies are surfaced.
Build a theory of how large language models internally represent, store, and retrieve factual knowledge in response to varied natural language prompts.
Build a theory of the internal mechanisms by which large language models encode factual knowledge and generate consistent, correct answers across paraphrased prompts and diverse surface forms, including the influence of entity and relation aliases.
Build a theory of how exhaustive temporal relation annotation schemes can be constructed to maximize reliability and informativeness in natural language texts.
Build a theory of how to systematically include long-distance and non-verb-centered events in temporal relation annotation schemes for news articles, and how annotation guidelines and automation can be optimized to maximize inter-annotator agreement and minimize annotation ambiguity.
Build a theory of how subjective logic-based uncertainty quantification and adaptive weighting of training pairs reduce matching ambiguity and improve reliability in text-based person retrieval.
Build a theory of how uncertainty modeling can systematically enhance cross-modal retrieval performance.
Build a theory of how and why language models generate persistent hallucinations that are robust to prompt rephrasing and model sampling.
Build a theory of the cognitive and representational mechanisms in large language models that lead to question-level hallucinations, and how these mechanisms interact with model architecture, training data, and prompt structure.
Build a theory of how tunable biases applied to language model logits facilitate efficient and controllable text generation under energy-based frameworks.
Build a theory of how large language models can autonomously generate and select in-context demonstrations to improve their own zero-shot reasoning performance without human-provided labels.
Build a theory of how gradient-optimized, step-wise biases over autoregressive language model logits interact with energy-based constraint satisfaction and fluency preservation during controlled text generation.
Build a theory of how large language models can use self-generated outputs, filtered and selected by criteria such as self-consistency, semantic diversity, and repetitiveness, to construct in-context demonstrations that maximize reasoning accuracy in zero-shot settings.
Build a theory of how meta-reinforcement learning with actor, critic, and meta-critic modules enables the discovery of brain effective connectivity networks from noisy, small-sample fMRI time series data.
Build a theory of how meta-reinforcement learning enables robust causal discovery in high-noise, small-sample environments.
Build a theory of how machine unlearning can be leveraged to systematically and efficiently remove biases from trained deep neural networks.
Build a theory of how influence-function-based machine unlearning with small external counterfactual datasets can target and remove specific attribute-driven biases in deep neural networks without requiring access to the original training data.
Build a theory of how mixed prompt architectures facilitate efficient few-shot adaptation in large language models.
Build a theory of how using an infinite mixture of asymmetric Laplace distributions as a VAE decoder with CRPS loss modulates the trade-off between synthetic data fidelity and privacy preservation as a function of the KL-divergence weight parameter (beta) in tabular data.
Build a theory of how the interplay between textual and key-value prompts, inserted at various layers and positions in a frozen Transformer backbone, enables rapid and robust adaptation to novel product attributes with minimal labeled data in attribute value extraction tasks.
Build a theory of how nonparametric, distributional decoders in variational autoencoders influence the fidelity, diversity, and privacy of synthetic data generation across different data modalities.
Build a theory of why and how certain tokens in large language model sequences persist as influential across multiple attention steps.
Build a theory of the mechanisms at the intersection of model architecture, training dynamics, and data properties that cause specific tokens to repeatedly receive high attention and remain pivotal for future generations in transformer-based LLMs.
Build a theory of how self-supervised learning (SSL) approaches can be systematically adapted to maximize representation learning from remote sensing imagery with heterogeneous sensor characteristics and minimal labeled data.
Build a theory of how contrastive self-supervised learning (SSL) methods can be optimized for multispectral Landsat imagery, accounting for differences in spectral bands, spatial resolutions, and product processing levels, to yield robust and transferable representations.
Build a theory of how multi-view prompting affects structured prediction performance in generative language models.
Build a theory of how aggregating outputs from multiple element order prompts (multi-view prompting) impacts the accuracy, stability, and error resilience of aspect sentiment tuple prediction in generative models.
Build a theory of how ambiguity in natural language queries impacts the generation and evaluation of multiple valid outputs in semantic parsing tasks such as text-to-SQL.
Build a theory of how specific ambiguity types (column name, table name, join, precomputed aggregate) affect the probability that a text-to-SQL system will generate all semantically valid SQL interpretations within its top-k outputs, and how these effects interact with schema complexity.
Build a theory of optimal multimodal fusion for emotion recognition in conversations, accounting for the complementary and asynchronous nature of textual, audio, and visual cues.
Build a theory of how bidirectional multi-head cross-attention mechanisms can model and exploit complex, context-dependent correlations and mapping relationships between textual, audio, and visual modalities to improve emotion recognition in conversational settings.
Build a theory of how automatic prompt augmentation and selection methods can systematically improve the reasoning capabilities of large language models across a variety of task types.
Build a theory of how the interplay between automatic generation, pruning, and policy-gradient-based selection of chain-of-thought exemplars impacts LLM performance on arithmetic, commonsense, symbolic, and non-reasoning tasks, including the mechanisms by which these steps mitigate order, complexity, diversity, and style sensitivities.
Build a theory of how large language model-based intelligent tutoring systems (LLM-ITS) can be systematically designed and refined to maximize student learning outcomes through the application of learning science principles.
Build a theory of how the explicit integration of step-by-step scaffolding and dynamic conversational strategies, as operationalized in the CLASS framework, influences student cognitive gains, engagement, and motivation in LLM-powered intelligent tutoring systems across diverse subject domains.
Build a theory of how multi-modal neurons in large language models represent, integrate, and mediate cross-modal concepts.
Build a theory of how feed-forward network neurons in transformer-based multi-modal LLMs encode, maintain, and causally influence the mapping between specific visual features and corresponding textual concepts during image captioning and related tasks.
Build a theory of how reinforcement learning-driven adaptive subgraph selection from large-scale knowledge graphs can be optimized to select the most relevant and least noisy knowledge for each individual news item, thereby improving detection of sophisticated multimodal fake news.
Build a theory of how adaptive, context-sensitive knowledge selection mechanisms impact the effectiveness of multimodal fake news detection systems.
Build a theory of how linguistic structures and social factors interact to determine code-switching patterns in multilingual communities.
Build a theory of how reinforcement learning agents can guarantee long-term safety in environments with unknown, stochastic dynamics and only binary safety feedback.
Build a theory of how speaker identity, social context, and grammatical constraints jointly predict the occurrence, type, and placement of code-switching in bilingual conversations, with a focus on European and Indian multilingual contexts.

Table 11: Example *theory queries* automatically generated from scientific papers.

Example 1: Laws judged as different	
Theory A (Name)	Information Bottleneck Regulation Theory for Temporal Annotation
Theory A (Description)	Maximizing informativeness and reliability in temporal relation annotation requires regulating annotation scheme complexity such that only salient, information-contributing temporal relations are annotated, avoiding both over-annotation (which reduces reliability by diluting signal) and under-annotation (which reduces informativeness). A well-constructed temporal annotation scheme thus features an explicit information bottleneck protocol—selecting relations for annotation based on their contribution to narrative temporal structure.
Theory A (Law/Statement)	Annotation reliability and informativeness are maximized when only temporal relations that contribute to the overall event timeline or narrative progression are annotated, with non-contributory or redundant relations filtered out at annotation time.
Theory B (Name)	Dual Optimization Theory of Temporal Relation Annotation Schemes
Theory B (Description)	This theory posits that the construction of exhaustive temporal relation annotation schemes must simultaneously optimize two orthogonal axes: reliability (inter-annotator agreement, cognitive tractability) and informativeness (expressiveness, ability to support downstream temporal inference). Through formal modeling of annotation noise, task complexity, and linguistic event ambiguity, an optimal scheme can be constructed that locally maximizes both criteria for a given text domain and annotator population.
Theory B (Law/Statement)	The reliability and informativeness of a temporal relation annotation scheme are maximized when its relation taxonomy and guidelines are empirically adapted to the distribution of event types, temporal connectives, tenses, and discourse structure specific to the targeted text domain.
Core Claim A	Maximizing annotation reliability and informativeness is achieved by annotating only temporal relations that meaningfully contribute to the event timeline or narrative progression and filtering out non-contributory or redundant relations (an information-bottleneck style selection).
Core Claim B	Maximizing annotation reliability and informativeness is achieved by empirically adapting the relation taxonomy and guidelines to the domain-specific distribution of event types, temporal connectives, tenses, and discourse structures.
Similarities	Both assert that the design of the temporal relation annotation scheme determines reliability and informativeness [ESSENTIAL]. Both frame reliability and informativeness as jointly optimizable objectives that can be degraded by a poorly chosen scheme [ESSENTIAL]. Both imply that annotation complexity should be controlled for cognitive tractability and downstream usefulness [NON-ESSENTIAL].
Differences	A prescribes a selection and filtering rule that removes redundant or non-contributory relations via an information-bottleneck principle [ESSENTIAL]. B prescribes empirical, domain-specific adaptation of taxonomy and guidelines [ESSENTIAL]. A emphasizes reducing complexity by filtering redundancy, whereas B emphasizes adapting expressiveness to support inference [ESSENTIAL]. B explicitly invokes formal modeling of noise and local optimization, which A does not [NON-ESSENTIAL].
Reasoning	Both theories claim that annotation-scheme design governs reliability and informativeness, but they prescribe different essential mechanisms: A requires filtering to a minimal set of narrative-contributory relations, while B requires empirical, domain-specific adaptation. Because these mechanisms are not equivalent, the theories are not duplicates.
Judgment	Not duplicates
Example 2: Laws judged as duplicates	
Theory A (Name)	Hierarchical Distributed Semantic Attractor Theory (HDSAT)
Theory A (Description)	Large language models store and retrieve factual knowledge via a hierarchical system of distributed semantic attractors. Facts are represented as high-dimensional attractor basins in activation space, structured hierarchically from general to specific. Retrieval involves dynamic settling into the nearest attractor consistent with prompt context, enabling generalization but also blending or misfiring when attractors overlap.
Theory A (Law/Statement)	The degree of factual recall robustness to prompt rephrasing or noise depends on the density and overlap of attractor basins in activation space—facts with dense, well-separated attractors are robust, while shallow or overlapping attractors are vulnerable to paraphrase failure or hallucination.
Theory B (Name)	Distributed Redundant Memory Encoding Theory
Theory B (Description)	Factual knowledge in LLMs is stored redundantly and distributively across many weights and layers, forming high-dimensional attractor-like states such that multiple pathways can lead to successful recall or recombination, providing robustness but also permitting interference-based failure modes.
Theory B (Law/Statement)	Factual knowledge is not localized but stored as overlapping patterns of parameters and activations across many layers and neurons, allowing information to be resilient to noisy prompts or sparse ablation.
Core Claim A	Factual knowledge is encoded as hierarchical, high-dimensional semantic attractor basins, and recall robustness depends on attractor density and overlap.
Core Claim B	Factual knowledge is encoded as distributed, redundant, overlapping high-dimensional patterns whose geometry determines robustness and failure modes.
Similarities	Both propose that factual knowledge is represented as high-dimensional, distributed attractor-like patterns rather than localized parameters [ESSENTIAL]. Both link robustness of recall to the geometry and overlap of those representations [ESSENTIAL]. Both acknowledge that the same representation yields resilience and interference-based failures [NON-ESSENTIAL].
Differences	A posits a hierarchical organization of attractors as central [ESSENTIAL]. B emphasizes redundancy across weights and layers and resilience to sparse ablation [ESSENTIAL]. A focuses on paraphrase robustness, whereas B emphasizes ablation and implementation details [NON-ESSENTIAL].
Reasoning	The two theories share the same central claim that factual knowledge is stored in distributed, high-dimensional attractor-like representations and that overlap determines robustness versus failure. The remaining differences concern hierarchy and implementation detail and do not alter the core equivalence.
Judgment	Duplicates

Table 12: Two examples of the duplicate detection process. Theory A and Theory B serve as input to a language model prompt (GPT-5-MINI), which fills the output schema that includes core claims of each law, similarities, differences, and a final judgment as to whether the laws are substantially duplicates of each other.