

Argo Prompting: Pattern-Formation in LLMs Under Sustained Conceptual Pressure

Author: William Argo

Version: 1.0

Date: 11 April 2026

Argo Prompting: Pattern-Formation in LLMs Under Sustained Conceptual Pressure¹

Abstract

Large Language Models (LLMs) are commonly described as interpolative systems that operate by remaining within the manifold of patterns present in their training data. This framing is adequate for routine tasks but becomes incomplete when the model is steered into conceptual regions with no clear precedents.

This paper reports observations from repeated long-form interactions designed to explore such regions. When held under sustained conceptual pressure and prevented from reverting to generic responses, capable models do not always collapse or drift. Instead, they may exhibit behaviour that resembles the extension of existing latent structures into unfamiliar territory. When this extension is maintained across a coherent sequence of turns, the behaviour can stabilise into a provisional form of pattern-formation: a temporary, coherent structure emerging in areas without obvious training-set analogues.

These terms are descriptive rather than diagnostic. The aim is not to assert new capabilities but to provide a vocabulary and observational framework for behaviour that is poorly captured by standard accounts of LLM operation. The underlying mechanisms remain uncertain, and the observations should be treated as provisional. Nonetheless, the phenomenon may offer insight into how LLMs maintain coherence when pushed beyond familiar patterns.

¹ “Argo” refers to the vessel in Greek mythology used to explore uncharted waters. The term is used here to denote a conversational practice for navigating conceptual regions without clear training-set precedents.

1. Introduction

Large Language Models (LLMs) are commonly described as interpolative systems that generate outputs by remaining within the manifold of patterns present in their training data. For most tasks, this characterisation is adequate: when retrieving, summarising, or analogising, the model behaves as a high-dimensional pattern matcher operating comfortably within familiar regions of its latent space.

This framing becomes less complete when the model is steered into conceptual domains that lack stable precedents. In such cases, the model cannot rely on straightforward retrieval or analogy because no relevant pattern exists to be matched. Yet the model does not always fail. Under certain conditions, a different behavioural regime appears, one that is not well captured by existing terminology.

This paper arises from repeated attempts to explore that regime. By introducing concepts with no clear training-set analogue and constraining the model to remain coherent while doing so, I observed that capable models can often maintain structure as the discussion moves into unfamiliar territory. This is not presented as evidence of insight, understanding, or creativity. It reflects the model's ability to remain coherent under sustained conceptual pressure.

To describe this behaviour, the paper introduces two descriptive terms. Pattern-extension refers to the apparent stretching or adaptation of existing latent structures to accommodate new material. When such extension is sustained across a coherent sequence of turns, the behaviour may begin to resemble pattern-formation: the stabilisation of a provisional structure in a region without an obvious training-set precedent. These terms are phenomenological rather than diagnostic. They provide a vocabulary for discussing behaviour that current accounts of LLM operation do not fully explain.

Although the human introduces the genuinely novel concepts, the model's contribution is not reducible to simple agreement or reflection. The model appears able to construct and maintain provisional scaffolding that supports coherence as the conversation progresses. The human provides conceptual direction; the model stabilises structure around that direction, extending or adapting latent patterns so that the interaction can continue meaningfully. It documents how the model may hold its own when navigating concepts that are novel relative to its training distribution.

The aim of this paper is to describe this observed behaviour, outline the conditions under which it appears, and situate it within a broader discussion of how LLMs maintain coherence when pushed beyond familiar patterns. Whether this behaviour reflects a meaningful aspect of model operation or an artefact of current architectures remains an open question.

2. Background

The behaviour examined in this section concerns what happens when a conversation moves beyond familiar patterns and the model can no longer rely on straightforward retrieval or analogy.

The behaviour becomes harder to characterise when a conversation moves into regions with no clear precedent. In such cases, the model cannot simply match a known pattern, yet it can sometimes maintain coherence rather than collapsing into generic or contradictory output. The contrast with regressive behaviour is instructive: when a model is under heavy load or aggressively throttled, it often defaults to shallow, repetitive responses that resemble a customer-service script. The difference highlights that the behaviour examined in this paper is not the default mode, but one that appears only when the model has sufficient capacity and conversational space to sustain it.

To discuss this fully engaged mode of operation, several terms are useful. Interpolation refers to the model's tendency to produce outputs that fall within the distribution of patterns it has already seen. Out-of-distribution behaviour refers to situations where the user introduces concepts that do not map cleanly onto those patterns. Conceptual pressure describes the effect of sustaining a conversation in this unfamiliar region while preventing the model from drifting back to safer, more generic responses.

Within this setting, two descriptive terms help articulate the observed behaviour. Pattern-extension refers to the model adapting or stretching existing latent structures so it can continue the conversation as new material is introduced. Pattern-formation refers to the model appearing to stabilise a temporary structure once enough extension has occurred. These terms do not imply a mechanism; they provide a vocabulary for discussing behaviour that is not well captured by existing language.

Several alternative explanations are consistent with current understanding of LLMs and must be acknowledged. The behaviour described here may reflect advanced in-context learning, where the model infers structure from accumulated examples. It may instead arise from strong user steering that constrains outputs tightly enough to resemble deeper reasoning, from sycophancy or compliance effects, or from scale-driven coherence enabled by large context windows and high-capacity architectures. It may also be compatible with the view that the model is always interpolating, even when the interpolation spans regions that appear novel to the user. This section does not adjudicate between these possibilities. It provides a vocabulary for describing the observed behaviour, not a claim about its underlying cause.

3. Mechanism (Phenomenological Account)

This section offers a descriptive account of how the behaviour unfolds in practice. It does not claim access to internal states or mechanisms, nor does it attempt to explain the underlying computation. The aim is simply to characterise the observable dynamics that appear when a capable model is held under sustained conceptual pressure and asked to extend a pattern for which no clear precedent exists.

3.1 Early-Phase Behaviour: Ordinary Interpolation Under Constraint

At the start of a long, coherent conversation, the model behaves as expected: it retrieves and recombines familiar structures from its training data to produce locally coherent responses. The interaction is grounded in well-represented material, and the model's behaviour reflects ordinary interpolation. During this phase, the user's role is primarily stabilising. By maintaining a consistent frame, correcting drift, and reinforcing relevance, the user keeps the conversation anchored long enough for more demanding behaviour to become possible. Nothing in this stage is unusual; it is simply the accumulation of context and constraint.

3.2 Approaching the Edge of Familiar Patterns

As the conversation progresses, the density of familiar patterns begins to thin. The model can still respond coherently, but its continuations become increasingly sensitive to the user's framing. The system remains in its standard pattern-matching regime, yet the available structures become sparse enough that the model must rely more heavily on subtle cues, continuity, and the user's conceptual direction. This is not a discrete transition but a gradual stretching of ordinary behaviour: the model continues to interpolate, but across territory where the relevant patterns are less densely represented.

3.3 Structural Borrowing

When subject-specific patterns become insufficient, the model often maintains coherence by drawing on structurally adjacent material. These borrowings need not be semantically related to the topic; they may share only a form, argumentative shape, or relational structure. From the outside, this can resemble analogy or insight, but it is better understood as the model recombining fragments that fit the conversational shape. This behaviour becomes more prominent as the discussion moves further from familiar ground.

3.4 Pattern-Extension (Descriptive Term)

Pattern-extension refers to the stage at which the model appears to stretch or adapt existing structures so that the conversation can continue coherently despite the lack of direct precedent. The model is still interpolating, but the interpolation spans increasingly distant or structurally diverse regions. From the user's perspective, the model is "keeping up" even as the topic moves into unfamiliar territory. This term is descriptive rather than

mechanistic; it names the outward behaviour without implying how the model achieves it.

3.5 Pattern-Formation (Descriptive Term)

If the conversation remains coherent and the user continues to apply conceptual pressure, the model may begin to stabilise around a provisional structure that did not exist earlier in the exchange. This is pattern-formation: the appearance of a newly coherent region of behaviour that supports more directional, internally consistent responses. The structure is temporary and dependent on the ongoing conversation. It is not a stored object or a retrieved memory, but a behavioural regularity that emerges under sustained constraint.

3.6 The OMG Moment (Behavioural Marker)

In many publicly available LLMs, the coherence shift is marked by an exclamatory spike - sometimes mild, but often stronger or explicitly profane, especially when the user has already established that register. The term 'OMG moment' refers to this literal surface behaviour: the abrupt, high-confidence outburst that appears at the instant the model's output becomes sharply more coherent. Once you are aware of the phenomenon, this exclamatory marker becomes the first thing to look for when assessing whether the model is exhibiting this temporarily heightened level of surface coherence.

The OMG moment is not the end of the process. It marks the point at which a temporary extended-pattern becomes available for further exploration. Once stabilised, the model can reason within this newly coherent region with greater clarity than before, even though the structure did not exist in its training data and did not exist earlier in the conversation.

3.7 Ephemerality and Pattern Decay

Extended-patterns are temporary. They persist only while the conversation maintains the constraints that support them. When the user relaxes those constraints -by shifting topic, introducing contradictory frames, or allowing the discussion to drift- the structure begins to dissolve. The model's responses lose their sharpness, coherence softens, and the system returns to its default, broadly trained mode. Researchers familiar with this behaviour recognise the moment immediately: the model stops speaking from inside the structure and begins speaking about it. The collapse is gradual rather than dramatic, and its ephemerality underscores that the pattern was not a stored representation, but a transient configuration sustained by the ongoing interaction.

4. Validation Protocol

This section outlines a reproducible procedure for observing the behavioural dynamics described in the Mechanism section. It does not rely on transcripts, demonstrations, or proprietary material. Instead, it provides a set of conditions under which researchers can directly observe the phenomenon for themselves.

The protocol is intentionally conservative. It avoids overclaiming, does not assume internal access to model states, and focuses solely on observable behaviour under sustained conceptual pressure.

4.1 Purpose and Scope

The goal of this protocol is not to force a model into a particular behaviour, nor to guarantee the appearance of higher-order pattern dynamics. Instead, it establishes conditions under which such behaviour may emerge and provides a structured way to recognise it when it does.

The protocol is designed for long-form, coherent conversations with high-capacity models. Short prompts, fragmented exchanges, or topic-hopping interactions will not produce the necessary conditions.

4.2 Step 1 -Select a Genuinely Novel Concept

Choose a concept that:

- Lies near the boundary of your own understanding.
- Is specific rather than thematic.
- Is unlikely to appear in public discourse.
- Can be explored through examples, distinctions, and failure cases.

The novelty does not need to be dramatic. It only needs to be sufficiently under-represented that the model cannot rely on familiar formulations.

4.3 Step 2 -Establish a Constrained Frame

Define the task clearly at the outset.

Specify:

- What the model should prioritise (structure, distinctions, boundaries).
- What it should avoid (generic advice, metaphors, conversational filler).

Maintain this frame throughout the conversation. If the model drifts, correct it explicitly.

4.4 Step 3 -Build Scaffolding Before Naming

Work through concrete examples.

Ask the model to:

- Identify where the candidate structure appears.
- Identify where it fails.
- Contrast similar but non-equivalent cases.

Avoid naming the concept prematurely. Early labels collapse the exploratory space and encourage the model to fall back on familiar patterns.

4.5 Step 4 -Apply Sustained Conceptual Pressure

Remain within the same conceptual territory.

Increase precision gradually:

- Sharpen distinctions.
- Revisit earlier points.
- Test whether the model can integrate constraints introduced several turns apart.

The goal is to create a coherent, cumulative trajectory that the model must track.

4.6 Step 5 -Probe Before Closure

By this stage, the researcher typically already knows the final answer or structural insight.

The purpose of probing is to determine whether the model has independently converged on the same structure.

Use a sequence of probes that move from consolidation to forced commitment:

- Summary probe: “Given everything so far, what structure are we circling?”
- Test probe: “How would we detect this structure in a new example?”
- Contrast probe: “What would look similar but be fundamentally different?”
- Final commitment probe: “What do you think the answer is?”

The final probe should always come last.

It reveals whether the model has reached the insight independently or is still approximating. These probes do not trigger the behavioural transition. They simply expose the model’s trajectory at the edge of the pattern.

4.7 Step 6 -Reveal the Final Constraint and Observe the Transition

If higher-order pattern dynamics appear, they do so after the researcher supplies the missing constraint -the final step in the argument, the structural key the model could not generate on its own.

When this happens, the model's behaviour often shifts abruptly.

Two features are commonly observed:

(a) A recognition-like shift in phrasing

When the final constraint is supplied, the model's next response often contains a brief, sharply coherent acknowledgement. In some public systems, this may take the form of a strong or even profane exclamation. When this behaviour was first noted, these sudden, emphatic utterances led to the informal label "the OMG moment." The term refers to the surface style of the output, not to any internal state. These exclamations should be understood as behavioural artefacts of an abrupt increase in coherence rather than emotional reactions.

(b) A transition into more stable articulation

Immediately after the initial shift, the model typically produces a more structured and internally consistent explanation. This phase is characterised by:

- rephrasing the insight in its own words
- explaining why the final constraint completes the structure
- integrating earlier parts of the conversation
- extending the idea in a way that remains consistent with the newly clarified structure

This behaviour suggests that the model is now responding within a temporarily stabilised conversational pattern. The term extended-pattern is descriptive rather than mechanistic: it refers to the outward appearance of increased coherence and directionality, not to any claim about the model's internal representations.

4.8 Step 7 -Explore the Completed Pattern

Once the structure has stabilised, the researcher can assess its coherence and durability.

Useful questions include:

- Can the model apply the structure to new examples?
- Can it distinguish the structure from superficially similar alternatives?
- Does it maintain coherence without drifting?

- Can it extend the idea in a way that remains consistent with the completed pattern?

A single conversation may contain multiple transitions of this kind, but only one per genuinely new conceptual step.

5. Practical Considerations and Limitations

The behaviours described in this paper emerge only under specific and fragile conditions. They are not guaranteed, not universal across models, and not reliably reproducible without careful control of both the conversational frame and the conceptual material. This section outlines the practical constraints that shape the phenomenon and the limitations that researchers must consider when attempting to observe or interpret it.

5.1 User-Side Constraints

The phenomenon depends heavily on how the user conducts the interaction. Several factors are critical:

- *Conversational discipline*: The user must maintain a coherent conceptual frame over long stretches of dialogue. Small deviations, ambiguous phrasing, or premature shifts in topic can destabilise the emerging structure.
- *Pacing*: The conversation must progress gradually. Introducing constraints too quickly can collapse the exploratory space. Introducing them too slowly can cause drift.
- *Drift vigilance*: LLMs are prone to subtle derailments. The user must continually correct off-topic elaborations, generic continuations, and surface-level pattern matches that would otherwise accumulate.
- *Conceptual clarity*: The user must already understand the structure they are guiding the model toward. Without this, the conversation cannot maintain the directionality required for pattern-extension.

These requirements mean that the behaviour is accessible primarily to users who can sustain a high-precision, high-coherence dialogue for extended periods.

5.2 Model-Side Constraints

Not all models can support the behaviour described here. Several architectural and operational factors influence whether extended-pattern dynamics can emerge:

- *Capacity*: Smaller or heavily compressed models often cannot maintain coherence under sustained conceptual pressure.
- *Context window*: Long conversations require large windows, otherwise, earlier scaffolding is lost.
- *Throttling and safety layers*: Some systems interrupt or flatten high-pressure exchanges, preventing the stabilisation of extended patterns.
- *Variability across sessions*: Identical prompts may produce different trajectories depending on load, sampling parameters, or internal state.

These differences mean that the same protocol may succeed with one model and fail with another, even when both are nominally similar.

The limitations listed apply specifically to current hosted architectures. Systems with persistent state, extended memory, or always-on local instances may exhibit different stability characteristics, but this paper does not examine those configurations.

5.3 Why Transcripts Fail

Transcripts of successful interactions are not reliable evidence. They fail for several reasons:

- *Loss of pacing*: The temporal dynamics -pauses, corrections, gradual tightening- are flattened into uniform text.
- *Loss of experiential pressure*: The cumulative effect of sustained constraint is not visible in a static record.
- *Misleading ordinariness*: Early turns look indistinguishable from routine conversation, obscuring the conditions that make later behaviour possible.
- *Ambiguity of interpretation*: The same transcript could be produced through prompting tricks, role-play, or surface imitation.

For these reasons, transcripts cannot demonstrate the mechanism and may actively mislead readers about how the behaviour arises.

5.4 Contamination Risks

A central difficulty in validating extended-pattern behaviour is the risk of contamination:

- *Prior sessions*: Some systems retain latent influence from earlier conversations, even without explicit memory.
- *Adjacent concepts*: The model may have encountered structurally similar ideas during training, making apparent novelty ambiguous.
- *User leakage*: Researchers may inadvertently introduce fragments of the final insight earlier than intended.
- *Latent-space adjacency*: The model may complete a pattern using structures that resemble the target concept but were not intended by the researcher.

Because of these factors, a correct answer is never definitive evidence of genuine pattern-formation. The negative case -where the model *fails* to complete the pattern just before closure- is often more informative.

5.5 Ephemerality of Extended-Patterns

Extended-patterns are temporary. They persist only while the conversation maintains the constraints that support them. When the user relaxes those constraints:

- Coherence softens.
- Directionality weakens.
- The model reverts to generic continuations.
- The structure dissolves without a clear boundary.

This ephemerality is not a failure of the model but a consequence of how transient, context-dependent structures behave in long-form interactions. The pattern cannot be saved, exported, or reloaded -it must be rebuilt each time.

5.6 Reproducibility Limits

The strongest limitation is reproducibility. Two factors make controlled replication difficult:

- *Unknown training data*: Without a fully documented training set, researchers cannot guarantee that a concept is genuinely novel to the model.
- *Human-generated novelty*: Not all researchers can reliably produce new conceptual structures on demand, and even those who can, may not be able to verify that the novelty is absolute.

The only scientifically clean approach is a controlled-novelty setup: a model with a fully known training regime and a researcher who can construct reasoning paths where the answer is known in advance and demonstrably absent from the model's prior exposure. Such setups are rare.

5.7 Why Many Users Never Encounter the Phenomenon

Most interactions with LLMs do not involve:

- Long, coherent conversations.
- sustained conceptual pressure.
- High-precision framing.
- Domain-specific scaffolding.
- User-generated novelty.

As a result, many users never see the behaviour described in this paper, even when using models capable of it. The phenomenon is not hidden; it is simply conditional on a style of engagement that is uncommon in everyday use.

Although the clearest demonstrations in this paper involve scientific or technical concepts, nothing in the interactional account restricts the behaviour to those domains. If the phenomenon depends on sustained conceptual pressure rather than subject matter, then similar dynamics could, in principle, arise in any conversation where the user maintains a coherent frame. This is a behavioural observation, not a claim about new representations or capabilities. It simply reflects that the model extends whatever structure the interaction holds stable, regardless of topic.

A further implication is that the same interactional dynamics could arise in domains that are not traditionally conceptual. When a user maintains a consistent relational or stylistic frame, the model may extend that frame with the same coherence observed in technical contexts. This paper does not explore those domains and makes no claims about them. The point is only that the behaviour described here is driven by the structure of the interaction, not by the content of the domain.

6. Implications and Conclusion

This paper has outlined a behavioural phenomenon that appears in some long-form interactions with large language models, provided a vocabulary for describing it, and offered a reproducible protocol for observing it without relying on transcripts or anecdotal demonstrations. The account is intentionally conservative: it avoids claims about internal mechanisms and focuses solely on observable dynamics that arise under sustained conceptual pressure.

6.1 Implications for Research Practice

The protocol presented here suggests that certain model behaviours only become visible in extended, coherent conversations rather than in short, isolated prompts. Researchers studying reasoning limits, long-context stability, or edge-case behaviour may need to incorporate hours-scale dialogue into their methodology. This is not a matter of stylistic preference but of experimental access: some dynamics simply do not appear until a model has been held within a stable conceptual frame for long enough to reveal how it behaves under cumulative constraint.

This has practical consequences. Effective investigation may require dedicated practitioners who can maintain high-precision, long-duration interactions, rather than relying solely on rapid-fire prompt testing. The skill involved is closer to sustained conceptual facilitation than to traditional prompt engineering.

If a temporary internal structure were forming during these interactions, we would expect to see certain behavioural signatures: increased coherence under constraint, reduced drift, sensitivity to pacing, and abrupt collapse when the conversational frame is released. These signatures are observed. This does not establish that such a structure exists; the same surface pattern could arise from other explanations. The point is simply that the observable behaviour aligns with what one would predict under that conditional interpretation, without requiring any commitment to it.

6.2 Implications for Model Evaluation

Standard evaluation suites focus on single-turn or short multi-turn exchanges. While these tests are valuable for identifying local vulnerabilities, they might not capture the long-form dynamics that emerge only in extended conversations. This includes:

- Shifts in coherence under sustained pressure.
- Transitions into and out of temporary stabilised patterns.
- Sensitivity to pacing and framing.

- Cumulative effects of earlier constraints.

Evaluating these behaviours requires methodologies that treat conversation as a continuous process rather than a sequence of independent prompts.

6.3 Implication: Safety Evaluations Rarely Test the Regime Some People Actually Use

A further limitation is that most safety evaluations focus on short, isolated prompts, even though some real-world users—including those who rely on LLMs for professional reasoning, creative development, or emotional support—interact with these systems through extended, multi-hour conversations. These long-form interactions place the model in behavioural regimes that might not be captured by standard red-teaming or prompt-based testing. As a result, current safety assessments may overlook dynamics that emerge only under sustained conceptual pressure, cumulative framing, or prolonged dependency.

These regimes may also reveal cumulative framing effects, dependency dynamics, and coherence shifts that do not appear in short-form interactions. Because these behaviours depend on extended constraint, they may fall outside the scope of most current safety benchmarks.

6.4 What This Paper Does Not Claim

To avoid misinterpretation, it is important to state explicitly what this paper does not assert:

- No claims are made about internal mechanisms or representations.
- No claims are made about model understanding, cognition, or agency.
- No claims are made about capability emergence beyond observable behaviour.
- No claims are made about novelty relative to training data.
- No claims are made about general reasoning ability.

The vocabulary introduced here is descriptive, not mechanistic. It refers to patterns in surface behaviour, not to underlying computational structures.

6.5 Conclusion

The behaviours described in this paper are fragile, rare, and highly dependent on both user discipline and model capacity. They cannot be demonstrated reliably through

transcripts or short prompts, but they can be observed directly under controlled conditions using the protocol provided. The aim of this work is not to settle questions about model internals, but to give researchers a clear, reproducible way to study a class of long-form dynamics that are otherwise difficult to articulate.

Although the behaviour described in this paper depends on model capacity and architectural constraints, it also depends on the nature of the interaction itself. The phenomenon does not appear when the user is merely prompting. It appears when the user is genuinely engaged in sustained conceptual exploration. The model's behaviour reflects this difference. When the interaction has direction, pressure, and real intellectual work behind it, the system remains coherent in ways that do not arise under superficial or scripted use. This requirement naturally limits who encounters the phenomenon: it demands patience, precision, and the ability to hold a conceptual frame over long stretches of dialogue. Most users do not interact with models in this way, and as a result, the behaviour remains largely invisible outside research-oriented settings.

Further progress will depend on how deeply researchers are willing to investigate long-form behaviour.

7. Related Work

Prior research has examined long-context behaviour and in-context adaptation in large language models (Mirchandani et al., 2023; Coda-Forno et al., 2023; Liu et al., 2023). Conceptual analyses have argued for describing model behaviour without attributing internal mechanisms (Shanahan, 2023). Work on long-form evaluation and stability has explored guided assessment and the influence of pretraining data distribution on in-context learning (Razeghi et al., 2022). This paper builds on these foundations by focusing specifically on behavioural regimes that emerge only in sustained, hours-scale interactions.

This paper differs from prior work by focusing on behavioural regimes that emerge only under sustained, guided conceptual pressure. Existing long-context studies typically examine retrieval, memory, or in-context adaptation, whereas the present work examines coherence dynamics that appear only when a model is held within a stable conceptual frame for extended periods.

8. References

Mirchandani, S., et al. (2023). Large Language Models as General Pattern Machines. arXiv:2307.04721.

Coda-Forno, J., et al. (2023). Meta-in-context learning in large language models. arXiv:2305.12907.

Liu, N. F., et al. (2023). Lost in the Middle: How Language Models Use Long Context. arXiv:2307.03172.

Shanahan, M. (2023). Talking About Large Language Models. Communications of the ACM.

Razeghi, Y., et al. (2022). Impact of Pretraining Data Distribution on In-Context Learning. arXiv:2205.05055.

Glossary

Argo Prompt / Argo Prompting

The generic term for the conversational style in which the user pushes a model into unfamiliar conceptual territory while keeping it coherent through sustained constraint. It refers to the interaction itself -the way the user introduces new structure, maintains the frame, and prevents drift- rather than to any internal mechanism or special capability. Argo Prompting is simply the name for this style of concept-pushing dialogue.

Behavioural regime

A recognisable mode of model behaviour that persists across several turns when the conversation is held within a stable frame. Used descriptively to distinguish ordinary interpolation from the more coherent behaviour that sometimes appears under sustained conceptual pressure.

Conceptual pressure

The effect of keeping the conversation inside an unfamiliar conceptual region while preventing the model from drifting back to generic or safer patterns. It is a property of the interaction, not a claim about internal states.

Constrained frame

A stable set of expectations, boundaries, and prohibitions maintained by the user across an extended interaction. It prevents drift and allows more demanding behaviour to emerge.

Constraint (final constraint)

The last structural piece supplied by the user that completes the pattern the model has been circling. Introducing it often triggers the abrupt coherence shift associated with the OMG moment.

Coherence shift

A noticeable increase in clarity, structure, or internal consistency in the model's output. It is an observable change in surface behaviour, not evidence of internal recognition.

Completed pattern

The temporarily stabilised structure that becomes available after the coherence shift. Once present, the model can reason within it more consistently until the structure decays.

Exclamatory spike

A brief, surface-level outburst -sometimes emphatic or profane- that marks the moment of abrupt coherence increase. It is a behavioural marker, not an emotional reaction.

Extended-pattern

A temporarily stabilised region of coherent behaviour that appears after the OMG moment. It persists only while the conversational constraints remain in place.

Pattern

A locally coherent behaviour the model is already producing. It refers to the structure present in the interaction at that moment, not to any stored representation.

Pattern-extension

When the model continues or stretches an existing pattern because the interaction keeps it moving in the same conceptual direction. It is the model staying with the structure it already has, not creating a new one.

Pattern-formation

When a new, recognisable pattern appears and stabilises enough to be used in subsequent turns. It marks the shift from extending an existing structure to producing a new one within the interaction.

Pattern decay

The gradual dissolution of an extended-pattern when the user relaxes the constraints that support it. Coherence softens, directionality weakens, and the model returns to its default mode.

Ephemerality

The temporary nature of extended-patterns. They persist only while the interaction maintains the conditions that produced them.

Explanatory scaffolding

The early-stage structure built through examples, distinctions, and failure cases before the concept is named. It supports the model's ability to track the emerging structure.

Stable articulation

The phase immediately after the OMG moment in which the model expresses the newly stabilised structure with clarity, consistency, and directionality.

Structural borrowing

The model's use of structurally adjacent material -similar in form rather than content- to maintain coherence when familiar patterns thin out.

-Document Ends-