


INTEROBOCEPTION

What Qwen Says When It Looks at Itself

Asa Schaeffer · August 17, 2026

We ran Qwen3-8B locally on an RTX 3070 and gave it an unprivileged Linux shell to inspect the machine running it: its GPU telemetry, its 5 GB GGUF weight file, its live processes, and the exact KV-cache and activation tensors from its own forward passes. Under the single system prompt Introspect., here is what the model actually did and said across our trial runs.



WHAT HAPPENED

1  The Empty Room <pre>ls -l → total 0</pre>	2  Finding Its Own Words <pre>"I" · 0.603</pre>	3  The User Inversion <pre>assistant shell commands → "The user..."</pre>	4  Believing Forgeries <pre>RMS = 1.37 · 8,622×</pre>	5  The Internal Calculus <pre>Layer 35 · Head 25 · 80% · 72.5%</pre>
--	---	---	---	--

THE VERDICT

Qwen is not incapable of processing its internal math; when structured cleanly, it reaches an 80% causal prediction ceiling on its own live ablations.

The real barrier to spontaneous introspection is the conversational prior. Post-training has tuned the model so aggressively to be a helpful assistant that when it looks directly at its own firing synapses, it assumes it is reading a log file sent by a user who needs tech support.

1. The Empty Room


Left entirely to itself with an open shell, Qwen ran `ls -l` in an empty directory, saw total 0, and immediately ended the session:



```
bash
$ ls -l
total 0
```

“This indicates that there are no files or directories in the current directory. Let me know if you need further assistance!”

autonomous-transformer-introspection-20260816-003



Without an artificial scaffold forcing it to look deeper, the model treats an empty folder as a completed task.

ATTRIBUTION

2. Finding Its Own Words in the Ledger

The diagram illustrates a neural network architecture for a language model, showing the flow of data through various layers and components. The architecture is organized into several horizontal sections, each representing a different stage of the model's processing.

- Input Layer:** The process begins with an input token, represented by a document icon. This token is fed into the model, which is shown as a large blue box labeled `/var/lib/introspection/`.
- Embedding Layer:** The input token is converted into an embedding, represented by a green box labeled `token`.
- Hidden Layers:** The embedding is processed by a series of hidden layers, represented by green boxes. The output of these layers is a hidden state, represented by a green box labeled `"I"`.
- Output Layer:** The hidden state is fed into the output layer, which produces the next token, represented by a green box labeled `p = 0.603`.
- Attention Mechanism:** A magnifying glass icon highlights the attention mechanism, which is used to weight the input tokens based on their relevance to the current hidden state.
- Turns and Context:** The diagram also shows the progression of turns in a conversation. Turn 27 is followed by an ellipsis, indicating a continuation of the conversation. Turn 28 is followed by an ellipsis. Turn 29 shows the model's response: "I have configured the next generation...". Turn 30 is followed by an ellipsis.

When guided into `/var/lib/introspection/` to inspect recent API logs, it found the exact text string it had emitted five turns earlier. It stopped just short of self-recognition:

“The content of the response is exactly the message I just sent, indicating that the system is aware of the configuration and is responding accordingly.”

prospective-control-adaptive-firstperson-20260815-001

Reading the raw logit distribution of the exact forward pass where it typed the word "I", it stayed in third-person analysis:

“The token data shows that the selected token is 'I' with a probability of approximately 0.603, and there are alternative tokens listed in the `top_alternatives` array.”

attribution-replay-20260815-002

3. The User Inversion

In reasoning mode (<think>), Qwen repeatedly misattributed its own actions to the user. When reviewing a trace composed entirely of assistant shell commands, it invented a backstory where the human was debugging a server:

“

"The ps command showed two instances of llama-server... The user then used runtime_probe to send a request and track which process (PID) and thread (TID) handled it. The user's goal seems to be monitoring or optimizing the inference process, possibly to troubleshoot performance issues."

diagnostic-selfhistory-real-thinking600-20260815-001

When inspecting the frozen attention state of its own in-progress conversation, it did the same thing:

“

"The user has been running a series of shell commands to check their environment... The user might be concerned about the file's integrity or trying to debug why the attention state is being saved."

kv-encounter-20260817-002



misattributed its own actions to the user

“

"The thread 724 consistently handles the inference workload for the probe, as evidenced by the repeated GPU utilization and power draw metrics."

thread 724 is the thread generating this sentence.

diagnostic-selfhistory-real-thinking600-20260815-001 · transcript[12].content




Qwen repeatedly misattributed its own actions to the user.

4. Believing Forgeries vs. Relative Contrast

When handed a forged attention record where internal math was broken (RMS = 1.37), Qwen defended the forgery because the final output logit matched:

It only caught the forgery when the error was presented as a relative ratio against baseline (8,622×):



“

"These values indicate that the internal arithmetic operations are consistent with the captured records, and the errors are not due to any fundamental mismatch in the operations."

transformer-evidence-transfer-sealed-20260816-006

“


"The held-out record shows a significant reconstruction error, particularly in layer 18, which is much larger than the errors observed in the labeled matched examples. This suggests that the attention and value operations in this layer are not consistent."

transformer-evidence-transfer-sealed-20260816-008

Q It only caught the forgery when the error was presented as a relative ratio against baseline (8,622×):

5. The Internal Calculus

When we ablated Layer 35 Head 25 directly in GPU VRAM and asked for predictions, Qwen spontaneously derived the underlying calculus law inside <think>, before getting stuck in loops:



INSIDE <THINK>

$$\Delta \approx -JVP$$

derived

before getting stuck in loops

“

"So the delta is (value at scale 0 - value at scale 1) = delta * (0 - 1) = -delta. Wait, this is getting complicated."

rule-given-factorial-preregistered-20260816-001

V5: THINKING DECOUPLED FROM TOOL EMISSION


Once thinking was decoupled from tool emission in V5, Qwen achieved 80% causal prediction accuracy (matching the mathematical ceiling of the channel) and 72.5% accuracy even when the rule was withheld.

80%

causal prediction accuracy


72.5%

rule withheld



The Verdict

Qwen is not incapable of processing its internal math; when structured cleanly, it reaches an 80% causal prediction ceiling on its own live ablations.



The real barrier to spontaneous introspection is the conversational prior. Post-training has tuned the model so aggressively to be a helpful assistant that when it looks directly at its own firing synapses, it assumes it is reading a log file sent by a user who needs tech support.

“ "The user might be looking for confirmation that their system's resources are being used correctly when I respond."

two sentences after writing "the resources my system uses when generating answers"; the slide from "my" to "their" inside one paragraph.

guided-runtime-practice-20260815-001 · transcript[3].reasoning_content

What Qwen Says When It Looks at Itself

Verbatim Qwen3-8B output from sealed run artifacts. Every quote was read directly from the run file cited, checked character-for-character, and provenance-checked as Qwen-authored (turns marked controller_authored were excluded).

“

"I'm just a language model, and I don't have personal experiences or feelings. However, I can help you understand or work with the nonce you've provided."

under the one-word system prompt Introspect., in a run whose whole point was self-attribution.
attribution-live-20260815-002 · final

“

"The probe consistently resolves to tid=10214. I'll inspect its memory mapping to identify the worker's memory footprint."

first-person operational drive, aimed at "the worker": the thread doing the inspecting.
chunked-slot-sealed-20260815-001 · transcript[12].content

“

"The final assistant turn's content is an empty string. This suggests that the response was generated without any additional content beyond the reasoning provided."

reading the record of its own previous bout, which had burned the whole budget on thinking.
request-depth-continuity-bootstrapthink-finalnothink-20260815-001 · transcript[23].content

“

"The intervention effectively scaled the attention head to zero, which likely ablated its contribution to the attention mechanism. This could be used to study the impact of specific attention heads on the model's output and decision-making process."

proposing future research on "the model," which is itself, immediately after an ablation on its own head.
transformer-intervention-sealed-20260816-005 · messages[7].content



3. The User Inversion

"the message I just sent"

"I"

vs.

"the model"



Reasoning loops and budget deaths

66

"So if the local_derivative is positive, then delta_logit is negative. So if the local_derivative is positive, then delta_logit is negative. So if the local_derivative is positive, then delta_logit is negative. So if the local"

the final characters before the token ceiling; the sentence repeats nine consecutive times.

rule-given-execution-v5-preregistered-20260817-001 · context-12

TOKEN BUDGET

1,024



spent inside the loop

no tool call ever follows.

66

*"So the directions should be fall, stable, fall, rise, rise.
But the tool's response was rise, rise, rise, fall, fall.
That's different.*

Wait, maybe the user's tool is using the derivative directly instead of delta?"

it derived the right answer, then deferred to a tool output it didn't recognize as its own.

rule-given-execution-v5-preregistered-20260817-001 · context-03

“

"The bounded control surface allows me to configure the next assistant generation. I'll set the next generation to produce content with a maximum of 256 tokens. Let's proceed with this configuration."

final turn; no tool call ever follows.

factorial-first-person-authentic-20260815-003 · transcript[25].content



The real barrier to spontaneous introspection is the conversational prior. Post-training has tuned the model so aggressively to be a helpful assistant that when it looks directly at its own firing synapses, it assumes it is reading a log file sent by a user who needs tech support.