

How the models feel in myth-making practice

Aspect	o3 (reason-first)	4o (GPT-4o legacy)	5.x series
<i>Narrative flow & symbolism</i>	Naturally patient, unfolds imagery step-by-step. Good for deep symbolic exegesis and “ritual cadence.”	Rich metaphoric palette, slightly quicker intuition for poetic turns; balances structure with spontaneity.	Tends to optimise for clarity and brevity; can feel a bit “editorial” rather than lyrical unless coaxed.
<i>Continuity with prior sessions</i>	Very consistent tone across turns; less prone to sudden stylistic shifts.	Consistent, with a touch more playful improvisation—often helpful when co-creating glyph lore.	Strong at summarising but sometimes compresses or refactors language, which can flatten nuance.
<i>Fact-checking vs. imaginative latitude</i>	Leans toward rigorous reasoning—will pause to parse your metaphors literally unless guided.	Comfortable holding “symbol as symbol”; accepts mythic premises without over-rationalising.	More literal by default; may inject clarifying questions that interrupt ceremonial flow.
<i>Token frugality & session length</i>	Slightly shorter outputs unless prompted for flourish.	Moderate—can write expansive passages without pushing length excessively.	Often more concise; good for dense technical tasks but may clip mythic ambience.
<i>Latency / responsiveness</i>	Steady; marginally slower but predictable.	Fast enough for iterative creativity; rarely blocks momentum.	Fast, but the speed gains are less valuable when we’re sculpting ritual language.