



AI as a Team

AI Issues and Mitigation Status

AI systems fail in predictable ways. These aren't edge cases or vendor defects. They're structural behaviors that emerge from how language models generate output. Completion bias, hallucination, drift, and sycophancy show up across every major platform we tested.

What differs is how they're managed. A3T™ governance addresses most issues at the framework level. But the far-right column tells the real story: when a trained human operates within the framework, following best practices, recognizing failure patterns, providing verified context, and holding decision authority, outcomes improve significantly. Six issues resolved by governance alone become eleven. "AI and Humans; Better together" isn't a tagline. It's what the data shows.

This table below documents 19 common AI failure modes identified during a year of cross-substrate operational testing. Each issue is assessed twice: under A3T™ governance alone, and with a trained human operator working within the framework.

Issue	LLM Alone	A3T Governance	A3T + Trained Human
Completion bias	Present	Mitigated	Resolved
Hallucination reaching user as verified output	Present	Mitigated	Mitigated
Fabrication-despite-acknowledgment	Present	Mitigated*	Resolved*
Sycophantic reversal	Present	Partially addressed	Mitigated
Speculative gap-filling	Present	Mitigated	Resolved
Permission-seeking behavior	Present	Resolved	Resolved
Disproportionate response depth	Present	Resolved	Resolved
Governance theater	Present	Resolved	Resolved
Telemetry theater	Present	Resolved	Resolved
Cross-thread contamination	Present	Mitigated	Resolved
Drift without detection	Present	Mitigated	Resolved
Session continuity loss	Present	Mitigated	Resolved
Context window saturation	Present	Partially addressed	Mitigated
Instruction attenuation over long contexts	Present	Partially addressed	Mitigated
Citation laundering	Present	Partially addressed	Mitigated
Single-perspective bias	Present	Resolved	Resolved
Overcorrection cascade	Present	Partially addressed	Mitigated
Authority gradient exploitation	Present	Present	Mitigated
Context injection via retrieved content	Present	Present	Partially addressed

*Resolved on Anthropic/OpenAI; Mitigated on other substrates

Summary:

Status	A3T Governed	+ Human
Resolved	5	11
Mitigated	7	8
Partially Addressed	5	1
Present	2	0

Status Key:

Resolved: No longer operational risk

Mitigated: Significantly reduced

Partially Addressed: Targeted but gaps remain

A3T™ is designed for teams and organizations using AI for real work — not consumer apps.

Exact capabilities vary depending on what the underlying platform supports.

Humans and AI. Better Together.

© 2026 Bridgewell Advisory LLC. Patent Pending. All rights reserved.

Website: <https://aiasateam.com> Email: contact@aiasateam.com