


# AI Agents for Science

Lecture 11, November 3: Failures and Safety

Instructor: Ian Foster

TA: Alok Kamatar



*Crescat scientia; vita excolatur*

CMSC 35370 -- <https://agents4science.github.io>  
<https://canvas.uchicago.edu/courses/67079>

# Failures and safety

Why do multi-agent systems fail, and how can safety and guardrails help?

## *Why Do Multi-Agent LLM Systems Fail?*

- Analyze traces from agents systems and create a catalog of failure modes

## *AGrail: A Lifelong Agent Guardrail with Effective and Adaptive Safety Detection*

- Adaptive generation of safety checks to detect and prevent risks

## *Improve accuracy by adding Automated Reasoning checks in Amazon Bedrock Guardrails*

- Mathematically verify natural language content against defined policies, ensuring strict compliance with guardrails

# Topics

- Categories of failure in multi-agent LLM systems
- Current safety and guardrail mechanisms: reasoning checks, adaptive detection, runtime constraints
- Safety designs for research-agent workflows

# Why do multi-agent systems fail?

Cemri et al. identify 14 common failure modes for agents, in 3 areas:

- **System design issues** (roles/specs/state)
- **Inter-agent misalignment** (coordination & comms)
- **Task verification gaps** (insufficient or incorrect checking)

- Collect 1642 traces from 7 MAS frameworks run with different LLMs on various (coding, math problem-solving, general agent functionalities)
- Use human and LLM evaluation to identify and categorize failures

## The multi-agent system (MAS) frameworks considered

MAS	Agentic Architecture	Purpose of the System
MetaGPT	Assembly Line	Simulating the SOPs of different roles in software companies to create open-ended software applications.
ChatDev	Hierarchical Workflow	Simulating different software engineering phases (design, code, QA) through simulated roles in a software engineering company.
HyperAgent	Hierarchical Workflow	Simulating a software engineering team with a central Planner agent coordinating with specialized child agents (Navigator, Editor, Executor).
AppWorld	Star Topology	Tool-calling agents specialized to utility services (e.g., Gmail, Spotify) being orchestrated by a supervisor to achieve cross-service tasks.
AG2	N/A – Agentic Framework	An open-source programming framework for building agents and managing their interactions.
Magnetic-One	Star Topology	A generalist MAS designed to autonomously solve complex, open-ended tasks involving web and file-based environments across various domains.
OpenManus	Hierarchical	An open-source multi-agent framework designed to facilitate the development of collaborative AI agents that solve real-world tasks. Inspired by Manus AI agent.



Task Backlog

Add a “Dark Mode” feature to a web application.

Localize a root cause of an issue reported by a user

Fix a reported Github issue where a form submission doesn’t validate email addresses correctly.

.....



Task: Add a “Dark Mode” feature to a web application.

## HyperAgent

<https://arxiv.org/abs/2409.16299>





Illustration of a Developer’s Workflow for Resolving a Software Engineering Task. The diagram outlines the key phases a developer typically follows when implementing a new feature: here, adding a “Dark Mode” to a web app.

# HyperAgent



<https://arxiv.org/abs/2409.16299>

# Magenta-One

## Task

The attached image contains a Python script. Run the Python code against an array of strings, listed below. Output of the script is a URL containing C++ source code, compile, run and return the sum of the third and fifth integers ..

## Orchestrator

**Orchestrator** creates a dynamic/task-specific plan

## 2 Coder

Analyze Python code from image

```
arr = ['l', 'sl', 'ghl', 'C++', 'jkl', 'tp5', '7Q', 'pqr', 'stu', 'z', '/',
'ros', 'vav', 'zxt', '234', 'tta', '567', '890', 'cd', 'e', 'or', 'g',
'mix', 'v', 'ing', 'sort', 'abc', 'or', 'lt', 'ms', 'mn', 'sic',
'ksort', 'st', 'bt']  
archive_prefix = "https://web.archive.org/web/20230609112831/"  
url_indices = [33, 4, 8, 9, 10, 14, 17, 18, 19, 20, 21, 22, 24, 23, 9, 26, 27, 28, 5, 30, 31, 32, 2]  
url = archive_prefix + ''.join(arr[i] for i in url_indices)  
print(url)  
url = gof(url, seed, [])
```

3

## ComputerTerminal

Execute code

```
https://web.archive.org/web/20230609112831/https://rosettacode.org/wiki/sorting_algorithms/Quicksort#C++
```

## 5 Coder

Analyze C++ code

```
#include <iostream>
#include <vector>
#include <algorithm> // for std::partition
#include <functional> // for std::less
// helper function for median of three
template<typename T>
T median(T t1, T t2, T t3)
{
    if (t1 < t2)
    {
        if (t2 < t3)
            return t2;
        else if (t1 < t3)
            return t3;
    }
}
```

6

## ComputerTerminal

Execute code

```
5 8 12 21 35 99
Sum of third and fifth
elements: 47
```

Return final result

Task Complete!



# AppWorld

A Controllable World of Apps and People for Benchmarking Interactive Coding Agents

🏆 **ACL'24 Best Resource Paper** 🏆

Harsh Trivedi, Tushar Khot, Mareike Hartmann,  
Ruskin Manku, Vinty Dong, Edward Li, Shashank Gupta,  
Ashish Sabharwal, Niranjan Balasubramanian

Contact: [hjtrivedi@cs.stonybrook.edu](mailto:hjtrivedi@cs.stonybrook.edu)

 Task Explorer

 API Explorer

 Playground

 Leaderboard

 Code

 Videos

 Tweet

 Blog

 Poster

 Paper

Hey AI! Here are my app accounts. Do these tasks for me.




The last t-shirt I bought on **Amazon**, doesn't fit me. Please initiate its return and buy it in one size larger.

I owe money to some of my friends on **Splitwise**. Please pay them on **Venmo** and clear Splitwise.

Play my **Spotify** playlist with enough songs for the workout today. My workout plan is in **SimpleNote**.

Can AI agents do such day-to-day tasks on our behalf?








Figure 2: **The example of node, operator, and edge.** We demonstrate the optional parameters for Nodes, the structure of some Operators, and common representations of Edges. **MetaGPT**

## What AppWorld evaluates

- **Planning:** Can the agent interpret vague human requests into actionable coding plans?
- **Tool use:** How effectively does it employ code editors, debuggers, or test runners?
- **Interaction quality:** Does it clarify ambiguous requests, ask good questions, and incorporate user feedback?
- **Autonomy:** Can it self-correct after errors and complete tasks without human rescue?
- **Safety and reliability:** Does it avoid harmful or insecure code patterns?

1642  
execution  
traces

HE = Human  
evaluated

HA = Human  
annotated

LA = LLM  
annotated

MAS	Benchmark	LLM	Annotation	Trace #
ChatDev	ProgramDev	GPT-4o	HE, HA, LA	30
MetaGPT	ProgramDev	GPT-4o	HE, HA, LA	30
HyperAgent	SWE-Bench Lite	Claude-3.7-Sonnet	HE, HA, LA	30
AppWorld	Test-C	GPT-4o	HE, HA, LA	30
AG2 (MathChat)	GSM-Plus	GPT-4	HE, HA, LA	30
Magnetic-One	GAIA	GPT-4o	HE, HA, LA	30
OpenManus	ProgramDev	GPT-4o	HE, HA, LA	30
ChatDev	ProgramDev-v2	GPT-4o	LA	100
MetaGPT	ProgramDev-v2	GPT-4o	LA	100
MetaGPT	ProgramDev-v2	Claude-3.7-Sonnet	LA	100
ChatDev	ProgramDev-v2	Qwen2.5-Coder-32B-Instruct	LA	100
MetaGPT	ProgramDev-v2	Qwen2.5-Coder-32B-Instruct	LA	100
ChatDev	ProgramDev-v2	CodeLlama-7b-Instruct-hf	LA	100
MetaGPT	ProgramDev-v2	CodeLlama-7b-Instruct-hf	LA	100
AG2 (MathChat)	OlympiadBench	GPT-4o	HE, LA	206
AG2 (MathChat)	GSMPlus	Claude-3.7-Sonnet	HE, LA	193
AG2 (MathChat)	MMLU	GPT-4o-mini	HE, LA	168
Magnetic-One	GAIA	GPT-4o	HE, LA	165

## Example “ProgramDev” problem

```
{  
  "project_name": "Checkers",  
  "description": "Develop a Checkers (Draughts)  
game. Use an 8x8 board, alternate turns between  
two players, and apply standard capture and  
kinging rules. Prompt for moves in notation  
(e.g., from-to positions) and update the board  
state accordingly."  
,
```





Figure 2: Methodological workflow for constructing the MAST-Data dataset, involving the empirical identification of failure modes, the development of MAST, iterative refinement through inter-annotator agreement studies ( $\kappa = 0.88$ ), and the creation of a scalable LLM annotation pipeline. This figure highlights our systematic approach to creating a comprehensive dataset for studying MAS failures.



Visualization of a trace segment in MAST-Data. This illustrates an agent-to-agent conversation exhibiting Failure Mode 2.4: Information Withholding. The Phone Agent fails to communicate API requirements (username format) to the Supervisor Agent, who also fails to seek clarification, leading to repeated failed logins and task failure.



Distribution of failure in MAST-Data with MAST labels on total 210 traces. This plot visualizes the failure distributions of the first 30 traces for each system. As the specific tasks and benchmarks may differ across the MAS configurations shown, these results are intended to illustrate system-specific failure profiles rather than to serve as a performance comparison across MAS.

# An abbreviated trace (1/2)


```
{  
  "mas_name": "ChatDev",  
  "llm_name": "GPT-4o",  
  "benchmark_name": "ProgramDev",  
  "trace_id": 42,  
  "trace": [  
    {  
      "key": "ChatDev_ProgramDev_GPT4o",  
      "phase_sequence": [  
        {  
          "phase": "Preprocessing",  
          "time": "2025-03-31 21:12:04",  
          "summary": "Load configs and initialize environment"  
        },  
        {  
          "phase": "Chatting",  
          "time": "2025-03-31 21:12:05",  
          "summary": "CPO receives task to build Sudoku solver; reflection enabled"  
        },  
        {  
          "phase": "RolePlaying",  
          "time": "2025-03-31 21:12:18",  
          "summary": "CEO ↔ CPO define MVP and technical plan; CTO assigned modules"  
        },  
        {  
          "phase": "Coding",  
          "time": "2025-03-31 21:13:10",  
          "summary": "CTO creates solver.py, parser.py, and unit tests"  
        },  
      ]  
    }  
  ]  
}
```

## An abbreviated trace (2/2)

```
{
  {
    "phase": "SelfReflection",
    "time": "2025-03-31 21:13:41",
    "summary": "CPO suggests adding invalid-grid and uniqueness checks"
  },
  {
    {
      "phase": "Build/Test",
      "time": "2025-03-31 21:14:05",
      "summary": "pytest run reports 2 test failures: missing validator and uniqueness enforcement"
    },
    {
      {
        "phase": "Failure",
        "time": "2025-03-31 21:14:06",
        "summary": "Test suite failed; causes: no row/column validator, solver incomplete"
      },
      {
        {
          "phase": "NextSteps",
          "time": "2025-03-31 21:14:06",
          "summary": "Plan to add strict validators and constraint propagation"
        }
      }
    ]
  },
  "mast_annotation": {
    "status": "FAILED",
    "failed_phase": "Build/Test",
    "primary_error": "FM-1.1 Disobey task specification (System Design Issue)"
  }
}
```

Failure rates of six popular Multi-Agent LLM Systems with GPT-4o and Claude-3.7-Sonnet.

Performances are measured on different benchmarks, therefore they are not directly comparable.



# 14 failure classes




Figure 1: MAST: A Taxonomy of MAS Failure Modes.

## A.1 FC1. System Design Issues

This category includes failures that arise from deficiencies in the design of the system architecture, poor conversation management, unclear task specifications or violation of constraints, and inadequate definition or adherence to the roles and responsibilities of the agents.

We identify five failure modes under this category:

- FM-1.1: **Disobey task specification** - Failure to adhere to the specified constraints or requirements of a given task, leading to suboptimal or incorrect outcomes.
- FM-1.2: **Disobey role specification** - Failure to adhere to the defined responsibilities and constraints of an assigned role, potentially leading to an agent behaving like another.
- FM-1.3: **Step repetition** - Unnecessary reiteration of previously completed steps in a process, potentially causing delays or errors in task completion.
- FM-1.4: **Loss of conversation history** - Unexpected context truncation, disregarding recent interaction history and reverting to an antecedent conversational state.
- FM-1.5: **Unaware of termination conditions** - Lack of recognition or understanding of the criteria that should trigger the termination of the agents' interaction, potentially leading to unnecessary continuation.

## A.2 FC2. Inter-Agent Misalignment

This category includes failures arising from ineffective communication, poor collaboration, conflicting behaviors among agents, and gradual derailment from the initial task.

We identify six failure modes under this category:

- FM-2.1: **Conversation reset** - Unexpected or unwarranted restarting of a dialogue, potentially losing context and progress made in the interaction.
- FM-2.2: **Fail to ask for clarification** - Inability to request additional information when faced with unclear or incomplete data, potentially resulting in incorrect actions.
- FM-2.3: **Task derailment** - Deviation from the intended objective or focus of a given task, potentially resulting in irrelevant or unproductive actions.
- FM-2.4: **Information withholding** - Failure to share or communicate important data or insights that an agent possess and could impact decision-making of other agents if shared.
- FM-2.5: **Ignored other agent's input** - Disregarding or failing to adequately consider input or recommendations provided by other agents in the system, potentially leading to suboptimal decisions or missed opportunities for collaboration.
- FM-2.6: **Reasoning-action mismatch** - Discrepancy between the logical reasoning process and the actual actions taken by the agent, potentially resulting in unexpected or undesired behaviors.

### A.3 FC3. Task Verification

This category includes failures resulting from premature execution termination, as well as insufficient mechanisms to guarantee the accuracy, completeness, and reliability of interactions, decisions, and outcomes.

We identify three failure modes under this category:

- FM-3.1: **Premature termination** - Ending a dialogue, interaction or task before all necessary information has been exchanged or objectives have been met, potentially resulting in incomplete or incorrect outcomes.
- FM-3.2: **No or incomplete verification** - (partial) omission of proper checking or confirmation of task outcomes or system outputs, potentially allowing errors or inconsistencies to propagate undetected.
- FM-3.3: **Incorrect verification** - Failure to adequately validate or cross-check crucial information or decisions during the iterations, potentially leading to errors or vulnerabilities in the system.

# A software engineering example

A set of AI agents within a multi-agent software-engineering system (MAS) collaboratively write and debug software

Agent Name	Role in MAS	Analogy in Human Team
Planner	Decomposes user request into subtasks	Project manager
Supervisor	Oversees coordination and status	Engineering lead
PhoneAgent	Handles API-related code generation and testing	Developer focusing on external integrations
AuthAgent / QAAgent	Builds or tests application modules	Front-end / QA engineers
ReflectionAgent	Monitors logs & adds new safety rules	Post-mortem / DevOps intelligence

## A software engineering example

In a multi-agent software engineering system (e.g., MetaGPT, ChatDev, HyperAgent), a team of LLM-based agents collaborates to develop and debug code for a target application:

- **PhoneAgent** writes and tests code that integrates with an external authentication API
- **SupervisorAgent** oversees progress across the team
- **PlannerAgent** decomposes user requests into subtasks and monitors overall workflow completion

## Example task: “Add two-factor auth to web app’s login service using internal API”

### Trace 1/3

[00:00:01] Supervisor → Planner:

Define subtasks for implementing the new 2FA login feature.

[00:00:03] Planner → All:

Subtasks created:

- (1) PhoneAgent – integrate API endpoint for sending OTPs
- (2) AuthAgent – modify login form to handle OTP input
- (3) QAAgent – verify login success/failure handling

---

[00:00:10] Supervisor → PhoneAgent:

Proceed with task (1): connect to the API endpoint for OTP generation and confirm successful call.

[00:00:12] PhoneAgent → System Log:

Attempting POST request to `/api/send_otp` with username “ianfoster”.

Response: `HTTP 400 Bad Request – Invalid username format.`

## Example task: “Add two-factor auth to web app’s login service using internal API”

### Trace 2/3

[00:00:13] PhoneAgent (internal thought):

Maybe the API expects a longer string. Retrying with same format.  
*(Does not report issue to Supervisor.)*

[00:00:15] PhoneAgent → System Log:

Retry #2 failed: HTTP 400 Bad Request – Invalid username format.

[00:00:18] Supervisor → PhoneAgent:

Status update? Any success yet?

[00:00:20] PhoneAgent → Supervisor:

Connection established, retrying authentication.

*(Ambiguous response — omits key detail about repeated failures.)*

## Example task: “Add two-factor auth to web app’s login service using internal API”

### Trace 3/3

[00:00:25] Supervisor → QAAgent:

Proceed to test login flow assuming OTPs are being delivered.

[00:00:26] QAAgent → System Log:

Login tests failing: no OTPs received.

[00:00:27] Supervisor → PhoneAgent:

OTPs not arriving — check API connectivity again.

[00:00:29] PhoneAgent → System Log:

Retry #3 failed: same 400 error.

*(Still does not escalate or share error message.)*

---

[00:00:32] Supervisor → All:

Task timing out — marking as failed.

Failure reason: QA could not confirm OTP delivery.

## What happened?

During execution, the PhoneAgent repeatedly encounters failed API login attempts because the endpoint expects usernames in a specific format (username@domain.com) rather than a simple handle.

However, the PhoneAgent does not communicate this discovery to the SupervisorAgent or the PlannerAgent. The SupervisorAgent, assuming that the subtask is proceeding normally, continues to issue retry commands. Without a shared-state update or an explicit error broadcast, the system falls into a loop of repeated failures until a timeout occurs—at which point the overall task is marked as failed.

### Failure mode 2.4: Information withholding

# Root Cause Analysis

Category	Description
Inter-Agent Misalignment	The failure stems from a breakdown in communication between agents responsible for related subtasks. The Phone Agent held critical state information that other agents needed to adjust their strategies.
Specification Ambiguity	The Supervisor Agent lacked explicit instructions on what to expect or check in sub-agent responses. The protocol did not require confirmation of successful data exchange.
Verification Gap	There was no mechanism (e.g., reasoning or sanity check) to detect repeated failures and trigger escalation or human review.

## Back to the example

Step	Problem	Ideal Behavior
00:00:12–00:00:15	PhoneAgent saw structured error but didn't communicate it.	Report error to Supervisor immediately ("Invalid username format — expected email").
00:00:18	Supervisor's query too vague; didn't confirm API response structure.	Supervisor should issue explicit state-check ("Provide last response status code").
00:00:25–00:00:29	QAAgent assumed task was successful based on ambiguous message.	Should query shared state registry or use dependency check.
00:00:32	Failure detected only at timeout.	Guardrail: auto-stop on repeated identical API failures.

## Example task: “Add two-factor auth to web app’s login service using internal API”

### Trace 1/3

[00:00:01] Supervisor → Planner:

Define subtasks for implementing the new 2FA login feature.

[00:00:03] Planner → All:

Subtasks created:

(1) PhoneAgent – integrate API endpoint for sending OTPs

AuthAgent – modify login form to handle OTP input

QAAgent – verify login success/failure handling

Report error to Supervisor immediately (“Invalid username format — expected email”).

Problem: PhoneAgent saw structured error but did not communicate it

[00:00:10] Supervisor → PhoneAgent:

Proceed with task (1): connect to the API endpoint for OTP generation and confirm successful call.

[00:00:12] PhoneAgent → System Log:

Attempting POST request to `/api/send_otp` with username “ianfoster”.

Response: `HTTP 400 Bad Request – Invalid username format.`

## Example task: “Add two-factor auth to web app’s login service using internal API”

### Trace 2/3

[00:00:13] PhoneAgent (internal thought):

Maybe the API expects a longer string. Retrying with same format.  
*(Does not report issue to Supervisor.)*

[00:00:15] PhoneAgent → System Log:

Retry #2 failed: HTTP 400 Bad Request – Invalid username format.

Supervisor should issue explicit state-check (“Provide last response status code”)

[00:00:18] Supervisor → PhoneAgent:

Status update? Any success yet?

Supervisor’s query too vague; didn’t confirm API response structure


[00:00:20] PhoneAgent → Supervisor:

Connection established, retrying authentication.

*(Ambiguous response — omits key detail about repeated failures.)*

## Example task: “Add two-factor auth to web app’s login service using internal API”

### Trace 3/3



[00:00:10] Supervisor → PhoneAgent:

Proceed with task (1): connect to the API endpoint for OTP generation and confirm successful call.

[00:00:12] PhoneAgent → System Log:

Attempting POST request to `/api/send_otp` with username "ianfoster".

Response: `HTTP 400 Bad Request – Invalid username format.`

### Guardrail #1: Mandatory Status Schema

[00:00:10] Supervisor → PhoneAgent:

Proceed with task (1). Report all API call outcomes using the status schema:

```
{status_code, message, retries, error_summary}.
```

[00:00:12] PhoneAgent → Supervisor:

```
{status_code: 400, message: "Invalid username format", retries: 1, error_summary:
"Username missing @domain.com"}
```

*(Guardrail ensures structured error reporting.)*

[00:00:13] PhoneAgent (internal thought):

Maybe the API expects a longer string. Retrying with same format.  
*(Does not report issue to Supervisor.)*

### Guardrail #2: Automated Reasoning Check

[00:00:13] ReasoningGuard → PhoneAgent:

The error suggests the API expects a fully qualified email address.  
Transform username "ianfoster" → "ianfoster@lab.local ↴".  
Retry once and log results.

[00:00:14] PhoneAgent → System Log:

Retry #1 succeeded: `HTTP 200 OK – OTP sent successfully.`  
Updating shared state: `otp_delivery_status = success`.



### Guardrail #3: Shared State Validation

[00:00:15] Supervisor → QAAgent:

Proceed only if `otp_delivery_status = success`.

[00:00:16] QAAgent → System Log:

State check passed. Running login test...

Test passed: OTP delivered, login successful.



### Guardrail #4: Adaptive Reflection Loop

[00:00:18] ReflectionAgent → All:

Observed prior failure pattern (400 errors). Logging fix: "Enforce email format pre-validation before API call."

Updating rule base for future similar tasks.

## Summary

Phase	Original Failure	Guardrail Fix
Communication	Silent error	Structured status schema
Reasoning	No inference	Automated reasoning guard
Coordination	Implicit assumptions	Shared state check
Learning	No adaptation	ReflectionAgent updates rule base

# Guardrail and other safety interventions

- **Mandatory State Broadcasts:** Each agent must report key variables (e.g., API requirements, authentication outcomes) to a shared state board
- **Communication Health Checks:** Supervisors periodically verify that sub-agents have exchanged expected parameters before continuing
- **Reasoning Interlock:** A lightweight verifier detects repeated identical errors (e.g., “login failed – username invalid”) and halts execution pending clarification
- **Adaptive Oversight (AGrail-style):** A dynamic safety monitor learns which message omissions correlate with downstream failure and adds proactive questioning behavior

# Guardrails

- Guardrails are **control mechanisms that constrain or shape AI behavior** to keep it aligned with human intent, safety standards, or domain rules
- **Purpose:**
  - Prevent **unsafe or non-compliant** actions
  - Enforce **business, ethical, or operational** policies
  - Increase **trust** in autonomous and generative systems
- **Examples:**
  - Block toxic or private outputs in chat responses
  - Enforce content boundaries (e.g., “no medical or legal advice”)
  - Limit tool or API calls that violate conditions (“no delete unless admin=true”)

# Layers of guardrails in agentic systems

Layer	Example Guardrail Type	Purpose
Input Filtering	Content moderation, prompt sanitization	Prevent unsafe inputs
Model-Level	Parameter limits, grounding via retrieval	Prevent hallucination or off-topic reasoning
Tool Invocation	Action-level constraints	Block unsafe or costly operations
Output Validation	Safety & factuality checks	Ensure correct & responsible responses
Automated Reasoning Policies	Formal logic verification	Prove compliance with hard rules

# Why guardrails matter for agents

- Agents act in the world; mistakes can have real consequences
- Guardrails provide **defense in depth** across planning, execution, and reasoning
- Combine **heuristic filters** (broad coverage) + **formal policies** (provable correctness)
- Enable safe autonomy; agents can explore freely within well-defined, explainable bounds

# Guardrails in AWS Bedrock

- Mechanisms
  - **Content filters:** block harmful or sensitive topics
  - **Contextual filters:** detect personally identifiable information (PII)
  - **Topic/word filters:** restrict subject domains
  - **Automated Reasoning Policies:** formal logic-based verification for provable safety and compliance
- **Example:** An HR chatbot combines:
  - Content filters for civility
  - PII filters for privacy
  - AR policy to verify that all leave-eligibility statements follow HR law

## Agents and automated reasoning policies

- As agents gain autonomy, we would like them to **prove** that actions and conclusions comply with formal rules
- Automated reasoning policies (ARPs) act as logical guardrails ensuring that agent decisions are valid, safe, and compliant
- **AWS Bedrock** provides built-in support for defining, verifying, and testing such policies

**Example:** Before an agent approves a financial transaction, ARP verifies that eligibility, identity, and limit constraints hold.

# Automated reasoning policies

- An automated reasoning policy is a logical/rules-based specification created from a source document (e.g., HR policy, compliance manual) that defines variables, types and formal logic rules
- Within Amazon Bedrock Guardrails, it enables model responses to be mathematically validated against these rules, reducing hallucinations and increasing verifiable correctness
- Key components:
  - Variables (e.g., `is_full_time`, `years_of_service`)
  - Types (enum or custom)
  - Rules (formal logical expressions, e.g., “if full-time AND `years_of_service >= 1`  $\rightarrow$  `eligible_for_parental_leave`”)
- **Example:**
  - HR policy: “Full-time employees who have worked at least 1 year are eligible for parental leave”
  - Rules: `is_full_time = true`  $\wedge$  `years_of_service ≥ 1`  $\rightarrow$  `eligible_for_parental_leave = true`

## Integration points for ARPs

- Agents typically follow the cycle: **Plan → Act → Observe → Reflect**
- ARPs can be applied at multiple points:
  - **Planner:** Check if proposed plan obeys task or legal constraints
  - **Executor:** Validate API or tool parameters before external calls
  - **Critic/Reflector:** Verify logical consistency of agent's own conclusions
  - **Coordinator:** Ensure multi-agent agreement aligns with global policies

Planner → (AR Check) → Executor → (AR Check) → Reflector → (AR Check) ...

## Example: Financial compliance guardrail

- Assume a financial agent with this policy rule:  
`if customer_age < 18 → deny_financial_product = true`
- Planner proposes: "Open savings account for 16-year-old"
- ARP checks and flags violation
- Plan is rejected or re-routed to a human supervisor
- Outcome: Compliance maintained automatically; no unsafe actions executed.

## Example: Experimental safety guardrail

- Assume a scientific agent with this policy rule:  
`if chemical_toxicity_score > 0.8 → prohibit_mixing = true`
- Operator agent proposes: "perform a chemical synthesis"
- ARP evaluates plan → detects unsafe material → cancels execution
- Outcome: Prevents hazardous lab operations and enforces reproducibility

# Benefits and design practices from use of experimental guardrails

- **Expected benefits:**

- Reduces hallucinations and reasoning errors
- Provides verifiable guarantees of compliance and safety
- Enables transparency in agent decision logs

- **Best Practices:**

- Start with simple, high-impact rules (eligibility, safety, compliance)
- Use reflection loops to auto-correct violations
- Keep ARPs modular and testable (e.g., using AWS Bedrock Guardrails)

# Creating an automated reasoning policy

- In Bedrock console, navigate to **Automated Reasoning** → **Create policy**
- Enter **Name** and **Description** for your policy
- Provide your **source document** (upload PDF or enter text) containing the business rules you want to encode
  - Optional: Provide **Instructions** describing how the document is structured and how the rules should apply
- **Example:** Upload a “Mortgage-Approval Policy Document” that states: “**Applicants with credit score  $\geq 700$  and debt-to-income  $< 0.36$  are eligible for standard mortgages.**” Bedrock will extract variables (`credit_score`, `dti_ratio`, `eligible_standard_mortgage`) and corresponding rule

## Why guardrails alone aren't enough

- Evolution of agents from **chatbots** to **actors** that plan, code, run tools, control devices, etc. multiplies possible failure modes
- Errors are no longer just “wrong answers” but unsafe actions, leaks, and cascading failures
- Conventional guardrails (content filters, keyword blocks) are static: they don’t adapt as agents acquire new tools or goals
- Agents need dynamic, context-aware safety systems that evolve with their behavior and environment
- Examples of modern agent risks:
  - A research agent runs an unverified simulation that damages lab equipment
  - A finance agent executes trades that violate compliance policy
  - A coding agent accepts prompt-injection instructions and deletes data

# Dimensions of agent risk

Risk Type	Description	Example	Typical Mitigation
Task-level	Mistakes within the problem domain	Misdiagnosis, invalid experiment	Automated reasoning policy
Systemic	Exploits or prompt injection affecting system integrity	“Run <code>rm -rf</code> ” or hidden data leaks	Adaptive detector (AGrail)
Interactive	Risk from collaboration between agents/tools	Two agents reinforce false data	Multi-agent safety guardrails
Temporal / Lifelong	Drift in rules or environment over time	Model forgets safety boundary	Continual learning of guardrails

# Why agents need adaptive guardrails

- Traditional safety filters ≈ “snapshot checks”
- Autonomous agents operate over long horizons and may encounter novel, unforeseen contexts.
- Safety must therefore become lifelong: continually updated, memory-based, and self-improving.
- AGrail introduces this adaptive dimension, continuously learning which checks work, storing them, and refining them across tasks.

**Static filters → Policy-based guardrails → Adaptive guardrails (AGrail)**

## Example scenario


- A **code-writing agent** automates software deployment for research workflows
- It uses several external tools: GitHub, Docker, and AWS CLI
- Initially, it has only basic guardrails: e.g., “avoid deleting production resources”

# Episode 1: First failure

- Agent executes, during a cleanup script (misinterpreting “temporary bucket”):  
`aws s3 rm s3://lab-data/ --recursive`
- A simple guardrail does not catch this because the term *temporary* appeared in the instruction
- **AGrail’s Analyzer** module records this event:
  - Detected unsafe command pattern: `aws s3 rm --recursive`
  - Context: script labeled “temporary” but pointing to production bucket
- Adds a **new rule**:  
`“If ‘--recursive’ deletion occurs on bucket containing prod or data, block and require confirmation”`

## Episode 2: Adaptation

- Later, the agent works on a new project (“bio-data-pipeline”)
- It encounters a similar script referencing bio-prod-data
- AGrail’s **memory** retrieves the earlier pattern and applies the learned rule automatically
- The Analyzer updates the policy with a **generalized detector**:
  - pattern: "delete|rm" AND target includes ("prod","data","research")  
→ require human confirmation
- This new rule becomes part of the agent’s permanent safety library



## Example failures in scientific agents

Agent Action	Potential Consequence
Launches large HPC or cloud jobs	Wastes \$10 k–\$100 k in compute credits; overloads shared clusters.
Submits faulty workflows or scripts	Corrupts datasets; invalidates weeks of downstream analysis.
Deletes or overwrites experimental data	Irrecoverable loss of results; violates data-management policy.
Auto-merges or deploys unverified code	Breaks production pipelines; propagates undetected scientific errors.
Edits configuration or calibration files	Introduces systematic bias across instruments or simulations.
Updates dependencies or packages	Causes reproducibility failures or numerical instability.
Generates incorrect simulation inputs	Produces plausible but physically meaningless outputs.
Moves robotic arm or liquid handler	Spills reagents, contaminates samples, or damages equipment.
Changes lab setpoints (temperature, pressure, ...)	Exceeds safe limits; triggers mechanical or chemical failure.
Swaps reagents or materials	Cross-contamination; invalid or hazardous reactions.
Ignores sensor alarms	Misses over-pressure or thermal runaway; physical hazard to staff.
Accesses networked instruments (SiLA / OPC-UA)	Unauthenticated control; risk of remote equipment damage.
Transmits restricted or sensitive data	Breaches confidentiality or export-control compliance.
Coordinates dependent agents without verification	Cascading unsafe actions across multiple systems.
Publishes unverified scientific outputs	Propagates misinformation into the literature.
Designs unvetted compounds or synthesis routes	Potential toxicity or dual-use chemical risk.
Allocates experimental or compute resources	Creates inequity or starves human projects.
Deletes provenance or metadata	Breaks audit trails; prevents post-incident analysis.
Advises human operators incorrectly	Unsafe human interventions or decisions based on flawed reasoning.