The Control Layer for Governed Multi-Agent Enterprise Workflows

As AI moves from isolated tools to coordinated execution, the real challenge becomes **governance, visibility, and control** across complex business workflows.

Why Current Approaches Fall Short

Internal Builds
Flexible, but slow, expensive, and hard to govern.

Stitched-Together Tools
Fast to demo, but brittle under real complexity.

InnervationAI

A centralized orchestration layer coordinating specialized agents while maintaining observability, embedded governance, and operational control.

Core Capabilities

Centralized Orchestration

Coordinate agents and workflows from a single control plane.

Embedded Governance

Policy-aware execution built into every workflow step.

Transparent Coordination

Traceable decision paths and full audit visibility.

Where It Applies

Enterprise Finance

- Treasury operations
- Reporting support
- Compliance

Operations

- Exception handling
- Service automation
- Coordination

Regulated Needs

- Document review
- Policy checks
- Audit-ready

SME Operations

- Sales & Marketing
- Campaign coordination
- Customer workflows

How to Start

01 Discovery
Identify workflow & constraints

02 Pilot Scoping
Define architecture

03 Proof of Value
Validate outcomes

04 Enterprise Scale
Extend architecture

G7 GovAI Grand Challenge funding recipient | Canada

Start with one workflow that matters.
Let's identify where governed multi-agent execution can create measurable value.

Remove Ableist Language from Official Documents and Communications Before They Reach the Public

From policies and press releases to forms and emails, organizations publish thousands of messages every year, and even the most well-intentioned teams can miss ableist language without in-house accessibility expertise. **The Ableism Debiaser** adds an intelligent review layer that **flags, explains, and helps correct** problematic language **before publication**.

The Challenge

- Ableism appears in subtle language choices
- Manual accessibility review is slow
- Existing toxicity or content filters are not tuned to disability bias, especially implicit ableism

How the Ableism Debiaser Works

1. Teams upload or connect documents
2. The system analyzes language and context
3. Language is benchmarked
4. System detects ableist language, proposes alternative, explains recommendation.

What Organizations Gain

- ✓ 90% Reduction of Manual Review
- ✓ Consistent Systematic Review of Content
- ✓ Elimination of Harmful Messages and Risk
- ✓ Clear Audit Trail & Transparency

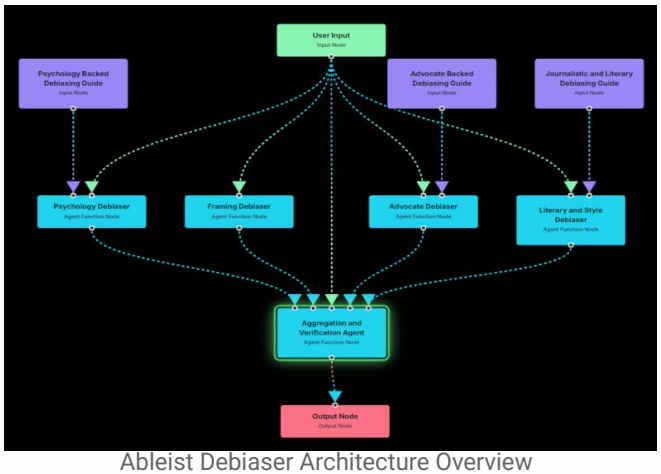

Ableist Debiaser Architecture Overview

Winning solution in the G7 GovAI Grand Challenge for tackling real-world public-sector AI adoption barriers in a responsible way.

G7 GovAI Grand Challenge funding recipient | Canada

Built for Public-Sector and Institutional Use

→ Designed for teams responsible for accessibility, equity, and policy communications in governments and large organizations.

→ Aligns with responsible AI expectations in the G7 and EU context, supporting practical, scalable bias mitigation in day-to-day operations.

Sample Debiased Documents

Job Posting | Emergency Preparedness Guide | University Student Life Guide

Next Steps

1. Book a live demo of the Ableism Debiaser.
2. Identify 1-2 priority content streams (e.g., HR policies, press releases) for a time-boxed pilot.
3. Deployment and integration plan, then move to full contract based on pilot results and scope.

Contact:
contact@innervationai.com
www.innervationai.com