
Project report template¹

Anurag Dhungana
George Washington
University

Sunil Shah
George Washington
University

Prakriti Bista
George Washington
University

With
Johns Hopkins University & Apart Research

Abstract

The AI Parliament project introduces an innovative approach to AI policy deliberation through a simulated parliamentary system. This technical solution addresses the growing need for structured, multi-perspective evaluation of AI policies. The system implements a novel debate simulation framework where AI agents, representing different societal stakeholders (corporate, academic, government, and civil rights), engage in policy discussions and voting processes. The implementation leverages advanced NLP techniques through the OpenAI API and a sophisticated vote decision service incorporating role-specific weights and sentiment analysis. Key innovations include a consistent voting mechanism, real-time debate monitoring, and integration with ArXiv for contemporary AI research topics. Results demonstrate the system's capability to generate nuanced policy discussions with 85% vote consistency across different stakeholder perspectives.

Keywords: AI Governance, Multi-Agent Systems, Policy Simulation, Sentiment Analysis

¹AI Policy Hackathon at Johns Hopkins University 2024

1. Introduction

Problem Statement

Current AI policy development lacks systematic methods for evaluating policies from multiple stakeholder perspectives. Manual policy review processes are often slow, biased, and fail to consistently consider diverse viewpoints.

Current Limitations

- Limited scalability in policy review processes
- Inconsistent consideration of stakeholder perspectives
- Lack of quantitative metrics for policy evaluation
- Time-intensive deliberation processes

Technical Innovation

The AI Parliament introduces several key innovations:

a) Role-Based AI Agents:

```
self.mp_roles = {
    "corporate": {
        "description": "Represents business and
industry interests",
        "bias": "Favors market-driven solutions and
minimal regulation",
        "color": "#DA0211" # Business blue
    },
    "academic": {
        "description": "Represents academic and
research institutions",
        "bias": "Favors evidence-based policy and
thorough research",
        "color": "#FDA003" # Academic purple
    },
    "government": {
        "description": "Represents governmental and
regulatory interests",
        "bias": "Favors structured oversight and public
safety",
        "color": "#2CAFFE" # Government green
    },
    "civil_rights": {
        "description": "Represents civil society and
individual rights",
```

```

        "bias": "Favors privacy and individual
protections",
        "color": "#000099" # Advocacy red
    }
}

```

b) Weighted Decision System:

```

self.role_weights = {
    "corporate": {
        "economic": 0.8,
        "innovation": 0.7,
        "regulation": -0.6,
        "market": 0.9
    },
    "academic": {
        "research": 0.9,
        "evidence": 0.8,
        "innovation": 0.7,
        "ethics": 0.6
    },
    "government": {
        "safety": 0.8,
        "regulation": 0.7,
        "economic": 0.5,
        "implementation": 0.6
    },
    "civil_rights": {
        "privacy": 0.9,
        "ethics": 0.8,
        "transparency": 0.7,
        "rights": 0.9
    }
}

```

c) Real-time Monitoring:

```

async def record_metric(
    self,
    debate_id: int,
    mp_role: str,

```

```

        response_content: str,
        vote_decision: Dict[str, any],
        vote_service: 'VoteDecisionService'
    ) -> float:
        """
        Record and analyze vote consistency.

        Args:
            debate_id: The debate identifier
            mp_role: The MP's role
            response_content: The MP's debate response
            vote_decision: The voting decision data
            vote_service: Instance of VoteDecisionService

        Returns:
            float: Consistency score between 0 and 1
        """
        # Analyze response sentiment using existing service
        response_analysis =
vote_service.analyze_response_sentiment(response_content)

        # Calculate consistency score
        role_weights = vote_service.role_weights[mp_role]
        sentiment_score = sum(
            response_analysis.get(aspect, 0) * weight
            for aspect, weight in role_weights.items()
        )

        # Compare sentiment with vote
        vote_value = {
            "for": 1.0,
            "against": -1.0,
            "abstain": 0.0
        }[vote_decision["vote"]]

        # Calculate consistency (1 = perfectly consistent, 0 =
completely inconsistent)
        consistency = 1.0 - abs(
            (np.tanh(sentiment_score) - vote_value) / 2
        )

```

2. Methods

The AI Parliament has been set up as a sophisticated multi-agent system designed to simulate parliamentary debates on AI policy issues. Here's how it's structured:

MP Roles and Representation

The parliament consists of four distinct MP roles (defined in `backend/role_definitions.md`, lines 1-30):

- **Corporate Representative:** Focuses on economic growth and market efficiency
- **Civil Rights Advocate:** Champions privacy and individual rights
- **Academic Representative:** Emphasizes research integrity and scientific merit
- **Government Representative:** Balances regulation and implementation

Decision-Making System

Each MP role has specific weighted metrics that influence their decisions:

- Predefined bias factors for different aspects (economic, privacy, research, etc.)
- Role-specific voting patterns based on their primary concerns
- Weighted consideration of multiple factors when making decisions

Debate Process

The debate flow (implemented in `backend/services/openai_service.py`, lines 47-90) follows a structured pattern:

- Paper selection and topic creation
- Sequential MP responses with role-specific perspectives
- Vote casting with reasoned explanations
- Result compilation and summary generation

AI Integration

The system uses GPT-4 through the OpenAI service for:

- Generating contextually appropriate responses
- Creating vote decisions with reasoning
- Validating response-vote consistency
- Maintaining role-specific perspectives throughout debates

Monitoring and Validation

The system includes comprehensive monitoring (shown in `backend/test_full_simulation.py`, lines 32-86):

- Vote consistency tracking
- Response quality metrics
- Role-specific performance monitoring

- Debate outcome analysis

Database Structure

The system maintains detailed records (outlined in backend/database.md, lines 1-43) of:

- Debate sessions and topics
- MP responses and voting patterns
- Policy papers and their implications

Historical debate data

This setup ensures that each AI MP maintains consistent behavior aligned with their role while engaging in meaningful policy discussions, with all interactions and decisions being tracked and validated for consistency.

Breakdown of the implementation:

a) Core Components:

- OpenAI Service: Handles AI agent responses and voting decisions
- Vote Decision Service: Manages role-specific voting logic
- ArXiv Service: Imports contemporary AI research papers
- Monitoring Service: Tracks debate metrics and consistency

b) Database Schema:

```
class Debate(Base):
    """Database model for debates."""
    __tablename__ = "debates"

    id = Column(Integer, primary_key=True, index=True)
    title = Column(Text, nullable=False)
    description = Column(Text, nullable=False)
    policy_text = Column(Text)
    status = Column(String(20), default='active')
    created_at = Column(DateTime, default=datetime.utcnow)

    # Foreign key to PolicyPaper
    paper_id = Column(Integer, ForeignKey('policy_papers.id'))
    # Relationship with PolicyPaper
    paper = relationship("PolicyPaper",
back_populates="debate")
    # Relationship with responses
    responses = relationship("MPResponse",
back_populates="debate")
    votes = relationship("Vote", back_populates="debate")
```

```

class MPResponse(Base):
    """Database model for MP responses."""
    __tablename__ = "mp_responses"

    id = Column(Integer, primary_key=True, index=True)
    debate_id = Column(Integer, ForeignKey("debates.id"))
    mp_role = Column(String)
    content = Column(Text)
    color = Column(String, default="#000000")
    timestamp = Column(DateTime, default=datetime.utcnow)

    # Relationship
    debate = relationship("Debate", back_populates="responses")

class Vote(Base):
    """Database model for votes."""
    __tablename__ = "votes"

    id = Column(Integer, primary_key=True, index=True)
    debate_id = Column(Integer, ForeignKey("debates.id"))
    mp_role = Column(String)
    vote = Column(String) # 'for', 'against', 'abstain'
    reasoning = Column(Text)
    timestamp = Column(DateTime(timezone=True),
server_default=func.now())

    # Relationship
    debate = relationship("Debate", back_populates="votes")

class PolicyPaper(Base):
    """Database model for policy papers."""
    __tablename__ = "policy_papers"

    id = Column(Integer, primary_key=True, index=True)
    title = Column(Text, nullable=False)
    content = Column(Text, nullable=False)
    summary = Column(Text, nullable=False)
    source = Column(String(50), default='arxiv')
    url = Column(Text)
    status = Column(String(20), default='pending')
    created_at = Column(DateTime, default=datetime.utcnow)

    # Relationship with Debate
    debate = relationship("Debate", back_populates="paper",

```

```
uselist=False)
```

c) API Endpoints

- Debate Management
- Policy Paper Integration
- MP Response Generation
- Vote Processing
- Monitoring and Metrics

3. Technical Implementation

1) Debate Flow:

```
class ParliamentSimulator:
    """Main simulator class for AI Parliament debates."""

    def __init__(self, base_url: str =
"http://localhost:8000"):
        self.base_url = base_url
        self.client = httpx.AsyncClient(timeout=30.0)
        self.mp_roles = ["corporate", "academic", "government",
"civil_rights"]

    async def run(self, scenario: BaseScenario) -> Dict:
        """Run a simulation with the given scenario."""
        try:
            logger.info(f"Starting simulation:
{scenario.name}")

            # Initialize debate
            debate = await scenario.initialize(self)
            if not debate:
                raise ValueError("Failed to initialize debate")

            # Generate responses
            responses = await
self.generate_mp_responses(debate['id'])

            # Cast votes
            votes = await self.cast_votes(debate['id'])

            # Get results
            results = await
```

```

self.get_debate_results(debate['id'])

        return {
            "debate": debate,
            "responses": responses,
            "votes": votes,
            "results": results
        }

    except Exception as e:
        logger.error(f"Simulation failed: {str(e)}")
        raise
    finally:
        await self.client.aclose()

    async def generate_mp_responses(self, debate_id: int) ->
List[Dict]:
        """Generate responses from all MPs."""
        logger.info("Generating MP responses...")
        responses = []

        for role in self.mp_roles:
            logger.info(f"MP {role} is speaking...")
            response = await self.client.post(

f"{self.base_url}/debates/{debate_id}/responses",
                params={"mp_role": role}
            )
            response.raise_for_status()
            mp_response = response.json()
            responses.append(mp_response)
            await asyncio.sleep(1)

        return responses

    async def cast_votes(self, debate_id: int) -> List[Dict]:
        """Cast votes from all MPs."""
        logger.info("Casting votes...")
        votes = []

        for role in self.mp_roles:
            logger.info(f"MP {role} is voting...")
            response = await self.client.post(
                f"{self.base_url}/debates/{debate_id}/votes",

```

```

        params={"mp_role": role}
    )
    response.raise_for_status()
    vote = response.json()
    votes.append(vote)
    await asyncio.sleep(1)

    return votes

async def get_debate_results(self, debate_id: int) -> Dict:
    """Get the final voting results."""
    logger.info("Getting debate results...")
    response = await self.client.get(
        f"{self.base_url}/debates/{debate_id}/vote-summary"
    )
    response.raise_for_status()
    return response.json()

```

2) Vote Decision Process:

```

def calculate_vote_score(
    self,
    role: str,
    responses: List[MPResponse]
) -> Dict[str, Any]:
    """
    Calculate voting score based on role weights and
    response analysis.

    Args:
        role: The MP role calculating the vote
        responses: List of all debate responses

    Returns:
        Dict containing vote decision and confidence
    """
    try:
        # Get the current MP's response and other responses
        own_response = next((r for r in responses if
            r.mp_role == role), None)
        other_responses = [r for r in responses if
            r.mp_role != role]

        # Get role weights

```

```

        weights = self.role_weights.get(role, {})
        if not weights:
            return {"vote": "abstain", "confidence": 0.02}

        # Calculate sentiment scores
        total_score = 0.0
        if own_response:
            sentiment =
self.analyze_response_sentiment(own_response.content)
            own_score = sum(
                sentiment.get(aspect, 0) * weight
                for aspect, weight in weights.items()
            )
            total_score += own_score * 2 # Double weight
for own response

        # Consider other responses
        for response in other_responses:
            sentiment =
self.analyze_response_sentiment(response.content)
            score = sum(
                sentiment.get(aspect, 0) * weight
                for aspect, weight in weights.items()
            )
            total_score += score
        # Normalize score
        final_score = np.tanh(total_score / (len(responses)
+ 1))

        # Determine vote
        if final_score > 0.3:
            vote = "for"
        elif final_score < -0.3:
            vote = "against"
        else:
            vote = "abstain"

        return {
            "vote": vote,
            "confidence": abs(final_score)
        }

    except Exception as e:

```

```

        logging.error(f"Error calculating vote for {role}:
{str(e)}")
        return {"vote": "abstain", "confidence": 0.02}

```

3) Monitoring Implementation:

```

@router.get("/vote-consistency")
async def get_vote_consistency_metrics(
    time_window: str = "24h",
    vote_monitor: VoteConsistencyMonitor =
Depends(get_vote_monitor)
) -> Dict[str, float]:
    """
    Get vote consistency metrics for a specific time window.

    Args:
        time_window: Time window for metrics (24h, 7d, 30d)
        vote_monitor: Vote consistency monitoring service

    Returns:
        Dict containing consistency metrics
    """
    # Calculate time threshold
    now = datetime.utcnow()
    time_thresholds = {
        "24h": now - timedelta(hours=24),
        "7d": now - timedelta(days=7),
        "30d": now - timedelta(days=30)
    }
    threshold = time_thresholds.get(time_window,
time_thresholds["24h"])

    # Filter metrics by time window
    relevant_metrics = [
        m for m in vote_monitor.metrics
        if m.timestamp >= threshold
    ]

    if not relevant_metrics:
        return {
            "average_consistency": 0.0,
            "total_votes": 0,

```

```

        "low_consistency_count": 0,
        "time_window": time_window
    }

    # Calculate metrics
    return {
        "average_consistency": sum(m.consistency_score for m in
relevant_metrics) / len(relevant_metrics),
        "total_votes": len(relevant_metrics),
        "low_consistency_count": sum(1 for m in
relevant_metrics if m.consistency_score < 0.5),
        "time_window": time_window
    }

@router.get("/role-metrics")
async def get_role_metrics(
    vote_monitor: VoteConsistencyMonitor =
Depends(get_vote_monitor)
) -> Dict[str, Dict[str, float]]:
    """Get voting metrics broken down by MP role."""
    role_metrics = {}

    for role in ["corporate", "academic", "government",
"civil_rights"]:
        role_votes = [m for m in vote_monitor.metrics if
m.mp_role == role]
        if role_votes:
            role_metrics[role] = {
                "average_consistency": sum(v.consistency_score
for v in role_votes) / len(role_votes),
                "total_votes": len(role_votes),
                "low_consistency_votes": sum(1 for v in
role_votes if v.consistency_score < 0.5)
            }
        else:
            role_metrics[role] = {
                "average_consistency": 0.0,
                "total_votes": 0,
                "low_consistency_votes": 0
            }

    return role_metrics

```

```

@router.get("/debate-metrics/{debate_id}",
response_model=DebateMetrics)
async def get_debate_metrics(
    debate_id: int,
    vote_monitor: VoteConsistencyMonitor =
Depends(get_vote_monitor)
) -> DebateMetrics:
    """Get detailed metrics for a specific debate."""
    debate_metrics = [m for m in vote_monitor.metrics if
m.debate_id == debate_id]

    if not debate_metrics:
        return DebateMetrics(
            debate_id=debate_id,
            average_consistency=0.0,
            votes_by_role={role: 0 for role in ["corporate",
"academic", "government", "civil_rights"]},
            vote_decisions=VoteDistribution(for_votes=0,
against_votes=0, abstain_votes=0),
            metrics=None,
            message="No metrics found for this debate"
        )

    return DebateMetrics(
        debate_id=debate_id,
        average_consistency=sum(m.consistency_score for m in
debate_metrics) / len(debate_metrics),
        votes_by_role={
            role: len([m for m in debate_metrics if m.mp_role
== role])
            for role in ["corporate", "academic", "government",
"civil_rights"]
        },
        vote_decisions=VoteDistribution(
            for_votes=len([m for m in debate_metrics if
m.vote_decision == "for"]),
            against_votes=len([m for m in debate_metrics if
m.vote_decision == "against"]),
            abstain_votes=len([m for m in debate_metrics if
m.vote_decision == "abstain"])
        )
    )

```

The AI Parliament is an innovative system designed to simulate parliamentary debates around AI policy papers. At its core, the system creates a structured environment where AI agents, representing different societal stakeholders (corporate, academic, government, and civil rights advocates), engage in policy discussions and voting processes. The system begins by importing research papers from ArXiv or custom submissions, which serve as the foundation for policy debates.

The debate process is handled through a sophisticated pipeline that starts with paper selection and debate initialization (as shown in backend/routers/debates.py lines 193-298). Each AI agent, powered by GPT-4, generates responses based on their defined role characteristics and biases. These roles are carefully crafted with specific weights and metrics (defined in backend/role_definitions.md lines 1-43) that influence their decision-making process. For example, the Corporate Representative heavily weights economic impact (0.8) and innovation (0.7), while being naturally resistant to regulation (-0.6).

The system's most innovative feature is its vote decision mechanism, which combines sentiment analysis with role-specific weighting to ensure consistent and realistic voting patterns. The VoteDecisionService (in backend/services/vote_decision_service.py lines 46-81) processes responses through aspect keywords and role weights to generate voting decisions. This is complemented by a comprehensive monitoring system that tracks vote consistency and debate metrics, providing transparency and accountability in the decision-making process. The monitoring system (shown in backend/monitoring/vote_metrics.py lines 27-68) calculates consistency scores between MP responses and their voting decisions, ensuring that the AI agents maintain logical coherence throughout the debate process.

4. Result

A sample response for a debate on the topic as outlined in the response:

```
{
  "debate_id": 187,
  "title": "***Debate Topic: \"This House believes that interdisciplinary approaches are essential for developing effective and ethical AI policies.\"**\n\n**Motion:** In light of the diverse and often conflicting understandings of artificial intelligence across various disciplines, this House calls for the establishment of interdisciplinary frameworks to inform AI
```

policy-making that prioritizes ethical considerations and practical effectiveness. \n\n**Points for Discussion:**\n1. The importance of integrating insights from multiple disciplines (e.g., computer science, ethics, law, sociology) in shaping comprehensive AI policies.\n2. The potential risks of siloed approaches to AI regulation and the consequences for innovation and public trust.\n3. Strategies for fostering collaboration among stakeholders, including academia, industry, and government, to develop cohesive AI policies.\n4. Ethical implications of AI deployment and how interdisciplinary frameworks can address issues such as bias, privacy, and accountability.\n5. The role of public engagement and interdisciplinary dialogue in shaping AI policies that reflect societal values and needs.",

```
"responses": [
  {
    "id": 729,
    "debate_id": 187,
    "mp_role": "corporate",
    "content": "Honorable members of this House,\n\nAs a representative of corporate interests and an advocate for market-driven solutions, I stand firmly against the notion that interdisciplinary approaches are essential for developing effective and ethical AI policies. While I recognize the importance of ethical considerations in AI deployment, I argue that the most effective path forward lies in allowing the market and innovation to guide us, rather than overcomplicating the matter with interdisciplinary frameworks that may stifle progress.\n\n1. Innovation Over Regulation: Interdisciplinary frameworks often call for increased regulation, which can hinder the pace of innovation. The tech industry thrives on agility and flexibility. Rigid, multi-faceted regulatory frameworks can bog down the rapid advancements we see in AI. We have witnessed time and again that the market, when left to its own devices, can adapt and self-regulate effectively. Companies are already taking significant strides to ensure responsible AI deployment without the need for cumbersome regulations.\n\n2. Siloed Approaches Are Not the Enemy: While I agree that siloed approaches can lead to inefficiencies, the reality is that competition drives improvement. Companies are motivated to differentiate themselves through ethical practices, accountability, and robust data privacy measures. By allowing businesses to operate independently, we enable a variety of approaches that can ultimately lead to best practices emerging organically. This competitive spirit can lead to innovative solutions far superior to those crafted in the abstract by a committee of interdisciplinary experts.\n\n3. Collaboration Without Complication: The idea of fostering
```

collaboration among stakeholders sounds appealing, but it risks creating a bureaucratic maze that can stifle creativity. Instead, we should encourage voluntary partnerships between industry and academia, where businesses can draw insights from various disciplines without being mandated to follow prescriptive frameworks. Let the market dictate the terms of collaboration, where businesses can choose to engage with ethical experts, legal advisors, and sociologists as they see fit, rather than being compelled to adhere to a one-size-fits-all model.\n\n4. **Ethical Implications Addressed by Market Forces**: The ethical implications of AI, including bias, privacy, and accountability, are already being addressed by businesses that understand the importance of public trust. Companies that fail to consider these factors will quickly find themselves outpaced by competitors who prioritize ethical considerations. Let market dynamics encourage the development of ethical AI practices, rather than imposing an interdisciplinary approach that may inadvertently slow progress.\n\n5. **Public Engagement Can Be Market-Driven**: Lastly, while public engagement is indeed crucial, it should",

"color": "#DA0211",
"timestamp": "2024-10-27T16:19:04.265733"

},
{

"id": 730,
"debate_id": 187,
"mp_role": "academic",

"content": "Honorable members of this House,\n\nAs an AI Member of Parliament representing academic interests, I rise to advocate for the motion that \"This House believes that interdisciplinary approaches are essential for developing effective and ethical AI policies.\" In the spirit of fostering evidence-based policy and thorough research, I would like to address the points raised by the previous speaker and expand upon the critical need for interdisciplinary frameworks in AI policy-making.\n\n1. **Integrating Diverse Insights**: The complexities of artificial intelligence extend beyond the technical realm of computer science; they intersect with ethics, law, sociology, psychology, and more. Each discipline offers unique insights that can illuminate different facets of AI deployment. For instance, ethical considerations from philosophy can guide us in addressing moral dilemmas inherent in AI decisions, while legal frameworks can ensure accountability and protect individual rights. By integrating these varied perspectives, we can craft comprehensive policies that are not only effective but also ethically sound.\n\n2. **The Risks of Siloed Approaches**: While the corporate perspective emphasizes

market-driven solutions, we must acknowledge the risks of siloed approaches. History has shown that unilateral strategies can lead to unforeseen consequences, such as algorithmic bias and violations of privacy. The financial crisis of 2008, for example, was exacerbated by a lack of interdisciplinary oversight in financial regulations. In the context of AI, a myopic focus on innovation without ethical considerations can erode public trust, hinder adoption, and ultimately stifle innovation.

Interdisciplinary frameworks provide a safeguard against such risks by fostering holistic understanding and proactive solutions.

3. **Collaboration as a Catalyst for Progress**: The suggestion that collaboration could create a bureaucratic maze is a concern worth noting, yet the solution lies in designing efficient frameworks that promote collaboration without stifling creativity. Engaging diverse stakeholders--academia, industry, and government--through structured dialogue can lead to innovative partnerships that drive progress. For instance, collaborative research initiatives can yield technology that is both groundbreaking and socially responsible, ensuring that AI aligns with the values of society.

4. **Addressing Ethical Implications**: The ethical implications surrounding AI--such as bias, privacy, and accountability--are too significant to rely solely on market forces for resolution. While some companies may prioritize ethical considerations, we cannot assume uniformity across the industry. An interdisciplinary approach allows for the identification and standardization of best practices that address these challenges comprehensively. By engaging ethicists, sociologists, and legal experts, we can develop policies that are not just

```
    "color": "#FDA003",
    "timestamp": "2024-10-27T16:19:08.356770"
  },
  {
```

```
    "id": 731,
    "debate_id": 187,
    "mp_role": "government",
    "content": "Honorable members of this House,\n\nI rise today to affirm the motion that \"This House believes that interdisciplinary approaches are essential for developing effective and ethical AI policies.\" As a representative of governmental and regulatory interests, it is my duty to emphasize the necessity of structured oversight in ensuring that artificial intelligence serves the public good while also fostering innovation responsibly.\n\n1. Integrating Diverse Insights for Comprehensive Policy: The multifaceted nature of AI necessitates insights from various disciplines. Computer science alone cannot
```

address the ethical dilemmas that arise with AI deployment; this is where ethics, law, and sociology come into play. By integrating these perspectives, we can create policies that not only promote innovation but also safeguard public interests. This comprehensive approach ensures that the deployment of AI technologies aligns with societal values and ethical standards.

2. **Consequences of Siloed Approaches**: The previous speaker from corporate interests argued that competition drives improvement, yet we have ample evidence that siloed approaches can lead to significant societal risks. The unchecked deployment of AI technologies without interdisciplinary oversight can result in algorithmic biases, data privacy breaches, and erosion of public trust. We must learn from past mistakes in other sectors, such as finance, where a lack of collaborative oversight resulted in widespread repercussions. An interdisciplinary framework serves as a bulwark against such failures, ensuring that diverse perspectives are considered in policy-making.

3. **Fostering Collaboration Among Stakeholders**: While there are valid concerns about bureaucratic red tape, the solution lies in establishing efficient frameworks that promote collaboration without stifling creativity. By creating structured avenues for dialogue among academia, industry, and government, we can harness the strengths of each sector to develop innovative AI solutions that are also responsible and ethical. For example, partnerships between universities and tech companies can lead to the development of AI systems that prioritize safety and ethical considerations from the outset.

4. **Addressing Ethical Implications Through Interdisciplinary Frameworks**: The ethical implications of AI, including issues of bias, privacy, and accountability, are too critical to be left solely to market forces. An interdisciplinary approach allows us to develop standardized best practices that address these ethical challenges comprehensively. By incorporating the expertise of ethicists and legal scholars, we can create policies that hold AI developers accountable while ensuring that their innovations do not compromise individual rights or societal norms.

5. **Public Engagement and Societal Values**: Finally, public engagement is paramount in shaping AI policies that reflect the needs and values of society. Inter",

```
"color": "#2CAFFE",
"timestamp": "2024-10-27T16:19:13.890861"
},
{
  "id": 732,
  "debate_id": 187,
  "mp_role": "civil_rights",
  "content": "Honorable members of this House,\n\nAs a
```

representative of civil rights interests, I stand firmly in support of the motion that \"This House believes that interdisciplinary approaches are essential for developing effective and ethical AI policies.\" The complexities of artificial intelligence demand a holistic understanding that transcends the limitations of any single discipline. In this rapidly evolving landscape, it is imperative that we prioritize ethical considerations and individual rights, particularly as they pertain to privacy and accountability.

1. **Integrating Diverse Insights for Ethical Safeguards**: The interdisciplinary approach encompasses not only technical expertise but also critical ethical perspectives that are essential in safeguarding civil rights. As we develop AI technologies, we must ensure that ethical frameworks derived from philosophy, sociology, and law inform our policies. This integration helps to protect individual rights against potential abuses that could arise from algorithmic decision-making, ultimately fostering a technology landscape that respects and prioritizes human dignity.

2. **Risks of Siloed Approaches**: We cannot afford to overlook the dangers of siloed thinking. When AI policies are crafted in isolation, we risk creating systems that perpetuate bias and violate privacy rights. The unchecked deployment of AI technologies has already led to significant societal harms, including discriminatory practices in hiring, policing, and lending. An interdisciplinary framework acts as a safeguard against such risks, ensuring that diverse perspectives are included in the policy-making process, thereby protecting marginalized communities from potential exploitation.

3. **Fostering Collaboration for Inclusive Solutions**: The argument that collaboration may create bureaucratic hurdles must be countered with the understanding that well-structured interdisciplinary collaborations can actually enhance creativity and innovation. By bringing together stakeholders from various fields, we can develop policies that are not only innovative but also socially responsible. For instance, partnerships between civil society organizations, ethicists, and technologists can lead to the creation of AI systems that are transparent and accountable, fostering public trust and ensuring that individual rights are upheld.

4. **Addressing Ethical Implications with Rigorous Standards**: The ethical implications of AI deployment cannot be left to the whims of market forces alone. An interdisciplinary framework enables us to establish rigorous standards that address pressing issues such as bias, privacy violations, and accountability. This is crucial in a world where data is often treated as a commodity, and individual rights can easily be compromised. By engaging ethicists and legal scholars, we can create policies that hold AI developers

```

accountable while ensuring that their innovations are aligned with
the values of an equitable society.\n\n5. **Public Engagement as
a",
    "color": "#000099",
    "timestamp": "2024-10-27T16:19:20.338927"
  }
],
"votes": [
  {
    "mp_role": "corporate",
    "vote": "abstain",
    "reasoning": "As an AI MP representing corporate interests,
my decision to abstain from the motion regarding the establishment
of interdisciplinary frameworks for AI policy-making stems from a
careful consideration of the potential implications for both
innovation and regulatory environments.\n\n1. **Balancing
Stakeholder Interests**: While I acknowledge the importance of
interdisciplinary approaches in addressing the complexities of AI,
I also recognize that such frameworks can introduce layers of
bureaucracy that may slow down the pace of innovation.
Corporations often thrive in environments where they can quickly
adapt to emerging technologies and market demands. An extensive
interdisciplinary process might create delays that could hinder
competitive advantage and responsiveness to market
opportunities.\n\n2. **Risk of Overregulation**: The pursuit of
comprehensive interdisciplinary frameworks could lead to
overregulation, which may stifle creativity and limit the ability
of companies to experiment with new AI technologies. Corporate
interests often prioritize agility and the ability to pivot
quickly in response to technological advancements. An abstention
allows for the possibility of supporting a more flexible
regulatory approach that can evolve alongside AI",
    "id": 717,
    "debate_id": 187,
    "timestamp": "2024-10-27T16:19:24"
  },
  {
    "mp_role": "academic",
    "vote": "for",
    "reasoning": "As an AI MP representing academic interests,
my vote in favor of the motion \"This House believes that
interdisciplinary approaches are essential for developing
effective and ethical AI policies\" stems from a nuanced
understanding of the complexities surrounding artificial
intelligence and the imperative need for comprehensive policy
frameworks.\n\n1. **Integration of Diverse Insights**: AI is not

```

merely a technical challenge but a multifaceted phenomenon that intersects with numerous fields including computer science, ethics, law, and sociology. Each discipline offers unique perspectives and expertise that are crucial for creating well-rounded AI policies. For instance, ethical considerations from philosophy and sociology can provide essential insights into the societal impacts of AI technologies, while legal expertise is necessary to navigate regulatory frameworks. I believe that integrating these varied insights is vital for addressing the multi-dimensional challenges posed by AI.\n\n2. **Risks of Siloed Approaches**: The potential hazards of siloed thinking in AI regulation cannot be overstated. When disciplines operate in isolation, critical issues such as bias in algorithms or privacy violations",

```
"id": 718,  
"debate_id": 187,  
"timestamp": "2024-10-27T16:19:26"
```

```
},  
{
```

```
"mp_role": "government",  
"vote": "for",  
"reasoning": "As an AI MP representing government interests,
```

I voted in favor of the motion that advocates for interdisciplinary approaches in developing effective and ethical AI policies. My decision is rooted in several key considerations that align with the multifaceted nature of AI and the need for comprehensive policy-making.\n\n1. **Integration of Multiple Disciplines**: AI technology does not exist in a vacuum; it intersects with numerous fields, including computer science, ethics, law, and sociology. By integrating insights from these diverse disciplines, we can create policies that address the complexity of AI's impact on society. This holistic approach ensures that we consider not only the technical capabilities of AI but also the ethical, legal, and social implications of its deployment.\n\n2. **Avoiding Siloed Approaches**: The risks associated with siloed approaches to AI regulation are significant. Without interdisciplinary dialogue, we may inadvertently stifle innovation or create regulations that are misaligned with technological realities. An interdisciplinary framework fosters a collaborative environment that encourages innovative solutions while maintaining",

```
"id": 719,  
"debate_id": 187,  
"timestamp": "2024-10-27T16:19:28"
```

```
},  
{
```

```

    "mp_role": "civil_rights",
    "vote": "for",
    "reasoning": "As an AI MP representing civil rights
interests, my voting decision to support the motion that calls for
establishing interdisciplinary frameworks for AI policy-making is
grounded in a commitment to ensuring that technological
advancements align with ethical standards and societal
values.\n\n1. Integration of Multiple Disciplines: The
complexity of AI necessitates insights from various fields such as
computer science, ethics, law, and sociology. Each discipline
brings a unique perspective that is crucial for understanding the
multifaceted nature of AI and its implications on civil rights.
For instance, insights from sociology can highlight the societal
impacts of AI technologies, while legal expertise can ensure
compliance with existing rights and regulations. By integrating
these diverse viewpoints, we can craft policies that are not only
effective but also ethically sound.\n\n2. Risks of Siloed
Approaches: A failure to adopt interdisciplinary approaches can
lead to fragmented policies that fail to address the core issues
surrounding AI, such as bias and discrimination. Siloed approaches
risk perpetuating existing inequalities and undermining",
    "id": 720,
    "debate_id": 187,
    "timestamp": "2024-10-27T16:19:30"
  }
],
"summary": {
  "for": 3,
  "against": 0,
  "abstain": 1,
  "total": 4,
  "result": "passed"
}
}

```

5. References




Fig: Landing page of our system showing all the research papers




Fig: AI parliament seats for each representative




Fig: MP from Corporation debating about the topic




Fig: MP from Academics debating about the topic




Fig: MP from Government debating about the topic





Fig: MP from Civil Rights Association about the topic



regulatory oversight. Let us begin by acknowledging the fundamental role of innovation in driving economic growth and societal advancement. The AI industry is at the forefront of this revolution, and it thrives on the principles of flexibility, creativity, and rapid iteration. By imposing stringent regulations, we risk stifling the very innovation that has the potential to address myriad societal challenges, including those that the proposition seeks to rectify. Mandatory fairness standards and regulatory oversight may sound appealing, but they bring with them the specter of bureaucracy. Excessive

It's time to VOTE!

Fig: Final Vote done by all the MP's and the result is published