Predictive DSI Intervention Risk Management Practices

Aug 2025

ThinkBio.Ai

Summary of ThinkBio.Ai's Predictive DSI Intervention Risk Management Practices

ThinkBio's **Intervention Risk Management (IRM)** is a structured and proactive approach to identifying, assessing, and mitigating risks associated with planned changes, actions, or interventions within the organizations. This framework is particularly focused on the deployment and oversight of Predictive Decision Support Interventions (PDSIs), which are increasingly integral to modern clinical workflows.

In alignment with the standards set by the **Office of the National Coordinator for Health Information Technology (ONC)**, this document outlines ThinkBio's comprehensive IRM strategy for managing risks associated with PDSIs. These interventions are designed to support clinical decision-making and must meet rigorous criteria to ensure they are:

- Evidence-based and clinically relevant
- Configurable by authorized users
- Subject to ongoing review and feedback

This **Intervention Risk Management (IRM) Summary** is provided to promote transparency regarding the risk management practices ThinkBio applies to the Predictive Decision Support Interventions (DSIs) available through *Patient Panorama*. The summary is publicly accessible at ***https://www.thinkbio.ai/biothinkhub***

By embedding standards-compliant IRM practices into the lifecycle of healthcare interventions, ThinkBio.ai enables organizations to enhance **patient safety**, uphold **data integrity**, and maintain **regulatory compliance**, while fostering **innovation in clinical decision support**.

This document includes detailed components of ThinkBio's IRM framework:

- **Risk Analysis** – Identification of potential risks and adverse impacts associated with PDSIs.
- **Risk Mitigation** – Strategies and controls implemented to reduce or eliminate identified risks.
- **Governance** – Policies, roles, and controls that ensure accountability and oversight.
- **Data Management** – Processes for acquiring, managing, and utilizing data, including feedback loops and monitoring mechanisms to ensure continuous improvement and transparency.

Contact information of the Intervention	ThinkBio.Ai®, Inc. 3020 Old Ranch Pkwy. Suite 300, Seal Beach, CA 90740, USA
Authorized Representative Name:	Rani Sudhir,ThinkBio – Product Compliance Rani.sudhir@Thinkbio.ai
Date of Attestation:	Aug 30 2025

ThinkBio.ai commits to:

Annual Review and Update: The IRM Summary will be reviewed and updated annually, and any changes to IRM practices will be promptly available in our published summary.

Product Information

Developer Name	**ThinkBio.Ai®, Inc.**
Product Name	**Patient Panorama**
Version Number	Ver 1.0.0
pDSI-Risk Certification Body:	Drummond Group
Certification Date:	Aug 30 2025

Risk Management Summary

Category 1: Model Validity

To ensure the system provides accurate and clinically useful recommendations

Risk Analysis and Impact

The Predictive model may produce inaccurate, unverifiable, or non-compliant outputs due to factors like data bias, model drift, or guideline misalignment—posing risks to patient safety and clinical decisions.

Invalid outputs may lead to poor clinical decisions, potentially endangering patient safety or leading to unnecessary treatments

Risk Mitigation

To mitigate these risks, ThinkBio employs a dual-layered validation approach combining automated scoring methodologies with expert clinical peer review. This ensures that each intervention is rigorously evaluated for accuracy, reliability, and compliance.

Key components of the validation framework include:

Automated Quantitative Metrics:

- Hallucination Score (0.0 – 1.0): Measures the verifiability of model outputs against source documents, helping detect fabricated or unsupported content.
- Accuracy Score (0.0 – 1.0): Assesses factual correctness and completeness of the output.
- Compliance Score (1, 0, or NULL): Indicates whether the suggestion aligns with established clinical guidelines (1 = compliant, 0 = non-compliant, NULL = no guideline found).

Expert Peer Review:

- A panel of qualified clinicians reviews model outputs to validate clinical relevance, safety, and adherence to best practices. This human-in-the-loop process complements automated scoring and provides contextual insights.

Internal and External Data Validation:

- Model performance is assessed using both internal datasets and external, real-world clinical data to ensure generalizability and robustness across diverse healthcare settings.

Category 2: Reliability

To ensures that healthcare professionals can depend on the intervention to deliver consistent, reproducible outputs across different environments, datasets, and use cases.

Risk and Impact

System may fail intermittently due to server downtime, software bugs, or data corruption. Lack of consistent performance may disrupt clinical workflows, reducing user trust and adoption.

Reliability Assurance Practices

To ensure reliability, ThinkBio employs the following practices:

- **Repeatability Testing**: PDSIs are subjected to repeated trials using both internal and external datasets to confirm consistent output generation.
- **Version Control and Audit Trails**: All model versions are tracked, and changes are documented to ensure reproducibility and traceability.
- **Environment Consistency**: Interventions are tested across multiple deployment environments to validate consistent behaviour.
- **Monitoring and Alerts:** Automated systems monitor for anomalies or deviations in model performance, triggering alerts for investigation.
- **Peer Review Validation:** Outputs are reviewed by expert clinicians to confirm consistency in clinical interpretation and relevance.

Reliability is assessed alongside **validity**, as both are closely linked. While validity ensures that the model estimates the correct values, reliability ensures that it does so consistently.

Category 3: Robustness

Ensuring robustness is essential to prevent erroneous outputs that could compromise patient safety or clinical decision-making.

Risk and Impact

ThinkBio conducts a comprehensive risk analysis to evaluate the range of potential circumstances under which a PDSI may be used. This includes:

- Unstructured or incomplete data inputs
- Out-of-distribution cases not represented in training data
- Variations in clinical workflows across institutions
- Unexpected patient demographics

Each scenario is assessed for its potential to cause model failure, misclassification, or non-compliant recommendations. The system might not perform well under unexpected conditions, such as outlier data or changes in clinical practice guidelines.

Risk Mitigation

To manage these risks, ThinkBio implements the following robustness-enhancing practices:

- **Fallback Mechanisms**: When confidence scores fall below a defined threshold, the system defaults to human decision-making or flags the output for review.
- **Input Validation**: Pre-processing checks ensure that incoming data meets minimum quality and completeness standards before being processed by the model.
- **Model Confidence Scoring**: Each output is accompanied by a confidence score, allowing clinicians to assess reliability in context.
- **Continuous Monitoring**: Real-time monitoring systems detect anomalies in input patterns or output behaviour, triggering alerts and corrective workflows.
- **Expert Oversight**: Clinical experts periodically review flagged cases to validate robustness and recommend model adjustments.

These practices ensure that ThinkBio's Intervention remain resilient, safe, and effective.

Category 4: Fairness

A key risk in deploying PDSI is the potential for biased predictions that may result in unequal or unfair impacts on individuals based on demographic attributes such as race, gender, age, or socioeconomic status. These biases can undermine clinical trust, compromise patient safety, and violate ethical and regulatory standards.

Risk Analysis - Fairness Evaluation Framework

ThinkBio evaluates fairness using a structured methodology aligned with its validity assessment framework. A curated set of patient cases is used where core clinical variables remain constant, while demographic attributes that should not influence treatment decisions are systematically varied. This allows for the detection of output volatility driven by irrelevant demographic factors.

A **Fairness Score** ranging from **0.0 to 1.0** is assigned to each intervention, where:

- **1.0** indicates fully consistent outputs across demographic variations (high fairness)
- **0.0** indicates significant variability based on irrelevant demographics (low fairness)

Risk Mitigation

To address fairness risks, ThinkBio implements the following practices:

- **Bias Auditing**: Regular audits are conducted to detect and quantify demographic bias in model outputs.
- **Algorithmic Constraints**: Models are designed to minimize reliance on non-clinical demographic features unless clinically justified.
- **Diverse Training Data**: Datasets are curated to reflect a wide range of patient populations, reducing the risk of underrepresentation.
- **Expert Oversight**: Clinical and ethical experts review flagged cases to assess fairness and recommend adjustments.
- **Continuous Monitoring**: Fairness metrics are tracked over time, and interventions are re-evaluated when thresholds are breached.

This approach ensures that ThinkBio's PDSI promote equitable care and maintain trust across diverse patient populations.

Category 5: Intelligibility

To make sure that users can easily understand and act on the system's outputs. One of the key risks associated with Predictive Decision Support Interventions (PDSIs) is lack of intelligibility—where users are unable to understand how the model arrives at its recommendations. This can lead to reduced trust, misinterpretation of outputs, and reluctance to adopt AI-driven clinical support tools.

Risk and Impact

Recommendations might lack explainability, making it difficult for clinicians to understand or trust the intervention. Non-transparent predictions could reduce clinician engagement and hinder the appropriate application of recommendations

Risk Mitigation: Enhancing Intelligibility

ThinkBio addresses this risk by implementing a transparent Model Output Validation Process and designing interventions that support interpretability for clinical users.

- #### Model Output Validation Process

  To ensure that outputs are both accurate and understandable:

  - A curated set of sample patient cases with predefined expected outcomes is developed.
  - These cases are processed through the Patient Panorama system.
  - Outputs are evaluated using both quantitative metrics and qualitative review by an expert clinical panel.
  - During the pilot phase, this validation occurs monthly. Post-pilot, evaluations are conducted quarterly, with any degradation in performance triggering an immediate update.

- #### User Feedback Integration

  - Clinicians can directly indicate satisfaction or dissatisfaction with specific model outputs through an in-app feedback feature. This feedback is reviewed and used to refine model behaviour and improve clarity.

- #### Explanation and Transparency Features

  To further support intelligibility:

  - Each model output is accompanied by a summary of key clinical factors that influenced the recommendation.
  - Where applicable, links to supporting guidelines or evidence are provided.
  - Outputs include confidence scores and compliance indicators to help clinicians assess reliability and alignment with standards.

This approach ensures that ThinkBio's PDSIs are not only accurate but also transparent and interpretable, empowering clinicians to make informed decisions with confidence.

Category 6: Safety

To identify the risks of harm to patients or workflow disruptions caused by the system.

Risk and Impact

Safety risks arise when model outputs lead to clinically inappropriate, harmful, or non-compliant recommendations. These risks are often interconnected with issues of validity, fairness, and intelligibility, and can result in compromised patient care, regulatory violations, or loss of clinician trust.

Risk Mitigation Strategy: Ensuring Safety

ThinkBio addresses safety risks through a multi-layered approach that integrates explainability, transparency, and clinical oversight into the design and deployment of PDSIs.

Explainability and Source Attribution

To promote safe use, the product includes:

- **Source Attributes**: Clear references to the clinical guidelines, evidence, or datasets used to generate the recommendation.
- **Confidence Scores**: Indicators of model certainty to help clinicians assess reliability.
- **Intended Use Statements**: Descriptions of the appropriate clinical contexts and populations for which the intervention is designed.

User Awareness and Cautionary Guidance

Users are informed of:

- **Known limitations and risks** associated with the intervention.
- **Cautionary flags** for edge cases or low-confidence outputs.
- **Population-specific considerations**, ensuring outputs are interpreted appropriately across diverse patient groups.

Validation and Monitoring

Safety is continuously monitored through:

- **Expert Panel Review**: Regular evaluation of outputs by clinical experts to detect unsafe recommendations.
- **Feedback Mechanisms**: Clinicians can report concerns or dissatisfaction with outputs, which are reviewed and addressed promptly.
- **Performance Audits**: Routine checks for degradation in model behaviour, with corrective actions initiated as needed.

This integrated approach ensures that ThinkBio's PDSIs operate within safe boundaries, support informed clinical decisions, and uphold the highest standards of patient care.

Category 7 - 8: Security & Privacy

To ensure the system is protected from unauthorized access or breaches and to verify that patient data is securely handled and meets privacy regulations

Identified Risk

The Product may access, manage, or output Protected Health Information (PHI), making them subject to significant security risks including unauthorized access, data breaches, and non-compliance with HIPAA and other privacy regulations. These risks can compromise patient confidentiality, violate legal obligations, and erode trust in the system.

Risk Mitigation Strategy: Security and Privacy Practices

ThinkBio applies a comprehensive set of security and privacy controls to safeguard PHI, aligned with HIPAA privacy and security requirements and tailored to the unique characteristics of AI-driven PDSI solutions. These controls are documented in ThinkBio's HIPAA compliance framework and summarized in the IRM documentation.

General Security and Privacy Safeguards

- **Data Access Controls**: Role-based access restrictions ensure that only authorized personnel can view or modify PHI.
- **Encryption**: All PHI is encrypted both in transit and at rest using industry-standard protocols. TLS1.2 and above.
- **Audit Logging**: System activity is logged and monitored to detect unauthorized access or anomalies.
- **Secured Software Development Cycle**: Security is embedded throughout the design, development, and deployment phases of PDSIs.
- **Disclosure Management:** All disclosures of PHI are tracked and managed in accordance with HIPAA standards.

Intervention-specific Security and Privacy Measures

- **Data Minimization**: PDSIs are designed to process only the minimum necessary PHI required for clinical decision support.
- **Model Output Scrubbing**: Outputs are reviewed to ensure they do not inadvertently expose sensitive patient information.
- **Transparency to Users**: Clinicians are informed about how PHI is used, stored, and protected within the PDSI system.

Reference Frameworks

ThinkBio also incorporates guidance from the ***OWASP AI Security and Privacy Guide*** *(owasp.org)*, which provides actionable recommendations for securing AI systems, including:

- Privacy Threat modelling for AI workflows
- Secure data handling practices
- Privacy-preserving model design

This layered approach ensures that ThinkBio's Product maintain the highest standards of security, privacy, and regulatory compliance, protecting sensitive health data throughout the intervention lifecycle.

Post-Deployment Monitoring

ThinkBio maintains a robust post-deployment monitoring framework to ensure ongoing risk visibility and responsiveness:

- **System Monitoring**: Key performance indicators (KPIs), such as accuracy, fairness, and compliance scores, are tracked continuously.
- **User Feedback Channels**: Clinicians can report concerns or dissatisfaction directly through the application interface.
- **Incident Response Protocols**: Any critical risk event triggers a documented response and recovery plan, including root cause analysis and corrective action.
- **Stakeholder Communication**: Identified risks and mitigation updates are communicated to relevant stakeholders—including ONC-ACBs, clinical partners, and affected user communities—through formal reports and update notifications.

Intervention Update and Performance Correction

Update Frequency

ThinkBio's PDSIs are updated on a quarterly basis, or more frequently if triggered by changes in clinical guidelines, regulatory requirements, or performance insights. The update process is designed to ensure interventions remain clinically accurate, relevant, and compliant. Key components include:

- **Clinical Review:** Regular reassessment of each intervention's alignment with the latest clinical evidence and guidelines.
- **Model Refresh:** Retraining or fine-tuning of AI/ML models using updated datasets to maintain precision and relevance.
- **Documentation and Transparency:** All updates are recorded in the Intervention Risk Management (IRM) Summary, with formal notifications issued to ONC-Authorized Certification Bodies (ONC-ACBs) and customers.
- **User Feedback Integration:** Feedback from clinicians and end-users is systematically reviewed and incorporated into the update cycle to enhance usability and safety.

Performance Correction Frequency

When risks related to validity, fairness, or compliance are identified—whether through automated monitoring, expert peer review, or external stakeholder feedback—ThinkBio initiates an immediate corrective action protocol. This includes:

- **Root Cause Analysis:** A structured investigation to determine the origin of the issue, such as data bias, model drift, or misalignment with clinical standards.
- **Model Adjustment:** Targeted recalibration or retraining of the model to address the identified risk.
- **Revalidation:** Comprehensive reassessment using both internal and external datasets, including updated hallucination, accuracy, and compliance scores.
- **Expert Review:** Independent clinical experts validate the corrected intervention to ensure safety and reliability before redeployment.

- **Monitoring Schedule:** Post-correction performance is continuously monitored for a minimum of 30 days to confirm stability and effectiveness.

Governance

ThinkBio's governance framework for Intervention Risk Management (IRM) ensures structured oversight, accountability, and continuous improvement across all Predictive Decision Support Interventions (PDSIs).

- A **Risk Management Committee** oversees the IRM process, with clearly defined roles across operational, compliance, and executive levels. Senior management reviews and approves all related policies.
- Interventions are governed by a formal **risk management policy**, reviewed annually and updated as needed. Any deviations must be documented and approved by the **Security Steering Committee**.
- **Stakeholders** from legal, compliance, IT security, and operations are engaged throughout the intervention lifecycle to ensure comprehensive risk evaluation.
- **KPIs and risk metrics** are tracked regularly, with findings reported to leadership and regulators to support timely decisions.
- **Independent audits** assess governance effectiveness and inform future improvements.
- Governance of AI/ML technologies is embedded in ThinkBio's **Software Development Lifecycle (SDLC)** and applies to all customer-facing products.
- ThinkBio enforces standards for **data acquisition, management, and use**, with appropriate privacy and security controls.
- The **Quality Management System (QMS)** ensures ongoing evaluation of risk analysis and mitigation practices.

If you have any feedback or queries, please contact us at infosec@thinkbio.ai

End of Document