Name of the person/entity proposing comments	Indian Association of Investment Professionals (CFA Society India)
Category (Listed Entity / MII / Market Intermediary / Professional / Investor / Academician / General Public etc.)	Association of Investment Professionals
Email id	advocacy@iaipm.org
Contact no.	91 8080566484 Shreenivas Kunte, CFA, CIPM

Sr. No.	Proposal	Confirmity to Proposal	Comments	Rationale
1 (A)	Market participants using AI/ML models should have an internal team with adequate skills, expertise and experience to monitor and oversee the performance, controls, testing, efficacy, and security of the algorithms deployed throughout their lifecycle as well as maintain auditability and explain ability/interpretability of AI/ML based models. This shall also include documentation of model development, validation, model versioning and ability to do replay for diagnosis etc.	Strongly Agree	We strongly agree with the need for differentiated treatment between core and ancillary AI tasks, ensuring governance without overburdening non-critical processes:	
			1. Critical Models Require In-House Understanding: For AI/ML models that directly influence investment decisions (e.g., what to buy/sell and how much), internal teams must possess sufficient intelligence and oversight. This ensures accountability and informed decision-making.	1). For AI/ML models that are integral to the investment decision-making process—such as determining asset selection, allocation, or timing—it is imperative that the regulated entity retains in-house expertise and control. This ensures that the investment team can fully understand, validate, and take responsibility for the model's functioning and output, thereby upholding the principles of accountability and sound governance. 2). It is important to distinguish between models that directly impact investment outcomes and those that serve auxiliary functions, such as data processing, formatting, or report automation. Applying a uniform level of oversight to both categories could lead to excessive compliance burden. A differentiated approach allows for stringent oversight where risks are material, while providing flexibility for lower-risk, efficiency-enhancing tools.
			2. Segregation of Critical vs. Ancillary Functions is Sensible: We emphasise differentiating core investment models from those used for ancillary or support tasks (e.g., data extraction, report generation). This helps avoid overregulation of low-risk functions while maintaining rigour where it's most needed. One suggestion is to explicitly state that the bulk of governance, information security, monitoring, and data retention efforts should focus on high impact applications—such as those influencing investment decisions or affecting a large number of investors—rather than on robotic automation or tasks aimed at improving personal efficiency.	2). It is important to distinguish between models that directly impact investment outcomes and those that serve auxiliary functions, such as data processing, formatting, or report automation. Applying a uniform level of oversight to both categories could lead to excessive compliance burden. A differentiated approach allows for stringent oversight where risks are material, while providing flexibility for lower-risk, efficiency-enhancing tools.
			3.**Audit Trail is Essential for Transparency:** Maintaining an audit log of inputs, outputs, and tools used is vital for accountability, especially when AI tools are involved in generating content or analysis for investors.	3). The implementation of AI/ML tools, particularly in investment and client-facing functions, must be accompanied by robust audit trails. Documenting the inputs, model logic, outputs, and the technology used not only ensures traceability and transparency but also facilitates learning from errors and enables proper accountability when discrepancies arise.
1 (B)	The team should implement appropriate risk controls measures and governance frameworks to oversee and challenge the outcomes derived from the AI/ML models (especially during market stress). The team should assess and manage potential risks on a continuous basis to ensure that AI/ML models function in a robust and resilient way. The robustness of AI/ML systems can be reinforced by careful training, and retraining, of ML models with datasets large enough to capture non-linear relationships and tail events in the data.	Partially Agree	We partially agree with the proposal with following key points and observations:	
			Need for Clarity and Specificity: While the intent of the proposal is appreciated, we emphasized that the language is too broad and open-ended. Terms like "non-linear relationships" and "tail events" are not clearly defined, which may lead to subjective interpretation and inconsistent implementation across entities.	While the underlying intent of the proposal—to ensure robust risk control and governance of AI/ML models—is commendable, its current articulation lacks precision. Terms such as "non-linear relationships" and "tail events" are inherently vague and undefined, which can result in subjective interpretations and uneven implementation across regulated entities.
			Events like COVID-19 or sudden geopolitical shocks (e.g., trade tariffs) are not predictable, and expecting AI/ML models to capture such events in training data may be unrealistic. Therefore, expectations around tail-event modeling must be realistic and contextual. Instead of relying fully on AI outputs, entities should define operational thresholds or ranges. (eg., model outputs are trusted only within a certain range). Human oversight and judgment should be layered on top of model outputs to prevent blind reliance, especially during abnormal market conditions.	Given the unpredictable nature of extreme events—such as the COVID-19 pandemic or sudden geopolitical disruptions like trade tariffs—it is unrealistic to assume that AI/ML models can be pre-trained to accurately capture such occurrences. Therefore, expectations around tail-event modeling must be framed in a more practical and contextual manner. A more balanced approach would involve establishing operational thresholds or confidence intervals, beyond which model outputs should not be solely relied upon. In such scenarios, human oversight and judgment must supplement model-driven insights, particularly during periods of market stress or abnormal conditions.
			Risk controls and robustness checks should be tailored to the type of ML algorithm and use case. There's a need for more granular direction on what size, structure, and diversity of datasets are considered sufficient for model robustness. We also suggest that the proposal should be accompanied by guidance notes, templates, or illustrative examples to help regulated entities implement robust risk frameworks in a consistent and auditable manner. There's a need for more granular direction on what size, structure, and diversity of datasets are considered sufficient for model robustness.	Risk control mechanisms and robustness validation should be proportionate to the nature of the algorithm and its application. The proposal would benefit from clearer guidance on appropriate dataset parameters—including minimum data size, structural integrity, and diversity—required for effective model training and retraining. To enable consistent and auditable implementation across the industry, the proposal should be supplemented with illustrative examples, use-case templates, or good-practice notes. This would help entities better interpret expectations and implement controls in line with regulatory intent.
1 (C)	The team should establish procedures for exception and error handling related to AI/ML based systems. The team should also establish back-up/fall back plans in the event an AI based application fails (e.g. due to technical issue or an unexpected disruption) to ensure that the relevant function is carried out through an alternative process.	Strongly Agree	We strongly agree with the proposal w.r.to establish back up/fall back plans in the event an AI based application fails.	
1 (D)	There should be a designated senior management, having appropriate technical knowledge and experience, responsible for the oversight of the model development, validation, ongoing testing, deployment, monitoring and controls of AI/ML based models.	Agree	We agree with the proposal that senior management with appropriate technical knowledge and experience should be responsible for the oversight of model development, validation, ongoing testing, deployment, monitoring, and control of AI/ML-based models. In addition, organisations deploying ML-based solutions should establish an independent model governance and regulation team to ensure robust oversight. Every model refresh or update—including the addition of new variables—must be accompanied by thorough documentation that explains the interpretability and relevance of each variable. Regular data quality control (QC) reports should be generated, including month-on-month (MoM) statistical analysis of model inputs, to monitor shifts in data behaviour. Furthermore, variables that repeatedly trigger alerts and are no longer justifiable should be evaluated for removal to maintain model integrity and prevent noise inflation.	Effective oversight of AI/ML-based models is critical to ensure their reliability, fairness, and alignment with organisational objectives. While assigning responsibility to senior management with the requisite technical expertise is essential, it must be complemented by a dedicated and independent model governance function. Such a structure strengthens accountability, mitigates bias, and reduces operational risk. Additionally, as models evolve—through updates or the introduction of new variables—transparent documentation becomes vital for maintaining interpretability and ensuring stakeholders understand the model's decision logic. Periodic data quality control reports, including month-on-month statistical analyses, help in identifying data drift or anomalies early. Variables that consistently generate false positives or lack clear justification should be reviewed and removed to preserve model performance and avoid unnecessary noise. These measures collectively reinforce model integrity, audit readiness, and investor confidence.
1 (E).	Market participants shall understand their reliance on and manage their relationship with third-party service providers/vendors of AI and ML, including monitoring providers' performance and conducting oversight. Market participants should have a clear service level agreement and contract in place with third-party vendors clarifying the scope of the outsourced functions, performance indicators and clearly determining their rights and remedies for poor performance by vendors. However, AI and ML services provided by third-party vendors are deemed to be provided by the market participants, who shall be responsible for ensuring compliance with all applicable laws, rules and regulations.	Strongly Agree	We strongly agree with the regulator's intent that market participants should remain accountable for the actions and performance of third-party AI/ML service providers. This aligns with existing regulatory principles in India.	We fully support the regulator's intent to hold market participants accountable for the actions and performance of third-party AI/ML providers, in line with existing regulatory frameworks in India. However, to ensure practical implementation and avoid unintended consequences, it's important to differentiate between mission-critical AI use cases—such as investment decision-making—and ancillary or low-risk applications like data summarization or marketing content. Imposing uniform compliance obligations, including strict SLAs, across all AI/ML usage may hinder innovation and limit adoption, particularly in early-stage or incidental applications. A proportionate, risk-based approach that distinguishes formal outsourcing from incidental tool usage is essential. Clear guidance on scope and thresholds will help compliance teams apply appropriate controls without unnecessarily restricting beneficial AI applications. Ultimately, while regulated entities must remain responsible, the regulatory framework should be flexible enough to support innovation, learning, and responsible adoption.
			Need for Segregation of Use Cases: A distinction should be made between mission-critical use cases (e.g., investment decision algorithms) and ancillary uses (e.g., data summarization or sales content generation). The same level of compliance and oversight may not be feasible or necessary for non-core AI applications.	
			Compliance Burden Must Be Proportionate': If all AI/ML use cases—regardless of criticality—are subject to stringent service-level agreements (SLAs), it may hinder innovation or practical adoption of such tools. A graded approach based on risk and criticality is needed. Ensuring SLA-type arrangements with widely used tools (e.g., ChatGPT or Google AI for summarizing data) is not realistic. The regulatory language should reflect what is practically enforceable and distinguish between formal outsourcing and incidental usage. Since the AI/ML regulatory landscape is evolving. The language should allow room for learning and error, especially in early-stage adoption, and not result in over-penalization for non-material lapses.	
			Ambiguity in language could lead internal compliance teams to blanket-ban or restrict AI tools across functions due to fear of non-compliance. Clarification on the scope and thresholds will help teams implement balanced controls. Despite practical concerns, there's clear consensus that the ultimate responsibility for regulatory compliance remains with the regulated entity, even when services are outsourced. This principle is non-negotiable.	
1 (F).	Since AI/ML applications can learn from live data and their model behaviour may hence change after deployment, market participants should conduct periodic reviews and on-going monitoring to ensure that the applications continue to perform as intended. Further, market participants shall share accuracy results of AI/ML models with SEBI on periodic basis.	Agree	There is broad agreement that due to the evolving nature of AI/ML models—especially those that learn from live data—ongoing performance monitoring and periodic reviews are essential to ensure models continue functioning as intended after deployment.	Given the dynamic nature of AI/ML models—especially those that continuously learn from live data—it is essential that market participants conduct ongoing performance monitoring and periodic reviews to ensure models continue operating as intended. However, the term "accuracy" is often too narrow and context-dependent to serve as a universal metric. We recommend using broader, more adaptable terms such as "effectiveness" or "stability," which better capture the real-world performance of models over time and across varying market conditions. Evaluation metrics should be tailored to the specific purpose of each model (e.g., directional correctness for market forecasting rather than absolute precision). Rather than mandating raw accuracy scores, a more practical approach would be to require reporting of performance trends or effectiveness within the regulatory domain, such as during audits. This would support incremental compliance, foster early-stage innovation, and prevent over-penalization for non-material lapses—ensuring that regulatory expectations remain both enforceable and innovation-friendly.
			Suggestion to Use Broader Terms Like "Effectiveness" or "Stability": We recommend replacing "accuracy" with broader, more adaptable terms such as effectiveness, stability, or performance consistency, which better reflect how models behave over time and under different market conditions. Metrics used to evaluate models should be customized based on the model's purpose. For example, a model predicting market trends may focus on directional correctness rather than absolute precision.	
			Rather than raw accuracy scores, we suggest that reporting performance trends or effectiveness within the regulatory domain (e.g., audit submissions), allowing for incremental compliance and feedback over time. Given the inherent subjectivity and evolving nature of AI/ML models, regulatory language should allow flexibility for early-stage adoption and avoid rigid benchmarks that could stifle innovation or penalize models unfairly.	
1 (G).	Market participants should clearly define data governance norms which inter-alia shall include data ownership, access controls, encryption mechanism, rights etc. Any requests for unmasking of data shall be recorded.	Strongly Agree		
1 (H).	AI/ML based systems and its use/test cases shall be subjected to independent auditing (team that has no role in development) mechanisms to ensure transparency and fairness. Audit findings shall be communicated to SEBI to enable proactive monitoring and supervisory oversight.	Strongly Agree		
1 (I).	While devising AI/ML based applications, market participants should provide for users' autonomy and agency in decision-making processes and develop AI models that are sensitive to diverse cultural backgrounds and values.	Strongly Agree		
1 (J).	Market participants should ensure responsible and ethical outcomes in usage of AI/ML against clearly defined rules and practices.	Strongly Agree		
1 (K).	Market participants should retain and adequately secure logs for AI/ML systems with full verbosity so that it is possible to chronologically reconstruct the occurrence of events.	Strongly Agree		
1 (L)	Market participants should have control to switch to manual feedback or auto feedback from time to time basis.	Strongly Agree		
1 (M).	The AI/ML models should operate in a way that complies with existing legal and regulatory obligations.	Strongly Agree		
2 (A).	Market participants using AI/ML models for business operations that may directly impact their customers/clients should disclose the same to the respective customers/clients to foster trust, transparency and accountability. Following is a non-exhaustive list of such operations:	Strongly Agree	We strongly agree with the proposal to disclose the use of AI/ML in client facing interactions and decision-making processes is both appropriate and necessary. Such disclosures are not burdensome but promote transparency, allowing end-users to clearly understand whether they are engaging with a machine or a human, or whether a report was generated autonomously or reviewed by a qualified individual. This clarity fosters trust, ensures informed interpretation of AI-generated content, and aligns with principles of responsible AI deployment. Given the increasing integration of AI in financial services, such transparency is essential to maintain accountability and user confidence.	We strongly agree with the proposal to disclose the use of AI/ML in client facing interactions and decision-making processes is both appropriate and necessary. Such disclosures are not burdensome but promote transparency, allowing end-users to clearly understand whether they are engaging with a machine or a human, or whether a report was generated autonomously or reviewed by a qualified individual. This clarity fosters trust, ensures informed interpretation of AI-generated content, and aligns with principles of responsible AI deployment. Given the increasing integration of AI in financial services, such transparency is essential to maintain accountability and user confidence.
	i. Selection of trading algorithms/Algorithmic trading (including High frequency trading)			
	ii. Asset Management/Portfolio Management			
	iii. Advisory and support services			
2 (B).	Further, non-exhaustive list of disclosure of information to investors for usage of AI and ML applications is given below:	Partially Agree	**Concern with Use of the Term "Accuracy" (Points i & iii):** The term "accuracy" is context-dependent and can be defined differently based on the model's purpose. Recommending accuracy disclosures may lead to confusion, especially when no standard metric is specified. Replacing "accuracy" with broader, more suitable terms like effectiveness or stability is preferred.	The proposal's intent to improve transparency through AI/ML-related disclosures is well understood and broadly supported; however, its current form raises practical concerns. Terms like "accuracy" and "data quality" are context-specific and difficult to standardize, making them unsuitable for investor-facing disclosures without clear definitions. Since responsible data practices are already embedded within regulatory expectations, requiring public disclosures of such technical parameters may add limited value while increasing complexity. Instead, disclosures should focus on clearly stating where AI/ML is used and whether any fees are associated, ensuring investors are informed without overwhelming them with technical jargon. A more standardised and focused set of disclosure metrics, aligned with use-case relevance, will ensure comparability and meaningful transparency without stifling innovation.
	i. Product features, purpose, risks involved, limitations and accuracy results of the model.		**Skepticism Over Disclosure of Data Quality (Point ii):** Disclosing data quality (accuracy, completeness, relevance) to investors is seen as redundant and impractical. Adequate data quality is already an operational assumption; poor-quality data would not be knowingly used. This requirement adds limited value since quality checks are inherent to responsible AI/ML usage.	
	ii. Fees/Charges to be levied, if applicable		**Need for Specificity and Standardization in Disclosures:** The proposal is currently too open-ended, leading to subjective interpretation. Suggest including a minimum, standardised set of disclosure metrics to ensure comparability across entities. This would enhance investor understanding and reduce ambiguity or potential misuse.	
	iii. Information about the quality of data that is used to make AI/ML driven decisions including its accuracy, completeness		**Disclosure of AI Use and Charges (Point ii):** Disclosing whether AI/ML has been used and any AI-related fees or charges is seen as fair and necessary. However, it was noted that AI implementation may reduce costs, so disclosures should clarify if fees are justified.	
2 (C).	The language used in the disclosures should be comprehensible to customers/clients. This will help facilitate customers/clients to understand the service and products that are being offered/sold and allow them to make informed decisions.	Strongly Agree		
2 (D).	Investor grievance mechanism for AI/ML systems shall be in line with existing regulatory framework of SEBI	Strongly Agree		
3 (A).	The market participants should adequately test and monitor the AI/ML based models to validate their results on a continuous basis.	Strongly Agree		
3 (B).	The testing should be conducted in an environment that is segregated from the live environment prior to deployment to ensure that AI/ML models behave as expected in stressed and unstressed market conditions.	Strongly Agree		
3 (C).	In addition to the existing methods of testing, market participants should perform shadow testing with live traffic of AI/ML models to ensure quality and performance before deployment in production environment.	Strongly Agree		
3 (D).	Market participants should maintain proper documentation of all the models and store input and output data for at least 5 years. Market participants should also maintain proper documentation explaining the logic of AI/ML models to ensure that the outcomes produced are explainable, traceable and repeatable.	Agree	WE agree with the proposal along with the suggestion that the record keeping should not limit to 5 years and the market participant should not only store data, they should also store the incident of that AI/ML code (positive or negative). They should also save the exact model, logic, variant or version.	
3 (E).	The behaviour of AI/ML model may change in an unforeseen manner as more data is processed over time. Market participants should think beyond the existing testing methods that may be used for traditional algorithms and ensure the AI/ML models are monitored continuously as the algorithms adjust and transform. Therefore, it is not enough for the AI/ML models to be tested thoroughly before deployment; they need to be continuously monitored throughout their deployment to ensure that the model does not behave in inexplicable ways owing to a subtle shift in the operating conditions or due to excessive noise.	Strongly Agree		
4 (A).	AI/ML based models should be fair. Specifically, they should not favour or discriminate one group of clients/customers over another.	Strongly Agree		
4 (B).	The behaviour of AI/ML model may change in an unforeseen manner as more data is processed over time. Market participants should think beyond the existing testing methods that may be used for traditional algorithms and ensure the AI/ML models are monitored continuously as the algorithms adjust and transform. Therefore, it is not enough for the AI/ML models to be tested thoroughly before deployment; they need to be continuously monitored throughout their deployment to ensure that the model does not behave in inexplicable ways owing to a subtle shift in the operating conditions or due to excessive noise.	Strongly Agree		
4 (C).	Market participants should implement appropriate processes and controls to identify and remove biases from data sets. Further, specific training courses to raise awareness amongst their data scientists (and/ or other relevant staff) of potential data biases may be conducted.	Strongly Agree		
5 (A).	Since the AI/ML systems are dependent on collection and processing of data, Market participants should have a clear policy for data security, cyber security and data privacy for the usage of AI/ML based models.	Strongly Agree		
5 (B).	Collection, usage, processing of investors' personal data, security measures etc. should be in compliance with applicable laws.	Strongly Agree		
5 (C).	Information about technical glitches, data breaches shall be communicated to SEBI and other relevant authorities, as applicable in line with existing regulatory and legal framework.	Strongly Agree		