AI USAGE POLICY

SECTION I – FOUNDATION

Art. 1 – Purpose, Scope, and Applicability

(1) This AI Usage Policy governs the deployment and permissible utilization of artificial intelligence systems integrated within the PeopleBotApp platform (the "Platform"). The Policy establishes technical capabilities, operational boundaries, and compliance obligations for all users accessing AI-powered features.

(2) The provisions apply to every person or organization that registers for, accesses, or utilizes AI-assisted functionality, including candidate screening recommendations, content generation tools, data analysis features, and workflow automation. This encompasses Business Customers, Individual Users, and all Authorized Users operating under delegated permissions.

(3) Where interpretation requires clarification beyond this Policy, recourse shall be made to the hierarchical documentation structure in the Terms & Conditions. Documents rank in descending precedence: Terms & Conditions, Data Processing Addendum, AI Usage Policy, Privacy Policy, Service Level Agreement, and Acceptable Use Policy. In conflict, the higher-tier document prevails unless the lower explicitly modifies a specific clause.

(4) This Policy addresses AI systems deployed exclusively within the Platform and does not extend to third-party applications or external tools that customers may integrate. Governance of third-party AI functionalities remains subject to those vendors' respective policies. Regulatory references in this Policy are provided for general orientation purposes only and apply solely to the extent that the Customer's use of the Platform triggers the applicability of the respective legal regime.

(5) In the event of any inconsistency between this Policy and the applicable contractual terms governing the Customer's use of the Platform, the contractual terms shall prevail. Nothing in this Policy is intended to, or shall be construed as, creating any rights or obligations enforceable by any third party.

Art. 2 – AI System Classification and Operational Framework

(1) The Platform incorporates artificial intelligence from OpenAI (GPT series models) and Anthropic (Claude variants). These systems function as co-pilot and recommendation engines designed to augment, not replace, human decision-making. The AI analyzes inputs, identifies patterns, and generates suggestions to inform recruitment and employment workflows.

(2) This co-pilot model classifies the Platform outside high-risk automated decision-making systems under emerging frameworks, including the EU AI Act. The

architecture preserves Customer autonomy consistent with employment law in the United States and India.

SECTION II – PERMITTED AND PROHIBITED USES

Art. 3 – Approved Use Cases

(1) The Platform's AI functionality supports recruitment and employment administration through specifically defined applications. Customers may utilize AI assistance for resume screening with candidate ranking based on qualifications, job description drafting, and content generation for employment communications, hiring trend analysis, diversity metrics assessment, and administrative workflow recommendations, including interview scheduling suggestions.

Art. 4 – Permissible and Prohibited Use Cases

(1) The AI features are intended exclusively for professional recruitment, talent management, and HR analytics. Any use of the Platform to generate non-professional content, personal advice, or data unrelated to the employment lifecycle is strictly prohibited.

(2) The Platform's AI features shall not be used to execute employment decisions that are solely automated within the meaning of applicable employment and data protection laws. All adverse employment actions must be subject to meaningful human review in accordance with the Human Oversight Framework set out in Article 13.

(3) Users shall not: (a) attempt to reverse-engineer or "de-bias" the underlying models; (b) bypass safety filters or access restricted API parameters; or (c) use the AI Output to train, fine-tune, or "distill" competing AI systems without the Provider's express written consent.

(4) The Platform must not be used to generate or disseminate hateful, sexually explicit, politically biased, or defamatory content. Any attempt to use the AI for mass-surveillance, unauthorized social scoring, or deceptive practices (e.g., creating deepfakes of candidates) constitutes a material breach of this Policy.

(5) The Provider reserves the right to monitor usage patterns for signs of Policy violations. Evidence of prohibited use will result in immediate suspension of AI features and may lead to contract termination for cause without further liability to the Provider.

Art. 5 – Data Input Requirements and Restrictions

(1) Permitted data inputs comprise candidate resumes and curriculum vitae, job descriptions and position specifications, employee records for which the Customer possesses lawful processing authority, and company information relevant to

recruitment requirements. All uploaded data must be accurate, current, and obtained through lawful collection practices with appropriate notices to data subjects.

(2) The Platform prohibits processing of protected health information as defined under the Health Insurance Portability and Accountability Act. The system is not designed for HIPAA compliance, and Customers must not upload medical records, disability-related health data, or protected health information, absent an alternative lawful basis unrelated to AI processing. Social security numbers, financial account credentials, biometric identifiers (facial recognition data, fingerprints, retinal scans), and payment card information represent prohibited data inputs except where specifically required for legitimate employment verification and processed through dedicated secure channels.

(3) Personal data of individuals under eighteen years of age may not be uploaded without verifiable parental or legal guardian consent obtained in compliance with applicable child data protection laws. The Platform is not designed for processing children's information, and Customers should refrain from such data collection except where essential to lawful employment of minors meeting minimum working age requirements in the relevant jurisdiction.

(4) Customers must ensure data quality and completeness for AI processing effectiveness. Outdated information, incomplete candidate profiles, poorly formatted documents, and inaccurate job specifications may degrade AI recommendation quality. The Provider disclaims responsibility for suboptimal AI outputs resulting from deficient input data quality or Customer failure to maintain current information.

Art. 6 – Customer Data Protection and Ownership

(1) The Customer retains complete ownership of all data uploaded to the Platform, including candidate resumes, employee records, job descriptions, company information, and any other content submitted through the account interface. This ownership extends to AI-generated outputs derived from Customer inputs, subject to the intellectual property limitations specified in Article 24 of this Policy regarding copyright eligibility for machine-generated content.

(2) The Provider commits that Customer data shall not be utilized for training, improving, or developing artificial intelligence models under any circumstances. This prohibition applies regardless of whether data would be anonymized, aggregated, or de-identified before training use. Customer data processes exclusively for service delivery purposes and terminates upon conclusion of the Subscription Term as defined in the Terms & Conditions.

(3) All data transmissions between Customer systems and the Platform employ Transport Layer Security protocol (TLS 1.3 or higher) encryption standards. Data at rest within Provider infrastructure maintains encryption using Advanced Encryption Standard with 256-bit keys (AES-256). Access to Customer data is restricted to

authorized personnel operating under strict confidentiality obligations and role-based access controls limiting exposure to the minimum necessary for service provision.

(4) AI processing operates temporarily, facilitating real-time recommendation generation but does not persist in Provider storage systems following session completion. Third-party AI providers (OpenAI, Anthropic) process data transiently under data handling commitments prohibiting retention beyond the duration required for output generation. The Provider maintains no permanent archive of individual AI queries, candidate-specific recommendations, or employment decisions absent explicit Customer instruction through Platform archival features.

Art. 7 – Data Retention and Deletion Procedures

(1) Upon subscription termination or cancellation, the Customer receives a thirty-day retrieval window during which account access permits data export through self-service dashboard functions. Machine-readable export formats include CSV for structured candidate data, JSON for system configurations, and PDF for generated documents. Customers bear sole responsibility for completing data retrieval within this period.

(2) Subject to the mandatory record-keeping obligations defined in Article 21 of this Policy, the Provider shall initiate the permanent deletion of all Customer data from active systems within sixty days following the expiration of the retrieval window. This deletion encompasses production databases and backup archives. However, specific metadata and documentation required for regulatory compliance, bias audits, and evidence of human oversight shall be sequestered and retained for the durations specified in Article 21 (4), after which they shall be subject to final forensic destruction.

Art. 8 – Third-Party AI Provider Security and Sub-Processor Management

(1) The Platform's AI capabilities utilize OpenAI and Anthropic integrations under data handling practices prohibiting the use of customer data for model training. Data transmitted to these providers is limited strictly to content necessary for generating recommendations and processing outputs requested by Customer actions.

(2) The Provider publishes a current sub-processor list accessible through account settings under the Data Processing section, identifying all third-party entities with potential access to Customer data during service delivery. Material changes to the sub-processor roster, including the addition of new AI model providers or infrastructure vendors handling personal data, require thirty days' advance notice to active customers through email notification and Platform announcements.

(3) Enterprise-tier customers retain objection rights permitting service termination without penalty where new sub-processors create unacceptable data protection risks, subject to the procedures established in the Data Processing Addendum. The

Provider conducts periodic security assessments of critical sub-processors and requires contractual commitments to security standards equivalent to those maintained for Provider-operated systems.

(4) Sub-processor agreements incorporate data protection obligations consistent with GDPR Standard Contractual Clauses and India Digital Personal Data Protection Act transfer requirements, where applicable to Customer data jurisdiction.

SECTION III – AI TRANSPARENCY AND LIMITATIONS

Art. 9 – AI Recommendation Generation Methodology

(1) The Platform's artificial intelligence processes user inputs through natural language understanding models that analyze textual content, identify relevant patterns, and generate structured recommendations based on learned associations from training data. For candidate screening applications, the AI evaluates resume content against job description requirements, assessing qualifications through keyword matching, skills alignment, experience relevance, and structural formatting quality.

(2) Factors considered in recommendation generation include demonstrated competencies matching position specifications, years of experience in relevant domains, educational credentials aligned with role requirements, and professional achievements indicating capability. The AI explicitly excludes protected characteristics from processing logic, including race, ethnicity, gender, age, disability status, religion, national origin, genetic information, and other classifications protected under applicable anti-discrimination law.

(3) AI systems exhibit inherent technical limitations. Certain contextual, cultural, linguistic, formatting, or data representation factors may affect the quality or relevance of AI-generated recommendations.

Art. 10 – Nature of AI Outputs and Risk Allocation

(1) Artificial intelligence outputs generated through the Platform are probabilistic in nature and do not constitute statements of fact, guarantees, or professional advice. AI systems operate by identifying statistical patterns and associations and may produce errors, omissions, or outputs lacking factual accuracy (“hallucinations”).

(2) The Provider makes no representations or warranties regarding the accuracy, completeness, originality, legality, or fitness for any particular purpose of AI-generated outputs.

(3) Output quality is materially dependent on the quality, completeness, structure, and accuracy of data inputs provided by the Customer. The Provider bears no responsibility for outcomes resulting from deficient, outdated, incomplete, or misleading inputs.

(4) AI models are subject to temporal and contextual limitations, including training data cut-off dates and evolving regulatory or industry standards. Customers remain solely responsible for verifying that AI-generated outputs remain suitable for their intended use and compliant with applicable law.

Art. 11 – Explainability Features and Transparency Tools

(1) The Platform provides recommendation rationale displays indicating primary factors contributing to AI-generated candidate rankings, job description suggestions, or workflow recommendations, where technically feasible. Confidence scores accompany outputs to communicate the AI's assessed reliability level, categorized as high confidence (strong pattern match with training data), medium confidence (partial alignment with typical patterns), or low confidence (limited basis for recommendation requiring additional human scrutiny).

(2) Customers may request general explanations for specific AI outputs through support channels, receiving information about input factors considered, matching logic applied, and limitations affecting the recommendation. Proprietary algorithm details, model weights, training data specifics, and other trade secrets remain confidential and are not disclosed in explainability responses.

(3) Enterprise-tier customers can request advanced explainability reports providing detailed technical analysis of AI decision pathways for specific use cases, subject to execution of appropriate confidentiality agreements and potential additional fees based on complexity. Such reports may include factor importance rankings, alternative outcome scenarios, and sensitivity analysis showing how input variations affect recommendations.

SECTION IV – BIAS PREVENTION AND FAIRNESS

Art. 12 – Provider Bias Mitigation Practices

(1) The Provider implements rigorous bias mitigation through diverse data selection and fairness testing. However, Customer acknowledges that AI systems are probabilistic and may reflect historical biases or produce "hallucinations." To manage this, the Provider supplies confidence scores and explainability tools to assist in interpreting AI recommendations.

(2) As the ultimate decision-maker, the Customer is solely responsible for conducting its own "Bias Audits" (e.g., NYC Local Law 144) and ensuring that AI outputs align with their internal DEI (Diversity, Equity, and Inclusion) policies and local employment laws.

(3) Should the Customer detect patterns suggesting algorithmic bias, they must immediately notify the Provider. Both parties agree to cooperate in investigating such claims and, if necessary, adjusting system configurations or retraining specific parameters to mitigate identified risks.

Art. 13 – Comprehensive Human Oversight Framework

(1) The Platform's artificial intelligence functions exclusively as a "co-pilot" support tool designed to augment human analysis. All AI-generated rankings, summaries, and assessments are strictly advisory in nature and shall not constitute final, legally binding, or operationally self-executing determinations.

(2) The Customer is strictly prohibited from making final adverse employment decisions—including, but not limited to, rejection of candidates, termination, or denial of promotion—based solely on AI-generated outputs. No decision affecting a data subject's legal status or livelihood shall be executed without prior meaningful human intervention.

(3) Human review must be substantive and performed by personnel possessing the requisite expertise in human resources and applicable labor regulations. Reviewers are mandated to independently evaluate qualitative factors that elude algorithmic processing, such as organizational culture alignment, interpersonal dynamics, and nuanced professional achievements. Mere "rubber-stamping" or perfunctory approval of AI recommendations shall be deemed a material breach of this Policy.

(4) The Customer maintains absolute discretionary authority to override, modify, or disregard any AI-generated recommendation. The Provider shall not impose any technical or operational penalties where human decisions consistently diverge from AI rankings, recognizing that such divergence is a legitimate exercise of professional judgment. Interview decisions remain entirely within human discretion. AI cannot compel the Customer to interview only top-ranked candidates or exclude individuals receiving lower algorithmic scores.

(5) AI outputs categorized as "low confidence" or those involving candidates with non-traditional backgrounds and irregular data profiles shall automatically trigger a requirement for heightened scrutiny. In such instances, the review must be escalated to senior management to mitigate the risk of algorithmic misinterpretation or disparate impact.

Art. 14 – Allocation of Legal Responsibility

(1) The Customer bears sole legal responsibility for all employment decisions and outcomes informed, assisted, or supported by AI-generated outputs, including compliance with applicable employment, anti-discrimination, and data protection laws.

(2) The Provider does not act as an employer, co-employer, agent, or decision-maker and assumes no responsibility for the Customer's hiring, termination, promotion, or other employment-related actions.

(3) The Customer remains responsible for implementing appropriate internal procedures, human oversight mechanisms, and documentation sufficient to demonstrate lawful decision-making in regulatory audits, administrative proceedings, or judicial disputes.

(4) Any assistance, tooling, audit logging, or explainability features provided by the Platform are offered solely as compliance-support mechanisms and do not shift or mitigate the Customer's legal responsibility.

SECTION V – CANDIDATE RIGHTS AND TRANSPARENCY

Art. 15 – Candidate Notification Obligations

(1) Where the Customer utilizes AI features of the Platform in connection with the processing of candidate applications or employee information, the Customer shall provide clear and timely notice to affected individuals that artificial intelligence is used to assist in the analysis of their data. Such notice shall be provided at or prior to the point of data collection and shall describe, in general terms, the role of AI in the relevant employment process.

(2) For positions located in New York City, the Customer must provide notification at least ten business days before using AI to screen applications or evaluate candidates. The notice shall be posted prominently on the careers page accessible to job applicants. Job postings for NYC positions must include disclosure that AI technology assists in the screening process. Required notice elements include confirmation that AI is used, description of the data categories analyzed by the system, and specification of job qualifications or requirements the AI considers during evaluation.

(3) New York City customers must conduct independent bias audits of AI tools within one year before use and annually thereafter. Audit summary results become publicly available on the Customer's careers page. The Provider's own bias testing does not satisfy the Customer's independent audit obligation under NY Local Law 144.

(4) Customers processing applications from candidates subject to Equal Employment Opportunity Commission jurisdiction must inform applicants that AI assists in application review. The notification includes contact information where candidates may submit questions about AI use in the hiring process. Candidates with disabilities receive information about requesting reasonable accommodations in the screening process to ensure the AI evaluation does not create barriers to employment opportunities.

(5) For candidates located in India or Indian citizens applying from any location, explicit consent becomes mandatory before AI processing of personal data. The consent notice must explain that AI is used for application screening or evaluation, identify types of data processed, such as resume content and professional qualifications, specify the purpose of processing as candidate matching and ranking, and inform candidates of their right to withdraw consent. Consent requires clear affirmative action - pre-checked boxes or buried terms of service clauses do not satisfy the India DPDP Act requirement for explicit, informed, and freely-given consent.

(6) Where the Customer processes applications from European Union residents or candidates protected under GDPR, notification must occur when applications are submitted. The notice explains the logic involved in automated processing, describes the significance of AI use in the decision-making process, and outlines envisaged consequences for the candidate. Customers inform candidates of the right to object to automated processing per Article 22.

(7) The Provider offers sample notification templates in the account dashboard Resources section. The Platform includes a Candidate Consent Tracking feature for recording when a notification was provided and consent obtained. These tools are provided as courtesy features to assist Customer compliance efforts. They do not constitute legal advice, and the Customer bears sole responsibility for ensuring notifications satisfy applicable law in relevant jurisdictions.

Art. 16 – Right to Human Review and Objection to Automated Processing

(1) Candidates may request human-only review of their application without AI assistance where such rights exist under applicable law. Laws conferring this right include GDPR Article 22 (European Union), India Digital Personal Data Protection Act Section 16, and certain US state regulations. The Customer must provide a clear, accessible mechanism for candidates to submit human review requests. Acceptable mechanisms include a dedicated email address, a web form on the careers page, or a telephone hotline with documented call handling procedures.

(2) The Platform provides a "Human Review Flag" feature enabling Customers to exclude specific candidates from AI processing. Once a candidate record receives the human review flag, the system prevents AI analysis of that individual's application materials. The flagged candidate's profile remains accessible only for manual review by authorized personnel. No AI recommendations, rankings, or assessments are generated for flagged candidates.

(3) Qualified personnel must conduct a thorough human review of flagged candidates without relying on AI-generated insights, recommendations, or preliminary assessments. The review proceeds as if AI functionality did not exist for that particular candidate. Documentation of the human-only review process becomes part of the candidate's record for potential regulatory inquiry or discrimination claim defense.

(4) The Customer must honor valid objections to automated processing and may not penalize candidates for exercising their rights. Rejecting candidates specifically because they requested human review constitutes retaliation prohibited under various employment protection statutes. When candidates validly object to AI processing, the Customer either proceeds with human-only evaluation or withdraws the application from consideration entirely with appropriate notice to the candidate.

(5) If receiving regulatory inquiries, candidate complaints, or legal demands regarding AI disclosures or human review rights, the Customer shall notify legal@peoplebotapp.com within five business days. The Provider may offer technical

assistance in responding to such inquiries, but does not assume legal responsibility for the Customer's compliance with candidate notification obligations.

Art. 17 – Candidate Explanation Rights and Transparency

(1) Candidates possess the right to request an explanation of how AI was used in evaluating their application under certain regulatory frameworks. The Customer must provide a meaningful explanation using the Platform's explainability tools described in Article 11. Explanations may not consist of generic boilerplate language about AI use generally, but must address the specific candidate's application.

(2) Adequate explanations typically include the factors AI considered during evaluation, how the candidate's qualifications and experience matched job requirements as assessed by the system, the AI's confidence level in its recommendation, and whether the human reviewer accepted, modified, or overrode the AI suggestion in making the final decision. The Customer need not disclose proprietary algorithm details, model weights, or training data specifics to satisfy explanation obligations.

(3) Explanation requests require a response within the timeframe mandated by applicable law. GDPR typically requires a response within one month of the request. The Indian DPDP Act establishes similar timelines. US state laws vary, but generally expect a prompt response. The Customer maintains records of explanation requests received and responses provided.

(4) The Provider does not communicate directly with candidates regarding individual applications or employment decisions. Any explanations provided using the Platform’s tools are generated for the Customer’s internal use and may be adapted or supplemented by the Customer prior to communication with the candidate.

SECTION VI – LEGAL COMPLIANCE

Art. 18 – United States Employment Law Compliance

(1) Customers operating in the United States must comply with federal employment statutes, including Title VII of the Civil Rights Act, the Americans with Disabilities Act, Age Discrimination in Employment Act, and the Equal Pay Act. AI use may not facilitate discrimination against individuals based on protected characteristics. Employment decisions informed by AI recommendations remain subject to all applicable anti-discrimination requirements.

(2) The Customer shall adhere to Equal Employment Opportunity Commission guidance regarding AI and automated systems in employment decision-making. This includes conducting adverse impact analyses where legally required. Customers must ensure AI tools do not screen out individuals with disabilities who could perform job functions with reasonable accommodation.

(3) Where AI features perform functions analogous to background checks or consumer reports, the Customer complies with Fair Credit Reporting Act requirements. This includes providing required disclosures to candidates, obtaining written authorization before AI processing, and following adverse action procedures when AI-informed decisions result in candidate rejection.

(4) Customers in New York City comply with Local Law 144 requirements for automated employment decision tools. Required actions include conducting independent bias audits annually, providing candidate notification at least ten days before AI use, and publishing audit summaries on publicly accessible careers pages. California customers comply with the California Consumer Privacy Act and the California Privacy Rights Act, including disclosure obligations for automated decision-making and consumer opt-out mechanisms.

(5) The Customer maintains compliance documentation for applicable statute of limitations periods. Employment discrimination claims typically permit filing within 180 to 300 days of the alleged violation at the federal level. State claim periods vary from one to four years, depending on jurisdiction and claim type.

Art. 19 – India Digital Personal Data Protection Act 2023 Compliance

(1) Customers operating in India or processing personal data of individuals located in India must comply with the Digital Personal Data Protection Act 2023. Explicit, informed, and freely-given consent becomes mandatory before uploading candidate or employee personal data to the Platform for AI processing. Consent requests must be separate from other authorizations and presented in clear, plain language.

(2) Consent notices explain that AI technology processes personal data for employment screening or evaluation purposes. The notice identifies data categories processed, such as educational qualifications and professional experience. Candidates receive information about their right to withdraw consent at any time. Consent withdrawal requires simple procedures accessible through the same mechanism used to provide initial consent.

(3) Data principals possess rights under DPDP Act Sections 11 through 14, including access to personal data processed by AI, correction of inaccuracies in data inputs or processing results, erasure subject to legal retention obligations, and grievance redressal through the designated officer. The Customer provides mechanisms for data principals to exercise these rights and processes requests within statutory timeframes.

(4) Personal data processing adheres to purpose limitation principles under the DPDP Act Section 4. Data collected for specific employment purposes may not be repurposed for unrelated uses without fresh consent. The Customer collects only data necessary for AI-assisted HR functions, avoiding excessive or irrelevant information collection.

(5) Processing personal data of individuals under eighteen years requires verifiable parental or guardian consent per DPDP Act Section 9. The Platform is not designed

for children's data. Customers should refrain from uploading minors information except where essential to lawful employment, meeting minimum working age requirements.

(6) Where Indian candidate data is transferred to US-based AI providers for processing, explicit consent for cross-border transfer becomes necessary under the DPDP Act Section 16. The Customer obtains separate consent specifically authorizing international data transfer and informs candidates that their information will be processed outside India.

Art. 20 – Record Keeping and Documentation Requirements

(1) The Customer shall maintain comprehensive records demonstrating compliance with human oversight obligations for all AI-assisted determinations. Such documentation must identify the specific personnel who performed the review, the precise date of the assessment, and a substantive rationale for the final decision. The records must explicitly reflect whether the reviewer accepted, modified, or entirely overrode the AI-generated recommendation, ensuring a clear chain of accountability.

(2) Where mandated by applicable jurisdiction (including, but not limited to, New York City Local Law 144), the Customer is responsible for conducting and documenting annual independent bias audits of the AI tools. This documentation must encompass the evaluation methodology, the specific data sets analyzed, and any corrective measures implemented to mitigate identified disparate impacts or statistical skews.

(3) The Customer shall maintain a centralized archive of all notices provided to data subjects regarding the use of AI in employment processes. These records must include the date and method of communication, the specific content of the notice provided, and evidence of explicit consent where required by prevailing regulations, such as the Digital Personal Data Protection Act (DPDP Act) of India.

(4) Documentation and associated data must be retained in accordance with the statutory limitation periods of the relevant jurisdiction. For operations within the United States, retention typically spans one to four years, depending on federal and state labor laws. For operations subject to the India DPDP Act, records shall be maintained for a minimum of three years or until the completion of the processing purpose, whichever period is longer.

(5) Automated audit logs and system records generated by the Platform are intended to support internal tracking and compliance workflows. Such records do not, by themselves, constitute legal determinations or conclusive evidence of compliance and should be assessed in conjunction with the Customer’s internal policies, procedures, and documentation.

SECTION VII – INCIDENT REPORTING

Art. 21 – Reporting AI Errors and Issues

(1) Customers report AI-related issues through designated channels based on incident severity. Critical issues, including suspected discriminatory outputs, material AI errors affecting employment decisions, or security concerns related to AI processing require immediate notification to legal@peoplebotapp.com.

(2) Reportable incidents include AI-generated recommendations exhibiting discriminatory patterns or disparate impact on protected groups, factual errors or hallucinations in AI outputs that could affect employment decisions, repeated failures or performance degradation in specific AI features, and candidate complaints regarding AI use in application processing. Security vulnerabilities related to AI data handling and unauthorized access to AI processing functions require immediate escalation.

(3) Incident reports include the date and time when the issue occurred, a description of the unexpected behavior or problematic output observed, and the impact on employment operations or candidate processing. Where possible, without exposing personal data, include relevant input data that triggered the error. Screenshots, error messages, or system logs assist the investigation. Specify whether the incident affected a single candidate or multiple individuals.

(4) The Provider acknowledges critical incident reports within two hours during business hours. Bias-related reports receive acknowledgment within twenty-four hours. General error reports follow standard support response timeframes based on the Customer's subscription tier. Acknowledgment confirms receipt but does not guarantee immediate resolution or acceptance of the Customer's characterization of the incident.

Art. 22 – Provider Incident Response and Remediation

(1) Upon receiving incident reports, the Provider investigates to assess severity, identify root causes, and determine the impact scope. Critical incidents involving confirmed discriminatory outputs or material AI errors may result in immediate temporary disablement of affected features pending investigation completion. The investigation timeline varies based on complexity but typically concludes within five to fourteen business days.

(2) Remediation approaches depend on the nature and severity. Immediate actions include disabling problematic AI features when safety or discrimination risks are confirmed. Short-term measures provide workarounds or manual alternatives to affected functionality. Long-term remediation involves model retraining, algorithm adjustment, additional fairness testing, and enhanced monitoring. The Provider implements fixes, balancing speed with thoroughness to prevent recurrence.

(3) The Provider keeps reporting to Customers informed of investigation progress and expected resolution timelines through email updates. Upon investigation completion, Customers receive summary findings and remediation steps implemented. Where incidents affect multiple customers, the Provider may issue broader notifications through Platform announcements or quarterly transparency reporting.

(4) If the Provider discovers material AI errors or bias affecting Customer data through internal monitoring, proactive notification occurs within seventy-two hours of discovery. Notification describes the issue nature, identifies affected data or outputs by date range and scope, and recommends corrective actions the Customer should consider. Where feasible, the Provider offers reprocessing or output correction to remediate affected recommendations.

(5) Upon receiving Provider incident notification, the Customer reviews potentially affected employment decisions and takes appropriate corrective action if discrimination or a material error occurred. This may include re-evaluating candidates previously rejected based on flawed AI recommendations, documenting the error and remediation in employment records, and notifying affected individuals where legally required or ethically appropriate.

SECTION VIII – INTELLECTUAL PROPERTY

Art. 23 – Ownership of AI-Generated Outputs

(1) The Customer retains complete ownership of all AI-generated content produced through the Platform, including job descriptions, candidate summaries, policy drafts, interview questions, email templates, and any other materials created by AI features in response to Customer inputs. No additional fees, royalties, or usage restrictions apply to the Customer's commercialization or redistribution of AI outputs beyond the subscription fees specified in the Terms & Conditions.

(2) Ownership rights transfer to the Customer immediately upon generation, subject to important limitations regarding intellectual property eligibility under current law. AI-generated content may not qualify for copyright protection in certain jurisdictions. United States Copyright Office guidance indicates that works created by artificial intelligence without sufficient human authorship may fall into the public domain. The Customer acknowledges this uncertainty and assumes all risks regarding the IP status of AI-generated materials.

(3) The Provider makes no warranty that AI outputs are original, unique, or free from resemblance to existing copyrighted works. AI models generate content based on patterns learned from training data, potentially producing outputs that inadvertently mirror existing materials. Before publishing AI-generated content externally, presenting it as proprietary work, or using it for commercial purposes, the Customer should conduct reasonable searches to verify that no substantial similarity exists with third-party intellectual property.

(4) While the Customer owns AI outputs, underlying AI technology, including models, algorithms, training data, and processing logic, remains the exclusive property of the Provider and third-party AI providers. The ownership grant extends solely to the specific content generated, not to any rights in the AI systems themselves or their operational methodologies.

SECTION IX – SYSTEM UPDATES

Art. 24 – AI Model Changes and Version Control

(1) The Provider implements AI model updates to improve performance, enhance accuracy, or expand functionality. Minor updates, including bug fixes, performance optimization,s and incremental improvements, are deployed without advance notice to customers. These routine enhancements maintain service quality and address technical issues discovered through monitoring.

(2) Material updates involving new AI models, significant algorithm changes, feature additions, or removals require thirty days' advance notice to active customers. Notification occurs through email to the registered account address and prominent disclosure within the Platform interface. Material update notices describe expected changes to output quality, processing characteristics, or feature availability.

(3) Emergency updates addressing critical security vulnerabilities, confirmed algorithmic bias or compliance requirements deploy immediately without prior notice. The Provider notifies affected customers of emergency updates within twenty-four hours following implementation, explaining the issue addressed and changes made. Emergency updates receive priority over standard release schedules when customer safety, data security, or legal compliance necessitates immediate action.

(4) The Provider makes no guarantee that AI outputs remain consistent across different model versions. Updates may alter recommendation styles, confidence scoring patterns, or output formatting. Customers relying on specific AI behavior characteristics should test new versions promptly following material update notifications and report compatibility issues through support channels.

Art. 25 – Beta Features and Experimental Functionality

(1) The Provider occasionally releases beta features or experimental AI capabilities to gather customer feedback and assess real-world performance before general availability. Beta functionality appears clearly marked as "Beta," "Experimental," or "Preview" within the Platform interface. Visual indicators distinguish beta features from production-ready capabilities.

(2) Beta features remain disabled by default and require explicit customer activation through account settings. The Platform displays disclosure notices explaining beta status, limitations, and potential risks before permitting activation. Customers must

affirmatively acknowledge understanding that beta features may exhibit reduced reliability or unexpected behavior.

(3) Beta functionality is provided "as-is" without warranties of any kind regarding accuracy, reliability, completeness, or fitness for particular purposes. The Provider makes no commitments regarding beta feature performance, continued availability, or eventual release as production functionality. Beta features may be modified substantially, withdrawn entirely, or discontinued without advance notice.

(4) Beta features are expressly excluded from Service Level Agreement commitments specified in the Terms & Conditions. Outages, errors, poor performance, or unexpected behavior in beta functionality do not constitute SLA breaches and create no entitlement to service credits, refunds, or other remedies. Customers should not rely on beta features for mission-critical employment decisions or time-sensitive recruitment activities.

(5) Use of beta features implies consent to provide feedback regarding functionality, usability, and performance. The Provider may collect usage analytics, error reports and performance metrics from beta feature operation to inform development decisions. Customer feedback submitted regarding beta features grants the Provider perpetual license to utilize suggestions for product improvement without compensation or attribution obligations.

SECTION X – POLICY GOVERNANCE

Art. 26 – Policy Modifications and Amendments

(1) The Provider reserves the right to modify this AI Usage Policy to reflect regulatory developments in artificial intelligence governance, technological enhancements to Platform capabilities, or evolving industry standards for responsible AI deployment in employment contexts. Policy updates may address emerging legal requirements, incorporate lessons learned from incident investigations, or respond to customer feedback regarding clarity and usability.

(2) Material modifications affecting Customer obligations, AI system functionality or compliance requirements receive at least thirty days' advance notice before the effective date. The Provider communicates policy changes through email notification to the registered account address and prominent disclosure within the Platform interface accessible upon login. Updated Policy versions display a clear indication of the revision date and summarize substantive changes in an accompanying change log.

(3) Minor updates correcting typographical errors, clarifying existing provisions without altering substantive requirements, or updating contact information and administrative details may be implemented without advance notice. The current Policy version remains accessible through the account dashboard and the Provider's website, displaying effective date and version number for reference.

(4) Continued utilization of AI features following the effective date of Policy amendments constitutes acceptance of revised terms. Customers who object to material changes may discontinue AI feature usage or terminate their subscription in accordance with cancellation procedures specified in the Terms & Conditions. The Provider maintains archived versions of prior Policy iterations for customer reference upon request.

Art. 27 – Training Resources and User Education

(1) The Provider maintains self-paced training modules accessible through the account dashboard covering AI capabilities and limitations, prohibited uses and compliance requirements, human oversight obligations, bias awareness and mitigation strategies, incident reporting procedures, and candidate rights and notification requirements. Training content updates periodically to reflect Policy changes, regulatory developments, and Platform enhancements.

(2) Educational resources include usage guidelines organized by common employment scenarios such as resume screening, job description creation, and candidate communication. Best practices documentation addresses regulatory compliance topics relevant to the USA and India markets. The Provider publishes regulatory update bulletins when significant AI-related legislation or guidance emerges affecting employment practices.

(3) Enterprise-tier customers receive access to live webinars conducted quarterly, addressing advanced AI topics, regulatory change,s and new feature demonstrations. Webinar recordings become available in the resource library for on-demand viewing. Enterprise accounts may schedule consultations with compliance specialists to discuss jurisdiction-specific requirements or complex use cases.

(4) While training resources are provided to support Customer success, completion of training modules is recommended but not mandatory except where specifically required by enterprise service agreements. The Customer bears ultimate responsibility for ensuring personnel using AI features possess adequate knowledge of Policy requirements and applicable employment law.

Art. 28 – Designated Contact Channels

(1) Customers direct inquiries, support requests, and compliance communications through designated channels based on subject matter. General questions regarding AI functionality, feature operation, or account administration should be routed to support@peoplebotapp.com or through the support ticket system accessible within the Platform interface.

(2) Critical AI incidents, including suspected discriminatory outputs, material processing errors, or security vulnerabilities, require immediate notification to support@peoplebotapp.com. Suspected algorithmic bias, fairness concerns, or pattern discrimination warrant reporting to legal@peoplebotapp.com for investigation by specialized personnel.

(3) Policy interpretation questions, compliance inquiries, or requests for legal documentation should be routed to legal@peoplebotapp.com.

(4) Contact information for designated officers and response timeframes for specific inquiry types are maintained in the Support section of the account dashboard. The Provider updates contact channels as organizational structure evolves and notifies customers of significant changes to reporting procedures or responsible personnel.