Considerations for the Use of Artificial Intelligence in the Creation of Lay Summaries of Clinical Trial Results

I. Introduction

The landscape of medicinal product research and development is continually evolving as new technologies reshape traditional practices. One such technological advancement is the use of Artificial Intelligence (AI) in generating lay summaries (LS) of clinical trial results. Lay summaries are crucial for increasing transparency and ensuring that trial results are accessible and understandable to patients and the public. As AI technology progresses, it presents both opportunities and challenges in the context of LS.

Our collaboration is composed of clinical trial transparency industry experts with diverse backgrounds in medical writing, technology, clinical operations, plain language, and patient engagement. This work group has considered how to leverage AI while maintaining high standards by ensuring that all AI-generated content is reviewed, revised, and validated by knowledgeable experts. Through this balanced approach, it can be ensured that LS are both accurate and accessible, fostering greater trust and understanding between the clinical research community and the public.

In this document, AI will be used to refer primarily to large language models (LLMs) that generate text. Commonly known LLMs include ChatGPT, Gemini, Claude, and Llama.

Additionally, lay summaries (LS) of clinical trial results are also known as lay language summaries (LLS) of clinical trial results, plain language summaries (PLS) of trial results, or trial results summaries (TRS).

A. Background

Lay summaries are essential for making clinical research results more transparent and accessible to non-scientists, addressing the traditional barrier of complex scientific language. AI has the potential to streamline the drafting of LS, saving time and resources. However, overreliance on AI to generate these summaries without appropriate human oversight can lead to inaccuracies or misinterpretations. This is especially relevant when using data from sources like ClinicalTrials.gov, which may lack the context to appropriately develop an LS.

AI's use in health and medical enterprises is increasingly subject to regulatory oversight, with the United States and EU developing frameworks aimed at ensuring data privacy, accuracy, and ethics in AI applications, such as the U.S. Blueprint for an AI Bill of Rights, the NIST AI Risk Management Framework, the EU Artificial Intelligence Act, and the EU Ethics Guidelines for Trustworthy AI. In parallel, guidelines from organizations on AI-generated medical writing like ICMJE, AMWA, EMWA, and IMPP emphasize the need for transparency, accuracy, and human oversight. In turn, organizations are developing and deploying AI use cases and respective policy documents, including tools for drafting scientific and public- or patient-facing documents.

As regulatory frameworks and guidelines and AI technology evolve, stakeholders must stay informed and adopt practices to ensure compliance and maintain high quality LS. At the time this document was developed, guidelines for using AI in medical information publication or disclosure were limited and lacked actionable recommendations, and no guidelines existed for the responsible use of AI in creating LS or other patient-friendly clinical research information.

This document was developed over an extended consultative process involving more than 15 organizations from the US and EU including industry, academia, and a non-profit patient-focused organization. The work group involved in creating this document is committed to ensuring accurate and appropriate use of AI in creating LS. This document aligns with recommendations from the Good Lay Summary Practice guidance document and other LS practices broadly accepted as industry standards.

This document was initially drafted using AI technology with the goal of testing the recommendations and considerations we have developed. Authors from the work group drafted select sections of the initial draft based on an outline the work group previously discussed and approved. After the first draft was created, the draft underwent multiple iterations of work group review and revision. This included incorporating feedback from a public comment period that informed the final contents of this document. [Placeholder: brief summary of the number and type of organizations that provided public comment feedback and the key comments/topics mentioned/revised]. Both human review and AI were also used to review and revise drafts for consistent voice, neutral language, any missing information, spelling, and grammar.

B. Opportunities and Challenges

AI has the potential to create efficiency in the creation of LS and enhance health literacy by making information more accessible to patients and the public. When used effectively, AI can streamline development, allowing for quicker delivery of clear, patient-friendly content.

While AI integration into LS processes offers many opportunities, it also presents challenges, including the risk of inaccuracies, biases, and ethical concerns. Human review is essential to ensure accuracy and clarity, as AI cannot fully comprehend the nuances of scientific data, cultural contexts, or the emotional tone needed for high-quality LS. By combining AI's efficiency with expert oversight, we can ensure the public and patients receive timely, accessible information promoting greater equity in healthcare. Responsible AI use is critical—with technology serving to complement, not replace, human expertise.

C. Application and Scope

The Good Lay Summary Practice guidance remains the accepted industry standard for creating and delivering high quality LS. This document can be used in addition to established standards and processes when seeking to use AI in the authoring stage of an established process.

We support a balanced approach that involves using AI for initial drafting and revisions, while maintaining transparency about its role and ensuring human oversight. This process mitigates risks, supports compliance with regulatory standards, and ensures summaries are accurate and patient friendly. As AI evolves, ongoing review and adaptation of these practices will be essential to meet emerging regulations and technological advances.

This document serves as a set of considerations for study teams and medical writers considering the use of AI in LS, emphasizing the need for human review and transparency. While focused on LS, many of the principles can also apply to other public- and patient-facing communications, such as the plain language protocol synopsis. By maintaining high standards and leveraging AI with expert oversight, we aim to produce understandable, accessible, and trustworthy LS that enhance patient and public understanding and overall trust in clinical research.

II. Areas for Concern:

A. Human Involvement

Overreliance on AI in creating LS for clinical trial results raises concerns, as AI lacks the nuanced understanding and contextual knowledge human experts provide. Without proper oversight, AI-generated LS may misrepresent complex data or miss critical details leading to inaccurate summaries. This concern was observed in 2023 in a large-scale instance of publicly posted, AI-generated LS that lacked proper human oversight. These lay summaries were eventually removed from the public domain after significant concerns were raised regarding their accuracy. To ensure accuracy, AI should complement, not replace, human expertise, with professionals reviewing and refining content. Additionally, involving patient and public representatives in the LS development process ensures summaries are both scientifically accurate and accessible to their intended audience. Their feedback enhances the relevance of LS, capturing perspectives that AI alone cannot address and fostering trust between researchers and the public.

B. Disclosure of AI use

Transparency about AI's role in developing LS is essential for maintaining public confidence. Undisclosed AI involvement may create skepticism or erode trust in the accuracy of the information. Given the evolving understanding of AI's capabilities, understandable communication about its use helps address potential misconceptions and fosters a more informed and trusting relationship between the public and the research community.

Per Chapter 4, Article 50, Paragraph 4 of the EU AI Act: *Deployers of an AI system that generates or manipulates text which is published with the purpose of informing the public on matters of public interest shall disclose that the text has been artificially generated or manipulated.*

C. Involvement of Research Sponsor

In most cases, the research sponsor(s) responsible for the disclosure of clinical trial results will develop an LS, either internally or outsourced with oversight including review and approval responsibilities. However, given the public availability of many clinical trial results and the increase in organizations who aim to make research results more accessible, there may be instances where research sponsors are not the drivers of LS development.

Sponsor staff and subject matter experts possess a deep understanding of a study's design, objectives, endpoints, and statistical analysis, providing critical insights that ensure an LS aligns with scientific nuances and addresses patient interests. Their involvement is vital to accurately interpret complex data and present it in a way that is both clear and accessible to the public. While tools and practices aimed at improving accessibility can promote equity, the absence of sponsor oversight increases the risk of misinterpretation or omission of critical details.

D. Misinformation and Disinformation

Misinformation refers to unintentional inaccuracies or errors in information, which can occur when AI misinterprets data or lacks the context needed to fully understand scientific nuances. Disinformation, on the other hand, involves the deliberate distortion of facts to mislead. While AI may unintentionally spread misinformation due to its limitations, there is also a risk that AI could be manipulated and generate content that leads to disinformation. This risk is particularly a concern if using AI tools that are open-source, or that draw from or are trained on public data.

E. Bias and Cultural Sensitivity

AI systems are heavily influenced by the data they are trained on and the prompts provided by users. If this training data or user inputs contain biases—whether intentional or not—AI can reproduce and even amplify those biases, resulting in skewed or unfair summaries that compromise the objectivity and fairness of information shared with patients and the public. Furthermore, because AI models are trained on large datasets that may not fully reflect the diverse cultural backgrounds and values of all readers, the generated content can lack cultural awareness and sensitivity. This may lead to misunderstandings or inaccuracies in handling language nuances and culturally specific references, diminishing engagement, trust, and the overall effectiveness of lay summaries for certain communities.

F. Consistency

AI-generated content can vary widely in tone, style, and accuracy depending on the data and algorithms used, which can lead to discrepancies across communications. Inconsistent messaging may confuse readers and reduce the clarity of critical information. Variations in language, phrasing, or emphasis may also result in misunderstandings.

G. Rapid technological change

As AI technologies and algorithms evolve quickly, using outdated models can lead to inaccuracies and inconsistencies in the generated content. This rapid pace of change may also make it challenging to keep AI tools aligned with the latest standards and best practices. This could increase the risk that LS may not meet current regulatory or quality expectations.

H. Trust

Readers depend on these LS for understandable and accurate information. Purely AI-generated content can sometimes appear impersonal, inconsistent, or biased, which can undermine confidence. The opaque nature of AI decision making—often referred to as the "black box" problem—further complicates trust, as it obscures how conclusions are reached. Without transparency and reliability in AI outputs, the credibility of LS can be significantly compromised.

I. Data privacy

Clinical study data sets contain detailed sensitive personal information about the patients from the clinical study. Even after a given level of anonymization, there remains a risk of re-identification, compromising patient confidentiality. This risk threatens patient trust and also poses reputational risks and legal risks (under privacy regulations such as the EU General Data Protection Regulation (GDPR), US HIPAA, and Health Canada's Privacy Act.) To ensure data privacy is always maintained, the input study data should be aggregated to a level that is no longer considered to be individual patient-level data and cannot be associated with an individual human person. AI LS writers must remember not to feed real individual patient level data into open AI models for risks of breaching various global privacy laws as well as patient trust.

III. Recommendations*:*

A. Overview

AI is a transformative tool that can significantly enhance human productivity in developing LS. However, its capabilities should be supplemented by human judgment, ensuring that critical decisions are not left solely to machine-generated outputs. Human involvement is crucial for addressing areas where AI might fall short, such as understanding the nuance of patient needs, scientific interpretation, and the complexities of compliance. Therefore, while AI can aid humans in drafting and organizing information, humans must retain ultimate control over the content and ensure its accuracy and compliance.

For example, AI might excel in creating initial drafts, identifying trends, or simplifying technical language. However, humans are essential for verifying contextual accuracy, ensuring the absence of bias, and aligning outputs with both the specific goals of the study and regulatory standards. AI, in this sense, serves as a powerful augmentative tool rather than a replacement for human expertise.

By fostering collaboration between humans and AI, organizations can maximize the efficiency and accuracy of LS, reducing the time spent on repetitive tasks while maintaining the integrity of content.

B. Suggested Additions to Process Flow

The below process flow is to be used during the authoring stage of an existing LS development process. We are not suggesting changes from the best practices and overall process as laid out in the Good Lay Summary Practice guidanceSummary Practice Guidance that has been largely adopted.

To effectively integrate AI into the development process for LS, it's essential to establish a clear process flow delineating where AI can be used and where human involvement is indispensable.

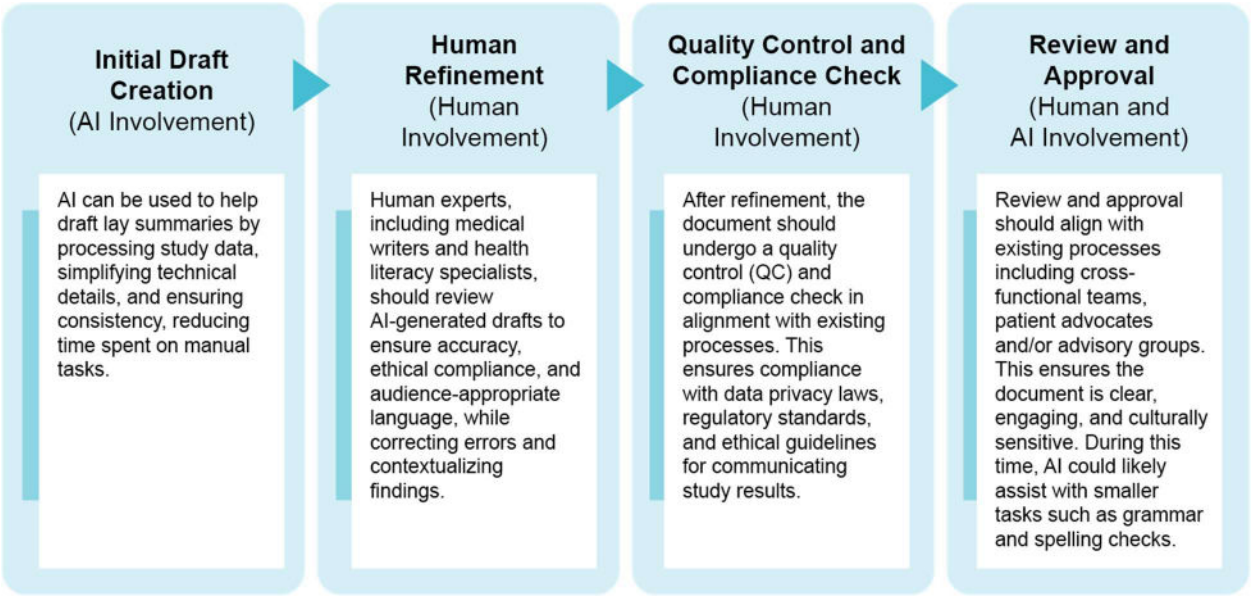

Throughout this process, maintaining clear documentation of AI's role in generating and refining content is essential, including any instances where human corrections or interventions were made. This documentation serves as a quality assurance measure and can be valuable in regulatory reviews.

C. What Humans to Involve and How They Should Be Involved with AI

The effectiveness of AI in generating LS is contingent upon the expertise of the humans involved in its training, oversight, and in carrying out the generation and revisions of LS. To ensure adherence to best practices and maintain quality and accountability, all reviewers and approvers recommended by the GLSP should retain their essential roles in the LS process, even when AI tools are integrated. While standard operating

procedures and resourcing at various organizations developing LS using AI may vary, stakeholders possessing the following knowledge and experience may play critical roles at various stages:

- **AI expertise**: Deep AI knowledge is required to design and train the AI systems used in drafting content to ensure the system is properly trained on relevant inputs and datasets, such as clinical study protocols, regulatory documents, LS templates, preferred terminology glossaries, and health literacy guidelines.
- **Health literacy & plain language**: Knowledge and skillsets in health literacy and plain language writing should be leveraged to help train the AI on simplifying complex medical language into terms that are accessible and understandable to the general public, including guiding AI on which terminologies, explanations, and formatting best align with the needs of the public and patients.
- **Data privacy**: To ensure that AI systems use sensitive personal information safely and in accordance with approved data use laws and policies (such as GDPR or HIPAA), data privacy expertise must be incorporated in training and monitoring the AI's use of patient data. Following the completion of AI training, data privacy can be monitored through standard LS review procedures.
- **Legal and compliance**: Individuals with knowledge of relevant legal and compliance standards may be involved in AI training and oversight. They review the content for compliance with local and international laws and guidelines, particularly concerning data privacy and AI regulation.
- **AI use:** Knowledge in the use of AI to generate LS and in LS development standards are needed to use AI both to create the initial draft LS and if used subsequently in LS draft refinement.
- **LS and medical writing**: Once AI generates a draft, LS writing knowledge is needed to ensure factual accuracy and alignment with the study's key findings by scrutinizing the AI’s interpretation of clinical results to ensure that no critical scientific nuances are absent in the LS.

By clearly defining roles and responsibilities across AI training, content accuracy, health literacy, and legal compliance, organizations can create a robust framework for integrating AI into existing LS practices while ensuring the highest standards of quality and ethical responsibility.

D. Considerations When Using AI

Implementing AI for LS requires a structured approach to ensure responsible and effective use. Large Language Models (LLMs), while capable of generating human-like text, have limitations in producing patient-friendly content. They may provide factually incorrect or inconsistent information, struggle with complex medical concepts or rare conditions, and fail to capture the appropriate tone for LS. Additionally, without proper guidance, LLMs can generate biased or insensitive content. Recognizing these limitations is key to ensuring AI-generated summaries are accurate and sensitive to public and patient needs.

When implementing AI for LS, several key considerations must be accounted for to ensure ethical, effective, and patient-centered outcomes:

Context: The specific clinical trial or healthcare setting in which the LS will be used must be carefully considered. AI should be trained in relevant, domain-specific data to ensure accuracy and appropriateness.

Audience: Reader demographics, health literacy levels, and cultural backgrounds should inform the AI's output. Customization options may be necessary to address diverse patient populations effectively.

Length: The ideal length of LS may vary depending on the complexity of the information and the preferences of the target audience. AI should be capable of producing content of varying lengths while maintaining clarity and completeness.

Template: Standardized templates can help ensure consistency and compliance with regulatory requirements. AI should be trained to work within these templates while allowing for necessary flexibility. AI should also be trained to use a glossary for preferred terminology within a particular document or set of documents.

Data inputs: The quality and comprehensiveness of data inputs are crucial for generating accurate and relevant LS. Key data sources may include:

- Tables, Figures, and Listings (TFLs) from clinical trial results
- Clinical Study Protocols (CSP)
- Clinical Study Reports (CSR)
- Informed Consent Forms (ICF)
- Other public or patient-facing documents
- Glossaries of medical terms and plain language equivalents

Prompt Engineering: A critical component of using AI effectively is prompt engineering, which guides the AI in creating accurate, understandable, and public- and patient-appropriate content. For each LS document to be drafted multiple prompts should be provided to the AI for drafting individual sections and for clear context setting. Specific instructions on tone and style, and guidelines for simplifying complex concepts should be provided. These prompts help the AI strike the right tone, ensure consistency with approved medical terminology, address potential biases, and promote inclusivity. By including reminders to provide necessary context and caveats, prompt engineering can help ensure that AI-generated content is both informative and patient-friendly. Please see Appendix A for components of good prompts and example prompts.

Governance: Robust AI governance is essential for overseeing AI in LS, requiring collaboration between medical writers, statisticians, legal experts, and patient advocates. Implementing AI is an iterative process that requires thorough initial testing and continuous improvement. Key elements include developing standards for AI use, comprehensive pre-launch suite testing, ongoing performance monitoring, and regular audits to ensure compliance with evolving regulations and best practices. Please see Appendix B for additional considerations.

Disclosure: Transparency in the use of AI for LS is vital to maintaining trust and ethical standards. This requires clear disclosure of AI involvement, explanation of human oversight, compliance with regulations like the EU AI Act, and acknowledgment of sponsor or patient community involvement in the process. Please see Appendix C for additional considerations and example statements of disclosure.

Advanced AI Architectures: Leveraging AI most effectively may require more advanced architecture, such as AI agent networks. Agent networks employ multiple AI agents, each with a specialized role such as a medical fact-checker, readability optimizer, and bias and sensitivity detector. Orchestrator agents can also be integrated into the architecture to coordinate the work of specialized agents, like a project manager, while humans continue to provide expert oversight and intervention at key points.

By carefully addressing these considerations, and through continuous learning, organizations can harness the potential of AI to enhance LS processes. Regular monitoring and updates to processes and AI models with the latest medical and regulatory information will likely be essential to mitigate associated risks and maintain the highest standards of accuracy, clarity, and ethical LS practice.

IV. Helpful tools and resources

Leverage existing tools and resources and develop additional, use-specific comprehensive resources to guide the development and use of AI for LS creation. The following tools, resources, and topics should be addressed.

A. Quality control checklists for content verification:

- Verification of medical facts and statistics against source documents (e.g., clinical study reports, published literature)
- Consistency checks with approved messaging and terminology
- Assessment of readability and health literacy levels
- Evaluation of cultural sensitivity and inclusivity
- Identification of potential biases or misleading statements
- Compliance with regulatory requirements and internal guidelines

B. AI model evaluation tools

- Bias detection and mitigation algorithms
- Model explainability tools to understand AI decision-making processes
- Performance benchmarking tools to compare AI outputs against human-generated content
- Annotation tools for providing feedback on AI-generated content

C. Data privacy tools

- Data anonymization and de-identification tools
- Secure file transfer protocols for sensitive information

- Access control systems to limit data exposure
- Encryption tools for data at rest and in transit
- Privacy impact assessment templates

D. Collaboration platforms

- Implement secure platforms for collaboration between AI systems and human experts
- Version control systems to track changes and approvals
- Annotation tools for providing feedback on AI-generated content
- Project management software to coordinate review and approval processes

E. Training resources for staff involved in using AI

- Develop comprehensive training materials for staff involved in AI-assisted LS creation
- E-learning modules on AI capabilities and limitations
- Good practices guides for human-AI collaboration
- Regular workshops and webinars on emerging AI technologies and ethical considerations

F. Additional Resources

- Good Lay Summary Practice Guidance (GLSP)
- International Society for Medical Publication Professionals (ISMPP) position statement and call to action on artificial intelligence
- European Medicines Agency (EMA) Artificial Intelligence Workplan
- 4 principles for safe and responsible use of LLMs (EMA)
- Guiding principles on the use of large language models in regulatory science and for medicines regulatory activities (EMA)

V. Conclusion

Incorporating AI into the creation of lay summaries presents both opportunities and challenges, underscoring the need for thorough planning and careful implementation. While AI can streamline certain aspects of the process, its output must always be guided by human expertise to ensure accuracy, sensitivity, and compliance. Successful implementation will be an ongoing process that requires continuous monitoring, evaluation, and refinement. Ultimately, integrating AI into LS development necessitates balancing innovation with accountability, ensuring that each summary meets the highest standards of transparency, ensuring trust and clarity for patients and the public.

Appendix A – Prompt Engineering Considerations

Components of good prompts:

- Clear context setting (e.g., "You are writing a lay summary for a clinical trial on [condition] for people with a 6th-grade reading level.")
- Specific instructions on tone and style (e.g., "Use a compassionate and encouraging tone while maintaining factual accuracy.")
- Guidelines for simplifying complex concepts (e.g., "Explain [medical term] in simple language a non-expert can understand.")
- Reminders to include necessary context and caveats (e.g., "Ensure to mention that these results may not apply to all patients and individual responses may vary.")

Example Prompts:

“Please create a lay summary of clinical trial results for a new diabetes medication. Your audience is the general public, including patients with type 2 diabetes, who have a 6th-grade reading level. Use a compassionate and encouraging tone while maintaining factual accuracy. Simplify complex medical terms but include them in parentheses after the simplified explanation. Ensure you mention the study’s limitations and that results may not apply to all patients. Structure the summary with understandable headings and bullet points for easy readability.”

“Please write a 3-paragraph explanation for why this trial: [trial name and NCT number from publicly available website] is being done. In the first paragraph please explain the condition, in the second paragraph please explain the study drug and why it is being developed, and in the third paragraph please discuss the trial design and restate the hypothesis for the final sentence. Please write the entire explanation at a 6th-grade reading level.”

“You are tasked with creating a lay summary of clinical trial results for a new diabetes medication. Your audience is the general public, including patients with type 2 diabetes, who have a 6th-grade reading level.

Here are the clinical trial results you will be summarizing: [insert documentation if within LLM capabilities/applicable].

Follow these guidelines to create your summary:

1. Use a compassionate and encouraging tone throughout the summary. Be warm and supportive but maintain factual accuracy.

2. Write at a 6th-grade reading level. Use simple words and short sentences. Avoid jargon or complex medical terminology.

3. Structure your summary with the following headings:

- What was the study about?

- What did the study find?

- What does this mean for me?

- What are the next steps?

4. Under each heading, use bullet points to present information clearly and concisely.

5. When introducing medical terms or concepts, first provide a simple explanation, then include the technical term in parentheses. For example: "sugar in the blood (glucose)".

6. Mention the study's limitations and clearly state that the results may not apply to all patients.

7. Begin your summary with a brief overview of the study's purpose (2-3 sentences).

Write your complete summary inside <summary> tags. Ensure that your summary is factually accurate based on the provided clinical trial results, while being easy to understand for the target audience."

Appendix B – Considerations for AI Governance

Effective governance is crucial when implementing AI for plain language summaries. A well-structured governance framework ensures that the use of AI aligns with organizational goals, regulatory requirements, and ethical standards. Key components of governance should include:

- Internal collaboration & standards development/implementation
  - Establish a cross-functional team including medical writers, statisticians, legal experts, patient advocates, and AI specialists.
  - Develop clear guidelines and standard operating procedures (SOPs) for AI use in patient communications.
  - Implement a review and approval process involving subject matter experts to validate AI-generated content.
  - Create a feedback loop to continuously improve AI performance based on human expert input.

- Initial testing
  - Develop a comprehensive test suite covering various scenarios, e.g., study phase, design, endpoints, safety data sets, patient populations
  - Conduct A/B testing comparing AI-generated content with human-written content for patient preference and understanding
  - Implement a feedback loop incorporating input from patients, healthcare providers, and subject matter experts
  - Regularly update and retrain AI models based on new data, feedback, and evolving best practices
  - Testing process example:
    1. Generate initial content using AI
    2. Review by humans for accuracy, readability, and health literacy levels using validated tools
    3. Incorporate public and patient involvement for feedback on understandability and relevance
    4. Iterate based on feedback, making necessary adjustments to prompts or AI models
    5. Repeat steps 1-5 until satisfactory results are achieved
    6. Implement in a limited rollout and monitor performance

    Scale implementation based on successful performance metrics

- Ongoing monitoring given AI's continuous learning
  - Implement a phased rollout, starting with low-risk applications and gradually expanding to more complex tasks.
  - Establish key performance indicators (KPIs) to measure the accuracy, readability, and effectiveness of AI-generated communications.
  - Conduct regular audits to assess AI performance.
  - Implement a system for ongoing monitoring of AI outputs, including random sampling and human expert review.

    - Develop protocols for addressing and correcting any errors or biases identified in AI-generated content.
    - Stay informed about advancements in AI technology and update systems accordingly to maintain state-of-the-art performance.

- Regulatory compliance
    - Ensure compliance with relevant regulations, such as the EU AI Act, GDPR, and FDA guidelines.
    - Maintain detailed documentation of AI training data, algorithms, and decision-making processes for regulatory audits.
    - Establish a process for staying updated on evolving regulations and adjusting AI systems and governance practices accordingly.

- Ethical considerations
    - Develop an ethical framework for AI use in patient communications, addressing issues such as bias, privacy, and transparency.
    - Implement safeguards to protect patient data and ensure confidentiality throughout the AI-assisted communication process.
    - Regularly assess the ethical implications of AI use and make necessary adjustments to maintain alignment with organizational values and societal expectations.

- Training and education
    - Provide comprehensive training for staff involved in AI-assisted patient communication processes.
    - Develop resources to help team members understand AI capabilities, limitations, and best practices for collaboration between humans and AI systems.

- Continuous improvement
    - Establish a process for collecting and analyzing feedback from patients, healthcare providers, and other stakeholders on AI-generated communications.
    - Use insights gained from feedback and performance monitoring to refine AI models and improve the quality of patient communications over time.

Appendix C – Considerations for AI disclosure

Transparency regarding the use of generative AI in creating patient communications is essential for maintaining trust, ethical standards, and regulatory compliance. Proper disclosure practices should address the following aspects:

- Where and when should the use of AI be disclosed and to what extent
  - Include a clear statement about AI involvement in the creation of the document, typically in the introduction or a dedicated section.
  - Disclose the extent of AI use, such as whether it was used for initial drafting, language simplification, or fact-checking.
  - Consider including a brief explanation of how AI was used in conjunction with human expertise to ensure accuracy and relevance.
  - Make the disclosure easily understandable for the target audience, avoiding technical jargon.

- AI regulation compliance
  - Ensure that disclosure practices align with the requirements of the EU AI Act or similar, applicable regulations.
  - Provide information on the AI system's purpose, capabilities, and limitations as required by applicable laws.
  - Include contact information for inquiries about the AI system or its outputs.

- Disclosure of sponsor or other human involvement:
  - Clearly state the level of involvement of the study sponsor and medical experts in reviewing and approving the LS.
  - Acknowledge any public or patient community involvement in the development or review of the LS.
  - If there was limited or no human involvement, this should also be disclosed transparently.

Example disclosure statements to include in LS:

- AI involvement disclosure
  - "This summary was initially drafted using artificial intelligence (AI) technology. After the first draft was created, it was reviewed, revised, and approved by qualified medical professionals to ensure accuracy, clarity, and relevance."

- Extent of AI use
  - "Artificial intelligence was used to assist in simplifying complex medical language and organizing information in this summary. All content has been verified and approved by the study team and patient representatives."

- Sponsor involvement
  - "The study sponsor, [Sponsor Name], has reviewed this AI-assisted summary to ensure its accuracy and alignment with the clinical trial results."

- Public and patient involvement
    - “Members of the public, patients, and patient advocates were also involved in the review of this summary to help ensure it is understandable and relevant."

- AI regulation compliance
    - "This document was created with the assistance of an AI system developed by [Company Name]. The system is designed to simplify medical language and organize information for LS. For more information about the AI system used, please contact [contact information]."

Appendix D – Good and Bad Examples of AI Use in LS creation

Good example:

> An AI-generated lay summary that accurately simplifies complex trial results, uses appropriate health literacy levels, and includes clear context for findings. The summary was reviewed according to standard LS review procedures, with their input incorporated to enhance accuracy, clarity, and relevance.

Lesson learned: Collaborative review processes involving diverse stakeholders can significantly improve the quality and patient-centeredness of AI-generated content.

Bad example:

> An AI-generated document that misinterprets statistical findings, leading to overly optimistic statements about treatment efficacy. There was not proper human oversight in the review of the document according to standard LS procedures, resulting in the distribution of misleading information to the public and patients.

Lesson learned: Implement multiple layers of expert review, including statistical verification, to catch and correct potential misinterpretations by AI systems.

Appendix E – Example of Advanced AI Architecture for LS Creation

- Advanced AI Architectures: For more complex LS tasks, a more advanced architecture might be needed, such as AI agent networks. These could include:
    - Multiple AI agents with specialized roles (e.g., medical fact-checker, readability optimizer, bias detector)
    - Orchestrator agents to coordinate the work of specialized agents
    - Human-in-the-loop systems for expert oversight and intervention at key points

- Example architecture:
    - Agent 1: Initial drafter using trial data
    - Agent 2: Medical accuracy checker
    - Agent 3: Readability and health literacy optimizer
    - Agent 4: Bias and sensitivity reviewer
    - Orchestrator: Coordinates the workflow and integrates outputs
    - Human Expert: Reviews and approves final output

- Rigorous Testing and Iterations: Implementing AI for LS is an iterative process that requires thorough testing and continuous improvement:
    - Develop a comprehensive test suite covering various scenarios and patient populations
    - Conduct A/B testing comparing AI-generated content with human-written content for patient and public preference and understanding
    - Implement a feedback loop incorporating input from patients, healthcare providers, and subject matter experts
    - Regularly update and retrain AI models based on new data, feedback, and evolving best practices
    - Conduct periodic audits to ensure ongoing compliance with regulatory requirements and ethical standards