A Symbiotic AI Approach for Trustable and Interactive Microbiome Data Interpretation

Francesco Greco*, Giuseppe Desolda and Antonio Piccinno

University of Bari "Aldo Moro", Department of Computer Science, Via Orabona 4, 70125, Bari, Italy

Abstract

The transition from raw genomic data to actionable clinical insights remains a primary challenge in microbiome-based diagnostics. While modern information systems can effectively identify taxonomic anomalies and provide quantitative visualizations (e.g., phyla abundance ratios), the resulting reports often present a "semantic gap" for patients and General Practitioners, who may lack specialized expertise. This position paper presents a Symbiotic AI (SAI) approach to bridge this gap by integrating Large Language Models (LLMs) with existing visualizations within a trustworthy and reproducible pipeline. Aligning with the AI2VIS4BigData Reference Model, our approach moves beyond static reporting by utilizing an LLM-based *Interpretative Layer*. To ensure reproducibility, we propose a semantic intermediate mapping that structures clinical data into a machine-readable format before the agent processes it. To ensure trust, the system employs an internal grounding process that synthesizes narrative explanations from validated clinical snippets and laboratory thresholds. Moreover, an interactive symbiosis is facilitated through a conversational interface linked to existing genomic visualizations. This interface enables a bi-directional exchange: the AI provides context-aware summaries of dysbiosis patterns, while the practitioner can conduct further inquiries and provide corrective feedback via the UI. The system captures these human-in-the-loop interactions in a vector-based expert knowledge base, enabling it to refine future explanations based on real-world clinical experience without compromising model stability. The proposed approach emphasizes traceability and interpretative clarity and may serve as an exemplar for regulatory-aware intelligent visualization systems.

Keywords

Microbiome, Symbiotic AI, LLM, Visual Analytics, Clinical Decision Support

1. Introduction

The human gut microbiome is increasingly recognized as a critical determinant of systemic health, playing a fundamental role in the pathogenesis of functional gastrointestinal disorders such as Irritable Bowel Syndrome (IBS) [1]. Recent advances in high-throughput sequencing and bioinformatic pipelines have enabled the generation of highly accurate microbial profiles. These complex reports provide a wealth of quantitative data, including taxonomic abundances, diversity indices, and metabolic markers. However, the transition from raw genomic artifacts to actionable clinical insights is hindered by a profound semantic gap for untrained people, such as patients seeking clarity on their well-being and General Practitioners (GPs), who either lack specialized knowledge of the gut microbiome and innovative analysis methods or do not have time to process complex reports during quick check-up office appointments. For such stakeholders, the dense array of bar charts and statistical indices of a microbiome report remains semantically opaque. On the other hand, specialists still need detailed, exhaustive information for accurate analyses and diagnoses.

While laboratory reports typically provide static descriptions of individual biomarkers, they often fail to synthesize these disparate data points into a holistic diagnostic narrative, making it challenging to correlate the data and consider all significant values (e.g., anomalies in stool color, abundance of specific phyla) simultaneously. Furthermore, clinical microbiome science is characterized by an abundance of recent scientific literature yielding experimental results that highlight the relevance of contextual factors (demographic, dietary, clinical, or lifestyle factors [2, 3]). Incorporating these

Proceedings of TRI-IVIS 2026, the First International Workshop on Trustable, Reproducible, and Intelligent Information Visualization Systems, June 2026, Venice, Italy.

francesco.greco@uniba.it (F. Greco); giuseppe.desolda@uniba.it (G. Desolda); antonio.piccinno@uniba.it (A. Piccinno)

0000-0003-2730-7697 (F. Greco); 0000-0001-9894-2116 (G. Desolda); 0000-0003-1561-7073 (A. Piccinno)

emerging scientific nuances into a stable, validated knowledge base within a traditional information system is a slow, regulatory-heavy process. Therefore, even specialists can find it hard to consider all relevant contextual information at once and provide an accurate diagnosis.

In this position paper, we propose a Symbiotic AI (SAI) approach that bridges this gap by integrating Large Language Models (LLMs) as *Semantic Orchestrators* within a trustworthy pipeline. The proposed approach adheres to the AI2VIS4BigData Reference Model [4] by decoupling analytical processing from interpretative synthesis.

2. Related Work

2.1. Genomic Data Visualization

The visual representation of genomic and microbiome data is a fundamental prerequisite for translating raw bioinformatics outputs into actionable clinical insights.

Traditional methods are designed to provide experts with detailed visualizations and provide advanced statistical computations, such as taxonomic profiling and biodiversity indices [5, 6, 7]. A commercial example is *Microbiome Insights* [8], which provides advanced scientific visualizations for practitioners, including heatmaps showing bacterial abundances, multivariate graphs for comparing samples, clustering diagrams, and microbial correlation networks. However, such approaches do not fit the use by a general customer with low health literacy (i.e., the patient) [9].

Specialized frameworks such as *phyloseq* [7] and *MicrobiomeAnalyst* [6] are significant examples that provide reproducible pipelines for statistical and graphical analysis of microbiome census data. Furthermore, recent platforms such as *MicrobioSee* [5] and *Biovis Diagnosik* [10] (see Fig. 1) have introduced comprehensive web-based visualization toolkits tailored for genomics data reports. There are also tools for microbiome analysis reports targeted at patients with relatively low digital literacy, such as *Viome* [11], which not only provides summary scores (e.g., inflammation activity, metabolic fitness) but also personalized dietary recommendations. Within the *GenDAI* European project[1], several authors have worked towards improving clinical metagenomic diagnostics. These efforts led to the development of a dynamic web platform for visualizing and interactively managing microbiome clinical reports, which also enables dynamic PDF report generation [12]. This is further supported by a cloud-based architecture that can provide reproducible, policy-compliant, and trustworthy metagenomic reporting by embedding data provenance-by-design and policy-as-code enforcement [13]. Moreover, Krause and colleagues defined a comprehensive and ethically grounded framework that leverages Genomic Foundation AI Models for biomarker discovery, user-centered interactive reporting tools, and a long-term data management system [14].

As highlighted by user-centered design studies in clinical contexts, standard static reports or conventional data dashboards typically impose a high cognitive load on general practitioners and patients [9]. This is largely due to their reliance on complex scientific nomenclature and non-contextualized numerical thresholds. The current challenge lies in developing a interactive interfaces that move beyond the generic representation of taxonomic data (e.g., phyla abundance, enterotypes, and metabolic scores) towards adaptive and progressive disclosure techniques that support immediate clinical interpretation. In the context of genomic applications, Artificial Intelligence (AI) is essential for clinical diagnosis, from interpreting raw genomic data to high-level clinical reasoning. The AI2VIS4BigData Reference Model [4] provides a comprehensive framework for developing visualizations based on AI systems that process big data, such as genomic data, covering every step from data processing to user interaction.

2.2. LLMs in Healthcare for Clinical Diagnoses

With technological advances in AI of recent years, particularly Large Language Models (LLMs), practitioners in the healthcare domain can now be properly supported by these technologies in diagnostic

[1] https://project-gendai.eu/

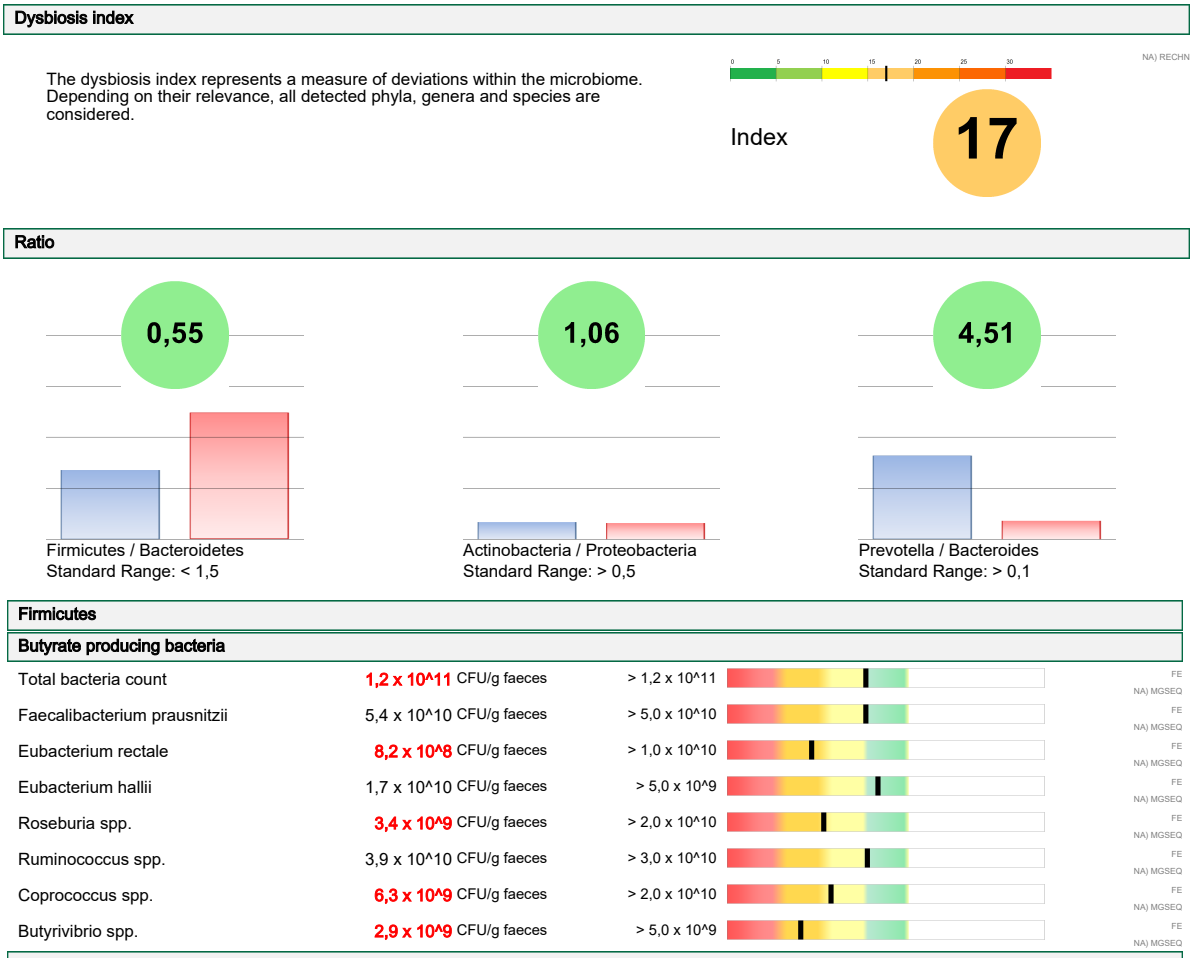

Figure 1: Microbiome analysis detail extracted from a sample report by *Biovis Diagnostik*, describing summary indices (Dysbiosis index, major ratios) and Firmicutes distribution [10]

processes [15, 16, 17, 18]. LLMs can provide an adaptive interpretative layer that translates high-dimensional data into natural-language summaries, thereby improving user understanding and reducing clinician burden [19, 20]. LLMs can also serve as a valuable source of information and education for patients, potentially enhancing communication with practitioners [21]. Moreover, the possibility of providing LLMs with external knowledge—which can be referred to as Retrieval Augmented Generation (RAG)—makes these AI models particularly fitting for incorporating new knowledge into their reasoning context, ultimately providing more accurate outputs [22, 23].

2.3. Symbiotic AI (SAI) and Human-in-the-loop vs. Human-on-the-loop

Modern visualization frameworks and LLMs have the potential to improve user efficiency in healthcare applications and in genomic data analysis. However, using these technologies without proper orchestration is insufficient to bridge the clinical semantic gap and convey deep medical insights. Symbiotic Artificial Intelligence (SAI) approaches can be used to integrate intelligent agents directly within the analytical and decision-making workflow. SAI systems can be defined as *Human-Centered AI systems powered by a symbiosis of human intelligence and artificial intelligence* [24]. This means that SAI systems benefit from human-AI collaboration, which should augment the abilities of both the user and the AI system. In the domain of healthcare, this can mean that the expert user (e.g., the practitioner), who benefits from AI suggestions to increase their efficiency, can instruct the system when it is wrong, eventually improving its future performance [25]. This bi-directional communication can also be conceived as a continuous negotiation of the user's and the AI's mental models, according to the Mutual

Theory of Mind [18].

When dealing with Human-Centered AI systems, we can define two interaction paradigms: *Human-in-the-Loop* (HITL) and *Human-on-the-Loop* (HOTL) that define the degree and nature of human agency within an automated system [24]. A HITL approach establishes a bidirectional relationship in which the human is an active participant in the decision-making process; the AI is not treated as an autonomous "oracle" but as a supportive collaborator that learns from continuous human expertise. In contrast, the HOTL paradigm prioritizes AI autonomy for high-efficiency tasks, positioning the human as a supervisor who monitors the system's behavior and intervenes only to correct errors or handle exceptions through specialized interfaces. While HITL focuses on co-creation and constant interaction, HOTL focuses on oversight and intervention.

3. Proposed Architecture: The Symbiotic Interpretative Framework

The tool we propose in this work is based on previous work conducted on the *Anonymized Project* and is structured as a layered interpretative information system for analyzing microbiome data with a Symbiotic AI approach. We map our architecture to the AI2VIS4BigData reference model [4] as follows:

1. **Data Management & Curation** – A Semantic Intermediate Mapping (SIM) layer transforms raw genomic artifacts into grounded semantic tokens using JSON-LD formatting.
2. **Analytics** – An LLM Orchestrator and a Similarity Engine map current profiles to historical expert cases and laboratory standards.
3. **Interaction & Perception** – The UI provides a multi-resolution narrative that adapts to different user roles through semantic zooming and bidirectional visual linking.
4. **Insight & Effectuation** – An Expert Knowledge Base (EKB) captures and redistributes expert *Contextual Exceptions* via Human-On-The-Loop (HOTL) oversight.

3.1. Data Management & Curation: The SIM Layer

To bridge the gap between low-level genomic analysis and upper-level clinical reasoning, we implement a Semantic Intermediate Mapping (SIM) layer. This layer prepares and curates raw data for downstream analytics and serves as the ground truth anchor for the top-level Large Language Model (LLM).

The raw output of the microbiome analysis—consisting of stool parameters (color, consistency, pH), metabolome values, phyla abundances (e.g., *Firmicutes, Bacteroidetes*), phyla ratios, and other core indeces (Biodiversity, Shannon, Butyrate, Inflammation, Dysbiosis)—is transformed into a structured JSON-LD (Linked Data) schema:

$$\mathcal{T} : \mathbf{D}_{raw} \rightarrow \mathbf{S}_{LD}$$

where $\mathbf{D}_{raw}$ represents the high-dimensional genomic artifacts and $\mathbf{S}_{LD}$ is the resulting JSON-LD schema. This schema maps numerical values to their respective laboratory reference intervals and several textual descriptions previously validated by experts, which can be considered Gold Standards. These Gold standard descriptions can, e.g., summarize what a Phylum is (*"Firmicutes is an abundant phylum in the gut, includes butyrate-producing bacteria that are important for intestinal balance"*) or describe what an off-range value means (*"Stool color: whitish – Possible biliary obstruction, carefully monitoring needed."*).

This formal curation prevents the LLM from interacting directly with raw data, rather providing the model with a solid, curated knowledge base in textual and structured format, which helps prevent hallucinations and increases traceability.

3.2. Analytics: The Semantic Orchestrator and Contextual Exceptions

The Analytics layer in the reference model typically handles the generation of models and insights. The proposed tool re-envisions this layer as a Semantic Orchestrator.

The core analytical intelligence resides in the LLM-based *synthesis agent*, which retrieves Gold Standard snippets and expert annotations to construct a context-aware narrative for interpreting the

microbiome analysis report at a high level. Simultaneously, this layer uses a similarity engine to retrieve relevant clinical exceptions for a patient by searching the Expert Knowledge Base (EKB). We can refer to these as *Contextual Exceptions*.

3.2.1. Contextual Exceptions

When a specialist identifies a clinical nuance that challenges or expands upon the "Gold Standard" static description (e.g., a specific dietary confounding factor [26]), they can submit a **Contextual Exception**. For example, while a low Shannon index is universally flagged as a potential sign of dysbiosis (associated with metabolic syndrome, IBS, and reduced resilience to pathogens), in elite athletes or individuals who are undergoing excessive training, a transient decrease in microbial diversity due to physiological stress can be considered normal [1]. Similarly, the Firmicutes/Bacteroidetes (F/B) ratio is highly sensitive to lifestyle factors: while an increased F/B ratio is normally considered as a hallmark of an "obese-type" microbiome (associated with increased energy harvest from the diet), in a lean patient with a high-carbohydrate, high-fiber intake, a high F/B ratio may be a neutral marker of high fermentation activity [3]. Finally, low diversity (lower Shannon index) and higher Proteobacteria abundance are associated with antibiotic use and can thus be taken as a sign of temporary recovery if that is the case [2].

3.2.2. Anchor-Weighted Similarity Matching

To ensure that a contextual exception is retrieved only when it is clinically relevant, we must define a distance metric between two clinical profiles. First of all, let P denote a **Patient Profile**, defined as a vector in an n-dimensional space $\mathbb{R}^n$, where n represents the total number of features in the comprehensive microbiome report (including patient demographics, R4DQ results describing IBS symptoms [27], taxonomic abundances, metabolic indices, Bristol scale results describing stool characteristics [28]).

- $P_{new} = [v_{1,new}, v_{2,new}, \dots, v_{n,new}]$ represents the **current Patient Profile** under analysis.
- $P_{ex} = [v_{1,ex}, v_{2,ex}, \dots, v_{n,ex}]$ represents a **historical Exception Profile** stored in the EKB, which carries a specialist's contextual annotation.
- v_i denotes the **normalized value** of the i-th feature (e.g., the abundance percentage of *Proteobacteria* or the numeric Dysbiosis Index).

We then define the EKB containing M contextual exceptions defined by practitioners as:

$$\mathscr{K}_i = \langle P_{ex}, T_{ex}, A_{ex} \rangle$$

$$\text{EKB} = (\mathscr{K}_1, \mathscr{K}_2, \dots, \mathscr{K}_M)$$

where P_{ex} is the full microbial profile (the "whole report"), T_{ex} is the textual description, and

$$A_{ex} \subseteq \{1, \dots, n_a\},\ n_a < n$$

is the set of *Anchors*—the specific feature indices identified by the expert as the primary clinical drivers for the exception (e.g., Shannon index, Beta-glucuronidase levels).

The similarity between a new patient profile P_{new} and a historical exception P_{ex} is determined by a **Local Weighted Distance**. Unlike global weighting schemes, our approach utilizes the Anchors A_{ex} to dynamically inform the importance of features during retrieval:

$$d(P_{new}, P_{ex}) = \sqrt{\sum_{i=1}^{n} w_i (v_{i,new} - v_{i,ex})^2}$$

where the weight w_i is a function of the expert's anchors:

$$w_i = \begin{cases} \alpha & \text{if } i \in A_{ex} \\ \beta & \text{if } i \notin A_{ex} \end{cases}$$

with $\alpha \gg \beta$. This structure ensures that, while the *entire* report is considered (to maintain context), the features the expert anchored their reasoning upon are given clinical precedence.

At the point of expert intervention, a Specialist establishes a persistent link by creating a tuple $\langle P_{ex}, T_{ex}, A_{ex} \rangle$ that binds the qualitative insight (T_{ex}) to the quantitative evidence (P_{ex}) through the expert's rational anchors (A_{ex}).

Upon the ingestion of a new profile P_{new}, the Semantic Orchestrator performs a similarity scan across the EKB. For every candidate exception, the system dynamically recalculates the feature weights w_i based on the stored anchor set A_{ex}.

An exception is retrieved and presented in the UI as a "Contextual Insight" only if the distance falls below a defined fidelity threshold ϵ, which is calibrated to ensure high-fidelity matching, preventing the cross-contamination of expert advice across profiles that are only superficially similar:

$$\text{Display Exception} \iff d_w(P_{new}, P_{ex}) < \epsilon$$

At the end, the LLM then synthesizes the retrieved T_{ex} into the patient's overview, providing an expandable visual bridge that highlights the specific bar charts or indices corresponding to A_{ex} within the current report to justify the association. This enables the GP to verify the specialist's thought process—e.g., *"This insight is shown because your patient's Beta-glucuronidase and pH levels match a historical profile annotated by a specialist. Click here to compare the full profiles"*.

3.3. Interaction & Perception: Multi-Resolution Narrative UI

The Interaction and Perception layer is implemented within the proposed architecture as a dashboard that uses semantic zooming to support three distinct stakeholder resolutions. The LLM acts as an adaptive interface, modulating its output based on the stakeholder's profile:

- **Low Resolution (Patient):** The engine prioritizes understandability by providing high-level, qualitative summaries focusing on lifestyle and general well-being. Technical taxonomic data is abstracted into functional summaries (e.g., "Bacteria responsible for fiber digestion").
- **Medium Resolution (General Practitioner):** The narrative shifts toward actionable decision support. It highlights correlations across different cards of the report, linking summary indices (e.g., Dysbiosis Index) to specific microbial drivers, to assist the GP in determining whether a specialist referral is required.
- **High Resolution (Specialist):** The engine provides a high-density, anomalies-first synthesis. It bypasses basic definitions and focuses on the intersection of the entire dataset, identifying non-obvious patterns in the high-dimensional microbial landscape.

Furthermore, we devise a bidirectional link between the narrative text and the visual bar charts to ensure the user's perception of the AI's summary is always anchored in quantitative visual evidence. Specifically, when the LLM generates a narrative summary, key entities (e.g., Inflammation Score) serve as interactive triggers: clicking them automatically highlights and scrolls to the corresponding bar chart in the report (e.g., such as the ones depicted in Fig. 1).

A **conversational interface** is always present to allow the user to interact with the report by: (i) asking a generic question regarding the report, (ii) clicking on a specific entry and receiving a detailed explanation in the chat about what it means and what the value indicates (e.g., whether it is outside reference intervals), or (iii) asking a follow-up question to receive further insights. Apart from the report itself, the LLM engine is prompted about the user role (i.e., patient, GP, or specialist) to adapt its default level of detail and scientific terminology.

3.4. Insight & Effectuation: Expert Knowledge Capture

The Insight and Effectuation layer of the AI2VIS4BigData model employs a symbiotic feedback loop between the human user and the analytical engine. GPs and specialists can access the *contextual*

exceptions, which appear as non-intrusive, expandable pop-ups. These can provide them with insights from other specialists.

Specialists have access to the contextual exceptions submission functionality. Before submitting it, the practitioner can use the conversational interface to validate or refine their theses with the LLM. When a contextual exception is provided, this interaction is captured as a knowledge tuple and fed back into the Analytics layer (the EKB). This ensures that the system's intelligence is not static but evolves with time as experts use the tool and provide insights for specific clinical cases.

In fact, while gold-standard descriptions constitute an initial knowledge base that covers common knowledge and basic cases, medical knowledge is non-static and characterized by a long tail of contextual nuances that are too numerous and variable to be codified in a rigid expert system. On the other hand, an LLM-mediated exception layer allows the system to keep up with clinical discovery without compromising the integrity of the validated laboratory core. Moreover, this approach maintains the gold-standard knowledge base as a stable, regulated anchor to ensure clinical safety and regulatory compliance (e.g., adherence to the IVDR 2017/746 standards [29]).

HOTL Curation To secure the system against data poisoning, we adopt a Human-on-the-Loop (HOTL) curation model for the feedback loop. While any authorized specialist can submit a contextual exception, these are marked as "Provisional" until a senior curator periodically reviews and promotes them to the global EKB. This prevents subjective or conflicting advice from degrading the system's reliability and ensures it can provide a documented audit trail of how and why its interpretative logic has been updated over time, maintaining regulatory compliance.

4. Discussions and Future Work

4.1. Managing Expert Conflict

The primary challenge of the proposed approach is the potential for expert conflict. Specifically, two specialists may disagree and provide different diagnoses for similar cases. As a result, a GP (or another specialist) would be exposed to not only the AI summary but also to two additional contrasting suggestions. The solution we propose is that another senior practitioner acts as a curator and verifies all contextual exceptions before they are considered by the system for future diagnoses. However, this approach bears problems of (i) efficiency, as the curator becomes the bottleneck in the system, (ii) biased data, as all contextual exceptions are exclusively the ones approved by the curator, based on their knowledge (and beliefs), and (iii) responsibility overload, as the curator becomes accountable for any medical suggestions to be considered by the AI.

To address this problem, future work will investigate a "Specialist Consensus" mechanism to vet exceptions and explore the legal implications of AI-mediated medical annotations under the EU AI Act [30].

4.2. Towards a Holistic Clinical Reasoning Engine

While the proposed tool focuses on the interpretative synthesis of microbiome reports, the ultimate goal is to transition toward holistic clinical reasoning. This extension would involve integrating a *Phenotypic context Layer* that augments the microbial profile with additional patient data, including age, gender, weight/BMI, historical clinical records, lifestyle, and dietary habits.

To manage the "noise" inherent in subjective patient reporting, we propose the ingestion of validated clinical instruments into the similarity engine. For example, dietary patterns can be captured via Food Frequency Questionnaires (FFQ) [31], gastrointestinal symptom clusters can be categorized using the *Rome IV Criteria* [27], and physical activity levels would be quantified via the *International Physical Activity Questionnaire* (IPAQ) [32]. By mapping this quantitative information onto objective clinical scales, they can be integrated into the proposed weighted distance formula d_w as additional contextual data.

To further support the analysis of the resulting high-dimensional data into a holistic reasoning process, the tool should possess research capabilities to retrieve literature from medical databases (e.g., PubMed) and identify experimental findings specific to a patient's unique clinical context. This is especially needed given the sheer volume of new medical studies and practitioners' cognitive limitations, exacerbated by high-pressure medical contexts [18].

Technical Challenges The transition to a holistic engine introduces significant technical challenges. First of all, with a larger number of features, the so-called "curse of dimensionality" threatens the efficacy of Euclidean distance metrics; *hierarchical similarity matching*, in which microbial and phenotypic data are processed in nested layers, will be investigated in future work.

Secondly, transforming unstructured medical notes of previous medical records into structured ontologies for direct comparison between similar cases is not trivial and needs thorough experimentation (e.g., applying *Named Entity Recognition* pipelines).

Finally, managing Protected Health Information (PHI) within an LLM-mediated environment necessitates highly private data processing, rendering commercial tools such as GPT or Gemini unusable in the first place. Therefore, it is strictly necessary to explore locally deployed inference models and apply appropriate anonymization techniques to ensure compliance with global data protection regulations (e.g., GDPR) [33].

5. Conclusions

In this position paper, we presented an approach for high-level interpretation of microbiome data for clinical diagnosis that combines deterministic algorithms and LLMs in a Symbiotic AI tool. Various users (patients, GPs, and specialists) can leverage natural-language interaction combined with existing visualization techniques to obtain insights at different levels. Notably, with the proposed approach, specialists can extend the AI model's knowledge to provide colleagues with diagnoses grounded both in validated literature and their own insights.

While this position paper aims to present a novel approach at the crossroads of Symbiotic AI, medicine/bioinformatics, and information visualization, the next imminent step is to implement (and validate with users) the proposed approach in a web platform to support diagnosis in current microbiome analysis reports.

Acknowledgments

Funded by the European Union

This work has been supported by the EU Marie Skłodowska-Curie Actions (MSCA) GenDAI - Genomic applications for laboratory Diagnostics supported by Artificial Intelligence (Grant agreement ID: 101182801, DOI: 10.3030/101182801). This project is funded by the European Union, MSCA SE, Grant Agreement No 101182801.

Declaration on Generative AI

During the writing of this paper, the author(s) used *Grammarly Pro* to fix grammatical errors and *Gemini Pro* to improve text quality. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

[1] W. Barton, N. C. Penney, O. Cronin, I. Garcia-Perez, M. G. Molloy, E. Holmes, F. Shanahan, P. D. Cotter, O. O'Sullivan, The microbiome of professional athletes differs from that of more sedentary subjects in composition and particularly at the functional metabolic level, Gut 67

(2018) 625–633. URL: https://gut.bmj.com/content/67/4/625. doi:10.1136/gutjnl-2016-313627. arXiv:https://gut.bmj.com/content/67/4/625.full.pdf.

[2] S. A. Kelly, J. Nzakizwanayo, A. M. Rodgers, L. Zhao, R. Weiser, I. A. Tekko, H. O. McCarthy, R. J. Ingram, B. V. Jones, R. F. Donnelly, B. F. Gilmore, Antibiotic therapy and the gut microbiome: Investigating the effect of delivery route on gut pathogens, ACS Infect. Dis. 7 (2021) 1283–1296.

[3] F. Magne, M. Gotteland, L. Gauthier, A. Zazueta, S. Pesoa, P. Navarrete, R. Balamurugan, The firmicutes/bacteroidetes ratio: A relevant marker of gut dysbiosis in obese patients?, Nutrients 12 (2020). URL: https://www.mdpi.com/2072-6643/12/5/1474. doi:10.3390/nu12051474.

[4] T. Reis, M. X. Bornschlegl, M. L. Hemmje, Ai2vis4bigdata: A reference model for ai-based big data analysis and visualization, in: T. Reis, M. X. Bornschlegl, M. Angelini, M. L. Hemmje (Eds.), Advanced Visual Interfaces. Supporting Artificial Intelligence and Big Data Applications, Springer International Publishing, Cham, 2021, pp. 1–18.

[5] J. Li, Y. Sang, S. Zeng, S. Mo, Z. Zhang, S. He, X. Li, G. Su, J. Liao, C. Jiang, Microbiosee: A web-based visualization toolkit for multi-omics of microbiology, Frontiers in Genetics Volume 13 - 2022 (2022). URL: https://www.frontiersin.org/journals/genetics/articles/10.3389/fgene.2022.853612. doi:10.3389/fgene.2022.853612.

[6] A. Dhariwal, J. Chong, S. Habib, I. L. King, L. B. Agellon, J. Xia, Microbiomeanalyst: a web-based tool for comprehensive statistical, visual and meta-analysis of microbiome data, Nucleic Acids Research 45 (2017) W180–W188. URL: https://doi.org/10.1093/nar/gkx295. doi:10.1093/nar/gkx295. arXiv:https://academic.oup.com/nar/article-pdf/45/W1/W180/18137250/gkx295.pdf.

[7] P. J. McMurdie, S. Holmes, phyloseq: An r package for reproducible interactive analysis and graphics of microbiome census data, PLOS ONE 8 (2013) 1–11. URL: https://doi.org/10.1371/journal.pone.0061217. doi:10.1371/journal.pone.0061217.

[8] Microbiome Insights Inc., Microbiome analytics and visual reporting, https://info.microbiomeinsights.com/en/microbiome-study-guide, 2024. Commercial tool/service.

[9] S. S. Khairat, A. Dukkipati, H. A. Lauria, T. Bice, D. Travers, S. S. Carson, The impact of visualization dashboards on quality of care and clinician satisfaction: Integrative literature review, JMIR Human Factors 5 (2018) e22. URL: https://doi.org/10.2196/humanfactors.9328. doi:10.2196/humanfactors.9328.

[10] Biovis Diagnostik MVZ GmbH, Mikrobiom maxi refert: Molecular genetic microbiome analysis, 2022. Diagnostic report sample. Limburg, Germany.

[11] Viome Inc., Personalized microbiome health reports, 2024. URL: https://viome.co.za/Blog/Article/guide-your-viome-results, commercial direct-to-consumer microbiome analysis service.

[12] P. Buono, P. Tamla, T. Krause, F. De Luzi, F. Monti, M. Mecella, Revisiting data visualizations in diagnostic reports, in: 2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2025, pp. 7200–7206. doi:10.1109/BIBM66473.2025.11356146.

[13] P. Tamla, T. Krause, M. Hemmje, F. Monti, F. De Luzi, M. Mecella, B. Andrade, P. Buono, A. Molinari, The gendai cloud-native infrastructure and data stewardship for clinical metagenomic diagnostics, in: 2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2025, pp. 7229–7236. doi:10.1109/BIBM66473.2025.11356811.

[14] T. Krause, P. Tamla, A. Leoni, F. Monti, F. De Luzi, J. F. Gerald, B. Andrade, H. Afli, M. Mecella, P. Buono, A. Molinari, M. Hemmje, From reads to reports: A vision for a gfm-powered genomic diagnostic platform, in: 2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2025, pp. 7207–7213. doi:10.1109/BIBM66473.2025.11356528.

[15] S. Liu, A. P. Wright, B. L. Patterson, J. P. Wanderer, R. W. Turer, S. D. Nelson, A. B. McCoy, D. F. Sittig, A. Wright, Using ai-generated suggestions from chatgpt to optimize clinical decision support, Journal of the American Medical Informatics Association 30 (2023) 1237–1245. URL: https://doi.org/10.1093/jamia/ocad072. doi:10.1093/jamia/ocad072. arXiv:https://academic.oup.com/jamia/article-pdf/30/7/1237/50634330/ocad072.pdf.

[16] P. G. Brodeur, T. A. Buckley, Z. Kanjee, E. Goh, E. B. Ling, P. Jain, S. Cabral, R.-E. Abdulnour, A. D. Haimovich, J. A. Freed, A. Olson, D. J. Morgan, J. Hom, R. Gallo, L. G. McCoy, H. Mombini, C. Lucas,

M. Fotoohi, M. Gwiazdon, D. Restifo, D. Restrepo, E. Horvitz, J. Chen, A. K. Manrai, A. Rodman, Superhuman performance of a large language model on the reasoning tasks of a physician, 2025. URL: https://arxiv.org/abs/2412.10849. arXiv:2412.10849.

[17] S. Maleki Varnosfaderani, M. Forouzanfar, The role of ai in hospitals and clinics: Transforming healthcare in the 21st century, Bioengineering 11 (2024). URL: https://www.mdpi.com/2306-5354/11/4/337. doi:10.3390/bioengineering11040337.

[18] C. J. Greengrass, Transforming clinical reasoning–the role of ai in supporting human cognitive limitations, Frontiers in Digital Health 7 (2026) 1715440. URL: https://doi.org/10.3389/fdgth.2025.1715440. doi:10.3389/fdgth.2025.1715440.

[19] M. van Driel, N. Blok, J. A. van den Brand, D. van de Sande, M. de Vries, B. Eijlers, F. Smits, J. J. Visser, D. Gommers, C. Verhoef, M. E. van Genderen, D. J. Grünhagen, D. E. Hilling, Leveraging gpt-4 enables patient comprehension of radiology reports, European Journal of Radiology 187 (2025) 112111. URL: https://www.sciencedirect.com/science/article/pii/S0720048X25001974. doi:https://doi.org/10.1016/j.ejrad.2025.112111.

[20] S. J. Shah, A. Nair, K. Murtagh, S. P. Ma, K. Vogt, D. Clutter, L. Sheikh, H. Schmidt, M. Smith, A. Lakhotia, L. Bullock, A. Bhasin, M. A. Pfeffer, C. Sharp, S. Lin, P. Garcia, Clinician perspectives on ai-generated drafts of patient test result explanations, JAMA Network Open 8 (2025) e2528794–e2528794. URL: https://doi.org/10.1001/jamanetworkopen.2025.28794. doi:10.1001/jamanetworkopen.2025.28794.

[21] S. Aydin, M. Karabacak, V. Vlachos, K. Margetis, Large language models in patient education: a scoping review of applications in medicine, Frontiers in Medicine 11 (2024) 1477898. URL: https://doi.org/10.3389/fmed.2024.1477898. doi:10.3389/fmed.2024.1477898.

[22] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive nlp tasks, in: Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20, Curran Associates Inc., Red Hook, NY, USA, 2020, pp. 9459–9474. doi:10.5555/3495724.3496517.

[23] N. Bölücü, J. Irons, C. Lee, B. Jin, M. Rybinski, H. Yang, A. Duenser, S. Wan, Using a human-ai teaming approach to create and curate scientific datasets with the scilire system, 2026. URL: https://arxiv.org/abs/2603.12638. arXiv:2603.12638.

[24] G. Desolda, A. Esposito, R. Lanzilotti, A. Piccinno, M. F. Costabile, From human-centered to symbiotic artificial intelligence: a focus on medical applications, Multimedia Tools and Applications 84 (2025) 38983–38983. URL: https://doi.org/10.1007/s11042-025-21004-9. doi:10.1007/s11042-025-21004-9.

[25] A. Esposito, M. Calvano, A. Curci, F. Greco, R. Lanzilotti, A. Piccinno, Explanation-driven interventions for artificial intelligence model customization, in: C. Santoro, A. Schmidt, M. Matera, A. Bellucci (Eds.), End-User Development, Springer Nature Switzerland, Cham, 2025, pp. 161–170.

[26] H. J. Flint, K. P. Scott, P. Louis, S. H. Duncan, The role of the gut microbiota in nutrition and health, Nature Reviews Gastroenterology & Hepatology 9 (2012) 577–589.

[27] O. S. Palsson, W. E. Whitehead, M. A. L. van Tilburg, L. Chang, W. Chey, M. D. Crowell, L. Keefer, A. J. Lembo, H. P. Parkman, S. S. C. Rao, A. Sperber, B. Spiegel, J. Tack, S. Vanner, L. S. Walker, P. Whorwell, Y. Yang, Development and validation of the rome IV diagnostic questionnaire for adults, Gastroenterology 150 (2016) 1481–1491.

[28] S. J. Lewis, K. W. Heaton, Stool form scale as a useful guide to intestinal transit time, Scandinavian Journal of Gastroenterology 32 (1997) 920–924. URL: https://doi.org/10.3109/00365529709011203. doi:10.3109/00365529709011203, pMID: 9299672.

[29] European Parliament and Council of the European Union, Regulation (EU) 2017/746 on in vitro diagnostic medical devices, 2017. URL: https://europa.eu, oJ L 117, 5.5.2017, p. 176–332.

[30] European Data Protection Supervisor, AI Act Regulation (EU) 2024/1689 – Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU)

2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance), Publications Office of the European Union, 2025. doi:doi/10.2804/4225375.

[31] Q. Cui, Y. Xia, Q. Wu, Q. Chang, K. Niu, Y. Zhao, Validity of the food frequency questionnaire for adults in nutritional epidemiological studies: A systematic review and meta-analysis, Critical Reviews in Food Science and Nutrition 63 (2023) 1670–1688. URL: https://doi.org/10.1080/10408398.2021.1966737. doi:10.1080/10408398.2021.1966737, pMID: 34520300.

[32] C. L. Craig, A. L. Marshall, M. Sjöström, A. E. Bauman, M. L. Booth, B. E. Ainsworth, M. Pratt, U. Ekelund, A. Yngve, J. F. Sallis, P. Oja, International physical activity questionnaire: 12-country reliability and validity, Medicine and Science in Sports and Exercise 35 (2003) 1381–1395. doi:10.1249/01.MSS.0000078924.61453.FB.

[33] European Parliament, Council of the European Union, Regulation (EU) 2016/679 of the European Parliament and of the Council, 2016. URL: https://data.europa.eu/eli/reg/2016/679/oj.