REPLY: Generative AI-Powered Virtual Assistant for Guideline-Directed Medical Therapy Optimization

We read with interest the letter by Dr Bareeqa and colleagues. The criticisms raised regarding our paper are addressed below.

First, concerns regarding the absence of technical architecture, validation workflow, and safety framework are unfounded and can be found in the *JACC: Advances* Methods Companion,[1] which describes these elements comprehensively. That paper explicitly details the system architecture (retrieval-augmented generation [RAG]-based, guideline-anchored large language model), the multistage validation process based on real-world cases with iterative expert refinement, and the embedded safety logic with red flags and physician oversight. Second, the claim of inadequate predeployment validation is inaccurate. The Methods Companion outlines a structured validation process involving hundreds of real-world cases, comparison with cardiologist judgment, and iterative refinements to ensure guideline-aligned performance, demonstrating rigorous preclinical evaluation.

Third, the critique of RAG is generic and not applicable to the system under study. SIRIO-HF is a constrained, guideline-grounded clinical decision support system, with structured inputs and safeguards, so applying limitations of unconstrained RAG systems reflects a misunderstanding of the architecture.

Fourth, the authors have misinterpreted the trial. ASSIST-HF SIRIO did not evaluate "AI vs clinicians," but rather a supervised decision support tool with mandatory human-in-the-loop governance, in which every recommendation generated by the virtual assistant (VA) was reviewed, modified if needed, and ultimately authorized by a cardiologist, who retained full decision-making authority.

Fifth, the link between AI use and increased visit frequency is misinterpreted. More frequent contact, enabled by the AI model to address specialist capacity limits, was an intentional part of the care design, not a confounder. The study evaluated a scalable, supervised workflow combining frequent monitoring and AI support, rather than isolating AI's independent effect.

Finally, the statement regarding "zero disagreements" is wrong. The *JACC* Brief Report[2] clearly states that cardiologist-VA agreement on guideline-directed medical therapy optimization was 100% in the early phase of the trial and 93.4% at the end of treatment. Disagreements were few and mainly reflected perceived "overcaution" in VA recommendation, such as extreme caution in the presence of borderline potassium values. The Methods Companion paper[1] explicitly states that all VA-clinician disagreements were prospectively logged and reviewed within a structured governance framework, with continuous monitoring and embedded disagreement analysis for safety.

The concerns of Bareeqa et al can be addressed by a complete consideration of the Methods Companion and an understanding of the supervised, workflow-based intervention. ASSIST-HF SIRIO does not assess autonomous AI, but a physician-supervised model designed to reduce specialist input, enable more frequent visits, and improve scalability and standardization of care.

Eliano Pio Navarese, MD, PhD[a,b,c]
Joshua H. Leader, MB ChB, BSc[d,e]
Sogol Koolaji, MD, MPH[d,f]
Dean J. Kereiakes, MD[g]
***Diana A. Gorog, MD, PhD**[d,e,f,h,i]

*Faculty of Medicine
National Heart and Lung Institute
Imperial College
Dovehouse Street
London SW3 6LY, United Kingdom
E-mail: d.gorog@imperial.ac.uk

From the [a]Link Campus University, Rome, Italy; [b]SIRIO MEDICINE Research Network, Bydgoszcz, Poland; [c]Nicolaus Copernicus University, Bydgoszcz, Poland; [d]East and North Hertfordshire NHS Trust, Hertfordshire, United Kingdom; [e]National Heart and Lung Institute, Imperial College, London, United Kingdom; [f]School of Life and Medical Sciences, Postgraduate Medical School, University of Hertfordshire, Hertfordshire, United Kingdom; [g]The Christ Hospital and Lindner Research Center, Cincinnati, Ohio, USA; [h]School of Cardiovascular and Metabolic Medicine & Sciences, Kings College London, London, United Kingdom; and the [i]Royal Brompton and Harefield Hospitals, Guy's and St. Thomas' NHS Foundation Trust, London, United Kingdom.

The authors have reported that they have no relationships relevant to the contents of this paper to disclose.
The author attests they are in compliance with human studies committees and animal welfare regulations of the author's institution and Food and Drug Administration guidelines, including patient consent where appropriate. For more information, visit the Author Center.

REFERENCES

1. Navarese EP, Leader JH, Markides RIL, et al. Design and architecture of a generative-AI-supported, nonphysician-delivered model for GDMT optimization in HFrEF: the ASSIST-HF trial. *JACC Adv*. 2026;5:102588.

2. Navarese EP, Leader JH, Markides RIL, et al. AI-guided GDMT optimization after HFrEF hospitalization: the ASSIST-HF SIRIO randomized pilot trial. *J Am Coll Cardiol*. 2026. https://doi.org/10.1016/j.jacc.2025.12.066. Epub ahead of print.