Qeios

Peer Review

Review of: "A Claim-Governed Framework for Conservative Molecular Candidate Prioritization in Schizophrenia"

André Mateus Bertolino[1]

1. Instituto Federal de Educação Ciência e Tecnologia do Paraná, Curitiba, Brazil

This manuscript presents a timely and potentially valuable framework for addressing a recurrent problem in psychiatric genomics: the inappropriate escalation of heterogeneous evidence into claims of causality, biomarker status, therapeutic relevance, or biological validation. The distinction between genetic association, transcriptomic context, regulatory annotation, literature co-mention, and experimental validation is scientifically important. The proposed claim ladder and “claim firewall” are clear, intuitive, and potentially useful for researchers seeking a more transparent and conservative approach to molecular candidate prioritization.

A particular strength of the study is its explicit recognition that negative or inconclusive findings can still be scientifically informative. The decision not to promote any candidate to L2, and to retain SNAP91, OLR1, and GRM1 only within limited interpretive roles, demonstrates commendable restraint. The manuscript is also generally well-organized, clearly written, and appropriately cautious about the limitations of bulk transcriptomics, literature metadata, cross-disorder recurrence, and circuit-level annotations.

However, the principal limitation of the current version is that the reported audit cannot yet be independently evaluated. The manuscript states that 406 candidates were processed, five reached L1, and none reached L2, but it does not provide the complete candidate matrix, route-specific thresholds, source-level evidence, query logs, code, or freeze manifest. Without these materials, readers cannot determine how the initial universe was constructed, why individual candidates passed or failed, or

 doi.org/10.32388/M78X83

whether the classifications can be reproduced. These resources should therefore be deposited in a persistent public repository and linked directly in the manuscript.

The operational definition of the evidence ladder also requires further development. The authors should specify what constitutes an “admissible evidence route,” how independence between routes is established, how conflicting or missing evidence is handled, and which thresholds distinguish L0, L1, and L2. The five L1 candidates should all be identified and their evidence profiles reported, even if only three are discussed in detail. Candidate-level explanations for SNAP91, OLR1, and GRM1 would also help demonstrate how the framework functions in practice rather than only at the conceptual level.

The framework would be further strengthened through validation using prespecified positive and negative controls, sensitivity analyses under alternative thresholds, and comparison with at least one established prioritization strategy. Such analyses would help determine whether the conservative design effectively limits inferential inflation without systematically excluding well-supported candidates. The expression “negative by design” should also be clarified, since the absence of L2 candidates is scientifically meaningful only if the thresholds were specified independently of the observed outcome.

Overall, this is a thoughtful and promising methodological contribution that addresses an important weakness in computational neuropsychiatric research. Its conceptual foundation deserves consideration, and the author’s restraint in interpreting weak evidence is a notable strength. I encourage a major revision focused on operational detail, complete reproducibility, and empirical validation. With these additions, the framework could become a useful and transferable tool for improving transparency and claim discipline in molecular candidate-prioritization studies.

Declarations

Potential competing interests: No potential competing interests to declare.