

# Toward Continuous Neurocognitive Monitoring: Integrating Speech AI with Relational Graph Transformers for Rare Neurological Diseases

RAQUEL NOREL, IBM Research, USA

MICHELE MERLER, IBM Research, USA

PAVITRA MODI, IBM Software, Canada and University of New Brunswick, Canada

Patients with rare neurological diseases report cognitive symptoms—“brain fog”—invisible to traditional tests. We propose continuous neurocognitive monitoring via smartphone speech analysis integrated with Relational Graph Transformer (RELGT) architectures. Proof-of-concept in phenylketonuria (PKU) shows speech-derived “Proficiency in Verbal Discourse” correlates with blood phenylalanine ( $\rho = -0.50$ ,  $p < 0.005$ ) but not standard cognitive tests (all  $|r| < 0.35$ ). RELGT could overcome information bottlenecks in heterogeneous medical data (speech, labs, assessments), enabling predictive alerts weeks before decompensation. Key challenges: multi-disease validation, clinical workflow integration, equitable multilingual deployment. Success would transform episodic neurology into continuous personalized monitoring for millions globally.

Additional Key Words and Phrases: Digital biomarkers, speech analysis, rare diseases, graph transformers, phenylketonuria

## 1 The Monitoring Gap

Patients with rare neurological diseases consistently report cognitive difficulties undetectable by clinical assessments [2]. In phenylketonuria (PKU), adults describe “brain fog” and working memory deficits [7] yet score normally on neuropsychological tests [10]. Metabolic decompensation progresses undetected for weeks between quarterly blood tests.

A key challenge in PKU research is how to bridge subjective patient experience with objective, actionable biomarkers. Traditional assessment methods fail short in several ways: episodic clinic visits miss the everyday fluctuations that patient reports, standard neurocognitive tests lack ecological validity and do not capture how PKU affects real-world functioning, and the data that is collected often remains siloed across different formats and sources. This makes it difficult for clinicians to understand how cognitive load, mood, fatigue, and metabolic changes interact over time, and prevents the development of continuous, meaningful indicators of a patient’s true neurological state.

Natural language processing quantifies spontaneous speech as a neurocognitive biomarker [4, 9]; Relational Graph Transformers (RELGT) can integrate heterogeneous medical data at scale [3]. **We envision smartphone-based speech analysis integrated with medical databases via RELGT, enabling continuous neurocognitive monitoring—transforming reactive episodic care into proactive precision neurology. Figure 1 illustrates the envisioned schematic shift.**

---

Authors’ Contact Information: Raquel Norel, IBM Research, Yorktown Heights, NY, USA, [RNorel@gmail.com]; Michele Merler, IBM Research, Yorktown Heights, NY, USA, [mimerler@us.ibm.com]; Pavitra Modi, IBM Software, [Markham], Canada and University of New Brunswick, Fredericton, NB, Canada, [Pavitra.Modi@ibm.com/pmodi@unb.ca].



Fig. 1. Contrasting Traditional and Digital Approaches to Neurocognitive Monitoring. Panel A illustrates limitations of traditional episodic assessment, clinic-based testing, and siloed data sources that result in reactive care. Panel B presents an envisioned solution using smartphone-based speech analysis integrated with medical databases via RELGT, enabling continuous monitoring and proactive precision neurology through real-time data capture in remote, natural environments.

## 2 Vision: Integrated Framework

**Component 1: Speech Biomarkers.** Spontaneous speech engages executive control, semantic retrieval, working memory, and pragmatic language—vulnerable in frontal-subcortical dysfunction [1]. AI analysis of 60-second narratives yields features capturing discourse complexity, coherence, and syntax. Smartphone-based capture occurs continuously in natural environments.

**Component 2: RELGT Architecture.** Medical data is mostly stored and represented as relational entities: patients→tests→treatments→symptoms. RELGT [3] can handle heterogeneous graphs where nodes encode data types (speech, labs, medications) and edges capture relationships (temporal, causal). Unlike GNNs that limited by information bottlenecks, RELGT employs hybrid attention, enabling direct integration across multiple hops.

**Component 3: Predictive Alerts.** Learning nonlinear cross-modal relationships identifies baseline deviations and generates risk-stratified alerts. Example: Week 3 speech decline triggers an assessment revealing metabolic elevation before scheduled Week 8 blood test — intervention 4+ weeks earlier than with traditional approaches.

**Proof-of-Concept:** PKU analysis ( $n = 42$  vs 41 controls) [11] identified 23 linguistic features aggregated into “Proficiency in Verbal Discourse” that:

- Correlated with phenylalanine ( $\rho = -0.50$ ,  $p < 0.005$ ) and tyrosine ( $\rho = 0.44$ )
- Showed zero correlation with WAIS-IV (all  $r < 0.17$ ,  $p > 0.1$ )
- Captured complexity, detail, coherence, context, emotion (ChatGPT-4 analysis)

Despite a 40% scoring from speech biomarkers that was clinically significant for working memory deficits and a 45% reporting for neurocognitive burden, standard tests administered to the same individuals resulted in “normal” outcomes. Continuous speech analysis has therefore capability to bridge lived experience and assessment measurements.

### 3 Key Research Challenges

#### 3.1 Challenge 1: Multi-Disease Validation

A major challenge lies in ensuring that speech-based models generalize across disorders with fundamentally different pathophysiology. Parkinson’s disease involves hypophonia and speech fluctuations tied to medication timing. Huntington’s disease reflects CAG-repeat-driven degeneration and progressive motor-cognitive decline. Wilson’s disease presents with dysarthria linked to copper accumulation. Each condition affects speech through different mechanisms, yet all share speech as an accessible, real-time window into brain function. Research must explore cross-disease learning to determine which speech patterns are disease-specific, which are shared across conditions, and how to build models that remain robust despite this biological diversity. Success can be achieved by the development of biomarkers with high correlation levels across diseases, in end-to-end speech based models with high precision across conditions, or ideally both.

#### 3.2 Challenge 2: Scalable Integration

Another key challenge consists in integrating millions of heterogeneous speech and clinical records while accounting for temporal evolution over extended time periods (potentially decades) and missing/incomplete data. Large-scale modeling requires architectures that can handle millions of nodes, irregular sampling, and long-term patient trajectories without sacrificing interpretability. Research directions include developing hybrid attention mechanisms optimized for massive graphs, learning individualized temporal trajectories, applying federated learning to protect patient privacy, and designing explainable attention weights that reveal which signals drive predictions. Scalable, transparent systems will be essential for deploying speech-based monitoring across real-world clinical populations.

#### 3.3 Challenge 3: Clinical Workflow

Digital biomarkers only succeed when they fit naturally into clinical routines and do not add burden for patients or clinicians. Many promising technologies fail at this step, not because they lack accuracy, but because they disrupt workflows or generate unmanageable alert volumes. Research must focus on seamless integration with clinical systems through standards like FHIR, the design of risk-adaptive sampling strategies that minimize patient effort, careful calibration of alert thresholds to avoid overload, and studies that evaluate clinician trust and usability. Ensuring low-friction adoption is essential for turning speech-based biomarkers into real clinical tools.

#### 3.4 Challenge 4: Health Equity

The final challenge is ensuring that speech-based biomarkers do not widen existing health disparities. Models trained on narrow populations risk failing for speakers of different languages, dialects, or with limited access to technology. Research must therefore prioritize multilingual datasets spanning several languages and domain adaptation [6]/transfer learning [12] techniques for underrepresented languages, develop alternative interfaces

for users with speech or motor limitations, support on-device processing for those with limited connectivity, and conduct continuous bias auditing to identify and correct systematic errors. Equity-focused design is essential for ensuring that these tools benefit all patients, not just those represented in early datasets.

#### 4 Towards Precision Neurology

Solving the above challenges would shift neurological care from episodic snapshots to continuous monitoring, from population norms to individualized trajectories, and from reactive responses to predictive insight. Over 50 million individuals are affected by neurodegenerative diseases [5]; hundreds of thousands have neurometabolic conditions [8], yet many of their symptoms remain invisible to current monitoring tools.

Speech provides a natural window into brain functions through non-intrusive devices that patients already naturally interact with every day. With the convergence of advanced models such as relational graph transformers, widespread smartphone access, and increasingly rich medical databases, this vision could become achievable. **We call on the community to tackle those challenges—transforming how we monitor, predict, and preserve brain health for millions globally.**

#### References

- [1] Michael P. Alexander. 2006. Impairments of procedures for implementing complex language are due to disruption of frontal attention processes. *Journal of the International Neuropsychological Society* 12, 2 (2006), 236–247. doi:10.1017/S1355617706060309
- [2] Deborah A. Bilder, J. Kay Noel, Erin R. Baker, William Irish, Yao Chen, Minna J. Merilainen, Satish Prasad, and Barbara J. Winslow. 2016. Systematic review and meta-analysis of neuropsychiatric symptoms and executive functioning in adults with phenylketonuria. *Developmental Neuropsychology* 41, 4 (2016), 245–260. doi:10.1080/87565641.2016.1243109
- [3] Vijay Prakash Dwivedi, Sri Jaladi, Yangyi Shen, Federico López, Charilaos I. Kanatsoulis, Rishi Puri, Matthias Fey, and Jure Leskovec. 2025. Relational Graph Transformer for relational databases. arXiv:2505.10960 [cs.LG] <https://arxiv.org/abs/2505.10960>
- [4] Elif Eyigöz, Sachin Mathur, Mar Santamaria, Guillermo Cecchi, and Melissa Naylor. 2020. Linguistic markers predict onset of Alzheimer’s disease. *EClinicalMedicine* 28 (2020), 100583. doi:10.1016/j.eclinm.2020.100583
- [5] Valery L. Feigin, Emma Nichols, Tahiya Alam, Marlena S. Bannick, Ettore Beghi, Natacha Blake, Charles L. Culpepper, E. Ray Dorsey, Alexis Elbaz, Richard G. Ellenbogen, et al. 2019. Global, regional, and national burden of neurological disorders, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016. *The Lancet Neurology* 18, 5 (2019), 459–480. doi:10.1016/S1474-4422(18)30499-X
- [6] Hao Guan and Mingxia Liu. 2022. Domain Adaptation for Medical Image Analysis: A Survey. *IEEE Transactions on Biomedical Engineering* 69, 3 (2022), 1173–1185. doi:10.1109/TBME.2021.3117407
- [7] Laura McWhirter, H. Smyth, I. Hoeritzauer, A. Couturier, J. Stone, and A. Carson. 2023. What is brain fog? *Journal of Neurology, Neurosurgery & Psychiatry* 94, 4 (2023), 321–325. doi:10.1136/jnnp-2022-329683
- [8] S. Nguengang Wakap, D.M. Lambert, A. Olry, C. Rodwell, C. Gueydan, V. Lanneau, D. Murphy, Y. Le Cam, and A. Rath. 2020. Estimating cumulative point prevalence of rare diseases: analysis of the Orphanet database. *European Journal of Human Genetics* 28, 2 (2020), 165–173. doi:10.1038/s41431-019-0508-0
- [9] Raquel Norel, Carla Agurto, S. Heisig, J.J. Rice, and G. Cecchi. 2020. Speech-based characterization of dopamine replacement therapy in people with Parkinson’s disease. *npj Parkinson’s Disease* 6, 1 (2020), 12. doi:10.1038/s41531-020-0113-5
- [10] Cristina Romani, Andrew Olson, Lynne Aitkenhead, Lucy Baker, Dhanesha Patel, Francjan J. van Spronsen, Anita MacDonald, An-nemiek M.J. van Wegberg, and Stephan C.J. Huijbregts. 2022. Meta-analyses of cognitive functions in early-treated adults with phenylketonuria. *Neuroscience & Biobehavioral Reviews* 143 (2022), 104925. doi:10.1016/j.neubiorev.2022.104925
- [11] Susan E. Waisbren, Raquel Norel, Carla Agurto, Shifali Singh, Zoe A. Connor, Marina G. Ebrahim, and Guillermo A. Cecchi. 2025. Beyond neuropsychological tests: AI speech analysis in PKU. *Journal of Inherited Metabolic Disease* 48, 1 (2025), e12831. doi:10.1002/jimd.12831
- [12] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. 2021. A Comprehensive Survey on Transfer Learning. *Proc. IEEE* 109, 1 (2021), 43–76. doi:10.1109/JPROC.2020.3004555