Integrating behavioural health tracking in human genetics research

Nelson B. Freimer[1,*], **David C. Mohr**[2]

[1]Depression Grand Challenge, Center for Neurobehavioral Genetics, Jane and Terry Semel Institute for Neuroscience and Human Behavior, UCLA, Los Angeles, CA, USA.

[2]Center for Behavioral Intervention Technologies, Northwestern University, Chicago, IL, USA.

Abstract

Internet-connected devices could transform our understanding of the causes of behavioural variation and its impact on health and disease, in particular for neuropsychiatric disorders.

Genome-wide association studies (GWAS) have to date identified hundreds of loci that contribute to the risk of developing major psychiatric disorders[1], a record of success comparable to that of other common diseases. Virtually all these associations are to categorical diagnoses, formed based on symptoms elicited by self-report questionnaires or standardized interviews. In other areas of medicine, by contrast, disease-associated loci are complemented by an even greater number of associations to disease-related quantitative phenotypes (endophenotypes), measured on a large scale over decades. For example, associations between serum lipoprotein levels and both common and rare genetic variants have added enormously to our understanding of the biology underlying cardiovascular and metabolic disorders[2]. These associations could be identified because lipid phenotypes are available for millions of individuals, a fact due to both the longstanding recognition that these phenotypes are mechanistically related to disease risk and the ease and low cost of obtaining them uniformly on a large scale.

The value of data sets that attempt to comprehensively represent endophenotypes across multiple domains is increasingly evident, as genetics studies begin to leverage the vast amount of information generated longitudinally on the scale of entire health systems. The explosion of GWAS discoveries from the UK Biobank illustrates this development, which has also fostered entirely new paradigms for genetic discovery; for example, the phenome-wide association study (PheWAS), the systematic search of phenotype databases for measures associated with a given variant. So far, however, this transformation has largely bypassed psychiatry, which currently lacks large cohorts phenotyped for quantitative traits. This lack reflects primarily the difficulty, given available assessment tools, in obtaining

meaningful quantitative assays of human behaviour that are both scalable to genetic studies and relevant to clinical disorders. As with analyses of psychiatric diagnoses, studies of quantitative behavioural traits still rely heavily on subjective and unreliable participant self-report measures or on assessments with limited scalability, such as those from neuroimaging.

Behavioural health tracking

Networked sensors in smartphones and other devices used by billions of people worldwide are embedded in the fabric of our lives; a wide range of large-scale and objective behavioural data are thereby transmitted passively, effortlessly and continuously[3].

Combining behavioural health tracking with geno-typing or sequencing efforts of even a fraction of the population using connected devices could enable genetic discovery research in human behavioural traits on a previously unimaginable scale. It will be feasible to elucidate the genetic architecture of longitudinal phenotypes across multiple domains including sleep, circadian rhythms, activity, social behaviour, attention and cognition, each of which is fundamental to our understanding of neuropsychiatric disorders (Supplementary Figure 1). This step will advance the field in two main ways. First, it will enable better specification of existing diagnostic categories. Second, it will foster the delineation and measurement of endophenotypes on a phenome-wide scale. The anticipated use of behavioural tracking to assess sleep-related phenotypes illustrates both opportunities (Supplementary Box 1).

Implementing large-scale phenotyping studies based on sensor-derived data rests on three main steps: data collection; deriving quantitative features from sensor signals; and the analyses of candidate endophenotypes. First, the number of devices from which large volumes of data can be acquired nearly instantaneously and continuously from multiple sensors is increasingly rapidly. The most ubiquitous sensor data relevant to behaviour derive from smartphones, which detect location (GPS), accelerometry, light and sound, as well as text-based and image-based information from Internet activities such as social media usage. A variety of databases (for example, weather or map sites) provide environmental information which can be integrated with and enhance the interpretability of behavioural data. Wearables (for example, watches) are used by a smaller but rapidly growing proportion of the population and target specific behaviours (such as activity) as well as physiological parameters. Second, both unsupervised and supervised machine learning methods can be used to develop features from sensor data that can be employed to construct phenotypes. Unsupervised methods identify potentially meaningful patterns in the data. For example, data from GPS sensors can be clustered to identify features that are associated with clinical syndromes, such as the number of locations a person has visited over a given interval, the time spent in each location and the circadian rhythm of movement, all of which are correlated with major depression[3]. Supervised methods, by contrast, require labelling of those data. For example, by periodically launching a brief question on an individual's phone it may be possible to ascertain an aspect of their mental state (such as mood) in relation to GPS-derived features, such as time spent in given locations. Once developed, the algorithms for collecting such information can be run on people's devices with minimal user input.

Third, given the vast number of potential phenotypes that the sensor data will provide, efficient means of prioritizing between measures will be essential. As most human traits are heritable, applying heritability analyses to sensor data will help differentiate signal from noise; inferred phenotypes that represent measurement noise will not be heritable. A critical step in the identification of the most useful endophenotypes from behavioural health tracking will be the demonstration of the relationship between these measures and known clinically relevant end points. It remains unclear how best to perform this validation, and the development of new machine learning methods for this purpose is an area of intense research activity.

Phenotyping studies based on behavioural health tracking will not only have the advantages of scalability and objectivity as noted above, but also permit analyses that are not currently feasible. Most importantly, implementation of this strategy will enable the field to reincorporate the delineation of longitudinal trajectories as a core element in efforts to understand psychopathology. During the early 20th century, when most patients suffering from severe mental illness experienced long-term hospitalization, the observation of their different trajectories became the foundation of the diagnostic classification systems for these disorders that remains predominant. In particular, the division of psychotic disorders into episodic and chronic forms led to the delineation of schizophrenia and bipolar disorder as distinct syndromes[4].

Recent large-scale studies have highlighted overlapping genetic contributions to these disorders[1], but rely primarily on cross-sectional assessments. Behavioural health tracking will enable longitudinal assessment of phenotypic domains that may be critical for characterizing different disease trajectories and, therefore, for genetic analyses aimed at identifying biologically meaningful subtypes of disorders or predicting clinical outcomes.

Behavioural health tracking will also greatly accelerate PheWAS discoveries, by generating quantitative phenotypes across multiple domains, assessed uniformly in large populations. These data will be particularly valuable in health systems, where the phenotype search also includes wide-ranging electronic health record information. Finally, implementation of behavioural health tracking may accelerate the translation of genetic discoveries into both mechanistic understanding and new treatments. It can provide the first suite of phenotypes model systems — a comparability that derives in large part from the concomitant development of continuous, that are directly comparable between humans and animal automated tracking of behaviours in animal studies[5].

Remaining challenges

Behavioural health tracking has so far achieved few replicated results, as most studies have been too small to overcome the heterogeneity of the principal data; across multiple such studies (typically with <50 participants), the signal is swamped by the many possible sources of variability in sensor data. This variability stems in part from interindividual differences in how people use devices, many of which are irrelevant to the problems the field aims to address; patterns of phone use may vary based on a person's age, or whether they live in an urban or rural environment, or a northern versus southern climate. There is also

variability associated with the sensors themselves; different phones have different sensors and restrict access to sensor data in different ways. Finally, the substance and structure of the data will evolve as devices, sensors and user interaction patterns with their devices change.

While larger study sizes will lead to the development of algorithms that more dependably generate reliable phenotypes from sensor data, the impact of behavioural health tracking in elucidating the genetic architecture of behavioural traits will depend even more on the potential for data sharing on the largest possible scale. For such sharing to occur, the field must develop standards to increase the comparability of data across different sensors, platforms for obtaining behavioural features from sensors, algorithms for deriving endophenotypes from these features and common strategies for managing privacy and security. Roadmaps for developing standards for digital clinical research, broadly, are already available, for example those proposed by the Clinical Trials Transformation Initiative (CTTI). Given the rapid proliferation of proprietary approaches, realizing the potential of sensor-based technologies for understanding human behaviour will require the field to coalesce swiftly to outline and adopt standards.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

References

1. Sullivan PF et al. Psychiatric genomics: an update and an agenda. Am. J. Psychiatry 175, 15–27 (2018). [PubMed: 28969442]

2. Khera AV & Kathiresan S Genetics of coronary artery disease: discovery, biology and clinical translation. Nat. Rev. Genet 18, 331–344 (2017). [PubMed: 28286336]

3. Mohr DC, Zhang M & Schueller SM Personal sensing: understanding mental health using ubiquitous sensors and machine learning. Annu. Rev. Clin. Psychol 13, 23–47 (2017). [PubMed: 28375728]

4. Goodwin FK & Jamison KR Manic–Depressive Illness: Bipolar Disorders and Recurrent Depression 2nd edn (Oxford Univ. Press, 2007).

5. Hong W et al. Automated measurement of mouse social behaviors using depth sensing, video tracking, and machine learning. Proc. Natl Acad. Sci. USA 112, E5351–E5360 (2015). [PubMed: 26354123]