

PHENOMICS: THE SYSTEMATIC STUDY OF PHENOTYPES ON A GENOME-WIDE SCALE

R. M. BILDER,^{a,b,e,*} F. W. SABB,^{a,e} T. D. CANNON,^{a,b,e}
E. D. LONDON,^{a,c,e} J. D. JENTSCH,^{b,e}
D. STOTT PARKER,^d R. A. POLDRACK,^{b,e} C. EVANS^{a,e}
AND N. B. FREIMER^{a,e}

^aDepartment of Psychiatry and Biobehavioral Sciences, David Geffen School of Medicine at University of California, Los Angeles

^bDepartment of Psychology, University of California, Los Angeles College of Letters and Science

^cDepartment of Pharmacology, David Geffen School of Medicine at University of California, Los Angeles

^dDepartment of Computer Science, University of California, Los Angeles Henry Samueli School of Engineering and Applied Science

^eJane and Terry Semel Institute for Neuroscience and Human Behavior, Room C8–849, 740 Westwood Plaza, Los Angeles, CA 90095, USA

Abstract—Phenomics is an emerging transdiscipline dedicated to the systematic study of phenotypes on a genome-wide scale. New methods for high-throughput genotyping have changed the priority for biomedical research to phenotyping, but the human phenome is vast and its dimensionality remains unknown. Phenomics research strategies capable of linking genetic variation to public health concerns need to prioritize development of mechanistic frameworks that relate neural systems functioning to human behavior. New approaches to phenotype definition will benefit from crossing neuropsychiatric syndromal boundaries, and defining phenotypic features across multiple levels of expression from proteome to syndrome. The demand for high throughput phenotyping may stimulate a migration from conventional laboratory to web-based assessment of behavior, and this offers the promise of *dynamic phenotyping*—the iterative refinement of phenotype assays based on prior genotype-phenotype associations. Phenotypes that can be studied across species may provide greatest traction, particularly given rapid development in transgenic modeling. Phenomics research demands vertically integrated research teams, novel analytic strategies and informatics infrastructure to help manage complexity. The Consortium for Neuropsychiatric Phenomics at UCLA has been supported by the National Institutes of Health Roadmap Initiative to illustrate these principles, and is developing applications that may help investigators assemble, visualize, and ultimately test multi-level phenomics hypotheses. As the transdiscipline of phenomics matures, and work is extended to large-scale international collaborations, there is promise that systematic new knowledge bases will help fulfill the

*Correspondence to: R. M. Bilder, Jane and Terry Semel Institute for Neuroscience and Human Behavior, Room C8–849, 740 Westwood Plaza, Los Angeles, CA 90095, USA. Tel: +1-310-825-9474; fax: +1-310-825-2850.

E-mail address: rbilder@mednet.ucla.edu (R. M. Bilder).

Abbreviations: ADHD, attention deficit/hyperactivity disorder; CNP, Consortium for Neuropsychiatric Phenomics; *Gpr6*, G protein-coupled receptor 6 gene; GWAS, genome wide association studies; LA2K, study of 2000 community volunteers in Los Angeles; SNP, single nucleotide polymorphism.

0306-4522/09 \$ - see front matter © 2009 IBRO. Published by Elsevier Ltd. All rights reserved.
doi:10.1016/j.neuroscience.2009.01.027

promise of personalized medicine and the rational diagnosis and treatment of neuropsychiatric syndromes. © 2009 IBRO. Published by Elsevier Ltd. All rights reserved.

Key words: phenotype, genetics, genomics, informatics, cognition, psychiatry.

The nominal completion of the human genome project has fostered enormous enthusiasm and carried with it the promise of biomedical breakthroughs and a new era of personalized medicine in which genetic profiles serve as bases for rational diagnosis and treatment. Almost a decade into the post-genomic era, there is an emerging consensus that some of the problems are even harder than originally feared, and nowhere is this sense of sobriety clearer than in neuropsychiatric research. David Goldstein recently remarked: “There is absolutely no question that for the whole hope of personalized medicine, the news has been just about as bleak as it could be” (Wade, 2008). Others have expressed more optimism (Maher et al., 2008) based on recent reports of positive findings from genome-wide association studies (GWAS) of schizophrenia and bipolar disorder (Ferreira et al., 2008; O’Donovan et al., 2008). But there remain obvious challenges both in identifying relevant genetic variants, and once these variants are identified, determining what roles these variants may play in neuropsychiatric illness.

Freimer and Sabatti (2003) suggested that “The Human Phenome Project” is now an imperative to follow on and fulfill the promise of the human genome project. Given the continued decreases in cost and increasing availability of high-throughput genotyping platforms, it is today even clearer that phenotyping comprises the key rate- and cost-limiting factor in human genetics. Beyond the time and cost needed for phenotyping, the field faces an even grander challenge: what phenotypes should we be studying? The human “phenome,” which Freimer and Sabatti referred to as the “. . . manifold human phenotypes from molecule to mind. . .,” is a big place. The human genome, with only 3 billion bases, selected from a pool of only four nucleic acids, organized in a neat one-dimensional sequence, pales in comparison to the human phenome, which contains an unknown number of elements, many of which are characterized by enormous inter-individual variation that is at best only partially understood, and for which the dimensionality remains unknown. Phenomics—operationally defined as the systematic study of phenotypes on a genome-wide scale—is critically important to provide traction for biomedical advances in the post-genomic era (Bilder, 2008).

“Systematic study” implies not only the assessment of phenotypes in a well-operationalized fashion, but further that the strategy for studying phenotypes is defined in relevant biological contexts, with appreciation that any given study’s capacity to sample the human phenome is limited to a minute fraction of the phenomic state-space (all the possible kinds of phenotypes and their individual variants). “On a genome-wide scale” implies that examination of any specific phenotype involves considering the likely complex genetic contributions to that manifestation from the entire genome, and acknowledges that many of the putative “candidate genes” suggested by prior research may turn out to be false-positive signals.

In practical terms, phenomics demands broad scientific expertise in genetics, molecular biology, cell biology, systems biology, and higher levels of phenotypic expression, and these experts must be capable of communicating effectively and collaboratively designing and executing translational research projects on a large scale. For neuropsychiatric phenomics, the systems-level experts must include experts with knowledge about neural systems, cognitive systems, and neuropsychiatric symptoms and syndromes. Further expertise in mathematical modeling, statistics, and information sciences is needed to confront the many novel data analytic challenges, and enable the definition, visualization, and testing of complex multi-level hypotheses. Phenomics is therefore best defined as a “transdiscipline,” as it engages contributing disciplines interactively and synthetically to generate a new, emergent discipline.

NEUROPSYCHIATRIC PHENOMICS: REFRAMING MODELS OF STRUCTURE-FUNCTION RELATIONS

Discerning the genetic contributions to any complex human illness is fraught with challenges and demands broad phenomics scope, but neuropsychiatric genetics confronts unique obstacles. Among the most intriguing of the distinctions involves the mind-brain problem itself. While we might consider the emergent properties of any organ to involve qualitatively similar unknowns, the emergence of complex human behavior from brain function remains the epitome of scientific challenges. For example, while cardiovascular function involves myriad unresolved mysteries, already multiscale bioengineering models have been developed for the heart at molecular genetic, cellular, tissue, and whole organ levels, enabling prediction of whole organ consequences of genetic variation (Crampin et al., 2004). Similar efforts have helped model pulmonary function, kidney function and more (see for example, the International Union of Physiological Sciences Physiome Project; Hunter, 2004; Hunter and Borg, 2003). In contrast, there remain major gaps in explanations of how psychological processes relate to brain function, and this comprises a non-trivial limitation to modeling the biological bases of neuropsychiatric syndromes. Contemporary theorists have written extensively about the likely form of this relation, and increasingly call attention to the *emergent* properties of mental states that

are not simply reducible to the interactions of more elementary component processes (Tononi and Edelman, 1998; Edelman, 2003; Seth et al., 2006). Kendler (2008) has written eloquently about the need for iterative evolution of our concepts about structure-function relations, and highlights that our desideratum is more likely an *implementation* rather than a reductionist *replacement* solution. While these theoretical models are appealing, mechanistic models ultimately will require more detail about exactly what emerges from what, and will likely require some unexpected reframing of current concepts about psychological functions, to effectively advance phenomics research.

It is further critical to recognize that previous methods for carving brain function “at its joints” are likely to have generated divisions that align poorly with those that can be genetically determined. Much of what we have learned about brain-behavior relations comes from models that are of questionable value to understanding the genetic roots of brain function and dysfunction. To take an extreme example, the discipline of clinical neuropsychology has learned much from the study of individuals with brain injury, stroke, or neoplasm, but it is obvious that genetic variation is extremely unlikely to have similar effects on brain function. Thus the conventional catalogs of neuropsychological “domains” (i.e. language, memory, executive functions and so on) should not be expected to align well with the cognitive functions affected by genetic variation.

A more subtle point is that there is not a compelling reason to expect that decomposition of behavior, based on analyses within the behavioral level alone, will provide a useful target for genetic analysis. It might be hoped, for example, that the “normal” factorial structure of cognitive abilities revealed through decades of cognitive testing will be reflected in genetic variants. We can certainly seek genetic correlates of such constructs as “intelligence” or “memory” (see Sabb et al., 2009), and since most of the constructs have reasonably high heritability, it is clear that these are related to genetic variation (or epigenetic variation). It remains unclear to what degree any of the relations identified will be specific. The “generalist gene hypothesis” is based on findings that a single set of genes affects most cognitive abilities. Indeed the genetic correlations (i.e. the extent to which genetic effects on one trait are correlated with genetic effects on another trait) are approximately 0.80 among diverse verbal, spatial, memory, mathematical and reading abilities (Butcher et al., 2006; Kovas and Plomin, 2006). But how many genes must be involved to explain the heritability of these traits, which is generally estimated to be approximately 40%? Initial work using multi-stage genome-wide scanning revealed four candidate single nucleotide polymorphisms (SNPs), each of which explained only ~0.2% of phenotypic variance, and together these four SNPs accounted for only ~0.8% of variance in the cognitive phenotypes (Butcher et al., 2005). Given that these specific analyses compared extreme groups (those high and low in mental ability, thus perhaps not reflecting fully the quantitative trait distribution), and used what are now outmoded genotyping methods (they

used DNA pooling and only examined 10,000 SNPs), these estimates of shared variance might be considered pessimistic. But so far there is little evidence from more recent GWAS of substantially more robust genetic associations with similar phenotypes. This might suggest that at least several hundred genes will ultimately be associated with cognitive ability in general, and that the majority of these will account for less than 1% of phenotypic variance.

Similar high genetic correlations may be found for a range of other neuropsychiatric syndromes (Kendler et al., 2003), and help to explain the high “co-morbidity” of these syndromes. The challenge this poses for gene discovery at the syndromal level may be further exacerbated by the application of a taxonomic classification system that has reified invalid distinctions. It is already recognized that few of the diagnostic categories in the Diagnostic and Statistical Manual for Mental Disorders, 4th Edition (DSM-IV) represent valid classes even when examined in terms of the symptoms used to define the syndromes (Haslam and Kim, 2002; Haslam, 2003), and it is hoped that the Diagnostic and Statistical Manual for Mental Disorders, 5th Edition (DSM-V) may include indices of the continuous trait dimensions that will better reflect symptomatic variation (Kraemer, 2007). Even if some syndromal classes are verified to represent true “taxa” or “latent classes” based on the covariance structure of the symptoms used to define the syndromes, this is no guarantee that these syndromes will have a clearer genetic basis. For example, there is evidence that a “melancholic” subtype of depression may represent a valid taxon distinct from other forms of depression and from healthy groups (Haslam and Beck, 1994; Ambrosini et al., 2002), yet more detailed epidemiological data offer less support for the distinctiveness of this subgroup (Kessing, 2007), and assessment of familial risk reveals a pattern more consistent with a quantitative trait than genetically distinctive subgroup (Kendler, 1997). Thus empiric evidence suggesting discrete syndromal classes may be helpful, but is not sufficient to assert a clearer genetic basis.

The high genetic correlations for behavioral phenotypes may be an expected consequence of attempting to identify relations across biological scales that involve emergent functional consequences. Given that the level of gene effects must be translated through gene expression of proteins, to the functional roles played by those proteins in cellular systems and signaling pathways, to the functioning of those cells and signaling pathways in integrated neural circuits, and then ultimately make the brain-to-behavior traversal, there is no shortage of opportunities for specific effects to be obscured. Fisher (2006) has commented eloquently that “. . . the ability to undertake genetic analyses while employing only the most basic abstract concept of ‘the gene’, and without any understanding of molecular pathways, has become both a blessing and a curse, particularly in studies of the brain.” The phenomics strategy explicitly calls for efforts to redefine phenotypes as multi-level combinations of measures that may offer more realistic constraints on the mechanistic paths leading from genome to syndrome.

PHENOTYPE PRIORITIZATION FOR PHENOMICS RESEARCH

The phenomics strategy can be seen as an extension of the endophenotype approach that embraces multi-level modeling. This strategy acknowledges that many of the putative endophenotypes or intermediate phenotypes being investigated in biological psychiatry today may not possess much simpler genetic architecture than do the highest level syndromal phenotypes (Flint and Munafò, 2007). While this strategy does not replace many elements of previously suggested qualities important for phenotype prioritization (Gottesman and Gould, 2003; Bearden and Freimer, 2006), it shifts the emphases (see Table 1).

For example, heritability remains important, but the suggestion that a valuable endophenotype should associate with illness and show familial cosegregation with illness may be less critical, depending on how “illness” is defined. The critical point is that the most valuable phenotypes for phenomics research may well be those that cross boundaries of the current diagnostic taxonomy. Indeed, if it is true that substantial genetic correlations exist for neuropsychiatric syndromes, it may be most fruitful to study phenotypes that are shared across multiple diagnostic syndromes, and that by limiting research to more narrowly defined diagnostic groups, the genetic signals from the strongest genetic contributions may be lost amid other less important diagnosis-specific “noise.” It is further unclear how helpful conventional heritability statistics are to prioritizing phenotypes, given that many well-defined behavioral traits possess heritabilities exceeding 40% (Sabb et al., 2008), and higher heritability does not seem to assure a simpler genetic architecture. Since the principal value of identifying high heritability is to assure that the phenotypic trait is likely to be meaningfully related to genetic variation (i.e. that it possesses some evidence of *genetic validity*), it may also be valuable in phenotype selection to have evidence that the phenotype is significantly associated with known genetic variants. Further, to leverage the power of GWAS using quantitative trait loci, it is helpful if the genetic variant is common in human populations (although it should be recognized that with the advent of ever-increased capacity for genotyping, it may soon be possible for phenotypes associated with rare variants to be detected). It is also helpful if the genetic variant is associated with known functional effects at translational or transcriptional levels, in order to foster greater traction in molecular biology research.

Among the desiderata for phenotype selection, it is difficult to overemphasize the importance of fundamental *measurement properties*. Given that the upper limit on validity will be imposed by reliability, it is critical for research to focus on those measures that show adequate internal consistency (helping assure that a meaningfully coherent construct is being measured in the first place). If that is true, it is further important to identify those that are relatively stable over time and organismic states, or if there is some fluctuation that this fluctuation itself is a part of the phenotypic assay. The importance of phenotypic stability

Table 1. Considerations for phenotype selection in phenomics research

Genetic validity
(a) Moderately high heritability; association with related disorders in family members
(b) Relate to known common functional genetic variation
Measurement properties
(a) Reliability (nomothetic validity, internal consistency); test–retest reliability (at least within a particular state, and preferably across states in illnesses which have an episodic pattern); and other desirable psychometric properties (e.g. discriminating power across a broad range of individual differences)
(b) Have good concurrent validity (convergent and divergent validity) with respect to hypothesized endophenotypes
Relevance for human applications
(a) Related to higher level constructs in humans, enabling links to clinical effectiveness and outcomes
(b) Suitable for application in clinical trials
(c) Associated with neuropsychiatric morbidity
(d) Cross-disorders relevance—the phenotype construct has been identified as important in more than one diagnostic syndrome
Relevance for translational investigation
(a) Homologies of expression across species enabling both basic and clinical investigation
(b) Latent endophenotype constructs relate to neural systems models sufficiently well studied that physiologically plausible manipulations can be effected and measured using currently available techniques
(c) Implicated neural systems already are targets of known partially effective treatments for existing disorders, or these systems are the targets of new chemical entities
Knowledge of environmental effects
(a) Known environmental effects on phenotype expression can usefully constrain modeling of phenotypic variance explained by genotypic variance, and may be particularly valuable in model organisms where environments can be strictly controlled (epigenetic mechanisms are seen as an area of rapid expansion in knowledge and systems in which these can be studied effectively may be prioritized)

has been emphasized by others; in brief it is logical to propose that our fixed genetic inheritance will be most easily associated with features that are stable. But there also may be important phenotypic features that have instability as their signature, in which case it is precisely this variability that needs to be assessed. Finally, it is valuable to assure that the sensitivity of phenotype measurement is strong across widely varying levels of phenotype expression. For example, a cognitive phenotype is likely best measured by a test that shows strong measurement properties at both higher and lower levels of that ability, and if the measurement is biased towards either higher or lower levels of ability, then meaningful genetic effects may never be detected not because these do not exist but rather because the test is too insensitive.

A critical aspect of understanding the measurement properties of specific phenotypes is determining how amenable these are to *high throughput phenotyping*. Given the changing finances of genomics research, it has become clear that the rate-limiting step in advancing knowledge about human disease has shifted dramatically from genetics to phenomics. Phenotyping now far exceeds the time and financial costs of genotyping. Until “high throughput phenotyping” methods are developed, accrual of knowledge in phenomics will be thwarted. Development of detailed costing models is advocated to identify the most cost-effective methods for addressing key phenomic targets. High throughput collection of behavioral data is virtually impossible using traditional laboratory-based methods. For example, a “comprehensive” examination of cognitive abilities may require 10 to 20 h, and briefer assessments routinely fail to provide broad phenotypic coverage, or fail to specify adequately the phenotypes of interest, or both.

There are two major hopes in advancing high throughput phenotyping. One hope is based on leveraging ad-

vances in modern psychometric theory. Using item response theory and computerized adaptive testing algorithms, it is now routinely possible to increase efficiency in psychological assessment by a factor of two. Thus one can either specify the construct with double the precision, or decrease testing time for the same measurement precision in half the time of the original test. There are multiple challenges in using these methods, including: (a) the construct itself needs to be well operationalized in advance; and (b) relatively large samples (i.e. 500 or more with complete response sets) are suggested to apply these methods. A second hope is based on the incredible growth in use of the Internet, with more than 100 million people in the United States using the web daily. If we can provide well-structured phenotyping tools, enabling widespread access for individuals to provide meaningful data about themselves, the results could rapidly revolutionize behavioral genomics. A paradigm shift that could accelerate discovery would involve *dynamic phenotyping*, by which we mean the iterative refinement of phenotype assays based on prior genotype-phenotype associations. There remains considerable skepticism among scientists about both sampling biases and validity of data collected using web-based rather than conventional laboratory-based methods, despite demonstrations to the contrary (Krantz et al., 1997; Buchanan and Smith, 1999; Krantz and Dalal, 2000; Andersson et al., 2006; Chiasson et al., 2006; Cunningham et al., 2006; Graham et al., 2006). While much work remains to address these concerns in a compelling manner, the potential increases in throughput may dramatically outweigh, and even provide solutions to the sampling and quality control issues. Given that conventional laboratory studies running for 5 years often have difficulty ascertaining and examining a few thousand individuals, the opportunities for sub-sampling and data cleaning with several hundred thousand individuals may appear increas-

ingly attractive, particularly as we increasingly recognize the value of samples with tens of thousands of individuals.

There are some features of phenotypes that may be less important to advance genetics research, but may be prioritized because of their *relevance to other human applications*. Biomedical research ultimately seeks to alleviate human suffering, and thus those phenotypes that can be meaningfully related to clinical morbidity, or the clinical effectiveness and outcomes of treatments, may be important targets for study. Those phenotypes that are suitable for application in clinical trials may be particularly valuable given the possible opportunity to link phenotypic variation to putative effects at the levels of cellular systems or signaling pathways. Consistent with and extending the points made earlier about the value of cross-disorder phenotypes, those phenotypes that are relevant across multiple categories of neuropsychiatric illness may have even broader public health significance than those that do not. For example, finding the genetic bases of “normal” variation in memory or attentional control may be of greater value than identifying genetic associations with Alzheimer’s disease or attention deficit/hyperactivity disorder (ADHD).

The phenomics strategy also prioritizes phenotypes that are *relevant for translational investigation*, and thus those phenotypes that cross species boundaries. This is not intended to discount the importance of uniquely human phenotypes, and focusing on these is an alternate approach with great potential but different goals. For example, it may be possible to gain traction on the genetic bases of language by examining differences in gene expression between human and chimpanzee (Oldham and Geschwind, 2006). But the phenomics approach suggests that the “low hanging fruit” in neuropsychiatric genetics may be picked most readily by focusing on phenotypes that are generally well conserved phylogenetically. The primary rationale for this is that by examining phylogenetically conserved phenotypes, we increase the opportunity to gain traction on intervening biology using basic science models.

Of particular value today are transgenic models that can help illuminate the cellular systems and signaling pathways affected by genetic variants, and therefore a priority may be given to phenotypes that can be studied in both mouse and human. For phenotypes that have reasonable homologies from mouse to man, there is a high likelihood that there will be even more robust homologies in larger rodents and non-human primates, enabling more detailed analysis of higher level neuropsychopharmacological effects and neural systems level phenotypes.

Linked to prioritization of cross-species phenotypes is the idea that phenotypes selected for study may be most useful if embedded in plausible mechanistic hypotheses. The better fleshed out the biological mechanisms, the higher the likelihood that meaningful connections will be established with other biological knowledge. Thus, other factors being equal, a phenotype for which there exists already a set of mechanistic models may be preferable to one that does not possess similar evidence. While all such

mechanistic models are so far incomplete, working within a framework that includes relevant empiric science may help link findings to existing evidence and either further develop more effective models, or help prune these and better specify superior models for the future.

A critical point to consider in any phenotype prioritization effort is that examining “criteria” does not enable generation of some figure-of-merit for phenotype selection. In the Consortium for Neuropsychiatric Phenomics (CNP), we had started with the aim of generating a phenotype selection algorithm, and on more careful consideration recognized that at best it might be possible to generate a phenotype “profiling tool” capable of characterizing the strengths and weaknesses of a given phenotype. Some of the features noted above are likely to conflict directly with others (for example, validation with respect to molecular targets may run counter to validation with respect to outcomes).

There is further a substantial risk of phenotype reification that can actively interfere with discovery. For example, a phenotype prioritized because it is influenced reliably by a known molecular entity (a drug) might be seen as valuable by virtue of links to signaling pathways in a mechanistic model, but this might be misleading if the mechanistic model is not well understood. A case in point is the phenotype involving induction of catalepsy in rodents, used for decades as a screening test for antipsychotic drug development. This may have canalized drug discovery towards agents more likely to produce extrapyramidal symptoms than antipsychotic efficacy.

Finally, the phenomics strategy emphasizes that phenotype selection should weigh cautiously evidence from candidate gene strategies. There may be value in research that uses a “bottom-up” approach, starting with functional genetic variation, through identification of proteins affected, through identification of signaling and other cellular processes affected, through neural systems function to behavior. But there should also be appropriate concern that this approach can undermine discovery of novel gene-phenome correlations. It may be an impediment to discovery to be tied to known genomic “hot spots,” many of which were identified using phenotypes that are acknowledged to be suboptimal, and also may be subject to various other methodological problems (including population stratification and linkage disequilibrium) that may have led to false-positive identification of regions that are distant from the driving functional genomic regions of greater interest.

MANAGING COMPLEXITY IN THE HUMAN PHENOME

To facilitate this work, we have adopted a simplified seven-layer schema to reflect some of the key traversals across levels of inquiry and biological scales that are important in the complete representation of a phenomics hypothesis from genome to syndrome (Fig. 1). While experts representing a specific disciplinary perspective may argue coherently about the validity of these “levels,” we have found this helpful to foster transdisciplinary communication. The

Fig. 1. Seven-level schematic illustrating relevant components of mechanistic “genome to syndrome” hypotheses for phenomics research.

bottom few layers reflect basic biological “dogma” that is generic to all phenomics sub-disciplines (i.e. that genes “code” for proteins, that proteins exert their biological effects via incorporation in cellular systems and signaling pathways). Above this level, the schema is customized for neuropsychiatric research, by focusing on the assembly of cellular systems and signaling systems into neural systems, suggesting that activity within neural systems leads to the emergence of cognitive functions (which we define broadly to encompass a complete range of mental faculties and operations, including mood, affective, mnemonic, linguistic, and thinking abilities), and that these cognitive phenotypes form the basis of clinically observable symptoms, which in turn are used to define neuropsychiatric syndromes.

Given this simple framework, some crude calculations may offer perspective on the volume of the phenomics search space. Acknowledging that this schema is oversimplified reinforces the assertion that resulting computations are probably underestimates. But assuming only seven transformations across levels are needed to generate a plausible genome-to-syndrome hypothesis, we can estimate the impact of both pleiotropy and polygenicity on phenomics research programs. Consider the case of a modest fivefold expansion across levels (i.e. a given gene influences five proteins, a given protein influences five cellular systems and signaling pathways, and so on). Ascending through seven levels would generate 15,625 effects at the “syndromal” level. A 10-fold expansion at each level would generate more than a million effects after promulgation across seven levels. Similarly, if we imagine the spectrum of plausible polygenic contributions to high level phenotypic traits (i.e. that a syndrome is identified by multiple symptoms, that each symptom reflects variation in multiple cognitive phenotypes, which are in turn the product of action across multiple neural systems, and so on), it is easy to imagine that hundreds if not thousands of genetic variations contribute to complex neuropsychiatric syndromal phenotypes.

Comparable complexity is revealed by posing the question slightly differently. Imagine our goal is to determine how much variance is shared between observations at each level of analysis. Assume further that the amount of variance shared is constant across levels. For a single gene to account for 25% of variance in a complex syndrome would require that a feature on each level must share at least 80% variance with a feature on the next level. This scenario is likely unrealistic, given existing data showing that this exceeds the reliability of virtually all of the higher phenotypic features, and that genotype may explain only ~20% of variance at the level of the transcript (Flint and Munafò, 2007). Still optimistic but possibly realistic is identifying between-level associations equivalent to 50% shared variance. Through seven levels, this would result in 1.6% shared variance between a specific genotype and syndromal phenotype (i.e. among the strongest results so far obtained in GWAS research). Best supported by data are relations across levels that show shared variance of approximately 20% (consistent both with the Flint and Munafò review, and further with typical correlations among various phenotypes in neuropsychiatric research in the range of 0.4 to 0.5). In this scenario, a genetic variant would share only 0.01% with a phenotypic variant, or in other words, some 5000 genetic variants would be needed to explain the overall heritability of approximately 50% that have been identified through family-based studies of high level personality traits and neuropsychiatric diagnostic phenotypes.

There are so far insufficient empiric data to draw firm conclusions on the actual complexity we face in neuropsychiatric genetics, but if the results of recent GWAS in both neuropsychiatry and other complex disorders are a useful guide, it appears that even the strongest signals being detected are explaining less than 1% of phenotypic variance, and thus small contributions from many genetic variants appear more the rule than the exception. As noted above, the findings of Plomin and his group that four

genetic variants together explained only 0.8% variance in intelligence, suggests that at least hundreds of genes are likely involved to explain heritability of 40%. Such estimates ignore the possibility that epigenetic factors account for some currently unknown proportion of this variance.

The promise of phenomics rests in part on the concept that mechanistic biological hypotheses linking genome to syndrome can help set constraints on these massive webs of associations. There is hope that elements of convergence, following self-organizing principles that have been discerned in other disciplines as diverse as chemistry and economics may be applicable also to the emergence of multiple aspects of gene expression and higher level phenotypes (Kauffman, 1993; Nykter et al., 2008). Considerable work already has advanced the idea that the emergent properties of cognition from neural systems activity may depend on self-organizing neural networks (Tononi et al., 1994, 1999; Sporns et al., 2005). Similar progress is being made applying information science and network principles to the study of gene networks and other biological systems (Sridhar et al., 2007; Zheng et al., 2007; Centler et al., 2008). A major question remains: at what level of a multi-level mechanistic network is it most fruitful to attempt identifying such phenotypic convergence? Is it most fruitful to focus on the transcriptome, the proteome, the signalome, or elsewhere?

NARROWING THE SEARCH SPACE FOR PHENOMICS RESEARCH: EXAMPLES FROM THE CNP

Within the CNP, which is one of nine Interdisciplinary Research Consortia supported by the National Institutes of Health Roadmap Initiative starting in 2007 (<http://nihroadmap.nih.gov/interdisciplinary/index.asp>), we have prioritized investigation of two cognitive phenotypes for translational research. We refer to these phenotypes as *response inhibition mechanisms*, and *memory mechanisms*. Phenotype selection benefited from considering the criteria outlined above, together with input from two teams that considered catalogs of phenotypes worthy of investigation both across psychiatric disorders (cross-disorders workgroup) and across different species (cross-species workgroup). In the course of considering a diversity of possible phenotypic targets, one important question emerged: At what level of phenotypic expression would it be best to focus attention?

The need to narrow focus in phenomics research is prompted by the sheer magnitude of possible associations as summarized above. A key tenet of the phenomics strategy is to maximize agnosticism in recognition of the many unknowns we face (including the possibility that many “candidate genes” reflect false-positive findings, and that many phenotypes are ill-defined). But given extensive pleiotropy together with the likelihood that neuropsychiatric traits are polygenic—perhaps massively polygenic—it is clearly implausible to conduct research that will provide an unbiased sample of the human phenome (i.e. phenome-wide association).

We determined that for the purposes of neuropsychiatric phenomics (i.e. establishing a research program dedicated to linking genomic to neuropsychiatric syndromal levels of description), a focus on *cognitive phenotypes* would be most fruitful. The key rationale for focusing on this level of analysis is that cognitive phenotypes may be meaningfully related to higher level symptoms and syndromes on the one hand, and to underlying neural systems activity on the other hand. In this way, cognitive phenotypes were seen as offering a potential mediating link from the study of behavioral phenotypes to their brain bases, and interposing a bottleneck to facilitate transdisciplinary research capable of marrying clinical to basic research in neuropsychiatry. Many cognitive phenotypes further manifest many of the desirable properties alluded to previously, including: (a) genetic validity (with heritability of most well-defined cognitive phenotypes approximating 50%); (b) reasonable measurement properties (with internal consistency indices approximating 0.90 and test-retest reliability of many measures approximating 0.80); (c) relevance for translational research (since for at least some cognitive phenotypes, there exist non-human homologs or analogs).

Among the critical linkages, cognitive phenotypes may offer valuable tools for bridging from human research to basic research via the neural systems level. Current neuroimaging methods particularly have enabled interrogation of neural systems activation effects associated with cognitive manipulations, thus forging a relational link that bridges observable behavior to underlying biology. Some imaging methods even permit insight into cellular and signaling correlates of cognitive phenomena. The capacity to elicit neural systems responses to cognitive probes in humans can be actively complemented by the study of similar cognitive processes in lower-level biological systems in non-human species. Thus we conceptualize the vertical research strategy as operating both top-down in humans, from the level of syndromes through symptoms and cognitive phenotypes to neural systems (complemented by GWAS), and bottom-up in mice, given the facility with which we can create new transgenic models and determine the effects of the genetic manipulations on molecular expression, cellular systems and signaling pathways, and neural systems.

The CNP’s selection of response inhibition and memory mechanisms as focal phenotypes was based on the criteria noted above, and our capacity to examine these in integrated top-down and bottom-up research, given the specific areas of expertise possessed by members of an extensive group of collaborators. The research team of the CNP includes more than 50 investigators representing dozens of disciplines including genetics, statistical genetics, molecular biology, neurobiology, systems neuroscience, cognitive neuroscience, neuropsychopharmacology, clinical neurosciences (including psychiatry, neurology, psychology, and neuropsychology), public health, statistics, psychometrics, and computer science. We have established research teams to address both the response inhibition and memory mechanisms themes using the top down approaches (human GWAS studies with approximately 10 h of cognitive

phenotyping per individual, and additional neuroimaging investigations in a subgroup). These are complemented by bottom-up studies in which we are creating transgenic models based on candidate genes of putative relevance to these themes. For example, one set of studies in the response inhibition theme focuses on a bacterial artificial chromosome transgenic model that permits cell-specific alteration in expression of the *G protein-coupled receptor 6* gene *Gpr6* within medium spiny neurons in the striatum. These studies are following up on exciting findings that *Gpr6*^{-/-} mice show increased initiation of responses under instrumental conditioning schedules, but decreased capacity to withhold the same responses under leaner schedules (Lobo et al., 2007). In the memory mechanisms theme, ongoing work is following up on findings examining transgenic manipulations of *Disc1* and *Dysbindin*, which have shown a range of cognitive, neurophysiological, and neuroanatomic deficits paralleling those observed in patients with neuropsychiatric illnesses such as schizophrenia and bipolar disorder (Cannon et al., 2005; Li et al., 2007).

Another unique feature of this work, dictated by the phenomics strategy, is that the GWAS is principally targeting a group of community volunteers rather than specific diagnostic groups. This approach is based on the idea that quantitative traits representing these phenotypes are likely to be most powerfully identified in broad population samples, rather than in groups selected for diagnostic phenotypes, which are biased towards the extremes of the relevant phenotype distributions. To the extent that individuals with neuropsychiatric disorders represent extreme values on these quantitative traits, the relevant phenotyping assays will show restricted range and thus reduced power to detect associations. The CNP is therefore focusing its GWAS assessments of response inhibition and memory phenotypes on 2000 community individuals (recruited from the Los Angeles metropolitan region, and thus referred to as the “LA2K” study), and will also study smaller samples of 100 individuals in each of three diagnostic groups (including people with diagnoses of schizophrenia, bipolar I disorder, and/or ADHD). The 300 patients are included primarily to characterize their performance on response inhibition and memory phenotype assays, rather than to help identify novel genetic associations. If genetic associations with the cognitive phenotypes are identified in the LA2K sample, we can then determine the extent to which these genetic variants occur in the patient groups and if these appear useful in explaining deviant scores. This strategy is based on the premise that the genetic bases for quantitative traits such as memory and response inhibition are likely to show strong overlap between individuals who have psychiatric diagnoses and those who do not. If this is true, findings of genetic association in the non-diagnosed group will be informative for the diagnosed group. If this premise is wrong, and the phenotypic deviation shown among diagnosed individuals is due to “syndrome-specific” genetic factors, then the information derived from the LA2K will be a helpful start on the path to determining the unique genetic sources of these effects.

DATA ANALYTIC STRATEGIES FOR PHENOMICS RESEARCH

The sheer scope of phenomics research in the GWAS era raises data analytic challenges on an unprecedented scale. It has become difficult even to represent, much less comprehend the hypotheses that we are beginning to interrogate. This complexity places a new burden on the field to develop data mining and informatics strategies that are capable of identifying meaningful associations across multiple levels, and involving hundreds of thousands of variables.

Given that analysis of genome-wide datasets has only been widespread over the last few years, there have been rapid advances in development of novel analytic methods. Most effort so far, however, has involved the analysis of an individual phenotype, usually a high level diagnostic phenotype, in a “case-control” study design. In these designs, the simplest possible analytic model involves direct comparison of allelic frequencies at each one of several hundred thousand (or more) SNPs, possibly supplemented by fine mapping studies or other approaches to replication and/or verification of candidate genetic signals. Given that this approach involves hundreds of thousands of statistical tests, control for false-positive results is of paramount importance, and a variety of methods to determine false discovery rates have been suggested (Kang and Zuo, 2007; Wakefield, 2008). The basic analytic strategy, however, has widely remained one of conducting many independent tests, while acknowledging that the degree of independence remains uncertain. The importance of moving beyond the “massively univariate” perspective to consider gene×gene and gene×environment interaction effects is difficult to overestimate, as is highlighted in recent work (Burdick et al., 2008). Furthermore, recent reports suggest that important phenotypic variation may be related more strongly to either copy number variations or rare variants that may be missed using currently standard SNP arrays (Lencz et al., 2007; Sebat et al., 2007).

Efforts have been made to reduce dimensionality on the genomic level through identification of haplotypes, singular value decomposition and independent component analysis of gene expression data, and other novel analytic strategies (Li et al., 2005; Yuan and Li, 2007; Biswas et al., 2008; Madi et al., 2008). Advances in these methods may particularly aid in analysis of gene×gene interaction effects, which may be of great importance for complex polygenic phenotypes (Musani et al., 2007). Similar efforts have been made to reduce dimensionality at the phenotype level, again using singular value decomposition methods, clustering methods, or both (e.g. factor mixture modeling) (Muthen et al., 2006). The proliferation of new methods to advance dimension reduction at genotypic and phenotypic levels is exciting and will likely be helpful to advance phenomics research. Additional methods will likely be important, however, that are capable of both multi-level modeling of phenotypic constructs (in contrast to most dimension reduction strategies that are applied only within a single level of phenotypic observation), and

that are capable of identifying the best combinations of genotype and phenotype that relate to each other.

“Genomic convergence” is a term coined to describe a multifactorial, multistep approach combining gene expression and genomic data to identify and prioritize targets (Hauser et al., 2003). We suggest that extending this strategy to include higher level phenomic data are a logical, but not trivial, next step. Similar conclusions have been drawn, and some interesting new approaches are emerging to enable superior multivariate, multilevel analysis and visualization of genotypic and phenotypic data together (Beyene et al., 2007). Given the challenges of reducing scope to aid the pragmatics of both comprehension and computation, methods are needed to efficiently prune multilevel models. Sparse canonical correlation analysis, so far applied to genomic and gene expression data (Parkhomenko et al., 2007), is appealing in its aim to maximize shared variance between genotype and phenotype, rather than maximize variance explained within each level considered separately and then examine associations between the factors that emerge, which has been the dominant model. Major questions remain about the most effective ways to prune such models; ultimately we believe biological hypotheses will enable rational constraints to be applied without jeopardizing discovery.

Implicit in this prioritization of multivariate gene-phenotype association is the likelihood that the traditional concepts of heritability may be usefully revised. Adopting the conventional definition of broad-sense heritability (H^2) as the ratio of genetic variance (G^2) to phenotypic variance (P^2), Wikipedia suggests that: “Estimating heritability is not a simple process, since only P can be observed or measured directly” (accessed 29 October 2008). While questions remain about how much genotypic variance can be observed or measured directly, current studies assessing hundreds of thousands of markers already may provide some estimates, and future methods will surely come closer. A component of the phenomics research program is to advance development of models that will estimate heritability directly from genomic data rather than from phenotype inheritance patterns.

*OMICS NEED INFORMATICS

Informatics strategies are helping address the new and profound challenge of identifying meaningful signals from GWAS. Many aspects of informatics development require *ontologies* that provide sufficient identification of the concepts under study, and their possible inter-relations. The Gene Ontology project (<http://www.geneontology.org>) is serving a useful role already by providing a coherent framework within which investigators can identify gene products by cellular component, biological process and molecular function. The “Entrez” system provided by the National Center for BiInformatics already contains, in addition to the widely used literature repository (PubMed) and a literature-based database containing annotations specifically relevant for functional genomics (Online Mendelian Inheritance in Man), a rich and steadily growing series of major

databases covering nucleotides (15 databases), proteins (six databases); structure (three databases); complete genomes (five databases), expression (three databases), and chemicals (three databases). Publications now routinely utilize these bioinformatics resources in creative ways to constrain analysis, identify replication signals, or interpret results. Despite the rapid growth and utility of these knowledge resources, representation of higher-level phenomic information remains at best incomplete.

The lack of informatics resources for higher-level phenotype data currently poses a limit on phenomics research. To address one segment of this large gap, we have been working on developing ontologies and annotation databases for neuropsychiatric phenomics research, and particularly for the domain of cognitive phenotypes, which has many unique challenges (Bilder et al., in press). This has involved the building of controlled vocabularies for cognitive concepts and tasks, symptoms, and syndromes. This work already has revealed some striking idiosyncrasies in patterns of term use in the biomedical literature. For example, a recent study aimed to identify published literature on heritability for the concept “cognitive control,” which has enjoyed dramatic increase in use over the last decade (Sabb et al., 2008). In identifying published work on “cognitive control” it was found, however, that this new construct was actually measured using exactly the *same* tests that had previously been used to index other constructs (“working memory,” “task switching/set shifting,” “response inhibition,” or “response selection”). This instability in the use of cognitive concept labels poses a challenge that needs to be addressed by anchoring cognitive concept use to the measurement level, much as the labeling of SNPs has benefited from adoption of the Reference SNP ID system (rs-numbers) that can be related to physical genomic sequence data. In contrast, gene naming conventions remain elusive and creativity abounds (Seringhaus et al., 2008).

Efforts to develop a useful taxonomy of cognitive tasks face challenges. Given that the goal of developing a task taxonomy is better anchoring and specification of the more fluid cognitive concepts, organizing cognitive concepts by “function” would involve circular logic and could be counterproductive, possibly leading to reification of inaccurate concepts. An alternative that may be less subject to bias is development of task taxonomies that are more like the familiar phylogenetic taxonomies of species. The analogies between task development and evolution of species may have additional merits. Most published tasks are clearly “descended” from “ancestors,” and often involve either (attempted) direct replication or subtle modification of the parent task. Less frequently, a task paradigm is changed substantially; we suggested that the term “task speciation” may help identify those events in task evolution marking differences significant enough that the results of the new task cannot be (or should not be) “mated” with those of the previously existing task to produce a “viable” meta-analytic result. Determinations about whether data should be pooled or not remain somewhat subjective, and different decisions may reflect differing goals of different

investigators. But this kind of specification would nevertheless mark a significant advance over the current lack of conventions, and could facilitate development of guidelines on those features of cognitive test paradigms which when changed should prompt consideration of task specification. Cognitive ontology development is considered in greater detail elsewhere (Bilder et al., *in press*), and is the primary aim of an ongoing project, which we hope will benefit from collaborative network inputs (see <http://www.cognitiveatlas.org>).

Assuming that it is possible to develop ontologies representing higher phenomic levels, considerable additional effort is necessary to develop the methods for representing the complex, multi-level hypotheses capable of relating genome to syndrome. The CNP includes a project—Hypothesis Web Development for Phenomics Research (RL1LM009833)—that aims to design tools to help develop, visualize, and share multilevel phenomics hypotheses. Currently scientific hypotheses tend to be articulated using a string of assertions that take the form: “X is associated with Y” [citation]; “Y is associated with Z” [citation]; “suggesting that X may influence Z, mediated by Y.” In addition to the logical leaps and failures to prove causality in such assertions, it is further conspicuous that citations are used to assert the “truth” of evidentiary components in these hypotheses, as though these were binary, when instead we are aware that the reported relations are only partially true, typically having exceeded some arbitrary threshold of significance. Given that the threshold for significance is usually based on disconfirming the null hypothesis of zero association, interpreting such findings as “truths” (implying 100% confidence) is at least partially misleading.

This highlights the potential value of conditioning assertions in terms of quantitative effect sizes. When two alternative hypotheses are presented as chains of assertions with binary “truth” values assigned to every link, there is no way to judge the relative merits of the hypotheses. In contrast, if links are supported by quantitative estimates documenting the strength of each association, it is possible to identify the stronger chain. Extending this rationale to even more complex, partially overlapping networks of associations, the effect size statistics can be used further to select (and prune) branches that will maximize overall

hypothesis strength. It is hoped that the “background and significance sections of the future” may replace current practice with graphical “hypothesis webs” supported by such quantitative, empiric annotation. These efforts may benefit from newly developed strategies for meta-analytic structural equation modeling, which enables analyses to be pooled at the level of study results, rather than at the individual subject level (Furlow and Beretvas, 2005; Riley et al., 2007; Cheung, 2008). While the same caveats apply to data pooling in these analyses as they do to any meta-analytic pooling of findings, the challenges are magnified as the complexity of models increases. Developing these methods offers the promise, however, that one day literature searching can be converted to knowledge mining with dynamic extraction of concept maps based on strength of association, and for discovery to be realized based on exploratory interrogation of maps with which investigators may previously have been unfamiliar.

Some tools are already freely available on-line to advance the hypothesis web concept (see Table 2). PubGraph (<http://www.pubgraph.org>) was developed to generate graphical concept webs where nodes in the graphs are concepts defined by PubMed queries, and edges connecting the nodes are annotated with Jaccard coefficients revealing the strength of association between concept pairs. PubAtlas (<http://www.pubatlas.org>) has further developed this literature querying facility to generate heat maps of concept co-occurrence matrices, and enable the historical changes in these maps to be plotted or played as movies, enabling researchers to visualize the dynamic growth of concept relations over time, and then select the relevant literature that underlies trends. PubAtlas further contains a series of pre-defined lexica (e.g. cognitive concepts, cognitive tasks, neuroanatomy), organized as series of PubMed queries, so that it is possible to “blast” any arbitrary search term against these lexica to find intersections with any term in the pre-defined set. PubBrain (<http://www.pubbrain.org>) takes concept-queries, intersects these with a modified version of the NeuroNames lexicon (neuroanatomic terms from the Foundational Model of Anatomy), and then projects the intersections on a three-dimensional probabilistic atlas of the brain (Shattuck et al., 2008). The Phenowiki (<http://www.phenowiki.org>) was developed to enable collaborative quantitative annotation of cognitive phenotype concepts (Sabb et al., 2008). The Phe-

Table 2. Web-applications developed to advance neuropsychiatric phenomics research

Application	URL	Description
PubGraph	http://www.pubgraph.org	Creates graphs of PubMed query results, where nodes are concepts defined by each query and edges are annotated with association statistics
PubBrain	http://www.pubbrain.org	Generates co-occurrence statistics of PubMed query with neuroanatomic lexicon, and projects weighted results on a 3D probabilistic atlas of the human brain
PubAtlas	http://www.pubatlas.org	Generates co-occurrence statistics among PubMed queries, generates heat maps and history plots, provides pre-defined lexica for “blast” queries relevant to neuropsychiatric phenomics research
Phenowiki	http://www.phenowiki.org	Provides Wiki-like text annotation for cognitive concepts and tests, and collaborative database entry of quantitative effect size, heritability, psychometric, and sample statistics
Cognitive Atlas	http://www.cognitiveatlas.org	Provides collaborative annotation for cognitive concepts and tests, including specification of assertions with evidence for and evidence against, and confidence voting (under development)

nowiki combines some features of a Wiki (text annotation about cognitive concepts and tasks), with additional features of a relational database. These include quantitative data fields for empiric information about group difference effect sizes, effect sizes for associations between pairs of variables, study sample sizes and some demographic characteristics, along with psychometric properties for psychological test variables. Each of these applications offers partial solutions to a very large problem; it is hoped that further development and integration of these and similar tools will one day enable integrated development, visualization, and testing of complex multilevel hypotheses.

CONCLUSIONS AND FUTURE DIRECTIONS

Phenomics is an emerging transdiscipline that aims to leverage breakthroughs in genome-wide genotyping, burgeoning knowledge in the clinical and cognitive neurosciences, and unimagined developments in information and computer sciences, to yield traction on biomedical problems of enormous complexity. The phenomics perspective suggests that systematic study of multiple phenotypes, across multiple biological scales, will be important even if ever-larger sample sizes succeed in revealing robust genetic associations with high level syndromal phenotypes. The need for this perspective seems particularly compelling in its application to neuropsychiatric syndromes. As we move into the “post-GWAS era,” it is likely that declining costs and new methods will enable ever-finer mapping of genetic sequences, detecting rare variants and copy number variations not revealed by most current platforms. While so far it remains unclear to what extent epigenetic factors may help account for phenotypic variance not explained by genomic data, the theoretical potential is vast (Mehler, 2008), and some major initiatives are aiming to develop epigenomics more fully (<http://nihroadmap.nih.gov/epigenomics/initiatives.asp>), making it possible that epigenome-wide analysis will become a staple of biological characterization. There remain major challenges in identifying and effectively studying human population samples sufficiently large, and capable of addressing challenges of genomic and phenomic diversity. Given the limited throughput of existing laboratory-based phenotyping methods, web-based ascertainment and phenotyping strategies may represent the most rational way forward, but considerable work is needed to translate what is now technically feasible into pragmatic deployments. While still at early stages in the Human Phenome Project, there remains great promise that phenomics can help us realize the vision of personalized medicine and rational neuropsychiatric diagnosis and treatment.

Acknowledgments—This work was supported by the Consortium for Neuropsychiatric Phenomics and its component awards (UL1DE019580, RL1MH083268, RL1MH083269, RL1DA024853, RL1MH083270, RL1LM009833, PL1MH083271, PL1NS062410), and the Cognitive Atlas (R01MH082795).

REFERENCES

- Ambrosini PJ, Bennett DS, Cleland CM, Haslam N (2002) Taxonicity of adolescent melancholia: a categorical or dimensional construct? *J Psychiatr Res* 36:247–256.
- Andersson G, Westoo J, Johansson L, Carlbring P (2006) Cognitive bias via the Internet: a comparison of web-based and standard emotional Stroop tasks in social phobia. *Cogn Behav Ther* 35: 55–62.
- Bearden C, Freimer N (2006) Endophenotypes for psychiatric disorders: ready for primetime? *Trends Genet* 22:306–313.
- Beyene J, Tritchler D, Bull SB, Cartier KC, Jonasdottir G, Kraja AT, Li N, Nock NL, Parkhomenko E, Rao JS, Stein CM, Sutradhar R, Waaijenborg S, Wang KS, Wang Y, Wolkow P (2007) Multivariate analysis of complex gene expression and clinical phenotypes with genetic marker data. *Genet Epidemiol* 31 (Suppl 1):S103–S109.
- Bilder RM (2008) Phenomics: building scaffolds for biological hypotheses in the post-genomic era. *Biol Psychiatry* 63:439–440.
- Bilder RM, Sabb FW, Parker DS, Kalar D, Chu WW, Fox J, Freimer NB, Poldrack RA (in press) Cognitive ontologies for neuropsychiatric phenomics research. *Cognitive Neuropsychiatry*, in press.
- Biswas S, Storey JD, Akey JM (2008) Mapping gene expression quantitative trait loci by singular value decomposition and independent component analysis. *BMC Bioinformatics*. doi:10.1186/1471-2105-9-244.
- Buchanan T, Smith JL (1999) Research on the Internet: validation of a worldwide web mediated personality scale. *Behav Res Methods Instrum Comput* 31:565–571.
- Burdick KE, Kamiya A, Hodgkinson CA, Lencz T, DeRosse P, Ishizuka K, Elashvili S, Arai H, Goldman D, Sawa A, Malhotra AK (2008) Elucidating the relationship between DISC1, NDEL1 and NDE1 and the risk for schizophrenia: evidence of epistasis and competitive binding. *Hum Mol Genet* 17:2462–2473.
- Butcher LM, Kennedy JK, Plomin R (2006) Generalist genes and cognitive neuroscience. *Curr Opin Neurobiol* 16:145–151.
- Butcher LM, Meaburn E, Knight J, Sham PC, Schalkwyk LC, Craig IW, Plomin R (2005) SNPs, microarrays and pooled DNA: identification of four loci associated with mild mental impairment in a sample of 6000 children. *Hum Mol Genet* 14:1315–1325.
- Cannon TD, Hennah W, van Erp TGM, Thompson PM, Lonnqvist J, Huttunen M, Gasperoni T, Tuulio-Henriksson A, Pirkola T, Toga AW, Kaprio J, Mazziotta J, Peltonen L (2005) DISC1/TRAX haplotypes associate with schizophrenia, reduced prefrontal gray matter, and impaired short- and long-term memory. *Arch Gen Psychiatry* 62:1205–1213.
- Centler F, Kaleta C, di Fenizio PS, Dittrich P (2008) Computing chemical organizations in biological networks. *Bioinformatics* 24:1611–1618.
- Cheung MW (2008) A model for integrating fixed-, random-, and mixed-effects meta-analyses into structural equation modeling. *Psychol Methods* 13:182–202.
- Chiasson MA, Parsons JT, Tesoriero JM, Carballo-Dieguez A, Hirshfield S, Remien RH (2006) HIV behavioral research online. *J Urban Health* 83:73–85.
- Crampin EJ, Halstead M, Hunter P, Nielsen P, Noble D, Smith N, Tawhai M (2004) Computational physiology and the physiome project. *Exp Physiol* 89:1–26.
- Cunningham JA, Selby P, van Mierlo T (2006) Integrated online services for smokers and drinkers? Use of the Check Your Drinking Assessment Screener by participants of the Stop Smoking Center. *Nicotine Tobacco Res* 8 (Suppl 1):S21–25.
- Edelman GM (2003) Naturalizing consciousness: a theoretical framework. *Proc Natl Acad Sci U S A* 100:5520–5524.
- Ferreira MA, O'Donovan MC, Meng YA, Jones IR, Ruderfer DM, Jones L, Fan J, Kirov G, Perlis RH, Green EK, Smoller JW, Gozeva D, Stone J, Nikolov I, Chambert K, Hamshere ML, Nimgaonkar VL, Moskvina V, Thase ME, Caesar S, Sachs GS, Franklin J, Gordon

- Smith K, Ardlie KG, Gabriel SB, Fraser C, Blumenstiel B, Defelice M, Breen G, Gill M, Morris DW, Elkin A, Muir WJ, McGhee KA, Williamson R, MacIntyre DJ, MacLean AW, St CD, Robinson M, Van Beck M, Pereira AC, Kandaswamy R, McQuillin A, Collier DA, Bass NJ, Young AH, Lawrence J, Ferrier IN, Anjorin A, Farmer A, Curtis D, Scolnick EM, McGuffin P, Daly MJ, Corvin AP, Holmans PA, Blackwood DH, Gurling HM, Owen MJ, Purcell SM, Sklar P, Craddock N (2008) Collaborative genome-wide association analysis supports a role for ANK3 and CACNA1C in bipolar disorder. *Nat Genet* 40:1056–1058.
- Fisher SE (2006) Tangled webs: tracing the connections between genes and cognition. *Cognition* 101:270–297.
- Flint J, Munafò MR (2007) The endophenotype concept in psychiatric genetics. *Psychol Med* 37:163–180.
- Freimer N, Sabatti C (2003) The human genome project. *Nat Genet* 34:15–21.
- Furlow CF, Beretvas SN (2005) Meta-analytic methods of pooling correlation matrices for structural equation modeling under different patterns of missing data. *Psychol Methods* 10:227–254.
- Gottesman II, Gould TD (2003) The endophenotype concept in psychiatry: etymology and strategic intentions. *Am J Psychiatry* 160:636–645.
- Graham AL, Bock BC, Cobb NK, Niaura R, Abrams DB (2006) Characteristics of smokers reached and recruited to an Internet smoking cessation trial: a case of denominators. *Nicotine Tobacco Res* 8 (Suppl 1):S43–S48.
- Haslam N (2003) Categorical versus dimensional models of mental disorder: the taxometric evidence. *Aust N Z J Psychiatry* 37: 696–704.
- Haslam N, Beck AT (1994) Subtyping major depression: a taxometric analysis. *J Abnorm Psychol* 103:686–692.
- Haslam N, Kim HC (2002) Categories and continua: a review of taxometric research. *Genet Soc Gen Psych Monogr* 128:271–320.
- Hauser MA, Li YJ, Takeuchi S, Walters R, Noureddine M, Maready M, Darden T, Hulette C, Martin E, Hauser E, Xu H, Schmechel D, Stenger JE, Dietrich F, Vance J (2003) Genomic convergence: identifying candidate genes for Parkinson's disease by combining serial analysis of gene expression and genetic linkage. *Hum Mol Genet* 12:671–677.
- Hunter PJ (2004) The IUPS Physiome Project: a framework for computational physiology. *Prog Biophys Mol Biol* 85:551–569.
- Hunter PJ, Borg TK (2003) Integration from proteins to organs: the Physiome Project. *Nat Rev Mol Cell Biol* 4:237–243.
- Kang G, Zuo Y (2007) Entropy-based joint analysis for two-stage genome-wide association studies. *J Hum Genet* 52:747–756.
- Kauffman SA (1993) *The origins of order: self organization and selection in evolution*. New York: Oxford University Press.
- Kendler KS (1997) The diagnostic validity of melancholic major depression in a population-based sample of female twins. *Arch Gen Psychiatry* 54:299–304.
- Kendler KS (2008) Explanatory models for psychiatric illness. *Am J Psychiatry* 165:695–702.
- Kendler KS, Prescott CA, Myers J, Neale MC (2003) The structure of genetic and environmental risk factors for common psychiatric and substance use disorders in men and women. *Arch Gen Psychiatry* 60:929–937.
- Kessing LV (2007) Epidemiology of subtypes of depression. *Acta Psychiatr Scand Suppl* (433)85–89.
- Kovas Y, Plomin R (2006) Generalist genes: implications for the cognitive sciences. *Trends Cogn Sci* 10:198–203.
- Kraemer HC (2007) DSM categories and dimensions in clinical and research contexts. *Int J Methods Psychiatr Res* 16:S8–S15.
- Krantz JH, Ballard J, Scher J (1997) Comparing the results of laboratory and worldwide web samples on the determinants of female attractiveness. *Behav Res Methods Instrum Comput* 29:264–269.
- Krantz JH, Dalal R (2000) Validity of web-based psychological research. In: *Psychological experiments on the Internet* (Birnbaum MH, ed), pp 35–60. New York: Academic Press.
- Lencz T, Morgan TV, Athanasiou M, Dain B, Reed CR, Kane JM, Kucherlapati R, Malhotra AK (2007) Converging evidence for a pseudoautosomal cytokine receptor gene locus in schizophrenia. *Mol Psychiatry* 12:572–580.
- Li L, Jiang W, Li X, Moser KL, Guo Z, Du L, Wang Q, Topol EJ, Rao S (2005) A robust hybrid between genetic algorithm and support vector machine for extracting an optimal feature gene subset. *Genomics* 85:16–23.
- Li W, Zhou Y, Jentsch JD, Brown RA, Tian X, Ehninger D, Hennes W, Peltonen L, Lonnqvist J, Huttunen MO, Kaprio J, Trachtenberg JT, Silva AJ, Cannon TD (2007) Specific developmental disruption of disrupted-in-schizophrenia-1 function results in schizophrenia-related phenotypes in mice. *Proc Natl Acad Sci U S A* 104: 18280–18285.
- Lobo MK, Cui Y, Ostlund SB, Balleine BW, Yang XW (2007) Genetic control of instrumental conditioning by striatopallidal neuron-specific S1P receptor Gpr6. *Nat Neurosci* 10:1395–1397.
- Madi A, Friedman Y, Roth D, Regev T, Bransburg-Zabary S, Ben JE (2008) Genome holography: deciphering function-form motifs from gene expression data. *PLoS ONE* 3:e2708.
- Maher BS, Riley BP, Kendler KS (2008) Psychiatric genetics gets a boost. *Nat Genet* 40:1042–1044.
- Mehler MF (2008) Epigenetic principles and mechanisms underlying nervous system functions in health and disease. *Prog Neurobiol*, in press.
- Musani SK, Shriner D, Liu N, Feng R, Coffey CS, Yi N, Tiwari HK, Allison DB (2007) Detection of gene×gene interactions in genome-wide association studies of human population data. *Hum Hered* 63:67–84.
- Muthen B, Asparouhov T, Rebollo I (2006) Advances in behavioral genetics modeling using Mplus: applications of factor mixture modeling to twin data. *Twin Res Hum Genet* 9:313–324.
- Nykter M, Price ND, Aldana M, Ramsey SA, Kauffman SA, Hood LE, Yli-Harja O, Shmulevich I (2008) Gene expression dynamics in the macrophage exhibit criticality. *Proc Natl Acad Sci U S A* 105: 1897–1900.
- O'Donovan MC, Craddock N, Norton N, Williams H, Peirce T, Moskvina V, Nikolov I, Hamshere M, Carroll L, Georgieva L, Dwyer S, Holmans P, Marchini JL, Spencer CC, Howie B, Leung HT, Hartmann AM, Moller HJ, Morris DW, Shi Y, Feng G, Hoffmann P, Propping P, Vasilescu C, Maier W, Rietschel M, Zammit S, Schumacher J, Quinn EM, Schulze TG, Williams NM, Giegling I, Iwata N, Ikeda M, Darvasi A, Shifman S, He L, Duan J, Sanders AR, Levinson DF, Gejman PV, Buccola NG, Mowry BJ, Freedman R, Amin F, Black DW, Silverman JM, Byerley WF, Cloninger CR, Cichon S, Nothen MM, Gill M, Corvin A, Rujescu D, Kirov G, Owen MJ (2008) Identification of loci associated with schizophrenia by genome-wide association and follow-up. *Nat Genet*, in press.
- Oldham MC, Geschwind DH (2006) Deconstructing language by comparative gene expression: from neurobiology to microarray. *Genes Brain Behav* 5 (Suppl 1):54–63.
- Parkhomenko E, Trichtler D, Beyene J (2007) Genome-wide sparse canonical correlation of gene expression with genotypes. *BMC Proc* 1:S119.
- Riley RD, Abrams KR, Sutton AJ, Lambert PC, Thompson JR (2007) Bivariate random-effects meta-analysis and the estimation of between-study correlation. *BMC Med Res Methodol* 7:3.
- Sabb FW, Bearden CE, Glahn DC, Parker DS, Freimer N, Bilder RM (2008) A collaborative knowledge base for cognitive phenomics. *Mol Psychiatry* 13:350–360.
- Sabb FW, Burggren AC, Higier R, Fox J, He J, Parker DS, Poldrack RA, Chu WW, Cannon TD, Freimer NB, Bilder RM (2009) Challenges in phenotype definition in the whole-genome era: genetic influences on memory and intelligence in the healthy population. *Neuroscience* 164:88–107.

- Sebat J, Lakshmi B, Malhotra D, Troge J, Lese-Martin C, Walsh T, Yamrom B, Yoon S, Krasnitz A, Kendall J, Leotta A, Pai D, Zhang R, Lee YH, Hicks J, Spence SJ, Lee AT, Puura K, Lehtimäki T, Ledbetter D, Gregersen PK, Bregman J, Sutcliffe JS, Jobanputra V, Chung W, Warburton D, King MC, Skuse D, Geschwind DH, Gilliam TC, Ye K, Wigler M (2007) Strong association of de novo copy number mutations with autism. *Science* 316:445–449.
- Seringhaus MR, Cayting PD, Gerstein MB (2008) Uncovering trends in gene naming. *Genome Biol* 9:401.
- Seth AK, Izhikevich E, Reeke GN, Edelman GM (2006) Theories and measures of consciousness: an extended framework. *Proc Natl Acad Sci U S A* 103:10799–10804.
- Shattuck DW, Mirza M, Adisetiyo V, Hojatkashani C, Salamon G, Narr KL, Poldrack RA, Bilder RM, Toga AW (2008) Construction of a 3D probabilistic atlas of human cortical structures. *Neuroimage* 39:1064–1080.
- Sporns O, Tononi G, Kötter R (2005) The human connectome: A structural description of the human brain. *PLoS Comput Biol* 1:e42.
- Sridhar S, Dhamdhere K, Blalock G, Halperin E, Ravi R, Schwartz R (2007) Algorithms for efficient near-perfect phylogenetic tree reconstruction in theory and practice. *IEEE/ACM Trans Comput Biol Bioinform* 4:561–571.
- Tononi G, Edelman GM (1998) Consciousness and complexity. *Science* 282:1846–1851.
- Tononi G, Sporns O, Edelman GM (1994) A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proc Natl Acad Sci U S A* 91:5033–5037.
- Tononi G, Sporns O, Edelman GM (1999) Measures of degeneracy and redundancy in biological networks. *Proc Natl Acad Sci U S A* 96:3257–3262.
- Wade N (2008) A dissenting voice as the genome is sifted to fight disease. *New York: New York Times*.
- Wakefield J (2008) Reporting and interpretation in genome-wide association studies. *Int J Epidemiol* 37:641–653.
- Yuan S, Li KC (2007) Context-dependent clustering for dynamic cellular state modeling of microarray gene expression. *Bioinformatics* 23:3039–3047.
- Zheng Y, Yeh CW, Yang CD, Jang SS, Chu IM (2007) On the local optimal solutions of metabolic regulatory networks using information guided genetic algorithm approach and clustering analysis. *J Biotechnol* 131:159–167.

(Accepted 14 January 2009)
(Available online 20 January 2009)