



Software Application Profile

Software Application Profile: PHESANT: a tool for performing automated phenome scans in UK Biobank

Louise AC Millard,^{1,2*} Neil M Davies,¹ Tom R Gaunt,¹
George Davey Smith¹ and Kate Tilling¹

¹MRC Integrative Epidemiology Unit (IEU), School of Social and Community Medicine and ²Intelligent Systems Laboratory, Department of Computer Science, University of Bristol, Bristol, UK

*Corresponding author. MRC Integrative Epidemiology Unit (IEU), School of Social and Community Medicine, Oakfield House, Oakfield Grove, University of Bristol, Bristol BS8 2BN, UK. E-mail: louise.millard@bristol.ac.uk

Editorial decision 22 August 2017; Accepted 30 August 2017

Abstract

Motivation: Epidemiological cohorts typically contain a diverse set of phenotypes such that automation of phenome scans is non-trivial, because they require highly heterogeneous models. For this reason, phenome scans have to date tended to use a smaller homogeneous set of phenotypes that can be analysed in a consistent fashion. We present PHESANT (PHEnome Scan ANALysis Tool), a software package for performing comprehensive phenome scans in UK Biobank.

General features: PHESANT tests the association of a specified trait with all continuous, integer and categorical variables in UK Biobank, or a specified subset. PHESANT uses a novel rule-based algorithm to determine how to appropriately test each trait, then performs the analyses and produces plots and summary tables.

Implementation: The PHESANT phenome scan is implemented in R. PHESANT includes a novel Javascript D3.js visualization and accompanying Java code that converts the phenome scan results to the required JavaScript Object Notation (JSON) format.

Availability: PHESANT is available on GitHub at [<https://github.com/MRCIEU/PHESANT>]. Git tag v0.5 corresponds to the version presented here.

Introduction

Phenome scans test the association of a trait of interest with a comprehensive array of phenotypes (the ‘phenome’). Types of phenome scans include phenome-wide association studies (pheWAS),¹ Mendelian randomization-pheWAS (MR-pheWAS)² and environment-wide association studies (EnWAS).^{3,4} PheWAS seek to investigate the association of a genetic variant with a set of phenotypic

traits.^{1,5} A recent extension to pheWAS, MR-pheWAS, uses Mendelian randomization (MR) in a pheWAS framework in order to search for the causal effects of a particular exposure.² EnWAS seek to test the associations of a trait of interest with a set of other phenotypes.³

Epidemiological cohorts usually contain a large number of diverse phenotypes, such that testing the association of these phenotypes with another trait in an automated way

is non-trivial. For this reason, researchers wishing to perform a phenome scan will typically specify a homogeneous subset of traits, in order to automate the tests of association across these traits in a consistent way. For instance, pheWAS initially started using international classification of disease (ICD) codes from electronic health records, where each disease code could be treated as a binary variable and a consistent test performed.⁵ However, restricting the set of phenotypes provides only a partial view of associations with a trait of interest, and reduces the potential to identify novel associations.

In this paper we present PHEsANT (PHENome Scan ANalysis Tool), a parallelizable tool for phenome scans in UK Biobank, a prospective cohort of over 500 000 men and women in the UK aged between 37 and 73 years.⁶ This cohort includes genetic data and a large and diverse range of data from blood, analyses of urine and saliva samples, clinical assessments, record linkage and health and lifestyle questionnaires. The diversity of traits available, coupled with the large sample size, provide an opportunity to identify novel associations with phenome scans.

Implementation

PHEsANT is implemented in R and requires the following R packages: `optparse`, `MASS`, `lmtest`, `nnet`, `forestplot` and `data.table` (see GitHub repository for package versions). PHEsANT takes one data file as input containing the set of phenotypes, the trait of interest [which may be a single nucleotide polymorphism (SNP), a genetic score or a phenotypic trait depending on whether a pheWAS, MR-pheWAS or EnWAS is being performed] and confounders (the trait of interest and confounders can alternatively be provided in separate files if this is preferred). PHEsANT also makes use of two data files (included in PHEsANT) that contain information about the variables in the UK Biobank cohort: (i) a data coding information file, and (ii) a variable information file. These files have been set up for the example we describe in the usage section, but can be changed as needed for each particular phenome scan. For more information on the PHEsANT data and information files, see the documentation in the GitHub repository. In the rest of this section we describe the variable processing flow used in PHEsANT.

Automated processing flow to determine variable coding

In order to test the association of the trait of interest with the diverse range of phenotypes in UK Biobank in an automated manner, we developed a rule-based system to determine the appropriate coding of each variable and hence

test of association to use. These rules are shown in [Figure 1](#) and described in full in [Supplementary section S1](#) (available as [Supplementary data](#) at *IJE* online). The decision rules start with the variable field type as specified by UK Biobank at: [<http://biobank.ctsu.ox.ac.uk/showcase/list.cgi>], either continuous, integer, categorical (single) or categorical (multiple), and categorize each variable as one of four data types: continuous, ordered categorical, unordered categorical and binary. The categorical (single) field type refers to categorical fields (including binary) where each participant can only have one value. For example, by questionnaire participants were asked ‘How would you describe your usual walking pace?’ with options including ‘slow’, ‘average’ and ‘brisk’ (field ID = 924; see [Supplementary Figure 1](#), available as [Supplementary data](#) at *IJE* online). In contrast, categorical (multiple) fields can have multiple values per participant. For example, by questionnaire participants were asked what types of bread they ate the previous day (field ID = 20091; see [Supplementary Figure 2](#), available as [Supplementary data](#) at *IJE* online) and could, for instance, select both white and wholemeal options. Where a field is measured at several time points we use the first occurrence only (see [Supplementary section S2](#) for details, available as [Supplementary data](#) at *IJE* online). Continuous and integer variables may have more than one measurement at this first measured time point (typically to improve the estimate of a measurement). For instance, spirometry was measured two or three times a few moments apart—see for example field 3062 [<http://biobank.ctsu.ox.ac.uk/showcase/field.cgi?id=3062>]). When this is the case we take the mean to create a single value per participant (see [Supplementary section S2](#) for details).

Variables with the continuous field type are usually assigned to the continuous data type. In this case, the variable is transformed to a normal distribution using an inverse normal rank transformation. In a minority of cases, continuous fields are assigned to the ordered categorical data type (or binary if there are only two distinct values). For example, field 100022 [<http://biobank.ctsu.ox.ac.uk/showcase/field.cgi?id=100022>] contains the estimated alcohol intake based on responses to the ‘diet by 24-hour recall’ questionnaire. A large proportion of the participants have a zero value for this field, because they consume no alcohol. It is not possible to inverse normal rank transform this variable because, where a large number of participants have the same value, the rank assigned in this transformation is random among these, and this would add noise to the data. Instead, we transform this variable into three categories with roughly the same number of participants in each (placing split points between distinct values), and treat it as an ordered categorical variable (see algorithm in [Supplementary section S3](#), available as [Supplementary data](#)

at *IJE* online). Variables with the integer field type are usually treated exactly the same as the continuous variables. In a minority of cases, where there are 20 or fewer distinct values we treat this variable as ordered categorical (or binary if there are only two distinct values).

Categorical (single) variables may be assigned to the binary, ordered categorical or unordered categorical data types. UK Biobank consistently assigns negative values to categories denoting missingness (such as ‘Preferred not to answer’ and ‘Do not know’), and so we recode negative values to ‘NA’. UK Biobank defines ‘data codes’ to which one or more fields are assigned, and these define the set of categorical values for these fields and their corresponding numerical values. The PHESANT data-coding information file specifies whether a data code of a categorical (single) field defines an ordered or unordered category structure, and we use this information to assign each non-binary categorical (single) field as either an ordered or an unordered categorical data type.

Each categorical (multiple) variable is converted to a set of binary variables, each denoting whether a participant has a given value of this variable. For example, for the variable describing the bread eaten yesterday (field ID = 20091; [Supplementary Figure 2](#)), with values ‘white’, ‘mixed’, ‘wholemeal’, ‘seeded’ and ‘other’, we generate five binary variables: $white = \{true, false\}$, $wholemeal = \{true, false\}$ and so forth. Categorical (multiple) fields have an added complexity because when a person has no value in this field, this may be because: (i) the field values are incomplete—they do not contain all possible values (e.g. a participant who does not eat bread cannot choose any option above); or (ii) because the data are missing (e.g. because a participant did not answer this particular question). This affects who we assign as, for instance, $white = false$, either: (i) all people who selected a value other than ‘white’; (ii) all people who responded to this questionnaire and did not select ‘white’; or (iii) all people who did not select ‘white’ including those who did not respond to the questionnaire (see [Supplementary Figure 3](#) for illustration, available as [Supplementary data](#) at *IJE* online). In this case the second option might be preferred, such that we are comparing those who ate white bread with those who responded to the questionnaire but did not eat white bread. This decision is variable-specific, and can be specified in the PHESANT variable information file (the variable information file we used for the example described in the Usage section is available in the GitHub repository). For more details see [Supplementary material section S1](#), available as [Supplementary data](#) at *IJE* online.

Some categorical multiple fields include negative numerical values for particular categories denoting missingness (such as ‘Do not know’). We exclude all participants with a missing value from the *false* value of the generated binary variable, because we cannot know if they do or do

not pertain to the *true* value of this binary variable. For example, consider field 41228 [<http://biobank.ctsu.ox.ac.uk/showcase/field.cgi?id=41228>] describing the type of medical professional who conducted the delivery of a participant’s child, and a participant who has given birth twice and has values ‘midwife’ and ‘not known’ in this field. The generated binary variable for midwife includes this participant in the set of participants corresponding to $midwife = true$ because we know that on at least one occasion a midwife conducted the delivery. However, we cannot be certain that a hospital doctor did not conduct a delivery for this participant because the ‘not known’ value could refer to a ‘hospital doctor’. Hence, the generated ‘hospital doctor’ binary variable would not include this participant in the set of participants corresponding to $hospital_doctor = false$, because this is not known.

Tests of association with trait of interest

The association of each phenotype, having been appropriately coded and assigned one of the four data types (continuous, ordered categorical, unordered categorical and binary), is tested with the trait of interest as follows. The phenotype and the trait of interest are the dependent and independent variables of the regression, respectively. All regressions are adjusted for age at recruitment and sex (and also genotype chip when the trait of interest is genetic, derived from the genotype measurement batch), or for all variables in the confounder file if one is specified. We test the association with the transformed variables of the continuous data type using linear regression (`lm` R function). Ordered categorical, unordered categorical and binary variables are tested using ordered logistic regression (`polr` R function), multinomial logistic regression (`multinom` R function) and binomial regression (`glm` R function with family parameter as binomial), respectively. We do not test phenotypes where the sample size is fewer than 500, which typically occurs for a minority of fields such as follow-up questions on a subsample (e.g. field 22148 [<http://biobank.ctsu.ox.ac.uk/showcase/field.cgi?id=22148>]). We do not test unordered categorical variables with more than 1000 categories (above the default maximum of the `multinom` function), which occurs once in our usage example, field 132 [<http://biobank.ctsu.ox.ac.uk/showcase/field.cgi?id=132>].

Customizing a phenome scan with PHESANT

PHESANT allows researchers to easily customize the phenome scan by changing settings in the data coding and variable information files. This includes:

- changing the numerical values underlying a variable (such as recoding a value to missing) or the ordering of values for ordered categorical variables;
- assigning a default value to categorical (single) variables where this is not explicitly coded in the variable (see [Supplementary section S1](#));
- changing fields from the categorical (single) to the categorical (multiple) field type, as this may be more appropriate for a small number of fields;
- specifying which variables should be excluded a priori from the phenome scan;
- specifying which fields in the phenome dataset are essentially the same phenotype as the trait of interest [e.g. weight and body mass index (BMI)] such that, after the phenome scan is run, the results of association between these fields and the trait of interest are used for validation only, rather than being included in the results and adding to the multiple testing burden.

Further customization of a phenome scan using PHESANT is possible in two ways.

- A user can manipulate the data before running PHESANT, for example aggregating continuous values using the median rather than the mean or aggregating a field across multiple time points.
- A user can save the derived phenotypes from PHESANT and then perform their own customized analysis with these variables.

Results output

PHESANT ranks the results by *P*-value and outputs a results file, a QQ-plot, and forest plots for the continuous, ordered categorical and binary results that are below the Bonferroni corrected *P*-value threshold [a forest plot for the categorical unordered results is not generated because we have no overall estimate (and confidence interval) for these models, as we use a likelihood ratio test to generate a model *P*-value].

PHESANT-viz: a web-based visualization for phenome scans

Reviewing the results from phenome scans can be challenging due to the number and complexity of phenotypes. As part of PHESANT we have also developed PHESANT-viz, a D3 Javascript visualization that displays phenome scan results as an interactive graph, using the hierarchical field category structure defined by UK Biobank and available at [<http://biobank.ctsu.ox.ac.uk/showcase/label.cgi>]. PHESANT includes a Java program to convert the phenome scan results to the JavaScript Object Notation (JSON) format required for PHESANT-viz. We provide the PHESANT-viz of our usage example below.

Usage

To demonstrate PHESANT, we have performed a MR-pheWAS to search for the causal effects of BMI in UK Biobank (previously performed in a smaller cohort with continuous phenotypes only).² Such an analysis is predicated on the Mendelian randomization principle that genetic variants can be used as instrumental variables to estimate causal effects of the phenotype they proxy for on downstream outcomes.⁷ In the current context this would be a screening exercise to identify associations for detailed follow-up. This analysis is preliminary and for example only, having been run on a non-random subsample of 114 963 participants (containing the UK BILEVE samples selected on smoking status; see [Supplementary section S4](#) for details, available as [Supplementary data](#) at *IJE* online) for which genetic data are currently available in UK Biobank (a final analysis will be subsequently published upon release of the full 500 000 sample with genetic data).

We created an allele score from 96 genetic variants previously found to be associated with BMI, in a recent genome-wide association study (GWAS) meta-analysis.⁸ The score was calculated as a sum of the number of BMI-increasing alleles, weighted by the effect size as reported in⁸ (see [Supplementary table 1](#), available as [Supplementary data](#) at *IJE* online). Hence, a higher genetic score corresponds to a tendency towards higher BMI (F-statistic = 1979). We used PHESANT to test the association of the BMI genetic score with the 290 integer, 1030 continuous, 658 categorical (single) and 99 categorical (multiple) fields available in UK Biobank at the current time (excluding 66 fields a priori, see [Supplementary Table 2](#), available as [Supplementary data](#) at *IJE* online). [Supplementary Figure 4](#) shows the number of variables reaching each stage of our variable processing flow (available as [Supplementary data](#) at *IJE* online).

[Figure 2](#) shows the QQ plot of our MR-pheWAS results (see full results ranking in [Supplementary Data File](#), and forest plots in [Supplementary Figure 5](#), available as [Supplementary data](#) at *IJE* online). Of the 10 624 tests performed (excluding 87 phenotypes tested but specified a priori as being aspects of the same essential phenotype as BMI), 87 were associated at a Bonferroni corrected *P*-value threshold of 4.71×10^{-6} (0.05/10624). We detected several known effects of BMI, for example with hypertension⁹ (fields 41204 value I10, and 4079), diabetes¹⁰ (field 2443) and age at puberty in both sexes¹¹ (fields 2714, 2375 and 2385). For instance, a one standard deviation (SD) increase in BMI allele score was associated with a 1.08-fold [95% confidence interval (CI): 1.06, 1.11] higher odds of being diagnosed with hypertension in hospital (field 41204 value I10), and a 0.015 SD (95% CI: 0.010, 0.021) higher diastolic blood pressure (field 4079). We also detected a number of potentially causal




Figure 2. QQ plot of preliminary MR-pheWAS analysis seeking to identify the causal effects of BMI. Dashed line: Bonferroni corrected threshold $P=4.71 \times 10^{-6}$ ($P=0.05$ corrected for 10 624 tests). Dotted line: actual = expected.

associations that were previously unknown. For example, participants with a genetic propensity to higher BMI were less likely to perceive themselves as a nervous person (field 1970) or to call themselves tense or ‘highly strung’ (field 1990).

The PHESANT-viz of these preliminary results can be found at [<http://datamining.org.uk/PHESANT/>] or within the PHESANT package. When this analysis is performed using the full 500 000 sample, the power to detect associations will increase, and it is also likely that the number of tests will increase as fewer variables will be filtered out in the variable processing steps (due to a small sample size). This preliminary analysis took approximately 105 h (using a 1 core Intel E5–2670).

Discussion

PHESANT enables researchers to perform a comprehensive phenome scan in UK Biobank, including a pheWAS, MR-pheWAS or EnWAS. Whereas GWAS have been highly successful at identifying novel associations, we are currently only just beginning to explore the phenome in a hypothesis-free manner.^{1,2,12} In contrast to hypothesis-driven analyses, phenome scans allow exploration across hypotheses without strong priors, and should help to avoid publication bias as analyses are pre-specified and all results, not just the most ‘statistically significant’, are published together.

When undertaking a phenome scan there are several important considerations. First, phenome scans are a screening exercise to identify potentially interesting associations

that should then be analysed more rigorously—and as such the association estimates should be interpreted with caution. The strength of the strongest associations identified may be inflated due to the winner’s curse. Second, it is important to consider the number of tests performed when examining the strength of identified associations. In our usage example we used a conservative Bonferroni corrected threshold to identify potentially interesting associations, which, although reducing the type I error rate, is likely to increase the type II error rate. Third, interpretation of potentially hundreds of results is challenging because the correlated structure of phenotypes means that associations between the trait of interest and each phenotype are not independent. The strongest associations should not be viewed in isolation but alongside the results of related variables for which an association may not have been identified, and to do this we provide a novel visualization approach, PHESANT-viz.

We note the following limitations and areas of future work. It is possible that in some cases our automated rule-based method may deal with variables inappropriately. For example, field 132 [<http://biobank.ctsu.ox.ac.uk/showcase/field.cgi?id=132>], which describes participant’s jobs, is treated as an unordered categorical variable but was removed from our phenome scan because it has more than 1000 categories. In this case, it may be preferable to use the hierarchical structure defined by UK Biobank to combine categories of related jobs. The aim of PHESANT is to provide a ‘broad-brush’ approach for phenome scans, such that users should follow up their phenome scan with a comprehensive analysis of the identified associations, using the most appropriate representation and regression model for each particular trait. For example, instead of using an inverse normal rank transformation for continuous fields (as performed by PHESANT), a multivariable linear regression model, or other model with appropriate error distribution for non-normal variables (such as zero-inflated models), can be used.

PHESANT uses the first time point of a field where multiple time points are available. In future work we will investigate how to automate the analysis including data from multiple time points, and the potential gains that this would give.¹³ Currently, when continuous variables are converted to the ordered categorical data type we arbitrarily chose to generate three categories, and in future work will investigate whether a larger number of categories are beneficial, and whether the optimal number of categories can be calculated from the distribution of the variable. PHESANT is specifically designed for use in UK Biobank, but a cohort-independent tool for phenome scans would be highly valuable. Hence, in the future we will aim to adapt PHESANT for use in a general setting. Finally, we intend

to integrate PHEASANT with MR-base¹⁴ to enable automated construction of genetic instrumental variables to use in MR-pheWAS.

The large number of participants combined with the extensive range of phenotypes available in UK Biobank provide a great opportunity to comprehensively search for novel (potentially causal) associations in a hypothesis-free manner. We are aware of only one very small phenome scan that has been performed in UK Biobank to date,¹⁵ and a recent novel Bayesian approach of self-reported diagnoses and hospital episodes.¹⁶ PHEASANT allows users to initiate multiple ‘jobs’ where each scans a subset of phenotypes, such that phenome scans are easily parallelizable whether running on a single multi-core machine or a multi-node high performance computing cluster. To our knowledge, PHEASANT is the first open-source package to automate phenome scans across diverse sets of phenotypes.

Supplementary Data

Supplementary data are available at *IJE* online.

Funding

This work was supported by the University of Bristol and UK Medical Research Council [grant numbers MC_UU_12013/1, MC_UU_12013/8 and MC_UU_12013/9].

Conflict of interest: T.R.G. holds research funding from Sanofi and GlaxoSmithKline for unrelated research.

References

- Denny JC, Bastarache L, Roden DM. Phenome-wide association studies as a tool to advance precision medicine. *Annu Rev Genomics Hum Genet* 2016;17:353–73.
- Millard LAC, Davies NM, Timpson NJ, Tilling K, Flach PA, Davey Smith G. MR-PheWAS: hypothesis prioritization among potential causal effects of body mass index on many outcomes, using Mendelian randomization. *Sci Rep* 2015;5:16645.
- Tzoulaki I, Patel CJ, Okamura T *et al.* A nutrient-wide association study on blood pressure. *Circulation* 2012;126:2456–64.
- Patel CJ, Bhattacharya J, Butte AJ. An environment-wide association study (EWAS) on type 2 diabetes mellitus. *PLoS One* 2010;5:e10746.
- Denny JC, Ritchie MD, Basford MA *et al.* PheWAS: demonstrating the feasibility of a phenome-wide scan to discover gene-disease associations. *Bioinformatics* 2010;26:1205–10.
- Allen N, Sudlow C, Downey P *et al.* UK Biobank: Current status and what it means for epidemiology. *Health Policy Technol* 2012;1:123–26.
- Davey Smith G, Ebrahim S. ‘Mendelian randomization’: can genetic epidemiology contribute to understanding environmental determinants of disease? *Int J Epidemiol* 2003;32:1–22.
- Locke AE, Kahali B, Berndt SI *et al.* Genetic studies of body mass index yield new insights for obesity biology. *Nature* 2015;518:197–206.
- Timpson NJ, Harbord R, Davey Smith G, Zacho J, Tybjaerg-Hansen A, Nordestgaard BG. Does greater adiposity increase blood pressure and hypertension risk? Mendelian randomization using the FTO/MC4R genotype. *Hypertension* 2009;54:84–90.
- Holmes MV, Lange LA, Palmer T *et al.* Causal effects of body mass index on cardiometabolic traits and events: a Mendelian randomization analysis. *Am J Hum Genet* 2014;94:198–208.
- Mumby HS, Elks CE, Li S *et al.* Mendelian randomisation study of childhood BMI and early menarche. *J Obes* 2011;2011:180729.
- VanderWeele TJ. Outcome-wide epidemiology. *Epidemiology* 2017;28:399–402.
- Cao Y, Rajan SS, Wei P. Mendelian randomization analysis of a time-varying exposure for binary disease outcomes using functional data analysis methods. *Genet Epidemiol* 2016;40:744–55.
- Hemani G, Zheng J, Wade KH *et al.* MR-Base: a platform for systematic causal inference across the phenome using billions of genetic associations. *bioRxiv* 2016; doi: <https://doi.org/10.1101/078972>.
- Emdin CA, Khera AV, Natarajan P *et al.* Phenotypic characterization of genetically lowered human lipoprotein (a) levels. *J Am Coll Cardiol* 2016;68:2761–72.
- Cortes A, Dendrou C, Motyer A *et al.* Bayesian analysis of genetic association across tree-structured routine healthcare data in the UK Biobank. *bioRxiv* 2017; <https://doi.org/10.1101/105122>.