

Author's Response To Reviewer Comments

Close

We appreciate the valuable comments of the anonymous reviewers which were helpful to improve the quality of the paper. The changes to our manuscript that reflect reviewers' comments are highlighted in 'red' in the following document.

1. Revision on Comments of Reviewer 1
2. Revision on Comments of Reviewer 2
3. Revision on Comments of Reviewer 3

[Revision on Comments of Reviewer 1]

[R1: Comment 1]

Generally well written and logical flow. Some minor errors (e.g. "an SNP" rather than "a SNP") and some headers could be improved for readability (e.g. "Testing" is vague; this section really only touches upon Run time).

[Response to Comment 1]

We correct the identified mistakes in our manuscript, fixing all instances of "an SNP" to "a SNP," as well as renaming the "Testing" section to "Runtime Analysis."

[R1: Comment 2]

Figure 1- Displaying a single Manhattan plot for "PheWAS Summary Statistics" is not very intuitive. It makes me think of a single GWAS rather than a phenome-wide set of GWAS run on a Biobank. Perhaps revise the image.

[Response to Comment 2]

Figure 1 has been revised to include a Manhattan plot corresponding to a PheWAS instead of a GWAS, with phenotypes instead of genes appearing along the x-axis of the plot.

[R1: Comment 3]

Is the disease-disease network only applicable to case/control studies? Could there be an extension to quantitative traits, and if so, would that be pertinent for discoveries?

[Response to Comment 3]

Disease-disease networks are indeed applicable to both binary and continuous phenotypes. In the case of quantitative traits, a PheWAS would be performed between genetic variants as independent variables and the continuous value of each phenotype as the outcome variable. Considering the associations between variants and such diseases may provide additional nuance into the strength of links between phenotypes. We have revised the Discussion and Conclusions section to include a mention of NETMAGE's applicability to quantitative traits.

[In the revised manuscript]

(6. Discussion and Conclusions)

...to facilitate large-scale genetic analysis of the human diseasome. While the UKBB data used for our case study consisted solely of binary phenotypes, NETMAGE can also be applied to quantitative traits. Indeed, in such a situation, the continuous value of the quantitative phenotype serves as the outcome variable in the PheWAS. This process provides a more detailed degree of association between the trait and genetic variants, suggesting a link between the variant and the severity of the phenotype as

opposed to its presence or absence.

[R1: Comment 4]

The authors refer to "SNPs" throughout to define genetic variation. If the summary statistics contains another type of variation (e.g. indels), are those associations still used? If so, I would suggest using a more generic term to define the genetic variation.

[Response to Comment 4]

We thank the reviewer for identifying this over-simplification in our manuscript. Indeed, NETMAGE can use data involving any sort of genetic variation to generate a corresponding DDN. We replace the term "SNP" in our manuscript with "variant" or "genetic variant" as appropriate.

[In the revised manuscript]

(Abstract)

...Using summary statistics from a phenome-wide association study (PheWAS), we can generate a corresponding DDN where edges represent shared genetic variants between diseases...

...Users can search the map by various attributes and select nodes to view related phenotypes, associated variants, and various network statistics.

(1. Background)

...a phenome-wide association study (PheWAS) can be used to calculate a multitude of associations between phenotypes and genetic variants, such as single-nucleotide polymorphisms (SNPs) in an unbiased manner...

...In particular, a DDN that uses its edges to represent variants can be generated as a proxy to highlight potential shared genetic influences for diseases. Analyzing the topology of these genetics-based DDNs can provide insight into how inherited factors may drive the onset of disease complications...

(2. Purpose of the work)

...The network-based visualization of associations between variants and phenotypes can provide researchers...

... In particular, the resulting DDN is a projection of an undirected bipartite network of phenotypes and genetic variants, where nodes serve as diseases and edges serve as sets of common associated variants. Users can filter their input data by p-value and by minor allele frequency (MAF) to manipulate the rarity and significance of variants being used to generate the network. Furthermore, they can select nodes within the DDN to view information such as connected phenotypes, shared variants, and network statistics...

... users can follow up with phenotypic data in their corresponding EHRs to evaluate the predictive ability of genetics-based DDNs with respect to disease co-occurrences...

(Table 1)

...Allow users to search and create subsets of any produced networks by disease, by genetic variant, or by other network statistics.

(3. Implementation)

...Each row should correspond to a genetic variant, and the user can provide p-value...

...This file represents a dictionary of phenotype-to-variant mappings, where each phenotype serves as a key and each variant, p-value, MAF triplet serves as a value in a set...

...Based upon the p-value and MAF thresholds provided by the user, phenotype-variant mappings will be filtered to provide a final file containing a list of relevant variants for each disease...

... The weight of the edge is equal to the number of associated variants shared between the two phenotypes...

... Each row provides a distinct phenotype and a list of its associated variants...

... users can search the map for relevant phenotypes based upon any attributed defined, such as phenotype name, phenotype ID, variant name, node degree, and other parameters. In particular, the "search by variant" option allows users to find shared genetic variants between diseases.

...searchability of DDNs by both phenotype and genetic variant...

...providing a map of phenotypes to variants...

(Figure 2)

...Additionally, associated variants, connected phenotypes, and ...

(6. Discussion and Conclusions)

...Furthermore, we hope to enhance NETMAGE to allow for the automated construction of gene-based DDNs from variant-based data by including variant-to-gene mapping as a part of the website. We will also allow users the option to create variant-variant networks that depict edges between genetic variants based upon shared associations with phenotypes ...

[R1: Comment 5]

The discussion seems underdeveloped. Discussion of limitations rather than only future work would be helpful.

[Response to Comment 5]

We have revised the "Discussion" section to include a paragraph on current limitations of NETMAGE. The Discussion has also been extended to include a description of DDNs generated from continuous traits.

[In the revised manuscript]

(6. Discussion and Conclusions)

...The goal of this software is to improve the ease of visualization of genetic associations across diseases and to facilitate large-scale genetic analysis of the human diseaseome. While the UKBB data used for our case study consisted of entirely binary phenotypes, NETMAGE is also applicable to quantitative traits. Indeed, in such a situation, the continuous value of the quantitative phenotype, such as a laboratory test measurement like A1C level, is used as the outcome variable in the PheWAS. This process provides a more detailed degree of association between the severity of the trait and genetic variants, as compared to the identification of associations between a presence or absence of the trait with variants. A key point to note regarding NETMAGE is that the output DDNs will provide only as much information as the input data. Indeed, NETMAGE is an exploratory tool intended to help visualize connections between diseases. Including summary PheWAS data that provides insight into the statistical associations between phenotypes will yield an associative map but will tell us nothing about causality. Associations identified through PheWAS are often spurious, so any sort of analyses performed on these data must take this information into consideration. Nevertheless, these kinds of associative visualizations are still useful for the study of disease and may help identify connections between phenotypes and genetic variants, generate new hypotheses, and suggest future experiments that can be conducted. For a visualization that gives stronger insight into the causal connections between traits, one could potentially input the results of a Mendelian Randomization experiment. Several future directions exist for NETMAGE. First is the inclusion of directionality in the network – as of now, DDNs produced by NETMAGE give no indication regarding the direction of association between phenotypes. Using beta values for the association between phenotypes and genetic variants would be a useful inclusion, aiding in clinical interpretation of the network. We will also allow for the concurrent selection of multiple nodes within the DDN. The current NETMAGE user interface allows only one node to be selected at a time. The ability to select multiple nodes will allow clinicians to quickly identify if two phenotypes are associated in the network. We also hope to enhance NETMAGE to allow for the construction of gene-based DDNs from variant-based data by including variant-to-gene mapping as a part of the website. Finally, we will allow users to create variant-variant networks instead of disease-disease networks, which depict the connections between genetic variants (for instance, SNPs) based upon associations with phenotypes. Ultimately, NETMAGE will give researchers and clinicians insight into the underlying genetic architecture of disease complications...

[R1: Comment 6]

Case study-- The authors could improve the interpretability/discussion of the UKB PheWAS example. This is one of my largest concerns because the author state that the tool can help researchers and

clinicians get insight into the underlying genetic architecture of disease complications; however, the case study part of the manuscript is quite technical and could be challenging to interpret for someone without network experience; e.g. Table 2.

[Response to Comment 6]

We very much appreciate the reviewer's comments regarding interpretability – we have edited Table 2 to simply list hub diseases identified through network centrality measures. Phenotypes identified according to multiple centrality measures are depicted in bold in the table. Furthermore, to aid with interpretability when exploring the DDN, we have included a new hyperlink in the "Information Pane" when a phenotype is selected. This link directs the user to a new window which depicts a histogram of diseases connected to the phenotype of interest, sorted in order of number of shared variants. This new feature should aid users in visualizing the significance of disease connections to a phenotype in the DDN, allowing for improved interpretability.

[In the revised manuscript]

(3. Implementation)

...Clicking on a node will highlight the node and all its first-degree neighbors. A variety of default attributes will be presented on the right side of the webpage as part of an "Information Pane." The user can also define other custom attributes, and these will be included in the Information Pane as well. If the user inputs data that include rsID-formatted SNPs, then NETMAGE will automatically hyperlink each SNP's ID to its corresponding dbSNP profile, allowing for further exploration of the variant's information. To aid with interpretation and visualization of disease associations, a hyperlink to a histogram of disease connections is also included in the Information Pane. For each phenotype, this histogram depicts first-degree disease neighbors sorted in order of the number of shared variants...

[R1: Comment 7]

Additionally, more details should be provided on the underlying summary statistics used (e.g. some details can be found on the About page of the HRC-imputed UKB PheWeb page: <https://pheweb.org/UKB-SAIGE/about>).

[Response to Comment 7]

We thank the reviewer for pointing us to this clearer description of our input dataset. We include additional details about our data in the "Case Study" section of the manuscript

[In the revised manuscript]

(4. Case Study)

...These data corresponded to 1,403 binary phenotypes expressed in terms of PheCodes. All 400,000 British individuals of European ancestry in the dataset were imputed using the Haplotype Reference Consortium panel, yielding 28 million imputed SNPs.¹¹ SAIGE¹⁶, a generalized mixed model association test that uses the saddlepoint approximation to account for case-control imbalance, was used to generate summary statistics for each SNP, providing p-values of association between every SNP and every phenotype. This analysis was adjusted for genetic relatedness, sex, birth year, and the first four principal components.¹¹ All genomic positions are on GRCh37.¹¹ Phenotypes that had a case count lower than 200 were dropped to keep more relevant diseases, yielding a total of 1075 traits for consideration...

[R1: Comment 8]

The authors list additional filtering that they performed on the summary statistics, but it appears that some details are missing. For instance, how many traits remain after the case count filtering is applied? Also, what is used as a reference for the LD-pruning in PLINK?

[Response to Comment 8]

We have revised the description of our filtration steps to include how many traits were included after case count filtering, as well as a mention of the reference panel used for LD-pruning.

[In the revised manuscript]

(4. Case Study)

...SAIGE was used to generate summary statistics for each variant, providing p-values of association between every variant and every phenotype. Phenotypes that had a case count lower than 200 were dropped to keep more relevant diseases, yielding a total of 1075 traits. Data were also filtered in order to select significant associated common variants, based upon the following thresholds: maximum p-

value threshold of 5×10^{-8} , minimum MAF of 0.01, and LD-pruning through PLINK using the quality-controlled UKBB genetic data itself as our reference panel, with an R^2 of 0.2 and 250 kilobases for maximum search length. Removing nodes with degree 0 after the previously described filtration steps yielded a final network of 232 nodes and 2375 edges...

[R1: Comment 9]

Run time-- I am wondering why Table 3 (run time for subsets of the UKBB data) ends at 1000 phenotypes. It would be interesting to see the run time that is close to case example (e.g. possibly adding a column for the total number of phenotypes used in the UKBB DDN). Additionally, this section gives the impression that run time only depend on the number of phenotypes? I would assume that run time should also depend on the number of variants that were tested.

[Response to Comment 9]

We clarify in the text of our runtime section that increasing the number of variants under consideration will increase the runtime. We also include a new row in Table 3 that includes the runtime for the UKBB DDN case study for both the Fruchterman-Reingold and Force Atlas 2 layouts.

[In the revised manuscript]

(5. Runtime Analysis)

... This behavior makes sense, as runtime depends on not only the number of phenotypes included in the input data, but also the number of variants being tested. Indeed, assuming that each additional phenotype added to the network will include multiple associated variants, the inclusion of nodes will tend to exponentially increase the number of edges assuming a low clustering coefficient in the network...

[R1: Comment 10]

It is nice that on each page the authors have allowed users to download a pdf of the image and also the data behind the image (e.g. edge-map, node-map, etc.). The zoom-in feature for the visualization is also useful, as is the short video tutorial.

[Response to Comment 10]

We thank the reviewer for their comments regarding the software and website.

[R1: Comment 11]

I think that the search bar would be more user-friendly if suggestions automatically came up when the user begins to type.

[Response to Comment 11]

Auto-completion for any sort of categorical variable (e.g. phenotype ID, associated SNP ID, and category) has been implemented in the NETMAGE tool. Now, as a user begins to type, NETMAGE will refer to all possible values in the input data and provide suggestions that the user can search.

[In the revised manuscript]

(3. Implementation)

...Users can search the map for relevant phenotypes based upon any attribute defined, such as phenotype name, phenotype ID, variant name, node degree, and other parameters. In particular, the "search by variant" option allows users to find shared genetic variants between diseases. The custom attributes provided by the user are also automatically incorporated into the search dropdown menu. Any categorical variables, such as disease name, disease category, or variant name, will include an auto-completion dropdown menu that dynamically updates as users type out their query terms...

[R1: Comment 12]

Additionally, displaying the list of "associated SNPs" in a (sortable and/or searchable) table (with some annotations, such as chr, position, closest gene, consequence, rather than just rsID) could be a neater and more informative way to show these data, rather than simply as it appears currently as a list in the "information pane".

[Response to Comment 12]

The inclusion of a dynamically updating table of SNP information for each phenotype is challenging to

include in the current version of NETMAGE, particularly since users may not be uploading data purely corresponding to SNPs. Instead, we have revised the Information Pane's presentation of associated variant information to present SNP information in a more useful manner. If the user's input data includes variant IDs that are formatted in terms of rsIDs, then the variants will automatically be hyperlinked to their profiles on dbSNP. Otherwise, the list of associated variants stays as is. This behavior allows users to delve into the details of a SNP of interest. In the future, we will offer the ability to view a table of annotated SNP information for each phenotype based upon Annovar/VEP. We will also allow users to download a text file of associated variants for a phenotype of interest, including links to dbSNP if appropriate as well as Annovar annotations. We have raised a ticket on GitHub for this update, and it can be found at the following link: <https://github.com/dokyoonkimlab/netmage/issues/20>. For now, if users wish to download a list of variants associated with a phenotype, they can download the "Node Map" file to see the genetic associations for their desired disease.

[In the revised manuscript]

(3. Implementation)

...clicking on a node will highlight the node and all its first-degree neighbors. A variety of default attributes will be presented on the right side of the webpage as part of an "Information Pane." The user can also define other custom attributes, and these will be included in the Information Pane as well. If the user inputs data that include rsID-formatted SNPs, then NETMAGE will automatically hyperlink each SNP's ID to its corresponding dbSNP profile, allowing for further exploration of the variant's information. To aid with interpretation and visualization of disease associations, a hyperlink to a histogram of disease connections is also included in the Information Pane. For each phenotype, this histogram depicts first-degree disease neighbors sorted in order of the number of shared variants...

[R1: Comment 13]

My comment on interpretability for researchers and clinicians comes up again: I am not sure how useful/interpretable some of the search categories are for users to intuitively draw insights; for instance, number of triangles, page range, etc. I think the authors should really focus on the intuitiveness for the target audience so that the tool can have more impact.

[Response to Comment 13]

We appreciate the reviewer's concern regarding interpretability in the presentation of the network information for the DDN. We prefer to keep all network statistics in the visualization, as it is unclear what piece of information might be most useful for users to consider. However, we include "information" icons next to each network statistic term that can be hovered over to provide a brief description of the utility of the variable. We hope that this inclusion helps clear up confusion surrounding network analysis in the DDN.

[Revision on Comments of Reviewer 2]

[R2: Comment 1]

I tried the web interface Human-Disease Phenotype Map (<https://hdpm.biomedinfolab.com>), which utilizes NETMAGE. I found that sometimes it takes some time for the network to appear. While the network is loaded, only the gray empty space with the side panel is shown. I recommend the authors to show the progress bar while loading the network, especially when it is first loaded, to avoid users to think that their web browser is frozen.

[Response to Comment 1]

The NETMAGE website has been updated to include a loading circle as the network is being generated.

[R2: Comment 2]

In the Search bar, it is not always trivial to guess what to enter, especially for Phenotype Name, Associated SNPs, and category. Auto-completion features for these variables will significantly facilitate users' convenience.

[Response to Comment 2]

Auto-completion for any sort of categorical variable (e.g. phenotype ID, associated SNP ID, and category) has been implemented in the NETMAGE tool. Now, as a user begins to type, NETMAGE will

refer to all possible values in the input data and provide suggestions that the user can search.

[In the revised manuscript]

(3. Implementation)

...Users can search the map for relevant phenotypes based upon any attribute defined, such as phenotype name, phenotype ID, variant name, node degree, and other parameters. In particular, the "search by variant" option allows users to find shared genetic variants between diseases. The custom attributes provided by the user are also automatically incorporated into the search dropdown menu. Any categorical variables, such as disease name, disease category, or variant name, will include an auto-completion dropdown menu that dynamically updates as users type out their query terms...

[R2: Comment 3]

Meaning of edges is somewhat unclear to me. Are the existence and the weights of edges purely based on the number of shared SNPs or are they based on any statistical methods? When the weights of edges are calculated, are the marginal counts taken into account? The same number of shared SNPs can have different meanings when the disease to which this edge is connected has a small number of associated SNPs vs. a large number of associated SNPs. How is this factor considered?

[Response to Comment 3]

We very much appreciate this point noted by the reviewer. In the baseline version of the DDN, as described in Step 3 of the "Implementation" section, "The weight of the edge is equal to the number of associated variants shared between the two phenotypes." However, as the reviewer mentions, the degree of the phenotypes in question can have a clear impact on the significance of an edge between two diseases. To address this discrepancy, we have incorporated a "Marginalize edges" checkbox in the website. Users can specify if they want their edge weights to simply represent the number of shared variants between two diseases, or by selecting the checkbox, if they want the weight of the edge to be marginalized by the number of variants associated with each of the parent phenotypes.

[In the revised manuscript]

(3. Implementation)

...This file is used to generate an edge map and a node map. The edge map establishes all links in the network – each row corresponds to an edge from a source to a target. Depending on the user's choice, the weight of the edge equals either the number of associated variants shared between the two phenotypes, or the marginalized fraction of variants (the number of variants that constitute the edge divided by the union of the individual sets of variants for both phenotypes). In addition, the node map represents a list of all nodes in the network...

[R2: Comment 4]

The network generated by the Human-Disease Phenotype Map (<https://hdpm.biomedinfolab.com>) is usually huge and complex with a large number of edges. As a result, it is often not straightforward to understand the generated network. This is partially relevant to the fact that the network layout is static, i.e., locations of nodes remain the same regardless of which subnetworks are chosen. If the network layout is optimized for each subnetwork, it should be much easier for users to understand the network architecture. Given this, I recommend the authors to consider updating the network layout interactively when a subnetwork is selected.

[Response to Comment 4]

We appreciate the reviewer's suggestion to dynamically update the network layout when a subset of the network is chosen. This feature will require a considerable amount of work in terms of the structure of our code and will be handled in a future version of the software. We have raised a ticket on GitHub for this issue, and it can be found at this link: <https://github.com/dokyoonkimlab/netmage/issues/19>

[R2: Comment 5]

When a subnetwork is chosen, the "Information Pane" appears. In this pane, it might be helpful for users if the authors provide some quick help link for each network score, e.g., how to interpret PageRank scores, etc.

[Response to Comment 5]

We appreciate the reviewer's suggestion for how to improve the interpretability of resulting networks. We have edited the NETMAGE Information Pane so that for each network statistic, an information icon can be hovered over that provides a brief description of the statistic's purpose.

[R2: Comment 6]

In the "Information Pane", a long list of SNPs is provided for "Associated SNPs" but it is not easy to use this list. I recommend the authors to make it downloadable as a table so that users can do downstream analysis. In addition, it will significantly facilitate users' convenience if each SNP ID is chosen, it brings the user to the relevant database, e.g., dbSNP. In this way, users can easily check where it is located in the sense of chromosome, gene, exon/intron/promoter/intergenic, etc. Alternatively, the authors can consider to use a quick information table (SNP ID, gene name, exon/intron/promoter/intergenic) instead of simply providing as a list.

[Response to Comment 6]

We have revised the Information Pane's presentation of associated variant information to present SNP information in a more useful manner. If the user's input data includes variant IDs that are formatted in terms of rsIDs, then the variants will automatically be hyperlinked to their profiles on dbSNP. Otherwise, the list of associated variants stays as is. This behavior allows users to delve into the details of a SNP of interest. In the future, we will offer the ability to view a table of annotated SNP information for each phenotype based upon Annovar/VEP. We will also allow users to download a text file of associated variants for a phenotype of interest, including links to dbSNP if appropriate as well as Annovar annotations. We have raised a ticket on GitHub for this update, and it can be found at the following link: <https://github.com/dokyoonkimlab/netmage/issues/20>. For now, if users wish to download a list of variants associated with a phenotype, they can download the "Node Map" file to see the genetic associations for their desired disease.

[In the revised manuscript]

(3. Implementation)

...clicking on a node will highlight the node and all its first-degree neighbors. A variety of default attributes will be presented on the right side of the webpage as part of an "Information Pane." The user can also define other custom attributes, and these will be included in the Information Pane as well. If the user inputs data that include rsID-formatted SNPs, then NETMAGE will automatically hyperlink each SNP's ID to its corresponding dbSNP profile, allowing for further exploration of the variant's information. To aid with interpretation and visualization of disease associations, a hyperlink to a histogram of disease connections is also included in the Information Pane. For each phenotype, this histogram depicts first-degree disease neighbors sorted in order of the number of shared variants...

[Revision on Comments of Reviewer 3]

[R3: Comment 1]

A DDN based on true genetic associations is useful for understanding complex disease comorbidities and their shared genetic etiology (pleiotropy). An interactive web tool to explore such a complex networked information could be highly useful for the proposed purposes of this tool. However, the EHR/Biobank PheWAS associations data are statistical in nature and commonly with small effect sizes. The reported genetic associations often are not well understood at the mechanistic level, and many genetic associations are spurious. Although certain positive findings can be observed from the disease network generated by NETMAGE, it's of concern the general usability of the current implementation of the tool in order to facilitate novel applications in drug design and personalized medicine, which requires the genetic associations to best represent the underlying true causal mechanism. Further work is needed to verify the genetic associations reported from PheWAS to minimize the impact of spurious associations.

[Response to Comment 1]

We appreciate the reviewer's comments regarding the implications of NETMAGE. The applicability of the data that go into the software will dictate the applicability of the resulting DDN. Indeed, with the results of a PheWAS, NETMAGE will be able to produce only an associative map of disease connections. We include a more thorough discussion of association vs. causation in our "Discussion and Conclusions" section.

[In the revised manuscript]

(6. Discussion and Conclusions)

...A key point to note regarding NETMAGE is that the output DDNs will provide only as much information as the input data. Indeed, NETMAGE is an exploratory tool intended to help visualize connections

between diseases. Including summary PheWAS data that provides insight into the statistical associations between phenotypes will yield an associative map but will tell us nothing about causality. Associations identified through PheWAS are often spurious, so any sort of analyses performed on these data must take this information into consideration. Nevertheless, these kinds of associative visualizations are still useful for the study of disease and may help identify connections between phenotypes and genetic variants, as well as suggest future experiments that can be conducted. For a visualization that gives stronger insight into the causal connections between traits, one could potentially input the results of a Mendelian Randomization experiment...

[R3: Comment 2]

Network edges based on SNPs without considering the linkage disequilibrium (LD) between SNPs is misleading and could miss a significant portion of associations that should be linked between diseases if the LD correlations are considered. When construct the network using NETMAGE, the LD correlation between SNPs should be considered.

[Response to Comment 2]

We thank the reviewer for identifying this gap in our software. While we had included a feature to account for an input LD file in our back-end software, we had failed to include it in the NETMAGE website itself. The option for LD pruning according to an input LD file is now incorporated into the NETMAGE website.

[In the revised manuscript]

(3. Implementation)

...Each row provides a distinct phenotype and a list of its associated variants. If input data have not already been pruned for linkage disequilibrium (LD), then users can provide an LD-mapping file that gives mappings between each variant to blocks of LD. NETMAGE will then clump SNPs according to their specified LD blocks, ensuring that associations that should be linking phenotypes together are present in the map. Users can also provide an input disease category mapping file so that each row of the node map now represents the disease and its category...

[R3: Comment 3]

For the reported DDN and its statistics to be relevant to true disease - disease relationships, the quality of disease diagnosis using Phecode should be considered. Phecodes are based on ICD codes that are known to be noisy. The accuracy of ICD can be as low as only 50%. Ignoring this limitation and treating disease diagnoses from Phecodes as gold standards or as precise and accurate may result in irrelevant and misleading findings.

[Response to Comment 3]

This point regarding the appropriateness of Phecodes is extremely relevant. We include a small description of the limitations of Phecodes in our "Case Study" section.

[In the revised manuscript]

(4. Case Study)

...and rs780094's association with diabetes and lipid metabolism. One potential issue in terms of the conclusions that can be drawn from our UKBB DDN is the use of "PheCodes" as a method of defining phenotypes. PheCodes are defined according to ICD codes, but the accuracy of these codes for disease diagnosis is known to be questionable. Given such inaccuracies, users must be wary when treating PheCode or ICD-based diagnoses as a gold standard, as doing so may lead to inaccurate conclusions...

[R3: Comment 4]

Phecodes are hierarchical. For example, parent codes are three digits (008), and each additional digit after decimal point indicates a subset of ICD codes of the parent code (008.5 and 008.52). So here a code 008.52 implies 008.5 also 008. What's the impact of this hierarchy to the NETMAGE network and the inferences to be made based on the network?

[Response to Comment 4]

We agree with the reviewer that hierarchy between phenotypes may influence resulting DDNs. In our case study, the data we make use of includes mostly upper hierarchy phenotypes. More detailed hierarchical phenotypes are absent from the data. Users should be careful to avoid extensive hierarchical structure in their input data when generating DDNs through NETMAGE. We include a description of this

facet in the "Case Study" section of our manuscript.

[In the revised manuscript]
(4. Case Study)

...Another aspect of the use of PheCodes for phenotype definitions is their hierarchical nature. Digits that appear after decimal points correspond to subsets of phenotypes compared to the parent code that appears before the decimal. In our case study, the data we make use of include mostly upper hierarchy phenotypes. More detailed hierarchical phenotypes are for the most part absent from our network. Users should be careful about including extensive hierarchical structure in their input data when generating DDNs through NETMAGE. Including phenotypes that are essentially identical to one another will introduce unnecessary nodes and edges in the network, in the process clouding more significant disease connections...

[R3: Comment 5]

On Page 9, you said "Out of the 2189 edges for which phi correlations could be calculated, 1811 (82.73%) appeared in the DDN. This behavior suggests that our genetic associations identified by our PheWAS results serve as a reasonable approximation of disease co-occurrences". This is expected because both phi correlation and PheWAS analyses were performed on the same dataset. If a pair of disease highly co-occur in the dataset, you would expect a strong correlation on their genetic associations analyzed on the same dataset. However, it may not be generalizable that the genetic associations from PheWAS are a reasonable approximation to disease co-occurrences.

[Response to Comment 5]

We thank the reviewer for pointing out this flaw in our analysis of the DDN. We remove this analysis, and instead include a paragraph in our "Case Study" section that describes comparison to external EHR comorbidities as an area of future work.

[In the revised manuscript]
(4. Case Study)

...and rs780094's association with diabetes and lipid metabolism. In terms of future work for this case study, it would be interesting to compare the edges in our DDN with known disease comorbidities. We can take disease occurrence data from an external electronic health record and evaluate phi correlations between all pairs of phenotypes. Comparison of these co-occurrences to the genetic associations in our PheWAS may give us an indication if the DDN is a reasonable representation of disease connections...

[R3: Comment 6]

The disease-SNP relationships from the PheWAS analysis result are bipartite. Even though NETMAGE focuses on the projected disease-disease network, the information about how specific SNPs link to their corresponding disease pairs is important. For example, in your UKBB-based network (<https://hdpm.biomedinfolab.com/ddn/ukbb>), when a specific disease is selected, a subgraph of the selected disease and other disease linked to the selected one are showing, but only a lump of SNPs without linking to their specific disease pair is provided. This is not helpful.

[Response to Comment 6]

We appreciate this comment from the reviewer, and we agree that selecting a single disease in our DDN does not provide insight into the links between variants and their corresponding disease pairs. For this purpose, we recommend that the user makes use of the "Search by SNP" feature to identify in which disease pairs the variant is involved. As a future extension of NETMAGE, we will offer the ability to generate variant-variant or gene-gene networks, which will make it easier to visualize how variants connect to diseases.

[R3: Comment 7]

Also annotating those SNPs their genetic context could be very useful for users to quickly grasp the nature of the genetic associations in the subgraph.

[Response to Comment 7]

We have revised the Information Pane's presentation of associated variant information to present SNP information in a more useful manner. If the user's input data includes variant IDs that are formatted in

terms of rsIDs, then the variants will automatically be hyperlinked to their profiles on dbSNP. Otherwise, the list of associated variants stays as is. This behavior allows users to delve into the details of a SNP of interest. In the future, we will offer the ability to view a table of annotated SNP information for each phenotype based upon Annovar/VEP. We will also allow users to download a text file of associated variants for a phenotype of interest, including links to dbSNP if appropriate as well as Annovar annotations. We have raised a ticket on GitHub for this update, and it can be found at the following link: <https://github.com/dokyoonkimlab/netmage/issues/20>. For now, if users wish to download a list of variants associated with a phenotype, they can download the "Node Map" file to see the genetic associations for their desired disease.

[In the revised manuscript]

(3. Implementation)

...clicking on a node will highlight the node and all its first-degree neighbors. A variety of default attributes will be presented on the right side of the webpage as part of an "Information Pane." The user can also define other custom attributes, and these will be included in the Information Pane as well. If the user inputs data that include rsID-formatted SNPs, then NETMAGE will automatically hyperlink each SNP's ID to its corresponding dbSNP profile, allowing for further exploration of the variant's information. To aid with interpretation and visualization of disease associations, a hyperlink to a histogram of disease connections is also included in the Information Pane. For each phenotype, this histogram depicts first-degree disease neighbors sorted in order of the number of shared variants...

Close