

Systems biology

# Towards future directions in data-integrative supervised prediction of human aging-related genes

Qi Li<sup>1</sup>, Khalique Newaz <sup>1,2</sup> and Tijana Milenković<sup>1,\*</sup>

<sup>1</sup>Department of Computer Science and Engineering, Lucy Family Institute for Data & Society, and Eck Institute for Global Health (EIGH), University of Notre Dame, Notre Dame, IN 46556, USA and <sup>2</sup>Center for Data and Computing in Natural Sciences (CDCS), Institute for Computational Systems Biology, Universität Hamburg, Hamburg 20146, Germany

\*To whom correspondence should be addressed.

Associate Editor: Marieke Lydia Kuijjer

Received on May 25, 2022; revised on September 23, 2022; editorial decision on October 13, 2022; accepted on October 31, 2022

## Abstract

**Motivation:** Identification of human genes involved in the aging process is critical due to the incidence of many diseases with age. A state-of-the-art approach for this purpose infers a weighted dynamic aging-specific subnetwork by mapping gene expression (GE) levels at different ages onto the protein–protein interaction network (PPIN). Then, it analyzes this subnetwork in a supervised manner by training a predictive model to learn how network topologies of known aging- versus non-aging-related genes change across ages. Finally, it uses the trained model to predict novel aging-related gene candidates. However, the best current subnetwork resulting from this approach still yields sub-optimal prediction accuracy. This could be because it was inferred using outdated GE and PPIN data. Here, we evaluate whether analyzing a weighted dynamic aging-specific subnetwork inferred from newer GE and PPIN data improves prediction accuracy upon analyzing the best current subnetwork inferred from outdated data.

**Results:** Unexpectedly, we find that not to be the case. To understand this, we perform aging-related pathway and Gene Ontology term enrichment analyses. We find that the suboptimal prediction accuracy, regardless of which GE or PPIN data is used, may be caused by the current knowledge about which genes are aging-related being incomplete, or by the current methods for inferring or analyzing an aging-specific subnetwork being unable to capture all of the aging-related knowledge. These findings can potentially guide future directions towards improving supervised prediction of aging-related genes via -omics data integration.

**Availability and implementation:** All data and code are available at zenodo, DOI: 10.5281/zenodo.6995045.

**Contact:** tmilenko@nd.edu

**Supplementary information:** [Supplementary data](#) are available at *Bioinformatics Advances* online.

## 1 Introduction

### 1.1 Motivation and background

Human aging is a biological process associated with increased susceptibility to chronic disorders, such as cancer, cardiovascular, Parkinson's and Alzheimer's disease (Ferrucci *et al.*, 2020; Uyar *et al.*, 2020). The aging process has a strong genetic basis, e.g. genomic instability or DNA somatic mutations (Rodríguez-Rodero *et al.*, 2011). So, to understand it, treat associated diseases, and improve life quality for the elderly, it is critical to identify 'hallmark' genes that drive the aging-related molecular mechanisms (Liguori *et al.*, 2018).

Traditionally, such 'hallmark' genes have been identified by *wet lab* experiments (Bologna *et al.*, 2014; Paschos *et al.*, 2012). These efforts have yielded valuable public knowledge about which genes are related to the aging process (Berchtold *et al.*, 2008; Jia *et al.*, 2018; Lu *et al.*, 2004; Tacutu *et al.*, 2018; Simpson *et al.*, 2011).

Yet, such knowledge is limited, because *wet lab* experimentation is difficult due to the ethical constraints and long life span of the human species (Emanuel *et al.*, 2000). So, by benefiting from the existing *wet lab* experimental aging-related knowledge plus wealth of recent -omics data, one can *computationally* predict novel aging-related gene candidates as those that share -omics 'signatures' with the known *wet lab* experimental aging-related genes (Fabris *et al.*, 2017). Unlike with *wet lab* experiments, computational analyses of the aging process can be done on a large, systems-level scale.

Such computational prediction of aging-related genes is typically carried out by supervised classification using gene expression (GE) data or protein–protein interaction network (PPIN) data (Fabris *et al.*, 2017). In particular, GE-based approaches predict a gene as aging-related based on whether its expression level varies with age (Berchtold *et al.*, 2008; Holtman *et al.*, 2015; Jia *et al.*, 2018; Lu *et al.*, 2004; Simpson *et al.*, 2011). While GE-based approaches capture aging-specific information (i.e. the changes of GE with age),

they ignore complex interactions between genes' protein products, which ultimately carry out all biological processes (Chen et al., 2014), including aging. This is why PPIN-based approaches have been introduced, which predict a gene as aging-related if its network topology is 'similar enough' to the network topology of known aging-related genes (Fabris et al., 2016; Fang et al., 2013; Freitas et al., 2011). While PPIN-based approaches capture interactions between proteins, their interactions span different conditions, such as diseases or tissues. In our case, they span different ages, meaning that they do not capture aging-specific information. In other words, their considered PPIN data are context-unspecific (i.e. aging-unspecific). To address the above drawbacks, more recent studies have focused on predicting aging-related genes using both aging-specific GE data and entire context-unspecific PPIN data (Elhesha et al., 2019; Faisal and Milenković, 2014; Kerepesi et al., 2018; Li et al., 2021; Li and Milenković, 2021).

Our group has pioneered this research direction by inferring an aging-specific subnetwork that captures both aging-specific information from GE data and interactions from PPIN data through a series of studies/methods (Faisal and Milenković, 2014; Li and Milenković, 2019, 2021; Li et al., 2021; Newaz and Milenković, 2020). Our most recent finding is that inferring an aging-specific subnetwork that is both weighted and dynamic is superior (in terms of quality of aging-related gene predictions made from it) than inferring an aging-specific subnetwork that is unweighted or static as well as using entire context-unspecific PPIN data (Li et al., 2021).

Intuitively, to infer such a weighted dynamic aging-specific subnetwork, we used network propagation (Komurov et al., 2010; Leiserson et al., 2015) to map gene expression levels from GE data onto a context-unspecific PPIN via random walks or diffusion. This resulted in assigning each interaction in the PPIN with an age-specific weight for each age present in the GE data. Such weighted interactions at a given age form a weighted age-specific subnetwork snapshot. The collection of age-specific snapshots in the increasing order of ages present in the GE data forms a weighted dynamic aging-specific subnetwork. We believe that such a subnetwork results in higher-quality aging-related predictions than an unweighted or static subnetwork because the amounts of aging-specific information that different interactions carry are captured by their weight differences (i) within each age, unlike in an unweighted subnetwork and (ii) across different ages, unlike in a static subnetwork.

Briefly, we make aging-related gene predictions from a subnetwork as follows. Via several iterations (Li and Milenković, 2019, 2021; Li et al., 2021), we developed a comprehensive framework that uses a variety of network features and classifiers in cross-validation. The framework relies on established knowledge, primarily from GenAge (Tacutu et al., 2018), about which genes are aging- versus non-aging-related; this knowledge are genes' ground truth labels. The framework trains a predictive model (feature-classifier combination) on a part of the genes to learn network feature differences between the known aging- versus non-aging-related genes from the training data. Then, it tests the model on the remaining genes, by examining how well it distinguishes between the known aging- versus non-aging-related genes from the testing data. That is, the model predicts each gene from the testing data as either aging-related or not. Then, the model's accuracy, i.e. the quality of its predictions (whether the genes' predicted and ground truth labels match) is evaluated via the area under precision-recall curve (AUPR), precision, recall and F-score. Finally, the most accurate of all predictive models is selected for the subnetwork of interest.

Nonetheless, even this newest (weighted and dynamic) aging-specific subnetwork (Li et al., 2021), which is the state-of-the-art, yields suboptimal accuracy. Our postulations for this finding, which we explore in this study, are:

1. The GE data [from an over a decade old study (Berchtold et al., 2008)] and/or the PPIN data [from HPRD, which has not been updated for about a decade (Prasad et al., 2009)] that we used to infer our weighted and dynamic aging-specific subnetwork are outdated.
2. The current aging-related knowledge, i.e. ground truth labels from GenAge, might be incomplete or otherwise noisy and thus cannot reliably guide computational prediction of novel aging-related knowledge.
3. The methods used to infer an aging-specific subnetwork or to predict aging-related genes from an inferred subnetwork (Li and Milenković, 2019, 2021, Li et al., 2021) cannot capture all of the current aging-related knowledge.

Note that the reason why we used the older Berchtold GE and HPRD PPIN data in all of our previous studies is because our group started the research direction of inferring and analyzing aging-specific subnetworks in 2012, with the first paper published in 2014 (Faisal and Milenković, 2014). The only aging-specific GE data for human that encompassed multiple ages and enough samples for each age at the time was that curated by Berchtold et al. via a microarray technology. Similarly, HPRD was a state-of-the-art human PPIN data at that time. Since then, we proposed a series of more advanced methods for inferring or analyzing aging-specific subnetworks (Li and Milenković, 2019; 2021; Li et al., 2021; Newaz and Milenković, 2020). For fairness of evaluation of every new method against previous ones, we kept the input (i.e. GE and PPIN) data the same. More recently, newer aging-specific GE data with enough ages and enough samples per age, produced using a next-generation sequencing technology (RNA-Seq), have become available, i.e. GE data from the Genotype-Tissue Expression (GTEx) project (GTEx Consortium, 2015). Also, newer, regularly updated PPIN data exist in the BioGRID database (Oughtred et al., 2021).

In this study, we primarily (although not only) examine whether using the newer GTEx GE or BioGRID PPIN data to infer a weighted dynamic aging-specific subnetwork will improve the quality of predicted aging-related genes compared to using the older Berchtold GE and HPRD PPIN data.

## 1.2 Our study and contribution

We summarize our study in Figure 1. To examine which of the two data components (GE data or PPIN data) might lead to better prediction accuracy, we construct four weighted dynamic aging-specific subnetworks by varying one data component at a time (Table 1). To construct and analyze (make predictions from) the four subnetworks, we rely on our well-established framework discussed above (Li and Milenković, 2021, Li et al., 2021).

Postulation 1 would be validated if using GTEx-BioGRID is better than using all of Berchtold-HPRD, Berchtold-BioGRID and GTEx-HPRD. Shockingly, we find this not to be the case. That is, generally, neither using newer PPIN data nor using newer GE data helps. To attempt to understand this finding, we perform several analyses, including the following.

All four subnetworks perform statistically significantly better than expected by chance, i.e. each subnetwork captures some of the aging-related knowledge. So, we aim to test whether the different subnetworks yield complementary or duplicated aging-related gene predictions. We find that the subnetworks are overall complementary to each other, with relatively low (although statistically significant) average pairwise overlaps of 34%. Hence, even though newer GE or PPIN data do not necessarily improve upon their older counterparts, they do capture aging-related knowledge that older data cannot.

However, our subnetworks are imperfect regardless of what data they are inferred from: the highest AUPR, precision, recall and F-score over all subnetworks are 50, 67, 45 and 52%, respectively. To understand this, we rely on existing data on human proteins' memberships in aging-related KEGG pathways (Kanehisa et al., 2021) and their Gene Ontology (GO) term annotations (Gene Ontology Consortium, 2021). We test the enrichment of four groups of aging- and non-aging-related genes in the aging-related pathways and GO terms. These four gene groups are presented in Table 2.

Predicted-Aging genes, being supported by at least one of our subnetworks as well as GenAge, are our positive control and should be enriched in aging-related pathways or GO terms. Similarly,

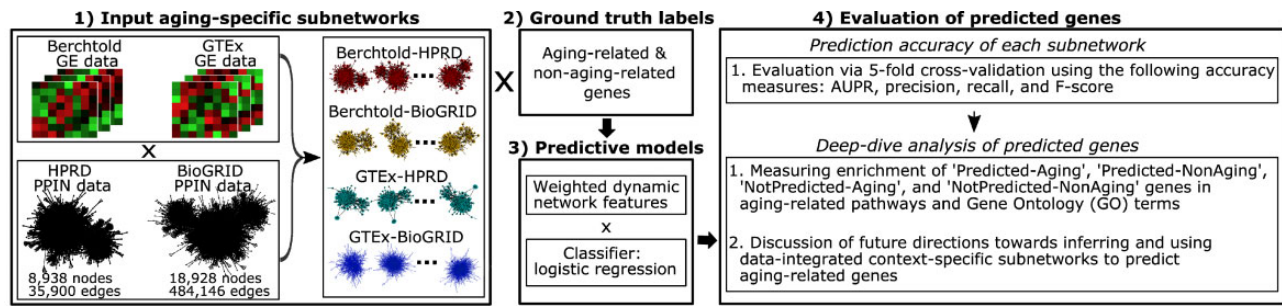


Fig. 1. Summary of our study. See the text for details

**Table 1.** Our four weighted dynamic aging-specific subnetworks (in bold) constructed from the four combinations of GE and PPIN data

PPIN\GE data	Berchtold (old)	GTEx (new)
HPRD (old)	<b>Berchtold-HPRD</b>	<b>GTEx-HPRD</b>
BioGRID (new)	<b>Berchtold-BioGRID</b>	<b>GTEx-BioGRID</b>

**Table 2.** The four considered groups of aging- and non-aging-related genes

	Ground truth label		
		Aging-related	Non-aging-related
Predicted as aging-related by $\geq 1$ subnetwork?	Yes	Predicted–Aging	Predicted–Non Aging
	No	NotPredicted–Aging	NotPredicted–NonAging

Note: The gene groups are named based on a combination of their predicted label and their ground truth label. Note that ‘Predicted’ are those genes predicted by at least one of the four subnetworks, while ‘NotPredicted’ are those genes not predicted as aging-related by any subnetwork.

NotPredicted–NonAging genes, being supported by neither our subnetworks nor GenAge, are our negative control and should not be enriched in any aging-related pathways and GO terms. Indeed, this is what we find. Then, we examine the enrichment of Predicted–NonAging and NotPredicted–Aging genes in the aging-related pathways and GO terms, for the following reasons.

Our suboptimal precision means that a portion of each subnetwork’s predictions are currently not known to be aging-related; these correspond to Predicted–NonAging genes. If these genes that are currently missing from GenAge are enriched in aging-related pathways or GO terms and are thus more similar to the positive control genes than the negative control genes, then this could imply that postulation 2 holds, i.e. that GenAge is incomplete. In this case, Predicted–NonAging genes could be considered as highly ranked candidates for future wet lab validation. Indeed, we find this to be the case—Predicted–NonAging genes are statistically significantly enriched in multiple aging-related pathways.

Our suboptimal recall means that each subnetwork fails to predict a portion of known aging-related genes from GenAge; these correspond to NotPredicted–Aging genes. If these genes are enriched in aging-related pathways or GO terms, it could mean that the aging process has at least two distinct network ‘signatures’: one signature shared between those genes from GenAge that can be captured by our subnetworks, and a different signature (or signatures) that the remaining genes from GenAge have but that our subnetworks cannot recognize. This would correspond to postulation 3, i.e. network methods used to infer a weighted dynamic aging-specific

subnetwork or to predict aging-related genes from an inferred subnetwork being unable to capture all of the current aging-related knowledge. Otherwise, if NotPredicted–Aging genes are not enriched in aging-related pathways or GO terms, this would mean that there might be discrepancies between GenAge aging-related ground truth data and aging-related pathway and GO term data, possibly because GenAge is noisy or because the different databases capture complementary aging-related knowledge. We find the former to be the case—NotPredicted–Aging genes are statistically significantly enriched in multiple aging-related pathways or GO terms.

## 2 Methods

### 2.1 Data

#### 2.1.1 Considered aging-specific GE data

The two considered aging-specific GE datasets are from Berchtold (Berchtold *et al.*, 2008) and GTEx (GTEx Consortium, 2015). Berchtold GE data encompasses 173 postmortem samples in brain from 55 individuals, which span 37 different ages from 20 to 99 years. GTEx GE data (version 8) encompasses 17 382 postmortem samples across 54 tissues from 175 individuals, which span six age groups: 20–29, 30–39, 40–49, 50–59, 60–69 and 70–79. Of the 54 tissues from 31 organs, 13 tissues are from the brain. These are tissues that we focus on, for fair comparison to Berchtold data that is solely brain-related. The 13 tissues encompass 2642 samples. See Supplementary Section S1.1 for more detail.

#### 2.1.2 Considered entire context-unspecific PPIN data

The two considered entire context-unspecific PPINs are from HPRD (Berchtold *et al.*, 2008) and BioGRID (version 4.4.197, the latest version at the time we started this study) (Oughtred *et al.*, 2021). We consider only the largest connected component of the two PPINs. The one from HPRD has 8938 proteins and 35 900 interactions. The largest connected component of human physical interactions from BioGRID has 18 928 proteins and 484 146 interactions.

#### 2.1.3 Inference of weighted dynamic aging-specific subnetworks

We infer multiple weighted dynamic aging-specific subnetworks, depending on which GE and PPIN data are used. Berchtold-HPRD is the benchmark subnetwork from our previous study (Li *et al.*, 2021). Berchtold-HPRD was inferred using a method called NetWalk (Komurov *et al.*, 2010), which was proved to be the best of two state-of-the-art network propagation algorithms for inferring aging-specific subnetworks in our previous study (Newaz and Milenković, 2020). For fairness, we use NetWalk to infer all other considered networks.

Namely, for an age (in case of Berchtold GE data) or age group (in case of GTEx GE data), NetWalk propagates GE levels onto the interactions in an entire context-unspecific PPIN via a biased random walk, and outputs a PPIN with two directed age-specific weights assigned to each interaction. That is, an interaction ( $u, v$ ) has a weight from  $u$  to  $v$  and another weight from  $v$  to  $u$ , see original publication (Komurov *et al.*, 2010) for methodological detail about NetWalk. The smaller of the two weights is selected for each

interaction, yielding a weighted static age-specific subnetwork snapshot for the given age (group). This way, we create one static weighted subnetwork snapshot for each of the considered ages or age groups.

Given all  $x$  weighted static age-specific subnetwork snapshots for the  $x$  considered ages or age groups, we obtain the final dynamic weighted aging-specific subnetwork (for a given combination of GE and PPIN data) using an established two-step approach (Li et al., 2021). First, to make the weights of the interactions comparable across snapshots, the weights are normalized over all  $x$  snapshots. Second, given the  $x$  snapshots with normalized interaction weights,  $x-1$  ‘differential’ snapshots are created; a ‘differential’ snapshot is created for each pair of consecutive normalized-weight snapshots  $i$  and  $i+1$  ( $i = 1, 2, 3, \dots, x-1$ ), where for each interaction, its ‘differential’ weight  $w_{i,i+1}$  is defined as the relative weight change between the snapshots  $i$  (i.e.  $w_i$ ) and  $i+1$  (i.e.  $w_{i+1}$ ):  $w_{i,i+1} = \frac{w_{i+1} - w_i}{w_{i+1} + w_i} \times 100$ . The collection of the  $x-1$  age-specific ‘differential’ snapshots in the increasing order of ages or age groups forms the final weighted dynamic aging-specific subnetwork (for a given combination of GE and PPIN data), see Supplementary Figure S12.

Because we aim to evaluate whether subnetworks inferred using newer GE data or newer PPIN data (or both) outperform the benchmark subnetwork (Berchtold-HPRD) that is inferred using both older GE data and older PPIN data, we need to systematically test which component (GE data or PPIN data) leads to potential performance improvement. To do so, in addition to Berchtold-HPRD, we infer three additional aging-specific subnetworks (also using NetWalk) by varying one data component at a time. That is, compared to Berchtold-HPRD, by changing the PPIN data component, we infer Berchtold-BioGRID; by changing the GE data component, we infer GTEx-HPRD; by changing both the GE and the PPIN data component, we infer GTEx-BioGRID. In addition, to examine whether grouping samples into individual (specific) ages versus age groups yields performance difference (Section 3.2), we group samples in Berchtold GE data into six age groups that match those in GTEx GE data (see above). Note that of the 173 samples in Berchtold GE data, we focus on all 119 samples that belong to one of the six age groups. Then, we combine this modified, age group-based Berchtold GE data with each of HPRD and BioGRID PPIN data using NetWalk to infer two additional subnetworks, referred to as Berchtold-HPRD-6 and Berchtold-BioGRID-6, respectively. Table 3 shows sizes of the six subnetworks.

Note that network propagation requires, for each age (group), expression levels for all genes in an entire context-unspecific PPIN. Also, because GE data encompasses multiple samples for each age (group), the expression level of a gene at a given age (group) needs to be somehow combined or normalized across these samples. We achieve these tasks as follows.

Berchtold GE data was curated using microarray technology (Affymetrix Hg-U133plus 2.0), containing expression information for 54 675 probes. Because nodes in our PPIN data have gene IDs, to integrate GE with PPIN data, we need to map probe IDs to gene IDs. To do this, we mimic our previous study (Newaz and Milenković, 2020). That is, by using the DAVID tool (Huang et al., 2009), we are able to map 48 724 of the 54 675 probes to 21 441 unique gene IDs. Next, while there are multiple samples (i.e. expression values) relevant to a gene at a given age, we need to obtain a

single score value for each gene at each age. We do this by assigning to a given gene the average expression value of all of its samples relevant to the given age. Moreover, we use the MAS 5.0 BioConductor package (Gentleman et al., 2004) that calculates the  $P$ -value of whether a gene that is present in Berchtold GE data is statistically significantly expressed (i.e. active) at a given age. This enables us to assign expression values to those genes that are present in the considered PPIN but not in Berchtold GE data at a given age. We assume such genes to not be active at a given age. This is why we assign them the average expression value of all non-active genes that are present in Berchtold GE data at the given age. For methodological details, see Supplementary Section S1.1.1 of Faisal and Milenković (2014).

GTEx GE data was curated using RNA-Seq technology. There are many methods proposed for normalizing RNA-Seq data. We use the Trimmed Mean of the M-values (TMM) (Robinson and Oshlack, 2010) because (i) TMM has been widely used for this purpose; and (ii) it was suggested by Evans et al. (2018) and Zhao et al. (2021) that TMM typically performs better than other existing normalization methods. In particular, TMM assumes that most of the genes are not differentially expressed, and the batch effect across multiple tissues or genotypes can be removed. Briefly, TMM works as follows. First, the log<sub>2</sub> fold change and absolute expression value of a sample against the reference sample are calculated. Then, those genes that have high fold changes or have large absolute expression values are trimmed, and the weighted average fold change per sample can be calculated. Finally, the read counts of genes are normalized by the weighted average fold change and the total number of samples in the GE data. For methodological details about TMM, see Robinson and Oshlack (2010). We use a BioConductor package called edgeR (Robinson et al., 2010) to obtain normalized GE data via TMM. Note that genes in GTEx GE data are represented as Ensembl IDs, and we convert these IDs to gene symbols using the DAVID tool. We remove from consideration a gene if its Ensembl ID shares more than one gene symbol. Unlike microarray-based GE data, we cannot obtain  $P$ -value for whether a gene is statistically significantly expressed at a given age group due to the lack of reference data. So, for those genes that are present in the considered PPINs but not in GTEx GE data at a given age group, we assign these genes an expression value of zero.

#### 2.1.4 Aging- and non-aging-related gene labels

For fairness in comparing the considered subnetworks, we consider all 8756 genes that are present in all of the subnetworks. Supervised classification requires ground truth labels, i.e. knowledge of which of the considered genes are aging- versus non-aging-related.

To define aging-related gene labels, as established (Li et al., 2021), we rely on a confident ground truth data source, GenAge (Tacutu et al., 2018). Human genes included in GenAge are sequence orthologs of aging-related genes in model species. All aging-related genes in model species are experimentally validated. Their human homologs are included in GenAge if they are supported by multiple pieces of literature and have aging-related phenotypical evidence in the human species. Of all 307 genes in GenAge, 277 of them are among the 8756 considered genes. We denote these 277 genes as *aging-related genes*.

Next, we define non-aging-related genes from the 8,756 – 277 = 8,479 genes that are not present in GenAge. To ensure that our non-aging-related labels are as confident as possible, we also rely on five other aging-related ground truth datasets curated by Jia et al. (2018), Lu et al. (2004), Berchtold et al. (2008) and Simpson et al. (2011); note that there are two gene sets from Jia et al. (2018), for genes that are upregulated and downregulated with age, respectively. That is, we remove all genes that are present in any of the five aging-related ground truth datasets from the 8479 genes. This leaves 4282 genes that we denote as *non-aging-related genes*. We refer to the 277 + 4282 = 4559 aging- and non-aging-related genes as *ground truth labeled genes*.

**Table 3.** The number of snapshots and size for each of the six considered weighted dynamic aging-specific subnetworks

Subnetwork	No. of snapshots	No. of nodes	No. of edges
Berchtold-HPRD	37	8 938	35 900
Berchtold-BioGRID	37	18 928	484 146
GTEx-HPRD	6	8 938	35 900
GTEx-BioGRID	6	18 928	484 146
Berchtold-HPRD-6	6	8 938	35 900
Berchtold-BioGRID-6	6	18 928	484 146

### 2.1.5 Aging-related pathways and GO terms

We focus on five confident aging-related pathways, i.e. adenosine monophosphate-activated kinase (AMPK), Sirtuin 1 (SIRT1), mammalian target of rapamycin (mTOR), phosphatidylinositol-3-kinase and protein kinase B (PI3K-AKT), and neurotrophin. Their relatedness to the aging process and lifespan was examined by numerous studies and was systematically reviewed by Yu *et al.* (2021), as follows.

- AMPK, SIRT1 and mTOR pathways are some of the key pathways that describe molecular mechanisms of the aging process. In particular, AMPK is considered as aging-related because it was found to control cell survival, growth, death and autophagy, and it can regulate cellular homeostasis and resistance to stress (Morgunova and Klebanov, 2019). SIRT1 is considered as aging-related because it was found to play an important role in improving oxidative stress resistance of cells and inhibiting cell death (Brunet *et al.*, 2004), and hence is involved in many age-related diseases (Zhao *et al.*, 2020). mTOR is considered aging-related because it functions as a sensor of intracellular energy and a central regulator of biological processes, including aging (Arriola Apelo and Lamming, 2016; Di Francesco *et al.*, 2018).
- PI3K-AKT pathway is considered as aging-related because it was identified as the signaling pathway of several aging-related diseases: diabetic encephalopathy (Wang *et al.*, 2018) and cancer (Porta *et al.*, 2014).
- Neurotrophin pathway is considered as aging-related because it was identified as the key signaling pathway of aging-related Parkinson's disease (Paudel *et al.*, 2020).

We obtain the genes for these five pathways from the KEGG database (Kanehisa *et al.*, 2021). We summarize the size of each pathway as in KEGG and as in our ground truth aging-related data in Table 4. In addition, we consider aging-related GO terms (Supplementary Section S1.2).

## 3 Results

As already discussed, we consider four weighted dynamic aging-specific subnetworks: Berchtold-HPRD, Berchtold-BioGRID, GTEX-HPRD and GTEX-BioGRID. We test whether GTEX-BioGRID, which is inferred using both newer GE and PPIN data, outperforms the remaining three subnetworks that are inferred using older GE data or older PPIN data (or both). To test this, first, for each subnetwork, we consider nine predictive models (i.e. feature-classifier combinations). We train and test each predictive model via 5-fold cross-validation in terms of prediction accuracy (AUPR, precision, recall and F-score averaged over the five folds). Then, we select the best predictive model that yields the highest AUPR, to give each subnetwork the best-case advantage (Section 3.1). For fairness, we force the gene sets that are randomly split into the training and testing data for cross-validation to be same for all predictive models across all four subnetworks (see Section 2.1.4 for details). Second, we compare the selected best predictive models to evaluate which subnetwork results in the highest

**Table 4.** The sizes of (numbers of genes in) the five aging-related pathways as available in KEGG and in our ground truth data

Pathway	Size in KEGG	Size in our ground truth data
AMPK	113	38
SIRT1	89	36
mTOR	165	41
PI3K-AKT	354	179
Neurotrophin	119	57

*Note:* That is, the latter is the number of genes from a given KEGG pathway that are present in our 4559 ground truth labeled genes.

prediction accuracy (Section 3.2). Third, we analyze whether the four subnetworks are predicting redundant or complementary aging-related genes (Section 3.3). Finally, we present a deep-dive analysis in terms of enrichment of the predicted genes in aging-related pathways and GO terms (Section 3.4).

### 3.1 Selecting the best predictive model for each subnetwork

The considered nine predictive models per subnetwork are combinations of nine features and one classifier. We consider the nine best features among 30 features evaluated in our previous study (Li *et al.*, 2021) (Supplementary Section S1.3). We use logistic regression as the classifier for all predictive models because it consistently performed the best among nine prominent classifiers evaluated in our previous studies (Li and Milenković, 2019, 2021; Li *et al.*, 2021).

For each of the four subnetworks, all nine predictive models perform statistically significantly better than expected by chance (adjusted  $P$ -values  $< 0.05$ ), with respect to all four accuracy measures (Supplementary Tables S1–S4 and Supplementary Figs S1–S4). For each subnetwork, the best-performing predictive model brings sufficient (although not always statistically significant) improvement compared to the remaining eight predictive models. This is why we choose the best-performing predictive model for a given subnetwork.

Overall, it is typically different predictive models that are selected for the different subnetworks. That is, the four subnetworks yield three distinct best predictive models. The same best predictive model is selected only for the two HPRD-based subnetworks (Berchtold-HPRD and GTEX-HPRD).

### 3.2 Comparing prediction accuracy of the four subnetworks

Given the selected best predictive model for each of the four subnetworks, we compare the prediction accuracy of the four subnetworks. With this, we aim to test postulation 1: whether using GTEX-BioGRID that is inferred from both newer GE and PPIN data would outperform all other three subnetworks that are inferred from older GE data or older PPIN data (or both).

First, we ask whether using newer BioGRID PPIN data is better than using its older HPRD counterpart when the GE data is fixed, i.e. whether Berchtold-BioGRID improves upon Berchtold-HPRD and whether GTEX-BioGRID improves upon GTEX-HPRD. We find neither of the two to hold (Fig. 2): Berchtold-BioGRID performs marginally worse than Berchtold-HPRD in terms of all four prediction accuracy measures, and GTEX-BioGRID performs marginally worse than GTEX-HPRD in terms of all four prediction accuracy measures.

Second, we ask whether using newer GTEX GE data is better than using its older Berchtold counterpart when the PPIN data is fixed, i.e. whether GTEX-HPRD improves upon Berchtold-HPRD and whether GTEX-BioGRID improves upon Berchtold-BioGRID. We find that these do not necessarily hold (Fig. 2). Namely, GTEX-HPRD performs marginally worse than Berchtold-HPRD in terms of precision and F-score, although marginally better in terms of AUPR and recall. GTEX-BioGRID performs marginally worse than Berchtold-BioGRID for all four prediction accuracy measures.

Overall, GTEX-BioGRID that is inferred from both newer GE and PPIN data does not outperform any subnetwork inferred from at least one older data component. That is, postulation 1 does not hold. In fact, Berchtold-HPRD that is inferred from both older GE and PPIN data marginally outperforms all other subnetworks in terms of precision and F-score. Meanwhile, Berchtold-HPRD is the second-best subnetwork in terms of AUPR and recall.

### 3.3 Examining prediction overlaps between the subnetworks

Thus far, we have found that using both newer GE and PPIN data does not necessarily improve the prediction accuracy upon using older GE data or older PPIN data (or both); also, all four

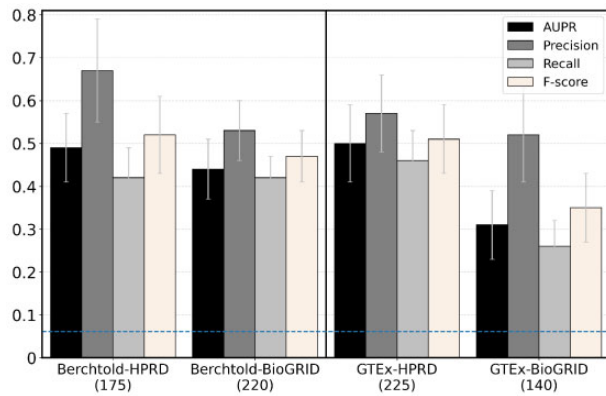


Fig. 2. The prediction accuracy in terms of AUPR, precision, recall and F-score of the four weighted dynamic aging-specific subnetworks, each under its best predictive model. The number below each subnetwork name represents the number of genes that are predicted as aging-related by the corresponding subnetwork. The horizontally dashed line indicates the prediction accuracy scores expected by chance, i.e. the fraction of all genes in the ground truth data that are labeled as aging-related

subnetworks yield prediction accuracies that are statistically significantly better than at random. Given this, we ask whether our four subnetworks are predicting complementary or redundant aging-related knowledge by examining overlaps of their predictions. Namely, we split predicted aging-related genes of each subnetwork into two groups: true positives (i.e. genes predicted as aging-related and labeled as aging-related in the ground truth data) and novel predictions (i.e. genes predicted as aging-related but are currently labeled as non-aging-related in the ground truth data); the former is more confident than the latter to actually be aging-related.

For true positive predictions, we find that all pairwise overlaps are statistically significantly high (i.e. adjusted  $P$ -values  $< 0.05$ ) but in reality still quite complementary. Namely, the largest, smallest and average Jaccard indices of all pairwise overlaps are 76.3, 35.7 and 51.3%, respectively (Fig. 3). Moreover, each of Berchtold-HPRD, Berchtold-BioGRID, GTEX-HPRD and GTEX-BioGRID predicts as aging-related some of the genes from GenAge that none of the other three subnetworks predict. These statistics are 5 (out of 175) predictions for Berchtold-HPRD, 13 (out of 220) predictions for Berchtold-BioGRID, 10 (out of 225) predictions for GTEX-HPRD, and 4 (out of 140) predictions for GTEX-BioGRID.

So, even though using newer GE or PPIN data does not improve upon using their older counterparts, the former does capture known aging-related knowledge that using older data does not. The opposite also holds. This means that the different data, i.e. subnetworks, yield complementary predictions.

Results for novel predictions are qualitatively similar to the results for true positive predictions, i.e. all overlaps are statistically significantly high but in reality still quite complementary (Supplementary Fig. S6). Namely, the largest, smallest and average Jaccard indices of all pairwise overlaps are 42.2, 5.9 and 16.9%, respectively. Moreover, each of Berchtold-HPRD, Berchtold-BioGRID, GTEX-HPRD and GTEX-BioGRID make some novel predictions that none of the other three subnetworks make. These statistics are 10 (out of 175) predictions for Berchtold-HPRD, 60 (out of 220) predictions for Berchtold-BioGRID, 39 (out of 225) predictions for GTEX-HPRD and 35 (out of 140) predictions for GTEX-BioGRID.

### 3.4 Validating predictions using aging-related pathways and GO terms

Although all four subnetworks yield statistically significantly high prediction accuracies, their accuracies are suboptimal, i.e. the highest accuracy score is ‘only’ 67%. To understand such suboptimal performance, we rely on existing data on human proteins’ memberships in aging-related KEGG pathways and their aging-related GO term annotations. As typically done, we analyze whether our

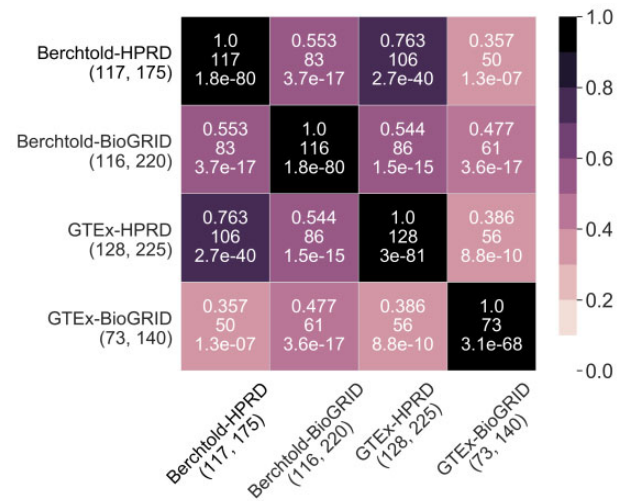


Fig. 3. Pairwise overlaps in terms of Jaccard indices of true positives for each pair of the four considered subnetworks. By true positives, we mean genes that are predicted as aging-related and are also present in GenAge. The two numbers in the parenthesis below each subnetwork name represent the number of true positives and the number of all predicted aging-related genes for the given subnetwork, respectively. In a cell, corresponding to a pair of subnetworks, the three numbers represent the Jaccard index (top), the raw number of genes in the overlap (middle) and the adjusted  $P$ -value indicating whether the overlap is statistically significantly high. The color shades are driven by Jaccard indices, where a darker color means a higher Jaccard index. For overlaps of novel predictions (rather than true positives), see Supplementary Figure S6

predictions are statistically significantly enriched in aging-related pathways or GO terms. To do so, we split the genes from the ground truth data into four groups (Predicted–Aging, Predicted–NonAging, NotPredicted–Aging and NotPredicted–NonAging genes, as defined in Table 2) based on whether their predicted labels agree with their ground truth labels (Section 1.2). Then, we test the enrichment of each gene group in aging-related pathways and GO terms. We use all five established aging-related pathways (Yu et al., 2021) and all 13 aging-related GO terms that annotate sufficiently many (at least three aging- or non-aging-related) of our ground truth labeled genes.

We find that the Predicted–Aging genes are significantly enriched (adjusted  $P$ -values  $< 0.05$ ) in all aging-related pathways (Fig. 4) and over half of the aging-related GO terms (Supplementary Fig. S7). This is expected, as Predicted–Aging genes are supported as aging-related by both our subnetworks and the ground truth data and are thus our positive control (i.e. genes that are the most confident to be aging-related).

On the hand, NotPredicted–NonAging genes are not significantly enriched (adjusted  $P$ -values = 1.0) in any aging-related pathway or GO term (Fig. 4 and Supplementary Fig. S7). This is also expected, as NotPredicted–NonAging genes are not supported as aging-related by either our subnetworks or the ground truth data and are thus our negative control.

Next, we look into Predicted–NonAging, the gene group that causes suboptimal precision. We find this group of genes to be significantly enriched (adjusted  $P$ -values  $< 0.05$ ) in three aging-related pathways (AMPK, PI3K-AKT and Neurotrophin, Fig. 4). That is, statistically significantly many Predicted–NonAging genes are linked to the aging process. In other words, postulation 2 seems to hold, i.e. GenAge appears to be incomplete. So, the novel predictions from our subnetworks could potentially guide future wet lab experiments to expand the current GenAge aging-related ground truth data.

Finally, we look into NotPredicted–Aging, the gene group that causes suboptimal recall. We find this group of genes to be significantly enriched in an aging-related pathway (PI3K-AKT, Fig. 4) and two aging-related GO terms (GO:0007568—aging and GO:0008340—determination of adult lifespan, Supplementary Fig. S7). That is, these GenAge genes that are not captured by our subnetworks are indeed

AMPK(38)	0.085 16	0.024 6	0.011 1	0.002 6
	2.8e-13	0.023	0.56	1
SIRT(36)	0.097 18	0.008 2	0.023 2	0.002 7
	1.9e-16	0.76	0.13	1
mTOR(41)	0.072 14	0.015 4	0.0 0	0.004 14
	2.7e-10	0.24	1	1
PI3K-AKT(179)	0.134 41	0.047 18	0.045 10	0.024 78
	7.5e-22	0.007	0.00012	1
Neurotrophin(57)	0.155 30	0.026 7	0.018 2	0.003 9
	7.1e-28	0.044	0.24	1
	Predicted-Aging (167)	Predicted-NonAging (222)	NotPredicted-Aging (54)	NotPredicted-NonAging (3147)

Fig. 4. Enrichments of the four gene groups (x-axis) in the five aging-related pathways (y-axis). The number below a pathway name or a gene group name represents the gene count in the pathway or gene group. In each cell, the three numbers represent the overlap size as measured by the Jaccard index (top), the raw number of genes in the overlap (middle) and the adjusted  $P$ -value indicating whether the overlap size is statistically significantly high, i.e. whether the given gene group is statistically significantly enriched in the given pathway. The adjusted  $P$ -values below 0.05 are highlighted in red. Note that there are 4559 genes in the ground truth data. However, the total number of genes over these four gene groups is 3590. This is because for NotPredicted–Aging and NotPredicted–NonAging, we only consider those genes that are not predicted as aging-related by any of the four subnetworks. Analogous results for aging-related GO terms are shown in [Supplementary Figure S7](#)

aging-related. This suggests that there may be multiple network ‘signatures’ of aging-related genes: one that is shared between GenAge aging-related genes and our subnetworks, and other signature(s) that are not recognized by our subnetworks. In other words, postulation 3 seems to hold, i.e. our subnetworks or network features might have failed to recognize all network ‘signatures’ of GenAge aging-related genes.

In summary, these analyses show that the enrichments of Predicted–Aging, Predicted–NonAging and NotPredicted–Aging genes in aging-related pathways and GO terms are statistically significantly high, i.e. these are confident aging-related genes. On the other hand, the enrichments of NotPredicted–NonAging genes are not statistically significant, i.e. these are confident non-aging-related genes.

## 4 Discussion

In this study, we systematically evaluate whether GTEx-BioGRID that is inferred from both newer GE and PPIN data, improves the prediction accuracy upon its counterparts inferred from at least one older data component. We discuss our findings and provide future directions in this section.

### 4.1 The choice of best predictive model matters

The choice of predictive model for a given subnetwork matters (Section 3.1). The fact that no model performs the best for all

subnetworks stresses the need to test multiple models, to give each subnetwork the best-case advantage.

Also, given the fact that the four subnetworks are predicting complementary aging-related knowledge (Section 3.3), it might be worth to pursue development of a novel network feature or an ensemble learning approach that would integrate the complementary aspect of the four subnetworks.

### 4.2 Higher-quality PPIN or GE data might be needed

When GE data is fixed, subnetworks inferred using newer BioGRID PPIN data never outperform subnetworks inferred using older HPRD PPIN data (Section 3.2). The superiority of using HPRD PPIN data may be due to the following. HPRD is manually curated (i.e. ‘read, analyzed and interpreted by expert biologists’ according to the HPRD website). As such, it contains interactions that are confident ([Lazareva et al., 2021](#)). Nonetheless, these interactions encompass fewer genes than newer PPIN data. On the other hand, PPIN databases that are not manually curated, including BioGRID, could potentially introduce technical or literature biases when newer interactions are added ([Lazareva et al., 2021](#)). That is, those proteins that are used as ‘baits’ for discovering new interactions ([Lazareva et al., 2021](#)), have known biological functions, are related to common diseases, etc., along with their interactions, may be studied more than the remaining proteins and their interactions. Henceforth, newer, regularly updated PPIN databases such as BioGRID might continuously increase their bias towards uncovering biology primarily about the aforementioned types of proteins and their interactions.

When PPIN data is fixed, subnetworks inferred using newer GTEx GE data do not necessarily outperform subnetworks inferred using older Berchtold GE data (Section 3.2). This may be because of a variation between Berchtold and GTEx GE data: the samples in Berchtold GE data are grouped across 37 specific ages in the [20–99] range, while the samples in GTEx GE data are grouped across six age groups (20–29, 30–39, 40–49, 50–59, 60–69, 70–79). So, the resulting Berchtold-based subnetworks have 37 snapshots while the GTEx-based subnetworks have six snapshots. To test whether the differences in the number or nature of the snapshots between Berchtold-based and GTEx-based subnetworks yield differences in their accuracy, we further infer two subnetworks, as follows. First, we group samples in Berchtold GE data into the same six age groups as in GTEx GE data. Then, we integrate the modified Berchtold GE data with each of HPRD and BioGRID PPIN data, to infer two new subnetworks, each with six snapshots; we refer to these as Berchtold-HPRD-6 and Berchtold-BioGRID-6, respectively. This way, both the number and nature of snapshots are matched between Berchtold-HPRD-6 and GTEx-HPRD, as well as between Berchtold-BioGRID-6 and GTEx-BioGRID. Thus, comparison may be more fair between Berchtold-HPRD-6 and GTEx-HPRD than between Berchtold-HPRD and GTEx-HPRD. Similarly, comparison may be more fair between Berchtold-BioGRID-6 and GTEx-BioGRID than between Berchtold-BioGRID and GTEx-BioGRID.

When we compare Berchtold-HPRD-6 and GTEx-HPRD, as well as Berchtold-BioGRID-6 and GTEx-BioGRID, i.e. even when the age groups match between Berchtold and GTEx GE data, using newer GTEx GE and BioGRID PPIN data still does not improve accuracy compared to using at least one older data component ([Supplementary Fig. S5](#)). That is, GTEx-BioGRID is still inferior among all four subnetworks with six age groups (i.e. snapshots). In fact, Berchtold-HPRD-6 inferred using both older GE and PPIN data performs marginally better than the other three six-snapshot subnetworks. Therefore, the difference in the nature of Berchtold and GTEx GE data does not account for why newer data does not help.

A side observation of this analysis is as follows. By comparing Berchtold-HPRD versus Berchtold-HPRD-6 (both of which happen to be the best-performing networks in their respective evaluations), we can evaluate the effect of using specific ages versus using age groups, respectively, to construct aging-specific subnetworks. We find that the two subnetworks yield almost indistinguishable prediction accuracies ([Supplementary Fig. S5](#)). Thus, the aforementioned

effect seems to be minimal, i.e. it does not seem to matter whether specific ages or age groups are used. The implications of this result are as follows. When in GE data the samples are already provided as being grouped across age groups, without having information about their specific ages, one clearly has to use age groups. On the other hand, when in GE data the samples' specific ages are known, one could use the specific ages directly or they could instead first group the ages into age groups and then use the groups. For the GE data we have evaluated, the choice between the two does not seem to matter in terms of accuracy. But the same might not necessarily hold for different GE data. So, ideally, one should empirically evaluate whether it is better to use specific ages or somehow form age groups. However, it is unclear how exactly to define age groups, e.g. whether one should define an age group (i.e. construct an age group-specific snapshot) for every 5, 10 or more years of the human lifespan, as well as whether age groups (i.e. snapshots) should span age intervals of equal length or of differing lengths (Newaz et al., 2020). A possible way to address these challenges could be to start with the network snapshots corresponding to specific ages and then use algorithms such as SCOUT (Hulovatyy and Milenković, 2016) to computationally identify age groups. Namely, SCOUT is able to identify  $n$  time points (in our context, ages) in the entire time interval where network structure (with respect to some network structural property, in SCOUT's case, community or clustering structure) significantly changes. This corresponds to identifying  $n + 1$  temporal segments (in our context, age groups) that are separated by these  $n$  changing time points and the corresponding  $n + 1$  network snapshots. The intuition here is that specific ages (that are consecutive in GE data) which have similar community (or other network) structures are likely to have similar biological 'signatures' and hence should belong to the same age group. See the SCOUT paper (Hulovatyy and Milenković, 2016) for more details.

While it is still unclear why newer GTEx GE data does not outperform older Berchtold GE data, the former does offer opportunities that the latter does not. While Berchtold GE data only encompasses samples from the brain, GTEx GE data encompasses samples from 31 organs. This, combined with the fact that GTEx-based subnetworks perform significantly better than expected by chance and only marginally worse than Berchtold-based subnetworks, means that GTEx GE data allows for studying aging in tissues other than the brain.

Finally, we comment on an additional opportunity of GE data. GTEx GE data was curated using RNA-seq in different tissues. More recently, even newer biotechnology, single-cell RNA-seq, has allowed for curation of the Human Cell Atlas data that allows for even more detailed investigations of the human aging process—in different cell types (Uyar et al., 2020).

### 4.3 More complete aging-related ground truth data might be needed

The fact that Predicted-Aging and Predicted-NonAging genes are captured (predicted as aging-related) by our subnetworks implies that these two gene groups have similar network topologies (features). We aim to illustrate this by visualizing these genes' features in 2-dimensional (2D) vector space as described in Supplementary Section S1.5. Their features are indeed close in the 2D space, i.e. are topologically similar to each other in our subnetworks (Fig. 5 and Supplementary Figs S8–S11). The 2D visualization further supports postulation 2, i.e. GenAge might be incomplete. This may not be surprising, because human genes in GenAge are mostly orthologs of aging-related genes in model species. That is, GenAge encompasses primarily the aging-specific 'biology' that is common to the human species and model species. However, each species, especially human, has its unique aging-specific 'biology' that is not shared with other species (Bronikowski et al., 2011).

Thus, there is a need for more complete, human-specific knowledge on which genes are linked to aging. To the best of our understanding, the only such data come from GE studies such as the Berchtold study (Berchtold et al., 2008), GTEx project (Jia et al., 2018) or the Human Cell Atlas (Uyar et al., 2020). But we already

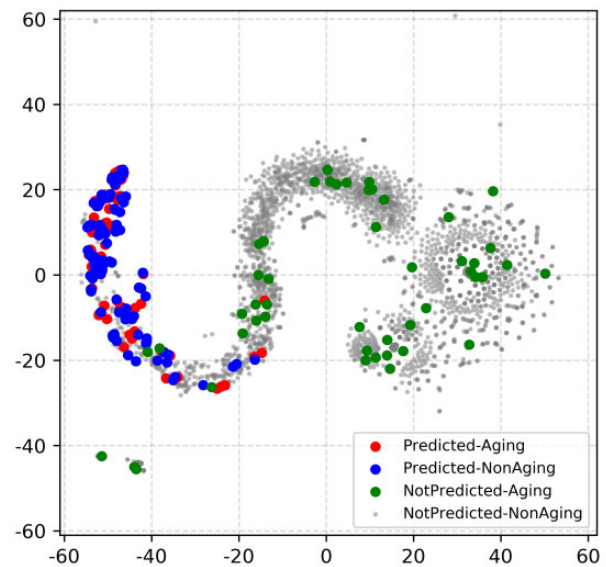


Fig. 5. Illustration of topological (dis)similarities between Predicted-Aging, Predicted-NonAging, NotPredicted-Aging and NotPredicted-NonAging genes in a given subnetwork by embedding their features into 2D space. This figure is for Berchtold-BioGRID; similar embedding trends hold for the other three subnetworks (Supplementary Figs S8, S10 and S11). This is a zoomed-in visualization with outliers removed for simplicity (see Supplementary Fig. S9 for the full visualization). When mapping features into 2D space, we have tested *t*SNE and PCA and selected the visualization with the clearest pattern. This figure is generated using *t*SNE

use such GE data to infer our subnetworks in the first place. Therefore, we cannot use that same data as the ground truth knowledge on which genes are linked to aging as well, as this would yield a circular argument. A potential solution for this challenge is to combine the genes in GenAge with those that are members of the aging-related pathways or are annotated by aging-related GO terms into a more comprehensive aging-related ground truth gene set. Another potential although imperfect solution might be to rely on predicted novel aging-related genes from existing computational studies, e.g. those predicted via network clustering (Hulovatyy and Milenković, 2016), network alignment (Faisal et al., 2015), unsupervised analyses of aging-specific subnetworks (Newaz and Milenković, 2020) or supervised analyses of aging-specific subnetworks. Note that these computationally predicted novel aging-related genes should not be used to validate prediction methods that share similar properties as the methods that generated the predictions, in order to avoid a circular argument. In this study, this would mean adding Predicted-NonAging genes to the ground truth knowledge, as these can be viewed as confident candidates for future wet lab validation.

### 4.4 Advanced algorithms for network inference or analysis might be needed

The fact that NotPredicted-Aging and NotPredicted-NonAging genes are predicted as *not* being aging-related implies that these genes' network features differ from network features of Predicted-Aging and Predicted-NonAging genes. This is not surprising for NotPredicted-NonAging genes, as these are our negative control genes (Section 3.4). What might appear surprising is that NotPredicted-Aging genes, which are aging-related according to the ground truth data, seem to have network features that are more similar to those of NotPredicted-NonAging genes (because NotPredicted-NonAging genes are not predicted as aging-related) than those of Predicted-Aging and Predicted-NonAging (which are predicted as aging-related). Indeed, we confirm the network feature (dis)similarities between the different gene groups in the 2D visualization (Fig. 5 and Supplementary Figs S8–S11). This further supports postulation 3 that our inferred subnetworks or predictive



models for their analyses might have failed to recognize all network ‘signatures’ of GenAge aging-related genes, i.e. that they have failed to identify as aging-related the ‘signatures’ of NotPredicted–Aging genes.

This result suggests a need for more advanced subnetwork inference or analysis methods. To infer our subnetworks, we relied on network propagation. A potential solution to improve subnetwork inference is to examine whether other types of integrative algorithms, such as kernel-, Bayesian- or non-negative matrix factorization-based methods (Gligorijević and Pržulj, 2015), would result in subnetworks that yield higher prediction accuracy or uncover complementary aspects of aging-related knowledge. A potential solution to improve network analysis is to use an ensemble learning approach that may be able to integrate complementary aspects of different subnetworks (Sagi and Rokach, 2018), as suggested in Section 4.1. Another potential solution is to use deep learning (e.g. graph convolutional networks) for dynamic network analysis (Pareja et al., 2020) to automatically learn and capture the differing types of network ‘signatures’ of genes of interest.

## 5 Conclusions

In this study, we have systematically evaluated whether using newer GE or PPIN data curated via more advanced biotechnologies to construct an aging-related subnetwork would improve the accuracy of predicting aging-related genes from the subnetworks. Unexpectedly, we have found that using newer GE or PPIN data does not help compared to using older data. In fact, the subnetwork inferred from both older Berchtold GE data and older HPRD PPIN data marginally outperforms the other subnetworks inferred from newer Berchtold GE data or newer HPRD PPIN data (or both). We have performed several analyses to try to explain this, without a clear answer. We have provided guidance on future directions towards the advancement of computational aging research, including a need for more complete aging-related ground truth data as well as more advanced subnetwork inference and analysis methods.

## Author contributions

Q.L. and T.M. designed the study. Q.L. carried out all computational experiments. Q.L. and T.M. analyzed the results. K.N. helped provide guidance on how to process Berchtold and GTEx GE data and a part of the code needed to pre-process the data. Q.L. drafted the initial paper, and all authors contributed to the writing of the final version. All authors approved the manuscript. T.M. supervised all aspects of the study.

## Funding

This work was supported by funds from the National Science Foundation [NSF CAREER CCF-1452795].

*Conflict of Interest:* none declared.

## References

Arriola Apelo,S.I. and Lamming,D.W. (2016) Rapamycin: an inhibitor of aging emerges from the soil of easter island. *GERONA*, 71, 841–849.

Berchtold,N.C. et al. (2008) Gene expression changes in the course of normal brain aging are sexually dimorphic. *Proc. Natl. Acad. Sci. USA*, 105, 15605–15610.

Bolignano,D. et al. (2014) The aging kidney revisited: a systematic review. *Ageing Res. Rev.*, 14, 65–80.

Bronikowski,A.M. et al. (2011) Aging in the natural world: comparative data reveal similar mortality patterns across primates. *Science*, 331, 1325–1328.

Brunet,A. et al. (2004) Stress-dependent regulation of FOXO transcription factors by the SIRT1 deacetylase. *Science*, 303, 2011–2015.

Chen,B. et al. (2014) Identifying protein complexes and functional modules—from static PPI networks to dynamic PPI networks. *Brief. Bioinform.*, 15, 177–194.

Di Francesco,A.D. et al. (2018) Intermittent mTOR inhibition reverses kidney aging in old rats. *J. Gerontol. A Biol. Sci. Med. Sci.*, 73, 843–844.

Elhesha,R. et al. (2019) Identification of co-evolving temporal networks. *BMC Genomics*, 20, 434.

Emanuel,E.J. et al. (2000) What makes clinical research ethical? *JAMA*, 283, 2701–2711.

Evans,C. et al. (2018) Selecting between-sample RNA-Seq normalization methods from the perspective of their assumptions. *Brief. Bioinform.*, 19, 776–792.

Fabris,F. et al. (2016) An extensive empirical comparison of probabilistic hierarchical classifiers in datasets of ageing-related genes. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, 13, 1045–1058.

Fabris,F. et al. (2017) A review of supervised machine learning applied to ageing research. *Biogerontology*, 18, 171–188.

Faisal,F.E. and Milenković,T. (2014) Dynamic networks reveal key players in aging. *Bioinformatics*, 30, 1721–1729.

Faisal,F.E. et al. (2015) Global network alignment in the context of aging. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, 12, 40–52.

Fang,Y. et al. (2013) Classifying aging genes into DNA repair or non-DNA repair-related categories. In: *International Conference on Intelligent Computing*. Berlin, Heidelberg: Springer, pp. 20–29.

Ferrucci,L. et al. (2020) Measuring biological aging in humans: a quest. *Ageing Cell*, 19, e13080.

Freitas,A.A. et al. (2011) A data mining approach for classifying DNA repair genes into ageing-related or non-ageing-related. *BMC Genomics*, 12, 27.

Gentleman,R.C. et al. (2004) Bioconductor: open software development for computational biology and bioinformatics. *Genome Biol.*, 5, R80–16.

Gligorijević,V. and Pržulj,N. (2015) Methods for biological data integration: perspectives and challenges. *J. R. Soc. Interface*, 12, 20150571.

Holtman,I.R. et al. (2015) Induction of a common microglia gene expression signature by aging and neurodegenerative conditions: a co-expression meta-analysis. *Acta Neuropathol. Commun.*, 3, 31.

Huang,D.W. et al. (2009) Extracting biological meaning from large gene lists with DAVID. *Curr. Protoc. Bioinform.*, 27, 13–11.

Hulovatyy,Y. and Milenković,T. (2016) SCOUT: simultaneous time segmentation and community detection in dynamic networks. *Sci. Rep.*, 6, 1–11.

Jia,K. et al. (2018) An analysis of aging-related genes derived from the genotype-tissue expression project (GTEx). *Cell Death Discov.*, 4, 26.

Kanehisa,M. et al. (2021) KEGG: integrating viruses and cellular organisms. *Nucleic Acids Res.*, 49, D545–D551.

Kerepesi,C. et al. (2018) Prediction and characterization of human ageing-related proteins by using machine learning. *Sci. Rep.*, 8, 4094.

Komurov,K. et al. (2010) Use of data-biased random walks on graphs for the retrieval of context-specific networks from genomic data. *PLoS Comput. Biol.*, 6, e1000889.

Lazareva,O. et al. (2021) On the limits of active module identification. *Brief. Bioinform.*, 22, bbab066.

Leiserson,M.D. et al. (2015) Pan-cancer network analysis identifies combinations of rare somatic mutations across pathways and protein complexes. *Nat. Genet.*, 47, 106–114.

Li,Q. and Milenković,T. (2019) Supervised prediction of aging-related genes from a context-specific protein interaction subnetwork. In: *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, San Diego, CA, USA, pp. 130–137.

Li,Q. and Milenković,T. (2021) Supervised prediction of aging-related genes from a context-specific protein interaction subnetwork. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, 19, 2484–2498.

Li,Q. et al. (2021) Improved supervised prediction of aging-related genes via weighted dynamic network analysis. *BMC Bioinformatics*, 22, 1–26.

Liguori,I. et al. (2018) Oxidative stress, aging, and diseases. *Clin. Interv. Aging.*, 13, 757–772.

Lu,T. et al. (2004) Gene regulation and DNA damage in the ageing human brain. *Nature*, 429, 883–891.

Morgunova,G.V. and Klebanov,A.A. (2019) Age-related AMP-activated protein kinase alterations: from cellular energetics to longevity. *Cell Biochem. Funct.*, 37, 169–176.

Newaz,K. and Milenković,T. (2020) Inference of a dynamic aging-related biological subnetwork via network propagation. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 19, 974–988.

Newaz,K. et al. (2020) Network analysis of synonymous codon usage. *Bioinformatics*, 36, 4876–4884.

Oughtred,R. et al. (2021) The BioGRID database: a comprehensive biomedical resource of curated protein, genetic, and chemical interactions. *Protein Sci.*, 30, 187–200.

- Pareja, A. et al. (2020) Evolvecgn: evolving graph convolutional networks for dynamic graphs. *Proc. AAAI Conf. Artif. Intell.*, **34**, 5363–5370.
- Paschos, G.K. et al. (2012) Obesity in mice with adipocyte-specific deletion of clock component ARNTL. *Nat. Med.*, **18**, 1768–1777.
- Paudel, Y.N. et al. (2020) Emerging neuroprotective effect of metformin in Parkinson's disease: a molecular crosstalk. *Pharmacol. Res.*, **152**, 104593.
- Porta, C. et al. (2014) Targeting PI3K/Akt/mTOR signaling in cancer. *Front. Oncol.*, **4**, 64.
- Prasad, T. et al. (2009) Human protein reference database—2009 update. *Nucleic Acids Res.*, **37**, D767–D772.
- Robinson, M.D. and Oshlack, A. (2010) A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol.*, **11**, R25–9.
- Robinson, M.D. et al. (2010) edgeR: a bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, **26**, 139–140.
- Rodríguez-Rodero, S. et al. (2011) Aging genetics and aging. *Aging Dis.*, **2**, 186–195.
- Sagi, O. and Rokach, L. (2018) Ensemble learning: a survey. *Wiley Interdiscipl. Rev. Data Min. Knowledge Discov.*, **8**, e1249.
- Simpson, J.E. et al.; MRC Cognitive Function and Ageing Neuropathology Study Group. (2011) Microarray analysis of the astrocyte transcriptome in the aging brain: relationship to Alzheimer's pathology and APOE genotype. *Neurobiol. Aging*, **32**, 1795–1807.
- Tacutu, R. et al. (2018) Human ageing genomic resources: new and updated databases. *Nucleic Acids Res.*, **46**, D1083–D1090.
- The Gene Ontology Consortium. (2021) The Gene Ontology resource: enriching a gold mine. *Nucleic Acids Res.*, **49**, D325–D334.
- The GTEx Consortium. et al. (2015) The Genotype-Tissue Expression (GTEx) pilot analysis: multitissue gene regulation in humans. *Science*, **348**, 648–660.
- Uyar, B. et al. (2020) Single-cell analyses of aging, inflammation and senescence. *Ageing Res. Rev.*, **64**, 101156.
- Wang, B. et al. (2018) Autophagy of macrophages is regulated by PI3k/Akt/mTOR signalling in the development of diabetic encephalopathy. *Aging (Albany NY)*, **10**, 2772–2782.
- Yu, M. et al. (2021) Key signaling pathways in aging and potential interventions for healthy aging. *Cells*, **10**, 660.
- Zhao, L. et al. (2020) Sirtuins and their biological relevance in aging and age-related diseases. *Aging Dis.*, **11**, 927–945.
- Zhao, Y. et al. (2021) TPM, FPKM, or normalized counts? A comparative study of quantification measures for the analysis of RNA-seq data from the NCI patient-derived models repository. *J. Transl. Med.*, **19**, 1–15.