

Predicting the pro-longevity or anti-longevity effect of model organism genes with enhanced Gaussian noise augmentation-based contrastive learning on protein–protein interaction networks

Ibrahim Alsaggaf ¹, Alex A. Freitas ² and Cen Wan ^{1,*}

¹School of Computing and Mathematical Sciences, Birkbeck, University of London, WC1E 7HX, London, UK

²School of Computing, University of Kent, CT2 7FS, Canterbury, Kent, UK

*To whom correspondence should be addressed. Tel: +44 02039261273; Email: cen.wan@bbk.ac.uk

Abstract

Ageing is a highly complex and important biological process that plays major roles in many diseases. Therefore, it is essential to better understand the molecular mechanisms of ageing-related genes. In this work, we proposed a novel enhanced Gaussian noise augmentation-based contrastive learning (EGsCL) framework to predict the pro-longevity or anti-longevity effect of four model organisms' ageing-related genes by exploiting protein–protein interaction (PPI) networks. The experimental results suggest that EGsCL successfully outperformed the conventional Gaussian noise augmentation-based contrastive learning methods and obtained state-of-the-art performance on three model organisms' predictive tasks when merely relying on PPI network data. In addition, we use EGsCL to predict 10 novel pro-/anti-longevity mouse genes and discuss the support for these predictions in the literature.

Introduction

Ageing is a highly complex biological process that involves many genes and biological pathways (1,2), and despite significant progress in ageing-biology research, the precise molecular mechanisms of ageing are still not well understood (2–4). In addition, ageing research is particularly important because ageing is a major driving factor for many diseases (5–7), and so a better understanding of the effects of ageing-related genes could lead to new therapies that would potentially extend not only the longevity but also the healthspan (period of health life) of individuals (6,8,9). With the help of Artificial Intelligence (more specifically, machine learning), research has been carried out to predict new ageing-related genes or biomarkers and to identify ageing-related biological pathways or processes (10,11). In this work, we focus on predicting the pro-longevity or anti-longevity effect of genes from four model organisms in ageing research (mouse, worm, fly and yeast). We cast this problem as a classification task from the perspective of supervised machine learning, where each instance (example) represents an ageing-related gene, each instance's class label indicates whether that gene has a pro-longevity or anti-longevity effect on the lifespan of an organism (12,13)—based on such class labels as recorded in the GenAge database (14). The predictive features are protein–protein interaction (PPI) network-based features.

PPI networks are a type of biologically meaningful and relevant features that have been widely used in multiple bioinformatics tasks like protein function prediction (15–17) and

disease–gene association prediction (18–20). PPI networks have also been used for ageing research. Freitas *et al.* (21) first exploited PPI networks as a type of features to classify DNA repair genes into ageing-related or non-ageing-related genes. Fang *et al.* (22) classified ageing-related genes into DNA repair or non-DNA repair-related genes using PPI networks-based features. This type of features were also used for predicting ageing-related genes for flies (23), mice (24) and humans (25). More recently, Magdaleno *et al.* (26) exploited PPI network features to predict ageing-related genes' dietary restriction associations, and Ribeiro *et al.* (27) used PPI network features to predict lifespan-extending chemical compounds for worms.

In this work, we propose a new contrastive learning-based framework to cope with PPI network features by developing two novel contrastive learning algorithms. In general, contrastive learning aims to learn a type of discriminative distribution where similar instances are pulled closer whilst different instances are pushed away. The conventional self-supervised contrastive learning methods like SimCLR (28) first create two views for each instance by using different data augmentation strategies. For each target instance, two views that are generated from that target instance are treated as positive views, and all other views that are not generated from that target instance are treated as negative views. Then SimCLR optimizes the network parameters to reduce the distance between two positive views, whilst enlarging the difference between positive views and negative views. The self-supervised learning paradigm was further extended to the supervised

Received: July 10, 2024. Revised: October 23, 2024. Editorial Decision: October 28, 2024. Accepted: October 29, 2024

© The Author(s) 2024. Published by Oxford University Press on behalf of NAR Genomics and Bioinformatics.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License

(<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

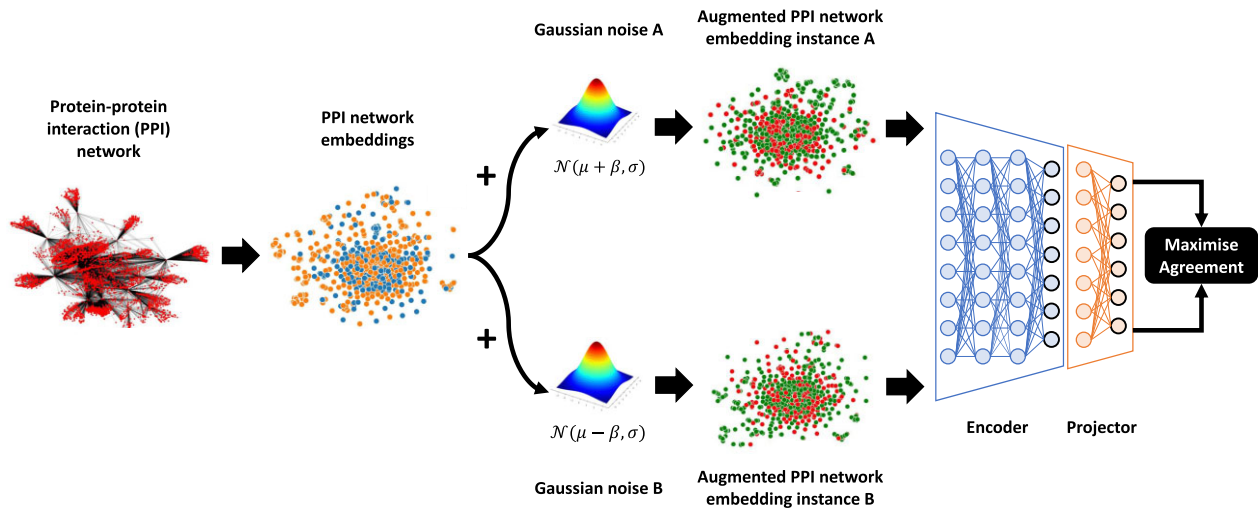


Figure 1. The flowchart for the proposed EGsCL framework based on PPI networks.

contrastive learning paradigm (29), where the definition of positive and negative views relies on the original labels of instances. For each target instance, views are considered positive if they are generated from those instances bearing the same class label as that target instance. Vice versa, the negative views are generated from the instances bearing different labels to that target instance. This supervised contrastive learning paradigm successfully demonstrated better predictive performance than the self-supervised contrastive learning paradigm.

Data augmentation plays a crucial role on contrastive learning and is usually considered domain-specific. For example, in the computer vision area, the mainstream augmentation methods (28–32) rely on spatial and colour transformations (e.g. random cropping and Gaussian blur) to create different views of original images. In the natural language processing area, text paraphrasing and word replacement (33) are usually used as augmentation methods. Several works introduced different data augmentation strategies for bioinformatics research. For example, Ciortan and DeFrance (34) and Wan *et al.* (35) used a type of random masking strategy to deal with single-cell RNA-seq expression profiles. Alsagaf *et al.* (36) and Xu *et al.* (37) adopted a noise-addition approach by randomly adding Gaussian noise vectors to gene expression profiles to create different views. In this work, we propose a new Gaussian noise-based data augmentation strategy that adopts a mean-shifting approach to enlarge the difference between views to improve the contrastive learning process.

The remainder of this paper is organized as follows. The materials and methods section introduces the newly proposed enhanced Gaussian noise augmentation-based contrastive learning (EGsCL) algorithms, followed by the results section and the discussion section, where a further analysis of the proposed algorithms was conducted. Finally, the conclusion section summarizes this paper’s major findings and mentions some future research directions.

Materials and methods

Enhanced Gaussian noise augmentation-based contrastive learning

In general, the proposed EGsCL framework learns a type of discriminative feature representations based on PPI net-

works. As shown in Figure 1, given a PPI network, EGsCL first extracts a type of PPI network embedding features using the well-known *node2vec* (38) method. Then the PPI network embedding features were used to create augmented instances (aka views) by using different Gaussian noises. For one d -dimensional PPI network embedding instance x in a given dataset, EGsCL randomly draws two Gaussian noises from two different Gaussian distributions, i.e. $\mathcal{N}(\mu + \beta, \sigma)$ and $\mathcal{N}(\mu - \beta, \sigma)$, where μ and σ denote the mean and standard deviation of the dataset, whilst β is a shifting hyperparameter that is used to manipulate the differences between those two Gaussian distributions. Those two Gaussian noises are then added with the values of x , leading to two different augmented PPI network embedding instances. After creating a pair of augmented instances for all individual PPI network embedding instances in the dataset, a new sample set that includes all those augmented instances is used as inputs for the contrastive learning networks consisting of an encoder and a projector. The contrastive learning networks optimize the parameters by adopting the conventional supervised or self-supervised contrastive learning strategies, i.e. minimizing the dissimilarity between each augmented instance and its corresponding positive augmented instance(s), whilst maximizing the dissimilarity to its corresponding negative augmented instances. To cope with the classification tasks in this work, we used the EGsCL-learned feature representations to train support vector machines to predict the pro-longevity or anti-longevity effect of different model organisms’ genes. The notations used in this paper are summarized in Table 1.

Algorithms 1 and S1 (in Supplementary File S1) show two different pseudocodes of the proposed EGsCL algorithms working with supervised and self-supervised contrastive learning loss functions, respectively. In Algorithm 1, supervised enhanced Gaussian noise augmentation-based contrastive learning (Sup-EGsCL) takes a training dataset \mathcal{X} and a corresponding class label set \mathcal{Y} as inputs and initialised five variables, i.e. a set of m -sized batches \mathcal{B} , an untrained encoder \mathcal{E} , an untrained projection head \mathcal{P} , a temperature hyperparameter τ and a mean-shift hyperparameter β . From lines 1 to 31, Sup-EGsCL processes each batch of training instances b in turns. It creates an empty variable \mathcal{L}_b to store the loss function value for b and an empty set \mathcal{S} to store the augmented in-

Table 1. The list of notations used in this paper

Notation	Description
x	A d -dimensional PPI network embedding instance.
$\mathcal{N}(\mu, \sigma)$	A Gaussian distribution, where μ and σ denote its mean and standard deviation.
β	A hyperparameter that is used to manipulate a Gaussian distribution by shifting its mean.
\mathcal{X}	A training dataset.
\mathcal{Y}	A set of class labels.
\mathcal{B}	A set of m -sized training batches.
\mathcal{E}	A contrastive learning encoder.
\mathcal{P}	A contrastive learning projection head.
τ	A temperature hyper-parameter.
b	A m -sized training batch.
\mathcal{S}	A set to store two different augmentations of each original instance.
z	A d -dimensional Gaussian noise.
\tilde{x}	An augmentation (i.e. view) of the original instance x .
\mathcal{L}_i^{SL}	The supervised contrastive loss function value for the i th instance.
\mathcal{H}_i^+	A set of projections of positive augmented instances w.r.t. \tilde{x}_i .
\mathcal{H}_i	A set of projections of all positive and negative augmented instances w.r.t. \tilde{x}_i .
$ \mathcal{H}_i^+ $	The number of positive augmented instances w.r.t. \tilde{x}_i .
$\mathcal{F}(\cdot)$	The cosine similarity.
$\mathcal{V}(\tilde{x})$	A variable that maps an augmented instance \tilde{x} to its original instance x .
\mathcal{L}_i^{SSL}	The self-supervised contrastive loss function value for the i th instance.

stances (a.k.a. views). For each training instance x_i in b (lines 4–11), two d -dimensional Gaussian noises, i.e. z_a and z_b , are randomly drawn from two different Gaussian distributions, i.e. $\mathcal{N}(\mu + \beta, \sigma)$ and $\mathcal{N}(\mu - \beta, \sigma)$, where μ and σ denote the mean and standard deviation of the training dataset \mathcal{X} . β is a hyperparameter that is used to adjust the differences between those two Gaussian distributions. Then z_a and z_b are added to x_i to create two different augmented instances, i.e. x_{ia} and x_{ib} (lines 7–8). Those two augmented instances are added to the set \mathcal{S} (lines 9–10). After obtaining the complete set \mathcal{S} that consists of all the augmented instances for the entire training dataset \mathcal{X} , from lines 12 to 28, Sup-EGsCL processes each augmented instance \tilde{x}_i in \mathcal{S} to compute the loss function value. It creates three empty variables, i.e. a variable \mathcal{L}_i^{SL} for storing the supervised loss function value for \tilde{x}_i , a set \mathcal{H}_i^+ for storing the projections of positive augmented instances with respect to \tilde{x}_i , and a set \mathcal{H}_i for storing the projections of all positive and negative augmented instances with respect to \tilde{x}_i . From lines 16 to 24, EGsCL defines the positive augmented instances according to the pre-defined class labels. Each augmented instance \tilde{x}_j in \mathcal{S} that is different from the target instance \tilde{x}_i is added to \mathcal{H}_i after getting its corresponding projection using the encoder \mathcal{E} and the projector \mathcal{P} (lines 17–19). Only the projections of those augmented instances bearing the same class label as \tilde{x}_i will be considered as positive augmented instances with respect to \tilde{x}_i and their projections will be added to \mathcal{H}_i^+ (lines 20–22). Vice versa, the negative augmented instances with respect to \tilde{x}_i are defined as those augmented instances bearing different class labels to \tilde{x}_i . After obtained the completed sets of \mathcal{H}_i^+ and \mathcal{H}_i , Sup-EGsCL creates the projection of the target instance \tilde{x}_i (line 25). Then Sup-EGsCL computes the loss function value \mathcal{L}_i^{SL} that will then be added to \mathcal{L}_b (lines 26–27). After processing all augmented instances in \mathcal{S} , the loss

function value \mathcal{L}_b will be normalised by $2m$ denoting the total number of augmented instances in \mathcal{S} , and both the encoder and the projector will be optimised (lines 29–30). The pseudocode will output a trained encoder \mathcal{E}^* after processing all batches (line 32). Equation 1 defines the supervised contrastive loss function for the target instance \tilde{x}_i , where $|\mathcal{H}_i^+|$ denotes the number of positive augmented instances w.r.t. \tilde{x}_i , j denotes the indices of the positive augmented instances, and k denotes the indices of all augmented instances except i . $\mathcal{F}(\cdot)$ denotes the cosine similarity and τ is a temperature hyper-parameter that controls the strength of penalty on positives and negatives.

Algorithm 1: Supervised enhanced Gaussian noise augmentation-based contrastive learning (Sup-EGsCL).

Input: a training dataset \mathcal{X} ;
a class labels set \mathcal{Y} for all individual instances x in \mathcal{X} ;
initialise a set of m -sized batches \mathcal{B} ;
initialise an untrained contrastive learning encoder \mathcal{E} ;
initialise an untrained contrastive learning projection head \mathcal{P} ;
initialise a temperature hyperparameter τ ;
initialise a mean-shift hyperparameter β .

Output: a trained contrastive learning encoder \mathcal{E}^* .

```

1  foreach  $b \in \mathcal{B}$  do
2  initialise an empty variable  $\mathcal{L}_b$  for the loss function value of batch  $b$ ;
3  initialise an empty set  $\mathcal{S}$  for the augmented instances of  $\mathcal{X}$ ;
4  foreach  $x_i \in b$  do
5   $z_a \sim \mathcal{N}(\mu + \beta, \sigma)$ ;
6   $z_b \sim \mathcal{N}(\mu - \beta, \sigma)$ ;
7   $x_{ia} \leftarrow x_i + z_a$ ;
8   $x_{ib} \leftarrow x_i + z_b$ ;
9   $\mathcal{S} \leftarrow \mathcal{S} \cup x_{ia}$ ;
10  $\mathcal{S} \leftarrow \mathcal{S} \cup x_{ib}$ ;
11 end
12 foreach  $\tilde{x}_i \in \mathcal{S}$  do
13 initialise an empty variable  $\mathcal{L}_i^{SL}$  for the loss function value for  $\tilde{x}_i$ ;
14 initialise an empty positive augmented instances projection set  $\mathcal{H}_i^+$  for  $\tilde{x}_i$ ;
15 initialise an empty projection set  $\mathcal{H}_i$  for all augmented instances except  $\tilde{x}_i$ ;
16 foreach  $\tilde{x}_j \in \mathcal{S}$  do
17   if  $\tilde{x}_i \neq \tilde{x}_j$  then
18      $h_j \leftarrow \mathcal{P}(\mathcal{E}(\tilde{x}_j))$ ;
19      $\mathcal{H}_i \leftarrow \mathcal{H}_i \cup h_j$ ;
20     if  $\mathcal{Y}(\tilde{x}_i) == \mathcal{Y}(\tilde{x}_j)$  then
21        $\mathcal{H}_i^+ \leftarrow \mathcal{H}_i^+ \cup h_j$ ;
22     end
23   end
24 end
25  $h_i \leftarrow \mathcal{P}(\mathcal{E}(\tilde{x}_i))$ ;
26  $\mathcal{L}_i^{SL} \leftarrow \text{LossFunction}(h_i, \mathcal{H}_i^+, \mathcal{H}_i, \tau)$ ;
27  $\mathcal{L}_b \leftarrow \mathcal{L}_b + \mathcal{L}_i^{SL}$ ;
28 end
29  $\mathcal{L}_b \leftarrow \mathcal{L}_b / 2m$ ;
30 optimize  $\mathcal{E}$  and  $\mathcal{P}$  by using  $\mathcal{L}_b$ ;
31 end
32 return  $\mathcal{E}^*$ 

```

Algorithm S1 shows the pseudocode of the self-supervised enhanced Gaussian noise augmentation-based contrastive learning (Self-EGsCL) method, which shares the same initialization and data augmentation process with the Sup-EGsCL method. The main difference between Algorithms 1 and S1 is the positive augmented instance selection strategy. As shown in lines 9 and 10, Self-EGsCL stores the original instance information for each augmented instance. For example, the value of variable $\mathcal{V}(x_{ia})$ is assigned as x_i , if x_{ia} is the augmented instance of x_i . In lines 22–24, for each augmented instance \tilde{x}_i , Self-EGsCL treated another augmented instance \tilde{x}_j as a positive augmented instance, if both \tilde{x}_i and \tilde{x}_j are generated by using the same original instance (i.e. $\mathcal{V}(\tilde{x}_i) == \mathcal{V}(\tilde{x}_j)$). All other augmented instances in \mathcal{S} are treated as negative augmented instances. Self-supervised EGsCL uses a similar loss function (Equation S1 in Supplementary File S1) as Sup-EGsCL. Because there is only one positive augmented sample w.r.t. one single target augmented instance (i.e. $|\mathcal{H}_i^+| = 1$), Self-EGsCL does not normalise the loss function value \mathcal{L}_i^{SSL} .

Table 2. Main characteristics of the created datasets

Model organisms		Mouse	Worm	Fly	Yeast
# Instances	Total	124	718	186	312
	Pro-longevity	80	239	117	34
	Anti-longevity	44	479	69	278
# Features	Embedding	128	128	128	128
	Binary	17438	16010	11535	5957
	Combined	17566	16138	11663	6085

Computational experiments

We evaluated the predictive performance of EGsCL using five different β values, i.e. 0.1, 0.2, 0.3, 0.4 and 0.5. We also compared EGsCL with the conventional Gaussian noise augmentation-based contrastive learning (GsCL) method, which also randomly draws two different Gaussian noises to create a pair of augmented instances for x , but from the same Gaussian distribution, i.e. $\mathcal{N}(\mu, \sigma)$. Therefore, GsCL is equivalent to the case when EGsCL's β value equals 0. We also compared with another GsCL variant with $\mathcal{N}(0, 1)$, which was used in (36) for cell type identification tasks. We used the well-known multi-layer perceptron (MLP) to create the encoder and the projection head of an EGsCL network. The encoder consists of three hidden layers and one output layer (i.e. the representation layer). The projection head consists of one hidden layer and one output layer. The ReLU activation function was used in both MLPs. We used Adam optimizer with a learning rate of 10^{-4} and a weight decay of 10^{-6} . The number of maximum training epochs was set to 1000. We set the value of τ to 0.1 for the supervised contrastive loss and 0.07 for the self-supervised contrastive loss. Due to the small number of instances, we set the batch size as the same as the number of training instances. The proposed EGsCL methods were implemented by PyTorch (39) and Scikit-learn (40).

We created 12 datasets in total using the ageing-related genes for four model organisms, i.e. mouse, worm, fly and yeast, as reported in the GenAge database (41). We generated three types of features based on the PPI networks deposited in the STRING database (version 12.0) (42). The first type of features is network embeddings learned by the well-known *node2vec* method (38) leading to a 128-dimensional vector for each individual protein included in the most informative combined score STRING PPI networks. The second type of features is binary PPI features, where the value of 1 denotes protein_a and protein_b have an interaction and the value of 0 means those two proteins do not have an interaction. The third type of features is the combination of both the network embedding and the binary PPI features. The characteristics of all 12 datasets are listed in Table 2. The numbers of instances for four different model organisms range between 124 and 718. The dimensionalities of binary features range between 5957 and 17 438 and the combined features range between 6085 and 17 566.

Each generated dataset was split into two subsets, i.e. 80% of the instances were used for conducting a 10-fold cross-validation, and the remaining 20% of the instances were used to create a validation set for conducting model selection during the contrastive learning process. For each fold of the cross validation, after every 5 training epochs, we froze the encoder \mathcal{E} and used it to transform the training folds, the validation set and the testing fold into the EGsCL feature representations. An Support Vector Machine (SVM) classifier was trained on

the transformed training folds and then predicted the labels of the transformed validation set. The best encoder was selected according to the highest validation set predictive accuracy. The corresponding SVM classifier was used to predict the predictive accuracy of the transformed testing fold. We measured the predictive performance using three well-known metrics, i.e. Matthews correlation coefficient (MCC), F1 score and average precision (AP) score, which were also used as model selection criteria when reporting corresponding metrics' values.

Results

EGsCL successfully improved the predictive performance of GsCL when using different types of PPI features to predict the pro-longevity or anti-longevity effect of four model organisms' genes

We first conducted pairwise comparisons between EGsCL and GsCL using supervised and self-supervised settings. In general, both Sup-EGsCL and Self-EGsCL outperformed Sup-GsCL and Self-GsCL, respectively. As shown in Table 3, when using the network embedding features to predict the longevity effects of mouse's genes, Sup-EGsCL with all different β values obtained higher MCC values and AP scores than Sup-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$, denoting by the double up arrows. The former with β values of 0.3 and 0.4 also obtained higher F1 scores than the latter. When using the binary PPI features, Sup-EGsCL with almost all β values except 0.1 obtained higher AP scores than Sup-GsCL. However, the latter obtained higher MCC values and F1 scores. When using the combined features, Sup-EGsCL with β values of 0.3 and 0.5 obtained higher MCC values and F1 scores than Sup-GsCL. The former with all β values also outperformed the latter due to higher AP scores. In terms of Self-EGsCL, when using the network embedding features and binary PPI features, it outperformed Self-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$ according to the higher MCC values, F1 and AP scores obtained with different β values, as denoted by the single up arrows. When using the combined features, Self-EGsCL with β values of 0.1 and 0.2 obtained higher MCC values than Self-GsCL. It also obtained higher AP scores with β values of 0.2 and 0.5, though Self-GsCL with $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$ obtained higher F1 scores.

When predicting the longevity effects of worm's genes using the network embedding features, Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ outperformed Sup-EGsCL, according to MCC values and F1 scores. However, Sup-EGsCL with all different β values obtained higher AP scores than Sup-GsCL with both $\mathcal{N}(\mu, \sigma)$ and $\mathcal{N}(0, 1)$. When using the binary PPI features, Sup-EGsCL obtained higher MCC values with β values of 0.2 and 0.3. It also obtained higher F1 scores with β values of 0.1 and 0.3. However, Sup-GsCL obtained higher AP scores. When using the combined features, Sup-EGsCL with different β values outperformed Sup-GsCL with both $\mathcal{N}(\mu, \sigma)$ and $\mathcal{N}(0, 1)$, according to the higher MCC values, F1 and AP scores. Analogously, as shown in Table 3, Self-EGsCL with almost all different β values using the network embedding features outperformed Self-GsCL with both $\mathcal{N}(\mu, \sigma)$ and $\mathcal{N}(0, 1)$, according to the higher MCC values and AP scores. It also obtained a higher F1 score with a β value of 0.1. When using the binary PPI features, Self-EGsCL with a β value of 0.2 obtained

Table 3. Predictive performance of Sup-EGsCL, Sup-GsCL, Self-EGsCL, Self-GsCL and the benchmark method.

Feature	Metrics	Mouse (<i>Mus musculus</i>)										Benchmark			
		Sup-EGsCL					Self-EGsCL						Self-GsCL		
		$\beta = 0.1$	$\beta = 0.2$	$\beta = 0.3$	$\beta = 0.4$	$\beta = 0.5$	$\mathcal{N}(\mu, \sigma)$	$\beta = 0.1$	$\beta = 0.2$	$\beta = 0.3$	$\beta = 0.4$		$\beta = 0.5$	$\mathcal{N}(0, 1)$	$\mathcal{N}(\mu, \sigma)$
Embeddings	MCC	0.309 ↑	0.366 ↑	0.380 ↑	0.397 ↑	<u>0.427</u> ↑	0.075	0.248	0.176	0.169	0.245 ↑	0.285 ↑	0.208	0.234	0.146
	F1	0.780	0.783	0.797 ↑	0.818 ↑	<u>0.789</u>	0.747	0.796	0.797 ↑	0.792 ↑	0.788	0.799 ↑	0.738	0.788	0.744
	AP	0.842 ↑	0.844 ↑	0.839 ↑	<u>0.847</u> ↑	0.836 ↑	0.774	0.826	0.837	0.820	0.828	0.811	0.844	0.791	0.764
	MCC	0.237	0.237	0.263	<u>0.263</u>	0.280	<u>0.373</u>	0.237	0.151	0.155	<u>0.212</u> ↑	0.175 ↑	0.142	0.157	0.325
	F1	0.794	0.800	0.800	0.800	0.800	0.821	0.821	0.816 ↑	0.809	0.811	0.811	0.815	0.756	0.811
Combined	AP	0.853	0.855 ↑	0.855 ↑	0.856 ↑	0.855 ↑	0.850	0.853	0.838 ↑	0.821	0.788	0.827 ↑	0.824	0.808	0.805
	MCC	0.237	0.278	0.329 ↑	<u>0.270</u>	0.402 ↑	0.309	0.237	0.367 ↑	0.343 ↑	0.254	0.234	0.271	0.334	0.371
	F1	0.787	0.792	0.806 ↑	0.796	<u>0.801</u> ↑	0.796	0.787	0.768	0.768	0.771	0.771	0.813	0.788	0.826
	AP	<u>0.860</u> ↑	0.837 ↑	<u>0.838</u> ↑	0.838 ↑	0.839 ↑	0.827	0.836	0.770	<u>0.826</u> ↑	0.788	0.783	0.794	0.798	0.813
							Worm (<i>Caenorhabditis elegans</i>)								
Embeddings	MCC	0.356	0.369	0.355	0.299	0.363	0.181	0.377	0.275 ↑	0.335 ↑	0.299 ↑	0.350 ↑	0.177	0.269	0.367
	F1	0.550	0.539	0.550	0.548	0.538	0.466	0.561	0.526 ↑	0.515	0.492	0.487	0.447	0.524	0.529
	AP	0.692 ↑	0.696 ↑	0.695 ↑	0.697 ↑	<u>0.698</u> ↑	0.483	<u>0.678</u>	0.593 ↑	0.597 ↑	0.590 ↑	0.593 ↑	0.500	0.587	0.685
	MCC	0.367	0.383 ↑	<u>0.387</u> ↑	0.348	0.350	0.295	0.374	0.308	<u>0.346</u> ↑	0.301	0.308	0.316	0.308	0.377
	F1	<u>0.566</u> ↑	0.534	0.559 ↑	0.553	0.551	0.551	0.555	0.496	0.503	0.494	0.499	0.520	0.517	0.530
Combined	AP	<u>0.639</u>	0.641	0.644	0.643	0.629	0.663	0.649	0.598	0.638 ↑	0.603	0.615 ↑	0.546	0.607	0.664
	MCC	0.344	0.352	0.338	0.347	<u>0.379</u> ↑	0.354	0.354	0.287 ↑	0.289 ↑	0.316 ↑	0.293 ↑	0.211	0.285	0.369
	F1	0.578	0.584 ↑	0.571	0.585 ↑	<u>0.599</u> ↑	0.544	0.579	0.496	0.493	0.549 ↑	0.516	0.492	0.516	0.529
	AP	0.640	0.649	0.649	<u>0.662</u> ↑	<u>0.651</u>	0.629	0.658	0.670 ↑	<u>0.675</u> ↑	0.670 ↑	0.665 ↑	0.562	0.586	0.647
							Fly (<i>Drosophila melanogaster</i>)								
Embeddings	MCC	0.260 ↑	0.231	<u>0.328</u> ↑	0.278 ↑	0.212	0.038	0.242	0.191	0.135	0.198 ↑	0.144	-0.052	0.194	0.134
	F1	0.747	0.747	<u>0.771</u> ↑	0.765 ↑	0.757	0.760	0.752	0.754	0.757	0.754	0.761	0.769	0.759	0.725
	AP	0.761	0.760	0.757	0.765	0.741	0.691	0.769	0.779 ↑	0.756	0.762	0.739	0.710	0.764	0.753
	MCC	0.157	0.157	0.157	0.157	0.157	0.147	<u>0.207</u>	0.014	0.038	0.021	0.113 ↑	0.015	0.101	0.270
	F1	0.756	0.756	0.756	0.756	0.756	0.736	0.774	0.733	0.745 ↑	0.725	0.738 ↑	0.727	0.733	0.769
Combined	AP	0.802	0.801	0.803	0.804	0.801	0.806	<u>0.802</u>	0.752 ↑	0.783 ↑	0.751 ↑	0.768 ↑	0.748	0.722	0.826
	MCC	0.283	0.267	0.283	0.275	0.292 ↑	0.275	0.283	0.172 ↑	0.197 ↑	0.241 ↑	0.244 ↑	0.116	0.094	0.230
	F1	0.771	0.771	<u>0.782</u>	0.776	<u>0.774</u>	0.781	0.782	0.768	0.765	0.762	0.759	0.767	0.777	0.760
	AP	0.802	0.806	<u>0.808</u>	0.811	0.804	<u>0.838</u>	<u>0.802</u>	0.744	0.708	0.689	0.689	0.749	0.732	0.821
							Yeast (<i>Saccharomyces cerevisiae</i>)								
Embeddings	MCC	0.099	0.154	0.074	0.082	0.114	0.219	0.016	0.034	0.040	0.026	0.004	0.152	0.023	0.274
	F1	0.130	0.163	0.083	0.090	0.130	<u>0.250</u>	0.130	0.153	0.107	0.090	0.073	0.220	0.127	0.297
	AP	<u>0.393</u> ↑	0.329	0.384 ↑	0.350	0.323	<u>0.277</u>	0.362	0.315	0.347	0.272	0.285	0.359	0.254	0.509
	MCC	0.040	0.103	0.103	0.095	0.103	0.165	0.082	0.010	0.019	0.024	0.073	0.066	0.173	0.034
	F1	0.050	0.100	0.100	0.100	0.100	<u>0.167</u>	0.090	0.040	0.040	0.040	0.040	0.126	0.247	0.050
Combined	AP	<u>0.469</u> ↑	0.430	0.448 ↑	0.418	0.417	<u>0.408</u>	0.435	0.393	0.391	0.357	0.380	0.393	0.374	0.397
	MCC	0.089	0.089	0.089	0.089	0.089	0.066	0.117	0.171	<u>0.180</u> ↑	0.171	0.171	0.112	0.171	0.034
	F1	0.100	0.100	0.100	0.100	0.100	0.090	<u>0.130</u>	0.180	<u>0.180</u>	0.180	0.180	0.150	0.180	0.050
	AP	0.379 ↑	0.367 ↑	0.357 ↑	0.385 ↑	0.374 ↑	0.321	<u>0.346</u>	0.306	0.310	0.325	0.301	0.263	0.385	0.402

↑: Higher values obtained by Sup-EGsCL compared with Sup-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$.
 †: Higher values obtained by Self-EGsCL compared with Self-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$.
 Double underline: the highest value between Sup-EGsCL and Sup-GsCL over all parameters.
 Underline: the highest value between Self-EGsCL and Self-GsCL over all parameters.
 Bold text: the overall highest value for the model organism.

a higher MCC value and a higher AP score than Self-GsCL, but the latter obtained a higher F1 score with $\mathcal{N}(0, 1)$. When using the combined features, Self-EGsCL with almost all different β values obtained higher MCC values and AP scores. It also obtained a higher F1 score than Self-GsCL with a β value of 0.3.

When using network embedding features to predict the longevity effects of fly's genes, Sup-EGsCL with β values of 0.3 and 0.4 obtained higher MCC values and F1 scores than Sup-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$. But Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ obtained a higher AP score. When using the binary PPI features, Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ outperformed Sup-EGsCL due to the higher MCC value and F1 score. Sup-GsCL with $\mathcal{N}(0, 1)$ also obtained a higher AP score than Sup-EGsCL. When using the combined features, Sup-EGsCL with a β value of 0.5 obtained a higher MCC value than Sup-GsCL. The former with a β value of 0.3 also obtained the same F1 score as the latter with $\mathcal{N}(\mu, \sigma)$. But Sup-GsCL with $\mathcal{N}(0, 1)$ obtained a higher AP score than Sup-EGsCL. In terms of Self-EGsCL, as shown in Table 3, according to MCC values, it outperformed Self-GsCL with a β value of 0.3 using the network embedding features. It also obtained higher AP scores with β values of 0.1 and 0.5, though Self-GsCL with $\mathcal{N}(0, 1)$ obtained a higher F1 score. When using the binary PPI features, Self-EGsCL outperformed Self-GsCL with different β values, according to the higher MCC values, F1 and AP scores. When using the combined features, Self-EGsCL with all different β values obtained higher MCC values, but Self-GsCL obtained higher F1 and AP scores with $\mathcal{N}(\mu, \sigma)$ and $\mathcal{N}(0, 1)$, respectively.

When predicting the longevity effects of yeast's genes, Sup-GsCL with $\mathcal{N}(0, 1)$ obtained higher MCC values and F1 scores than Sup-EGsCL using both the network embedding and binary PPI features. Sup-EGsCL with β values of 0.1 and 0.3 obtained higher AP scores than Sup-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$. It also obtained higher AP scores than Sup-GsCL when using the combined features with all different β values. However, Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ performed better due to the higher MCC value and F1 score. Analogously, when using the network embedding features, Self-GsCL with $\mathcal{N}(0, 1)$ outperformed Self-EGsCL, according to the higher MCC value and F1 score. However, Self-EGsCL with a β value of 0.5 obtained a higher AP score. When using the binary PPI features, Self-GsCL with $\mathcal{N}(\mu, \sigma)$ performed better than Self-EGsCL due to the higher MCC value and F1 score. Self-EGsCL with a β value of 0.1 obtained the same AP score as Self-GsCL with $\mathcal{N}(0, 1)$. When using the combined features, Self-EGsCL with a β value of 0.2 obtained a higher MCC value than Self-GsCL with both $\mathcal{N}(0, 1)$ and $\mathcal{N}(\mu, \sigma)$. Self-EGsCL with all different β values also obtained the same F1 scores as Self-GsCL with $\mathcal{N}(\mu, \sigma)$. However, the latter obtained a higher AP score than the former.

EGsCL successfully obtained state-of-the-art accuracy in predicting the pro-longevity or anti-longevity effect of three model organisms' genes using PPI network-based features

We further compared EGsCL with the benchmark method that uses raw PPI network features to train SVM classifiers. When predicting mouse genes' longevity effects using the network embedding features, both Sup-EGsCL and Self-EGsCL with all different β values obtained higher MCC values, F1

and AP scores than the benchmark method. Analogously, when working with the binary PPI features, both Sup-EGsCL and Self-EGsCL with almost all different β values obtained higher AP scores, though the benchmark obtained a higher MCC value. In addition, Self-EGsCL with a β value of 0.1 obtained a higher F1 score. When working with the combined features, Sup-EGsCL with a β value of 0.5 obtained a higher MCC value. It also obtained higher AP scores with all different β values, though the benchmark method obtained a higher F1 score. In terms of Self-EGsCL, it failed to obtain any higher MCC value and F1 score, but it obtained a higher AP score with a β value of 0.2.

When predicting worm genes' longevity effects using the network embedding features, Sup-EGsCL with all different β values obtained higher F1 and AP scores than the benchmark method, though the latter obtained a higher MCC value. When working with the binary PPI features, Sup-EGsCL with β values of 0.2 and 0.3 obtained higher MCC values. It also obtained higher F1 scores with all different β values, though the benchmark method obtained a higher AP score. When using the combined features, Sup-EGsCL with a β value of 0.5 obtained a higher MCC value. It also obtained higher F1 and AP scores with almost all different β values than the benchmark method. In terms of Self-EGsCL, it failed to obtain any higher MCC value, F1 and AP scores than the benchmark method using both the network embedding features and the binary PPI features. However, when working with the combined features, it obtained a higher F1 score with a β value of 0.3. It also obtained higher AP scores than the benchmark method with all different β values.

When predicting fly genes' longevity effects, both Sup-EGsCL and Self-EGsCL with almost all different β values obtained higher MCC values, F1 and AP scores than the benchmark method using the network embedding features. However, when using the binary PPI features, the latter obtained higher MCC value, F1 and AP scores. When working with the combined features, Sup-EGsCL with all different β values obtained higher MCC values and F1 scores, though the benchmark method obtained a higher AP score. In terms of Self-EGsCL, it obtained higher MCC values with β values of 0.3 and 0.4. It also obtained higher F1 scores with almost all β values, though the benchmark method obtained a higher AP score.

When predicting yeast genes' longevity effects using the network embedding features, the benchmark method outperformed both Sup-EGsCL and Self-EGsCL due to its higher MCC value, F1 and AP scores. However, when using the binary PPI features, Sup-EGsCL with almost all different β values obtained higher MCC values, F1 and AP scores, but Self-EGsCL failed to obtain higher F1 and AP scores than the benchmark method. When working with the combined features, Sup-EGsCL outperforms the benchmark method with all different β values due to the higher MCC values and F1 scores, though the latter obtained a higher AP score. Analogously, Self-EGsCL also obtained higher MCC values and F1 scores than the benchmark method with all different β values.

Sup-EGsCL is also the overall best method for predicting mouse, worm and fly genes' longevity effects. As denoted by the bold texts in Table 3, in terms of the mouse datasets, Sup-EGsCL with a β value of 0.5 obtained the overall highest MCC value (i.e. 0.427), whilst it also obtained the overall highest AP score (i.e. 0.860) with a β value of 0.1. The overall highest F1 score (i.e. 0.826) was obtained by the bench-

Table 4. New predictions about the pro-/anti-longevity effect of mouse genes and their homologous genes from human, fly and worm

Mouse Gene ID	Mouse Gene Name	Predicted Class	Predicted Probability	Homologous genes from Human (HS), Fly (DM) and Worm (CE)
Pofut1	Protein O-fucosyltransferase 1	Pro-longevity	87.8%	POFUT1 (HS), O-fut1 (DM), pfut-1 (CE)
Ints15	Integrator complex subunit 15	Pro-longevity	87.7%	INTS15 (HS), CG5274 (DM), Y56A3A.31 (CE)
Plod2	Procollagen lysine, 2-oxoglutarate 5-dioxygenase 2	Pro-longevity	87.7%	PLOD2 (HS), Plod (DM), let-268 (CE)
Arid3a	AT-rich interaction domain 3A	Pro-longevity	87.6%	ARID3A (HS), retn (DM), cfi-1 (CE)
Col3a1	Collagen, type III, alpha 1	Pro-longevity	87.3%	COL3A1 (HS)
Grk5	G protein-coupled receptor kinase 5	Anti-longevity	71.3%	GRK5 (HS), Gprk2 (DM), grk-1 (CE)
C2cd4b	C2 calcium-dependent domain containing 4B	Anti-longevity	70.5%	C2CD4B (HS)
Sstr3	Somatostatin receptor 3	Anti-longevity	69.6%	SSTR3 (HS), AstC-R1 (DM), npr-24 & npr-16 (CE)
Rab44	RAB44, member RAS oncogene family	Anti-longevity	69.5%	RAB44 (HS), rsef-1 (CE)
Ntsr1	Neurotensin receptor 1	Anti-longevity	69.5%	NTSR1 (HS)
‡ Apln	Apelin	Anti-longevity	70.2%	APLN (HS)

‡ This gene is predicted as an anti-longevity gene, but the literature suggests it is a pro-longevity gene.

mark method. Analogously, in terms of the worm datasets, Sup-EGsCL also obtained the overall highest MCC value (i.e. 0.387), F1 score (i.e. 0.599) and AP score (i.e. 0.698) with different β values. The overall highest MCC value (i.e. 0.328) and F1 score (i.e. 0.782) for the fly datasets were obtained by Sup-EGsCL with a β value of 0.3. Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ also obtained the same overall highest F1 score, whilst Sup-GsCL with $\mathcal{N}(0, 1)$ obtained the overall highest AP score (i.e. 0.838). In terms of the yeast datasets, the overall highest MCC value (i.e. 0.274), F1 score (i.e. 0.297) and AP score (i.e. 0.509) were all obtained by the benchmark method.

Sup-EGsCL successfully predicted novel mouse genes with the pro-/anti-longevity effect

We then used one of the trained Sup-EGsCL-based classifiers during the 10-fold cross-validation to predict the pro-/anti-longevity effect of all the mouse genes included in the STRING database. The pro-longevity genes are defined as those genes whose decreased expression reduces lifespan and/or their overexpression extends lifespan. Vice versa, the anti-longevity genes are defined as those genes whose overexpression reduces lifespan and/or their decreased expression extends lifespan (14).

We focus on predicting novel mouse genes for several reasons, as follows. First, the predictive models for mouse data are the most accurate models in general, across the models for the four organisms. Second, mice are much closer to humans than the other three model organisms investigated (with results for mice being more useful as evidenced from pre-clinical studies). Third, experiments with mice are much slower and more time consuming than experiments with the other three types of organisms investigated, so it is particularly important to use machine learning methods to prioritize mouse genes for further testing via wet-lab experiments.

Table 4 shows the top-ranked mouse genes that were most likely to bear pro-/anti-longevity labels according to their probabilities predicted by the trained Sup-EGsCL-based classifier. Those genes are considered potentially novel pro-/anti-longevity genes because they are not included in the GenAge database (and so, they are not in the datasets used to learn our Sup-EGsCL-based classifiers). The table also includes information about homologous genes from human, fly and worm

according to the Alliance of Genome Resources database (43) with the stringent homolog information deposit criterion. The complete list of mouse genes that are included in both STRING (42) and NCBI (44) databases with their predicted probabilities of bearing the pro-/anti-longevity effect is included in [Supplementary File S2](#). Other genes might also be considered potentially exhibiting a pro-/anti-longevity effect if their predicted probabilities are no less than a certain threshold, which can be specified by each researcher based on their research requirements.

For example, in order to identify the small sets of top-ranked genes reported in Table 4, we consider that a mouse gene is likely to have a pro-longevity effect if its corresponding predicted probability is no <85%; whilst a mouse gene is likely to have an anti-longevity effect if its corresponding probability is no <67%. We consider a somewhat smaller probability threshold for identifying potentially novel anti-longevity genes due to the fact that, overall, the degree of confidence (predicted probabilities) for the predicted anti-longevity genes is substantially smaller than the degree of confidence for the predicted pro-longevity genes.

Regarding the predicted pro-longevity genes in Table 4, there is support in the literature for their pro-longevity role, as follows. As the top-ranked pro-longevity gene, Pofut1 and its homologous genes from human, fly and worm play important roles in the well-known ageing-related notch pathway (45). It has been found in mice that this gene's deletion is linked to multiple muscle ageing-related phenotypes (46) and promotes colorectal cancer cell apoptosis (47). Ints15 is another top-ranked mouse gene predicted to have a pro-longevity effect. It is known to be related to RNA polymerase II—another well-known ageing-related factor in multiple species (48). Recent research on mice's Ints15 gene (49) also confirmed its crucial role in cell survival—the knockout of Ints15 induces cell apoptosis. Analogously, Plod2 and its corresponding homologous human genes play an important role in responses to hypoxia (50), which could extend the lifespan of mice (51). Arid3a and its homologous genes from human, fly and worm are another group of genes that are linked to RNA polymerase II-related transcription regulations. It has been revealed that the loss of Arid3a gene leads to defects in hematopoiesis (52)—a common pattern observed in aged individuals (53). Col3a1 and its human homolog are linked with type III collagen, which




Figure 2. 2D *t*-SNE visualizations of the training and testing datasets for fly genes using the network embedding features (**A** and **D**) and the feature representations learned by Sup-GsCL (**B** and **E**) and Sup-EGsCL; $\beta = 0.3$ (**C** and **F**), respectively.

plays a crucial role in normal collagen I fibrillogenesis in the cardiovascular system, and the deletion of *Col3a1* shortens the lifespan of mouse (54).

Among those predicted mouse genes that have an anti-longevity effect as shown in Table 4, *Grk5* regulates responses to inflammatory factors (55)—a key factor leading to senescence (56). Recent research in human and mouse has revealed that silencing the *Grk5* gene could suppress inflammatory factors (55). *C2cd4b* is linked with reactive oxygen species, which is a well-known ageing-related factor (57). The over-expression of *C2cd4b* leads to an increased risk of type 2 diabetes (58,59), but inhibition of *C2CD4B* expression prevents hyperglycemia-induced oxidative stress (60). *Sstr3* and its homologs are linked with the G protein-coupled receptor (GPCR) signalling pathway. It has been found that GPCRs play important roles in T-cell-related ageing processes (61), and the blockade of *SSTR3* in human cells can reduce T-cell responses (62). *Rab44* is also closely associated with immunosenescence. The knockout of *Rab44* in mice diminishes anaphylaxis (63), which is a process involving a large number of mast cells releasing a wide range of inflammatory mediators (64). *Ntsr1* has also been found to regulate apoptotic processes—the inhibition of *NTSR1* in human breast cancer cell lines leads to reduced ERK 1/2 phosphorylation (65), which induces apoptotic processes (66). However, among the top-ranked genes that are predicted to have an anti-longevity effect, *Apln* was actually found to be associated with the longevity effect, since accelerated senescence was observed in *Apln* knockout mice (67). This shows that of course even highly accurate models like our Sup-EGsCL-based classifiers can occasionally make wrong predictions; and so experiments measuring mouse lifespan need to be done, in future work, to determine whether the novel pro-/anti-longevity genes predicted in this work really have their predicted effect.

Discussion

Sup-EGsCL successfully learns discriminative feature representations based on network embedding features leading to better decision boundaries

We compared the raw network embedding features and two types of feature representations learned by Sup-EGsCL and Sup-GsCL, respectively. Figure 2 shows the 2D *t*-SNE visualization of the training and testing datasets for fly genes including the learned SVM decision boundaries. As shown in Figure 2A and D, when using the raw network embedding features, both the training and testing instances bearing different class labels are distributed in overlapping areas. The learned decision boundary also failed to distinguish the red and green dots denoting two different class labels. As shown in Figure 2B and E, Sup-GsCL failed to learn discriminative feature representations since the instances bearing different class labels were still distributed in the overlapping areas. Analogously, the learning SVM decision boundaries also failed to separate the majority of the red and green dots. In contrast, Sup-EGsCL with a β value of 0.3 shows better sample distributions. As shown in Figure 2C and F, both the training and testing instances are grouped into two separate areas, whilst the learned SVM decision boundaries successfully distinguished more red and green dots.

Augmentation with noises sampled from two different Gaussian distributions leads to higher predictive accuracy.

We further discussed the differences in augmentation approaches between EGsCL and GsCL. The former samples noises from two different Gaussian distributions, i.e. $\mathcal{N}(\mu + \beta, \sigma)$ and $\mathcal{N}(\mu - \beta, \sigma)$, whilst the latter samples two noises

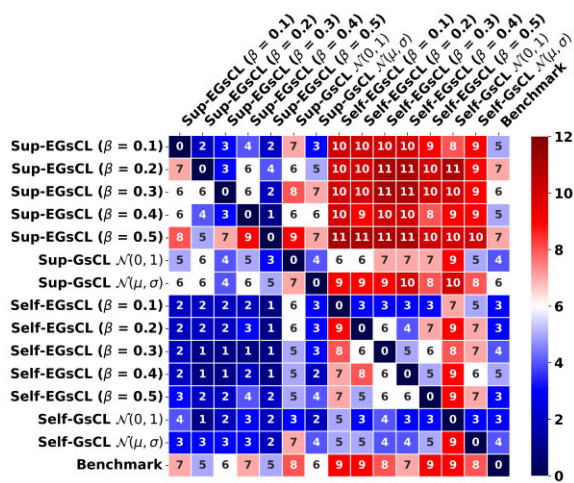


Figure 3. A heatmap showing the numbers of datasets where the methods on the rows obtained higher MCC values than the methods on the columns.

from one single Gaussian distribution, e.g. $\mathcal{N}(\mu, \sigma)$. In general, noises sampled from two different Gaussian distributions lead to higher predictive accuracy, compared with using noises sampled from one single Gaussian distribution. Figure 3 shows a heatmap for the pairwise comparisons between different methods according to their MCC values obtained by 12 datasets, i.e. 4 model organisms’ ageing-related genes described by 3 different feature types. Sup-EGsCL with both β values of 0.3 and 0.5 obtained higher MCC values in more datasets (i.e. 7 out of 12) than Sup-GsCL with $\mathcal{N}(\mu, \sigma)$, whilst Self-EGsCL with almost all different β values except 0.1 also obtained higher MCC values than Self-GsCL with $\mathcal{N}(\mu, \sigma)$ in more datasets. Sup-EGsCL with a β value of 0.4 obtained higher MCC values in the same number of datasets as Sup-GsCL with $\mathcal{N}(\mu, \sigma)$, which obtained higher MCC values in more datasets than Sup-EGsCL with β values of 0.1 and 0.2.

Supervised contrastive learning paradigm leads to higher predictive accuracy than self-supervised contrastive learning paradigm

In terms of the differences between supervised and self-supervised paradigms, the former leads to higher predictive accuracy for both EGsCL and GsCL methods. As shown in the top right area of Figure 3, Sup-EGsCL with all different β values obtained higher MCC values than Self-EGsCL with all different β values in the vast majority of the datasets. Analogously, Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ obtained higher MCC values than Self-GsCL with $\mathcal{N}(\mu, \sigma)$ in 8 out of 12 datasets, whilst Sup-GsCL with $\mathcal{N}(0, 1)$ also outperformed Self-GsCL with $\mathcal{N}(0, 1)$ in 9 out of 12 datasets.

In terms of the differences between two Gaussian distribution settings, i.e. $\mathcal{N}(\mu, \sigma)$ and $\mathcal{N}(0, 1)$, the former outperformed the latter using either supervised or self-supervised settings. As shown in Figure 3, Sup-GsCL with $\mathcal{N}(\mu, \sigma)$ obtained higher MCC values than Sup-GsCL with $\mathcal{N}(0, 1)$ in 7 out of 12 datasets, whilst Self-GsCL with $\mathcal{N}(\mu, \sigma)$ also outperformed Self-GsCL with $\mathcal{N}(0, 1)$ in 9 out of 12 datasets.

Conclusion

In summary, we proposed two new contrastive learning methods, i.e. Sup-EGsCL and Self-EGsCL, which successfully learn a type of discriminative representations based on protein-protein interaction network data, leading to state-of-the-art accuracy in predicting pro-longevity or anti-longevity effect of model organisms’ genes. In addition, we have used Sup-EGsCL to predict 10 novel pro-/anti-longevity mouse genes, and have discussed the support for these predictions in the literature. An interesting future research direction would be to propose new contrastive learning methods for other features like Gene Ontology terms or their corresponding hierarchy embeddings.

Data availability

The datasets used in this work and the pretrained encoders can be downloaded from <https://doi.org/10.5281/zenodo.12143797>. Source code is available at <https://doi.org/10.6084/m9.figshare.26227532> and at <https://github.com/ibrahimsaggaf/EGsCL>.

Supplementary data

Supplementary Data are available at NARGAB Online.

Acknowledgements

The authors acknowledge the support of the School of Computing and Mathematical Sciences and the Birkbeck GTA programme.

Funding

No external funding.

Conflicts of interest statement

None declared.

References

- Schmauck-Medina, T., Molière, A., Lautrup, S., Zhang, J., Chlopicki, S., Madsen, H.B., Cao, S., Soendenbroe, C., Mansell, E., Vestergaard, M.B., et al. (2022) New hallmarks of ageing: a 2022 Copenhagen ageing meeting summary. *Ageing*, 14, 6829–6839.
- de Magalhães, J.P. (2024) Distinguishing between driver and passenger mechanisms of ageing. *Nat. Genet.*, 56, 204–211.
- Gems, D. and de Magalhães, J.P. (2023) The hoverfly and the wasp: a critique of the hallmarks of ageing as a paradigm. *Ageing Res. Rev.*, 70, 101407.
- Rattan, S. (2024) Seven knowledge gaps in modern biogerontology. *Biogerontology*, 25, 1–8.
- Fuentealba, M., Linda, P., Thornton, J.M., Handan, M.D. and Fabian, D.K. (2021) Common genetic associations between age-related diseases. *Nat. Aging*, 1, 400–412.
- Li, Z., Zhang, Z., Ren, Y., Wang, Y., Fang, J., Yue, H., Ma, S. and Guan, F. (2021) Ageing and age-related diseases: from mechanisms to therapeutic strategies. *Biogerontology*, 22, 165–187.
- Magdaleno, G.D.V. and de Magalhães, J.P. (2021) Pleiotropy and disease interactors: the dual nature of genes linking ageing and ageing-related diseases. bioRxiv doi:

- <https://doi.org/10.1101/2024.05.20.594975>, 22 May 2024, preprint: not peer reviewed.
8. Parkhitko, A.A., Filine, E., Mohr, S.E., Moskalev, A. and Perrimon, N. (2020) Targeting metabolic pathways for extension of lifespan and healthspan across multiple species. *Aging Res. Rev.*, **64**, 101188.
 9. Parkhitko, A.A., Filine, E. and Tatar, M. (2023) Combinatorial interventions in aging. *Nat. Aging*, **3**, 1187–1200.
 10. Fabris, F., de Magalhães, J.P. and Freitas, A.A. (2017) A review of supervised machine learning applied to ageing research. *Biogerontology*, **18**, 171–188.
 11. Zhavoronkov, A., Mamoshina, P., Vanhaelen, Q., Scheibye-Knudsen, M., Moskalev, A. and Aliper, A. (2019) Artificial intelligence for aging and longevity research: recent advances and perspectives. *Aging Res. Rev.*, **49**, 49–66.
 12. Wan, C., Freitas, A.A. and de Magalhães, J.P. (2015) Predicting the pro-longevity or anti-longevity effect of model organism genes with new hierarchical feature selection methods. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, **12**, 262–275.
 13. Wan, C. (2019) Hierarchical feature selection for knowledge discovery: application of data mining to the biology of ageing. Springer Cham, Switzerland.
 14. de Magalhães, J.P., Abidi, Z., Santos, G.A.D., Avelar, R.A., Barardo, D., Chatsirisupachai, K., Clark, P., De-Souza, E.A., Johnson, E.J., Lopes, L., et al. (2024) Human Ageing Genomic Resources: updates on key databases in ageing research. *Nucleic Acids Res.*, **52**, D900–D908.
 15. Vazquez, A., Flammini, A., Maritan, A. and Vespignani, A. (2003) Global protein function prediction from protein-protein interaction networks. *Nat. Biotechnol.*, **21**, 697–700.
 16. Xiong, W., Liu, H., Guan, J. and Zhou, S. (2013) Protein function prediction by collective classification with explicit and implicit edges in protein-protein interaction networks. *BMC Bioinform.*, **14**(Suppl. 12), S4.
 17. Wan, C., Cozzetto, D., Fa, R. and Jones, D.T. (2019) Using Deep Maxout Neural Networks to improve the accuracy of function prediction from protein interaction networks. *PLoS One*, **14**, e0209958.
 18. Ortutay, C. and Vihinen, M. (2009) Identification of candidate disease genes by integrating Gene Ontologies and protein-interaction networks: case study of primary immunodeficiencies. *Nucleic Acids Res.*, **37**, 622–628.
 19. Navlakha, S. and Kingsford, C. (2010) The power of protein interaction networks for associating genes with diseases. *Bioinformatics*, **26**, 1057–1063.
 20. Guney, E. and Oliva, B. (2012) Exploiting protein-protein interaction networks for genome-wide disease-gene prioritization. *PLoS One*, **7**, e43557.
 21. Freitas, A.A., Vasieva, O. and de Magalhães, J.P. (2011) A data mining approach for classifying DNA repair genes into ageing-related or non-ageing-related. *BMC Genomics*, **12**, 27.
 22. Fang, Y., Wang, X., Michaelis, E.K. and Fang, J. (2013) Classifying aging genes into DNA repair or non-DNA repair-related categories. In: *Intelligent Computing Theories and Technology, Lecture Notes in Computer Science*. Nanning, China, pp. 20–29.
 23. Song, X., Zhou, Y.-C., Feng, K., Li, Y.-H. and Li, J.-H. (2012) Discovering aging-genes by topological features in *Drosophila melanogaster* protein-protein interaction network. In: *2012 IEEE 12th International Conference on Data Mining Workshops*. Brussels, Belgium, pp. 94–98.
 24. Feng, K., Song, X., Tan, F., Li, Y.-H., Zhou, Y.-C. and hui Li, J. (2012) Topological analysis and prediction of aging genes in *Mus musculus*. In: *2012 International Conference on Systems and Informatics (ICSAI)*. Yantai, China, pp. 2268–2271.
 25. Li, Y.-H., Zhang, G.-G. and Guo, Z. (2010) Computational prediction of aging genes in human. In: *2010 International Conference on Biomedical Engineering and Computer Science*. Wuhan, China, pp. 1–4.
 26. Magdaleno, G.D.V., Bespalov, V., Zheng, Y., Freitas, A.A. and Magalhaes, J.P.D. (2022) Machine learning-based predictions of dietary restriction associations across ageing-related genes. *BMC Bioinform.*, **23**, 10.
 27. Ribeiro, C., Farmer, C.K., de Magalhães, J.P. and Freitas, A.A. (2023) Predicting lifespan-extending chemical compounds for *C. elegans* with machine learning and biologically interpretable features. *Ageing*, **15**, 6073–6099.
 28. Chen, T., Kornblith, S., Norouzi, M. and Hinton, G. (2020) A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. PMLR, pp. 1597–1607.
 29. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C. and Krishnan, D. (2020) Supervised contrastive learning. In: *Advances in Neural Information Processing Systems*. Vancouver, British Columbia, Canada. Vol. **33**, pp. 18661–18673.
 30. Chen, T., Kornblith, S., Swersky, K., Norouzi, M. and Hinton, G. (2020) Big self-supervised models are strong semi-supervised learners. In: *Advances in Neural Information Processing Systems*. Vancouver, British Columbia, Canada. Vol. **33**, pp. 22243–22255.
 31. Kang, B., Li, Y., Xie, S., Yuan, Z. and Feng, J. (2021) Exploring balanced feature spaces for representation learning. In: *International Conference on Learning Representations*. Virtual.
 32. Robinson, J., Chuang, C.-Y., Sra, S. and Jegelka, S. (2020) Contrastive learning with hard negative samples. arXiv doi: <https://arxiv.org/abs/2010.04592>, 09 October 2020, preprint: not peer reviewed.
 33. Qu, Y., Shen, D., Shen, Y., Sajeev, S., Han, J. and Chen, W. (2020) Coda: Contrast-enhanced and diversity-promoting data augmentation for natural language understanding. arXiv doi: <https://arxiv.org/abs/2010.08670>, 16 October 2020, preprint: not peer reviewed.
 34. Ciortan, M. and Defrance, M. (2021) Contrastive self-supervised clustering of scRNA-seq data. *BMC Bioinform.*, **22**, 280.
 35. Wan, H., Chen, L. and Deng, M. (2022) scNAME: Neighborhood contrastive clustering with ancillary mask estimation for scRNA-seq data. *Bioinformatics*, **38**, 1575–1583.
 36. Alsaggaf, J., Buchan, D. and Wan, C. (2024) Improving cell type identification with Gaussian noise-augmented single-cell RNA-seq contrastive learning. *Brief. Funct. Genomics*, **23**, elad059.
 37. Xu, Y., Das, P. and McCord, R.P. (2022) SMILE: Mutual information learning for integration of single-cell omics data. *Bioinformatics*, **38**, 476–486.
 38. Grover, A. and Leskovec, J. (2016) node2vec: scalable feature learning for networks. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. San Francisco, California, USA.
 39. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019) PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: *Advances in Neural Information Processing Systems*. Curran Associates, Inc., Vancouver, British Columbia, Canada, pp. 8024–8035.
 40. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011) Scikit-learn: Machine learning in Python. *J. Machine Learn. Res.*, **12**, 2825–2830.
 41. Tacutu, R., Thornton, D., Johnson, E., Budovsky, A., Barardo, D., Craig, T., Diana, E., Lehmann, G., Toren, D., Wang, J., et al. (2018) Human Ageing Genomic Resources: new and updated databases. *Nucleic Acids Res.*, **46**, D1083–D1090.
 42. Szklarczyk, D., Kirsch, R., Koutrouli, M., Nastou, K., Mehryary, F., Hachilif, R., Gable, A.L., Fang, T., Doncheva, N.T., Pyysalo, S., et al. (2023) The STRING database in 2023: protein-protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic Acids Res.*, **51**, D638–D646.
 43. Alliance of Genome Resources Consortium (2024) Updates to the Alliance of Genome Resources central infrastructure. *Genetics*, **227**, iyae049.

44. NCBI Resource Coordinators (2016) Database resources of the National Center for Biotechnology Information. *Nucleic Acids Res.*, **44**, D7–D19.
45. Balistreri, C.R., Madonna, R., Melino, G. and Caruso, C. (2016) The emerging role of Notch pathway in ageing: Focus on the related mechanisms in age-related diseases. *Ageing Res. Rev.*, **29**, 50–65.
46. Zygmont, D.A., Singhal, N., Kim, M.-L., Cramer, M.L., Crowe, K.E., Xu, R., Jia, Y., Adair, J., Valenzuel, J. M.-P., Akaaboune, M., et al. (2017) Deletion of Pofut1 in mouse skeletal myofibers induces muscle aging-related phenotypes in cis and in trans. *Mol. Cell Biol.*, **37**, e00426-16.
47. Du, Y., Li, D., Li, N., Su, C., Yang, C., Lin, C., Chen, M., Wu, R., Li, X. and Hu, G. (2018) POFUT1 promotes colorectal cancer development through the activation of Notch1 signaling. *Cell Death Dis.*, **9**, 995.
48. Debès, C., Papadakis, A., Grönke, S., Özlem, Karalay, Tain, L.S., Mizi, A., Nakamura, S., Hahn, O., Weigelt, C., Josipovic, N., et al. (2023) Ageing-associated changes in transcriptional elongation influence longevity. *Nature*, **616**, 814–821.
49. Azuma, N., Yokoi, T., Tanaka, T., Matsuzaka, E., Saida, Y., Nishina, S., Terao, M., Takada, S., Fukami, M., Okamura, K., et al. (2023) Integrator complex subunit 15 controls mRNA splicing and is critical for eye development. *Human Mol. Genet.*, **32**, 2032–2045.
50. Rosell-García, T., Palomo-Álvarez, O. and Rodríguez-Pascual, F. (2019) A hierarchical network of hypoxia-inducible factor and SMAD proteins governs procollagen lysyl hydroxylase 2 induction by hypoxia and transforming growth factor β 1. HIF and SMAD signaling pathways induce PLOD2 expression. *J. Biol. Chem.*, **294**, 14308–14318.
51. Rogers, R.S., Wang, H., Durham, T.J., Stefely, J.A., Owiti, N.A., Markhard, A.L., Sandler, L., To, T.-L. and Mootha, V.K. (2023) Hypoxia extends lifespan and neurological function in a mouse model of aging. *PLoS Biol.*, **21**, e3002117.
52. Ratliff, M.L., Templeton, T.D., Ward, J.M. and Webb, C.F. (2014) The Bright side of hematopoiesis: regulatory roles of ARID3a/Bright in human and mouse hematopoiesis. *Front. Immunol.*, **5**, 113.
53. Ratliff, M.L., Garton, J., James, J.A. and Webb, C.F. (2020) ARID3a expression in human hematopoietic stem cells is associated with distinct gene patterns in aged individuals. *Immun. Ageing*, **17**, 24.
54. Liu, X., Wu, H., Byrne, M., Krane, S. and Jaenisch, R. (1997) Type III collagen is crucial for collagen I fibrillogenesis and for normal cardiovascular development. *Proc. Natl Acad. Sci. U.S.A.*, **94**, 1852–1856.
55. Toya, M., Akasaki, Y., Sueishi, T., Kurakazu, I., Kuwahara, M., Uchida, T., Tsutsui, T., Tsushima, H., Yamada, H., Lotz, M.K., et al. (2021) G protein-coupled receptor kinase 5 deletion suppresses synovial inflammation in a murine model of collagen antibody-induced arthritis. *Sci. Rep.*, **11**, 10481.
56. Li, X., Li, C., Zhang, W., Wang, Y., Qian, P. and Huang, H. (2023) Inflammation and aging: signaling pathways and intervention therapies. *Signal Transd. Target. Ther.*, **8**, 239.
57. Volpe, C.M.O., Villar-Delfino, P.H., dos Anjos, P.M.F. and Nogueira-Machado, J.A. (2018) Cellular death, reactive oxygen species (ROS) and diabetic complications. *Cell Death Dis.*, **9**, 119.
58. Zhang, M.J., Pisco, A. O., Darmanis, S. and Zou, J. (2021) Mouse aging cell atlas analysis reveals global and cell type-specific aging signatures. *Elife*, **13**, e62293.
59. Kycia, I., Wolford, B.N., Huyghe, J.R., Fuchsberger, C., Vadlamudi, S., Kursawe, R., Welch, R.P., d'Oliveira Albanus, R., Uyar, A., Khetan, S., et al. (2018) A common type 2 diabetes risk variant potentiates activity of an evolutionarily conserved islet stretch enhancer and increases C2CD4A and C2CD4B expression. *Am. J. Human Genet.*, **102**, 620–635.
60. Pietro, P.D., Abate, A.C., Prete, V., Damato, A., Venturini, E., Rusciano, M.R., Izzo, C., Visco, V., Ciccarelli, M., Vecchione, C., et al. (2024) C2CD4B evokes oxidative stress and vascular dysfunction via a PI3K/Akt/PKC α -signaling pathway. *Antioxidants*, **13**, 101.
61. Mittelbrunn, M. and Kroemer, G. (2021) Hallmarks of T cell aging. *Nat. Immunol.*, **22**, 687–698.
62. Zhang, B., Feng, H., Lin, H. and Li, R. (2024) Somatostatin-SSTR3-GSK3 modulates human T-cell responses by inhibiting OXPHOS. *Front. Immunol.*, **15**, 1322670.
63. Kadowaki, T., Yamaguchi, Y., Kido, M.A., Abe, T., Ogawa, K., Tokuhisa, M., Gao, W., Okamoto, K., Kiyonari, H. and Tsukuba, T. (2020) The large GTPase Rab44 regulates granule exocytosis in mast cells and IgE-mediated anaphylaxis. *Cell. Mol. Immunol.*, **17**, 1287–1289.
64. Pedicini, L., Smith, J., Savic, S. and McKeown, L. (2024) Rab46: a novel player in mast cell function. *Discov. Immunol.*, **3**, kyad028.
65. Heikal, Y., Woll, M.P., Fox, T., Seaton, K., Levenson, R. and Kester, M. (2011) Neurotensin receptor-1 inducible palmitoylation is required for efficient receptor-mediated mitogenic-signaling within structured membrane microdomains. *Cancer Biol. Ther.*, **12**, 427–435.
66. Kong, T., Liu, M., Ji, B., Bai, B., Cheng, B. and Wang, C. (2019) Role of the extracellular signal-regulated kinase 1/2 signaling pathway in ischemia-reperfusion injury. *Front. Physiol.*, **10**, 1038.
67. Zhou, Q., Chen, L., Tang, M., Guo, Y. and Li, L. (2018) Apelin/APJ system: a novel promising target for anti-aging intervention. *Clin. Chim. Acta*, **487**, 233–240.