



REVIEW

The Human Ageing Genomic Resources: online databases and tools for biogerontologists

João Pedro de Magalhães,¹ Arie Budovsky,²
Gilad Lehmann,² Joana Costa,³ Yang Li,⁴
Vadim Fraifeld² and George M. Church¹

¹Department of Genetics, Harvard Medical School, Boston, MA, USA, ²The Shraga Segal Department of Microbiology and Immunology, Center for Multidisciplinary Research on Aging, Ben-Gurion University of the Negev, Beer-Sheva, Israel, ³Department of Pathology, Brigham and Women's Hospital, Boston, MA, USA, ⁴Department of Mathematics, McGill University, Montreal, Quebec, Canada

Summary

Aging is a complex, challenging phenomenon that requires multiple, interdisciplinary approaches to unravel its puzzles. To assist basic research on aging, we developed the Human Ageing Genomic Resources (HAGR). This work provides an overview of the databases and tools in HAGR and describes how the gerontology research community can employ them. Several recent changes and improvements to HAGR are also presented. The two centrepieces in HAGR are GenAge and AnAge. GenAge is a gene database featuring genes associated with aging and longevity in model organisms, a curated database of genes potentially associated with human aging, and a list of genes tested for their association with human longevity. A myriad of biological data and information is included for hundreds of genes, making GenAge a reference for research that reflects our current understanding of the genetic basis of aging. GenAge can also serve as a platform for the systems biology of aging, and tools for the visualization of protein–protein interactions are also included. AnAge is a database of aging in animals, featuring over 4000 species, primarily assembled as a resource for comparative and evolutionary studies of aging. Longevity records, developmental and reproductive traits, taxonomic information, basic metabolic characteristics, and key observations related to aging are included in AnAge. Software is also available to aid researchers in the form of Perl modules to automate numerous tasks and as an SPSS script to

analyse demographic mortality data. The HAGR are available online at <http://genomics.senescence.info>.

Key words: bioinformatics; comparative biology; functional genomics; genetics of longevity; systems biology.

Introduction

Despite recent progress, human aging is a largely controversial process. Many age-related changes have been described, yet there are multiple and conflicting theories regarding what mechanism(s) drive such changes (de Magalhães, 2005). Moreover, we do not know why different species age at different paces, and there is still no proven intervention capable of delaying or postponing the human aging process (Olshansky *et al.*, 2002). As such, it is clear that aging is a complex, challenging phenomenon that requires extensive research using multiple, interdisciplinary approaches to unravel its puzzles.

To help researchers address the key questions in gerontology, we developed a collection of online resources entitled the Human Ageing Genomic Resources (HAGR). Previously, we described the databases and tools present in HAGR and how these resources were created and implemented (de Magalhães *et al.*, 2005a). Herein, our aim is to present a non-technical description of these resources to the gerontology community and provide an overview of how biogerontologists can use them. Numerous recent updates and improvements are also described. We hope this work will help readers better realize what is on offer at HAGR and how its contents can be useful for their own research. The HAGR are freely available online (no registration necessary) at <http://genomics.senescence.info> (Fig. 1).

GenAge: the aging gene database

Philosophy and overview of resources

It is undisputed that genetic factors influence aging. In a remarkable series of recent breakthroughs, a number of genes capable of altering the aging process as a whole – or at least to a large degree – have been identified in animal models and even a few in humans (Finch & Ruvkun, 2001; de Magalhães, 2005; Kenyon, 2005). Furthermore, multiple alleles have been examined for their association with human exceptional longevity (Vijg & Suh, 2005). This is a fascinating and important area of research, yet there are now so many genes being associated with aging and longevity that keeping track of them all is becoming increasingly more difficult. Moreover, it is necessary now to study not only individual genes but their interactions with each other and with

Correspondence

João Pedro de Magalhães, School of Biological Sciences, University of Liverpool, Biosciences Building, Crown Street, Liverpool L69 7ZB, UK.
Tel.: +44 1517954517; fax: +44 1517954408; e-mail: jp@senescence.info

Accepted for publication 19 October 2008

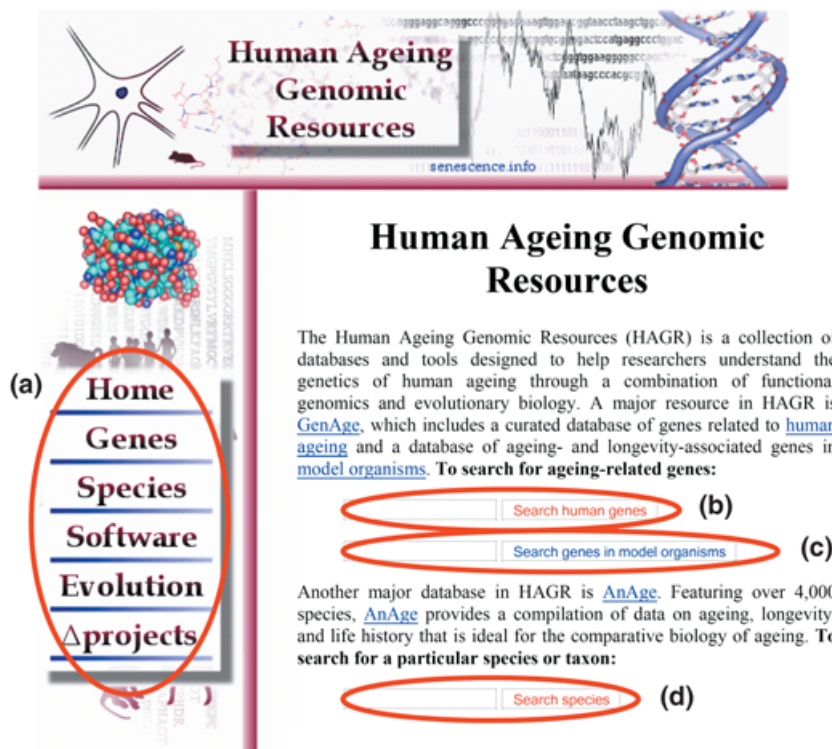


Fig. 1 From the main page of the Human Ageing Genomic Resources (HAGR; <http://genomics.senescence.info/>), users have hyperlinks to the various databases and tools (a). Users may also search directly the HAGR databases: GenAge human genes data set (b); GenAge model organisms' genes data set (c); and AnAge (d). HAGR pages are interconnected and feature hyperlinks to other relevant databases and tools as well as to the main menu.

Table 1 Species in GenAge model organisms

Organism (species)	Number of genes
Roundworm (<i>Caenorhabditis elegans</i>)	555
Baker's yeast (<i>Saccharomyces cerevisiae</i>)	87
Fruit fly (<i>Drosophila melanogaster</i>)	75
Mouse (<i>Mus musculus</i>)	68
Filamentous fungus (<i>Podospira anserina</i>)	2
Golden hamster (<i>Mesocricetus auratus</i>)	1
Fission yeast (<i>Schizosaccharomyces pombe</i>)	1

the environment, and how together genes give rise to a given phenotype: the so-called systems biology approach. To help researchers address these issues we created GenAge, a database of genes related to longevity and/or aging.

GenAge consists of several searchable data sets. Considering the extraordinary discoveries in the genetics of aging in model organisms, GenAge includes a data set of genes associated with longevity and/or aging in model organisms. We consider a given gene for inclusion in GenAge if genetic manipulations of the gene result in noticeable changes in the aging phenotype and/or longevity. Most genes in GenAge are from the four typical model organisms: mice, worms, fruit flies and yeast (Table 1). Strikingly, homologues of many genes – such as insulin receptors and sirtuins – have been shown to regulate aging in model organisms separated by large evolutionary distances (Kenyon, 2005; Liu *et al.*, 2005; Smith *et al.*, 2008). Moreover, we have shown that genes associated with aging and/or longevity in model organisms are evolutionary conserved in terms of having more homologues than predicted by chance (Budovsky *et al.*,

2007, 2008) and exhibiting slower molecular evolution rates (de Magalhães & Church, 2007). Therefore, it is now clear that at least some genes identified in model organisms may be relevant to human aging.

To allow researchers to focus specifically on human aging, GenAge features a data set of genes that may regulate aging in humans or that at least appear to be considerably associated with the human aging phenotype. This data set includes orthologues – derived from established databases, mainly In-Paranoid (O'Brien *et al.*, 2005) but also HomoloGene (<http://www.ncbi.nlm.nih.gov/homologene>) – of genes strongly associated with aging in model organisms. Also included are genes in which mutations result in segmental progeroid syndromes, such as the Werner's syndrome gene, as well as genes critical in pathways previously related to aging, such as the insulin/insulin-like signalling pathway (de Magalhães *et al.*, 2005a). The main reason why each gene is included in the database is detailed in each entry (Table 2). Even though we are sure that GenAge will continue to change in the future as a result of new breakthroughs, this data set represents the genes thus far most strongly associated with human aging, which opens several opportunities for research (see below).

Lastly, one of the latest additions to GenAge is a list of human genes analyzed for their possible association with human longevity (<http://genomics.senescence.info/genes/longevity.html>). All longevity association studies in humans we could find by the time of the latest update were added to this list. These include studies reporting negative results, which we see as essential since many genes display population-specific associations with longevity.

Table 2 Criteria used to select entries for inclusion in the GenAge human data set

Main reason for selection	Number of genes
Evidence directly linking the gene product to aging in humans	3
Evidence directly linking the gene product to aging in a mammalian model organism	27
Evidence directly linking the gene product to aging in a non-mammalian animal model	25
Evidence directly linking the gene product to aging in a cellular model system	14
Evidence directly linking the gene product to the regulation or control of genes previously linked to aging	29
Evidence linking the gene product to a pathway or mechanism linked to aging	64
Indirect or inconclusive evidence linking the gene product to aging or showing the gene product to be an effector of genes related to aging	99

As of this writing, build 14 of GenAge, updated on 20 July 2008, features 261 human genes and 789 genes in model organisms plus 258 longevity association experiments in humans. GenAge can be accessed at <http://genomics.senescence.info/genes/>.

Usage and applications

In the modern information age, even established experts in the field of aging may find it difficult to keep up-to-date with all the data being generated. As such, the first use of GenAge is as a reference database, and GenAge is in its essence an organized information source on the genetics of aging. In particular for the human data set, a wealth of biological data is provided for each entry, including relevant information in the context of biogerontology (see example below) to an extent that is not available in larger, more generic databases like Entrez Gene. Therefore, to learn about the involvement of a given gene in aging, a quick search in GenAge is the best place to start. GenAge features a sophisticated search engine and its user-friendly interface is easy to use by anyone with access to the Internet.

Searching GenAge can be done from its main page, which also includes hyperlinks to the human and models data sets. The established nomenclature of the human data set is that of the Human Genome Nomenclature Committee (HGNC), but genes' common names and frequently used aliases – obtained from HGNC and organism-specific databases (described below) – can also be searched. For example, searching 'ghbp' (HAGR search pages are case-insensitive) in the human data set will retrieve the growth hormone receptor (GHR) since GHR is frequently referred to as 'GHBP'. In GHR's entry (<http://genomics.senescence.info/genes/entry.php?hgnc=ghr>), apart from basic biological information, the potential link between GHR and aging is described and key bibliographical references are cited: *GHR* gene products have been linked to aging in mice – and this was the major reason for GHR's inclusion in GenAge – but not in humans, even though patients with mutations in *GHR* have been documented. Searching for genes in the model organisms' data set is similar with users having the added option of being able to search for genes in a given species and, in fact, we provide quick hyperlinks to the four major model organisms.

Although the models data set comprises all genes (to our knowledge) shown by the time of the latest update to statistically

increase longevity or alter the aging process in a noticeable way, in the human data set we try to evaluate whether a given intervention is affecting the aging process itself or not. For example, many mutations may increase longevity by decreasing the incidence of specific diseases, rather than by altering the basic process of aging (de Magalhães *et al.*, 2005a, 2005b). Therefore, the human data set is not merely an extension of the work conducted in model organisms and of its bibliography, but a manually selected list of the most pertinent human aging candidate genes, each presented with a higher annotation level. We cite studies on whether the functions of aging-associated genes in model organisms are conserved in their human orthologues. Likewise, we cite flaws in previous studies based on new published observations, although we have a neutral stance on conflicting findings from different research groups. Our policy is to cite all conflicting reports and let visitors make their own decisions on how to interpret them. By contrast, each entry in GenAge model organisms has only one reference: the first publication reporting an association of the gene with longevity or aging. Moreover, one of the latest enhancements in the human data set was the inclusion of Gene Ontology annotation. Gene Ontology terms and annotation files were obtained from the Gene Ontology Consortium website (<http://www.geneontology.org/>) and provide an additional layer of description for the gene products in a cellular context (Ashburner *et al.*, 2000).

Genes capable of delaying aging or modulating multiple age-related changes represent important clues in advancing our knowledge of aging. Therefore, a second use of GenAge is as a way to integrate different mechanisms of aging. Numerous authors – including ourselves (de Magalhães & Toussaint, 2004; Budovsky *et al.*, 2007, 2008) – have attempted to reconstruct the different pathways linked with aging and even study gene networks. In such works, GenAge is an adequate resource as it provides a framework for the functional genomics of aging. For example, Xue *et al.* (2007) used GenAge to construct a modular network of aging and obtain insights into aging, including the fact that genes connecting different modules are more likely to affect longevity and/or aging, an hypothesis the authors validated experimentally in worms (Xue *et al.*, 2007). In a sense, GenAge offers an overall view of what is presently known about the genetics of aging in model organisms and in humans that can be used for numerous studies, including in contemporary functional genomics and systems biology methods.

By providing a list of candidate genes, the genes in GenAge can serve as basis for gene expression and genetic association longevity studies, including human studies, or even for clinical studies of interventions hypothesized to affect aging. In fact, recent gene expression studies have used GenAge to focus on aging-associated genes (Chen *et al.*, 2008; Hardman & Ashcroft, 2008). Because researchers may have disparate opinions regarding the relevance of different model systems to understand human aging, an important tool to investigate the human data set as a whole is GenAge's browser (<http://genomics.senescence.info/genes/browser.php>). The browser makes it possible to (among other things) retrieve only those entries that pass certain criteria related to the annotation in GenAge, such as selection process and gene function. For example, users can retrieve only genes associated with aging in non-mammalian model organisms. Several criteria can be set, all of which must be passed for genes to be retrieved. Afterwards, users can select all or a subgroup of the genes retrieved, which can then be further analysed using simple statistical tools (<http://genomics.senescence.info/genes/analyse.php>). Cross-links between the different sections of GenAge make it easy to navigate between the different tools and data sets.

Another use of GenAge is for researchers to associate genes already under investigation with other, little-known genes, which can lead to new experimental designs. To do this, protein–protein interactions are one possible approach, and GenAge's human data set features 673 interactions, most of which manually curated obtained from the Human Protein Reference Database (HPRD) (Peri *et al.*, 2003). In fact, one of our earliest applications of GenAge involved finding novel genes that may be linked to aging by way of an analysis of protein–protein interactions. The principle being that proteins not previously thought to be related to aging which interact with a large number of proteins directly linked to aging might too be involved in aging and are thus promising candidates for future studies (de Magalhães & Toussaint, 2004; Budovsky *et al.*, 2007). Similar works are made easy with GenAge. Protein–protein interactions with one or more genes as query can be visualized (Fig. 2), or they can be downloaded for use with more advanced biological pathway analysis software.

GenAge can be helpful in more classical genetic studies of aging and longevity. For example, if a given chromosomal region is identified as suspect in influencing longevity, it is now possible to look up which GenAge genes are present in that region (<http://genomics.senescence.info/genes/map.php>). This new feature of GenAge may assist researchers trying to map genes influencing aging or longevity in human populations.

Although it was not our intention that GenAge became a bibliographic database, the manually included literature in the human data set (1956 references in the latest build) may prove a valuable resource. In addition to providing a searchable bibliographic database where visitors can search titles, authors and PubMed IDs (<http://genomics.senescence.info/genes/biblio.php>), and look up which references are most cited within GenAge, one novel feature in GenAge is bibliographic

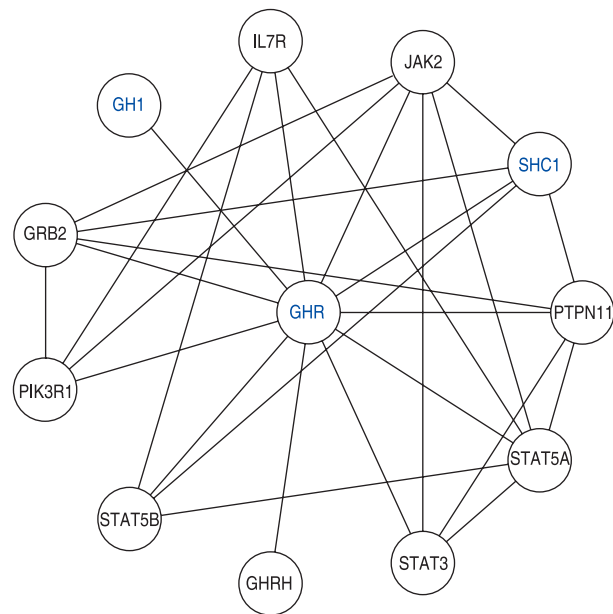


Fig. 2 Protein–protein interactions of the growth hormone receptor (GHR) with one degree of freedom as generated by Human Ageing Genomic Resources' Interactions Graphical Display (IGD). Entries directly linked to mammalian aging are highlighted in blue.

data-mining. The system uses the simplest method available, which is that of text co-occurrence (Stephens *et al.*, 2001): the principle being that genes or proteins mentioned in the same academic work have higher chances of being related. Consequently, it is now possible to search through the bibliography of GenAge to find genes referred to in the same work to generate putative functional links (http://genomics.senescence.info/genes/mine_biblio.php). To use this feature, users must select – using the browser or search pages – or type the GenAge entries they wish to look up, and then set the threshold defining how many references two genes must have in common for them to be considered linked. This process can be useful, for instance, to discover novel associations between genes of interest – which may serve as inspiration to develop novel experiments – or to be kept up-to-date with the growing literature in gerontology.

Of course, GenAge has its limitations and while we aim to include the most relevant information (in particular in the context of genetics and aging), not all data are available. For instance, searching 'GHR' in PubMed yields (at the time of writing) roughly 1000 results. It would be unreasonable to include them all in GenAge. GenAge is an integrative database, however, that draws together data obtained from multiple other sources, including OMIM, SWISS-PROT, Entrez Gene and HPRD (de Magalhães *et al.*, 2005a). Hyperlinks to these and other major databases are included to help researchers quickly locate additional information within other or broader biological scopes. Model organisms' genes also have relevant hyperlinks specific to each organism: mouse genes link to the Mouse Genome Informatics website (<http://www.informatics.jax.org/>),

fly genes link to FlyBase (<http://flybase.bio.indiana.edu/>), worm genes link to WormBase (<http://www.wormbase.org/>), and yeast genes link to the Saccharomyces Genome Database (<http://www.yeastgenome.org/>), in addition to hyperlinks to the AnAge database described below.

For the above reasons, the design and information in GenAge is prone to automated computational methods, and GenAge can be subject to a variety of data-mining algorithms. Since we included, if available, open reading frame, promoter and protein sequences for each entry, a variety of bioinformatics approaches may be used, such as studies of gene networks and transcriptional regulation. For example, we studied selection pressures on GenAge genes since humans and chimpanzees diverged to obtain novel insights, such as showing that aging-related genes evolve under stronger functional constraints than expected by chance (de Magalhães & Church, 2007). If researchers wish to apply sophisticated computational methods to data-mine GenAge, the whole database can be downloaded as tab-delimited file or as an XML file.

AnAge: the animal aging and longevity database

Philosophy and overview of resources

As argued by many others (Finch, 1990; Miller, 2001; Austad, 2005), if we could understand why different species age at different rates, we would be able to gather key insights into the mechanisms of aging. Comparative studies of aging are thus one established way of conducting research on aging, and multiple parameters can be compared between species (Finch, 1990; Frolkis & Muradian, 1991). With the comparative biology of aging as its primary goal, AnAge is a database of longevity and aging in animal species. Although commonly used model organisms such as yeast are included, AnAge focuses primarily on chordates and a greater attention is given to species evolutionary closer to humans such as primates and mammals (Table 3). Arguably, AnAge is the most extensive, complete and accurate database of longevity records available in any format. In addition, observations on physiological or pathological changes with age in animals are (where available) featured.

The main aim of AnAge is to provide data for comparative and evolutionary studies of aging that allow researchers to study how and why the aging phenotype and longevity vary across phylogeny and which factors influence such differences. Although demographic measurements of aging are included in AnAge,

these require detailed animal studies that are rarely available and thus represent only a small fraction of the data. Maximum lifespan as estimated from longevity records will continue to be the standard – even if not perfect – measure of aging (de Magalhães, 2006), and this is available for the large majority of AnAge entries. Recently, we incorporated the longevity information from FishBase (Froese & Pauly, 2008) and from Richard Weigl's excellent compilation of mammalian longevity records in captivity (Weigl, 2005) as well as data from mammalian aging expert Steven Austad, which includes personal correspondence with zoo personnel, field biologists and veterinarians.

Included in AnAge is also extensive data for a number of factors that might bias comparative studies of aging, such as body size, metabolic rates and developmental schedules, so authors can analyze these factors in their studies. Otherwise, correlations with longevity may merely reflect a correlation between the parameter under study and, for example, body size which in turn tends to correlate with longevity (Charnov, 1993). Indeed, we recently conducted a number of analyses of factors related to longevity (de Magalhães *et al.*, 2007). Succinctly, we showed that metabolic rate, when corrected for the effects of body size and phylogeny, does not correlate with longevity or aging rates. In contrast, developmental schedules such as age at maturity appear to be robust predictors of longevity and aging rates. This work highlights the importance of taking potentially confounding traits in consideration in aging studies that employ the comparative method. Consequently, a great effort has been made to include updated and extensive developmental and reproductive data in AnAge.

At the time of writing (build 10, updated on 18 April 2008), AnAge features 4122 entries, including 1331 mammals (Table 3). Our focus is on accuracy and quality, however, not quantity, and only species for which we have confidence in the data are featured. For instance, longevity records based on single or a few observations, particularly if not in captivity, are normally excluded, at least if much better data are available for similar species (e.g. for birds, bats and cetaceans AnAge features extensive records from animals in the wild since very few records in captivity are available). Likewise, although we use data from websites, these have to come from authoritative sources, such as zoos, parks and academic institutions. Unverified longevity records or those from unreliable sources may be cited as anecdotes in the observations section of each entry but do not become the established maximum lifespan for that species, as our quantitative data has to be consistent in order to serve as basis for comparative studies. AnAge is available at <http://genomics.senescence.info/species/>.

Table 3 Number of entries per taxa in AnAge

Taxa (scientific name)	Number of entries
Mammals (Mammalia)	1331
Birds (Aves)	1098
Reptiles (Reptilia)	539
Amphibians (Amphibia)	169
Fishes (Actinopterygii, Chondrichthyes, etc.)	962
Non-chordates	28

Usage and applications

Like GenAge, AnAge can be useful as an information source and Web portal by allowing researchers to learn more about the aging process of a particular species in a user-friendly environment. To use AnAge, its search engine – also accessible from AnAge's main page – is the place to start (<http://genomics.senescence.info/species/>).

senescence.info/species/query.php). Species and common names can be searched (see example below). In fact, terms at any taxonomic level can be searched and a recent addition to AnAge has been a number of entries at higher taxonomic levels. The taxonomy of AnAge follows that of the Integrated Taxonomic Information System (<http://www.itis.usda.gov>).

AnAge provides an up-to-date synthesis of what is known about longevity and aging of a given organism, which is particularly useful for those not familiar with it. Furthermore, species with unique aging phenotypes or of special interest to biogerontologists, such as species with negligible senescence, are featured. Again, like GenAge, AnAge has its limits and external hyperlinks point users to further sources of information, such as the Tree of Life (<http://tolweb.org>) and Animal Diversity Web (<http://animaldiversity.ummz.umich.edu>) websites. To obtain pictures of a given species the popular search engine Google (<http://images.google.com>) is used.

Like GenAge, AnAge features a browser (<http://genomics.senescence.info/species/browser.php>). AnAge's browser is somewhat different, though, as it allows users to browse through the taxonomy of species in AnAge. Once a given set of entries has been selected, it can be surveyed to gather simple descriptive statistics (<http://genomics.senescence.info/species/survey.php>). As in GenAge, it is easy to navigate between the different AnAge tools due to extensive cross-links. Several species of interest can be selected with a mouse click and the parameter of interest's average \pm standard deviation can be calculated. For instance, to search for turtles one would simply type 'testudines', the order to which turtles belong to. The list of 122 turtles in AnAge would then be displayed according to taxonomy or longevity, and one could then obtain the averages and standard deviations for longevity, adult body size, age at sexual maturity, or other trait present in AnAge. The phylogeny of the species selected can be inferred from their taxonomy using the Phylogenetic Tree Plotter (<http://genomics.senescence.info/species/ptp.php>), although its use is limited to fewer than 50 species and it should not be employed as a replacement for established phylogenies; integrating published phylogenies in AnAge is, however, in our future plans.

Given the size and scope of AnAge, it is a major resource to study the evolution of aging, longevity and life history. Indeed, several groups have taken advantage of AnAge to study the evolutionary events responsible for species differences in aging and longevity. For example, a number of studies have used AnAge to investigate how features of the mitochondrial DNA, such as mutation and substitution rates, relate to longevity and other life-history traits (Khaidakov *et al.*, 2006; Lehmann *et al.*, 2006, 2008; Moosmann & Behl, 2008; Nabholz *et al.*, 2008; Welch *et al.*, 2008). Others have used AnAge to compare experimentally determined parameters, such as cellular stress resistance (Harper *et al.*, 2007), rates of free radical production in mitochondria (Lambert *et al.*, 2007) and telomerase activity (Seluanov *et al.*, 2007), with species longevity.

Given the size of AnAge, downloading the entire database as a tab-delimited or XML file may be necessary for some analyses and is also possible.

Software, tutorials and additional tools

It is crucial that the information in HAGR is kept up-to-date and, thus far, AnAge and GenAge have been updated regularly. To update our databases we make use of computer scripts to automatically retrieve data from other online resources. As believers in open source and sharing, not only the databases but all these scripts, most of which written in Perl, are available for download as part of the Ageing Research Computational Tools (ARCT). Admittedly, a basic understanding of Perl is necessary for ARCT to be useful. This is by far the greatest weakness of the ARCT toolkit, which thus appears to be useful only to those in the community with a moderate-to-high level of computer knowledge. Still, it may be a valuable set of tools to those wishing to apply bioinformatics to study aging. ARCT is available for download at <http://genomics.senescence.info/software/>.

Also available is a script for use with SPSS for demographic analysis. Succinctly, this script allows users with the SPSS statistical package (SPSS Inc., Chicago, IL, USA) to determine whether the slope of the age-related mortality curve changes between two or more different cohorts, a method commonly employed to infer whether the rate of aging was altered or not, as detailed previously (de Magalhães *et al.*, 2005b).

The HAGR are part of the parent website [senescence.info](http://www.senescence.info) (<http://www.senescence.info>), which one of us (J.P.M.) started in 1997 as a mix of speculative essays, hypotheses and reviews. Primarily, [senescence.info](http://www.senescence.info) is an educational resource on the science of aging, including not only information on the biology of aging and advice for students but discussions on relevant issues to gerontology, such as the sociological implications of aging research and lifespan increase. Of interest, [senescence.info](http://www.senescence.info) features numerous tutorials regarding aging research aimed at students and newcomers to the field.

Lastly, a non-exhaustive list of people and companies working on aging and related disciplines is available at the WhosAge database (<http://whoswho.senescence.info/>). Presently, WhosAge features 171 researchers and 13 companies, including a brief description of the work of each company or individual, contact information and hyperlinks to websites (if available) and further sources of information such as hyperlinks to PubMed and Google searches. WhosAge is searchable – it employs the engine of GenAge and AnAge – and can be used to find a specific individual or company, or search for researchers working at a given institution or on a particular topic.

Concluding remarks

We designed the databases and tools in HAGR primarily so they can be useful to researchers working on the science of aging. The audience of HAGR ranges from beginners to experts in gerontology, from experimentalists to theorists, from computer illiterates to expert computational biologists. Generally, while GenAge and AnAge use simple and intuitive navigation tools, the ability to download the databases gives ample freedom to computational biologists to employ more powerful methods, as

exemplified in the numerous works that have already taken advantage of these data sets.

Because our aim has always been to develop HAGR in collaboration with the gerontology community, nearly all pages have a hyperlink to our contact information or to an easy-to-use feedback form. Visitors are thus encouraged to send us feedback on errors they find, propose enhancements, and suggest features they would like to see in HAGR – and many have already done so. Understanding aging is a collaborative effort. The HAGR will hopefully reflect that spirit and continue to improve not only due to our efforts but thanks to the feedback from the research community working on the science of aging. In the past 12 months, these online resources attracted roughly $15\,000 \pm 7500$ (SD) unique users per month. In the future, we hope to continually upgrade, update and expand the resources in HAGR, as well as develop new tools that can benefit the gerontology community. Ideally, we would like HAGR to become a common platform for biogerontologists and thus invite everyone in the community to visit, use and help us improve HAGR.

Acknowledgments

We are indebted to many people who have helped in this project, of which we highlight the assistance and opinions of Olivier Toussaint and Steven Austad. J. P. de Magalhães was supported by a NIH-NHGRI CEGS grant to G. Church. The work of V. Fraifeld's laboratory is partially supported by the European Union FP7 Health Research Grant number HEALTH-F4-2008-202047.

References

- Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, Harris MA, Hill DP, Issel-Tarver L, Kasarskis A, Lewis S, Matese JC, Richardson JE, Ringwald M, Rubin GM, Sherlock G. (2000) Gene ontology: tool for the unification of biology. The gene ontology consortium. *Nat. Genet.* **25**, 25–29.
- Austad SN (2005) Diverse aging rates in metazoans: targets for functional genomics. *Mech. Ageing Dev.* **126**, 43–49.
- Budovsky A, Abramovich A, Cohen R, Chalifa-Caspi V, Fraifeld V (2007) Longevity network: construction and implications. *Mech. Ageing Dev.* **128**, 117–124.
- Budovsky A, Tacutu R, Yanai H, Abramovich A, Wolfson M, Fraifeld V (2008) Common gene signature of cancer and longevity. *Mech. Ageing Dev.* doi:10.1016/j.mad.2008.04.002.
- Charnov EL (1993) *Life History Invariants: Some Explorations of Symmetry in Evolutionary Ecology*. Oxford, UK: Oxford University Press.
- Chen DP, Weber SC, Constantinou PS, Ferris TA, Lowe HJ, Butte AJ (2008) Novel integration of hospital electronic medical records and gene expression measurements to identify genetic markers of maturation. *Pac. Symp. Biocomput.* 243–254.
- Finch CE (1990) *Longevity, Senescence, and the Genome*. Chicago, IL: University of Chicago Press.
- Finch CE, Ruvkun G (2001) The genetics of aging. *Annu. Rev. Genomics Hum. Genet.* **2**, 435–462.
- Froese R, Pauly D (2008) *FishBase*. <http://www.fishbase.org> (version April 2008).
- Frolkis VV, Muradian KK (1991) *Life Span Prolongation*. Boca Raton, FL: CRC Press.
- Hardman MJ, Ashcroft GS (2008) Estrogen, not intrinsic aging, is the major regulator of delayed human wound healing in the elderly. *Genome Biol.* **9**, R80.
- Harper JM, Salmon AB, Leiser SF, Galecki AT, Miller RA (2007) Skin-derived fibroblasts from long-lived species are resistant to some, but not all, lethal stresses and to the mitochondrial inhibitor rotenone. *Ageing Cell* **6**, 1–13.
- Kenyon C (2005) The plasticity of aging: insights from long-lived mutants. *Cell* **120**, 449–460.
- Khaidakov M, Siegel ER, Shmookler Reis RJ (2006) Direct repeats in mitochondrial DNA and mammalian lifespan. *Mech. Ageing Dev.* **127**, 808–812.
- Lambert AJ, Boysen HM, Buckingham JA, Yang T, Podlutzky A, Austad SN, Kunz TH, Buffenstein R, Brand MD (2007) Low rates of hydrogen peroxide production by isolated heart mitochondria associate with long maximum lifespan in vertebrate homeotherms. *Ageing Cell* **6**, 607–618.
- Lehmann G, Budovsky A, Muradian KK, Fraifeld VE (2006) Mitochondrial genome anatomy and species-specific lifespan. *Rejuvenation Res.* **9**, 223–226.
- Lehmann G, Segal E, Muradian KK, Fraifeld VE (2008) Do mitochondrial DNA and metabolic rate complement each other in determination of the mammalian maximum longevity? *Rejuvenation Res.* **11**, 409–417.
- Liu X, Jiang N, Hughes B, Bigras E, Shoubridge E, Hekimi S (2005) Evolutionary conservation of the clk-1-dependent mechanism of longevity: loss of mclk1 increases cellular fitness and lifespan in mice. *Genes Dev.* **19**, 2424–2434.
- de Magalhães JP (2005) Open-minded scepticism: inferring the causal mechanisms of human ageing from genetic perturbations. *Ageing Res. Rev.* **4**, 1–22.
- de Magalhães JP (2006) Species selection in comparative studies of aging and antiaging research. In *Handbook of Models for Human Aging* (Conn PM, ed.). Burlington, MA: Elsevier Academic Press, pp. 9–20.
- de Magalhães JP, Toussaint O (2004) GenAge: a genomic and proteomic network map of human ageing. *FEBS Lett.* **571**, 243–247.
- de Magalhães JP, Church GM (2007) Analyses of human-chimpanzee orthologous gene pairs to explore evolutionary hypotheses of aging. *Mech. Ageing Dev.* **128**, 355–364.
- de Magalhães JP, Costa J, Toussaint O (2005a) HAGR: the Human Ageing Genomic Resources. *Nucleic Acids Res.* **33** (Database Issue), D537–543.
- de Magalhães JP, Cabral JA, Magalhaes D (2005b) The influence of genes on the aging process of mice: a statistical assessment of the genetics of aging. *Genetics* **169**, 265–274.
- de Magalhães JP, Costa J, Church GM (2007) An analysis of the relationship between metabolism, developmental schedules, and longevity using phylogenetic independent contrasts. *J. Gerontol. A Biol. Sci. Med. Sci.* **62**, 149–160.
- Miller RA (2001) A position paper on longevity genes. *Sci. Aging Knowledge Environ.* **2001**, vp6.
- Moosmann B, Behl C (2008) Mitochondrially encoded cysteine predicts animal lifespan. *Ageing Cell* **7**, 32–46.
- Nabholz B, Glemin S, Galtier N (2008) Strong variations of mitochondrial mutation rate across mammals – the longevity hypothesis. *Mol. Biol. Evol.* **25**, 120–130.
- O'Brien KP, Remm M, Sonnhammer EL (2005) Inparanoid: a comprehensive database of eukaryotic orthologs. *Nucleic Acids Res.* **33**, D476–D480.
- Olshansky SJ, Hayflick L, Carnes BA (2002) No truth to the fountain of youth. *Sci. Am.* **286**, 92–95.
- Peri S, Navarro JD, Amanchy R, Kristiansen TZ, Jonnalagadda CK, Surendranath V, Niranjana V, Muthusamy B, Gandhi TK, Gronborg M, Ibarrola N, Deshpande N, Shanker K, Shivashankar HN, Rashmi BP, Ramya MA, Zhao Z, Chandrika KN, Padma N, Harsha HC, Yatish AJ,

- Kavitha MP, Menezes M, Choudhury DR, Suresh S, Ghosh N, Saravana R, Chandran S, Krishna S, Joy M, Anand SK, Madavan V, Joseph A, Wong GW, Schiemann WP, Constantinescu SN, Huang L, Khosravi-Far R, Steen H, Tewari M, Ghaffari S, Blobel GC, Dang CV, Garcia JG, Pevsner J, Jensen ON, Roepstorff P, Deshpande KS, Chinnaiyan AM, Hamosh A, Chakravarti A, Pandey A. (2003) Development of human protein reference database as an initial platform for approaching systems biology in humans. *Genome Res.* **13**, 2363–2371.
- Seluanov A, Chen Z, Hine C, Sasahara TH, Ribeiro AA, Catania KC, Presgraves DC, Gorbunova V (2007) Telomerase activity coevolves with body mass not lifespan. *Aging Cell* **6**, 45–52.
- Smith ED, Tsuchiya M, Fox LA, Dang N, Hu D, Kerr EO, Johnston ED, Tchao BN, Pak DN, Welton KL, Promislow DE, Thomas JH, Kaeberlein M, Kennedy BK. (2008) Quantitative evidence for conserved longevity pathways between divergent eukaryotic species. *Genome Res.* **18**, 564–570.
- Stephens M, Palakal M, Mukhopadhyay S, Raje R, Mostafa J (2001) Detecting gene relations from Medline abstracts. *Pac. Symp. Biocomput.* 483–495.
- Vijg J, Suh Y (2005) Genetics of longevity and aging. *Annu. Rev. Med.* **56**, 193–212.
- Weigl R (2005) *Longevity of Mammals in Captivity; from the Living Collections of the World*. Stuttgart, Germany: Kleine Senckenberg-Reihe.
- Welch JJ, Bininda-Emonds OR, Bromham L (2008) Correlates of substitution rate variation in mammalian protein-coding sequences. *BMC Evol. Biol.* **8**, 53.
- Xue H, Xian B, Dong D, Xia K, Zhu S, Zhang Z, Hou L, Zhang Q, Zhang Y, Han JD (2007) A modular network model of aging. *Mol. Syst. Biol.* **3**, 147.