



HAL
open science

DetecDiv, a deep-learning platform for automated cell division tracking and replicative lifespan analysis

Théo Aspert, Didier Hentsch, Gilles Charvin

► **To cite this version:**

Théo Aspert, Didier Hentsch, Gilles Charvin. DetecDiv, a deep-learning platform for automated cell division tracking and replicative lifespan analysis. 2021. hal-03451298

HAL Id: hal-03451298

<https://hal.science/hal-03451298>

Preprint submitted on 26 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DetecDiv, a deep-learning platform for automated cell division tracking and replicative lifespan analysis

Théo Aspert^{1234*}, Didier Hentsch¹²³⁴ and Gilles Charvin^{1234*}

1) Department of Developmental Biology and Stem Cells, Institut de Génétique et de Biologie Moléculaire et Cellulaire, Illkirch, France

2) Centre National de la Recherche Scientifique, UMR7104, Illkirch, France

3) Institut National de la Santé et de la Recherche Médicale, U964, Illkirch, France

4) Université de Strasbourg, Illkirch, France

*: Corresponding authors: charvin@igbmc.fr, aspert@igbmc.fr

Abstract (65 words)

Automating the extraction of meaningful temporal information from sequences of microscopy images represents a major challenge to characterize dynamical biological processes. Here, we have developed DetecDiv, a microfluidic-based image acquisition platform combined with deep learning-based software for high-throughput single-cell division tracking. DetecDiv can reconstruct cellular replicative lifespans with an outstanding accuracy and provides comprehensive temporal cellular metrics using timeseries classification and image semantic segmentation.

Main

Yeast has a limited replicative lifespan (RLS, i.e., 20-30 divisions) before entering senescence and dying (Mortimer and Johnston 1959). Over the last decades, this simple unicellular has become a reference model for understanding the fundamental mechanisms that control longevity (Denoth-Lippuner, Julou, and Barral 2014; Janssens and Veenhoff 2016). Several independent mechanistic models have been proposed to explain entry into replicative senescence (Hughes and Gottschling 2012; Sinclair and Guarente 1997; Aguilaniu et al. 2003). In this context, whether there are multiple parallel causes responsible for senescence remains highly debated (Dillin, Gottschling, and Nyström 2014; C. He, Zhou, and Kennedy 2018). A crucial difficulty in solving this puzzle lies in the very labor-intensive nature of RLS assays (McCormick et al. 2015) and the limited information derived from them regarding the dynamics of senescence entry. The initial development of microfluidic systems for RLS assays has partially alleviated this problem by allowing continuous observation of individual cell divisions and relevant fluorescent cellular markers under the microscope from birth to death (Lee et al. 2012; Xie et al. 2012; Fehrmann et al. 2013). Recent efforts further increased data acquisition throughput (Jo et al. 2015; Liu, Young, and Acar 2015) and attempted to automate data analysis (Ghafari et al. 2021; Ghafari, Mailman, and Qin 2021). Yet, retrieving individual cellular lifespans from large sets of image sequences so far remained an insurmountable bottleneck to characterize senescence entry quantitatively or to screen large numbers of mutants and environmental conditions.

To overcome this challenge, we have developed DetecDiv, an integrated platform that combines high-throughput observation of cell divisions using a microfluidic device, a simple benchtop image acquisition system, and a deep learning-based image processing software with several image classification frameworks. The microscope was built using a rigid frame with inverted epifluorescence optics, transmission (bright field) illumination, a camera, and a motorized stage (Fig. S1A and B). The motorized frame carries the microfluidic device to trap individual cells and follow their successive divisions from birth to death (Fig. 1A, S1C and S1D). Even though its principle is similar to previously reported designs (Jo et al. 2015; Liu, Young, and Acar 2015; Crane et al. 2014), we have brought significant improvements to the trap geometry to improve cell retention, avoid replacement of a mother cell by its daughter, and prevent device clogging and any source of contamination (see Fig S1E-G and supplementary text for details). The device includes 16 independent chambers (with 2000 traps per chamber) to image different strains in parallel or to vary environmental conditions. Altogether, this system allows following the successive divisions and the entry into senescence of typically 30000 individual cells in parallel with a 5-min resolution (knowing

that there are ~500 traps per field of view using a 20x objective), i.e., about 1 to 2 orders of magnitude above the previously described techniques (Lee et al. 2012; Jo et al. 2015).

This image acquisition system generates a large amount of cell division data (on the Terabytes scale depending on the number of channels, frames, and fields of view), only a tiny part of which can be manually curated in a reasonable time. In particular, the determination of replicative lifespans requires counting successive cell divisions until death, hence, reviewing all images acquired for each cell in each field of view over time. In addition, automating the division counting process is complicated by the heterogeneity in cell fate (i.e. division times, cell shape), especially during the entry into senescence. To overcome this limitation, we have developed an image classification pipeline to count divisions and reconstruct the entire lifespan of individual cells dividing in the traps. For this, we have trained a GoogleNet convolutional neural network (CNN) (Szegedy et al. 2015) to determine the budding state of the trapped cells by assigning one of six possible classes (unbudded, small-budded, large-budded, dead, empty trap, clogged trap) to each frame, see Fig. S2 and S3A. In this framework, the alternation between the 'large budded' and the 'unbudded'/'small budded' states revealed the successive cell divisions, and the occurrence of the 'dead' class allowed us to reconstruct the cell's lifespan (Fig. 1C)

However, the assignment of the cellular state based on the CNN, which processes the images independently of one another, led to sporadic ambiguities and errors that compromise the accuracy of the distribution of division times (Fig. 1D) and *a fortiori* of the lifespans (Fig. 1C-E). These problems could be partially alleviated by post-processing the predictions made by the CNN (see "CNN+PP" in Fig. 1C-E and supplementary text for details). Yet, to improve the robustness of the method, we have combined the CNN with a long short-term memory network (LSTM) (Venugopalan et al. 2015; Hochreiter and Schmidhuber 1997), to take into account the time-dependencies between images (Fig. 1B). Thus, by providing full image sequences rather than individual images (Fig. 1B), we obtained an outstanding accuracy and recall for both division quantification and lifespan reconstruction (Fig. 1D, 1E and Movie M1). Comparing the predictions by the classifier to the manually annotated data ("ground truth") revealed a non-significant difference in the distributions (and an excellent correlation ($R^2=0.99$ for both divisions and lifespans, Fig 1D and 1E). To estimate the robustness of the classification model which was only trained on images of WT cells, we measured the large-scale RLS in classical longevity mutants. Remarkably, we recapitulated the increase (resp. decrease) in longevity observed in the *fab1Δ* (resp. *sir2Δ*) mutant (Defossez et al. 1999; Lin, Defossez, and Guarente 2000) and we could compute the related death rate with high accuracy (Fig 1F) (Morlot et al. 2019). Importantly, only 200 manually annotated lifespans were necessary to achieve robust RLS reconstruction. Thus, rapid user annotation of a small cohort of cells allows the model to be deployed on larger datasets in different genetic contexts and/or environments.

We then sought to apply other classification schemes of DetecDiv to further characterize the trajectories of the cells as they transition to senescence. First, we set up an LSTM sequence-to-sequence classifier to detect the onset of cell-cycle slowdown during entry into senescence (also referred to as the senescence entry point or SEP (Fehrmann et al. 2013)), see Fig. 2A and S4. We thus trained the classifier to assign a 'pre-SD' or 'post-SD' label (before and after cell-cycle slowdown, respectively) to each frame, using the sequence of cellular state probabilities (i.e., the output of the CNN/LSTM image classifier described in Fig. 1) as input. Using this method, we could successfully identify the transition to a slow

division mode (Fig. 2A) and recapitulate the evolution of average division times after aligning individual trajectories from that transition (Fig. S5).

Second, we used an encoder/decoder network based on a Resnet50 CNN (K. He et al. 2016) and the DeepLab v3+ architecture (Chen et al. 2018), see Fig. S6, to segment brightfield and fluorescence images of cells carrying a histone-Neogreen fusion (see Fig 2B, Movie M2 and supplementary text) (Chen et al. 2018). After training the model on ~1500 manually segmented brightfield images using three output classes (i.e., 'background,' 'mother cell,' 'other cell'), we obtained accurate mother cells contours (Fig. 2B and Fig. S7). This allowed us to quantify the cellular volume increase, as previously reported (Morlot et al. 2019). A similar training procedure with ~3000 fluorescence images yielded accurate nuclei contours (see Fig. 2C and S8). It successfully recapitulated the sharp burst in histone fluorescence that follows cell-cycle slowdown (Fig. 2C and S5) (Morlot et al. 2019), hence further validating our methodology.

Overall, DetecDiv provides an integrated system to acquire and track cell division events with high throughput, which unleashes the potential of microfluidic cell trapping devices to perform fully automated replicative lifespan analyses. The imaging system was designed to perform heavy duty image acquisition sequences (i.e, no filter wheel, fixed objective) to generate high throughput microscopy datasets. The hardware could be easily assembled from simple optical components -for a price of about one-third that of a commercial automated microscope. By processing temporal sequences of images rather than individual ones, our software demonstrated an outstanding accuracy that matches human capabilities for image classification yet with a much higher throughput. The robustness of the imaging pipeline benefited from improvements made in the design of the microfluidic device (see supplementary text). Therefore, our framework now overcomes all intrinsic technical limitations of conventional RLS assays and provides an unprecedented potential to perform large screens for players and environmental perturbations that dynamically control replicative longevity. More broadly, this work illustrates how temporal dependencies in image sequences can be exploited using a combined CNN and LSTM architecture to accurately reveal and quantify dynamic cellular processes. Despite large efforts to make deep-learning models available to the community of microscopists, very little work has attempted to fully exploit information encoded in image sequences. With its comprehensive set of generic classification schemes that can be fully user-parameterized, DetecDiv may be used well beyond the scope of the present study and applied to any biological context with complex temporal patterns (cellular differentiation, cell division, organelles dynamics, etc).

Methods

Strains

All strains used in this study are congenic to S288C (see Supplementary Table 1 for details). See supplementary methods for detailed protocols for cell culture.

Microfabrication and microfluidics

The designs were created on AutoCAD to produce chrome photomasks (jd-photodata, UK). The microfluidic master molds were then made by standard photolithography processes (see supplementary text for details).

The microfluidic device is composed of geometric microstructures that allow mother cells trapping and flushing of successive daughter cells (see Fig. S2 and supplementary text). The cell retention efficiency of the traps is 99% after the five first divisions. We designed a particle filter with a cutoff size of 5 μm to prevent dust particles or debris from clogging the chip. The microfluidic chips were fabricated with PDMS using standard methods (PDMS, Sylgard 184, Dow Chemical, USA, see supplementary text for detailed protocols). We connected the chip using PTFE tubing (1mm OD), and we used a peristaltic pump to ensure media replenishment (Ismatec, Switzerland). We used standard rich media supplemented with 2% dextrose (YPD). See supplementary methods for additional details.

Microscopy

The microscope was built from a modular microscope system with a motorized stage (ASI, USA, see the supplementary text for the detailed list of components), a 20x objective 0.45 (Nikon, Japan) lens, and an sCMOS camera (ORCA flash 4.0, Hamamatsu, Japan). A dual band filter (#59022, Chroma Technology, Germany) coupled with a two-channel LED system (DC4104 and LED4D067, Thorlabs, USA). Sample temperature was maintained at 30°C thanks to a heating system based on an Indium Thin Oxide coated glass and an infrared sensor coupled to an Arduino-based regulatory loop. Micromanager v2.0. (<https://micro-manager.org/>) was used to drive all hardware, including the camera, the light sources, and the stage and objective motors. We developed a custom autofocus routine to minimize the autofocus time (https://github.com/TAspert/DetecDiv_Hardware). The interval between two frames for all the experiments was 5 min. We could image approximately 80 fields of view (0.65mmx0.65mm) in brightfield and fluorescence (using a dual-band GFP-mCherry filter) with this interval.

Image processing

We developed Matlab software, DetecDiv, which provides different classification models : image classification, image sequence classification, time series classification, and pixel classification (semantic segmentation), see Fig. S9. DetecDiv was developed using Matlab, and additional toolboxes (TB), such as the Computer Vision TB, the Deep-learning TB, and the Image Processing TB. A graphical user interface was designed to facilitate the generation of the training sets. The DetecDiv software is available for download on github:<https://github.com/gcharvin/DetecDiv>

Image classification for division tracking and lifespan reconstruction:

DetecDiv was used to classify images into six classes after supervised training using a GoogleNet (Szegedy et al. 2015) network combined with an LSTM network (Hochreiter and Schmidhuber 1997). See supplementary text for details.

Image segmentation using pixel classification and fluorescence quantification.

DetecDiv was used to segment images using a pixel classification model called Deeplab v3+ (Chen et al. 2018), after supervised training based on 1000-3000 manually annotated images. See supplementary text for details.

Statistics All experiments have been replicated at least twice. Data are presented in Results and Figures as the mean \pm SEM (curves) or median. Group means were compared using the Two-sample t-test. A P value of < 0.05 was considered significant.

Computing

Image processing was performed on a computing server with 8 Intel Xeon E5-2620 processors and 8 co-processing GPU units (Nvidia Tesla K80), each of them with 12Go RAM. Under these conditions, the image classification of a single trap (roughly 60x60pixels) with 1000 frames took between 3 and 5s for the CNN/LSTM classifier. For image segmentation, it took about 30s to classify 1000 images.

Data Availability

The custom MATLAB software DetecDiv, used to analyze imaging data with deep-learning algorithms, is available on <https://github.com/gcharvin/DetecDiv>. Information regarding the microfluidic device and the custom imaging system are available on [https://github.com/TAspert/DetecDiv Hardware](https://github.com/TAspert/DetecDiv_Hardware)

Acknowledgments

We thank Audrey Matifas for constant technical support throughout this work, Sophie Quintin and Nacho Molina for carefully reading the manuscript. We are grateful to Olivier Tassy for insightful discussions. We thank Denis Fumagalli at the IGBMC Mediaprep facility for media preparation. We are grateful to the IT service for efficient support and providing the computing resources. We thank the Charvin lab members, Bertrand Vernay, Jerome Mutterer, Serge Taubert and the IGBMC imaging facility for discussions and technical support. This work was supported by the Agence Nationale pour la Recherche (T.A. and G.C.), the grant ANR-10-LABX-0030-INRT, a French State fund managed by the Agence Nationale de la Recherche under the frame program Investissements d'Avenir ANR-10-IDEX-0002-02.

Figures legends

Figure 1 - DetectDiv division tracking and RLS reconstruction pipeline

A) Left: Sketch of the microfluidic device, featuring 16 independent channels with 2000 individual cell traps in each; Top right: zoom on the cellular traps, scale bar, 20 μ m; Bottom right: brightfield image of a budding yeast cell in a trap, scale bar, 5 μ m.

B) Principles of the DetecDiv division tracking and lifespan reconstruction pipeline; Brightfield images are processed by a convolutional neural network (CNN) to extract representative image features. The sequence of image features is then processed by a long short-term memory network (LSTM) that assigns one of the 6 predefined classes ('unbud', 'small', 'large', 'dead', 'clog', 'empty'), taking into account the the time dependencies. Temporal oscillations between "large" and "small" or "large" and "unbudded" indicate the completion of the division cycle. The appearance of the "dead" class marks the end of the lifespan.

C) Comparison of the different methods used for 6 sample cells. The gray bars represent the ground truth data made from manually annotated image sequences. Colored lines indicate the corresponding predictions made by CNN/LSTM (orange), the CNN + post-processing (magenta),

and the CNN (blue) networks (see Methods and supplementary text for details). The red segments indicate the position of division events.

D) Left: histogram of division times representing ground truth data and predictions using different processing pipelines). The p-value indicates the results of a ranksum test comparing the predictions to the ground truth for the different pipeline variants. The total number of division events annotated in the ground truth or detected by the networks are indicated in the legend. Right: Scatter plot in log scale representing the correlation between ground-truth-calculated division times and those predicted by the CNN/LSTM network. R^2 represents the coefficient of correlation between the two datasets. Accuracy and recall are defined in the supplementary text.

E) Left: cumulative distribution showing the survival of cells as a function of the number of divisions (N=50 cells). The numbers in the legend indicate the median replicative lifespans. The p-value indicates the results from a statistical ranksum test. Right: Scatter plot showing the correlation of the replicative lifespans of individual cells obtained from the ground truth with that predicted by the CNN/LSTM architecture (N=50). Inset: same as the main plot, but for the CNN and CNN with post-processing pipelines. R^2 indicates the coefficient of correlation between the two datasets.

F) Replicative lifespans obtained using the CNN/LSTM network for longevity mutants (solid colored lines, genotype indicated). The shading represents the 95% confidence interval calculated using the Greenwood method (Pokhrel, Dyba, and Hakulinen 2008). The median RLS and the number of cells analyzed are indicated in the legend. The dashed lines with shading represent the hazard rate and its standard deviation estimated with a bootstrap test (N=100). Results from logrank tests (comparing WT and mutants distributions) are indicated on the left of the legend.

Figure 2- Deep learning-based measurement of the entry into senescence dynamics and semantic segmentation of cell and nuclear contours

- A) Left: Sketch depicting the detection of the time of slow division (SD) state. The temporal sequence of classes probabilities (ie, unbud, small, large, dead) is fed into an LSTM network that predicts the onset of cell-cycle slowdown by assigning one of the two predefined classes pre-SD or post-SD to each frame. Middle: correlogram showing the correlation between the onset of cell-cycle slowdown predicted by the LSTM network and the ground truth data, obtained as previously described (Fehrmann et al. 2013). The color code indicates the local density of the points using arbitrary units as indicated by the colorbar. Right: sample trajectories indicating the successive divisions of individual cells (red lines) along with the division times (color-coded as indicated).
- B) Left: Principles of semantic cell contours segmentation based on brightfield images ; Top and middle row) Individual brightfield images were processed by the DeeplabV3+ network that was trained to perform pixel classification using three predefined classes representing the background (black), the mother cell of interest (orange), or any other cell in the image (blue). Bottom row: overlay of brightfield images with segmented cellular contours . Middle: correlogram showing the correlation between individual cell area predicted by the segmentation pipeline and the ground truth data, obtained by manual annotation of the images. The color code indicates the local density of the points using arbitrary units. Right: same as in A), but representing the cell surface using a color-code.

- C) Left: Principles of semantic cell nuclei segmentation based on fluorescent images of cells expressing a histone-Neongreen fusion. The semantic segmentation network was trained to classify pixels between two predefined classes ('background' in black, 'nucleus' in green). Middle/Right: same as in B, but with nuclear surface.

References

- Aguilaniu, Hugo, Lena Gustafsson, Michel Rigoulet, and Thomas Nyström. 2003. "Asymmetric Inheritance of Oxidatively Damaged Proteins during Cytokinesis." *Science* 299 (5613): 1751–53.
- Chen, Liang-Chieh, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. 2018. "Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation." *arXiv [cs.CV]*. https://openaccess.thecvf.com/content_ECCV_2018/papers/Liang-Chieh_Chen_Encoder-Decoder_with_Atrous_ECCV_2018_paper.pdf.
- Crane, Matthew M., Ivan B. N. Clark, Elco Bakker, Stewart Smith, and Peter S. Swain. 2014. "A Microfluidic System for Studying Ageing and Dynamic Single-Cell Responses in Budding Yeast" 9 (6): e100042.
- Defossez, P. A., R. Prusty, M. Kaeberlein, S. J. Lin, P. Ferrigno, P. A. Silver, R. L. Keil, and L. Guarente. 1999. "Elimination of Replication Block Protein Fob1 Extends the Life Span of Yeast Mother Cells." *Molecular Cell* 3 (4): 447–55.
- Denoth-Lippuner, Annina, Thomas Julou, and Yves Barral. 2014. "Budding Yeast as a Model Organism to Study the Effects of Age." *FEMS Microbiology Reviews* 38 (2): 300–325.
- Dillin, Andrew, Daniel E. Gottschling, and Thomas Nyström. 2014. "The Good and the Bad of Being Connected: The Integrins of Aging." *Current Opinion in Cell Biology* 26 (February): 107–12.
- Fehrmann, Steffen, Camille Paoletti, Youlian Goulev, Andrei Ungureanu, Hugo Aguilaniu, and Gilles Charvin. 2013. "Aging Yeast Cells Undergo a Sharp Entry into Senescence Unrelated to the Loss of Mitochondrial Membrane Potential." *Cell Reports* 5 (6): 1589–99.
- Ghafari, Mehran, Justin Clark, Hao-Bo Guo, Ruofan Yu, Yu Sun, Weiwei Dang, and Hong Qin. 2021. "Complementary Performances of Convolutional and Capsule Neural Networks on Classifying Microfluidic Images of Dividing Yeast Cells." *PloS One* 16 (3): e0246988.
- Ghafari, Mehran, Daniel Mailman, and Hong Qin. 2021. "Application Note: μ Polar - An Interactive 2D Visualization Tool for Time-Series Microscopic Images." <https://doi.org/10.2139/ssrn.3827406>.
- He, Chong, Chuankai Zhou, and Brian K. Kennedy. 2018. "The Yeast Replicative Aging Model." *Biochimica et Biophysica Acta, Molecular Basis of Disease* 1864 (9 Pt A): 2690–96.
- He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep Residual Learning for Image Recognition." In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–78.
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. "Long Short-Term Memory." *Neural Computation* 9 (8): 1735–80.
- Hughes, Adam L., and Daniel E. Gottschling. 2012. "An Early Age Increase in Vacuolar pH Limits Mitochondrial Function and Lifespan in Yeast." *Nature* 492 (7428): 261–65.
- Janssens, Georges E., and Liesbeth M. Veenhoff. 2016. "Evidence for the Hallmarks of Human Aging in Replicatively Aging Yeast." *Microbial Cell Factories* 3 (7): 263–74.
- Jo, Myeong Chan, Wei Liu, Liang Gu, Weiwei Dang, and Lidong Qin. 2015. "High-Throughput Analysis of Yeast Replicative Aging Using a Microfluidic System." *Proceedings of the National Academy of Sciences of the United States of America* 112 (30): 9364–69.
- Lee, Sung Sik, Ima Avalos Vizcarra, Daphne H. E. W. Huberts, Luke P. Lee, and Matthias

- Heinemann. 2012. "Whole Lifespan Microscopic Observation of Budding Yeast Aging through a Microfluidic Dissection Platform." *Proceedings of the National Academy of Sciences* 109 (13): 4916–20.
- Lin, S. J., P. A. Defossez, and L. Guarente. 2000. "Requirement of NAD and SIR2 for Life-Span Extension by Calorie Restriction in *Saccharomyces Cerevisiae*." *Science* 289 (5487): 2126–28.
- Liu, Ping, Thomas Z. Young, and Murat Acar. 2015. "Yeast Replicator: A High-Throughput Multiplexed Microfluidics Platform for Automated Measurements of Single-Cell Aging." *Cell Reports* 13 (3): 634–44.
- McCormick, Mark A., Joe R. Delaney, Mitsuhiro Tsuchiya, Scott Tsuchiyama, Anna Shemorry, Sylvia Sim, Annie Chia-Zong Chou, et al. 2015. "A Comprehensive Analysis of Replicative Lifespan in 4,698 Single-Gene Deletion Strains Uncovers Conserved Mechanisms of Aging." *Cell Metabolism* 22 (5): 895–906.
- Morlot, Sandrine, Jia Song, Isabelle Léger-Silvestre, Audrey Matifas, Olivier Gadal, and Gilles Charvin. 2019. "Excessive rDNA Transcription Drives the Disruption in Nuclear Homeostasis during Entry into Senescence in Budding Yeast." *Cell Reports* 28 (2): 408–22.e4.
- Mortimer, R. K., and J. R. Johnston. 1959. "Life Span of Individual Yeast Cells." *Nature* 183 (4677): 1751–52.
- Pokhrel, Arun, Tadeusz Dyba, and Timo Hakulinen. 2008. "A Greenwood Formula for Standard Error of the Age-Standardised Relative Survival Ratio." *European Journal of Cancer* 44 (3): 441–47.
- Sinclair, D. A., and L. Guarente. 1997. "Extrachromosomal rDNA Circles--a Cause of Aging in Yeast." *Cell* 91 (7): 1033–42.
- Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. "Going Deeper with Convolutions." In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9.
- Venugopalan, Subhashini, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. "Sequence to Sequence -- Video to Text." In *2015 IEEE International Conference on Computer Vision (ICCV)*. IEEE. <https://doi.org/10.1109/iccv.2015.515>.
- Xie, Zhengwei, Yi Zhang, Ke Zou, Onn Brandman, Chunxiong Luo, Qi Ouyang, and Hao Li. 2012. "Molecular Phenotyping of Aging in Single Yeast Cells Using a Novel Microfluidic Device." *Aging Cell*, April. <http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=22498653&retmode=ref&cmd=prlinks>.

