Tracking Microbial Growth and Evolution at High-throughput

Doctoral Thesis

UNIVERSITY OF GOTHENBURG

Martin Zackrisson
Department of Chemistry and Molecular Biology
Faculty of Science
University of Gothenburg
Göteborg, Sweden 2017

Tracking Microbial Growth and Evolution at High-throughput
MARTIN ZACKRISSON

ISBN: 978-91-629-0266-7 (PDF)
ISBN: 978-91-629-0267-4 (Print)
Obtainable from http://hdl.handle.net/2077/53352

Cover image: Martin Zackrisson

Printed in Kållered, Sweden 2017
By Ineko AB

“Normal science, the activity in which most scientists inevitably spend almost all their time [...]”

Thomas Kuhn (1962)

Bortom storskalig mikrobiell tillväxt och evolution

Storskalighet är inom modern mikrobiologi inte bara ett slagord utan även ett aktivt forskningsområde. Genom att parallellisera försök kan nya sorters frågor ställas och tidigare frågeställningar nu undersökas med större noggrannhet. Om en forskare till exempel hoppas på att en sällsynt mutation ska uppstå skulle hen behöva väldigt mycket tid eller tur för att kunna hitta den med traditionella småskaliga metoder.

Mitt bidrag till forskningsområdet är en ny metod för att parallellt övervaka tillväxten i ett stort antal mikrobiella kolonier. Det i sig är visserligen inte nytt, men vi anser att kvaliteten på datan vi samlar in är högre än jämförbara tekniker samtidigt som kostnaden för systemet är relativt låg. Utvecklingen av denna plattform, Scan-o-matic, beskrivs i artikel ett och två.

En teknik är bara relevant om den används, och i artikel tre använder vi storskaligheten och mätkvaliteten för att med hjälp av ett avancerat avelssystem förstå vilka sorters interaktioner mellan olika gener som förklarar komplexa egenskaper hos jäst. I artikel fyra testar vi i mindre skala hur jäst kan selekteras för att tåla arsenit och hur stabil en sådan anpassning är efter att den uppnåtts om jästen tillåts leva utan arsenit för ett tag.

Abstract

In modern biology, large scale is not just a slogan but a very active area of research. By parallelizing trials, new kinds of questions can be asked and questions examined with greater accuracy. For example, if you hope that a rare mutation will occur, you need a lot of time or luck to find it if you do not massively parallelize the experiment. My contribution to this area is the development of a new method for monitoring growth in a large number of microbial colonies in parallel. In itself, this is not new, but we believe that the quality of the data we collect is higher than comparable technologies, while the cost of setting up the system is kept relatively low. The development of this platform, Scan-o-matic, is described in articles one and two.

However, a technique is only relevant if it is used, and in article three we are using the large scale and the quality of measurement to determine what types of interactions between genes explain complex traits in yeast growth using an advanced breeding system. In article four we test, on a smaller scale, how yeast can evolve to withstand arsenite and how stable such adaptation is after it has been achieved if the yeast is allowed to live without arsenic for a while.

1. High throughput phenotyping

The phenotype of an organism is the collection of traits that the organism has, its physical and behavioural characteristics (Churchill, 1974). Phenotyping is the act of recording such traits, and – perhaps a little confusingly – phenotyping many times just records a single trait or a small set of traits.

The throughput and quality of data from genomics have undergone explosive developments during the last decades (Koboldt *et. al.*, 2013; Mardis 2013). Phenotyping on large scale gained momentum a bit later (Houle *et. al.*, 2010; Warringer *et. al.*, 2003; Giaever *et. al.*, 2002; Winzeler *et. al.*, 1999; Costanzo *et. al.*, 2010; Kvitek *et. al.*, 2008; Bean *et. al.*, 2014; Lawless *et. al.*, 2010; Baryshnikova *et. al.*, 2010; Tong, *et. al.*, 2001; Collins *et. al.*, 2006; Narayanan *et. al.*, 2015; Hartman *et. al.*, 2015; Allen *et. al.*, 2003) and arguably the measurement quality has many times been sacrificed for throughput. Sometimes that throughput may warrant a lack of quality and the data may still be useful, but the lacking reproducibility of scientific findings in general is particularly aggravated in large-scale screening (throughput produced at the cost of number of replicates and measurement quality) and so all strides to limit the risk of interpreting noise patterns as results (Munafò *et. al.*, 2017) makes prioritizing throughput questionable. It is because of these aspects that there persists a strong interest in improving phenotyping so that it can work in tandem with genotyping to enrich our understanding of biology (Gegas *et. al.*, 2014; Lawless *et. al.*, 2010; Baryshnikova *et. al.*, 2010; Bean *et. al.*, 2014).

High throughput phenotyping is unfortunately a vague term and there is no consensus in sight on the expected volume or scalability for a platform to recognize itself as high throughput. As an example, the now aged but reliable liquid screening platform BioScreen C requires considerable manual labour to initiate 200 experiments yet is reported as high-throughput (Murakami *et. al.*, 2008). As a contrast, with the technology developed in Paper 1 (Scan-o-matic) and many other solid media techniques, about the same amount of work and time would yield 10 000 to 100 000 experiments. Throughput is also a matter of initial investment cost in equipment, maintenance and materials needed, and

ease of data-processing after the laboratory experiment has been done. Such statistics could easily be required to be reported when presenting a new methodology. Today, the discussion of the scientific value of new techniques very much pretends that scientific research exists in a post-scarcity world. It is my strong conviction that neglecting these aspects in the evaluation of the scientific contribution made by different techniques is hurting the scientific progress.

The focus of my thesis is the development and usage of a high-throughput platform that monitors colony growth of 6144 colonies in tandem with high accuracy at a low cost (papers 1 and 2). A colony is a blob of cells that propagate without any general movement, and in this case the colony rests on an agar surface. In principle two types of experiments are common, either the use of a strain collection (Giaever *et. al.*, 2002; Sopko *et. al.*, 2006; Li *et. al.*, 2011; Brachmann *et. al.*, 1998) to study how different genetic perturbations respond to an environment (Warringer *et. al.*, 2003; Bloom *et. al.*, 2013; Costanzo *et. al.*, 2010) (Paper 1, Paper 3) or the experimental evolution of multiple replicates of isogenic starting conditions (Cubillos *et. al.*, 2011; Parts *et. al.*, 2011) (Paper 4). The technique developed in Paper 1 and 2 was used in Paper 3 but has also been used in (Märtens *et. al.*, 2016; Yue *et. al.*, 2017; Vazquez *et. al.*, 2016) and more manuscripts in the making.

2. Genotypes and Phenotypes

In microbiology, if a mutation in a gene causes a cell not to propagate, i.e. to not form a colony, we say that the mutation was lethal. We only say this if cells without the mutation did indeed form colonies. It may seem trivial, but the point here is that when we are discussing genotype to phenotype connections, we discuss how mutations cause *changes* in phenotypes. We are doing specific comparisons with controls. Another point to make is that growth not only captures birth rates, but also encompasses death – it is the birth-rate contrasted to the death-rate (Sibly & Hone, 2002). If a mutation has an effect on the viability of the cell it will have consequences for death-rates. If a mutation changes the efficiency of nutrient uptake, metabolite conversions or waste disposal it will alter the energy budget of the cell and hence what resources may

be diverted to producing progeny (Barral *et. al.*, 1995; Warner, 1999). It is at least theoretically possible that all positive consequences are precisely and exactly countered by negative consequences. Though unlikely, it is important to realize that this is not the lack of effect. It is the lack of observable effect.

Consequences is a rather vague term, and when we are talking about microbial growth, it is really effects on fitness that we discuss. Fitness is the ability to propagate one's genome to future generations (Orr, 2009). Often it is described as a relative fitness, the ability to propagate one's genome in competition with all other genomes in the population. With a mitotically dividing unicellular microbe, the fitness can be broken down to the resilience of the cell, how good it is at surviving, for how many generations of cell divisions, and the rate of such divisions. It is also how these properties are inherited in the daughters. If a cell cheats at the cost of viability or fertility of its daughters or grand-daughters, it doesn't matter much that the first cell was much more efficient in budding off new cells.

I would argue, to paraphrase John Donne, that no gene is an island, and even that the idea of genes working in signalling pathways is a much too reductionist view to be beneficial given our current level of understanding about how the cell works. It was a worthwhile perspective when most of the inner workings of the cell were unknown, but perhaps not so much when we have reached a basic understanding of most of its fundamental processes. Instead, genes and the gene products, if considered as an information system is a highly connected network (Yu *et. al.*, 2008; Costanzo *et. al.*, 2010; Hartman *et. al.*, 2015). The roughly 6000 genes in S. cerevisiae have more than 90 000 unique physical interactions and more than 400 000 unique genetic interactions documented to date (Tyers, 2017). Therefore, even if a mutation has no direct effect on cell cycle progression, senescence, fertility, cellular integrity, protein folding, genome integrity and so on, the chance that it has no indirect effect on any of these processes or that the effects of positive and negative types perfectly cancel each other out is very slim. For example, when studying genes previously reported to have no effect, they were generally found to have fitness effects. Just very small effects that normally would be discarded by standard methods (Thatcher *et. al.*, 1998). Similar argument about the interconnectedness has been presented in the

context of disease (Boyle *et. al.*, 2017). This does not imply that the expectancy is for large effects all around, rather the opposite: the interconnectedness makes the system more robust (Hartman *et. al.*, 2015; Hartman *et. al.*, 2001). This was indicated by the lack of correlation between essential genes and the connectedness of these in the gene network (Yu *et. al.*, 2008). If this is true, then the typical null hypothesis, that a genetic perturbation would cause no change to a specific phenotype, may in a strict sense never be true. Instead, it is probably closer to the truth that all genetic perturbations affect all phenotypes (Fisher, 1930).

If the structure of the cell makes the integrated fitness effect of most genetic perturbations small, it doesn't mean they are insignificant, or that they are irrelevant. Neither does it mean that the minute effect isn't reliable, which is how they were found in (Thatcher *et. al.*, 1998). Small effects are enough to drive evolution as long as they are reliable and enough time is allowed to pass (Fisher, 1930), though they may not be commonly responsible for speciation (Hallam, 1978). This is a strong argument for high quality in phenotyping, which has been the focus of paper 1, especially when doing experimental evolution as in paper 4.

The whole argument can be reduced to this: cells are generally robust. Had most genetic changes had large effects, then life would have been very chaotic and brittle and it would have been very unlikely to survive for 3 billion years. I would say it is high time for science to model cells as complex, highly interconnected and robust too. Evidence from *Escherichia coli* on synthetic lethal interactions indicate the same: there are extensive redundancies and interconnectedness in contrast to the reductionistic view of the classical pathways (Côté *et. al.*, 2016). That study used death as observed phenotype, which is arguably a rather dramatic growth defect, so their results should really be considered to capture only a small fraction of the full interconnectedness and level of redundancies in the cell.

So far the discussion has been on genes that we stipulated had effects, but are there mutations that don't have an effect? Mutations outside of genes may affect the binding properties of enhancers, repressors and transcription factors

(Pennisi, 2012). Genetic changes inside genes may be synonymous, i.e. may not change the amino-acid of the translated protein, or non-synonymous and cause such a change. In the latter case it will necessarily cause change in 3D properties of the protein because no two amino acids have exactly the same geometry. It can also change the chemical properties, i.e. changing the amino acid from hydrophilic to hydrophobic. As an example, even synonymous changes have effects because the abundance of the different tRNA that translate to the same amino acid is never identical and because the mRNA may be targeted differently as it will change the structure of the RNA, the accessibility for expression and so on (Gartner *et. al.*, 2013).

The size of the genome also has an obvious relation to the cost of synthesizing a genome during the S-phase of the cell cycle as larger genomes require more building blocks to be copied. In other words, many types of mutations can be concluded to have phenotypes without even considering the possible cellular effects of the mutations.

3. Growth

Growth can refer to two different processes that are in part related, but not always. First, growth of an individual cell in size is the change of its volume. In yeast, around 500 genes have been linked to cells becoming abnormally small or large (Jorgensen *et. al.*, 2002). It is also a consequence of traversing the cell cycle, where the cellular growth may regulate progression through the cycle and hence population growth (Turner *et. al.*, 2012). In this latter case, cellular growth is related to the second meaning of growth, population growth. This is the process that increases the number of cells in the population rather than an increase in size of individual cells. The increase in number of cells in the population is typically the sought property in high throughput growth assays (Levy *et. al.*, 2012; Warringer *et. al.*, 2003; Giaever *et. al.*, 2002; Winzeler *et. al.*, 1999; Costanzo *et. al.*, 2010; Lawless *et. al.*, 2010; Collins *et. al.*, 2006; Banks *et. al.*, 2012). If cells are not directly and individually counted, the cellular growth can confound the measurements if changes in cell sizes are pronounced and non-random throughout the experiment.

However, we should also remember that growth is but one class of all the phenotypes that can be studied. Examples of a different kinds of phenotypic observation done with high throughput are metabolic suppression (Zlitni *et. al.*, 2013), gene expression (Nagalakshmi *et. al.*, 2008), and external metabolites (Allen *et. al.*, 2003) to mention three.

3.1 Cellular Perspective

Because the focus of my PhD is population growth, the increase in number of cells in the population, the cellular processes that affect cell sizes are mainly seen as confounding noise and bias. In liquid, there's an intricate correlation between size of individuals, their concentration and the optical density reported (Stevenson *et. al.*, 2016). However, in general, the fact that individual cells vary in size causes few measuring artifacts as long as the number of cells is not so small that the contribution of each cell is substantial and as long as the cells don't change size in a synchronized fashion.

3.1.1. Cell size by cell-cycle progression

The eukaryotic mitotic cell-cycle is the progression of events that a cell passes through and that culminates in mitosis, when one cell becomes two. Very briefly and as and overview, it is divided into the phases and progression G1 -> S -> G2 -> M (see figure 1) (Hartwell & Weinert, 1989). The G1 and G2 are gap-phases in which the cells prepares for and ensures it is ready for the coming phase. The other two phases are the S-phase, when new copies of the DNA are primarily made (Alabert & Groth, 2012), and the M-phase when the cell undergoes mitosis and becomes two cells (Nurse, 1994). During the G1 phase the cell becomes larger in preparation for copying the genome and eventually splitting to become two cells (Di Talia *et. al.*, 2007). Opposing this, in the M-phase when the daugher is budded off, the daughter will have taken part of the mother's cytosol and as a direct effect of this, the mother cell will have become smaller. Further, the daughter birth sizes account for part of the variability of the duration of their first G1-phase as they increase in size (Di Talia *et. al.*, 2007).

Figure 1: Eukaryotic Mitotic Cell-Cycle

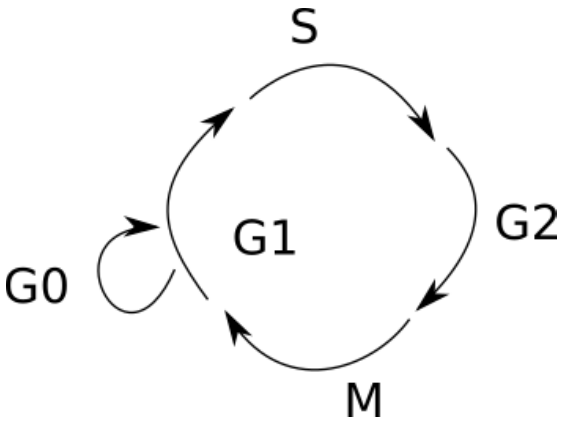

The eukaryotic mitotic cell cycle with the two gap-phases G1 and G2, the DNA synthesis phase S and the cell division mitotic phase M. G0 is the quiescent state when cells have exited the cell-cycle.

This very short summary of the cell-cycle illustrates that cells change in size as they progress through the cell-cycle. For this to be of importance to the measurement of colony growth, the distribution of where cells are in their cell-cycles must be structured. One case for this is if the properties of the pre-culture causes cells to enter cell-cycle arrest and exit G1 to G0 due to nutrient depletion (Barral *et. al.*, 1995). Most cells transferred to the new culture will then start with a bias to be in G1 even after taking into account the variation in individual lag-time. Though the resting G0 state may seem like a parenthesis, it is expected to be account for most biomass on earth (Gray *et. al.*, 2004).

3.1.2. Age

Age can be measured in chronological time but an alternative way is counting the number of cell-cycles a cell has gone through. Generally here I discuss the latter, reproductive age. Yeast cells do not only vary within the cell-cycle. They also vary with age over a sequence of cell-cycles. Cells increase in relative sizes up to third or their fourth cell-cycle (Mortimer & Johnston, 1958; Levy *et. al.*, 2012). They can reach at least the age of 50 divisions (Carter & Jagadish, 1978). Each division leaves a bud-scar, which it is reasonable to assume alters the optical properties of the cell. In a population with mixed ages one can assume that these processes negate each other, however in batch culture there is an

inherent synchronization with regard to cellular age. Inoculum is typically taken from late stationary phase cells of a preculture and the distribution of ages in the population at this stage is relatively skewed towards older cells. The founding cells of the new culture (N_0, denoting the number of individuals and subscript their generation) will give rise to roughly an equal amount of daughter cells (N_1). In this first cycle, there is a dramatic decrease in average cell age in the population thanks to all the N_1 cells being in their first cell-cycle. In the next division, both N_0 and N_1 will produce in total N_2 newborn individuals where (N_2 is approximately two times N_0), implying that 75% of the culture will be less than 2 cell-cycles old. Depending on the average age among the N_0 cells and on how long the population is growing at near exponential rates, population growth will imply rejuvenation of the average cell in the colony. This shift in cell age produces a second synchronization effect because, as noted above, cell size varies with cell age. It should be expected that the major synchronization effects occur during the first divisions as colonies enter their major, near exponential, growth phase and that the synchronization will persist as long as the cell count growth remains near exponential.

As yeast cells become old they stop dividing or do so with exceedingly low frequency (Levy *et. al.*, 2012). The cells that don't divide can persist for very long times as shown for instance by the Carlsberg lager yeast that was reanimated after more than a hundred years in a bottle (Walther *et. al.*, 2014). The resting cells that don't contribute to increasing the population size will act as a sort of dead-weight when calculating the doubling times and to lesser extent the yields of the population. Because non-dividing cells will tend to be older cells, they are expected to be more abundant when experiments are started as well as during late stationary phase.

3.1.3. Response to environment

Cells can change in size temporarily due to shifts in the environment. Osmolarity changes in the environment causes changes in cell size until the cell has managed to balance this by production of compatible solutes, but for yeasts, plants and other cell types that have cell walls, the cell size more often remains constant and the pressure on the wall is instead altered. This does not negate

compensatory measures by the cell to rebalance the pressure, nor the possibility of these having optical effects on the cells. Notably, sensing turgor pressure causes growth arrest (Warner, 1999). In the laboratory setting, changes in environmental properties are most dramatic at the start of the experiment when a small inoculum is transferred from the pre-culture to the experimental culture. In particular this is true if the pre-culture lacks a stressor that is present in the experimental medium, but even if both were created equal, the population in the pre-culture will have had time to modify the pre-culture so that there will be a substantial difference between the pre-culture and experimental conditions at the time of cell transfer (Allen *et. al.*, 2003). These effects may be less drastic in chemostats, especially if recording of data is initiated some time after a culture is introduced to the chemostat.

Because inoculation isn't a process, but a near instant act, all founding cells of a new culture will be perfectly synchronized with regard to the chronological time spent in the new medium. This implies that processes relating to response to the new medium will also be synchronized.

3.1.4. Time before first division

When seeding a new batch, i.e. when cells are taken from a stationary phase preculture and deposited on a nutrient surface, there is an initial lag phase (Gray *et. al.*, 2004; Dens *et. al.*, 2005). From bacteria it is known that a population of previously starved cells is heterogenous with respect to how they exit the lag phase. A subpopulation emerges smoothly while another fraction of the population does not and this affords complex behaviour of the colony during early phases of growth (Kaprelyants & Kell, 1996). The delay before onset of growth also has an inverse correlation to the initial amount of cells deposited, which is neither wholly a mathematical, nor an instrumental effect. Instead the effect has biological roots, and in some cases it can be countered by including supernatant of exponentially growing cells in the medium (Kaprelyants & Kell, 1996). Further, it is common to take cells from a stationary phase pre-culture, where it can be assumed that most cells are in a quiescent G0 state and need to exit this state before growth can occur (Gray *et. al.*, 2004). Here carbon availability plays a part (Gray *et. al.*, 2004). There exists at least one mutant, *gcs1*, with specific interaction to cold temperatures where exit from G0 is not

permissible (Gray *et. al.*, 2004). Variability of fraction of quiescent cells in the pre-culture will then be a strong determinant of measured lag time in the experimental culture.

This could be taken to mean that cells need to sense other cells to grow, but this extreme interpretation is obviously not true for species like *S. cerevisiae* where single streaked cells will form new colonies. However, it does suggest that the concentrations of growth factors may alter lag-phase properties and that the effect depends on inoculum size. Growth promoting or inhibitory factors may be secreted as the growth arresting α-factor in yeast (Chang & Herskowitz, 1990) or they can be membrane bound factors, in which case the cells need to be in contact to activate membrane bound receptors. When using large inoculums in the laboratory setting, the variability introduced by social signalling between cells is probably negligible. If the inoculum is heavily diluted as suggested in the Colonyzer toolkit (Lawless *et. al.*, 2010) as a way to decrease variability between replicates, there exists the possibility of the results only being valid to that specific design because the concentrations of substances in the inoculum may have substantial interaction effects with the resulting growth. Such effects could be tested by screening the deletion collection over a range of inoculum sizes and looking for both systematic trends in the recorded phenotypes over the data series as well as investigating if there are genes that are particularly dependent on inoculum size or have distinct modes of dependence. Similar concerns with dilution of growth modifying molecules exists when the medium is continuously exchanged as in chemostats and possibly to a lesser degree in general if cells are perturbed or shaken in a liquid culture as both processes dilute the immediate surroundings of the cell.

3.2. Colonial Perspective

For the purpose of this text, *colony* is limited to mean a population of microbes growing on a medium, forming a blob of cells. While the size of the colony could refer to the volume or even area covered by the cells in an image, here I use it to refer to the number of cells in the colony. There is assumed to be no migration to or from the colony or within the colony, the cells don't move, except for slowly being pushed outwards as the inner cells of the colony multiply and later when growth is mostly limited to the leading edge of the

colony (Pipe & Grimson, 2008) pushed in more complex patterns. Within this definition, populations on agar form colonies while the experimental setup in liquid screening tend to partially be aimed at disrupting colony formation by stirring or shaking (Warringer & Blomberg, 2003; Ziv *et. al.*, 2013). The liquid screening methods try to keep the yeast evenly dispersed in the liquid column, but note that this is not a necessary feature of liquid growth: left alone, yeast will start forming colonies on the bottom or at the surface of the container (Warringer & Blomberg, 2003). Screening species that are motile, like *E. coli* will still result in colonies on agar given that the experimental duration is short and movement speed is limited compared to population growth. In this no-migration regime, the population size can only change as a result of unequal birth and death rates. If the birth rate exceeds the death rate, the colony grows. If the birth rate is equal to the death rate, it maintains its size. This is typically observed during the lag and stationary phases of growth. Finally, if the death rates exceeds the birth rates, the colony decreases in size, or experiences negative growth.

It is important to understand that the the colony growth is the aggregated effect of a large number of cellular events and there may not even exist any cell that behaves like the average cell (Carter & Jagadish, 1978), i.e. the population description doesn't really say what individual cells are doing. If the colony is growing slowly, it might be that a small subpopulation grows only slightly hindered while the remaining major subpopulation shows no growth at all (Carter & Jagadish, 1978). In the slow growing colony, the average description will not be a good description of any cell. To obtain direct information with regards to this, single cell screening methods are needed, but if socal microbial effects are of any importance, then they run an extreme risk of lacking generalizable results as they tend to trap individual cells separated from each others (Chingozha *et. al.*, 2014; Reece *et. al.*, 2016). The example segues into the assertion that the colony is necessarily heterogeneous. With little or no motility internally except being pushed around by the budding process, mothers will tend to be in closer proximity to their daughters compared to any otherwise related or non-related cells in the colony. Put another way, the cell you are touching is probably a clone or near clone to you. However, if you can sense other cells through extra-cellular cues, but you can't touch these cells, following

from the previous argument, they are likely to be less related to you. They could be considered competition. These relations should be a stable and predictable aspect of colony-forming microbial growth and it seems likely to me that processes for responding to such differences would have evolved.

The colony is also structured with regards to the environment, because some cells will be internal to the colony while others will form the outer tiers of cells in the colony. The outer cells will be exposed to the nutrient substrate if at the bottom of the colony, the air if along the top hull and both if at the leading edge of the colony. These axis mentioned form gradients across the colony of nutrient availability, toxin exposure, light intensities and so forth, making the colony not only heterogeneous with regards to descent, but also with regards to the local environment (Pipe & Grimson, 2008).

3.2.1. Contrast to liquid

As stated above, liquid screens tend to employ shaking (Warringer & Blomberg, 2003) or stirring (Ziv *et. al.*, 2013) to homogenize the distribution of cells as well as well as environmental factors like nutrients and toxins, i.e. to counteract structure. This disrupts beneficial as well as growth detrimental gradients being formed by the microbes and the geometry of the container. In effect, the local environment of the population becomes larger due to dispersion. Shaking and stirring also work against cells sticking to each other, ideally making each cell free floating. As a result the environmental variables become more homogenous over all cells compared to solid screens. The shaking of micro-titre plates will in fact not result in a perfect mixture. Instead more cells will accumulate at the bottom by the gravitational pull, top by surface tension, and walls by adhesion (Warringer & Blomberg, 2003).

Contrasting solid and liquid screens, the former often have a rather ecologically realistic setting for microbes (Pipe & Grimson, 2008) and specifically for *S. cerevisiae* considering where it is typically found in nature (oaks, fruits, soil, insects and as human pathogens to mention a few) (Liti, 2015; Hittinger, 2013). An organism modifies its surroundings, takes out nutrients and deposits waste and through these actions some specific traits that are linked to such actions increase in fitness. This type of reasoning is the foundation of niche

construction theory (Laland *et. al.*, 1999). A conclusion that can be drawn from Laland is that selection pressures for primary properties of niche construction will be different and more focused than other pressures (Laland *et. al.*, 1999). With this in mind, disrupting colony formation in liquid screens may be problematic. It should also be noted that shaking and stirring is a mechanical stressor that can be assumed to not be a constant for all experiments even within the same machine and experimental protocol; if a strain, due to genetic manipulations, lacks in cell wall structure, plasma membrane fluidity or cytoskeleton integrity, one would expect this mechanical stress to be aggravated, as for example yeast *kre6* mutants (Roemer & Bussey, 1991). The stress will thus correlate positively with the hypothesis tested and not be easily separable from the true result of the intended experimental condition.

Focusing on *S. cerevisae*, alcohol production could be said to be one of the cornerstones of its niche construction (Buser *et. al.*, 2014). The dynamics of how local gradients of ethanol concentration are formed will differ with regard to solid versus liquid media. Specifically, solid media screens allows for colony formation that makes internal cells of the colony experience a different environment than the outer tiers of the colony. As niche construction is expected to modulate growth dynamics in a feedback pattern, any process that counters the mixing of liquid screens, like improved cell to cell adhesion, will not only interfere with the free floating expectation of the measurement technique, but also on its own be capable of changing the growth properties. A short list of examples with this potential would include any mutation that would make cells stay at the surface; stick to container walls; or stay at the bottom through buoyancy, cell surface stickiness or other processes. This would be an adaptation to the experimental setting (shaking and stirring) rather than the conditions tested (medium composition). Put more generally, if a laboratory setup is not attuned to the ecology of the organism, there is a risk that the results of the experiment will be confounded with effects of the setup.

4. Measuring growth

4.1. Noise and Bias

All observations of the state of the world are fraught with measurement errors. More technically each observation, *O(x)*, of a true state of *x*, *T(x)*, comes with a measurement error *E(x)*:

$$O(x) = T(x) + E(x)$$

The error function is unknown, but can be further separated into its two component functions: the random noise $E_N(x)$ and the systematic bias $E_B(x)$:

$$O(x) = T(x) + E_N(x) + E_B(x)$$

While the exact knowledge of both error functions inner descriptions is generally inaccessible, i.e. the measurement noise for the next observation of *x* can't be known beforehand, the output of each function can be analyzed given assumptions about *T(x)*. When observing the number of cells in a colony in stationary phase as an example, *T(x)* should be constant. In other words, given that repeated measurements of a colony over an extended period of time doesn't show any growth nor decline trends, the variations around the average observed colony size describe the distribution of $E_N(x)$ for that specific population size. The average noise magnitude does not have to be constant, nor does it have to be constant as a proportion to the population size. This estimation of the noise is only true given that the assumptions hold and as long as *T(x)*, $E_B(x)$ and $E_N(x)$ were actually constant for the period. If the culture has an extended lag phase before growth as well as a stationary phase after growth, properties of $E_N(x)$ can be investigated in relation to different values of *T(x)*.

The impact of noise can be reduced in *O(x)* by repeating experiments or by changing experimental setup (Brideau *et. al.*, 2003). In paper 1 both these have been applied, but the application of the latter isn't obvious to the reader because of how scientific papers typically don't characterize the exploratory journey in

developing new methodologies, but rather only the end product. The introduction of fixtures, transmissive calibration targets in each image, and precise control over power supply to the scanners in Scan-o-matic are examples of changing the setup to reduce noise and bias. Repeating experiments is helped with pinning robots as the cost of making more than one copy of the experiment becomes neglectable. Also worth noting: in a growth curve consecutive measurements are partial repeats of each other as long as sampling frequency is high. Each measurement is lending support to the next.

Bias on plate screening has been reported to behave as continuous gradients over rows as well as alternating with every second row (Brideau *et. al.*, 2003). We showed in paper 1 with examples of positional bias on colony growth dynamics on a shared nutrient agar, that the bias is more complex and dynamic than is captured by such approximations. Bias also exists between experimental runs as batch effects. Imagine for example that in casting two agar plates one got 5% more than the other. This will result in one plate being slightly thicker than the other and among other things there will be a batch effect on the dynamics of nutrient depletion when the colonies on the agar surface begin to grow.

The good part about bias is that it is systematic, which means it can be estimated and removed from the observations (Brideau *et. al.*, 2003; Baryshnikova *et. al.*, 2010), but the bad part about bias is also that it is systematic and because its removal will never be perfect, the remaining bias causes serious problems for statistical treatment of the results. So while normalization of spatial bias and batch effects are important tools, the most important tools are standardization and platform design, because these can limit the exposure to bias in the first place.

Evaluating attempts to correct for positional bias on test data, typically all positions of the plate are kept isogenic. This incorporates the bold assumption that there's no interaction between the growth properties of the colonies and the bias. If the properties of the outcome of random noise have been predetermined, the soundness of the bias estimator can be seen as the degree of correlation between the direct observations and the bias estimator. However, I would argue

that while correlation should be required to be high, too high correlation can also be a symptom of the bias estimator being too flexible and overfitting the data, thereby also compensating for the random noise in the data. In terms of an experimental setting, such a flexible bias estimator would not only remove the noise but also the otherwise observed true effects along with the bias. While the correlation is an important statistics to evaluate the bias removal, it seems to me that science tends to focus on picking out the extremes and so the more important question is how the outlier effects of bias removal behaves. In other words, if the normalization removes bias well in general, except in a few rare cases when it exaggerates bias, the end result may be that the normalization causes more erroneous findings to be reported than if the bias was left in the data.

Two primary methods of positional normalizing exist. Either through the use of the experiments themselves (Collins *et. al.*, 2006; Baryshnikova *et. al.*, 2010) or through the use of controls e.g. as in paper 1. Both methods face similar challenges with regard to the actual dynamics of the bias on the experiments. The controls can't be placed on top of the experiments, so interpolation or extrapolation is needed to estimate the bias at the experiment position. The more controls you have, the closer they can be to the experiments and hence be a better basis for the estimates. However, using more controls means taking up space that could have been used for experiments. There's also the issue that the controls may have a particular interaction with the bias factors that is not shared with the experiments and may therefore misrepresent the bias on the experiments. As an example, maybe the distribution of lag effects on the controls is non-representative of the experiments because the controls had been in the fridge for three months prior to the experiment. Using experiments, the issues are similar. Because each experiment will contain both the signal from the true result and from the bias, a rather drastic smoothing is needed. If bias is assumed as 100% local, you would need to subtract the entire value of the experiment from the observation and hence all measurements would be zero. Instead different average constructs are used. This may be fine if the dynamics of the bias are slow. If not, as we noted in paper 1, heavy smoothing is at risk of locally aggravating bias in a way prone to generate false positives and negatives rather than improving the results.

Perhaps the better solution is to use a low-pass filter on the experiments to guide the magnitude and variability of the bias estimator while using control positions to construct the actual bias estimator. This also allows for correlating the two and warning if the correlation is unexpectedly low. A similar principle was used in paper 1 when considering normalization by initial value, but the procedure is intricate and may reduce the quality of the data rather than improve them.

4.2. Sources of noise and bias

In high throughput screening, the systems need to be fairly complex to facilitate the throughput. But for every component or tier of complexity added to an experimental setup, another source of noise and bias is also introduced. Even with components whose sole purpose is to regulate the noise and bias in other components, this is true. For example, in the Scan-o-matic setup, power to the scanner is tightly regulated by a network connected power manager, in order to decrease bias resulting from the scanner lamp remaining lit after the scan near the parking position of the lamp and to decrease variance in sensor properties over time. However, while the power manager does exactly this, it requires the software to connect to the web-interface of the power manager to invoke this control. This means introduced noise from the response times of the router and from the power manager's web-interface itself. Further, it introduces two critical points of failure should either the power manager or the router not respond and because this correlates with the load on them which in turn correlates with the number of scanners connected to the same computer, there will be biases relating to the number of parallel scanners running on a computer.

The obvious first step is to look for robust technology that has little noise and bias to start with, but there are more things to consider. Measurement equipment typically has a range for which their resolution and accuracy is optimized and this range should be assumed to have both a lower and an upper bound where their fidelity drop off. The experimental design should be optimized to keep experiments within these ranges, or at least not to exceed them much. Another theoretical option is redundancies, but this is hard to achieve; if there are signs of malfunction, employ measures to regain nominal function. If nothing else, high-throughput methods need to detect and warn about issues. For example, if

a measurement device is deteriorating in quality in a high-throughput screening setup, it can take time and cost a lot of lost research before anyone notices. The high-throughput makes non-critical issues that never the less can cause substantial bias hard to detect for humans which mean that automatic detection and notifications about issues is needed.

Measurements are many times not continuous, but digitalized in discrete ways. An image contains pixels that represent an average of observed intensities in the corresponding area in the world. The average intensity of the pixel is also digitized, typically into 256 intensity categories. This is a great source of batch effects between images, because most pixels will fall somewhere between two neighbouring intensities and minute shifts in lighting or sensor properties will sometimes shift bulks of pixels from one digitized value to the next. This is especially problematic when the imaged scene is largely homogeneous in intensities such as a well mixed liquid medium or an agar surface. If lighting isn't strictly standardized the digitalization process can result in minute differences being systematically exaggerated due to rounding. Increasing the digitization precision by changing the number of pixel intensity values, also know as depth, from 8-bit/256 values to 16-bit/65536 values greatly diminishes these problems. Equally decreasing the area of each pixel by increasing the dots per inch (DPI) of the image is generally beneficial, though for scanners this typically results in longer image acquisition times which introduces two issues: prolonged light exposure and systematic difference in time of measurement in relation to where on the image the measurement was made.

While imaging colonies on a solid growth medium, aspects of the experimental setup will inadvertently affect the recorded intensities of the colony: agar coloration and light properties; the properties of the plastic containers such as casting imperfections and scratches; and scanner variability in image acquisition. In Scan-o-matic, we subtract the mean of the inter-quartile range (IQR) of the area surrounding the colony as a way to compensate for such variations. The particular mean type is important in combination with digitalization and IQR mean was used for its combination of stability and sensitivity (Mangat *et. al.*, 2014).

4.3. Accuracy and Precision

Precision is a description of the level of noise in a measurement system. However, the observed precision in a sample from such a system says very little about the truth of the measurements as it neglects bias. Accuracy on the other hand describes how close to the truth the output of the measurement system is. Having more information, either by increasing repetitions or having more wealth of information in the primary data, generally increases the precision, but not necessarily the accuracy.

Removing bias is a combination of platform design as to not include bias-sources in the first place and normalization to remove observed systematic trends. As an example, the coloration and transparency of the growth medium can bias population size estimates. In liquid culture it is highly problematic to truly remove this contribution to the measurement thanks to the dynamic properties of the growth medium. Typically the medium without inoculum is used in the hope that the cells and their growth do not affect the properties of the growth medium. In the case of a medium with Cu-ions, the coloration of the medium is dependent on pH, something that typically changes where yeast growth. For solid screens, it isn't possible to know exactly the properties of the medium below the colony, but it is possible to estimate its contribution from the colony surroundings.

4.4. Standardization

The controls must be kept as similar to the experiments as is possible. This is to avoid unaccounted for systematic differences between controls and experiments that may cause spatial bias or batch effects. The pre-culture setup of both should be shared and/or common as far as possible, but the choice of control for capturing biases should also be made to minimize the differences between the experiments and the control. This is important because each difference is a risk of a gene-environment interaction effect difference that may cause the controls to have a different bias, both in magnitudes and distribution, compared to the experiments.

There are problems inherent in standardization. While it increases the quality of the measurements the same lack of incorporated noise from inadvertently varying the experimental conditions increases the risk that the findings are private to the specific settings of the experiment. In other words, standardization puts the generality of the findings at risk if there is no supporting evidence for them.

4.5. Normalization

The use of controls has been critiqued on the basis that controls tend to be few and not spatially distributed to capture positional bias, that the bias may act differently on controls, that variability among controls is neglected, and because of the effect of potential outliers among the controls (Brideau *et. al.*, 2003). Several of the critiques stem from an assumption that controls are always few and only used for batch normalization. None of these assumptions are necessarily true. The issue that the control may not exhibit the same or at least not the same magnitude of bias as some or all of the experiments, I find the most valid. In part this can be countered by standardization of pre-cultures, but if there are gene-environment interaction effects specific to the control phenotypes that cause their bias to be distinct from the experiment bias, the normalization by controls will be problematic, possibly detrimental. The most obvious method of avoiding this, discussed above, is through experimental design and validating the bias estimator. The experimental evolution setup in paper 4 implies experiments are adapting to their adverse environment, but to be able to compare between different times during the evolution experiment, the controls have to remain the same. If the controls struggle to grow, the magnitude and possibly the positional dynamics of the bias could be expected to be dramatic, while the magnitude of the same bias if measured by the experiments could be expected to be smaller and with slower positional dynamics. Here is a gap in our understanding, and it would be interesting to see the results of normalization by the same reference strain of uniform plates that contain only non-adapted, semi-adapted and fully adapted experiments. It should be noted that this issue is probably worse for control based normalization, but also exists for experiment based normalization as the experiments are not expected to be homogenous with this regard, the described phenomena will confound bias estimates from heterogenous experiments.

The author of aforementioned critique of control based normalization also neglects to discuss potential problems with using experimental results as basis of normalization, which we'll tend to in the next paragraph:

The first issue with using experimental positions for normalization of batch or positional effects is that they are typically never randomized in high throughput screening. Instead, they tend to reflect chromosomal positions of gene deletions or similar structuring bias (Winzeler *et. al.*, 1999). Because of gene duplication and selection pressure to keep genes with epistatic interactions in proximity of each other as to not break up alleles in meiosis, proximate genes can be expected to show correlated phenotypes to a larger extent than randomly selected genes (Spellman & Rubin, 2002; Lercher *et. al.*, 2003; Petkow *et. al.*, 2005). Further, replicates are often placed next to each other as a constraint of the high-throughput and the robotics used. Both cause expectancies of mean shifts in phenotypes based on positions that are true results and not bias. This indicates that the use of experiments to normalize experiments should produce a wealth of false negatives.

Another challenge for normalization is the exact calculation of the normalization procedure and to determine the validity of such. For the following discussion, we assume the bias has been determined by interpolation from controls, but the same reasoning generally applies to bias estimated from the experiments. The question that remains is how to remove the bias from the experimental observations.

If bias and batch effect errors are growth dynamics factors, i.e. chiefly multiplicative, bias is removed by division:

$$R_{normalized} = \frac{T_{experiment}B_{experiment}}{T_{control}B_{control}} \approx \frac{T_{experiment}}{T_{control}},$$

or if considered on a log scale

$$log(R_{normalized}) \approx log(T_{experiment}) - log(T_{control}).$$

Where R signifies the relative observation, B is the bias and T is the true value free from bias. Random measurement noise is disregarded in these equations for clarity.

On the other hand, if bias is chiefly additive, such calculation is wrong and instead the direct phenotypic difference will remove the errors:

$$R_{normalized} = (T_{experiment} + B_{experiment}) - (T_{control} + B_{control}) \approx T_{experiment} - T_{control}$$

.

It is hard to rationalize why either of the two variants must be necessarily true for all types of phenotypes, not even why the the bias of a specific phenotype should be driven by either process exclusively in all environments or gene-environment combinations. These two examples are by no means exhaustive of the modes by which bias can affect results and here it would be interesting to see investigations into whether the use of correlation between positional bias estimates and observed experimental values can be helpful in quantifying the normalizability of the data.

I opted for using the $\log_2$ difference between the experiment and an interpolated value for the hypothetical reference if such reference had been placed at the exact same position as the experiment in Scan-o-matic. This ensured almost direct comparability, with some minor differences, to the methods utilized for BioScreen C data in our lab (e.g. Warringer *et. al.,* 2003) while countering several of the major critiques of using controls to normalize data (Brideau *et. al.*, 2003). Using log-scale difference has the advantage that it isn't as sensitive to errors in estimating the controls when the control value is small compared to using a ratio (Brideau *et. al.*, 2003). However one point that it fails to address, which other scores such as the Z-scores do, is to adjust the confidence in the value to some variability of the data (Brideau *et. al.*, 2003). It would be of interest to extend our platform's ability to normalize using different normalization methods.

4.6. Randomization

Randomization of experiments and their replicates as a method for countering bias can be characterized as incorporating the bias as noise in the measurements, i.e. countering the systematic aspect of bias but losing in the observed measurement precision. In that interpretation, this is a method geared towards producing elevated levels of false negatives in frequentist tests (Malo *et. al.*, 2006). But to make bad things worse, it rests upon a misconception about the output of random processes. The assumption about randomness is that it is homogenous while the output of random processes only become homogenous over a large number of repetitions, something humans often fail to realize (Poláček, 2017). This implies that using randomization on high throughput screening, with all the labour cost it applies, also means that one should expect a few experimental repetitions to be strongly correlated with the bias after randomization of positions if the number of repetitions is low. These would be producing false results that will be prone to pass statistical tests.

My intuition is that the beneficial effects of randomization outweigh the detrimental effects only when the number of repetitions of each experiment is large. In my work, I never got to investigate this.

4.7. Information content

The wealth of information in the data from which a growth phenotype is determined will affect the noise levels in the phenotypic estimates. There are two distinct aspects of information content relevant to high throughput screening: repetition/sampling frequency and the amount of information in each observation.

Repetition is a property of experimental design where the same hypothesis, or in the case of colony growth phenotyping, the same strain, is tested several times. Sampling frequency is also a form of partial repetition of the observations made about one and the same colony. The more frequently the colony is measured, the less time has elapsed between observations and thus the state of the colony is more similar in the two observations – the number of births and deaths that

happened between the two is small. Therefore, measurements in close temporal proximity are partial repetitions, and this is a cornerstone in the justification of applying growth curve smoothing. Given some basic assumptions about the true values of a colony size, deviations from local trends can be used to estimate the nature and magnitude of the noise function. This has been discussed above. It may sound like the higher the sampling frequency is the better, and in part this is true. However, measuring may at times be invasive. This is true for imaging, which with most of currently available technology needs elevated light intensities for image quality to be high. Visible light has been reported to affect yeast (Bodvard *et. al.*, 2013; Logg *et. al.*, 2009) and it is not uncommon for light emitters to emit light in the UV spectrum too. Light may also generate heat which also affects experiments. Consequently, there needs to be a balance between sampling frequency and the detrimental light effects when imaging. It is also in part a computational and storage problem in that the higher the sampling frequency, the more data is generated. In reality it becomes a very practical balance that needs to be met where wealth of information is balanced against the negative effects and how manageable that data volume is.

It has been argued that models fit data well and that simple models are just as good as more complex models (Buchanan *et. al.*, 1997). While the choice of model and its complexity may be of little consequence to the fit, I find the evaluation data in the mentioned article too sparse to be usable in evaluating if any of the models describes the actual growth dynamics of the experiment. The fit of the models is evaluated against the size or yield of the colonies. In this respect, the models may give acceptable approximations of the colony but this says very little about how well the model represents the growth dynamics. In paper 2 we illustrate this. The proper comparison here is the fit of the first derivative of the model to the first derivative of the observed population sizes. To obtain any kind of precision in this comparison, frequent sampling of the population size is needed. This makes model-fitting superfluous. Evaluating the fit of the derivatives even when the model is created from frequently sampled data shows the representation of the growth dynamics can be catastrophic even if the fit of the population sizes indicate that the model is a near perfect representation of the data. This point that was brought forward in paper 2 pertains particularly to underfitting the growth curve due non-standard growth

dynamics of the observed colony, for instance if a colony exhibits biphasic growth (Warringer *et. al.*, 2008).

Deviations from the growth norm would probably be the most interesting phenomena to the researchers, yet at the same time the cases that the models would be particularly inapt to handle since these growth patterns will violate the assumptions of the model. It could be argued that in such cases models can be updated, but the deviations aren't necessarily known beforehand and in a large set of curves, prone to be missed. If models are made more complete to encompass multiple modes of growth, they also become more prone to overfit simpler data.

In summary, holistic growth models are problematic and the the fit of a growth model is a bad quality indicator. To allow evaluation of the model in the experimental setting, something that is needed on each individual growth curve, the sampling frequency needs to be so high that the application of growth models serves little purpose. While the initial lag and final stationary phases tend to be extended in time with slow changes in growth dynamics, which lends them to accurate estimation from a few measurements, the other parts of the curve, such as acceleration and retardation in and out of near exponential growth as well as the properties of the near exponential growth, are highly dynamic. Correct estimation on even the simplest of curves requires measurement frequencies that again question the use of growth models entirely.

4.8. Imaging

When an image is captured of the world, the sensor – be it the human retina, camera film, a digital camera, or a scanner – is virtually a 2D array of elements that detect light. As a consequence of the geometry of the sensor and any lenses, different parts of the sensor receives light from different parts of the world. The size of the sensor in combination with its resolution (how tightly packed and small the sensing units are) and the lenses, will determine the view of the world and the detail with which it is recorded. The spectral sensitivity and fidelity of the sensor is also of outmost importance for the quality of the recorded image. When the data is recorded for computers, the general rule is to place it into a 2D array representation, with each element being a pixel. The light intensity may be

recorded once per pixel or into several channels for different parts of the light spectra using different sensors or filters. The information in each channel of each pixel can be digitized into discrete categories: 2-bit (has light or not) and 8-bit (256 different intensities) are the most common. The intensity can also be represented by a floating point number, which in its representation in computers generally comes with a fixed precision too such as 8, 16, 32 or 64-bit (Sonka *et. al.*, 2014). This precision, whether float or discrete, is the depth of a image. Resolution is the number of measurements per area unit. OD-readers that only outputs a single value is an example of extreme simplicity/low resolution imaging. This is a 1x1 pixel intensity image with float data.

Because light intensity is what the sensor records, the geometry of the sensor, light source and object imaged becomes important. One major distinction is whether the light is reflected off the surface of the object – the light source is ambient lighting or point/spot lights near the camera – or if light is shone through the object in transmissive mode. The main light source is placed on the opposing side of the object relative to the sensor and light that is not absorbed, reflected or refracted is detected.

In my work, the focus has been on scanners in the transmissive mode using 8-bit depth images. Unfortunately this constitutes rather coarse rounding off of intensity measurement into 256 discrete categories. Partly that issue is reduced by the fact that the data from transmissive scanning is more accurate and can be used to calculate growth rates reliably (French *et. al.*, 2016). Contrast this with where reflective imaging in cameras are used to decide which pixels have colonies or not (i.e. an 8-bit image is reduced to a final 2-bit depth image), and the improvement with regards to information content should be obvious.

4.9. Other ways to count

In continuous cultures, one common way to count success of competing strains is by sequencing samples of the culture, in particular barcoded strains where only the barcodes are sequenced. One issue with this method is that it introduces sampling errors – only a fraction of the population is sequenced, else measuring terminates the experiment. Additionally, if the sequencing library preparations include any PCR steps, then they are sources of bias as well (Aird *et. al.*, 2011).

Theoretical modelling suggests that frequency dependence is an intrinsic property of microbial growth, which complicates competition experimental designs if frequencies of strains are not continuously monitored, i.e. it isn't sufficient to compare to initial frequencies or valid to generalize the results without addressing the possibility of frequency dependence (Manhart *et. al.*, 2016). There's also experimental evidence for frequency-dependent selection (Turner *et. al.*, 1996). If two strains are pooled and one outcompetes the other, it can neither be generalized beyond the initial total starting concentration, nor their initial ratios, unless other supporting evidence is presented.

In machines like the FACS:s, capillary systems are used to detect individual cells as they pass in front of a sensor (Julius *et. al.*, 1972). But again, this requires subsampling the population or requires very small populations which can only be measured once.

4.10. Finding the interesting few

The principal methodological purpose of large scale screening is to act as a sieve into which a large number of tests are thrown, and out of which a smaller number of interesting data is retained for further validation. So far, I have discussed principles governing the construction of robust and accurate high throughput screens, but this leaves the question of where to draw the line: which results to consider positive and which to consider negative.

Perhaps the most common method is to determine the statistical significance via a multiple hypothesis corrected p-value from t-tests. There are many issues with such methods: the threshold is arbitrary, the tests tend to be underpowered to an unknown degree, and they don't answer any interesting research question, just to name a few. The outcomes of the t-test tend to mostly depend on the variance rather than effect size (Halsey *et. al.*, 2015), so in combination with high-throughput methods typically having few replicates, and hence lower precision, the method becomes prone to judge the results based on their biases and stochastic sampling and measurement errors, rather than the observed effects. Further, the combination of stringent but arbitrary thresholds and low power implies that many interesting potential finds are discarded while not

securing much trust for the results. All these dubious properties aside, few journals and reviewers will question t-test based filtering of results and so unfortunately it remains a very appealing method.

Another method is to simply disregard the noise observation in the measurements and rank the results based on their means. The mean is much less sensitive to stochastic properties caused by low number of replicates than the variance. It also directly reflects an important research question: Which experiments did better or worse than expected? However, as with t-test based methods, the threshold for which results should be regarded as positive and which negative is arbitrary. Nevertheless, compared to t-tests, this method of selecting which results to validate is much more rational. The p-value defender may question trusting results based only on magnitude, but as noted about generality of the results and how high-throughput methods are generally used: it is important to independently validate the findings. In contrast, repeating to verify a p-value is problematic as the risk of false negatives can be unexpectedly high (Goodman, 1992).

But is there nothing that can done about the process of selecting which experiments to validate to make it more rational? During the preparation of the manuscript for Scan-o-matic (paper 1), we had a novel suggestion that never made it into the final manuscript due to scope, focus and time:

If a large scale screen is done on a well studied organism like *Saccharomyces cerevisiae* and it is done by testing a large number of known mutations in a specific environment, some assumptions can be made about the expected outcome. First, it is expected that the stress to the cell in the test condition will disproportionately affect or change the function of some cellular pathways or functions compared to the average. Second, the more tightly correlated a gene is to the function of one of these pathways, the more the results from manipulated versions of the gene in the tested environment should deviate from the norm. Third, the expected norm result of a growth defect in an environment is the product of the environmental effect and the gene manipulation effect in the reference medium. These assumptions mean that the gene-by-environment interactions of the true positives should be enriched for the GO annotations

(Gene Ontology (Ashburner *et. al.*, 2000)) that reflect the mode of action for the stress invoked by the environment. There should also be a way of using mainly the effect sizes to sort the true positives from the negatives and the success of said sorting can be evaluated based on the resulting GO enrichments of any of the many tools and methods to do so (Huang *et. al.*, 2008). For this to work, there must be a wealth of trusted GO annotations about the genes of the organism. The more narrow mode of action that the stress of an environment has on the cell, the better it should work. It is expected that diffuse and general acting stresses, such as heat and to some extent osmotic, stress work less well than heavy metal ions or toxic compounds that target specific proteins in the cell.

At this point several questions remain, most prominently: *How would you do the test?* For all experiments, calculate the deviation from expectancy as well as the p-value for the probability of observing data at least as extreme as the results given that there is no difference to the expectancy. This places all genes in a two-dimensional space. It should be noted that the dimensions are not fully independent, but as pointed out before, the t-test largely depends on the variance rather than effect size. The simplest way to divide the space is to require both p-values to be less than a threshold **a** and deviations from expectancy to be larger than another threshold **m**. Testing a large number of combinations of thresholds for **a** and **m**, the GO annotations of the selected group for each combination can be evaluated and the combination of thresholds that produces the best selection of genes used. A direct way of evaluating GO annotations is simply counting the number of significantly enriched categories. Another is to sum up the total enrichment in all significantly enriched categories. Of course, the significance threshold for the enrichment is still arbitrary, but the total test much less so than the other two methods as the selection of positive result strains is guided by the information content within the positive results.

For the NaCl experiment in paper 1, the data indicated that roughly 600-1200 genes have an interaction with the environment depending on how the evaluation of the GO annotations is constructed. In figure 2, using signal to noise ratio rather than t-tests (both reflect the same basic properties in the data), in general the thresholds that performed best combined a very permissive

precision threshold (signal to noise around 0.2 and above) with a moderate **m** (a bit lower than 0.15 $\log_2$ deviation from expectancy). In other words, this suggests that the precision tests such as the t-test is helpful in sorting out the most noisy results and that there are many genes with moderate and small effects that are true, though since the magnitude of the effect is small some help from precision tests is needed to filter out false results in this lower range. Interestingly the use of the t-test is more or less inverted from its traditional use with highly stringent thresholds.

Figure 2: Information Guided Selection

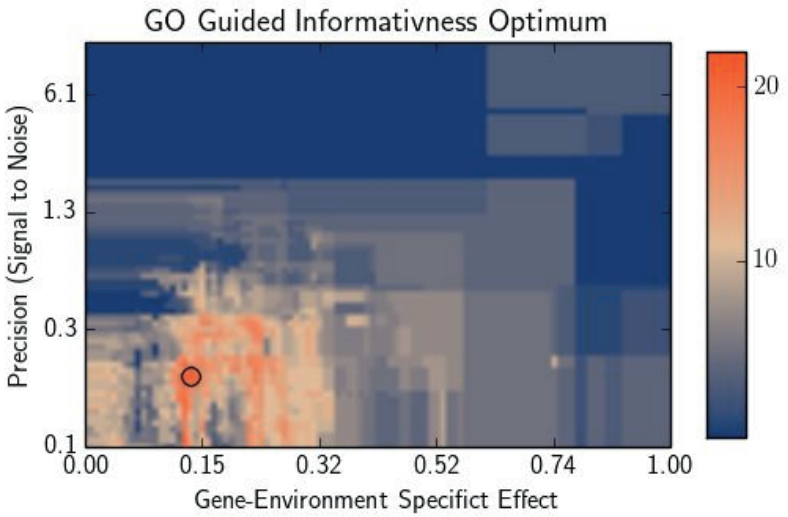

Number of enriched GO Slim terms, excluding terms with less than 10 genes, shown as color for different threshold combinations for relative effect size of the measurement (rate) on the X-axis and the quality of the measurement (signal to noise ratio) on the Y-axis. Circle indicates optimum combination.

There are some obvious issues with this method. First and foremost, it is quite probable that genes with moderate and smaller deviations from expectancy will not have been annotated to the same degree as genes with large deviations. In part this prediction rest on what tends to capture the interest of the researcher and be considered of value to the journal. It is also a consequence of that a lot of replicates (costly and time consuming) and/or a high quality experimental setup

are required in order to reliably detect small deviations. The optimization step will therefore be more restrictive towards including minor effect experiments than is actually warranted. The above NaCl example also hints at the opposite issue: if no gene truely has the expected effect in any environment, then it is impossible to enrich for these cellular processes by sub-selecting among all possible gene modifications. However even if this is true, in the current context, it would still be assumed that some genes have a positive effect while others a negative and if the annotation system correctly characterize this, the problem is of little concern. The current GO annotation system is partially flawed in this respect as many of the process descriptions follow the pattern of 'response to osmotic stress' though there are exceptions like 'negative regulation of translation in response to osmotic stress' (Ashburner, 2017). Further, if in each case all or nearly all genes can be reliably shown to deviate from the expectancy, however little, the related annotations become non-informative. This is taking it to an extreme, but if many categories start to include many or most genes the value of the annotation also drops.

It would be of great interest to see a standardized set of gene perturbations, such as the yeast deletion collection, be tested over a range of different types of environments to see if the prediction about specificity of the stress is reflected in the number of genes selected by the suggested method and to see if there are consistencies as to which combination of test parameters perform best.

4.11. Parametrization of growth

Measuring the size of population over time produces a growth curve. However, direct comparisons of such curves doesn't give easily understandable or comparable biological data. Normalizing for batch effects and spatial biases on growth curves is also tricky. Instead, in quantitative analysis of growth, these curves are parameterized, either directly on the data or by applying growth models as discussed above and extracting parameters from these models.

In general, any parameterisation that can reliably distinguish differences in growth dynamics between different strains is of value as it can be used to group strains based on growth behavior. Some of the used parameters go further and

have quite clear biological interpretations as well as arguable direct effect links to the fitness of the strains.

4.11.1. Lag

Observations about lag, the period of no or decreased growth before the main growth sets in, was first discovered by observing a correlation between growth rate estimates (yield divided by experiment duration) and the duration of the experiments. Max Müller drew attention to the inconsistency in 1895, and for bacteria Penfold collected various reasons for effects on lag, including number of founding cells, volume of inoculum (if volume is small), age of founding cells, and composition of growth medium even if cells are pre-cultured in the same medium (Penfold, 1914).

One could argue that from a cellular perspective, the lag time should reflect the time it takes before the cell fully enters the cell cycle or alternatively the time it takes until the first division of the cell. This, the time it takes for the cell to reach its first division, is the sum of the time that the cells need to adapt to a new environment and the time it takes to complete the first cell-cycle. Neither of these interpretations is readily parametrized from a population growth curve even if we consider an isogenic population and the average times of each of these phenomena. The lag time of the colony reflects neither of these events but a time in-between the two. Instead measurements of lag such as those used by Warringer *et. al.* (2003) will report times between the average time of adaptation and the average time of the first division (Buchanan *et. al.*, 1997). The latter limit also in part contains the growth rate. Therefore, not only will the measurement in part be calculated from the growth rate (Warringer & Blomberg, 2003), it will also contain part of that growth rate. If in truth the process of switching over from a resting stage is distinct from the first progression through the cell cycle it could be possible to correct the direct colony lag measurements with

$$L_c = L - f(R)$$

Where $\boldsymbol{L}$ is the primary measurement of lag and $\boldsymbol{f}$ is a corrective function that takes the growth rate $\boldsymbol{R}$ as parameter. A suitable function $\boldsymbol{f}$ would be determined

by fitting the data $\boldsymbol{L_c}$, $\boldsymbol{L}$ and $\boldsymbol{R}$ from simulations. There are, however, some problematic ifs: if there exists a suitable function $\boldsymbol{f}$ that reliably corrects the lag over a large range of growth parameter settings; if the simulations can be relied on to be biologically relevant – that is to say, will all adaptation times of all strain and environment combinations possible be distinctly separated from the first cell cycle? The finding of a suitable function is aggravated by the fact that newborn cells spend more time in G1 to regain normal cell size. Furthermore, if the culture is slow growing, as discussed above, the average growth rate may not be reflecting any subpopulation in the colony. With these issues, finding a suitable function may be impossible.

4.11.2. Rate

Rate is typically the slope of the growth curve at the steepest part of the curve. If the curve is on $\log_2$ scale, one divided by the rate is the population doubling time. Rate is directly linked to fitness given all other things equal: cells with faster progression through the cell-cycle will proliferate faster and gain in relative abundance. However, in a near exponential growth mode, a shortened cell cycle may incur future fitness defects that will have to be compensated for during the stationary phase. If nothing else as a consequence of resource allocation from maintenance towards reproduction (Kirkwood, 2005).

One concern is if the exponential growth is only a theoretical concept and not empirically supported. On a growth-curve with the y-axis on log-scale, an extended sloped linear segment of the growth curve would reflect exponential growth. The corresponding derivative should have a plateau rather than a hill or peak. This is not true in the experimental data that I've seen, however there are extended periods of growth for which the colony is experiencing near-exponential growth. This can be seen by extending the tangent of the growth curve from the steepest part of the curve; observations remain proximate to this tangent, in some cases over several hours. So it could be said that exponential growth is neither entirely a true nor a completely false description of the growth optimum reached by microbial colonies. The maximum growth rate parametrization of the growth curve will as a consequence of this tend to underestimate the true maximum and that tendency will increase with sparse data as well as with increased smoothing of the data.

For monomodal growth curves with only one growth phase, the estimation is rather straightforward. However, since there exists growth curves with more than one mode, such as the classic biphasic shift in glucose repression (Monod, 1949), defining which growth phase to compare with which is not straightforward. Still, finding the maximum growth rate is easy using the same strategies as with monomodal curves. The comparison between different experiments becomes highly problematic though if some curves are monomodal and some bimodal, or if not all curves agree on which of the two modes has the steepest part. This would imply making comparisons between biologically distinct features as if they were the same. A slightly less problematic issue is the loss of information. For multimodal growth curves, it is of interest to gain rate measurements from each mode. This is the argument for having the parametrization encompassing sufficient understanding of different types of growth curve dynamics to reliably extract only comparable data and to get rate data from all growth modes.

4.11.3. Yield

The concept of yield wants to reflect the efficiency with which cells have used the environment to produce biomass. In this perspective, it only matters to produce as many offspring as possible with as low footprint on the environment as possible. It is readily measurable as the number of cells produced throughout the experiment, the final population size minus the initial inoculum. Because the former value tends to be very large and the latter very small the precision lost by only measuring the final value can in some cases acceptable. For comparisons to be reliable, the colony size must have reached a near zero net growth stationary phase. Otherwise the results will not be comparable since they will be sensitive to the timing details of the experiment. When this will happen is difficult to know beforehand and that there will be a point in time at which all colonies on a plate are in stationary phase isn't necessarily true. If we pick apart what processes may cause the population to enter the stationary phase, it soon becomes evident that the yield is neither independent of the growth rate nor the lag.

There may be negative trade-offs between yield and growth rate. It has been stipulated that there is a trade-off between ATP synthesis rate and the amount of ATP extracted per nutrient source molecule in the medium. In bacterial studies, this was not visible when comparing distinct populations, but the cell count growth rate to yield had a negative correlation between repetitions of the same population (Novak *et. al.*, 2006). While the trade-off is likely, there are other processes that ties population growth rate to yield.

First and foremost, given all other things equal, two populations only differing in growth rate will have different accumulated footprint of their populations (integral of the growth curve) at the same yield. This implies that given that cells consume nutrients at a fixed or near fixed rate throughout the cell-cycle independent of speed, the total nutrient consumption needed by two population with different rates to reach a certain yield will differ as a consequence of there having been more cells alive for longer times in one compared to the other. The slower population will have extracted more nutrients to reach the same population size, i.e. a positive correlation between growth rate and yield. The argument doesn't need a key nutrient to be depleted entirely from the medium as causing the stationary phase phenotype, but only that the consumption rate of the population surpasses the diffusion rate from the surroundings in such a fashion that the the local concentrations are below a growth permissible threshold.

Second, if microbial growth has social components (MacLean, 2008; Crespi, 2001; West *et. al.*, 2006), for instance if the cells produce secreted growth promoting and/or inhibitory factors, the growth dynamics of the culture will affect the dynamics of the concentrations of these molecules. If the social behaviour signalling is dependent on cell-density or medium component concentration, the effect on growth yield could be substantial and correlated to the population growth dynamics. Social components could explain why yeast stops growing before nutrients are exhausted (Lafon-Lafourcade *et. al.*, 1979; Warringer *et. al.*, 2011). Independent of the explanation for not exhausting the environment, the phenomenon makes yield comparisons complicated as they are not easily transferrable to efficiency of energy metabolism.

Third, lag should be partially negatively correlated with yield based on the same argument as growth rate – different accumulated footprint of the population.

The complex entanglement with other aspects of growth is probably difficult, if not impossible, to reliably unravel. Since the effect of the dependencies qualifies as a bias rather than noise, stringency with significance tests is an irrelevant countermeasure to ensure lag and rate independent yield results. Instead, I would argue that larger deviations from expectancies must be required. Interestingly, there seem to be specific environment-dependent correlations between rate and yield (Warringer *et. al.*, 2011). Here it would be of interest to systematically study these relationships – how yield may depend on rate and lag – in order to give precise recommendations on the evaluation of yield results so as not to risk reporting the same rate finding twice: once as rate and once as yield. It would also be interesting to take these correlations apart and understand why different environments give different correlations in light of the above complications for the yield measurement.

5. Yeast

Yeast is often used as a shorthand for the baker's yeast *Saccharomyces cerevisiae*, though the term generally is used to describe a paraphyletic group of unicellular fungi that has the tendency to make their environment bubble or foam (Kurtzman *et. al.*, 2011). In the context of this dissertation it is generally used as shorthand for *S. cerevisiae*. It has a long history in human culture and may well be one of our oldest and most important domestications (Siccard & Legras, 2011). This species is of importance here because experimental parts of the included papers revolve around this organism. It should be said that though Scan-o-matic was designed around it, trials using *Escherichia coli* have been undertaken.

Some characteristics of *S. cerevisiae* are of importance with regard to the assumptions baked into the measurement platform I created in papers 1 and 2, as well as for papers 3 and 4 where it was used as experimental organism. It is unicellular, as noted before, and while there are two sexes, in the lab it mostly propagates asexually thanks to the inactivation of the HO gene (Jensen *et. al.*,

1983). This has implications of how novel mutations spread in the population, especially for the ability to acquire multiple beneficial mutations in a single genome. This is of importance in paper 4. It lacks taxis, which is helpful in screening them: the cells stay where they were born in the absence of external forces. They form colonies. It tends to form round cells (Kurtzman *et. al.*, 11), which implies that they can be approximated by small balls laying on the agar surface, an assumption used in paper 1. *S. cerevisiae* has another growth mode that it shares with many of the species in the true yeast clade: they can form pseudohyphae (Liu *et. al.*, 1996; Gimeno *et. al.*, 1992), which cause invasive growth linked to pathogenicity (Shepherd, 1988). Pseudohyphal cells become elongated and form a chain of cells linked together (Gimeno *et. al.*, 1992). In the context of my work on Scan-o-matic, any property of the studied organism that causes it to migrate is problematic as it disrupts the colony forming properties which is the fundamental assumption around which the platform is made. In particular, invasive growth would put cells inside the agar, which would give them other optic properties beside those caused by changed cell shape. As a contrast, *E. coli* has motility and accurate estimates of colony population sizes should therefore be assumed to deteriorate with increased migratory behavior.

As possibly all organisms do to some extent, *S. cerevisiae* shapes its local environment to its liking and does so strongly. One of the prominent aspects of its process to secure its local environment is to make it hostile to other organisms by producing alcohol. This also happens to be one of the main reasons for the human domestication of yeast (Libkind *et. al.*, 2011). But besides being of human economical and cultural interest, these behaviors have consequences for the screening of yeast as well as the understanding of the evolutionary forces acting on the organism. They have been discussed above with regard to liquid screens, mixing and stirring. Such forces could trigger further allocation of resources into niche construction than would otherwise be expected.

6. Toxicity and As(III)

In paper 4 we studied how resistance to arsenic can evolve, and found ultra-fast adaptations to mainly use the strategy of compound exclusion. Basically, if a compound is toxic inside the cell but never allowed to enter, it is no longer toxic. A similar but more costly strategy is to actively export the compound. Other alternatives are to metabolize the compound into less toxic compounds such as arsenite (V) to arsenite (III) conversion (Wysocki & Tamás, 2010) or changing the properties of the cell so that it can accommodate the compound better.

Arsenite is reported to causes oxidative stress in mammals. In yeasts, it is suggested that trivalent arsenical molecules can affect the redox state of the cell (Styblo *et. al.*, 1997). It can bind a range of proteins such as beta-tubilin, pyruvate kinase, and pyruvic acid dehydrogenase where it either impairs the function or leaves it relatively intact (Wysocki & Tamás, 2010). It can interfere with DNA repair, though evidence from sensitivity or increased double-stranded breaks is lacking (Wysocki & Tamás, 2010). The possible toxicity targets are as varied as abundant, which could be said to be a reflection of how life has evolved to depend on heavy metals and perturbing the default concentration has implications for proteins that have very little to do with each other.

7. Evolution

Evolution is at its heart infused with random processes and it is important to understand that it speaks of trends that only make sense when watched from a certain distance. The random element of evolution, genetic drift, is the redistribution of allele frequencies in a population due to chance. In my work, boom and bust cycles of colony growth makes timing of mutation events critical.

A typical colony is started with around 50k cells and at the end of the experiment, it has undergone roughly four doublings. The final population size is then around 800k cells. If a novel mutation with negligible effect appears among the 50k founder cells, the cell line will consist of 16 cells at the end of

the experiment when cells are transferred to start the next experimental cycle. If however the same mutation appears in the very last cell division, there will only be one out of 800k cells that has this mutation. Both of these may sound like impossibly low frequencies and the differences minute, but because 50k cells are taken out of the 800k cells to start the next batch, the chance of losing this novel neutral mutation is 94% in the latter case, but only 36% in the former.

If the example is modified so that there is a substantial beneficial growth effect due to the mutation, i.e. in the case when the population is exposed to an adverse environment, the non-mutated population will struggle to survive and only undergo say three doublings during the batch growth due to most of the cells' focus being redirected to detoxification and maintenance. If the mutation solves most of the toxicity by excluding arsenite and the mutation appears among one of the 50k founders, the mutated line would have managed five divisions given the same competition for resources as in the previous example, but because the bulk of the cells are not growing well, it can overshoot this by, let's say, managing ten doublings. The increased fitness will have caused the final population to be around 1k mutated cells among roughly 400k cells for the entire population. Because 50k cells are still transferred, chances of losing the mutation are virtually zero (less than 10^{-57}%). In fact, it is extremely likely that the next batch will start with many more than one mutated cell. If, however, the beneficial mutation appeared at the very end of the experiment, the fact that it increases fitness is irrelevant for its propagation, it will still only be one cell. However because the main population is struggling and only produced 400k cells, the odds of losing it has decreased to 87%.

So while selection is the skewed probability of propagation into future generation of some alleles, much like genetic drift, it is by no means an orderly process. The odds in the above examples come from the hypergeometric distribution and it is a simple matter of varying the parameters to see how the experimental setting affects the probability outcome. The second example reflects how the experiments of paper 4 were setup and so part of the swift adaptation can be explained by how evolution works on batch cultures.

This introduction focused on drift and selection, but there are other evolutionary processes. Migration in the laboratory setup of Scan-o-matic is of little concern as colonies are separated by sufficient distance and the primary study organism, yeast, does not have any motility. There is always the small risk that while transferring cells from one plate to another with the pinning robot, a cell will happen to fall off a pin and land elsewhere on the target plate as if moved by an occult hand. This unwanted migration event, or cross-contamination as it is commonly called, can happen and evolutionary experiments must of course be verified not to be explained by such experimental errors.

7.1. Change

Mutations happen in part because the DNA molecule is used. Either during the S-phase duplication of the genome or the transcription of individual genes (Alabert & Groth, 2012; (Rando & Verstrepen, 2007). Mutation rate sets the rate of change of the DNA and varies around 10^{-6} to 10^{-8} in most microbes (Rando & Verstrepen, 2007), but also within genomes (Coulondre & Miller, 1977). Generally mutation is considered a slow processes. The likelihood, as reflected by the mutation target size, to obtain a particular effect tends to be low and so the probability of many individuals getting mutations with similar effect at the same time becomes exceedingly small. Therefore the slow speed reflects the time it takes a novel allele to first appear and then to increase in frequency until it becomes relevant on the population level. In contrast epigenetic modifications are considered fast as they assume an underlying mechanism present in all cells waiting to be activated. The very dynamic and variable adhesion and flocculation properties of yeast is an example of epigenetically inherited cell states (Verstrepen & Klis, 2006).

Both papers 3 and 4 verified the reported observations by genome sequencing. In the case of paper 4, I contributed by testing the dynamics of losing these adaptations; if the strains were grown for 5, 10, 15, or more generations outside the test condition, how well would they then perform in the test condition? This uses precisely the discussed assumptions of speed, because mutations are considered practically irreversible while epigenetic effects more transient. Speed is also decided by the underlying complexity of the phenotype; the number of genes or nucleotide positions whose mutation state has a measureable

effect on the phenotype. In paper 4, the simulation experiments factor in both data in calculating the mutation rates for each of the substantial effect mutations that had previously been observed in the experimental evolution. Yet speed is not a sufficient argument and this is why sequencing is necessary.

There are ways that mutation-based adaptation can go faster than expected. In paper 4, we discuss mutation rates and while assuming basal mutation rates is parsimonious. Traits can have a capacity for accumulating variation if they are polygenic and rest on complex underlying structure. Under certain conditions this allows for accumulation of variation that is relatively silent, but given the correct environment or stress may alter the phenotype (Halfmann & Lindquist, 2010). There are also observations of regulation of local and global variation (Jablonka & Lamb, 2005; Metzgar & Wills, 2000) to modulate the need for change versus stability.

In the opposite direction, some epigenetic effects can be more long-lasting than others. One example can be seen in a methylation event with a frequency of 10^{-4} (and 10^{-2} in other direction) as exemplified in the methylation of the pyelonephritis-associated pili operon (Hernday *et. al.*, 2002).

7.2. Experimental evolution

I would like to start by making a distinction between artificial selection and experimental evolution. To illustrate the former, consider a researcher performing a selection experiment on *Drosophila melanogaster* and only allowing the individuals with the highest number of bristles to mate and form the next generation (Spiess & Wilke, 1984). This is artificial selection in two ways. First, the selection event correlates very poorly with the overall environment that the flies are exposed to; it is only invoked by the researcher from the outside and only at the end of each experimental cycle. Second, the selection is very binary given that the researcher makes no sorting mistakes, i.e. the probability of survival is zero up to the threshold where flies will be included in the next cycle and almost one after that (allowing for unwanted deaths). The *Zea mays* high oil content experiment that began in 1896 and was recently reported in (Moose *et. al.*, 2004) also has aspects of poor integration of the selection pressure, oil content, with the overall environment of the study

(living on a laboratory farm plot). Neither fly nor maize selection should be understood as fully separate from the non-experimental selection pressures (bristle counts have been shown to correlate with mating success (Markow *et. al.*, 1996), and oil concentration affects health of the crop at least with low oil content (Laurie *et. al.*, 2004)). However, because the oil content testing is destructive for the maize kernel tested, it is the neighbouring kernels that are either included or excluded in the next crop cycle. This introduces variation and makes the probability of survival given a certain oil content a smoother transition than the bristle counting example. Finally, in our experimental evolution experiment in paper 4, *Saccharomyces cerevisiae* is left to grow in a highly toxic environment all to itself and while the exactitude of the environment and the magnitude of the toxicity are unlikely to be encountered outside the lab, it is the integrated experience of this environment that is the selection. The final population is randomly sub-sampled for individuals to seed the next. With regards to probability of survival and to how integrated the selection is with the general environment, the artificial aspect of the selection is low. Arranging the experiments from high to low levels of artificialness in the selection scheme they come out: fly, maize, and finally yeast. All three experiments are however still experimental evolutions.

The introductory paragraphs to evolution had a numerical example of the probability of retaining near neutral and beneficial mutations in a boom and bust batch culture type of setting. In describing selection on standing genetic variation versus novel mutations, the latter is often characterized as a problematic paradigm as the beneficial mutations will be exceedingly rare and since the rare mutations will be very easily lost to the population due to genetic drift. As illustrated in my example, batch growth increases the odds of beneficial mutations being kept greatly as long as the sampling bottleneck for the next batch isn't too narrow or too wide.

It would be of interest to study the selection pressures for batch growth lifestyle on clonal organisms. With regards to niche construction and the primarily clonal propagation, I discussed the high probability of proximate cells to have identical or near identical genomes. This opens up the possibility of arguing for kin selection. Using Hamilton's rule $rB > C$ from the reasoning in Hamilton (1963)

and Queller (1992), where r is the relatedness, B is the beneficial effect to the receiving party and C is the cost on the altruistic party, the outcome in the stipulated scenario would reduce to $B > C$, since r is 1 or near 1. So while altruistic behaviour is controversial in general, as the equation suggests, it should not be surprising in this context if my characterization of colony structure is correct. It should be noted too that sacrificing for the greater good of the near isogenic microbial colony in this sense is quite parallel to cells showing altruism in a multicellular organism. The latter being a fundamental requirement for multicellularity and by no means controversial.

If kin selection can modulate saturation of the population size rather than nutrient limitations in the environment, this could be a way to maximize the evolvability of positive traits in the near isogenic populations by maximizing the ability of new adaptive mutations to establish themselves in the population. Another prediction is that if an experimental batch culture evolutionary setup includes very extended stationary phases with substantial cell turnover rate, this would increase the risk of beneficial mutations being lost due to drift. The hypothesis requires that the mutation doesn't affect survivability – that it only affects growth dynamics. These kin selection hypotheses would most easily be tested with computer simulations.

8. Main findings of included papers

Paper 1: Throughput doesn't negate quality

In paper 1 we found that it is possible to maintain very high quality measurements of colony sizes during growth while still having a high throughput. We introduced a novel method for removing spatial bias observed in solid media growth experiments and characterized the noise and bias in the data. We found that we could reliably translate the image data to population sizes. We also found that good results could be achieved by using ordinary consumer products. I developed the software and method.

Paper 2: Phasing the growth curve

We investigated the reliability of segmenting the growth curve into its fundamental parts without presupposing any overall sigmoid or multi-modal shape of the curve and using these segments as basis for extracting more data from each experiment. In this process we updated the curve smoothing for increased robustness. I developed the software and method.

Paper 3: Phasing the genomes

We found that through a well constructed mating system, the variation in traits could be teased apart into the contribution from additive, dominance, and multi-order epistasis. This was done for several environments and the underlying genetic structure was investigated with QTL analysis. I assisted in the phenotyping and primary analysis.

Paper 4: Change can be exceedingly fast

The most curious outcome of the study was that parallel evolutionary processes could more or less synchronously sweep independent populations with change in a single jolt of adaptation to each of them. Four parallel populations independently adapted fully to arsenite after less than twenty generations and the with very rapid approach to the adapted phenotype. And while these sweeps superficially thus appeared deterministic, the underlying genetic solutions to arsenic toxicity were different. I did the laboratory deadaptation.

))<>((

Included papers

1. Zackrisson, Martin, et al. "Scan-o-matic: high-resolution microbial phenomics at a massive scale." *G3: Genes, Genomes, Genetics* 6.9 (2016): 3003-3014.
2. Zackrisson, Martin, et al. "Phase-segmentation of the microbial growth curve" (2017) Manuscript in preparation
3. Hallin, Johan, et al. "Powerful decomposition of complex traits in a diploid model." *Nature communications* 7 (2016): 13311.
4. Gjuvsland, Arne B., et al. "Disentangling genetic and epigenetic determinants of ultrafast adaptation." *Molecular systems biology* 12.12 (2016): 892.

Papers not included

1. Sundqvist, Lisa, et al. "Directional genetic differentiation and relative migration." *Ecology and evolution* 6.11 (2016): 3461-3475.
2. Babazadeh, Roja, et al. "The Ashbya gossypii EF‐1α promoter of the ubiquitously used MX cassettes is toxic to Saccharomyces cerevisiae." *FEBS letters* 585.24 (2011): 3907-3913.
3. Fernandez-Ricaud, Luciano, et al. "PRECOG: a tool for automated extraction and visualization of fitness components in microbial growth phenomics." *BMC bioinformatics* 17.1 (2016): 249.

Thanks

I have a bunch of people to thank, but I'll do so without using their names. Because I'm afraid, actually I'm 100% sure, I might miss some. So here it goes: I want to thank them who ensured I kept reasonable hours and workload. Those who let me reign rather freely and redefine the whole purpose of my PhD. Who somehow thought I would actually complete that piece of software – I sure was doubting. My pet physicist, without whom there would have been many errors and bad math. Those that caused or assisted extra long fika breaks when things were rough or the sun was finally out after the long winter. Those sometimes very animated lunchroom discussions about statistics, science in general, tangential nerdy topics about some ancient military thingies of some war, and actually even eurovision song contest. Those who let me play my strange music up in the lab before I escaped all the experimentation. Those who sat next to me and patently let me interrupt them as if they were part of my dev-team. Those students who passed through or stayed for much longer times and in many ways helped form how it all turned out. Those who lent me a book they had written about an alternative reality at the lab. Those who helped me in setting up the system I was building, lending tools, cables, cutting things, installing all sorts of software, compromising the it security, sending hundred thousand mails in matter of seconds by accident. The administrators who helped though I refused to have neither mac nor windows.

References

Aird, Daniel, et al. "Analyzing and minimizing PCR amplification bias in Illumina sequencing libraries." *Genome biology* 12.2 (2011): R18.

Alabert, Constance, and Anja Groth. "Chromatin replication and epigenome maintenance." *Nature reviews. Molecular cell biology* 13.3 (2012): 153.

Allen, Jess, et al. "High-throughput classification of yeast mutants for functional genomics using metabolic footprinting." *Nature biotechnology* 21.6 (2003): 692.

Ashburner, Michael, et al. "Gene Ontology: tool for the unification of biology." *Nature genetics* 25.1 (2000): 25.

Ashburner, Michael. "Gene Ontology Consortium" Gene Ontology Consortium, Web. 29 Jul. 2017

Banks, A. P., C. Lawless, and D. A. Lydall. "A quantitative fitness analysis workflow." *Journal of visualized experiments: JoVE* 66 (2012).

Baryshnikova, Anastasia, et al. "Quantitative analysis of fitness and genetic interactions in yeast on a genome scale." *Nature methods* 7.12 (2010): 1017-1024.

Barral, Yves, Stefan Jentsch, and Carl Mann. "G1 cyclin turnover and nutrient uptake are controlled by a common pathway in yeast." *Genes & Development* 9.4 (1995): 399-409.

Bean, Gordon J., et al. "Development of ultra-high-density screening tools for microbial "omics"." *PLoS One* 9.1 (2014): e85177.

Bloom, Joshua S., et al. "Finding the sources of missing heritability in a yeast cross." *Nature* 494.7436 (2013): 234.

Bodvard, Kristofer, et al. "The yeast transcription factor Crz1 is activated by light in a Ca2+/calcineurin-dependent and PKA-independent manner." *PLoS One* 8.1 (2013): e53404.

Boyle, Evan A., Yang I. Li, and Jonathan K. Pritchard. "An Expanded View of Complex Traits: From Polygenic to Omnigenic." *Cell* 169.7 (2017): 1177-1186.

Brachmann, C. Baker, et al. "Designer deletion strains derived from Saccharomyces cerevisiae S288C: a useful set of strains and plasmids for PCR-mediated gene disruption and other applications." *YEAST-CHICHESTER-* 14 (1998): 115-132.

Brideau, Christine, et al. "Improved statistical methods for hit selection in high-throughput screening." *Journal of biomolecular screening* 8.6 (2003): 634-647.

Buchanan, R. L., R. C. Whiting, and W. C. Damert. "When is simple good enough: a comparison of the Gompertz, Baranyi, and three-phase linear models for fitting bacterial growth curves." *Food Microbiology* 14.4 (1997): 313-326.

Buser, Claudia C., et al. "Niche construction initiates the evolution of mutualistic interactions." *Ecology Letters* 17.10 (2014): 1257-1264.

Carter, B. L. A., and M. N. Jagadish. "The relationship between cell size and cell division in the yeast Saccharomyces cerevisiae." *Experimental cell research* 112.1 (1978): 15-24.

Chang, Fred, and Ira Herskowitz. "Identification of a gene necessary for cell cycle arrest by a negative growth factor of yeast: FAR1 is an inhibitor of a G1 cyclin, CLN2." *Cell* 63.5 (1990): 999-1011.

Chingozha, Loice, et al. "A generalizable, tunable microfluidic platform for delivering fast temporally varying chemical signals to probe single-cell response dynamics." *Analytical chemistry* 86.20 (2014): 10138-10147.

Churchill, Frederick B. "William Johannsen and the genotype concept." *Journal of the History of Biology* 7.1 (1974): 5-30.

Collins, Sean R., et al. "A strategy for extracting and analyzing large-scale quantitative epistatic interaction data." *Genome biology* 7.7 (2006): R63.

Costanzo, Michael, et al. "The genetic landscape of a cell." *science* 327.5964 (2010): 425-431.

Côté, Jean-Philippe, et al. "The genome-wide interaction network of nutrient stress genes in Escherichia coli." *mBio* 7.6 (2016): e01714-16.

Coulondre, Christine, and Jeffrey H. Miller. "Genetic studies of the lac repressor: IV. Mutagenic specificity in the lacI gene of Escherichia coli." *Journal of molecular biology* 117.3 (1977): 577-606.

Crespi, Bernard J. "The evolution of social behavior in microorganisms." *Trends in ecology & evolution* 16.4 (2001): 178-183.

Cubillos, Francisco A., et al. "Assessing the complex architecture of polygenic traits in diverged yeast populations." *Molecular ecology* 20.7 (2011): 1401-1413.

Dens, Els J., et al. "Cell division theory and individual-based modeling of microbial lag: Part I. The theory of cell division." *International journal of food microbiology* 101.3 (2005): 303-318.

Di Talia, Stefano, et al. "The effects of molecular noise and size control on variability in the budding yeast cell cycle." *Nature* 448.7156 (2007): 947.

Fisher, Ronald Aylmer. *The genetical theory of natural selection: a complete variorum edition*. Oxford University Press, 1930.

French, Shawn, et al. "A robust platform for chemical genomics in bacterial systems." *Molecular biology of the cell* 27.6 (2016): 1015-1025.

Gartner, Jared J., et al. "Whole-genome sequencing identifies a recurrent functional synonymous mutation in melanoma." *Proceedings of the National Academy of Sciences* 110.33 (2013): 13481-13486.

Gegas, Vasilis C., et al. "Challenges of crop phenomics in the post-genomics era." *Phenomics*. CRC Press, 2014. 142-171.

Giaever, Guri, et al. "Functional profiling of the Saccharomyces cerevisiae genome." *nature* 418.6896 (2002): 387.

Gimeno, Carlos J., et al. "Unipolar cell divisions in the yeast S. cerevisiae lead to filamentous growth: regulation by starvation and RAS." *Cell* 68.6 (1992): 1077-1090.

Goodman, Steven N. "A comment on replication, P-values and evidence." *Statistics in medicine* 11.7 (1992): 875-879.

Gray, Joseph V., et al. ""Sleeping beauty": quiescence in Saccharomyces cerevisiae." *Microbiology and Molecular Biology Reviews* 68.2 (2004): 187-206.

Halfmann, Randal, and Susan Lindquist. "Epigenetics in the extreme: prions and the inheritance of environmentally acquired traits." *Science* 330.6004 (2010): 629-632.
Hallam, Anthony. "How rare is phyletic gradualism and what is its evolutionary significance? Evidence from Jurassic bivalves." *Paleobiology* 4.1 (1978): 16-25.
Halsey, Lewis G., et al. "The fickle P value generates irreproducible results." *Nature methods* 12.3 (2015): 179-185.
Hamilton, W. D. "The genetical evolution of social behaviour. I." *Journal of Theoretical Biology* 7.1 (1964): 1-16.
Hartman, John L., Barbara Garvik, and Lee Hartwell. "Principles for the buffering of genetic variation." *Science* 291.5506 (2001): 1001-1004.
Hartman, John L., et al. "Yeast phenomics: an experimental approach for modeling gene interaction networks that buffer disease." *Genes* 6.1 (2015): 24-45.
Hartwell, Leland H., and Ted A. Weinert. "Checkpoints: controls that ensure the order of cell cycle events." *Science* 246.4930 (1989): 629-634.
Hernday, Aaron, et al. "Self-perpetuating epigenetic pili switches in bacteria." *Proceedings of the National Academy of Sciences* 99.suppl 4 (2002): 16470-16476.
Hittinger, Chris Todd. "Saccharomyces diversity and evolution: a budding model genus." *Trends in Genetics* 29.5 (2013): 309-317.
Houle, David, Diddahally R. Govindaraju, and Stig Omholt. "Phenomics: the next challenge." *Nature Reviews. Genetics* 11.12 (2010): 855.
Huang, Da Wei, Brad T. Sherman, and Richard A. Lempicki. "Bioinformatics enrichment tools: paths toward the comprehensive functional analysis of large gene lists." *Nucleic acids research* 37.1 (2008): 1-13.
Jablonka, Eva, and Marion J. Lamb. "Evolution in four dimensions: genetic, epigenetic, behavioral, and symbolic variation in the history of life. A Bradford Book." (2005).
Jensen, Robert, George F. Sprague, and Ira Herskowitz. "Regulation of yeast mating-type interconversion: feedback control of HO gene expression by the mating-type locus." *Proceedings of the National Academy of Sciences* 80.10 (1983): 3035-3039.
Jorgensen, Paul, et al. "Systematic identification of pathways that couple cell growth and division in yeast." *Science* 297.5580 (2002): 395-400.
Julius, M. H., T. Masuda, and L. A. Herzenberg. "Demonstration that antigen-binding cells are precursors of antibody-producing cells after purification with a fluorescence-activated cell sorter." *Proceedings of the National Academy of Sciences* 69.7 (1972): 1934-1938.
Kaprelyants, Arseny S., and Douglas B. Kell. "Do bacteria need to communicate with each other for growth?." *Trends in microbiology* 4.6 (1996): 237-242.
Kirkwood, Thomas BL. "Understanding the odd science of aging." *Cell* 120.4 (2005): 437-447.
Koboldt, Daniel C., et al. "The next-generation sequencing revolution and its impact on genomics." *Cell* 155.1 (2013): 27-38.
Kuhn, Thomas S., and David Hawkins. "The structure of scientific revolutions." *American Journal of Physics* 31.7 (1963): 554-555.
Kurtzman, Cletus, Jack W. Fell, and Teun Boekhout, eds. *The yeasts: a taxonomic study*. Elsevier, 2011.
Kvitek, Daniel J., Jessica L. Will, and Audrey P. Gasch. "Variations in stress sensitivity and genomic expression in diverse S. cerevisiae isolates." *PLoS genetics* 4.10 (2008): e1000223.
Lafon-Lafourcade, S., F. Larue, and P. Ribereau-Gayon. "Evidence for the existence of “survival factors” as an explanation for some peculiarities of yeast growth, especially in grape must of high sugar concentration." *Applied and environmental microbiology* 38.6 (1979): 1069-1073.
Laland, Kevin N., F. John Odling-Smee, and Marcus W. Feldman. "Evolutionary consequences of niche construction and their implications for ecology." *Proceedings of the National Academy of Sciences* 96.18 (1999): 10242-10247.
Laurie, Cathy C., et al. "The genetic architecture of response to long-term artificial selection for oil concentration in the maize kernel." *Genetics* 168.4 (2004): 2141-2155.
Lawless, Conor, et al. "Colonyzer: automated quantification of micro-organism growth characteristics on solid agar." *BMC bioinformatics* 11.1 (2010): 287.
Lercher, Martin J., Thomas Blumenthal, and Laurence D. Hurst. "Coexpression of neighboring genes in Caenorhabditis elegans is mostly due to operons and duplicate genes." *Genome research* 13.2 (2003): 238-243.
Levy, Sasha F., Naomi Ziv, and Mark L. Siegal. "Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant." *PLoS biology* 10.5 (2012): e1001325.
Li, Zhijian, et al. "Systematic exploration of essential yeast gene function with temperature-sensitive mutants." *Nature biotechnology* 29.4 (2011): 361-367.
Libkind, Diego, et al. "Microbe domestication and the identification of the wild genetic stock of lager-brewing yeast." *Proceedings of the National Academy of Sciences* 108.35 (2011): 14539-14544.
Liti, Gianni. "The Natural History of Model Organisms: The fascinating and secret wild life of the budding yeast S. cerevisiae." *Elife* 4 (2015): e05835.
Liu, Haoping, Cora Ann Styles, and Gerald R. Fink. "Saccharomyces cerevisiae S288C has a mutation in FL08, a gene required for filamentous growth." *Genetics* 144.3 (1996): 967-978.
Logg, Katarina, et al. "Investigations on light-induced stress in fluorescence microscopy using nuclear localization of the transcription factor Msn2p as a reporter." *FEMS yeast research* 9.6 (2009): 875-884.
MacLean, R. C. "The tragedy of the commons in microbial populations: insights from theoretical, comparative and experimental studies." *Heredity* 100.5 (2008): 471.
Malo, Nathalie, et al. "Statistical practice in high-throughput screening data analysis." *Nature biotechnology* 24.2 (2006): 167.
Mangat, Chand S., et al. "Rank ordering plate data facilitates data visualization and normalization in high-throughput screening." *Journal of biomolecular screening* 19.9 (2014): 1314-1320.

Manhart, Michael, Bharat V. Adkar, and Eugene I. Shakhnovich. "Tradeoffs between microbial growth phases lead to frequency-dependent and non-transitive selection." *arXiv preprint arXiv:1612.07908* (2016).
Mardis, Elaine R. "Next-generation sequencing platforms." *Annual review of analytical chemistry* 6 (2013): 287-303.
Markow, Therese Ann, David Bustoz, and Scott Pitnick. "Sexual selection and a secondary sexual character in twoDrosophilaspecies." *Animal Behaviour* 52.4 (1996): 759-766.
Metzgar, David, and Christopher Wills. "Evidence for the adaptive evolution of mutation rates." *Cell* 101.6 (2000): 581-584.
Moose, Stephen P., John W. Dudley, and Torbert R. Rocheford. "Maize selection passes the century mark: a unique resource for 21st century genomics." *Trends in plant science* 9.7 (2004): 358-364.
Monod, Jacques. "The growth of bacterial cultures." *Annual Reviews in Microbiology* 3.1 (1949): 371-394.
Mortimer, Robert K., and John R. Johnston. "Life span of individual yeast cells." *Nature* 183.4677 (1959): 1751-1752.
Munafò, Marcus R., et al. "A manifesto for reproducible science." *Nature Human Behaviour* 1 (2017): 0021.
Murakami, Christopher J., et al. "A method for high-throughput quantitative analysis of yeast chronological life span." *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 63.2 (2008): 113-121.
Märtens, Kaspar, et al. "Predicting quantitative traits from genome and phenome with near perfect accuracy." *Nature communications* 7 (2016).
Nagalakshmi, Ugrappa, et al. "The transcriptional landscape of the yeast genome defined by RNA sequencing." *Science* 320.5881 (2008): 1344-1349.
Narayanan, Siddharth, et al. "Quantitative Fitness Analysis Identifies exo1Δ and Other Suppressors or Enhancers of Telomere Defects in Schizosaccharomyces pombe." *PloS one* 10.7 (2015): e0132240.
Novak, Maja, et al. "Experimental tests for an evolutionary trade-off between growth rate and yield in E. coli." *The American Naturalist* 168.2 (2006): 242-251.
Nurse, Paul. "Ordering S phase and M phase in the cell cycle." *Cell* 79.4 (1994): 547-550.
Orr, H. Allen. "Fitness and its role in evolutionary genetics." *Nature reviews. Genetics* 10.8 (2009): 531.
Parts, Leopold, et al. "Revealing the genetic structure of a trait by sequencing a population under selection." *Genome research* 21.7 (2011): 1131-1138.
Penfold, William Jas. "On the nature of bacterial lag." *The Journal of hygiene* 14.2 (1914): 215.
Pennisi, Elizabeth. "ENCODE project writes eulogy for junk DNA." *Science* 337.6099 (2012): 1159-1161.
Petkov, Petko M., et al. "Evidence of a large-scale functional organization of mammalian chromosomes." *PLoS genetics* 1.3 (2005): e33.
Pipe, Leonie Z., and Malcolm J. Grimson. "Spatial-temporal modelling of bacterial colony growth on solid media." *Molecular BioSystems* 4.3 (2008): 192-198.
Poláček, Lukáš. "How to Shuffle songs?" *Labs*. Spotify, 09 Dec. 2014. Web. 02 Aug. 2017
Queller, David C. "A general model for kin selection." *Evolution* 46.2 (1992): 376-380.
Rando, Oliver J., and Kevin J. Verstrepen. "Timescales of genetic and epigenetic inheritance." *Cell* 128.4 (2007): 655-668.
Reece, Amy, et al. "Microfluidic techniques for high throughput single cell analysis." *Current opinion in biotechnology* 40 (2016): 90-96.
Roemer, Terry, and Howard Bussey. "Yeast beta-glucan synthesis: KRE6 encodes a predicted type II membrane protein required for glucan synthesis in vivo and for glucan synthase activity in vitro." *Proceedings of the National Academy of Sciences* 88.24 (1991): 11295-11299.
Shepherd, Maxwell G. "Morphogenetic transformation of fungi." *Curr Top Med Mycol* 2 (1988): 278-304.
Sibly, Richard M., and Jim Hone. "Population growth rate and its determinants: an overview." *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357.1425 (2002): 1153-1170.
Sicard, Delphine, and Jean-Luc Legras. "Bread, beer and wine: yeast domestication in the Saccharomyces sensu stricto complex." *Comptes rendus biologies* 334.3 (2011): 229-236.
Sonka, Milan, Vaclav Hlavac, and Roger Boyle. *Image processing, analysis, and machine vision*. Cengage Learning, 2014.
Sopko, Richelle, et al. "Mapping pathways and phenotypes by systematic gene overexpression." *Molecular cell* 21.3 (2006): 319-330.
Spellman, Paul T., and Gerald M. Rubin. "Evidence for large domains of similarly expressed genes in the Drosophila genome." *Journal of biology* 1.1 (2002): 5.
Spiess, Eliot B., and Charles M. Wilke. "Still another attempt to achieve assortative mating by disruptive selection in Drosophila." *Evolution* 38.3 (1984): 505-515.
Stevenson, Keiran, et al. "General calibration of microbial growth in microplate readers." *Scientific reports* 6 (2016): 38828.
Styblo, Miroslav, et al. "Comparative inhibition of yeast glutathione reductase by arsenicals and arsenothiols." *Chemical research in toxicology* 10.1 (1997): 27-33.
Thatcher, J. W., Janet M. Shaw, and W. J. Dickinson. "Marginal fitness contributions of nonessential genes in yeast." *Proceedings of the National Academy of Sciences* 95.1 (1998): 253-257.
Tong, Amy Hin Yan, et al. "Systematic genetic analysis with ordered arrays of yeast deletion mutants." *Science* 294.5550 (2001): 2364-2368.
Turner, Jonathan J., Jennifer C. Ewald, and Jan M. Skotheim. "Cell size control in yeast." *Current Biology* 22.9 (2012): R350-R359.
Turner, Paul E., Valeria Souza, and Richard E. Lenski. "Tests of ecological mechanisms promoting the stable coexistence of two bacterial genotypes." *Ecology* 77.7 (1996): 2119-2129.
Tyers, Mike. "BioGRID Database Statistics v3.4.150" *BioGRID*. TyersLab.com, 01 Jul. 2017, Web. 02 Aug. 2017
Vázquez-García, Ignacio, et al. "Background-dependent effects of selection on subclonal heterogeneity." *bioRxiv* (2016): 039859.
Verstrepen, Kevin J., and Frans M. Klis. "Flocculation, adhesion and biofilm formation in yeasts." *Molecular microbiology* 60.1 (2006): 5-15.

Walther, Andrea, Ana Hesselbart, and Jürgen Wendland. "Genome sequence of Saccharomyces carlsbergensis, the world's first pure culture lager yeast." *G3: Genes, Genomes, Genetics* 4.5 (2014): 783-793.
Warner, Jonathan R. "The economics of ribosome biosynthesis in yeast." *Trends in biochemical sciences* 24.11 (1999): 437-440.
Warringer, Jonas, and Anders Blomberg. "Automated screening in environmental arrays allows analysis of quantitative phenotypic profiles in Saccharomyces cerevisiae." *Yeast* 20.1 (2003): 53-67.
Warringer, Jonas, et al. "High-resolution yeast phenomics resolves different physiological features in the saline response." *Proceedings of the national academy of sciences* 100.26 (2003): 15724-15729.
Warringer, Jonas, et al. "Chemogenetic fingerprinting by analysis of cellular growth dynamics." *BMC chemical biology* 8.1 (2008): 3.
Warringer, Jonas, et al. "Trait variation in yeast is defined by population history." *PLoS genetics* 7.6 (2011): e1002111.
West, Stuart A., et al. "Social evolution theory for microorganisms." *Nature Reviews Microbiology* 4.8 (2006): 597-607.
Winzeler, Elizabeth A., et al. "Functional characterization of the S. cerevisiae genome by gene deletion and parallel analysis." *science* 285.5429 (1999): 901-906.
Wysocki, Robert, and Markus J. Tamás. "How Saccharomyces cerevisiae copes with toxic metals and metalloids." *FEMS microbiology reviews* 34.6 (2010): 925-951.
Yu, Haiyuan, et al. "High-quality binary protein interaction map of the yeast interactome network." *Science* 322.5898 (2008): 104-110.
Yue, Jia-Xing, et al. "Contrasting evolutionary genome dynamics between domesticated and wild yeasts." *Nature* 201: 7.
Ziv, Naomi, Nathan J. Brandt, and David Gresham. "The use of chemostats in microbial systems biology." *Journal of visualized experiments: JoVE* 80 (2013).
Zlitni, Soumaya, Lauren F. Ferruccio, and Eric D. Brown. "Metabolic suppression identifies new antibacterial inhibitors under nutrient limitation." *Nature chemical biology* 9.12 (2013): 796-804.