Genome-scale analysis of interactions between genetic perturbations and natural variation

Joseph J. Hale[1], Takeshi Matsui[2,#], Ilan Goldstein[1,#], Martin N. Mullis[1], Kevin R. Roy[3,4], Chris Ne Ville[1], Darach Miller[2], Charley Wang[1], Trevor Reynolds[1], Lars M. Steinmetz[3,4,5], Sasha F. Levy[2,%], Ian M. Ehrenreich[1,%]

[1]Molecular and Computational Biology Section, Department of Biological Sciences, University of Southern California, Los Angeles, CA 90089, USA

[2]SLAC National Accelerator Laboratory, Menlo Park, CA, 94025, USA

[3]Stanford Genome Technology Center, Stanford University, Palo Alto, California, USA

[4]Department of Genetics, Stanford University School of Medicine, Stanford, California, USA

[5]European Molecular Biology Laboratory, Genome Biology Unit, Heidelberg, Germany

[#]Equal contribution

[%]Correspondence to:
Ian M. Ehrenreich
Phone #: 213 – 821 – 5349
Email: ian.ehrenreich@usc.edu

Sasha F. Levy
Phone #: 646 - 417 - 3262
Email: sflevy@stanford.edu

Abstract

Interactions between genetic perturbations and segregating loci can cause perturbations to show different phenotypic effects across genetically distinct individuals. To study these interactions on a genome scale in many individuals, we used combinatorial DNA barcode sequencing to measure the fitness effects of 8,013 CRISPRi perturbations targeting 1,721 distinct genes in 169 yeast cross progeny (or segregants). We identified 457 genes whose perturbation has different effects across segregants. Several factors caused perturbations to show variable effects, including baseline segregant fitness, the mean effect of a perturbation across segregants, and interacting loci. We mapped 248 interacting loci and found four hub loci that interact with many different perturbations. Perturbations that interact with a given hub exhibit similar epistatic relationships with the hub and show enrichment for cellular processes that may mediate these interactions. These results suggest that an individual's response to perturbations is shaped by a network of perturbation-locus interactions that cannot be measured by approaches that examine perturbations or natural variation alone.

Introduction

Genetic perturbations can interact with genetic variants (or loci) segregating in a population[1]. These interactions cause perturbations to show different phenotypic effects in genetically distinct individuals (also known as genetic background effects)[2]. For example, disease mutations often exhibit incomplete penetrance and variable expressivity in humans[3]. Similarly, background effects could cause genome editing to have variable efficacy, as applied to human disease, productivity of crops or livestock, and engineering of cells and microorganisms for industrial applications. In an evolutionary context, the fitness effect of *de novo* mutations can depend on interactions with specific alleles present in an individual. These interactions can impact the dynamics of beneficial mutations spreading through a population[4,5] or deleterious mutations persisting within a population[6].

Interactions between genetic perturbations and loci also impact the phenotypic effects of loci. A perturbation may alter the degree to which an interacting locus influences a trait (magnitude epistasis) or which allele of that locus confers a higher trait value (sign epistasis)[7,8]. In some cases, a perturbation can uncover cryptic loci that do not typically show phenotypic effects[9,10], or mask loci that usually exhibit phenotypic effects[11]. The net effect of a perturbation can depend on epistasis with multiple loci that interact not only with the perturbation but also each other (higher-order epistasis)[12]. That is, the response of a given individual to a perturbation will depend on its genotype at all interacting loci[13].

While numerous examples of background effects have been reported[14–28], we lack a detailed understanding of interactions between genetic perturbations and genetic backgrounds. For example, we do not yet know what characteristics of perturbations and individuals' genotypes increase the likelihood of interactions, what types of epistasis are most likely to underlie background effects, and the properties of interaction networks between perturbations and loci. To answer these questions systematically, we require two things: perturbations at a genomic scale and extremely high-throughput

measurement of the phenotypic effects of these perturbations across many individuals. Studies of this scale have been technically challenging and cost prohibitive.

Here, we developed a scalable platform for assaying interactions between genetic backgrounds and genetic perturbations using a genomic double barcoding system[29,30] and inducible CRISPR interference (CRISPRi)[31–33]. We utilized 169 haploid progeny (segregants) from a cross between the BY4742 (BY) lab and 322134S (3S) clinical strains of *Saccharomyces cerevisiae*, each of which contained a barcode that marked a segregant[34,35]. We then integrated a genome-scale library of 8,013 gRNA RNAs (gRNAs), each marked with a barcode, at an adjacent site. Pooled competitions, double barcode sequencing, and lineage trajectory analyses enabled precise measurement of the fitness of ~2,900,000 double barcode combinations, representing ~875,000 segregant-gRNA combinations. This large dataset allowed us to identify 457 genes whose perturbation has different effects across segregants and to characterize factors that cause these background effects.

Results

Construction of a double barcoded library of segregants with inducible CRISPRi perturbations

We previously generated, genotyped, and barcoded 822 haploid, *MATα* BYx3S segregants carrying a genomic landing pad at the neutral *YBR209W* locus[34,35] (Supplemental Figure 1 and 2). These segregants are amenable to pooled, high-throughput phenotyping because they lack the main flocculin (*FLO11*)[36] and the primary transcriptional activator of cell-cell and cell-surface adhesion phenotypes (*FLO8*) present in this species[37]. In this study, we randomly selected 169 of these segregants to generate a genetically diverse panel (Supplemental Data 1 and 2). Most of these segregants were represented by a single barcode in the panel. To enable tests for reproducibility, 28 segregants were represented by 2-3 distinct barcodes.

We also designed a barcoded CRISPRi plasmid library that could be integrated into the genomic landing pad adjacent to the segregant barcode (Supplemental Figures 3 and 4). An inducible CRISPRi system[32] was chosen because it provides several key advantages. CRISPRi perturbations are only induced during phenotyping experiments, which helps prevent accumulation of *de novo* suppressor mutations during strain construction and culturing. The use of inducible CRISPR interference also allows essential genes to be perturbed, expanding the space of possible target genes. Finally, CRISPRi-mediated gene repression is orthogonal to a cell's endogenous transcriptional machinery and should perform similarly across strains. Each plasmid in this library contained an endonuclease-dead Cas9 (dCas9) fused with the Mxi1 repressor domain, an inducible single guide RNA (gRNA) targeting dCas9 to a specific promoter, a unique 20-nucleotide barcode, and a loxP site to enable chromosomal integration. This plasmid library targeted a total of 1,739 genes that have been reported to be essential under fermentative or respiratory growth conditions with 8,864 unique gRNAs, with an average of five distinct gRNAs per gene[33] (Supplemental Data 3, Supplemental Figure 5). 189 of these unique gRNAs were controls, which targeted intergenic and noncoding regions rather than promoters. Roughly 10 distinct barcodes were generated for each gRNA on

average, resulting in an estimated plasmid library size of ≥90,000. All gRNAs were under the control of a tetO-modified *RPR1* promoter, enabling gRNA induction by the addition of anhydrotetracycline (ATC) to liquid growth medium. ATC itself has no effect on yeast growth at the concentrations used in this assay[32,38,39].

We integrated this library of barcoded CRISPRi plasmids into each segregant and the BY parent using large-scale transformations and Cre/loxP-mediated site-specific recombination. Natural variation in transformability impeded integration of the library into certain segregants, as well as into the 3S parent (Supplemental Figure 6). Based on colony counts and sequencing, we estimated that ~10,000-15,000 integrants, or ~1.1-1.7 barcodes per gRNA, were recovered per successfully transformed strain (Supplemental Data 4). Transformants were grown to stationary phase and roughly 10^8-10^9 cells per strain were combined into an initial pool (T0).

Phenotyping of the segregant-gRNA pool

Plasmids were designed so that after integration, chromosomally-encoded segregant and plasmid barcodes would be 99 bp apart, enabling their co-amplification by PCR and sequencing within the same paired-end read (**Fig. 1A**, Supplemental Figure 3). Sequencing of barcode pairs revealed that 169 unique segregants (223 segregant barcodes) and 8,013 unique gRNAs (107,003 gRNA barcodes) targeting 1,721 distinct genes were present in T0. Each gRNA appeared in 107 segregants on average, with T0 containing ~2,951,000 double barcodes in total. On average, ~4.6 gRNAs targeted a given gene (Supplemental Data 5).

The T0 pool was split into three separate fitness assays: two experimental assays where the gRNA expression was induced with ATC and one control assay where no ATC was added. All assays were done in synthetic complete medium lacking uracil, a selection for integrated CRISPRi plasmids, and containing glucose, a fermentable carbon source. We performed the assays for 10 generations in serial batch culture, diluting 1:4 roughly every two generations, with a bottleneck size of ~2.5×10^{10} cells. We sequenced double barcodes, counted their frequencies over five time points, and used

barcode frequency trajectories to estimate the relative fitness of each lineage[40,41] (**Fig. 1B,** Supplemental Figures 7 - 10; Supplemental Table 1; Supplemental Data 6 - 8).

We found that fitness estimates for individual double barcode lineages were highly correlated between the replicate experimental flasks (ATC1 and ATC2, **Fig. 1C,** Spearman's correlation: 0.798). Correlation was lower between the experimental ATC2 flask and the control flask that lacked ATC entirely (CON, **Fig. 1D,** Spearman's correlation: 0.689). Any lineages carrying efficacious gRNAs are likely to have very different fitness estimates between ATC2 and CON. The reason that ATC2 and CON show a weaker, but still high correlation is likely because only a subset of gRNAs have strong effects on fitness. To identify these efficacious gRNAs, we used mixed effects linear models that accounted for both the baseline fitnesses of segregants and the mean fitness effects of gRNAs across segregants. Similar to prior work[32], ~32% of gRNAs had significant fitness effects under this model. The specificity of CRISPRi targeting is supported by 36 gRNAs in our library that directly target polymorphic sites. At these sites, impairment was observed almost exclusively in the presence of the allele that perfectly matches the gRNA (BY allele), with the impact of a SNP correlated with its distance from the PAM site (**Fig. 1E**, Supplemental Figure 11 and Supplemental Table 2). All gRNAs that bind near polymorphic sites were excluded from further analysis. As an additional validation, within the same flask, efficacious gRNAs targeting the same gene had correlated effect sizes (Spearman correlation = 0.263, pval = 6.0×10^{-10}; Supplemental Figure 12). Such a correlation was never observed across 1,000 permutations of the same data (Spearman correlations in permutations: mean = 0.002, min = -0.140 max = 0.160).

Because all genes targeted in this study have been previously annotated as essential under fermentative or respiratory growth conditions[33], we did not expect gRNAs to show positive effects. Thus, in addition to the ANOVA model, we employed a conservative effect size threshold to identify efficacious gRNAs. We only considered a gRNA further if its fitness effect was at least three standard deviations below a null distribution inferred from the data (**Fig. 1F**, Supplemental Data 9). Of the 8,013 CRISPRi perturbations in

this study, 1,635 met these criteria, targeting 817 distinct genes (Supplemental Figure 13). To validate this thresholding step, a single fitness estimate was obtained for each of these efficacious gRNAs by averaging all of its genotype and gRNA barcode replicates within an experiment. These average fitnesses were highly correlated between ATC1 and ATC2 (**Fig. 1G**, Spearman's correlation: 0.983), but not between ATC2 and CON (**Fig. 1H**, Spearman's correlation: 0.222).

Identification and characterization of background effects

We measured the fitness of each segregant-gRNA combination under both induced and uninduced conditions, making it possible to explicitly test for interactions between gRNAs and segregants using mixed effect linear models. At a False Discovery Rate of 0.05, 717 of the 1,635 efficacious gRNAs had significant background effects based on a gRNA-segregant interaction term (Supplemental Data 9). These results suggest that segregants do not show major differences in global CRISPRi efficacy, as ~56% of gRNAs had indistinguishable effects across segregants. gRNAs that showed background effects targeted 457 distinct genes, suggesting a large fraction of genetic perturbations interact with segregating loci (56% of all genes with efficacious gRNAs). Prior work across model systems has shown that 15%-32% of all genes show background effects[13]. However, excluding genes whose perturbation does not show measurable effects in any individual has been shown to increase estimates of the prevalence of background effects to as high as 74%-89%[13]. Thus, the proportion of genes that we identified with background effects is in line with other studies. For the remainder of this paper, we focus on a reduced data set, selecting a single gRNA for each of the 457 genes. Selections were made based on the significance of the background effect, prioritizing gRNAs with mapped loci (as described later).

We next identified factors that caused gRNAs to show variable effects across segregants. To compare fitness effects across diverse gRNAs, we first estimated the effect of each individual gRNA in each individual segregant using mixed effects linear models. Then, the mean effect of a gRNA was subtracted from these segregant-specific values. This assigned a single 'deviation value' to each segregant, representing the

difference between the effect of the gRNA in a specific segregant and its effect on average (**Fig. 2A**). This process was repeated for each gRNA, assigning a deviation value to every observed combination of gRNA and segregant (**Fig 2B**, Supplemental Data 10). We found a slight negative correlation between the deviation values (the effect of a gRNA in an individual segregant) and the baseline fitness of that segregant, with the same gRNA having a more detrimental effect in higher-fitness strains (R^2 = 0.157, pval < $1.0x10^{-100}$; **Fig. 2C**). In addition, the variance of a gRNA's deviation values correlated with its mean effect size, with more detrimental gRNAs showing larger variances (R^2 = 0.589, pval < $1.2x10^{-92}$; **Fig. 2D**, Supplemental Figure 14). These results indicate that the baseline fitness of a segregant (without any perturbation) and the mean effect of the gRNA can both impact background effects. Baseline fitness impacts the direction of a background effect in an individual segregant and mean gRNA effect impacts the magnitude of background effects across segregants.

We also determined the extent to which genetic factors segregating in the cross explain the deviations. Using replicated measurements among genotypes and gRNAs, we estimated broad sense heritability (H^2), which measures the total genetic contribution to a trait[42]. H^2 was 0.707 on average (sd = 0.06, min = 0.52, max = 0.87, **Fig. 3A**), implying that background effects identified here have a mostly genetic basis. All loci contributing to a background effect show epistasis with a perturbation (pairwise epistasis). However, some loci may interact not only with a perturbation but also each other (higher-order epistasis). The relative contribution of pairwise vs. higher-order epistasis underlying background effects can be estimated by comparing H^2 and narrow sense heritability (h^2), which measures only the additive genetic contribution to a trait[42]. Using genomic relatedness-based estimation, we found that h^2 was 0.176 on average (sd = 0.116, min = 0, max = 0.593). Thus, 75.1% (1 - 0.176/0.707) of the genetic basis of background effects could be attributed to higher-order epistasis on average. This measurement showed a skew towards higher-order epistasis, with around 93 genes (20.4%) having a $1 - h^2/H^2$ value above 0.9, and only a few genes (0.9%) having a value below 0.25 (**Fig. 3B**).

Hub loci control the phenotypic effects of many gRNAs

We next used linkage mapping to identify interacting loci that cause gRNAs to show variable effects across segregants. Specifically, we treated the deviations as trait values and mapped loci contributing to these deviations. By definition, loci identified in such scans show pairwise genetic interactions with gRNAs and play an additive role in the deviations. Because only 169 genotypes were present in our data set, statistical power was limited, and we were only able to detect loci for a subset of the gRNAs. Further, genetic variants that modify chromatin accessibility near gRNA targets may impact CRISPRi efficacy at those sites. Such variants were a minor contributor to our data set: we found only six genes with detected loci within 10 kb of their gRNA targets and excluded these cases from subsequent analyses (Supplemental Table 3). Across the remaining gRNA-targeted genes with background effects, we detected 248 loci at a significance threshold of $-\log_{10}(\text{pval}) \geq 4.2$ (**Fig. 3C,** upper panel, Supplemental Data 11). Of the targeted genes, 184 had a single mapped locus, 29 had two mapped loci, and two had three mapped loci.

Detected loci were not evenly distributed across the genome (**Fig. 3C,** lower panel). Across all loci, the BY allele produced higher deviation values roughly as often as the 3S allele did (**Fig. 3D**). On average, individual loci explained 20% of a deviation's total genetic basis (mean = 0.201, sd = 0.058, **Fig. 3E**). Of the 248 detected loci, 170 (68.5%) were grouped into one of four 'hubs' that interacted with a larger number of genes than expected by chance. Hubs on Chromosomes VII, X, XIII, and XIV showed linkage to 39, 30, 82, and 19 gRNA-targeted genes, respectively (**Fig. 3F,** Supplemental Table 4). The gRNA-targeted genes to which each hub showed linkage were largely distinct, with only 9 genes (5.3%) connected to multiple hubs. This result further supports the notion that the segregants do not show major differences in global CRISPRi efficacy. We also found four additional loci that are likely hubs but did not pass our detection threshold (Supplemental Figure 15).

In addition, we mapped three loci that contribute to the baseline fitness of the segregants (**Fig. 3C**). This low number of detected fitness loci was due to the small

number of segregants tested. The hubs on Chr X and Chr XIII each overlapped one of these fitness loci, implying that the same locus can affect both baseline segregant fitness and the effects of genetic perturbations. We found that 45% of all detected loci with effects on deviations were mapped to these two hubs. However, as we show below, these loci had different impacts on baseline fitness and CRISPRi perturbations.

Properties of hub loci that interact with many genetic perturbations

We found that gRNAs that target different genes and interact with the same hub show similar epistatic relationships with the hub. At the Chromosome X and XIII hubs, which both affect baseline fitness, the 3S allele has a beneficial effect on fitness in the absence of gRNA induction (**Figs. 4A** and **B**, mean difference = 0.097 and pval = $2.3x10^{-07}$ for Chr X, mean difference = 0.107 and pval = $7.1x10^{-09}$ for Chr XIII). However, when nearly all gRNAs are induced, the 3S allele is instead detrimental, although the magnitude of the effect varies depending on the gRNA (**Figs. 4A** and **B**, mean difference = -0.039 and pval = $3.6x10^{-18}$ for Chr X, mean difference = -0.043 and pval = $8.8x10^{-75}$ for Chr XIII). The only exceptions were three gRNAs that interact with Chr X but do not show sign epistasis, which may be due to them interacting with a different genetic variant than the other 27 gRNAs. Excluding these three cases, both the Chr X and XIII hubs show a similar pattern, where the 3S allele is beneficial in the absence of any gRNA activity, and only individuals carrying the 3S allele are strongly affected by the gRNA. This suggests that the causal genes underlying these two hubs may be impaired or nonfunctional in BY.

The hubs on Chromosome VII and XIV exhibit a different pattern, with the BY and 3S alleles only showing major differences after gRNA induction. In the absence of gRNA activity, the Chr VII hub shows no significant effect on fitness (**Fig. 4C**, mean difference = -0.011, pval = 0.567). However, when gRNAs interacting with this hub are induced, the 3S allele of the locus is significantly more beneficial than the BY allele (**Fig. 4C**, mean difference = 0.035, pval = $5.4x10^{-29}$). This is true for all 39 of the gRNAs that interact with the Chr VII locus (Supplemental Figure 16). At the Chr XIV hub, the 3S allele confers a slight, nonsignificant benefit in the absence of gRNA induction (**Fig. 4D**,

mean difference = 0.003, pval = 0.864), but a large benefit after induction of any of the 19 interacting gRNAs (**Fig. 4D**, mean difference = 0.039, pval = $6.7x10^{-17}$). In summary, at both the Chr VII and XIV hubs, the fitness of individuals carrying the 3S alleles of these loci was substantially less impacted than the fitness of individuals carrying the BY alleles. This suggests that 3S carries alleles of these two hubs that may impart higher robustness to certain perturbations [43].

We hypothesized that the consistent interactions between hubs and different gRNAs were due to common underlying mechanisms. Using the CellMap genetic interaction network[44] and gRNA-targeted genes to which a hub showed linkage, we identified multiple significant functional enrichments (**Fig. 4E** and **4F**). At p-values $\leq 1.0x10^{-06}$, three hubs (VII, X, and XIV) were enriched for DNA replication and repair and two hubs (VII and X) were enriched for mitosis. Several hubs also had unique enrichments: VII for protein turnover and cell polarity (p-values $\leq 1.0x10^{-05}$); X for glycosylation and MVB sorting and RIM signaling (p-values $\leq 1.0x10^{-04}$); XIII for rRNA/ncRNA processing, vesicle traffic, and ribosome biogenesis (p-values $\leq 1.0x10^{-06}$); and XIV for transcription (p-values $\leq 1.0x10^{-04}$). These data suggest that each hub likely acts through a distinct combination of cellular processes to influence the phenotypic effects of genetic perturbations. The Chr X and XIII hubs and the Chr VII and XIV hubs each show similar relationships between alleles, indicating that different mechanisms can produce similar allelic relationships at distinct hubs (**Fig. 4E** and **4F**).

Discussion

We found that the effect of a genetic perturbation depends on baseline fitness (i.e. the fitness of an individual in the absence of a perturbation), the mean fitness impact of the perturbation across individuals, and the specific alleles carried by an individual at interacting loci. Our findings do not appear to be driven by major differences in global CRISPRi efficacy, as most efficacious gRNAs do not show background effects and hubs interact with different sets of gRNAs. Further, variation in chromatin accessibility near gRNA target sites cannot explain our results because the vast majority of loci that contribute to deviations did not occur near gRNA target sites. Based on the heritability

explained by detected loci, we can infer that there are likely to be multiple interacting loci contributing to each background effect. The large number of perturbations in our data allowed us to determine the degree to which segregating loci interact with different perturbations. While some loci interact with only one perturbation, we identified several hub loci that interact with many perturbations, implying that responses to many different perturbations can have a shared genetic basis.

Perturbation-locus genetic interaction networks like the one described here have been historically difficult to map because they require introducing and quantitatively phenotyping a large number of genetic perturbations in a large number of individuals. These previously unmeasured networks shape how individuals respond to different perturbations. Features of these networks likely impact how genetic background shapes the effects of perturbations, leading to incomplete penetrance, variable expressivity, and differences in robustness to perturbations. They might also control which mutations are deleterious or beneficial, thereby influencing evolutionary trajectories. Similar to more well-described networks between genetic perturbations alone[45] or between natural variation alone[34,46], perturbation-locus networks appear to contain many nodes with few interactions and few nodes with many interactions. Further study is needed to determine how these different types of genetic networks relate to each other.

The measurable perturbation-locus network will depend on the genetic variation segregating among examined individuals, and this will necessarily be a subset of the full perturbation-locus network within a species. To simplify this problem, we focused on a systematic analysis of the perturbation-locus network between two haploid strains, BY and 3S. The current study could not fully map the network of interactions between perturbations and loci in the BYx3S cross due to the modest number of assayed segregants. This sample size was chosen to facilitate examination of a genome-scale gRNA library, with the expectation that many gRNAs would not be efficacious under our assay conditions. These results now enable follow-up studies with more segregants and fewer gRNAs. Such studies will facilitate the detection of many more loci that interact with perturbations and thereby comprehensive mapping of the network of interactions

between perturbations and loci in this cross. The ultimate goal of this line of study would be to predict the outcome of a perturbation in an unseen genetic background with potential lessons for therapeutic genome editing, crop and livestock improvement, and engineering of cells and microorganisms for industrial applications.

Acknowledgments

We thank Cara Hull, Zach Krieger, and Daniel Lusk for feedback on a manuscript draft. This work was supported by startup funds from the University of Southern California to I.M.E., grant R35GM130381 from the National Institutes of Health to I.M.E, and grants R01AI164530 and R01HG010378 from the National Institutes of Health to S.F.L.

Competing Interests

The authors declare no competing interests.

Author Contributions

J.J.H., T.M., I.G., M.N.M., K.R.R., S.F.L., and I.M.E. conceptualized this project. T.M., K.R.R., L.M.S., and S.F.L. generated the gRNA libraries. J.J.H., T.M., I.G., M.N.M., K.R.R., C.N.V., C.W., and T.R. performed the experiments. J.J.H. and I.M.E. analyzed the data. D.M. provided modified Fit-Seq code. J.J.H. generated the figures. J.J.H., T.M., I.G., S.F.L., and I.M.E. wrote and edited the manuscript.

Figures

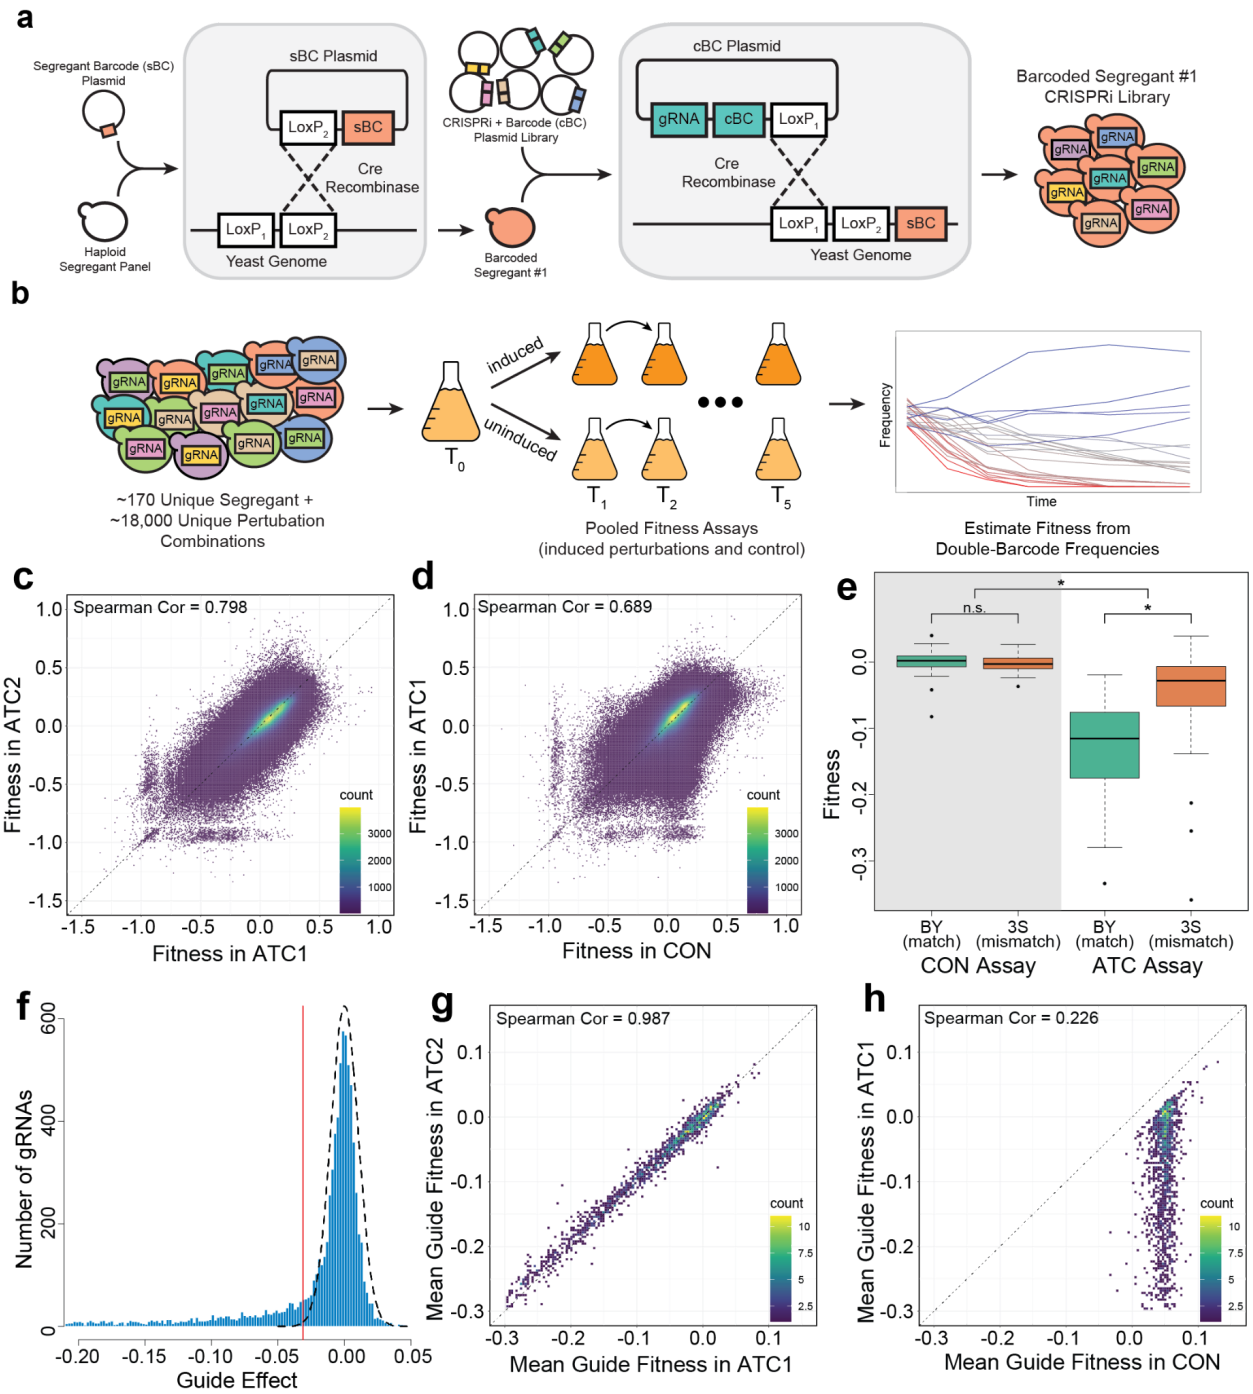

Figure 1: Generating and phenotyping a panel of double-barcoded segregants carrying CRISPRi perturbations. a) Design and generation of the haploid panel. b) Fitness assay and phenotyping of perturbed segregants. c) Correlation plot of estimated

fitness values for all lineages across the two fitness assay flasks (ATC1 and ATC2). d) Correlation plot of estimated fitness values for all lineages across one experimental fitness assay flask and one control flask (ATC1 and CON). e) Average fitness values of 36 gRNAs that directly overlap BY/3S polymorphisms, across both experimental and control flasks. gRNAs were designed for the BY genome, so segregants with 3S alleles have a mismatch at the gRNA binding site. Asterisks and 'n.s.' respectively indicate pvalues $<1x10^{-5}$ and >0.05 via t test. f) Distribution of all mean gRNA effects as calculated by the mixed effects linear model. The estimated null distribution is shown by the dotted line, and the mean effect threshold is shown by the vertical red line. g) Correlation plot of average fitness values for all gRNAs with mean effects across the two fitness assay flasks (ATC1 and ATC2). h) Correlation plot of average fitness values for all gRNAs with mean effects across one experimental fitness assay flask and one control flask (ATC2 and CON).

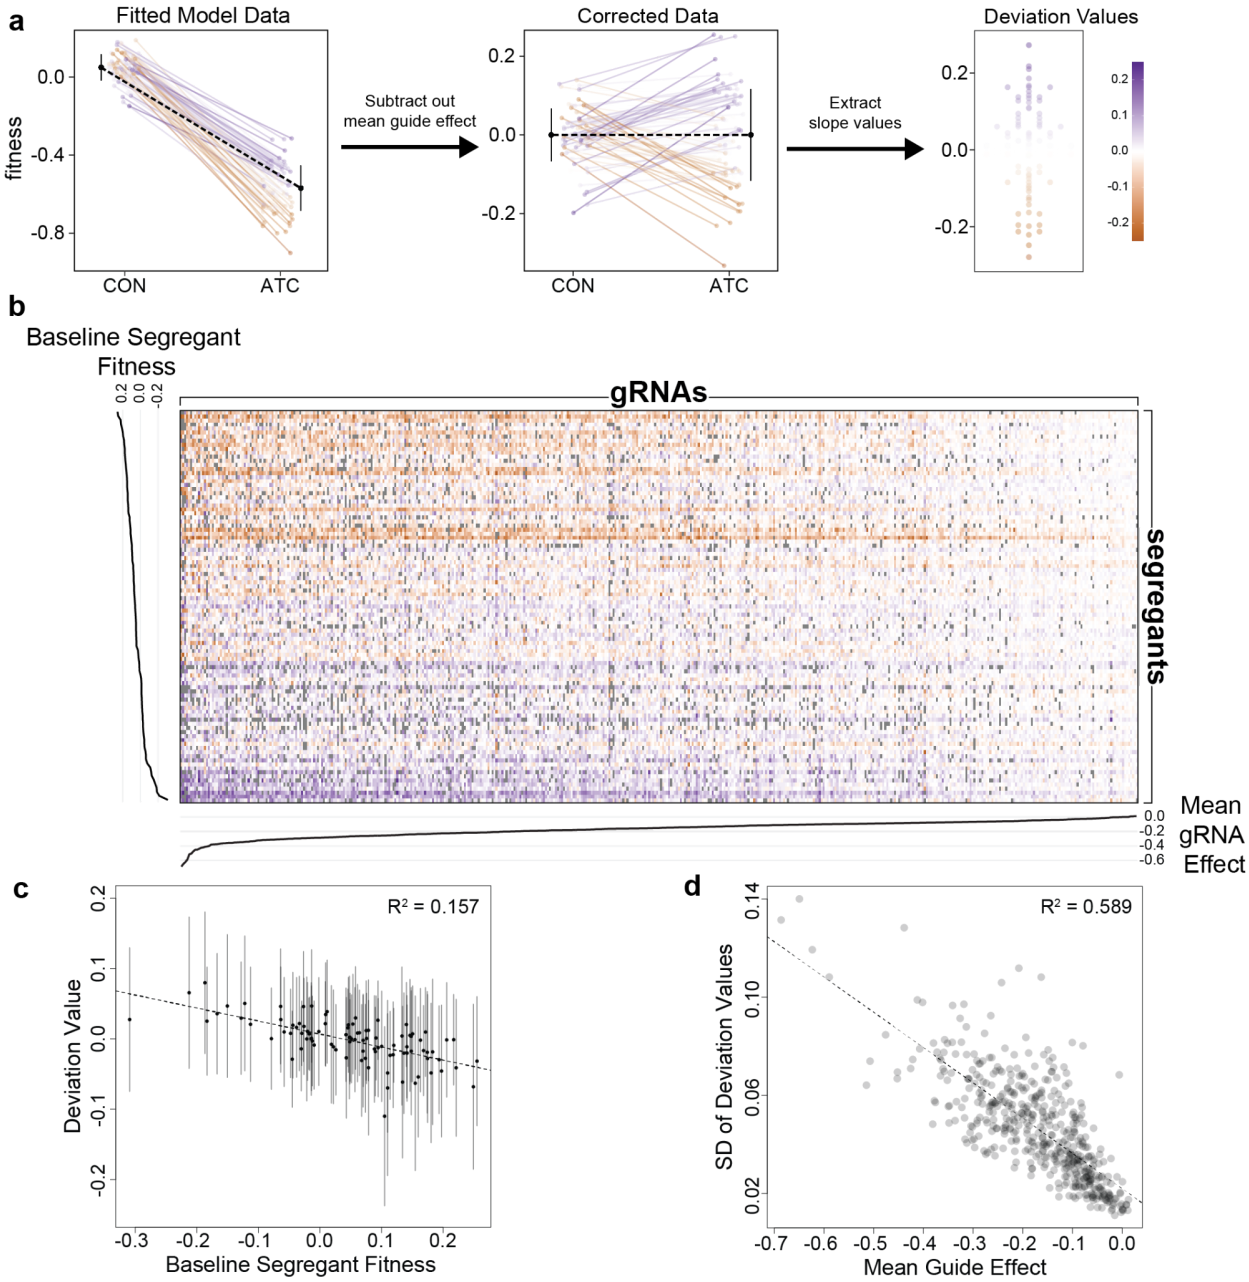

Figure 2: Calculating deviation values for all segregant-gRNA combinations. a) Schematic representation of how deviation values are calculated from fitness. b) Heatmap of all deviation values. Segregants are shown along the y-axis, ordered by mean fitness, and gRNAs are shown along the x-axis, ordered by mean effect size. Missing combinations are shown in gray. Only rows and columns with 30% or less missing data are shown. c) Relationship between baseline segregant fitness and deviation values for that genotype. Dots indicate average deviation value, and lines

indicate 2 standard deviations. d) Relationship between mean gRNA effect and the standard deviation (SD) of deviation values for that gRNA, with the best-fit line shown.

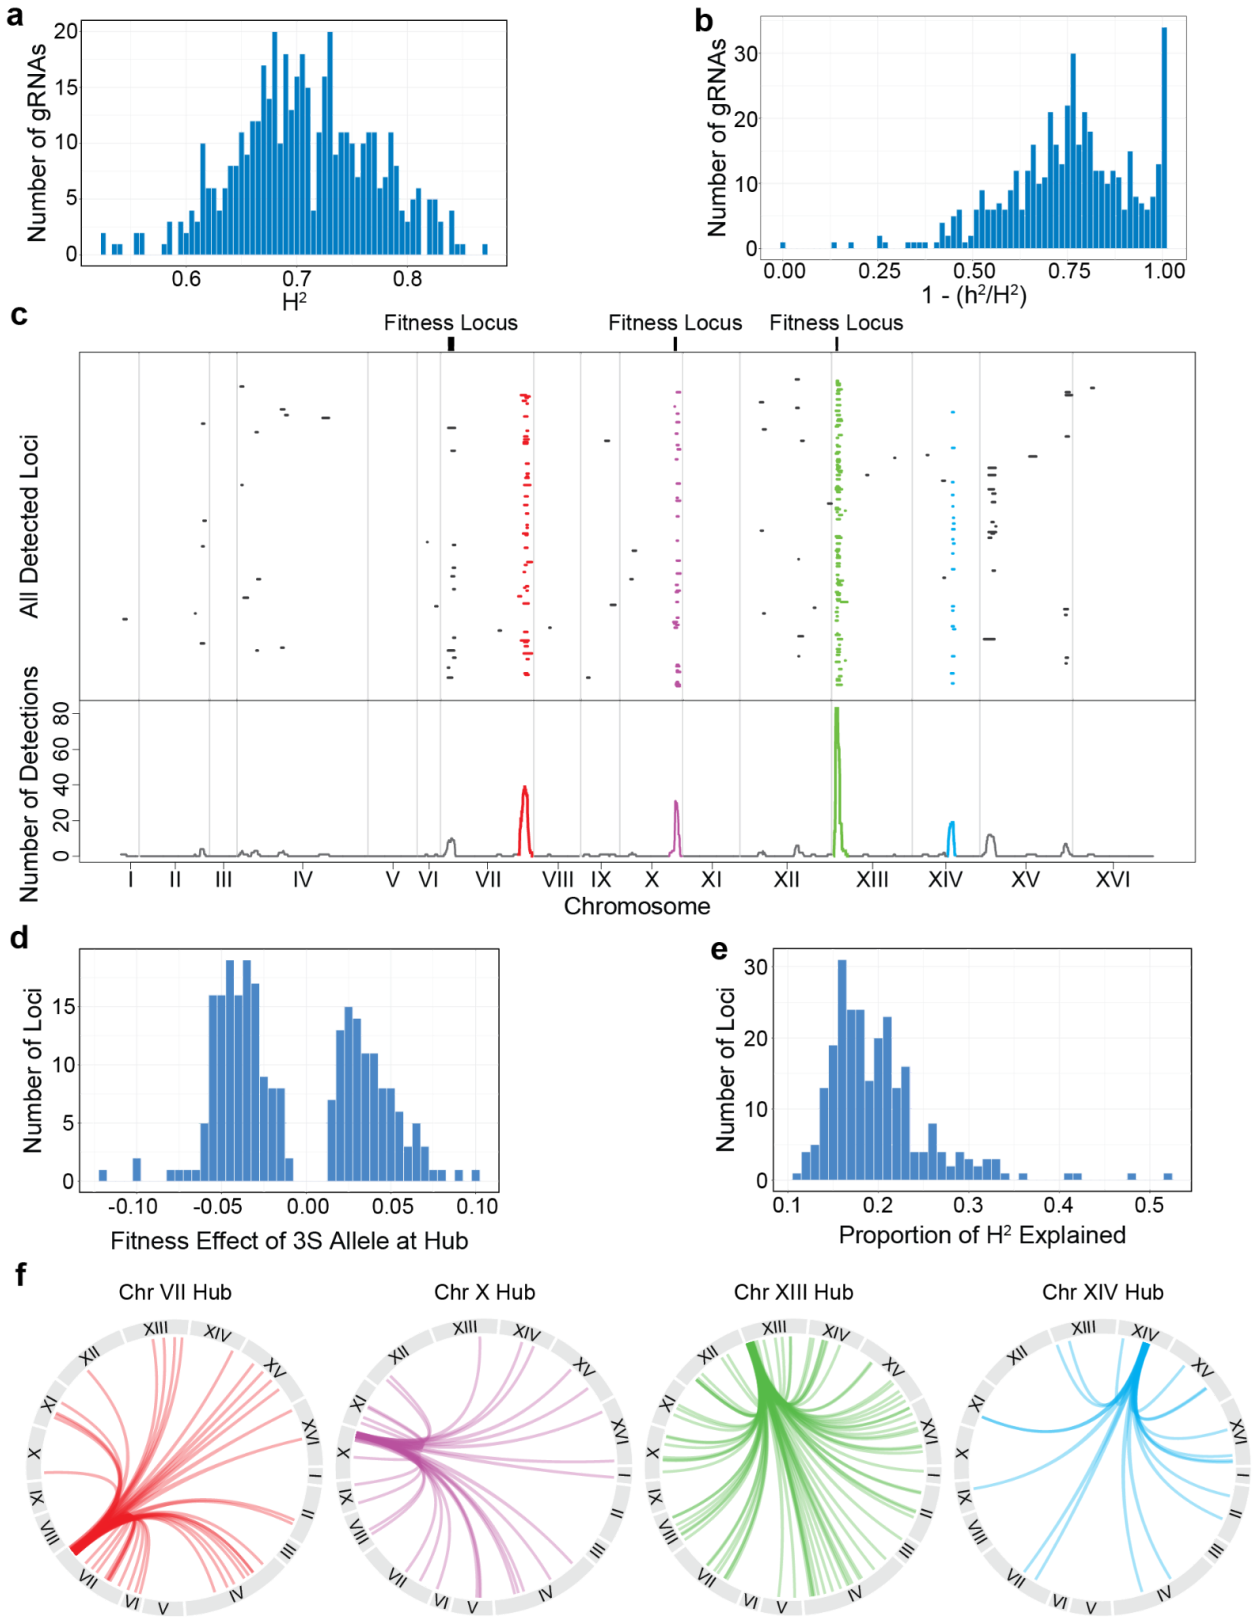

Figure 3: Using deviation values as traits in linkage mapping. a) Histogram of

broad sense heritability (H^2) values for all gRNAs with background effects. b) Histogram of the ratio between narrow sense heritability (h^2) and H^2 for all gRNAs with background effects, shown as ($1 - h^2/H^2$). c) Linkage mapping on deviation values. Upper panel: Visualization of all 2x-$\log_{10}$(pval) drops for the gRNAs with significant peaks. Each row is a unique gRNA. Locations of major fitness loci are indicated by black bars. Lower panel: Number of overlapping 2x-$\log_{10}$(pval) drops at each nucleotide position along the genome. Regions where more intervals overlap than expected by chance (hubs) are shown in unique colors. Vertical lines indicate chromosome boundaries. d) Histogram of all effect sizes for the detected loci. e) Histogram of the proportion of heritability explained by each detected locus, calculated as R^2/H^2. f) Interaction plots connecting gRNA location to the corresponding 2x-$\log_{10}$(pval) drop for each of the four hubs.

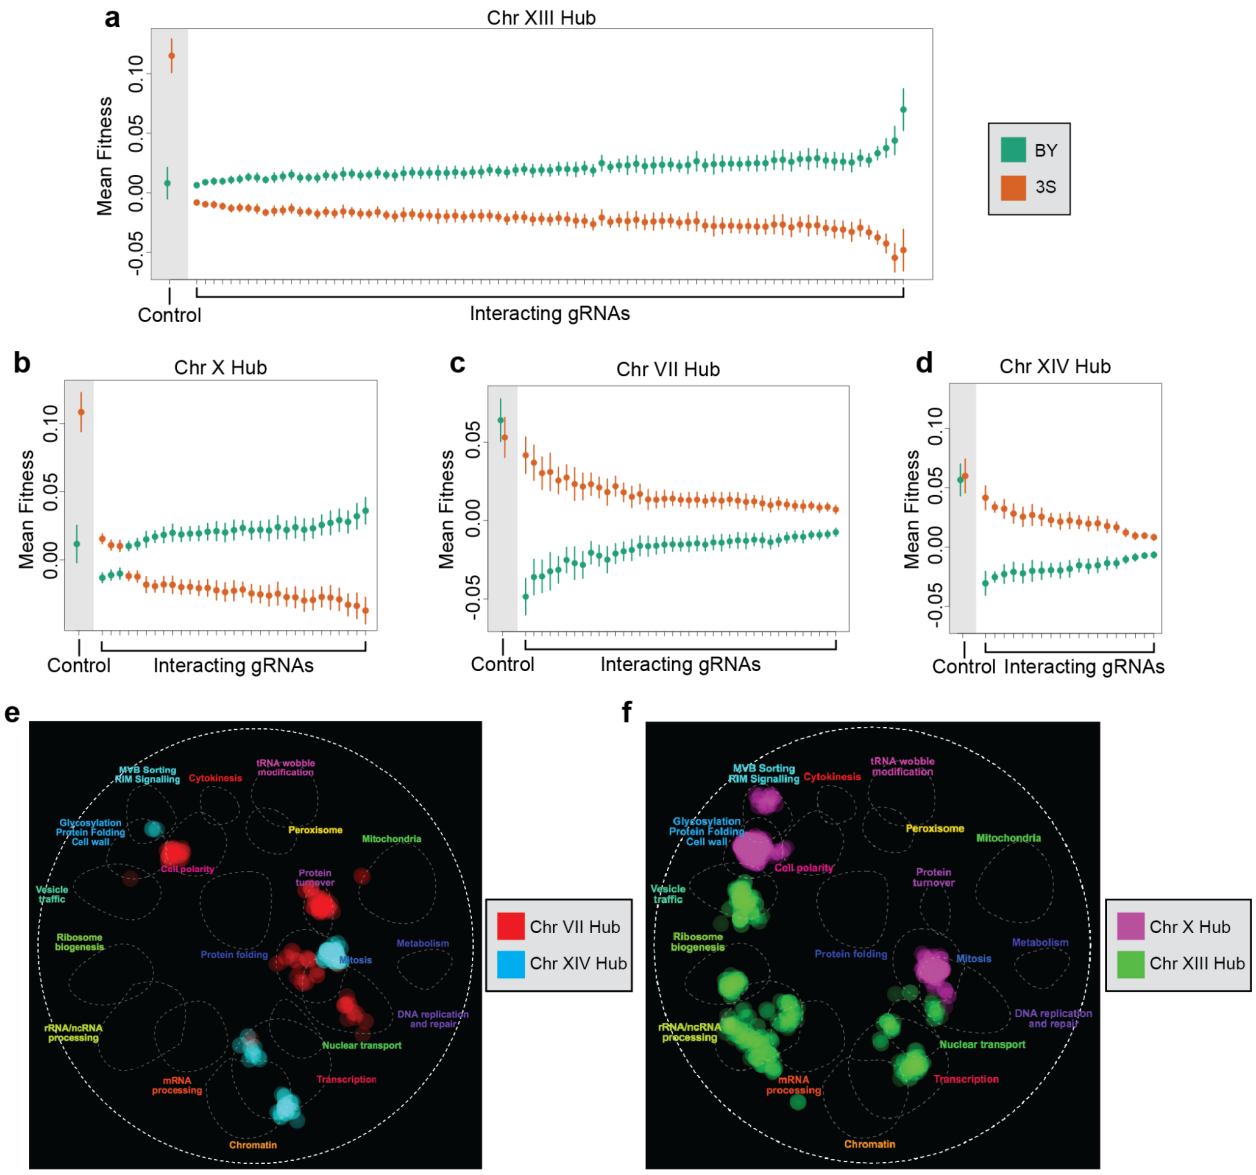

Figure 4: Features of and comparisons between hubs. a) Average effect of the Chr XIII hub. Genotypes with the BY allele are shown in green and genotypes with the 3S allele are shown in orange. Dots indicate the mean fitness and bars indicate one standard deviation. The effect of the hub in the absence of gRNA activity, calculated via control gRNAs, is shown in the gray box. The effect of the hub in the 83 interacting gRNAs is shown to the right. b) Average effect of the Chr X hub. c) Average effect of the Chr XIV hub. d) Average effect of the Chr VII hub. e) CellMap visualization of the interaction network for the Chr VII and XIV hubs. f) CellMap visualization of the interaction network for the Chr X and XIII hubs.

Methods

Generation of barcoded haploid segregants

The barcoded BYx3S *MATα* segregants in this study were previously reported[34,35]. In brief, all segregants were tetrad dissected from a cross between *fcy1*Δ *flo8*Δ *flo11*Δ *ura3*Δ versions of BY and 3S that both carried the same genomic landing pad at the neutral *YBR209W* locus. The *fcy1*Δ *ura3*Δ deletions provide counterselectable markers, while the *flo8*Δ *flo11*Δ deletions eliminate cell-cell adhesion phenotypes, such as flocculation. The genomic landing pad contains multiple partially-crippled loxP sites in close proximity, in addition to a galactose-inducible Cre recombinase. To barcode segregants, we individually transformed them with a plasmid library containing 20-nucleotide random barcodes, integrated plasmids into the landing pad by Cre/loxP-mediated recombination, and obtained random integrants. The barcodes present in the segregants, as well as segregant genotypes, were then determined via multiple Illumina HiSeq 2500 lanes[35]. Briefly, six HiSeq lanes were used to perform whole-genome sequencing on each individual segregant, with reads mapped against the S288c reference genome R64-2-1_20150113 using BWA[47]. SAMTOOLS[48] was used to convert alignments to pileups, from which high-confidence SNP lists were generated. Segregant barcodes were identified using a separate HiSeq 2500 lane, in which the barcode was PCR amplified from each segregant using specific primers. After being isolated from sequencing reads, barcodes were clustered using Bartender[49], with clusters making up greater than 5% of all reads from a segregant being assigned as its true barcode.

CRISPRi plasmid construction

The pE32 CRISPRi plasmid was constructed using plasmids pKR387, pKR482, and pKR919. An annotated plasmid map of pE32 is available in the supplementary information. To construct the pE32 CRISPRi plasmid, a C to T synonymous mutation was first introduced at amino acid position 1319 (Glycine, GGC -> GGT) of dCas9 on pKR387 to remove an AscI recognition site using the QuikChange II site-directed mutagenesis kit (Agilent #200523). The resulting product was transformed into NEB 10-beta cells using a standard heat shock protocol and transformants were selected on

LB agar plates with 100 µg/mL of carbenicillin. From this pKR387 plasmid with the mutated dCas9, a DNA fragment containing a *GPM1* promoter, *tetR* gene, *GPM1* terminator, *TEF* promoter, dCAS9, MxiI repressor domain, and *CYC1* terminator was amplified by PCR. The resulting product was inserted into linearized pKR482 using NEBuilder HiFi DNA assembly kit (NEB E5520S) and transformed into 10-beta cells. Transformants were selected on LB agar plates with carbenicillin, and the resulting plasmid was named pSL25. This pSL25 plasmid was linearized with AscI and ClaI and ligated with a double-stranded oligonucleotide containing a partially crippled Lox66 site using T4 DNA ligase (NEB M0202S). The ligation product was transformed into 10-beta cells, selected on LB agar plates with carbenicillin, and named pSL55. Finally, a DNA fragment containing *RPR1* promoter, tetO sequence, and a hammerhead ribozyme sequence was amplified from pKR919 by PCR. The hammerhead ribozyme was included because preliminary experiments found that it improves gRNA efficacy in CRISPRi. The PCR product was inserted into pSL55, which was linearized with SpeI and BspQI, using NEBuilder HiFi DNA assembly kit. The resulting product was transformed into 10-beta, selected on LB agar plates with carbenicillin, and named pE32.

Generation of the CRISPRi library

All CRISPRi gRNAs used in this paper were previously generated[32,33]. To generate the CRISPRi plasmid library for this paper, ~20,000 arrayed yeast strains[33], each containing a unique chromosomally-encoded gRNA, were plated onto YPD agar plates with 300 µg/mL of Hygromycin B in a 384-well format using Singer ROTOR. The yeast cells were grown overnight at 30°C. Although these yeast strains are genetically identical except for the 20 bp gRNA, heterogeneity in growth was observed across yeast strains. To minimize bias in gRNA frequency, equal amounts of each overnight colony were transferred into 200 µL of water using the ROTOR. Cells containing essential gRNAs targeting essential and non-essential genes were then pooled separately and spun down at 3000 rpm for 15 minutes prior to genomic DNA extraction. Genomic DNA was extracted using MasterPure yeast DNA purification kit (Lucigen MPY80200). A DNA fragment containing a gRNA with the tracrRNA scaffold sequence, *URA3* promoter, 5'

half of *URA3*, an artificial intron, and an I-SceI cut site was amplified from the pooled genomic yeast DNA. The reverse primer used to amplify this region contained a random 20 bp sequence, which was used to mark each gRNA with an unique barcode. To minimize PCR bias, the amount of DNA required for amplification was calculated such that each gRNA is represented by ~1,000 DNA molecules. A total of 330 ng and 270 ng of yeast genomic DNA was used as template for amplification of ~11,000 non-essential and ~9,000 essential gRNAs, respectively. The resulting product was inserted into pE32, which was linearized with BspQI and AscI, using NEBuilder HiFi DNA assembly kit. The resulting product was transformed into 10-beta cells and selected on LB agar plates with carbenicillin. To ensure high barcode complexity, ~110,000 and ~90,000 transformation colonies were scraped and pooled prior to plasmid extraction. While the CRISPRi library used in this study initially had a library size of ~20,000 gRNAs, a previously unknown design issue with the non-essential plasmid library affected growth on selective media, resulting in substantial depletion of non-essential gRNAs in the T0 pool. As a result, all non-essential gRNAs were excluded from analyses, which had a minor impact on our total data. Raw data from non-essential gRNAs is included in the Supplementary Information.

Linking barcodes to gRNAs in the CRISPRi library

The CRISPRi library was initially generated using fully randomized barcode sequences, necessitating a linkage step to connect known gRNAs with their barcodes. To link barcodes and gRNAs, the full plasmid library was sequenced using a PCR-free library preparation method and Oxford Nanopore MinION. The SQK-LSK109 protocol was used to prepare this sequencing library, with elution steps after bead cleanup performed at 37°C for 10 min to increase yield. Sequencing was performed on R9.4.1 Chemistry MIN-106 flow cells. Basecalling was performed on the USC Center for Advanced Research Computing Discovery cluster. We ran Guppy v6.0.1 with the configuration dna_r9.4.1_450bps_sup.cfg on a single node with 16 threads and a V100 GPU. After basecalling, we first mapped the gRNA portion of a read to its best match among the known gRNAs and then mapped the barcode portion of read to its best match among the known gRNA barcodes from barcode sequencing of the T0 pool. Fuzzy string

matching was done in Python using the process.extractOne() function in the rapidfuzz library[50].

Transformation of haploid strains with the CRISPRi library

All strains were individually transformed with the full CRISPRi library, using a modified version of the standard lithium acetate protocol[51] that scaled up all volumes by a factor of 10. Roughly 8.0×10^8 cells and 30 µg of plasmid library were used for each transformation. Immediately after transformation, cells were grown in 3.4 mL of YP + Galactose liquid media for 18-20 hours to induce integration of the plasmid while minimizing cell division. After induction, cultures were grown in SC - URA liquid media in order to select for cells that successfully integrated the plasmid. Since SC - URA is not immediately lethal to non-integrated cells, cultures were first grown in 60 mL of selective media for 48 hours. Next, 20 mL of the culture was combined with 40 mL fresh SC - URA to perform a second setback. This two-step enrichment was done to ensure that as many cells in the cultures were true integration events and to maintain the barcode and gRNA diversity among yeast transformants. After enrichment, 12.5 mL of the stationary phase culture was saved and frozen as a freezer stock.

Fitness Assay

One full freezer stock from each transformed segregant (roughly 1.25×10^8 cells) was used for the fitness assay, regardless of transformation efficiency. This value was chosen assuming a roughly uniform distribution of gRNA frequency within each segregant, leading to an expectation of $\geq 10^4$ cells per lineage. All freezer stocks were thawed and grown to stationary phase in SC - URA liquid media. Then, 10 mL from each culture (roughly 2×10^9 cells) was combined to generate a 2.5 L T0 pool. Multiple T0 cell pellets were frozen for fitness assays and DNA extraction. To begin the assay, 250 mL of the T0 pool (roughly 5×10^{10} cells) was thawed and inoculated into 750 mL fresh SC - URA liquid media. To induce gRNA expression, ATC was added to experimental flasks at a final concentration of 250 ng/mL. ATC was not added to the control flask. To generate the first time point (T1), flasks were grown for 24 hours at 21°C with shaking. The majority of the culture was harvested and saved as a frozen cell

pellet, while the remaining 250 mL was inoculated into 750 mL of fresh media and grown for another 24 hours. This process was repeated to generate the remaining time points (T2 through T7). The ATC1, ATC2, and CON flasks were grown, harvested, and diluted in parallel, under identical growth conditions.

Library preparation for barcode sequencing

DNA was extracted from frozen cell pellets using the Zymo Research Quick-DNA Fungal/Bacterial Midiprep Kit, with roughly 120 µg extracted from each time point. In order to separate the barcode region from the rest of the genome, extracted DNA was digested for 18 hours with the restriction enzyme I-SceI. The ~250 bp barcode region was then isolated from the genomic DNA via a standard AMPure XP bead cleanup protocol. Briefly, cleanup was performed by incubating DNA with 0.4X beads for 15 minutes, taking the supernatant and incubating with 0.8X beads for 15 minutes, then discarding the supernatant and washing DNA off the beads. 70% Ethanol washes were performed between each step. Total DNA after bead cleanup ranged from around 20 µg to 30 µg. The resulting DNA was first amplified with a 5-cycle PCR with Phusion polymerase, using 200 ng of input DNA and primers specific to the double barcode region. These primers were:

1)AATGATACGGCGACCACCGAGATCTACACNNXXXXNNACACTCTTTCCCTACACGAC

2)CAAGCAGAAGACGGCATACGAGATNNXXXXNNGTGACTGGAGTTCAGACGTGTGCTCTTCCGATCT

The N's in the above sequence represent the random unique molecular identifiers (UMIs) used to account for and eliminate PCR duplicates that could appear in later amplification steps. The X's represent Illumina multiplexing indices used to enable eventual pooling of multiple libraries onto a single sequencing lane. To ensure consistency of the PCR protocol, each library was split in half on this step, with each aliquot amplified with a different set of multiplexing indices. Final results and data output were similar across all index combinations for each library. All primer sequences are available in the Supplementary Information.

After the first PCR step, all products were pooled together and purified with Zymo Research DNA Clean & Concentrator-5 columns, using 150 µL of PCR product per column and eluting into 30 µL. The purified product was used to set up a second 26-cycle PCR with Phusion polymerase, using 15 µL of template. Universal Illumina primers P1 and P2 were used for this reaction. All PCR products from this reaction were pooled and put through another Clean & Concentrator-5 column, with 375 µL of product per column, and eluted into 35 µL. The cleaned product was re-pooled, and the 200-300 bp region was extracted via agarose gel electrophoresis. DNA was recovered from the gel using Zymo Research Zymoclean Gel DNA Recovery Kits and Clean & Concentrator-5 columns. The purity and concentration of the final library was assessed using Life Technologies Qubit 2.0 and Thermo Fisher NanoDrop ND-1000 Spectrophotometer. The structure, sequence, and diversity of each library was verified via Sanger sequencing.

Barcode sequencing

150 bp paired-end sequencing was performed for all time points, on either a NovaSeq 6000 or a HiSeq 4000 platform. Up to three multiplexed time points were included on each NovaSeq lane, with each time point being allocated roughly 800 million reads regardless of platform. Each sequencing lane was performed with 25% PhiX spike-in to ensure nucleotide diversity during the run, since the majority of the amplicons consisted of fixed bases. Sequencing data was analyzed with custom-written code in Python and R. For the initial processing of large files, USC's Center for Advanced Research Computing Discovery cluster was used. Reads were sorted based on several multiplexing indices unique to each time point. Forward reads were used to extract gRNA barcodes and reverse reads were used for genotype barcodes, with reads filtered out if the expected fixed nucleotides did not appear immediately downstream of either barcode. These downstream sequences are CCCGAGTCGCGATAA and TACCGTTCGTATAGG for the gRNA and segregant barcodes, respectively. Reads were also filtered out if the first 35 bases of either the forward or reverse read had an average Illumina quality score below 30.

To minimize the effects of sequencing errors and random mutations on barcode counts, very similar barcodes were clustered together and treated as the same sequence. Clustering was performed using both Bartender[49] software and the rapidfuzz package[50], and was done separately for genotype and gRNA barcodes. First, a list was generated of every unique barcode detected across all time points by the extraction method described above. Consensus sequences were identified from this input using Bartender, with the merging threshold disabled (-z -1) to avoid frequency-based clustering. Each consensus sequence was then compared against the entire list of verified reference barcodes using rapidfuzz, with similarity scores of 90 or higher considered a match. Reference lists for genotype barcodes were obtained through separate HiSeq 4000 lanes on pooled segregant barcodes[35], in addition to Sanger sequencing on individual segregants. Reference lists for gRNA barcodes were obtained from separate Nanopore sequencing lanes (see section “Plasmid Library Preparation” for more detail).

After clustering, the frequency of each double barcode in each time point was calculated. Barcode amplicons undergo PCR before being sequenced, and PCR duplicates can cause significant bias if they are not removed when calculating frequencies. To account for this, a random 4-mer UMI was added to every multiplexing index during the library preparation, for a total of 8 random nucleotides per fragment. These UMIs are added on an early step of PCR. Because the total number of possible 8-nucleotide combinations (4^8 = 65,536) is significantly higher than the expected number of reads for any individual double-barcode fragment, it is unlikely that the same DNA template will contain the same UMIs by chance. This can facilitate estimation of barcode frequencies by counting unique UMIs detected for each double barcode. However, UMI-containing primers were not fully removed from the sample by PCR cleanup kits due to their large size (> 50 bp), causing unique UMIs to be added in all stages of PCR, not just the first few cycles. This meant that not every PCR duplicate could be identified via UMIs and removed from our frequency measurements.

Removing PCR chimeras

The double barcode constructs in this study consist of two highly variable barcodes separated by a region of invariant nucleotides, which can enable the production of spurious double barcodes (chimeras) during PCR[52]. To quantify and account for chimeras, we chose ~10 individual segregants and extracted their barcodes in isolation from each other. These individually-prepared segregants were sequenced via the methods described above on the HiSeq 4000 platform. Since segregants were completely isolated during this library preparation, there was no opportunity for PCR chimeras to be generated. This allowed us to construct a high-confidence list of the true gRNA barcodes present in those 10 segregants, which in turn identified double barcodes from the fitness assay data that were likely chimeras. The number of chimeric double barcodes per segregant from the fitness assay data was consistently <10%. To more broadly correct these data for chimeras, a linear model was fit for the frequency of each known chimera as a function of the total frequency of the corresponding segregant and gRNA barcodes in that sequencing lane: *(frequency of chimera) ~ (total frequency of segregant barcode) + (total frequency of gRNA barcode)*. This linear model was fit separately in four early time points from the fitness assay, and the coefficients were averaged across the four models. The final model was used to correct every double barcode's calculated frequency in each time point for its expected number of chimeras, with corrected double barcode frequencies set to a minimum of zero.

Fitness estimation

Before estimating fitness, differences in coverage between time points (Supplemental Figure 7) were accounted for by normalizing read counts. This was done by dividing each read count by the total number of reads from that time point. All read counts were then multiplied by the same arbitrary value (450,000,000, roughly the number of reads present in T0). PyFitSeq[40] was then used to estimate fitness of each double barcode lineage, which consisted of normalized read counts from five separate time points. The maximum number of iterations was set to 200, and all default PyFitSeq settings were used otherwise. Lineages were removed from the data set if their log likelihood score was one interquartile range below the first quartile, or if <5 reads were present at the first time point (T0), as both cases led to inaccurate fitness estimates. Additionally,

segregants with <100 associated gRNAs were removed from the data set entirely, as were gRNAs with <2 associated segregants.

Identification and quantification of gRNA effects

The mean effects and background effects of gRNAs were identified using the fitness values. Because the distribution of fitnesses differed between experimental and control flasks, normalization was performed by first setting a conservative threshold for gRNAs that do not have effects in our data set. Neutral gRNAs were identified by averaging the fitness of all lineages containing that gRNA in both ATC and CON flasks. Any gRNAs where the average fitness in ATC was equal or higher than the average fitness in CON were assumed to be neutral and used for normalization. Normalization was then performed by taking the difference in median fitness among the neutral gRNAs between ATC and CON flasks, with both ATC flasks normalized to CON.

After both flasks were normalized, mixed effects linear models were used to determine if gRNAs had a significant mean effect across all segregants, and if any background-specific effects were present. These analyses were conducted separately for ATC1 and ATC2, in each case using comparison against the CON data. Mixed effects models were fit using the lme() function in the nlme package in R. This was done for each gRNA individually. The model *fitness ~ segregant + gRNA + error* was applied to all lineages from both flasks corresponding to a specific perturbation, with *segregant* treated as a random effect and *gRNA* treated as a fixed effect. All lineages from the control flask were assumed to lack the perturbation (*gRNA* = 0), and all lineages from the experimental flask were assumed to have the perturbation present (*gRNA* = 1). If the addition of a *genotype:gRNA* interaction term significantly improved model fit, the model *fitness ~ segregant + gRNA + segregant:gRNA + error* was used for that gRNA instead. Regardless of model, if the p-value for the *gRNA* term was significant after Benjamini-Hochberg multiple testing correction, the perturbation was considered to have a potential mean effect. Among the gRNAs with mean effects, if the p-value for the *gRNA:genotype* interaction term was also significant after Benjamini-Hochberg multiple testing correction, the gRNA was considered to have a background-dependent effect.

Mean gRNA effects were estimated by extracting the *gRNA* term coefficient from the appropriate linear model described above. This was done for all gRNAs, regardless of their effects. Analysis of the control data indicated that some gRNAs had leaky expression, having similar fitness effects across both experimental and control flasks. The linear models described above also removed these gRNAs from the data set, as the mean gRNA effect was determined by comparing the same gRNA between ATC and CON flasks. We found that the p-value threshold described above was not sufficiently stringent to identify efficacious gRNAs. To set a conservative threshold for efficacious gRNAs, we took the subset of gRNAs with neutral or beneficial effects and simulated a normal distribution from this data. Perturbations were only considered to have a mean effect if they were also at least three standard deviations below the median of this simulated distribution.

In the case that a background-specific effect was detected using the method described above, segregant-specific model coefficients were extracted using the predict() function in R. This provided an estimated gRNA effect for every segregant in which that gRNA occurs. The overall mean gRNA effect was then subtracted from these values, giving us segregant-specific deviation values. These represent the direction and magnitude of each segregant's response to a gRNA after correcting out the gRNA's mean effect. Segregants with insufficient data to obtain a coefficient were excluded from this calculation. gRNAs with deviation values for less than 35 segregants (~20%) were excluded from further analysis.

Linkage mapping

Linkage mapping was performed individually for each gRNA using deviation values as phenotypes. For each gRNA, we performed genome-wide scans using the lm() function in R, with the model lm(*deviations ~ locus*). Significance thresholds were determined by running 1,000 permutations in which a random gRNA was chosen, its deviations were randomly shuffled, and the lowest p-value was saved. From these 1,000 p-values, the 5th percentile was used as the significance cutoff for the respective linkage mapping

model. Peak detection was performed by taking 2x-$\log_{10}$(pval) drops for each locus that was both above the significance cutoff and at least 100 kb away from any other 2x-$\log_{10}$(pval) drop. Any loci for which the gRNA's binding location was <10 kb away from the peak marker were excluded from analysis. These could potentially represent polymorphisms that disrupt a gRNA target site or that impact gRNA binding by modifying local chromatin accessibility. In total, 70 gRNAs were removed this way, 44 of which were already known to directly overlap polymorphisms. Thus, local variation in chromatin did not have a major impact on our study.

Heritability Estimation

Reproducibility of guide effects across ATC1 and CON conditions was used to estimate broad-sense heritability. This was possible due to the replication present in the data set, with most gRNAs and many segregants each represented by multiple distinct barcodes. First, a mean fitness estimate for each segregant was obtained from the CON flask and control gRNAs. Then, for each gRNA, these segregant fitnesses were subtracted from each corresponding double barcode fitnesses from the ATC1 flask, providing a guide effect estimate for every individual double barcode. These estimates were used to fit the linear model *guide_effect ~ genotype*. Using this model, broad-sense heritability for each gRNA was determined by taking the sum of squares between genotypes and dividing by the total sum of squares. Narrow-sense heritability was estimated using the "sommer" package in R. The A.mat() function was used to generate an additive relationship matrix from the deviation values previously generated for each segregant-gRNA combination, and the mmer() and vpredict() functions were used to generate the narrow-sense heritability estimate from this matrix.

Analysis of hub loci

A significance threshold for detecting hubs was calculated by assuming a random distribution of intervals across the genome[53]. Specifically, the genome was split into 20 kb bins and a threshold was set at 0.05 / (number of bins). A Poisson distribution was used to determine the number of overlapping intervals per bin that would result in a p-value below this threshold. For this Poisson distribution, lambda was set as (total

length of all detected 2x-$\log_{10}$(pval) drops in bp) / (total length of the yeast genome in bp). Hubs were defined as adjacent loci with counts (number of overlapping intervals) above this threshold. The locus with the highest count was selected as the marker for that hub. In the case of ties, the genomic positions of the tied loci were averaged, and the locus closest to the average was used as the marker. All gRNA-targeted genes whose 2x-$\log_{10}$(pval) drops overlapped this marker were considered a part of that hub. The effect of each hub locus on deviation value in the context of interacting gRNAs was extracted from the linear model used for linkage mapping lm(*deviations ~ locus*), taking the coefficient of the *locus* term. Proportion of broad-sense heritability explained by the hub locus for each gRNA was also calculated this way, by taking the variance explained by the model and dividing it by the broad-sense heritability estimate for that gRNA. CellMap analysis was performed by uploading lists of genes comprising each hub to cellmap.org and downloading enrichment tables and visualizations[44].

References

1. Mullis, M. N., Matsui, T., Schell, R., Foree, R. & Ehrenreich, I. M. The complex underpinnings of genetic background effects. *Nat. Commun.* **9**, 3548 (2018).
2. Chandler, C. H., Chari, S. & Dworkin, I. Does your gene need a background check? How genetic background impacts the analysis of mutations, genes, and evolution. *Trends Genet* **29**, 358–66 (2013).
3. Nadeau, J. H. Modifier genes in mice and humans. *Nat Rev Genet* **2**, 165–74 (2001).
4. Weinreich, D. M., Delaney, N. F., DePristo, M. A. & Hartl, D. L. Darwinian Evolution Can Follow Only Very Few Mutational Paths to Fitter Proteins. *Science* **312**, 111–114 (2006).
5. Kryazhimskiy, S., Rice, D. P., Jerison, E. R. & Desai, M. M. Global Epistasis Makes Adaptation Predictable Despite Sequence-Level Stochasticity. *Science* **344**, 1519–1522 (2014).
6. Johnson, M. S., Martsul, A., Kryazhimskiy, S. & Desai, M. M. Higher-fitness yeast genotypes are less robust to deleterious mutations. *Science* **366**, 490–493 (2019).
7. Phillips, P. C. Epistasis—the essential role of gene interactions in the structure and evolution of genetic systems. *Nat. Rev. Genet.* **9**, 855–867 (2008).
8. de Visser, J. A. G. M., Cooper, T. F. & Elena, S. F. The causes of epistasis. *Proc. R. Soc. B Biol. Sci.* **278**, 3617–3624 (2011).
9. Rutherford, S. & Lindquist, S. HSP90 as a capacitor for morphological evolution. *Nature* **396**, 336–42 (1998).
10. Queitsch, C., Sangster, T. A. & Lindquist, S. Hsp90 as a capacitor of phenotypic variation. **417**, 7 (2002).
11. Jarosz, D. F. & Lindquist, S. Hsp90 and Environmental Stress Transform the Adaptive Value of Natural Genetic Variation. *Science* **330**, 1820–1824 (2010).
12. Taylor, M. B. & Ehrenreich, I. M. Higher-order genetic interactions and their contribution to complex traits. *Trends Genet.* **31**, 34–40 (2015).

13. Goldstein, I. & Ehrenreich, I. M. The complex role of genetic background in shaping the effects of spontaneous and induced mutations. *Yeast* **38**, 187–196 (2021).
14. Dworkin, I., Palsson, A., Birdsall, K. & Gibson, G. Evidence that Egfr Contributes to Cryptic Genetic Variation for Photoreceptor Determination in Natural Populations of Drosophila melanogaster. *Curr. Biol.* **13**, 1888–1893 (2003).
15. Chari, S. & Dworkin, I. The conditional nature of genetic interactions: the consequences of wild-type backgrounds on mutational interactions in a genome-wide modifier screen. *PLoS Genet* **9**, e1003661 (2013).
16. Chandler, C. H., Chari, S., Tack, D. & Dworkin, I. Causes and consequences of genetic background effects illuminated by integrative genomic analysis. *Genetics* **196**, 1321–36 (2014).
17. Chandler, C. H. *et al.* How well do you know your mutation? Complex effects of genetic background on expressivity, complementation, and ordering of allelic effects. *PLoS Genet* **13**, e1007075 (2017).
18. Taylor, M. B. & Ehrenreich, I. M. Genetic Interactions Involving Five or More Genes Contribute to a Complex Trait in Yeast. *PLoS Genet.* **10**, e1004324 (2014).
19. Taylor, M. B. & Ehrenreich, I. M. Transcriptional derepression uncovers cryptic higher-order genetic interactions. *PLoS Genet* **11**, e1005606 (2015).
20. Taylor, M. B., Phan, J., Lee, J. T., McCadden, M. & Ehrenreich, I. M. Diverse genetic architectures lead to the same cryptic phenotype in a yeast cross. *Nat. Commun.* **7**, 11669 (2016).
21. Lee, J. T., Taylor, M. B., Shen, A. & Ehrenreich, I. M. Multi-locus Genotypes Underlying Temperature Sensitivity in a Mutationally Induced Trait. *PLoS Genet* **12**, e1005929 (2016).
22. Lee, J. T., Coradini, A. L. V., Shen, A. & Ehrenreich, I. M. Layers of cryptic genetic variation underlie a yeast complex trait. *Genetics* **211**, 1469–1482 (2019).
23. Dowell, R. D. *et al.* Genotype to phenotype: a complex problem. *Science* **328**, 469 (2010).

24. Hou, J., Tan, G., Fink, G. R., Andrews, B. J. & Boone, C. Complex modifier landscape underlying genetic background effects. *Proc Natl Acad Sci U A* **116**, 5045–5054 (2019).
25. Paaby, A. B. *et al.* Wild worm embryogenesis harbors ubiquitous polygenic modifier variation. *Elife* **4**, (2015).
26. Galardini, M. *et al.* The impact of the genetic background on gene deletion phenotypes in Saccharomyces cerevisiae. *Mol Syst Biol* **15**, e8831 (2019).
27. Caudal, E. *et al.* Loss-of-function mutation survey revealed that genes with background-dependent fitness are rare and functionally related in yeast. *Proc. Natl. Acad. Sci.* **119**, e2204206119 (2022).
28. Schell, R. *et al.* Genetic basis of a spontaneous mutation's expressivity. *Genetics* **220**, iyac013 (2022).
29. Levy, S. F. *et al.* Quantitative evolutionary dynamics using high-resolution lineage tracking. *Nature* **519**, 181–186 (2015).
30. Schlecht, U., Liu, Z., Blundell, J. R., St.Onge, R. P. & Levy, S. F. A scalable double-barcode sequencing platform for characterization of dynamic protein-protein interactions. *Nat. Commun.* **8**, 15586 (2017).
31. Qi, L. S. *et al.* Repurposing CRISPR as an RNA-Guided Platform for Sequence-Specific Control of Gene Expression. *Cell* **152**, 1173–1183 (2013).
32. Smith, J. D. *et al.* Quantitative CRISPR interference screens in yeast identify chemical-genetic interactions and new rules for guide RNA design. *Genome Biol.* **17**, 45 (2016).
33. Smith, J. D. *et al.* A method for high-throughput production of sequence-verified DNA libraries and strain collections. *Mol. Syst. Biol.* **13**, 913 (2017).
34. Matsui, T. *et al.* The interplay of additivity, dominance, and epistasis on fitness in a diploid yeast cross. *Nat. Commun.* **13**, 1463 (2022).
35. Mullis, M. N. *et al.* Complex genetics cause and constrain fungal persistence in different

parts of the mammalian body. *Genetics* **222**, iyac138 (2022).

36. Lo, W. S. & Dranginis, A. M. FLO11, a yeast gene related to the STA genes, encodes a novel cell surface flocculin. *J Bacteriol* **178**, 7144–51 (1996).
37. Kobayashi, O., Suda, H., Ohtani, T. & Sone, H. Molecular cloning and analysis of the dominant flocculation gene FLO8 from Saccharomyces cerevisiae. *Mol. Gen. Genet. MGG* **251**, 707–715 (1996).
38. Bak, G. *et al.* On-off controllable RNA hybrid expression vector for yeast three-hybrid system. *BMB Rep.* **43**, 110–114 (2010).
39. Farzadfard, F., Perli, S. D. & Lu, T. K. Tunable and Multifunctional Eukaryotic Transcription Factors Based on CRISPR/Cas. *ACS Synth. Biol.* **2**, 604–613 (2013).
40. Li, F., Salit, M. L. & Levy, S. F. Unbiased Fitness Estimation of Pooled Barcode or Amplicon Sequencing Studies. *Cell Syst.* **7**, 521-525.e4 (2018).
41. Liu, X. *et al.* iSeq 2.0: A Modular and Interchangeable Toolkit for Interaction Screening in Yeast. *Cell Syst.* **8**, 338-344.e8 (2019).
42. Lynch, M. & Walsh, B. *Genetics and Analysis of Quantitative Traits*. (Sinauer Associates, 1998).
43. Stelling, J., Sauer, U., Szallasi, Z., Doyle, F. J. & Doyle, J. Robustness of Cellular Functions. *Cell* **118**, 675–685 (2004).
44. Usaj, M. *et al.* TheCellMap.org: A Web-Accessible Database for Visualizing and Mining the Global Yeast Genetic Interaction Network. *G3 GenesGenomesGenetics* **7**, 1539–1549 (2017).
45. Costanzo, M. *et al.* A global genetic interaction network maps a wiring diagram of cellular function. *Science* **353**, aaf1420 (2016).
46. Ba, A. N. N. *et al.* Barcoded Bulk QTL mapping reveals highly polygenic and epistatic architecture of complex traits in yeast. *eLife* https://elifesciences.org/articles/73983 (2022) doi:10.7554/eLife.73983.

47. Li, H. & Durbin, R. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* **25**, 1754–1760 (2009).
48. Li, H. *et al.* The Sequence Alignment/Map format and SAMtools. *Bioinformatics* **25**, 2078–2079 (2009).
49. Zhao, L., Liu, Z., Levy, S. F. & Wu, S. Bartender: a fast and accurate clustering algorithm to count barcode reads. *Bioinformatics* **34**, 739–747 (2018).
50. Bachmann, M. maxbachmann/RapidFuzz. (2023).
51. Gietz, R. D. & Schiestl, R. H. High-efficiency yeast transformation using the LiAc/SS carrier DNA/PEG method. *Nat. Protoc.* **2**, 31–34 (2007).
52. Omelina, E. S., Ivankin, A. V., Letiagina, A. E. & Pindyurin, A. V. Optimized PCR conditions minimizing the formation of chimeric DNA molecules from MPRA plasmid libraries. *BMC Genomics* **20**, 536 (2019).
53. Brem, R. B., Yvert, G., Clinton, R. & Kruglyak, L. Genetic Dissection of Transcriptional Regulation in Budding Yeast. *Science* **296**, 752–755 (2002).