UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Medicine 1919

The impact of selection on the genetic architecture of complex traits

TILMAN RÖNNEBURG

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2023

ISSN 1651-6206
ISBN 978-91-513-1751-9
URN urn:nbn:se:uu:diva-498417

Dissertation presented at Uppsala University to be publicly examined in C8:305, BMC, Husargatan 3, Uppsala, Friday, 8 September 2023 at 13:00 for the degree of Doctor of Philosophy (Faculty of Medicine). The examination will be conducted in English. Faculty examiner: Professor Dirk-Jan De Koning (Department of Animal Breeding and Genetics, Swedish University of Agricultural Sciences).

Abstract

Rönneburg, T. 2023. The impact of selection on the genetic architecture of complex traits. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Medicine* 1919. 42 pp. Uppsala: Acta Universitatis Upsaliensis. ISBN 978-91-513-1751-9.

Accurate dissection of highly polygenic traits is difficult, in part due to the power required to identify and characterise minor loci, but also due to the potential nonadditive interactions between the contributing genetic variations within a population. This is often further complicated by the genetic features of natural or agricultural populations, where a good understanding of the genetic architecture of a quantitative trait, e.g. the risk to develop a disorder, or growth-traits in farm animals, would be beneficial. The aim of this thesis is to contribute to a better understanding of the genetic architecture of quantitative traits. In order to do this, the three studies in this thesis make use of a large 18-generation intercross population created from a long running selection experiment on 56-day bodyweight in chicken, the Virginia weight lines.. Combining this population with a new, cost efficient approach to genotyping, we created a large, powerful dataset to explore multiple aspects of the quantitative trait in question, and how its genetic architecture has been shaped by artificial selection.

The first study describes the approach used to generate the dataset and uses the increased power and resolution for a comprehensive genome wide QTL scan, identifying multiple novel loci and mapping others at better resolution.

The second study leverages the same dataset to study the contribution of capacitating epistasis to the selection response. We identify multiple capacitors that explain a modest amount of the selection response, as well as dissect a previous interaction between two QTL into a larger epistatic network with multiple within and across chromosome interactions that explains a large fraction of the phenotypic variance and selection response.

In the third study, we make use of the outbred nature of the founders to investigate the contribution of still segregating variants to the selection response by adding a GWAS approach to the QTL mapping. We identify multiple novel loci that have not been identified by the QTL approach before, many of which likely still contribute to the selection response due to only segregating in one of the two founding lines. Overall, this thesis showcases the complexity of quantitative trait genetic architecture under selection, by identifying multiple novel loci and epistatic networks that contribute to the selection response in different ways, as well as highlights some of the benefits of combining multiple approaches with different assumptions.

Tilman Rönneburg, Department of Medical Biochemistry and Microbiology, Box 582, Uppsala University, SE-75123 Uppsala, Sweden.

ISSN 1651-6206
ISBN 978-91-513-1751-9
URN urn:nbn:se:uu:diva-498417 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-498417)

For Paul, and all the people that have worked with the Virginia lines before me

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I. **Rönneburg, T**., Zan, Y., Honaker, C. F., Siegel, P. B., Carlborg, Ö. (2023) Low-coverage sequencing in a deep intercross of the Virginia body weight lines provides insight to the polygenic genetic architecture of growth: novel loci revealed by increased power and improved genome-coverage. *Poultry Science* 102(5): 102203

II. **Rönneburg, T**., Pettersson, M.E., Honaker, C. F., Siegel, P. B., Carlborg, Ö. Mapping and dissecting capacitating epistasis in a population subjected to long-term, directional selection. *Manuscript in preparation*

III. **Rönneburg, T**., Ou JH., Pettersson, M.E., Honaker, C. F., Siegel, P. B., Carlborg, Ö. Within-line segregation as a contributor to long-term, single-trait selection-responses in the Virginia chicken lines. *Manuscript in preparation*

Contents

Abbreviations

AIL	Advanced Intercross Line
QTL	Quantatiative Trait Locus
GWAS	Genome Wide Association Study
HWS	High Weight Selected
LWS	Low Weight Selected
FDR	False Discovery Rate
vQTL	variance Quantitative Trait Locus
SNP	Single Nucleotide Polymorphism
ATAC-seq	Assay for Transposase-Accessible Chromatin using sequencing
RIL	Recombinant Inbred Line

Introduction

Quantitative traits, complex traits

Quantitative Traits

Genetics is the study of heredity, the passing of genetic information and associated phenotypic features or quantitative variation in such from progenitor to offspring.

The former are most commonly used to illustrate genetics, using examples such as flower-color in Gregor Mendels peas, (e.g Figure 1A), albinism, or coat colour and length in cats (1,2). Assessing the impact of this genetic variation on the phenotype in question (Color of the inflorescence, lack of pigmentation or e.g. length of the haircoat) is straightforward, and the line of inheritance easily traced. These traits are considered monogenic, meaning they are governed by genetic variation in a single region of the genome, and are inherited in a mendelian fashion. Beyond their use as demonstrators, in humans these traits often have medical considerations as rare, inheritable and often recessive disorders, such as cystic fibrosis, Huntington's disease and thalassaemia (3–5). While these are often debilitating for the affected, identification of the relevant regions in the genome via linkage mapping and studies in affected families, as well as prevention via screening for carriers are theoretically easy, barring ethical considerations (6,7). Other traits and their hereditary components are more difficult to assess, since their effect does not partition into qualitative categories. Taking height in humans as an example, assuming monogenic, mendelian inheritance, there would be only 'tall' and 'small' people, with every individual fitting neatly into one of the two categories. However, while there are tall and small people, there is also everything in between. Therefore, instead of qualitative categories, we measure height with a value on a continuous scale and consider it a quantitative trait. This holds true for many medical considerations as well, such as the risks that a given individual will suffer from cardiovascular disorders, type II diabetes or schizophrenia during their lifetime (8–10). Often, some of the phenotypic variance in a quantitative trait can be explained by environmental influences. Height for example depends on the availability of food and nutrients during the childhood of a given individual (11), but commonly they also have a heritable component. Tall people are more likely to have tall children (12), and

people whose parents suffer from cardiovascular disorders, type II diabetes or schizophrenia are more likely to have these diseases and disorders as well. While it is often difficult to estimate the fraction of phenotypic variance that is heritable and disentangle it from environmental effects, identifying the genomic regions responsible for the inheritable variation is even more difficult. Some of the wide variation that makes a trait quantitative is due to environmental influences, but it is commonly also due to the polygenicity of the trait: many different variants in the genome contributing to the phenotypic variation we observe. This leads to a staggering amount of genotype-combinations with unique contributions to the phenotype, with a set of four bi-allelic loci in a diploid organism already leading to $3^4 = 81$ unique genotypes. Given sufficient polygenicity, the phenotype distribution of any quantitative trait quickly resembles a normal distribution (e.g. figure 1B) with many different combinations leading to very similar phenotype values.

For human height, the upper bound of estimates assumes that more than 100.000 variants across the human population contribute to height (13). While some of these may have comparatively large effects, the vast majority of these variants will only explain a minute fraction of the total phenotypic variation. This feature, and the inverse relationship between effect size of variants and statistical power required to confidently assess their contributions, makes for one of the core challenges of investigating the genetic basis of quantitative traits.

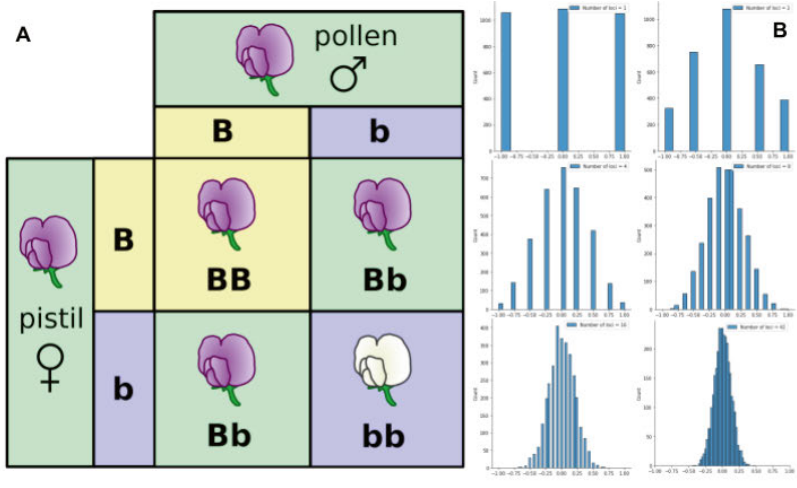

Figure 1. A: example punnett square for pea-flower colour, cross of two F_1 individuals, with purple corresponding to the dominant Allele (B) and white to the recessive allele (b). Madeleine Price Ball, CCO1.0 via Wikimedia Commons. **B** simple simulations to show how combinations of loci with discrete effects across a population rapidly resemble continuous phenotype distributions akin to a normal distribution with the number of loci contributing to the trait. From top left to bottom right: random draw of phenotypes for 1,2,4,8,16 and 42 equally contributing loci, 3200 individuals each.

Complex traits

Beyond the challenge of correctly identifying and characterising small effect size loci in an ideal, purely additive scenario, quantitative traits rarely conform to such assumptions. Biology is rife with examples of complex interactions on every sense of scale and dependencies, be it complex metabolic pathways (14), signalling cascades (15), protein-protein interactions (16), protein-genome interactions (17) or interactions leading to different 3-dimensional structures of the genome (18). It stands to reason that the statistical genetic architecture of quantitative traits underlying complex biology mirrors this complexity (19–24). These gene-by-gene interactions, or epistasis, describe the phenomenon that the effect of one variant is contingent upon, or modified by, another. Epistasis, particularly capacitating epistasis, where one variant enhances or suppresses the effect of another, has been posited as one of the mechanisms allowing for the continued selection response beyond the initial phenotypic range of a trait (20,21), as well as a potential explanation for the "missing heritability problem" (25), which posits that the regions and variants identified in QTL or GWAS studies often only explain a small fraction of the total heritability. While there is no shortage of biological or genetic explanatory mechanisms for the presence of epistasis, it has been surprisingly difficult to identify and confirm epistatic interactions when investigating quantitative traits. Part of this is due to the increased power requirements of trying to account for the interactions between two loci, or even higher order interactions, with combinatorial explosion quickly leading to a searchspace of testable combinations that would overwhelm the power of any dataset.

Quantitative Traits in Natural Populations

Studying Quantitative traits in natural populations is particularly difficult. This is predominantly due to environmental influences, population structure and unbalanced allele-frequencies.

Environmental influences are, on the one hand, problematic due to the size of their effects relative to the effect of individual loci in a polygenic trait, but polygenic, quantitative traits also represent many genetic components and biological systems that can be impacted by environmental factors, leading not only to larger environmental effects, but also a higher possibility for gene-by-environment interactions (22–24). Similarly, confounding of a significant fraction of the variance with population structure is more likely due to polygenicity, since many variants contributing means that a larger fraction of the genome is contributing to the trait, or being in the vicinity of variation that does. Beyond the fact that low frequency variants result in fewer observations, thus raising the overall power-requirements for detection, low allele frequencies in natural populations also often mask epistatic variance as additive. Assuming most variants are very rare, it is unlikely that all epistatic loci involved in an interaction are at balanced allele-frequencies, resulting in some genotype

combinations being very rare, particularly double-homozygotes, assuming Hardy-Weinberg Equilibrium. This results in a large fraction of the variance explained by the interaction becoming hidden, assuming the 'on' combination is very rare, or appearing as an additive effect on the capacitated locus, assuming it is common (26). In addition to that, similar to how linked minor QTL can appear as a single large QTL given limited resolution, local epistatic effects will appear as additive given insufficient resolution and linkage between the loci.

Similarly, having a wide range of low or very low minor-allele-frequency variants might also lead to these variants perfectly coinciding with the rarer of the 'off' or 'on' combinations of epistatic loci, covering the epistatic effect and appearing as an extra additive locus. This may not be a problem if the goal is only accounting for the variance or doing genomic prediction in the population studied, but it precludes identifying the causative regions or mechanisms and limits transfer of these findings to populations that only share partial ancestry with the study population.

Quantitative traits in experimental crosses

The aforementioned issues and the power requirements to alleviate them is often the reason for quantitative traits to be investigated in a more limited and controlled fashion. One way to do so, in non-human organisms, is by creating experimental populations. Experimental crosses can solve part of this problem by moving the power-requirements into the achievable range by minimising environmental effects, exaggerating the investigated phenotypes, balancing the allele frequencies and reducing the complexity of the investigated genetic architecture.

Experimental crosses allow for controlling environmental effects by keeping all individuals under the same conditions, as well as accounting for population structure by either selecting individuals very carefully, or alternatively create them via controlled mating, making population structure and kinship within the population a known variable.

The most common way to do so has been by making F_2 populations from a cross of individuals that differ in the investigated trait (e.g. individuals from two different populations, cultivars, breeds or inbred lines). This is done by crossing the founders (the F_0 generation) to create individuals that are heterozygous for the founder-genotypes (the F_1 generation), relying on recombination events in these individuals to create an F_2 population where individuals carry unique mosaics of founder genotypes. Using these mosaic individuals and their corresponding phenotype, one can investigate where in the genome genetic differentiation between the founders contributes to the difference in the investigated phenotype. When using inbred founders, quantitative trait analysis in F_2 populations only makes use of the phenotypic and genotypic variation between the founders instead of the variation within a population.

This can reduce the scope of enquiry and explanatory power somewhat, or be a boon due to larger between-founder-populations difference in phenotype than phenotypic variance within the populations. However, it also limits the resolution of the study to the number of recombinations that occurred in the F_1 generation. This enables investigations with lower requirements in terms of statistical power, at the cost of resolution. Beyond that tradeoff, lower resolution also comes with additional artefacts, such as so-called “ghost QTLs", apparent QTLs in the space between two real QTLs that are linked due given the limited number of recombinations, or linked QTL not being visible due to opposite effects. This has historically been less of an issue, since the number of genotype markers was the limiting factor for resolution. However, with the advent of affordable next-generation-sequencing, the number of individuals needed (and therefore recombination events) to attain that same state quickly becomes cost-prohibitive.

Advanced Intercross Populations

One way to alleviate this has been to increase the number of recombinations per individual, by generating individuals from subsequent crosses of the F_2 - generating an intercross line, or advanced intercross line (AIL) instead of creating more F_2 individuals, with individuals from subsequent generations inheriting recombinations from their progenitors, breaking up the mosaic further. While this enables much higher resolution, if the marker density allows, it comes with some of the problems inherent to more natural populations, e.g. environmental effects, family structure and unbalanced allele-frequencies, albeit more attenuated. Special care needs to be taken that environmental conditions are the same across the generations to decrease batch effects, and similarly to the breeding scheme, in order to avoid pronounced population structure. Some unbalanced allele-frequencies are unavoidable due to genetic drift in smaller than infinite-sized populations, but due to the balanced allele-frequencies in the starting population, this remains less of a problem than in natural populations.

Crosses from outbred founders

Another way to increase resolution for e.g. fine mapping is to utilise outbred founders for the initial cross. While this does not increase the number of recombinations, the within-founder-population genetic variance can still be used to infer founder-line origin of a given region, increasing the marker density and resolution, with no additional multiple testing burden beyond the increase in resolution. More importantly, the genetic variation retained within the line can be used to explain within-line phenotypic variance as well. Even if the experimental population is not powerful enough to bear the multiple testing burden a genome wide marker based association study would entail, they can

still be used for fine mapping. If the population is powerful enough for that, a genome wide association study (GWAS) opens up the possibility to investigate genotype-phenotype association beyond the difference between the founders of a cross, and also investigate QTL that are not fixed or highly divergent between the lines due to e.g. rarity, strongly deleterious effects in the homozygote or other other factors.

The Virginia weight lines

The Virginia weight lines are a bidirectional selection experiment for 56-day bodyweight in White Plymouth Rock chicken, initiated in 1957 by Paul B. Siegel at Virginia Polytechnic Institute and State University (27). These weight lines were created from a common, outbred stock that had previously been generated from 7 inbred lines of White Plymouth Rock chicken by selecting the 56 (48+8, female+male) individuals of the population with the highest and lowest body weight at 56-days of age and then continue to select the top 40-24% (due to variation of generation size between 150-250) for the respective extreme over the course of subsequent generations, resulting in a High Weight Selected (HWS), and a Low Weight Selected (LWS) line (28–30). The easily accessible phenotype and the polygenicity make bodyweight an excellent model for complex traits, while the choice of chicken as a model organism meant that the relatively fast generation time, modest requirements in space and effort to standardise the environment beyond what is already common in the poultry industry, as well as a sustained commercial interest in poultry breeding enabled running a large-scale selection experiment in a vertebrate model for over 60 years.

Artificial selection and breed history

Arguably, the genetic history of the studied population and the selection regime applied are both far from what one would expect from a natural population, be it red jungle-fowl, the undomesticated ancestor of chicken, or a human population. The utilised breed of chicken, the Plymouth Rock chicken, first emerged as a distinct breed in the mid-19th century (31) in North America, both in its white and barred form, having been created from a mixture of asian and european breeds (32).

Transport of the founding breeds to the Americas likely acted as a strong bottleneck on these founding populations, with the breed-formation only exacerbating the resulting reduction of genetic diversity of the founding individuals and the resulting recombinant individuals making up the new breed. Subsequently, in order to partially balance allele-frequencies, the starting population for the selection experiment was created from seven inbred lines derived from an existing outbred stock.

Each of these events, including the following selection in the high and low weight selected lines will likely have generated linkage disequilibrium between functional variants, creating the haplotypes that the subsequent selection will have acted upon. Throughout the artificial selection process, selection has acted upon these lines at different scales, and, due to the strong selection pressure, in a sequential fashion. Starting with an initial selection of recombined ancestral-breed chromosomes from the initial stock to form the lines, then combinations of said chromosomes within the lines, and then sequential sweeps on whichever haplotype confers the most extreme phenotype. Notably, this could have also facilitated the selection of variance-increasing loci such as epistatic capacitors, as the limited number of individuals chosen from each generation favours extreme phenotypes.

Taken together, these events reduced the genetic diversity and simplified the genetic architecture to the point of making it feasible to gain meaningful insights into said architecture beyond individual loci. The limited genetic diversity and particularly the strong continued selection without defined optima or any other fitness constraints beyond the outright lethal, also means that this is best seen as a caricature of selection and natural populations, highlighting characteristics and features of complex trait architecture, instead of a direct stand-in.

One of the first striking observations from the initial generations of selection in the virginia lines was the strong, unabated selection response beyond the initial phenotypic variance of the starting population. Contrary to expectations that the extreme selection pressure would rapidly exhaust the genetic variation, leading to inbred or near-inbred lines that plateau in regards to the phenotype after a few generations, the lines continued to respond to selection for more than 50 generations (30) Leading to a nine-fold difference in body weight at Generation 40 and 16-fold at generation 59 (33), With the lines showing strong genetic differentiation without exhausting genetic variation (34–36).

Creation of the AIL

In order to investigate how the genetic differentiation between the lines contributes to the strong phenotypic divergence between the lines, a cross was generated, using founders (n_{HWS}=29, n_{LWS}=30) from generation 40 with an initial intent of QTL mapping in the F_2 (37,38) this was subsequently continued into an Intercross line, with the aim of doing fine mapping in the F_8 and then F_{15} (33,34,39–41) , with individuals generated from up to generation 18.

Sequencing

Collecting all available individuals yielded a population of around 3k intercross line individuals, which was sequenced at very low coverage (~0.4X) using a tagmentation approach (42). This was partially done due to cost-efficiency of the selected library preparation and sequencing approach, but also to preserve options for downstream processing of the data, in particular with regards to different imputation approaches, and exploring within-line variation, in regards to GWAS approaches.

Summary of Manuscripts

This thesis consists of three projects, each attempting to evaluate how different aspects of a quantitative trait contribute to the selection response observed after the many generations of extreme artificial selection for bodyweight that the Virginia weight lines have been subjected to. At the heart of these three studies is not a new theory, such as e.g. a new statistical model to consolidate the different ways genetic variation can exert power over the phenotype investigated, but a gamble on statistical power: “What could we find (and prove) if we had a sufficiently powerful dataset?" and aims to find confident answers to more limited questions, e.g. “Which regions contribute how much to the selection response?" (Paper I), “What is the contribution of capacitating epistasis to the selection response?” (Manuscript II) And “What is the contribution of still segregating variants to the phenotype, and how much of it also contributes to the selection response?" (Manuscript III). Most of my time has been spent creating this dataset and solidifying our confidence in this dataset, as well as developing and improving the methods to both create and quality-control it. While this work is described in Paper I, it is the foundation of Manuscript II and III as well.

Paper I

Low-coverage sequencing in a deep intercross of the Virginia body weight lines provides insight to the polygenic genetic architecture of growth: novel loci revealed by increased power and improved genome-coverage.

This study is a classical QTL study on a nonclassical dataset, investigating the genetic basis of 56-day bodyweight in the Virginia weight lines. In order to maximise statistical power on a limited budget, we collected all available individuals from an existing population - the entire intercross line - instead of focussing on a single generation. In order to be able to genotype so many individuals we sequenced them at extremely low coverage (~0.25-0.4X) adapting a library preparation and imputation approach previously developed in our

research group (51) for use with individuals from a deep intercross line. This approach utilises a sliding window hidden-markov-model approach to impute recombination breakpoint positions and the founder haplotypes in between, using founder-line informative markers, combining information from high-coverage sequenced founder-individuals, pedigree information and the low-coverage sequencing data. The main advances here include procedures to select founder-line informative and still segregating but partially informative markers from each individual depending on their ancestry, as well as modifications accounting for the extreme variability in marker density across the genome from this sequencing approach, and the diminishing number of partially informative segregating markers due to further crosses to establish subsequent generations of the intercross line.

Using this approach, we obtained high-confidence genotypes across 99.3% of the chicken genome for 3327 individuals (Fig 2).

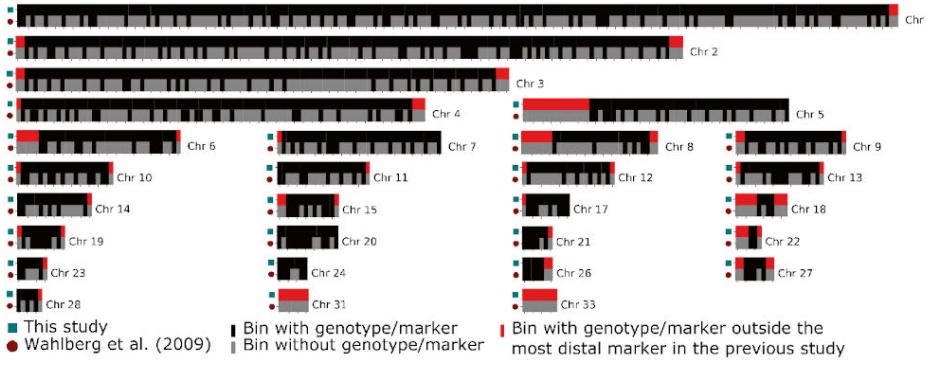

Figure 2. Genome coverage by the use of low-coverage sequencing data compared to an earlier F_2 genome scan (Wahlberg et al. 2009) based on 372 SNP and microsatellite markers. Black/grey colors indicate 1 Mb bins with/without genotypes and red highlights chromosome ends with new genotypes outside the outermost markers in Wahlberg et al. (2009). Figure and caption from Rönneburg et al. (2023) (Paper I) under CC BY 4.0

Comparing it to the most recent previous genome wide QTL scan on this population, I found many more QTL contributing to 56 day bodyweight: 12 instead of just two Genome-wide significant QTL ,and an additional 30 suggestive (i.e. passing 10% FDR threshold) QTL, with some of these completely novel, and others confirming previously suggestive QTL (Fig. 3).

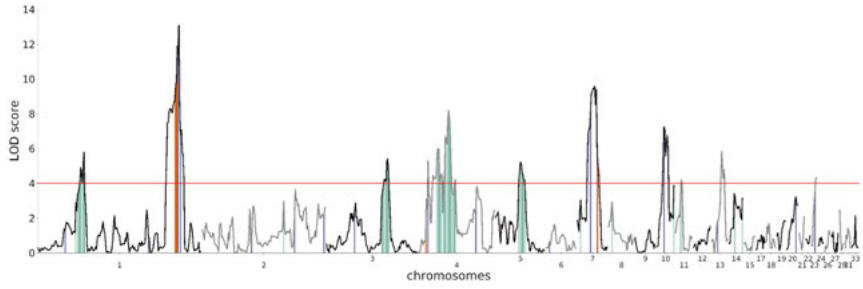

Figure 3. Genome-wide QTL scan for 56-day body weight in generations F2-F18 of the Virginia body weight lines AIL. The y-axis shows the statistical support for QTL as LOD scores and the x-axis the genomic location in Mb bins. The solid red horizontal lines shows the genome-wide significance thresholds. Red vertical segments show the most recent earlier reported associations in genome-wide scans in the F2 (Wahlberg et al. 2009), blue vertical segments indicate associations in the fine-mapping analyses in the F15 (Zan et al. 2017). Green vertical segments indicate new suggestive QTL without a previous association within 15 Mb. Figure and caption from Rönneburg et al. (2023) (Paper I) under CC BY 4.0

Most of these are due to the increased power of the dataset, but some of the QTL, e.g. on Chromosome 8. are novel because they are in regions previously not covered. Others are now identifiable due to the increased resolution, either in terms of marker density, or increased number of recombinations (e.g. on Chromosome 1 and 4, see Fig. 4). Taken together, these loci explain a large fraction of the difference between the founding lines (37% or 84% for genome wide significance and 10% FDR QTLs, respectively), but only a modest albeit still improved fraction of total phenotypic variance (8.3%). While the obvious main result of this project is the increased list of QTL contributing to the selection response, for me the key achievement is generating a powerful dataset from an eclectic collection of existing individuals, using extremely low coverage sequencing and bespoke imputation that makes use of the features of an AIL with known founders. Both because it is an example for how to leverage cheap sequencing to make use of existing individuals, which could be applicable to many otherwise languishing selection experiments, and because it serves as a foundation for subsequent studies into the nature of quantitative traits.

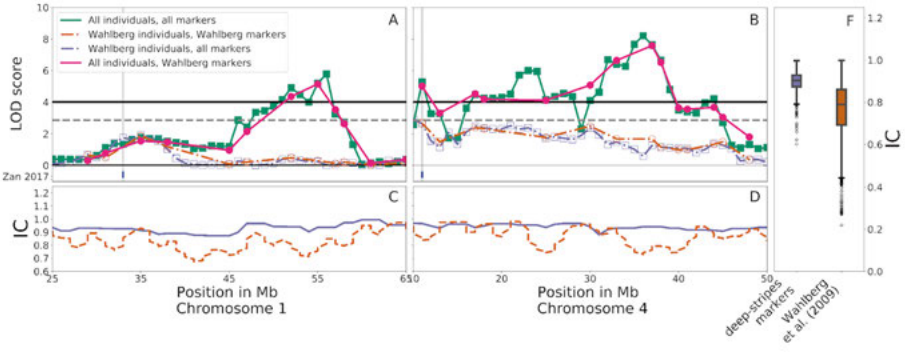

Figure 4: Comparison to previous marker-sets. Upper left and middle panels show the LOD score across selected peak regions on chromosome 1 (25-65 Mb) and chromosome 4 (10-50 Mb), using different selections of markers and individuals. Dashed/solid lines show the QTL scans using the F_2 individuals from Wahlberg et al. / All individuals, solid round / empty square markers show QTL scans using markers from Wahlberg et al. (2009) / low coverage data. Bottom left and middle panels display information content across the same regions on chromosome 1 and 4 for the markers used in Wahlberg et al. (orange, dashed) and the *deep-Stripes* markers (periwinkle, solid) across the individuals used by Wahlberg et al. (2009). The right panel summarises information content for the same sets of markers and individuals, but across the 30 largest chromosomes. Figure and caption from Rönneburg et al. (2023) (Paper I) under CC BY 4.0

Manuscript II

Mapping and dissecting capacitating epistasis in a population subjected to long-term, directional selection.

This study is an exploration of capacitating epistasis and its potential contribution to the selection response in a population that underwent long-term, divergent selection on a complex trait, using the dataset generated in the project described by Paper I. Interaction and interdependence of individual components is one of the hallmarks of biological systems, yet remains understudied as an explanatory mechanism in quantitative genetics. However, for this system and study population, the presence of epistatic interactions has been both shown and replicated (20,41) as well as posited as one of the explanatory mechanisms for the continued selection response in long-term selection experiments (21,28). Here, I make use of the size, power and resolution of the population dataset to look for capacitating hubs of epistatic networks with a clear directional effect - the epistatic interactions we deem most likely to contribute to the long term selection response. In order to do this, I combined two different approaches to identify potential loci. Firstly, assuming capacitative loci would be visible as marginal QTL in a sufficiently powerful additive scan, I took all additive QTL as candidate loci, stratifying the dataset on their genotype and ran the QTL scan in the two homozygous strata to see which QTL

impacted which others. I used the resulting data to build a directed graph with weighted edges for network analysis, identifying hub nodes and their subnetworks. This is a very sensitive approach, picking up a lot of other signals such as strong local haplotype effects, or epistasis that is not contributing to the selection response (Fig. 5).

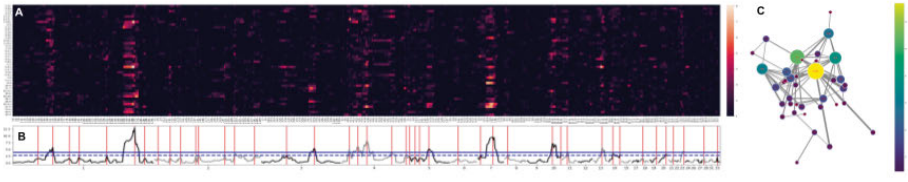

Figure. 5: Summary of the results from the genome-wide analysis to detect capacitor loci using stratification. A: Illustration of the differences in LOD scores between the homozygote high and low strata (Y-axis, color as in legend) across the evaluated 1Mb bins across the genome (X-axis). B: Genome-wide QTL scan for additive-effects (Rönneburg et al, 2022). Red vertical lines indicate the stratifying loci used in the capacitor scan (i.e. QTL detected at a 10% FDR threshold), blue horizontal lines indicate the genome (solid) and chromosome wide (dashed) 5% significance thresholds, respectively. C: Graph representation of all loci where a stratification resulted in a 4.01≥ LODscore difference between the high and low strata. Edgewidth is scaled with the difference between the strata, with nodes colored and sized by their degree weighted by the difference between the strata. Figure and caption from Manuscript II

In order to focus on capacitating loci with a strong directional effect, I implemented a variance-QTL scan using a Brown-Forsythe (43) approach, which is less sensitive, but better at singling out loci with a strong directional effect, since it only evaluates one locus at a time. For this I used a higher resolution haplotype imputation, since I was less concerned about novel loci, but rather identifying capacitive components of existing ones. Combining these approaches, I identified a set of 6 candidate loci across 4 chromosomes, capacitating a modest total 259g of phenotypic variance across their epistatic networks. Beyond this systematic approach, I also investigated a set of QTL in close vicinity on chromosome 4 that had shown some irregularities in ancillary investigations, together with a previously identified epistatic interaction between growth-associated QTL on chromosomes 4 and 7 (20). Previously thought to be two larger QTL, *Growth7* and *Growth9*, I show that when investigated with the resolution of a very large advanced intercross population, they decompose into a network where two capacitors on chromosome 4 and 7 release an effect for two capacitated loci on chromosome 4, explaining more than twice the effect of the purely additive model when accounting for the interaction (136.9g vs. 331.5g, Fig. 6). This provides not only a better estimate for how genetic components correspond to the selection response observed, but also a mechanistic example of how higher resolution and power enabled the dissection of previously thought to be large, additive QTL into a complex

network of multiple smaller, interacting QTL, coherent with previous assumptions about the polygenicity and complexity of a highly polygenic trait.

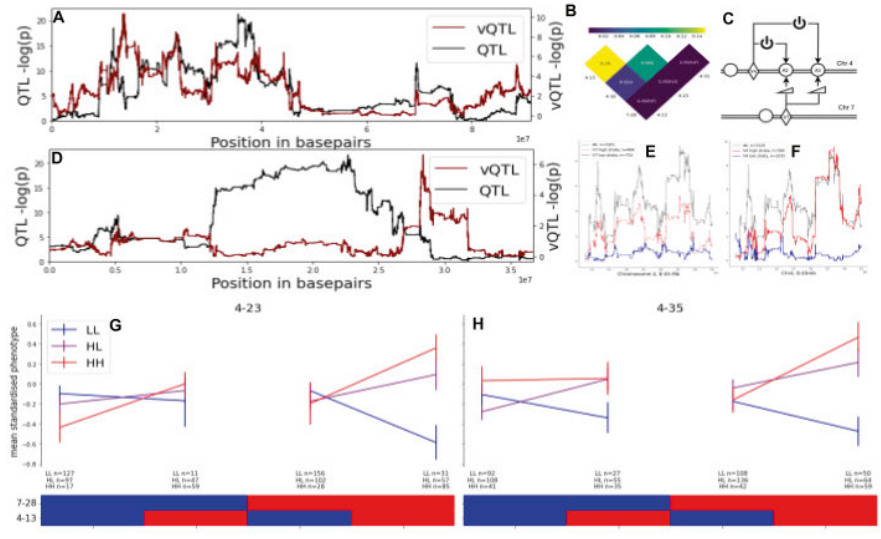

Figure. 6: Interaction between QTL on Chromosome four and seven. **Panel A:** Higher resolution QTL and vQTL scan of Chromosome 4. The black line indicates the negative log of pvalues for the QTL (left y-axis), the red line indicates the negative log of pvalues for the vQTL (right y-axis). **Panel B:** Linkage Disequilibrium between the investigated QTL, Pearsons r^2. **Panel C:** Schematic of the Interaction. **Panel D:** Higher resolution QTL and vQTL scan of Chromosome 7. The black line indicates the negative log of pvalues for the QTL (left y-axis), red line indicates the negative log of pvalues for the vQTL (right y-axis).**Panel E:** QTL-scan for Mb 8 to Mb 45 on chromosome 4 in the total (black), homozygous high (red) and low-weight (blue) individuals for vQTL 7-28. Y-axis indicates LOD score, X-axis indicates position in base pairs. **Panel F:** QTL-scan for Mb 8 to Mb 45 on chromosome 4 in the total (black), homozygous high (red) and low-weight (blue) individuals for vQTL 4-13. Y-axis indicates LOD score, X-axis indicates position in base pairs. **Panel G:** Mean standardised phenotypes for the different genotypes of 4-23 separated by the genotypes of the associated vQTL homozygotes (bottom, blue indicates low-weight homozygote, red indicates high-weight homozygotes).Y-axis indicates the mean standardised phenotype, colored lines the genotype at the 4-23 locus. (blue, purple and red indicating homozygous low-weight, heterozygous and homozygous high-weight genotypes, respectively). Vertical bars indicate the standard error of the mean. **Panel H:** Mean standardised phenotypes for the different genotypes of 4-35 separated by the genotypes of the associated vQTL homozygotes (bottom, blue indicates low-weight homozygote, red indicates high-weight homozygotes) . Y-axis indicates the mean standardised phenotype, colored lines the genotype at the 4-35 locus. (blue, purple and red indicating homozygous low-weight, heterozygous and homozygous high-weight genotypes, respectively). Vertical bars indicate the standard error of the mean. Figure and caption from Manuscript II

Manuscript III

Within-line segregation as a contributor to long-term, single-trait selection-responses in the Virginia chicken lines.

Manuscript III is an attempt to quantify the genetic contribution to both the selection response and more generally the phenotypic variance by regions out of scope for the QTL paradigm and line-cross design of the experimental population. While the founding lines of the AIL were generated from outbred founders, one of the primary assumptions of paper I and manuscript II (and most earlier work) is that, after 40 generations of intense selection which resulted in a nine-fold divergence in bodyweight, the important regions contributing to the selection response were fixed or had highly divergent allele-frequencies between the lines. While there were previous indications that the founding lines still retained significant variation, even in regions that were highly divergent between the lines (35), analysing this variation has so far been difficult, due to the higher power, marker density and number of recombination events required to do so. The imputation approach we developed and used for paper I and manuscript II uses segregating variants, but only uses them to assign founder-line-origin to a given region, and predict recombination breakpoints. Here, we supplemented this imputation approach with a marker based imputation using AlphaFamImpute (44) which imputed SNP-markers from the low coverage sequencing data to the high marker density of the founders sequenced at 30X coverage, using both sequencing-read and pedigree information. We then used the resulting data in a GWAS approach (Fig.7) to build a multi-locus model using a forward selection and backwards elimination (45,46) approach to retain all important markers at 5% FDR. Using these markers, and combining them with the information from the QTL approach and several summary statistics derived from both the study population and its founders, we classify how much of the explained phenotypic variance is explained by loci that are not overlapping identified QTL regions and whether or not these still contribute to the selection response (Tab 1). Strikingly, a large fraction (24 out of 40) retained loci do not overlap known QTL regions, and explain a large fraction of the total phenotypic variance (15%, a=557.5g). However, out of those, 14 (a=346.6g, 9.3% of total phenotypic variance) are fixed in at least one of the lines, and likely still contribute to the selection response. These findings indicate that a large fraction of the selection response is likely contained in rare alleles only segregating in one line, which

makes it difficult to detect them with a classical QTL approach, but also represent potential for a continued selection response after 40 generations of selection.

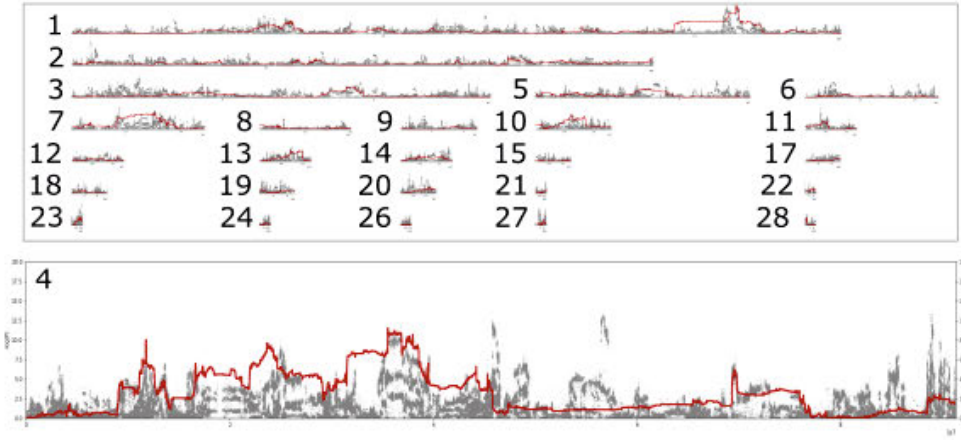

Figure 7: Overview of the 28 largest Autosomes. Grey Manhattan Plot indicates GWAS Pvalues, Red line indicates LODscore of the QTLscan. **Upper figure:** Thumbnails of all chromosomes. Larger figures for each chromosome can be found in the Supplementary. **Bottom figure:** Chromosome 4. Xaxis indicates basepairs, left y-axis is -log(P) of the GWAS (grey markers), right y-axis indicates LOD scores for the QTL-scan (red line). Figure and caption from Manuscript III

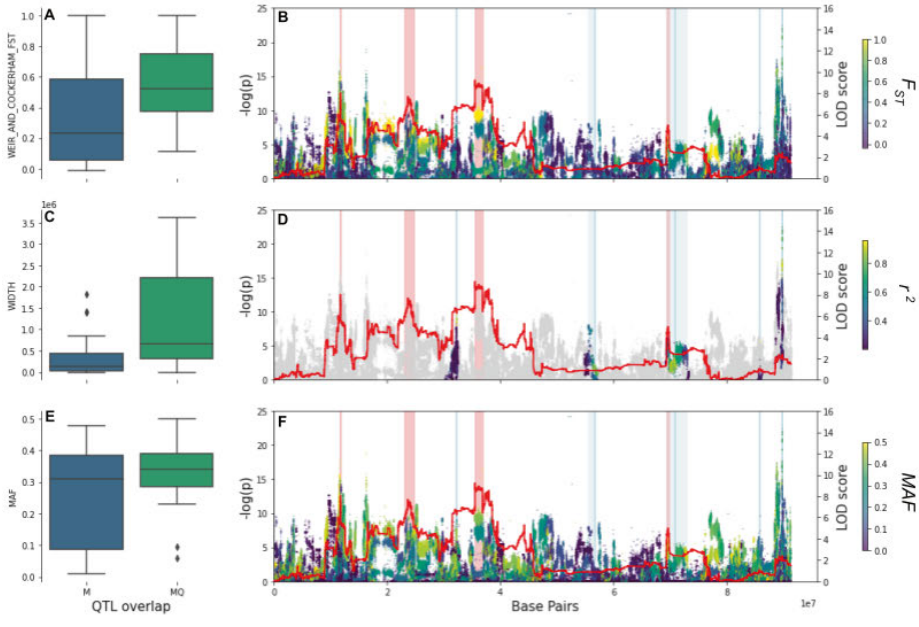

Figure 8. Characteristics of regions detected in QTL and GWAS analyses. **A** Boxplot illustrating the differentiation of selected markers between the founding lines, using Fst for markers that overlap a previously identified QTL (MQ, green) and those that do not (M, blue). **B** Illustration of GWAS and QTL regions on Chromosome 4. GWAS Manhattan plot colored by F_{st} (see colorbar on the right side, with 1 indicating total fixation between the founders, and 0 no differentiation). The red line indicates QTL lod score, with vertical red shading indicating the extent of peaks reaching significance at 10% FDR. Blue lines indicate the position of markers selected in the backwards elimination analysis, with blue shaded regions indicating their extent as determined by LD. **C** Boxplot illustrating the LD-extent of for markers that overlap a previously identified QTL (MQ, green) and those that do not (M, blue). **D** Illustration of GWAS and QTL region widths on Chromosome 4. GWAS Manhattan light gray, with peak-regions colored by the LD R2 to the selected marker (see colorbar on the right side) The red line indicates QTL lod score, with vertical red shading indicating the extent of peaks reaching significance at 10% FDR. Blue lines indicate the position of markers selected in the backwards elimination analysis, with blue shaded regions indicating their extent as determined by LD. **E** Boxplot of the Minor Allele Frequency for markers that overlap a previously identified QTL (MQ, green) and those that do not (M, blue). **F** Illustration of GWAS and QTL regions on Chromosome 4. GWAS Manhattan plot colored by the MAF (see colorbar on the right side). The red line indicates QTL lod score, with vertical red shading indicating the extent of peaks reaching significance at 10% FDR. Blue lines indicate the position of markers selected in the backwards elimination analysis, with blue shaded regions indicating their extent as determined by LD. Figure and caption from Manuscript III.

Table 1: Partitioning of total explained variance between different groups of markers as gram(percentage of phenotypic variance). Table taken from Manuscript III

fully segregating		247.7g (6.7%)
fixed in at least one line		642.7g (17,3%)
Overlapping QTL		332.8g (8.9%)
Not overlapping QTL	of which fully segregating	211.0g (5.7%)
	of which fixed in at least one line	346.6g (9.3%)
	total	557.5g (15.0%)
Total		890.4g (23.9%)

Discussion and Perspectives

Post mortem of a PhD

It is a common trope that PhD students start with high hopes and grand ambitions, only to wonder in the end whether they now know less about their area of research than when they started. This was certainly true for me. While this could be easily just attributed to my personal growth of understanding and overview of the field, I like to think that it is in part also because our research has increased the general knowledge of the field and its extent. Looking at it in hindsight, it is likely that I have underestimated the complexity of the quantitative trait we were investigating, but also that this would not have been as obvious without the research presented here.

More loci, more problems

I do not think that this aforementioned conclusion was foregone: With high quality genotypes for more than 3000 individuals from an 18-generation intercross line, following 40 generations of selection on a single trait, we had a very powerful tool for investigating quantitative traits, coming with a long history of experiments stretching into the 1950s, likely one of the largest and longest running selection experiments in vertebrates to this date, no less! However, the large body of previous experiments and their sometimes conflicting results is also a good indication for how complex the genetic architecture is. Another is the number of novel loci, in particularly novel suggestive loci we identified in the study described in Paper I. The presence of many suggestive loci with intermediate or minor effects indicates that the limit for identifying contributing loci is not the number of loci contributing in a meaningful way, but rather our power to detect them.
While this does not detract from these findings in itself, it may limit how authoritative our findings on the general nature of the genetic architecture are. More importantly, beyond that, they represent a bit of a two-edged sword: On the one hand, they represent a step forward in understanding which regions

contribute to the selection response and are one of the main points demonstrating that our methodology can be used to generate a large dataset on a relatively modest budget. On the other hand, the large number of novel loci recreates a previous conundrum that these studies were initially intended to solve: investigating epistasis is a costly endeavour when it comes to statistical power, and additional candidate loci increase the potential search space in an exponential manner, making the investigation into the genetic architecture of a trait more difficult.

Would adding more statistical power lead to more novel loci?

Given the number and effect sizes of QTL identified here, I think increasing statistical power would have rapidly diminishing returns and will likely not lead to a sufficient increase in explanatory power. Given the large selection pressure exerted, the vast majority of the selection response will likely be explained by several tens or low hundreds of loci, and not reach the polygenicity of a more neutral trait, like e.g. height in humans. This makes it unlikely to lead to the detection of many novel loci, even though there are probably many more loci that affect 56-day body weight. However, these are likely firmly within the domain of the infinitesimal model, and will have negligible individual effect sizes, affecting body weight more as a pleiotropic effect, given that some of the statistically significant or suggestive loci already verge on insignificance in a practical biological sense. It is likely impractical to quantify and locate each of these loci beyond accounting for them jointly as a sort of polygenic effect, though given that we already cover more than 30% of the genome with suggestive QTL, likely even that is unnecessary.

Benefits of added power and resolution

However, the added power and resolution helped dissect existing QTL. Given a sufficiently polygenic architecture, it is likely that many of the larger effect size QTL detected in an early generation population are composites, consisting of multiple smaller QTL in close vicinity. This can certainly help with fine mapping of causative regions, and in the case of local epistatic interactions increase the explanatory power of a region, as it has been for the complex region on chromosome 4. I think the breakdown of a large QTL into multiple additive and epistatic components delivers a fascinating insight into the nature of large effect size QTL that are just as likely to consist of multiple interacting components, than a single, high impact variant. With a dependency on multiple epistatic capacitors and capacitated QTL, it is also easy to see how some of these features could lead to replication failures of these seemingly “large effect QTL” even in closely related populations, or fine mapping approaches in different generations of an intercross. The major caveat is that the system on Chromosome 4 and 7 was discoverable due to the large effect size of the

QTL, as well as pronounced and directional epistasis involved. Similar processes probably contribute to the effect or lack thereof in a multitude of loci, albeit to a much smaller degree, and a comprehensive, systematic investigation of all possible contributing interactions is far outside the statistical power-budget of any experimental population, given the exponential nature of potential interactions between identified loci.

Limiting the search space for epistasis

When I attempted to get a formal and systematic overview of the contribution of epistasis to the selection response, I tried to alleviate the combinatorial explosion problem by first limiting the search for epistasis to identified marginal QTL, and the search for capacitor hubs to vQTL as well. Despite this, it is likely that the sub-network selection for each capacitor hub was either limited by our statistical power, or compromised by other features that would have a similar impact on the stratification results, such as local haplotype effects. This resulted in rather modest effect sizes capacitated by the capacitor loci I identified, though alternatively, this could also be explained by the some of the capacitors having divergent effects across the affected loci, meaning that they may be capacitative hub loci, but do not have a strong directional effect, indicating that my selection criteria for directional capacitation hubs was suboptimal.

In contrast, the region that likely benefited most from the increased resolution and power contained another epistatic network which indeed has a strong directional effect. While the individual nodes of this network indeed show up in the stratification scan, and the main capacitor on chromosome 4 features quite prominently in the vQTL scan, I identified the relevant region on chromosome 4 already earlier, when investigating the identified QTL for stability across generations of the intercross, hoping to find an explanation for why I can identify large effect loci that have evaded detection in the earlier F_2 scans (37,38) with this intercross line. Instead of a QTL that was driven by the later generations, due to e.g. recombination between two previously linked QTL of opposite effects, this QTL was predominantly driven by the earlier generations (F_2-F_5) disappearing in the later ones. Looking at recombination events and allele-frequencies across the generations, I could identify a progressive loss of the high allele in a location that was not the location of the QTL, but later turned out to be the location of the epistatic capacitor, with the disappearing HWS allele being the “on"-allele. This initial evidence of local epistasis, taken together with historical evidence of an interaction between chromosome 4 and 7 (20,41), allowed us to piece together a complex network with multiple interacting partners, where two capacitors jointly release an effect of more than 300 grams. I think this example is indicative of the general trend in my research trying to characterise the genetic architecture of 56-day body weight:

While we have the statistical power to make general statements about the general genetic architecture, as well as discover the specific genetic architecture of specific loci, we have neither the statistical power nor adequate models to make specific statements about the general genetic architecture of a trait.

Contribution to the selection response by rare variants.

One of the more interesting features of the GWAS approach is the large fraction of phenotypic variance explained by GWAS loci not covered in the QTL scan. While some of these loci are freely segregating in both lines and not contributing to the selection response, a larger and more interesting fraction (equivalent to 9.3% of total phenotypic variance) is segregating in only one of the founding lines, contributing to the selection response and phenotypic difference between the lines. These variants are present at low minor allele frequencies, explaining why they have been missed by the QTL approach.

These could be loci that have not yet been swept to fixation or high allele-frequency difference between the founders, representing some of the still selectable genetic variation that enabled the continuous selection response observed leading up to, and ongoing for generations in the selected lines after the formation of the Intercross line.

Alternatively, these loci could still be segregating due to the strong selection on the lines, not despite it. The strong rank-based selection pressure will have initially acted upon entire RIL-like chromosomes or haplotypes from the starting population, sweeping large haplotypes to fixation or near fixation, which could result in a rapid initial decrease in allele frequency even for alleles in line with the selection pressure, if they are in linkage disequilibrium with alleles of opposite effect, resulting in selection at a later stage, as well as a stronger effect of drift.

Similarly, the aggressive selection for bodyweight in the absence of moderating factors or fitness constraints has likely also selected for variants with strong deleterious effects and a pleiotropic effect on bodyweight, with homozygous lethality, cumulative or epistatic deleterious effects likely limiting the allele-frequency of some loci.

There is evidence for selection on traits that indirectly contribute to body weight, such as behavioural changes, hormone balance and the immune system, as well as for deleterious effects of the selection, particularly in the LWS line (30) , where one would expect deleterious variants to be more correlated with the direction of selection overall, putting a stronger constraint on deleterious alleles.

It is therefore interesting to note that a large majority of these loci (10 out of 13) are still segregating in the HWS, while being fixed in the low line. This difference in potentially selectable variation is mirrored in the difference in selection response observed in the selected lines after the creation of the intercross line, with the response in the LWS slowly attenuating (30) .

Taken together, it indicates that the genetic features seen in this population are likely more representative of strong, rapid selection events from maladaption to a distant phenotype optimum, due to e.g. a rapidly changing environment, novel diseases as well as hybridisation or domestication events, rather than modest selection pressure towards a near optimum. The overall picture painted by this is a genetic architecture more messy and diverse than what can be accurately captured by a QTL approach, which assumes fixation or at least strong divergence in allele-frequency between the lines.

Apparent Epistasis, Apparent additivity

That is not to say that the GWAS approach perfectly captures the underlying genetics. The population used here violates assumptions of both the QTL and GWAS approach, with both capturing overlapping, but distinct aspects of the genetic architecture, together with artefacts based on the different assumptions. Some of my more interesting thoughts about the general genetic architecture of 56-day bodyweight likely result from my attempts (and failures!) to compare the results of both approaches. Striking for me here was firstly the failure to retain the epistatic system on chromosome 4 within the set of final GWAS markers, secondly my inability to make the backwards elimination approach workable for the QTL genotypes. Here, even the very strongest QTL failed to be retained in a meaningful number of bootstraps, despite the existence of a clear GWAS equivalent, or a more traditional stepwise backwards approach retaining them just fine. Ultimately, this points to a much stronger sensitivity to bootstraps, which could result from much stronger reliance of the signal local haplotype effects and epistasis, or stronger confounding by residual kinship than the GWAS markers.

Working within the confines of a mostly additive framework such as a basic QTL or GWAS approach, it would be easy to afford the opinion that a GWAS approach is better at describing the genetic architecture, either because it can simply explain more of the phenotypic variance (here, 23.9% in contrast to 14.6% for the QTL). This is partially because of otherwise undetected variants, or underestimation of effect sizes when still segregating variants only incompletely correlate with the founder-haplotypes. Not accounting for segregating variants can also lead to artefacts, e.g. apparent epistasis, when founder-genotypes at multiple regions are needed to tag an underlying still segregating variant.

On the other hand, being able to explain more phenotypic variance with an additive model does not necessarily mean that this additive model accurately captures the underlying genetic architecture. Given the high density of markers at all allele-frequencies in this population, it is easy to explain most epistatic interactions with a lower-frequency variant that near-perfectly tags the relevant combination of alleles. While this distinction may not be relevant if the goal is simply to use variants from the very same population for breeding

or genomic prediction, it does limit our understanding of the underlying genetics, and how much of our findings are transferable to other populations that are only partially related.

Conclusions and Future works

No shortcuts

I've had modest success in improving our understanding of how selection on 56-day bodyweight has shaped the genetic architecture of this trait, and managed to untangle some specific loci or interactions that hopefully serve as examples for the different ways genetic variation can contribute to quantitative traits in general. But thinking about the more holistic modelling of this trait specifically or quantitative traits in general, I do not think that it is a problem to be solved by better and more sophisticated models. Beyond the comparatively simple approaches I used here, I have attempted to use multiple more advanced statistical models, developed by people with a much better grasp on the underlying statistics, such as the NOIA approach (47–49), and decided not to use them each time due to concerns about the reliability and interpretability of the results. While they account for some of the specifics encountered in complex genetic architectures and imperfect populations, they do not account for all of them, often resulting in less robust estimates when some of the assumptions are violated, requiring the same if not more caution and oversight in use and interpretation as simple models. Similarly, increasing statistical power of the study population has diminishing returns in explaining the genetics underlying a trait, as the power requirements of a comprehensive, systematic scan for epistatic interactions will overwhelm the budget of any reasonably structured scientific endeavour for most quantitative traits.

From my experience, there is likely no good, single way to model quantitative traits in a generic fashion, regardless of how sophisticated the model gets. And, at least for this system, sufficient statistical power to "do it all" is neither within reach, nor a sensible use of resources. There is likely room for improvement in the selection of potential capacitors and their networks, but no improvement or novel method will get around a significant amount of manual investigation, for the same reason there is no optimal way of modelling quantitative traits: biology, and by extension genetics, is a messy, complicated business, and the emergent properties of large interconnected systems make for novel complications each time. Likely, there are no shortcuts that let you see the larger picture before you spend enough time with the puzzle pieces.

Future directions

That is not to say that investigating quantitative traits and attempting to gain a comprehensive overview is a fool's errand, but doing so by trying to perfect a single approach is likely a labour of love as much as anything. While the current trend in, e.g. machine learning, leans towards increasingly large datasets and models, I believe interpretability of results and understanding of the general features of a system can more easily be found in the intersection of multiple, different lines of enquiry. Especially so when under the constraints of a budget. Such different lines of enquiry could be as simple as the overlap of multiple interpretable and robust statistical models with different assumptions or phenotypes, but may also include making use of existing public information. Similarly, genomic summary statistics, information about other genomic features or even datasets from different populations can be applied.

End of the line

For the Virginia lines, or more specifically the investigation of 56-day bodyweight in the AIL, I believe further success will require alternative sources of information. Even if the AIL had not yet been terminated, there would be little to gain from generating more generations, more individuals, with only miniscule advancements in understanding to be gained from additional power or resolution. Already, there are more candidate loci than there is budget or personnel to follow up on, and further understanding is better to be gained from other sources of information that help identify the function, mechanism or extent of a given region. For the latter, our group had some success with elucidating the likely ancestral breeds of the White Plymouth Rock chicken, and how they contribute to different regions across the genome of the selected lines (32), which may be useful for ancestry-guided fine mapping, or identification of more complex haplotypes that cannot be captured by the founding-line paradigm.

Still, while these are likely to increase our modelling of the trait, I think that genomic, particularly comparative or functional genomics approaches would serve better to further our understanding of individual loci, and therefore in sum, our understanding of the underlying genetics in general. Likely, the most accessible and useful data for these approaches is gene annotation data. It is difficult to use gene-function from e.g. orthologues in other species to narrow down the likely causative region, since body weight as a trait has many contributing biological systems. But beyond information about potential function, information about the gene-models also allows layering of other information over the existing data, such as impact of individual variants, or sequence conservation scores across orthologues of this gene or region from multiple related species.

Another approach, instead of being hampered by them, would be to utilise the individual biological components contributing to the trait by collecting nested or closely related phenotypes that pertain to them. While maybe not always feasible for existing populations, this has been done, at least in part, for the F_8 generation, where traits such as shank-length, other weight measurements and carcass traits as well as blood chemistry have been collected. The overlap between results from mapping multiple phenotypes may increase our confidence in the relevance of a given region, could be used to identify which biological systems are driving the emergence of specific QTL, as well as narrow down regions and potential candidate variation by limiting candidate genes to a specific subphenotype. Similarly, this would also enable comparison of gene content or gene ontology enrichment with a broader set of other selection experiments, possibly identifying convergently selected genes by e.g. overlapping with the results of a selection experiment for shank-length in mice (50). In the same fashion that one can look at the overlap between the results of experimental populations with a shared phenotype, knowing the ancestry of a given experimental population also allows for comparison with results from populations that share partial ancestry with this population, such as one or multiple overlapping founder-breeds. Particularly for a commercially relevant species such as chicken, there are likely multiple experimental or commercial populations with extensive data on a variety of carcass traits which share at least partial ancestry with the White Plymouth Rock chicken. Comparison here could allow for prioritisation of the QTL that are robust to differences in the genetic background, less likely to have strong deleterious effects or be results of the strong selection pressure or unique genetic makeup of the starting population.

Finding functional explanations for Epistasis

Another promising approach is to utilise annotation data to gain insight into the functional underpinnings of statistical epistasis, either for the sake of itself, or to gain confidence in the validity of a discovered interaction. Statistical epistasis can be vexing due to the many artefacts leading to apparent epistasis, such as higher-order linkage disequilibrium in the context of insufficient marker coverage, hidden population structure or mismatch between the utilised model and the genetic architecture, and the fact that the high power requirements mean that prioritisation of a given network usually involves manual decision making, as well as prohibits exhaustive testing for confounders. Given this, any additional evidence is a significant boon for gaining confidence in the identified interaction. While this additional evidence can also be driven by gene-annotations, such as products of candidate genes taking part in similar metabolic or signalling pathways, or evidence of protein-protein interactions in orthologues, derived from large scale yeast-two-hybrid scans, an-

other likely avenue for interactions are additional datasets detailing e.g. chromosome accessibility, such as ATAC-seq or Hi-C methodologies, with identification of e.g. topology associated domains in regions suspected of within-chromosome epistatic interactions such as on chromosome 4 being strong supporting lines of evidence.

Acknowledgements

It is difficult to express the deep gratitude I feel towards the people that have helped me along the way in small or large ways. Even were I inclined towards the excessively saccharine, I doubt I could do it justice, nor could I possibly mention everyone that deserves it. However, anything worth doing is worth doing badly, so I will attempt it nonetheless.

The result of a PhD, as I have been told, is not the thesis or scientific publication I produced, but the education I received, and the person it has shaped me into. As much as this is often considered a personal achievement, In my mind, it has always been a group effort, with many people mentoring me along the way. I'd like to therefore start in a loosely reverse chronological order here, by thanking my scientific mentors which have formed the supporting framework of my education:

Örjan, the time I have spent under your supervision has shaped me from a student into a scientist. I might have been a student of genetics and bioinformatics before, but I am a bioinformatician and geneticist now. Much of this improvement has been due to your tutelage. As much as I would like to joke that you have remade me in your image, I doubt I could ever emulate your glittering, razorsharp brilliance! I would like to think your careful and thorough mentoring has helped me grow into my own instead. The advice, guidance and opinions I have received from you are a priceless gift, and I will continue to benefit from them for the rest of my days.

Mats, thank you for stepping in when I needed help. I don't doubt that I would have collapsed on the home-stretch of my PhD, were it not for your comprehensive support, advice, patience and overall sunny disposition.

Ying and **Carina**, Thank you for having my back. I knew I could always count on you.

Åke, being a stalwart and stoic support during my master thesis. I sure did a lot of hair-raising things to that data, as well as a lot of hair-pulling, yet somehow you weathered all of this with a smile.

Malin, coming to Sweden was a big thing, starting out as a small thing. It was meant to be one semester of Erasmus, remember? Instead, I am still here, nearly seven years later. My success, be it present or future, can be traced back to the opportunities you have provided to me, as well as the multitude of professional and personal advice I have received in your care.

Alisandra, who taught me my first bioinformatics, bash and the wonders of python. I don't think I would have tried to become a bioinformatician, had I not idolised you and Tony as much as I did. So in a sense I am holding you responsible for any future back pain I might incur!

Wayne and Gregor, who taught me statistics and fieldwork, and instilled in me a healthy respect and appreciation for the rigours of data collection. While I have left the fieldwork behind (for now!), the ecology I have learned from you will forever contribute to my view on science.

Frau Schmidtholz, While it's been a long time, I have not forgotten how you fostered my interest in biology, but also the confidence that, maybe, maybe, this could be something I can do.

In a similar way, people say that discovery is made, metaphorically speaking, by standing on the shoulders of giants. This is true, and maybe for my research more obvious than for most. If I were to liken my supervisors to the framework that props up the house of my doctoral education, **Paul** and his collaborators are the bedrock, having laid the foundation of my research long before I was born, with the founding of the selection experiment that eventually resulted in the data I have been using. It has been a humbling and awe-inspiring experience to work with this data and look back upon an unbroken chain of research by so many people, reaching back more than half a century. In addition to that, I would also like to thank Paul and **Christa**, for being kind, thorough and patient collaborators who taught me a lot writing my first paper with me.

To stretch the house metaphor beyond its breaking point, a house does not consist of foundation and support beams alone:

While it is likely foolish to admit so in a public document, let me tell you: I am terrible at administrative work, terrified by bureaucracy, and downright lousy at planning ahead. That I have made it this far despite that is not due to some herculean effort on my part, but because I had help. Thank you, **Tanel**, **Veronica, Rehné** and **Malin & Malin** for your help and patience.

Daniela, I don't think I would have managed to finish my Master thesis in time without your help, which would have brought this whole complicated artifice tumbling down before it even started. You were kind, caring and helped me without hesitation, having a much bigger impact on the trajectory of my life than you probably realised. I think of this often, and try to follow your example.

There has also been a vast ensemble of colleagues and friends that have enriched my work and life throughout the time we have shared here. I cannot possibly mention everyone, but believe me that every friendly word, every shared laugh matters, irregardless how fleeting. They are the vast, luminescent tapestry that makes life joyful.

Thibaut, During the time we shared an office, you were probably one of the people i spent the most time with, and every moment of it has been better

for it. Thank you for the far ranging, deep, diverse and often ridiculous discussions and shared struggles. I hope France treats you well, but whenever you are ready to build your empire in the great pacific garbage patch, let me know.

Mark, stalwart and unflappable paragon of positivity. Despite the wide vortex of calamity that has surrounded us these past years, you have been the calm eye of the storm, with a ray of sunshine to boot! Stoicism, incredible positivity or just unshakable confidence? Who knows! Either way, you have been an inspiration, a fantastic colleague and a brilliant scientist, and I am sure you will be doing great things.

Yanjun, thank you for sharing your time, ideas and thoughts with me. If I had a big brother, I would have liked him to be like you.

Simon, thank you for the interesting discussions. Rest assured that next time I see you, I will be peppering you with questions.

Mette, thank you for your kindness, time and an open ear when I needed one badly.

Yunzhou, Claudio and **Yiwen**: Thank you for the good times!

Beyond these immediate colleagues, I would like to highlight how much the general community of the **D11:3 corridor** has helped me retain my sanity. If i were to mention everyone individually that has contributed to my wellbeing and joy, this thesis would be twice as long, and i would run out of adequate words to describe their positive qualities halfway through. This corridor is a veritable clown-car of excellence, and lunchtime there has always been one of the highlights of my day.

Beyond these Friends and colleagues attached to the worksphere or otherwise already mentioned here, i have been blessed with friends near and far, that have kept me tethered to reality and a life beyond science, despite my tendencies to become the human equivalent of an isolationist hermit crab: **Iris**, **Laura**, **Gaëtan**, **Matilda**, **Les**, **Ruben**, **Phillipp**, **Susanna**, **Yuval**, **Arielle**, **Liana**, **Carles**, **Moritz** as well as **Hugo**, **Mariona**, **Alessio** and the rest of the W&D. I will not write down individual paragraphs to commemorate their valorous contributions here. This is not an oversight, but a promise. Over the course of my PhD and particularly the pandemic I have been struggling very badly to stay in contact or generally devote time to social endeavours. Consider this omission a promise to rectify this, and a placeholder for future memories.

Finally, for my family: Mama, Papa, Vivi and Susmita. Any sentence I can think of here feels entirely inadequate, too derivative and meaningless. How could it not? It feels ridiculous to pick any given paragraph to lionise my family for a specific thing, when all I am, all my achievements, is and are founded and contingent on the love, care and unwavering support I have received from you my entire life. Please allow me to reach for the most generic of platitudes: **Mama**, **Papa**, **Vivi**; Danke für alles.

Susmita, this is a bit easier. Mostly because you gave me strict instructions (“keep it short, don't make it too cheesy”), but also because you already know what you mean to me, and I am looking forward to telling you time and time again.

References

1. Christensen AC. Cats as an Aid to Teaching Genetics. Genetics. 2000 Jul 1;155(3):999–1004.
2. Smith MU, Gericke NM. Mendel in the Modern Classroom. Sci Educ. 2015 Jan 1;24(1):151–72.
3. Cao A, Galanello R. Beta-thalassemia. Genet Med. 2010 Feb 1;12(2):61–76.
4. Cutting GR. Cystic fibrosis genetics: from molecular understanding to clinical application. Nat Rev Genet. 2015 Jan;16(1):45–56.
5. McColgan P, Tabrizi SJ. Huntington's disease: a clinical review. Eur J Neurol. 2018;25(1):24–34.
6. Lilani A. Ethical Issues and Policy Analysis for Genetic Testing: Huntington's Disease as a Paradigm for Diseases with a Late Onset. Hum Reprod Genet Ethics. 2005 Jul 22;11(2):28–34.
7. Cousens NE, Gaff CL, Metcalfe SA, Delatycki MB. Carrier screening for Beta-thalassaemia: a review of international practice. Eur J Hum Genet. 2010 Oct;18(10):1077–83.
8. Simonson MA, Wills AG, Keller MC, McQueen MB. Recent methods for polygenic analysis of genome-wide data implicate an important effect of common variants on cardiovascular disease risk. BMC Med Genet. 2011 Oct 26;12(1):146.
9. Loh PR, Bhatia G, Gusev A, Finucane HK, Bulik-Sullivan BK, Pollack SJ, et al. Contrasting genetic architectures of schizophrenia and other complex diseases using fast variance-components analysis. Nat Genet. 2015 Dec;47(12):1385–92.
10. Mahajan A, Taliun D, Thurner M, Robertson NR, Torres JM, Rayner NW, et al. Fine-mapping type 2 diabetes loci to single-variant resolution using high-density imputation and islet-specific epigenome maps. Nat Genet. 2018 Nov;50(11):1505–13.
11. Perkins JM, Subramanian SV, Davey Smith G, Özaltin E. Adult height, nutrition, and population health. Nutr Rev. 2016 Mar 1;74(3):149–65.
12. Visscher PM, McEVOY B, Yang J. From Galton to GWAS: quantitative genetics of human height. Genet Res. 2010 Dec;92(5–6):371–9.
13. Boyle EA, Li YI, Pritchard JK. An Expanded View of Complex Traits: From Polygenic to Omnigenic. Cell. 2017 Jun 15;169(7):1177–86.
14. Thauer RK. Citric-acid cycle, 50 years on. Eur J Biochem. 1988;176(3):497–508.
15. Wiese KE, Nusse R, van Amerongen R. Wnt signalling: conquering complexity. Development. 2018 Jun 26;145(12):dev165902.
16. Nooren IMA, Thornton JM. Diversity of protein–protein interactions. EMBO J. 2003 Jul 15;22(14):3486–92.
17. Blazeck J, Alper HS. Promoter engineering: Recent advances in controlling transcription at the most fundamental level. Biotechnol J. 2013;8(1):46–58.

18. Berlivet S, Paquette D, Dumouchel A, Langlais D, Dostie J, Kmita M. Clustering of Tissue-Specific Sub-TADs Accompanies the Regulation of HoxA Genes in Developing Limbs. PLOS Genet. 2013 Dec 26;9(12):e1004018.
19. Carlborg Ö, Haley CS. Epistasis: too often neglected in complex trait studies? Nat Rev Genet. 2004 Aug;5(8):618–25.
20. Carlborg Ö, Jacobsson L, Åhgren P, Siegel P, Andersson L. Epistasis and the release of genetic variation during long-term selection. Nat Genet. 2006 Apr;38(4):418–20.
21. Le Rouzic A, Carlborg Ö. Evolutionary potential of hidden genetic variation. Trends Ecol Evol. 2008 Jan 1;23(1):33–7.
22. Shen X, Pettersson M, Rönnegård L, Carlborg Ö. Inheritance Beyond Plain Heritability: Variance-Controlling Genes in Arabidopsis thaliana. PLOS Genet. 2012 Aug 2;8(8):e1002839.
23. Forsberg SKG, Bloom JS, Sadhu MJ, Kruglyak L, Carlborg Ö. Accounting for genetic interactions improves modeling of individual quantitative trait phenotypes in yeast. Nat Genet. 2017 Apr;49(4):497–503.
24. Zan Y, Carlborg Ö. Dynamic genetic architecture of yeast response to environmental perturbation shed light on origin of cryptic genetic variation. PLOS Genet. 2020 Nov 5;16(5):e1008801.
25. Eichler EE, Flint J, Gibson G, Kong A, Leal SM, Moore JH, et al. Missing heritability and strategies for finding the underlying causes of complex disease. Nat Rev Genet. 2010 Jun;11(6):446–50.
26. Hill WG, Goddard ME, Visscher PM. Data and Theory Point to Mainly Additive Genetic Variance for Complex Traits. PLOS Genet. 2008 Feb 29;4(2):e1000008.
27. Siegel PB. Selection for Body Weight at Eight Weeks of Age: 1. Short term Response and Heritabilities1. Poult Sci. 1962 May 1;41(3):954–62.
28. Dunnington EA, Siegel PB. Long-Term Divergent Selection for Eight-Week Body Weight in White Plymouth Rock Chickens. Poult Sci. 1996 Oct;75(10):1168–79.
29. Márquez GC, Siegel PB, Lewis RM. Genetic diversity and population structure in lines of chickens divergently selected for high and low 8-week body weight. Poult Sci. 2010 Dec 1;89(12):2580–8.
30. Dunnington EA, Honaker CF, McGilliard ML, Siegel PB. Phenotypic responses of chickens to long-term, bidirectional selection for juvenile body weight—Historical perspective. Poult Sci. 2013 Jul 1;92(7):1724–34.
31. Corbin FH. Plymouth Rocks: Their Origin, Characteristics, Requirements, Etc., with Special Reference to the Improved Strain. Hartford, CT: Press of the Case,Lockwood & Brainard Company,; 1879.
32. Guo Y, Lillie M, Zan Y, Beranger J, Martin A, Honaker CF, et al. A genomic inference of the White Plymouth Rock genealogy. Poult Sci. 2019 Nov 1;98(11):5272–80.
33. Zan Y, Sheng Z, Lillie M, Rönnegård L, Honaker CF, Siegel PB, et al. Artificial Selection Response due to Polygenic Adaptation from a Multilocus, Multiallelic Genetic Architecture. Mol Biol Evol. 2017 Oct;34(10):2678–89.
34. Sheng Z, Pettersson ME, Honaker CF, Siegel PB, Carlborg Ö. Standing genetic variation as a major contributor to adaptation in the Virginia chicken lines selection experiment. Genome Biol. 2015 Oct 1;16(1):219.
35. Lillie M, Sheng ZY, Honaker CF, Andersson L, Siegel PB, Carlborg Ö. Genomic signatures of 60 years of bidirectional selection for 8-week body weight in chickens. Poult Sci. 2018 Mar 1;97(3):781–90.

36. Lillie M, Honaker CF, Siegel PB, Carlborg Ö. Bidirectional Selection for Body Weight on Standing Genetic Variation in a Chicken Model. G3 GenesGenomesGenetics. 2019 Apr 1;9(4):1165–73.
37. Jacobsson L, Park HB, Wahlberg P, Fredriksson R, Perez-Enciso M, Siegel PB, et al. Many QTLs with minor additive effects are associated with a large difference in growth between two selection lines in chickens. Genet Res. 2005 Oct;86(2):115–25.
38. Wahlberg P, Carlborg Ö, Foglio M, Tordoir X, Syvänen AC, Lathrop M, et al. Genetic analysis of an F2 intercross between two chicken lines divergently selected for body-weight. BMC Genomics. 2009 May 27;10(1):248.
39. Johansson AM, Pettersson ME, Siegel PB, Carlborg Ö. Genome-Wide Effects of Long-Term Divergent Selection. PLOS Genet. 2010 Apr 11;6(11):e1001188.
40. Besnier F, Wahlberg P, Rönnegård L, Ek W, Andersson L, Siegel PB, et al. Fine mapping and replication of QTL in outbred chicken advanced intercross lines. Genet Sel Evol. 2011 Jan 17;43(1):3.
41. Pettersson M, Besnier F, Siegel PB, Carlborg Ö. Replication and Explorations of High-Order Epistasis Using a Large Advanced Intercross Line Pedigree. PLOS Genet. 2011 Jul 21;7(7):e1002180.
42. Picelli S, Björklund AK, Reinius B, Sagasser S, Winberg G, Sandberg R. Tn5 transposase and tagmentation procedures for massively scaled sequencing projects. Genome Res. 2014 Dec;24(12):2033–40.
43. Brown MB, Forsythe AB. Robust Tests for the Equality of Variances. J Am Stat Assoc. 1974 Jun 1;69(346):364–7.
44. Whalen A, Gorjanc G, Hickey JM. AlphaFamImpute: high-accuracy imputation in full-sib families from genotype-by-sequencing data. Bioinformatics. 2020 Aug 1;36(15):4369–71.
45. Abramovich F, Benjamini Y, Donoho DL, Johnstone IM. Adapting to unknown sparsity by controlling the false discovery rate. Ann Stat. 2006 Apr;34(2):584–653.
46. Zan Y. r/BE [Internet]. 2017 [cited 2023 Jan 11]. Available from: https://github.com/yanjunzan/BE
47. Álvarez-Castro JM, Carlborg O. A unified model for functional and statistical epistasis and its application in quantitative trait Loci analysis. Genetics. 2007 Jun;176(2):1151–67.
48. Le Rouzic A, Álvarez-Castro JM. Estimation of Genetic Effects and Genotype-Phenotype Maps. Evol Bioinforma. 2008 Jan 1;4:EBO.S756.
49. Álvarez-Castro JM, Rouzic AL, Carlborg Ö. How To Perform Meaningful Estimates of Genetic Effects. PLOS Genet. 2008 Feb 5;4(5):e1000062.
50. Castro JP, Yancoskie MN, Marchini M, Belohlavy S, Hiramatsu L, Kučka M, et al. An integrative genomic analysis of the Longshanks selection experiment for longer limbs in mice. Nordborg M, Wittkopp PJ, Rockman MV, Kelly J, editors. eLife. 2019 Jun 6;8:e42014.
51. Zan Y, Payen T, Lillie M, Honaker CF, Siegel PB, Carlborg Ö. Genotyping by low-coverage whole-genome sequencing in intercross pedigrees from outbred founders: a cost-efficient approach. Genet Sel Evol. 2019 Aug 14;51(1):44.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Medicine 1919

Editor: The Dean of the Faculty of Medicine

A doctoral dissertation from the Faculty of Medicine, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Medicine. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Medicine".)

Distribution: publications.uu.se
urn:nbn:se:uu:diva-498417

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2023