4 Tissue-specificity of *cis* regulatory variants

Most of the knowledge on the genetic basis of regulatory variation has been gained so far by examining whole-genome expression patterns in blood-derived cells or cell lines (lymphocytes, LCLs) (Cheung, Spielman et al. 2005; Stranger, Forrest et al. 2005; Goring, Curran et al. 2007). Blood-derived cell-types continue to be the easiest accessible source for large-scale transcript level profiling, however incorporating information from a variety of other tissues is essential. Both during development as well as throughout the process of cellular differentiation, some genes are expressed ubiquitously while others display tissue-specific characteristics (Myers, Gibbs et al. 2007; Schadt, Molony et al. 2008). Additionally, many phenotypes manifest themselves only in certain tissues (Nowak and Davies 2004; Oksenberg and Baranzini 2010). Given the key role of regulatory variation in shaping complex phenotypes of medical importance, it is of special interest to assess the extent of expression differences between tissues that can be attributed to differential regulatory control.

Promising advancement towards this goal has been made recently by several studies identifying and comparing eQTLs in multiple human tissues. Myers etal. were the first to explore genetic variation influencing normal human cortical expression (Myers, Gibbs et al. 2007). The authors estimate using expression and genotypic data from 193 samples that 58% of the transcriptome is cortically expressed and of the expressed transcripts, 21% have significant eQTLs. Little overlap can be found between this eQTL set and results from previous analyses on blood-derived cells. While differences between the compared studies with respect to the samples and genotyping platforms used explain some of the modest overlap, it is very likely that variants discovered in the cortical samples underlie brain specific control of gene expression. In conjunction with results from GWAS, brain specific eQTLs could help uncover the genetic basis of some neurologic disorders.

In a study on 400 human liver samples, Schadt etal. identified more than 6000 SNP – gene associations (Schadt, Molony et al. 2008). Many of the genes detected in this experiment had already been linked to a variety of complex diseases, expectedly given

the liver's essential role in many human metabolic processes. The same expression platform employed for the human liver cohort has also been used on a set of human blood and adipose tissue samples in another study (Emilsson, Thorleifsson et al. 2008). The authors evaluated the *cis* eQTL overlap in the three tissues, estimating ~30% sharing (Schadt, Molony et al. 2008). Efforts from our group have also contributed to the understanding of regulatory variation in a cell-type specific context. In a systematic study controlling for confounding associations due to different population samples or discrepant technological and statistical methods used, eQTLs were detected and compared across LCLs, fibroblasts and T-cells derived from the same 75 GenCord individuals (Dimas, Deutsch et al. 2009). The authors report that 69-80% of all discoveries (*cis* eQTLs) are cell-type specific, highlighting the need of sampling multiple tissue expression datasets in order to describe the full repertoire of regulatory variants.

Documenting cell-type specific regulatory variation is very important from the disease perspective. Integrating expression data with GWAS results can be informative for discovering genes and pathways whose disruption likely causes disease (Chen, Zhu et al. 2008; Nica and Dermitzakis 2008; Nica, Montgomery et al. 2010). However, this is only possible when the tissue of expression is relevant to the interrogated complex trait (Nica and Dermitzakis 2008). eQTLs discovered in LCLs have helped explain GWAS associations with childhood asthma (Moffatt, Kabesch et al. 2007) and Crohn's disease (Libioulle, Louis et al. 2007), two autoimmune inflammatory disorders. The adipose and blood cohorts analyzed by Emilsson etal. had been assessed for various phenotypes too, including obesity relevant traits. Notably, 50% of the *cis* signals were estimated as overlapping between the two cohorts, but a marked correlation with obesity-related traits was only observed for gene expression measured in adipose tissue (Emilsson, Thorleifsson et al. 2008). These observations certify the importance of integrating data from a relevant tissue when trying to interpret GWAS results using gene expression as an intermediate phenotype. Nevertheless, it is still unclear what the pattern of diminishing returns is across human tissues and what tissues could serve as highly informative in large cohorts. For example, LCLs have been useful in less expected cases enabling candidate gene discovery for associations with autism (Nishimura, Martin et al. 2007) or bipolar disorder (Iwamoto, Bundo et al. 2004).

In this chapter, I further explore the complexity of the human regulatory variation landscape in LCLs and two primary tissues (skin and fat) derived from the same subset of female twins from the UK Adult Twin registry (Spector and Williams 2006). In line with previous studies, I report extensive tissue-specificity of eQTLs using both a standard association method as well as a Bayesian factor analysis model. I describe the properties of eQTLs in each tissue and I propose that continuous estimates of statistical significance as well as the direct comparison of the magnitude of effect on the fold change in expression are essential properties that jointly provide a biologically realistic view of tissue-specificity.

4.1 Abundant eQTL discoveries per tissue

The pilot MuTHER samples were genotyped and profiled for gene expression in three tissues: LCLs, skin and fat. Normalization was performed separately in each tissue (Methods). The overlapping set of successfully genotyped samples with available expression data amounted to 156 individuals for LCL (30 MZ pairs, 37 DZ pairs, 22 singletons), 160 for skin (31 MZ pairs, 37 DZ pairs, 24 singletons) and 166 for fat (31 MZ pairs, 40 DZ pairs, 24 singletons). This final dataset was used for eQTL analysis (MZ and DZ pairs per tissue - Table 4.1).

The probes on the array were mapped to Ensembl gene IDs and only a confident subset was kept for analysis (27,499 probes mapping uniquely to 18,170 Ensembl genes). 865,544 SNPs passing quality check (Methods) were tested for associations with these probes.

		MZ pairs	DZ pairs
3 tissues	LCL-SKIN-FAT	28	30
2 tissues only	LCL-SKIN	1	2
	LCL-FAT	1	5
	SKIN-FAT	2	5
1 tissue only	LCL	0	0
	SKIN	0	0
	FAT	0	0
Total		32	42

Table 4.1. Successfully genotyped twin pairs (MZ and DZ) with available gene expression data.
Number of twin pairs per tissue sharing both genotypic and expression information.

The eQTL analysis was performed separately in each tissue. I considered only unrelated individuals at a time by separating twins from the same pair and thus performing two independent eQTL analyses per tissue. This study design, hereafter named Matched Co-Twin Analysis (MCTA), permits immediate replication and validation of eQTL discoveries. This is important and unique with respect to previous eQTL studies which do not go beyond reporting the most significant findings (Morley, Molony et al. 2004; Cheung, Spielman et al. 2005). Given the known inter-individual variability in gene expression levels and the multiple sources of variation that can contribute to this, replicating the genetic determinants of expression differences (eQTLs) is essential, much like in any GWAS exercise. Furthermore, in a multiple-tissue expression design like here, where one of the main goals is to assess the extent of eQTL tissue-specificity, it is very useful to contrast between-tissue to within-tissue variability of expression changes for properly assessing the tissue-dependent level of regulatory control (section 4.4).

Spearman Rank Correlation (SRC) was used to detect associations and I restricted the search to *cis* effects located within 1Mb on either side of a gene's transcription start site (TSS). Statistical significance was assessed at different thresholds using permutations (10,000 per gene) (Methods). An abundance of *cis* eQTLs was detected in each tissue, at a comparable rate to other studies of similar sample size (Stranger, Nica et al. 2007; Dimas, Deutsch et al. 2009). At a permutation significance level of 10^{-3}, roughly 18 genes are expected to have at least one significant association by chance. At this threshold level, I detect significant associations with 509, 238 and 462 genes in LCL, skin and fat respectively for the first subset of the twin cohort (Twin 1) (Table 4.2). Unless otherwise stated, the 10^{-3} permutation cut-off corresponding to an FDR rate of 3.5% in LCL and fat and 7.5 % in skin was henceforth chosen when exploring eQTL properties.

	LCL		SKIN		FAT	
Permutation Threshold	**Twin 1**	**Twin 2**	**Twin 1**	**Twin 2**	**Twin 1**	**Twin 2**
10^{-4}	296	360	123	125	303	304
10^{-3}	509	556	238	231	462	488
10^{-2}	1014	1059	605	676	982	1068

Table 4.2. *Cis* eQTL associations detected with SRC analysis. Significant discoveries (number of genes with eQTLs) are shown at different permutation thresholds for each tissue. Within each tissue, two independent eQTL analyses were performed after separating related individuals in two subsets (Twin1, Twin2).

Compared to LCL and fat, proportionally less eQTLs were detected in skin, at all levels of significance. This is likely due to lower power in skin, which is a more heterogeneous tissue and consists of a variety of cell-types (Sorrell and Caplan 2004; Leek and Storey 2007).

The MCTA study design allows replication of eQTL discoveries in each tissue. Replication was assessed using the mean value of the proportion of true positives (π_1) (see Methods and (Storey and Tibshirani 2003)) estimated from the exploration of significant eQTLs in the reciprocal co-twin. Specifically, significant SNP - gene combinations discovered in the first co-twin are tested in the second co-twin and the nominal SRC p-value distribution of the same initial associations is analyzed. The reciprocal test (SNP - gene associations discovered in co-twin 2 tested in co-twin 1) is also performed. The enrichment of low p-values from the distribution described above is used to estimate π_1. For each tissue, the mean π_1 of the two reciprocal tests is reported (Table 4.3). The discovered eQTLs appear robust as they replicate well between individuals of the two co-twin groups per tissue, with a mean proportion of true positives from 0.93 for skin to 0.98 for LCL and fat. I also checked the proportion of true positives specifically among the subset of genes that do not replicate in the co-twin at the same threshold. This too is high (π_1 = 0.84 for skin and 0.94 for LCL and fat), suggesting that exact overlap of genes at a given permutation threshold (PT) is an underestimate of eQTL replication due to winner's curse i.e. I see eQTLs in the co-twin that clearly replicate the initial findings, but at higher p-value and thus marginally not meeting the initial discovery threshold.

	SRC analysis			SRC-FA analysis		
	Twin 1	Twin 2	Replication (Mean π_1)	Twin 1	Twin 2	Replication (Mean π_1)
LCL	509	556	0.98	1068	1226	0.97
SKIN	238	231	0.93	534	544	0.95
FAT	462	488	0.98	1054	1072	0.97

Table 4.3. Replication of *cis* eQTL discoveries (number of significant genes per tissue at 10^{-3} permutation threshold). Results from both the Spearman Rank Correlation (SRC) and Factor Analysis (SRC-FA) are presented. Proportion of replicating signals calculated as the mean co-twin π_1 estimates from the p-value distribution of same SNP-gene associations in the reciprocal twin set

4.2 Substantial increase in number of eQTLs per tissue by Factor analysis

The observed variation in gene expression is not entirely due to genetic effects. Experimental noise and environmental conditions also affect transcript levels in a global manner. Therefore, it is desirable to remove the effects of such random variables and thus increase the power to detect eQTLs. For this purpose, factor analysis (FA) was employed on each tissue separately (Stegle, Parts et al. 2010). We corrected for global latent effects on all individuals in each tissue and fitted various parameters such as number of learned factors and proportion of variance explained, in order to maximize for replication of eQTLs per tissue between twin sets (Methods).

After performing standard SRC eQTL analysis on the factor-corrected expression data (SRC-FA), a substantial improvement in eQTL discovery at each of the standard permutation thresholds used was obtained (Table 4.4). The MCTA design is useful as it permits the validation of the new eQTL discoveries in the replication co-twin for each tissue separately. This is essential in order to verify that FA performs as expected by modelling environmental factors and not correcting out a vast proportion of genetic effects. The improvement in eQTL discovery with SRC-FA is considerable (twice as many eQTLs at 10^{-3} PT) and consistent in all three tissues. The high eQTL replication between twin sets persists after FA, with an additional improvement of true positives detection in skin: $\pi_1 = 0.95$ (Table 4.3).

	LCL		SKIN		FAT	
Permutation Threshold	**Twin 1**	**Twin 2**	**Twin 1**	**Twin 2**	**Twin 1**	**Twin 2**
10^{-4}	721	828	329	344	690	720
10^{-3}	1064	1220	532	542	1052	1070
10^{-2}	1839	1967	1103	1080	1732	1812

Table 4.4. *Cis* eQTL associations detected with SRC-FA analysis. Number of genes with a significant eQTL is shown for each co-twin analysis per tissue at different permutation thresholds.

I validated the results of the FA correction by investigating the eQTLs resulting from the SRC-FA analysis. As expected, FA recovers the majority of eQTLs discovered with the initial analysis (roughly 90% of LCL and fat and 80% of skin results) and allows the discovery of additional signals (Table 4.5).

	Twin 1			Twin 2		
	Total Std	**FA recovered (%)**	**Total FA**	**Total Std**	**FA recovered (%)**	**Total FA**
LCL	509	460 (90.37%)	1064	556	494 (88.85%)	1220
SKIN	238	189 (79.41%)	532	231	188 (81.39%)	542
FAT	462	421 (91.13%)	1052	488	436 (89.34%)	1070

Table 4.5. Recovery of SRC eQTLs (10^{-3} PT gene associations) with factor analysis correction. In each tissue and for both co-twins, 80-90% of eQTLs detected before correction (standard analysis - Std) are recovered with SRC-FA.

The additional eQTLs likely represent real effects that could not be detected initially due to low power. To test this hypothesis, the eQTLs revealed only after FA correction were tested in the uncorrected expression dataset The p-value distribution of the exact same SNP – gene combinations showed a highly significant enrichment of low values (Figure 4.1). In each tissue and for each co-twin subset, the estimated enrichment corresponded to a π_1 value of 0.99. This confirms that the vast majority of new eQTLs are real and would be picked up using the standard SRC pipeline if a larger sample size would be available.

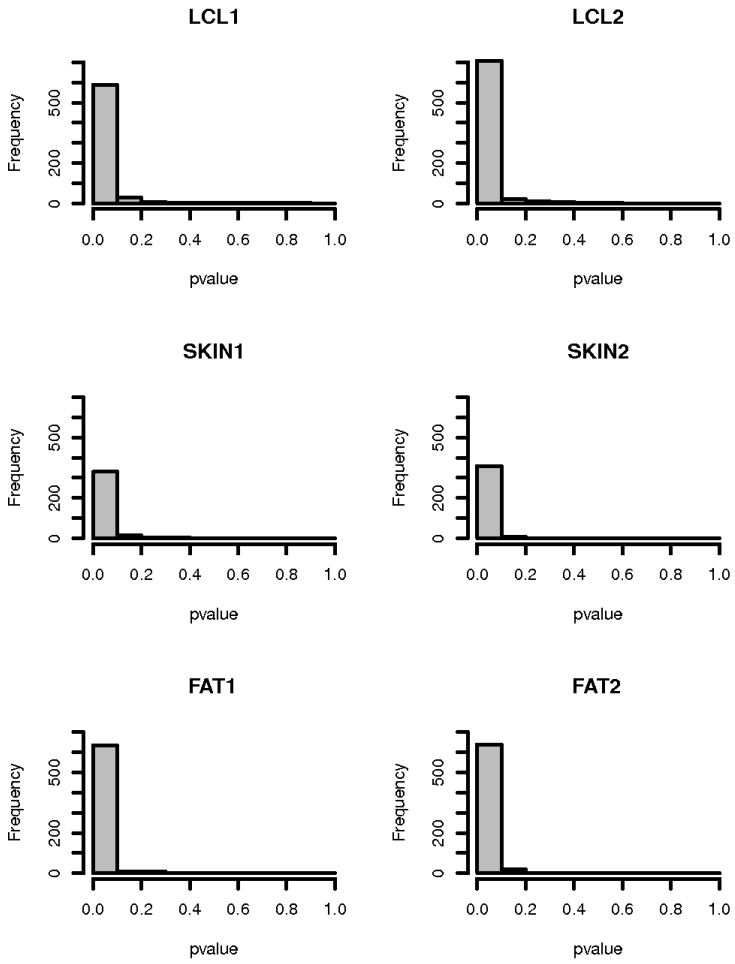

Figure 4.1. P-value distribution of *cis* eQTLs (10^{-3} PT) gained with FA correction in the uncorrected data. The significant overrepresentation of low p-values for the new eQTLs (π_1 = 0.99) shows that the signal existed in the uncorrected data but wasn't called significant due to low power. Result consistent in all tissues for both sets of co-twins (Twin 1 – left panel, Twin 2 – right panel).

4.3 eQTL properties across tissues

The eQTLs (10^{-3} permutation threshold) resulting from both SRC and SRC-FA analyses were compared across all three tissues. Initial direct tissue overlap of significant eQTLs supports an extensive level of tissue-specificity with very similar proportion in both detection methods employed.

A visual representation of the percentages of eQTLs found in only one tissue, shared in only two tissues and common in all three tissues for SRC and SRC-FA respectively can be seen in Figure 4.2.

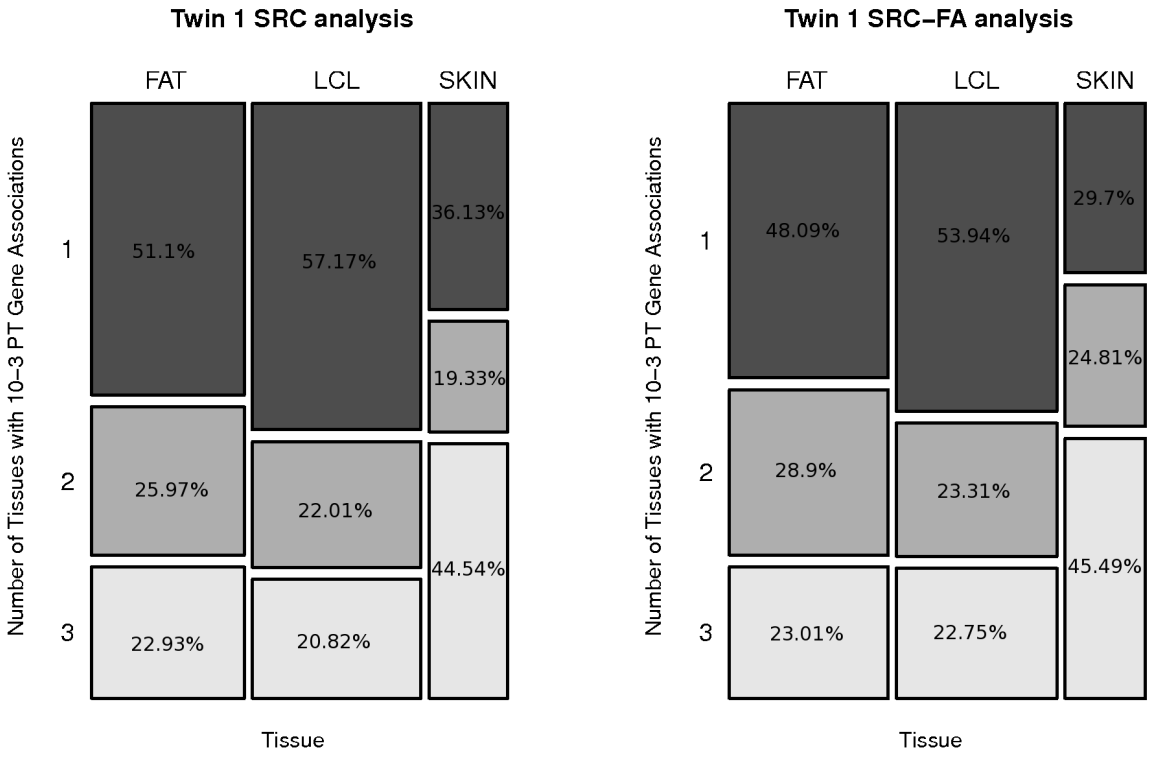

Figure 4.2. Percentage of eQTLs (10^{-3} PT) found only in one tissue, only in two tissues and in all three tissues with the SRC and SRC-FA analysis respectively. Both methods reveal similarly high extents of tissue-specificity. Skin specific eQTLs of smaller effects are harder to detect due to low power.

In the first co-twin set we discover 858 non-redundant eQTL genes at 10^{-3} PT in all three tissues (Table 4.6). Of these, 106 genes (12.35%) are shared across all tissues, 139 (16.2%) are shared in at least two tissues and 613 genes (71.44%) are detected in only one tissue. In skin, where we are least powered likely due to tissue heterogeneity and variety of cell-types, we detect proportionally fewer tissue-specific effects (10.02% of skin eQTLs are only present in skin at 10^{-3} PT).

		Twin 1			Twin 2	
		10^{-3} PT	% total	Overlap	10^{-3} PT	% total
3 tissues	LCL-SKIN-FAT	106	12.35	78	102	11.02
2 tissues only	LCL-SKIN	19	2.21	4	12	1.29
	LCL-FAT	93	10.84	52	107	11.56
	SKIN-FAT	27	3.15	11	26	2.81
1 tissue only	LCL	291	33.92	150	335	36.18
	SKIN	86	10.02	17	91	9.82
	FAT	236	27.5	103	253	27.32
Total significant	LCL	509		363	556	
	SKIN	238		132	231	
	FAT	462		304	488	
Union of total significant		858	100	563	926	100

Table 4.6. Tissue-shared and tissue-specific gene associations (10^{-3} PT), SRC analysis

SRC-FA results confirm the estimated ~30% of eQTLs to be shared in at least two tissues based on threshold eQTL discovery (Table 4.7).

		Twin 1			Twin 2	
		10^{-3} PT	% total	Overlap	10^{-3} PT	% total
3 tissues	LCL-SKIN-FAT	242	13.28	192	270	13.86
2 tissues only	LCL-SKIN	38	2.09	8	42	2.16
	LCL-FAT	210	11.53	84	232	11.91
	SKIN-FAT	94	5.16	28	70	3.59
1 tissue only	LCL	574	31.5	302	676	34.7
	SKIN	158	8.67	51	160	8.21
	FAT	506	27.77	221	498	25.56
Total significant	LCL	1064		781	1220	
	SKIN	532		338	542	
	FAT	1052		735	1070	
Union of total significant		1822	100	1312	1948	100

Table 4.7. Tissue-shared and tissue-specific gene associations (10^{-3} PT), SRC-FA analysis.

In the currently examined three tissues, shared eQTLs show the same allelic direction of effect (Figure 4.3), i.e. if one SNP allele predisposes to increased levels of expression of a gene, it will also tend to elevate the expression level of that gene in the other tissue. This is true for both eQTLs significant at 10^{-3} and 10^{-2} PT.

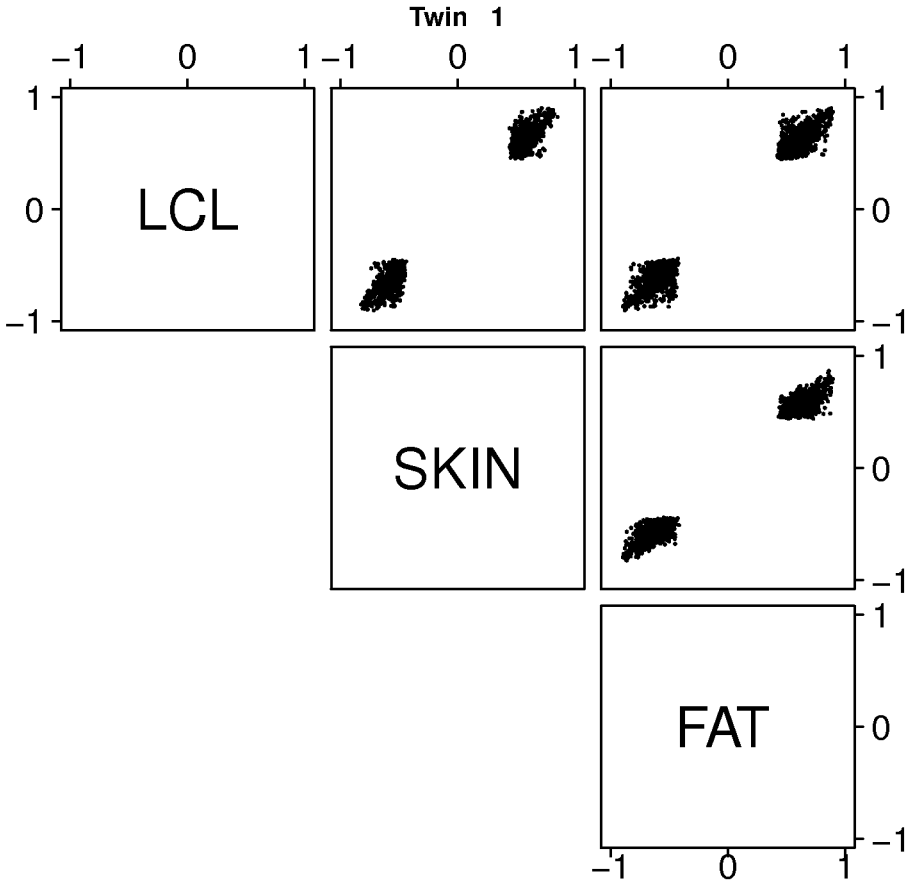

Figure 4.3. Shared eQTLs (10^{-2} PT, SRC) have the same direction of effect (SRC rho) across tissues

As reflected by the SRC correlation coefficient rho (Figure 4.4), eQTLs significant in one tissue explain a substantially higher fraction of gene expression variation in the tissue of discovery than in other tissues (same SNP-gene association), whereas shared effects at the same significance threshold (10^{-3} PT) have comparable variance explained by the SNP across tissues.

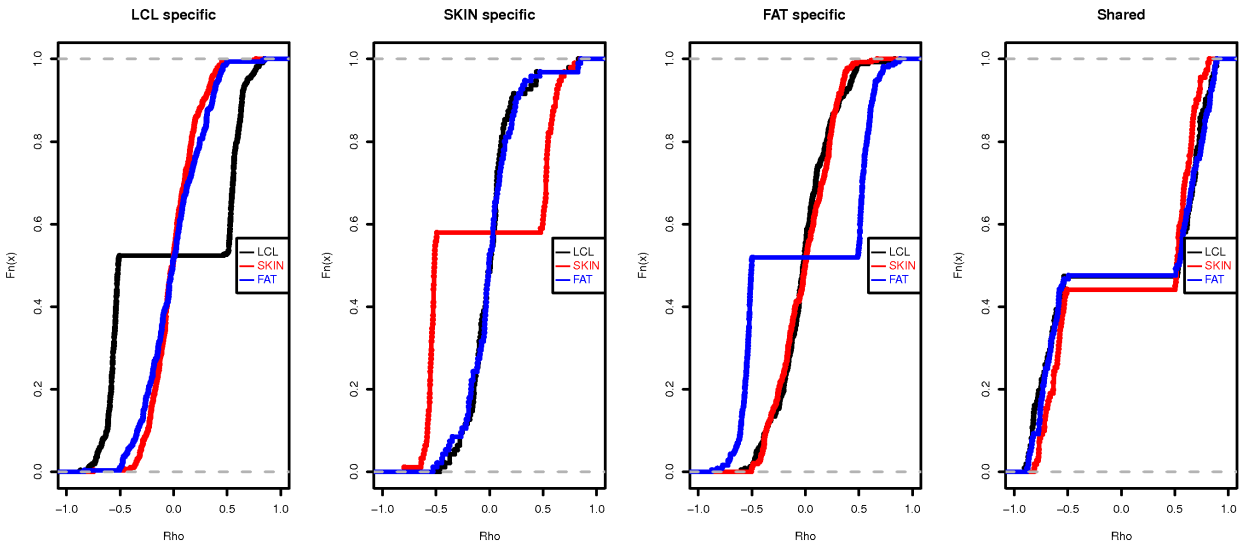

Figure 4.4. Cumulative SRC rho distribution across tissues for tissue-specific and shared eQTLs (10-3 PT, Twin 1). eQTLs discovered in one tissue only have distinctively higher variance in the tissue of discovery compared to shared effects.

In order to refine the expression association signals and describe independently acting eQTLs, I mapped them to recombination hotspot intervals and filtered subsequently by LD (Methods). I observe in all tissues that the majority of genes (90-95%) are controlled by single independent *cis* eQTLs with similar estimates from the standard and factor eQTL analysis. The finer comparison of eQTL effects requiring the sharing of both the gene and the genomic interval harboring the eQTL SNP yields similar counts of shared and specific effects (Table 4.8). The results are similar for SRC-FA. This suggests that the vast majority of shared genes also share regulatory variants across tissues.

		Twin 1			Twin 2	
		10^{-3} PT	% total	Overlap	10^{-3} PT	% total
3 tissues	LCL-SKIN-FAT	104	10.86	70	96	9.44
2 tissues only	LCL-SKIN	17	1.77	5	14	1.37
	LCL-FAT	90	9.39	49	103	10.13
	SKIN-FAT	30	3.13	12	26	2.56
1 tissue only	LCL	339	35.39	151	374	36.77
	SKIN	101	10.54	18	106	10.42
	FAT	277	28.91	100	298	29.3
Total significant	LCL	550		348	587	
	SKIN	252		128	242	
	FAT	501		302	523	
Union of total significant		958	100	565	1017	100

Table 4.8. Tissue-shared and tissue-specific interval-gene associations (10^{-3} PT), SRC analysis.

Furthermore, the genomic location of the independent eQTLs with respect to basic gene structure landmarks was investigated. Similar results to previous studies are observed (Dimas, Deutsch et al. 2009). As such, eQTLs cluster symmetrically around the TSS, with shared effects distributed more tightly compared to specific ones (Figure 4.5). The broader distribution of cell-type specific effects around the TSS suggests their role on tissue-specific enhancer elements. Independent eQTLs gained with FA correction were also investigated. It was found that they have the same pattern as the SRC eQTLs, supporting furthermore their likely biological role.

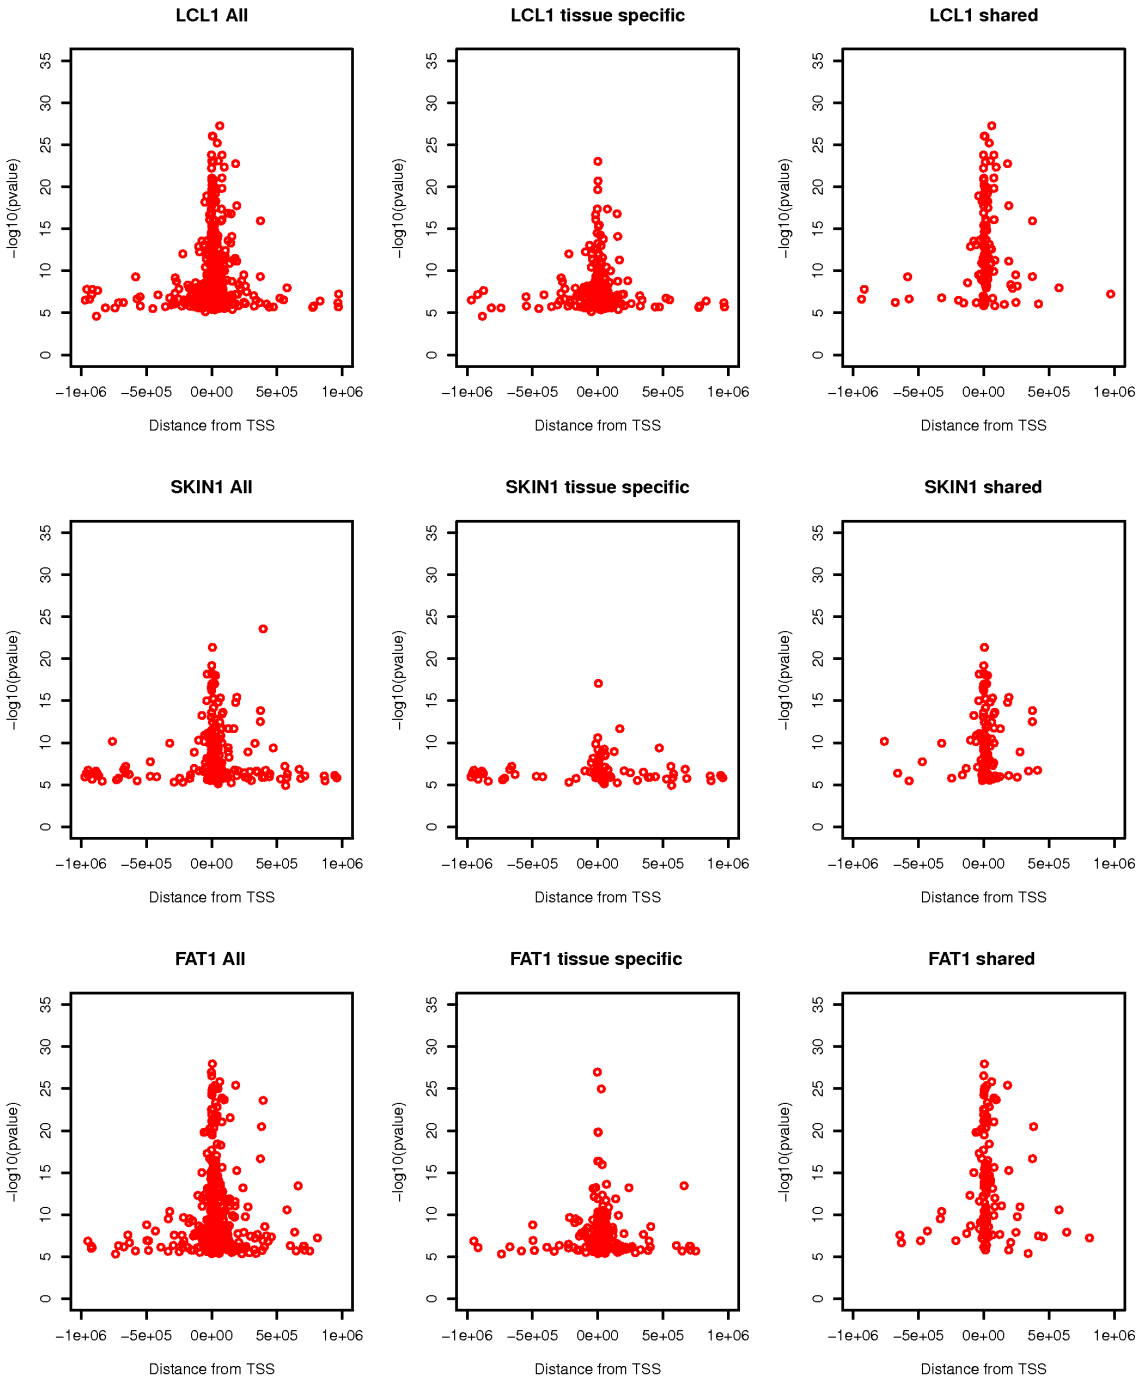

Figure 4.5. Distribution of independent *cis* eQTLs (10^{-3} PT, SRC) around the transcription start site (TSS). Data from co-twin 1 shown here; left panel displays all eQTLs, the middle panel includes only tissue-specific eQTLs while the right panel shows only eQTLs shared across all three tissues. Similar results are obtained for co-twin 2 and the independent eQTLs revealed by the SRC-FA analysis.

4.4 Alternative estimates of eQTL tissue-specificity

Thresholds are driven by statistical significance and overlaps at these levels are heavily dependent on power. In addition, eQTLs sharing statistical significance may still have notable effect differences on gene expression levels across tissues, with potentially different biological consequences. Given these caveats, I examined tissue-specificity in a continuous way using the estimate of significant low p-value enrichment (π_1). More specifically, I investigated the p-value distribution of significant SNP-gene pairs (10^{-3} PT) from a reference tissue in the other two tissues. The p-value distribution in the other two tissues suggests a high degree of tissue sharing (53 to 80%) both with the SRC and SRC-FA, varying slightly depending on the reference tissue in the comparison (Table 4.9). This indicates that we are still underpowered to detect eQTLs of smaller effects that would increase also the previous threshold-based estimates of tissue sharing. In any case, 29% of eQTLs (1-mean π_1) are expected to be exclusively tissue-specific.

#Twin 1			
Reference	**Secondary**	**SRC analysis π_1**	**SRC-FA analysis π_1**
LCL	SKIN	0.67	0.71
	FAT	0.73	0.77
SKIN	LCL	0.77	0.67
	FAT	0.72	0.84
FAT	LCL	0.63	0.72
	SKIN	0.73	0.78
#Twin 2			
Reference	**Secondary**	**SRC analysis π_1**	**SRC-FA analysis π_1**
LCL	SKIN	0.53	0.66
	FAT	0.73	0.75
SKIN	LCL	0.72	0.71
	FAT	0.8	0.84
FAT	LCL	0.69	0.58
	SKIN	0.81	0.76

Table 4.9. Continuous estimates of tissue sharing by enrichment of low p-values (π_1) of reference eQTLs (SNP-genes 10^{-3} PT) in the secondary tissues.

Tissue sharing should not just be the common presence of a statistically significant regulatory effect, but also the similar effect size (fold change in expression) of that variant across tissues. In this respect, I report the fold change as the difference between the gene expression means of the two homozygous genotypic classes. Within the same tissue, the two co-twin sets are only slightly different in their fold change estimates (0.94 Pearson's correlation of fold change between Twin 1 and Twin 2 in LCL, 0.80 in skin and 0.90 in fat – Figure 4.6). This difference in estimated effect size is much more apparent however between tissues (LCL eQTLs have a 0.65 and 0.72 fold change correlation with skin and fat eQTLs respectively). To a large extent, this is due to the tissue-specificity of eQTLs. However, shared eQTLs at the same threshold of significance don't always share the same effect size across tissues, suggesting additional possible hidden tissue-specific effects (LCL fold change correlation of 0.72 in skin and 0.77 in fat for shared eQTLs i.e. 20% difference in fold change magnitude between tissues compared to within tissue difference). This suggests that even statistically tissue shared eQTLs have additional dimensions of tissue-specificity and their mere discovery in multiple tissues does not guarantee similar magnitude of consequences.

The extent of these observations remains to be tested in *trans* in the better-powered full MuTHER dataset (N ~ 800 individuals). Here, an extension of the MCTA design will be most valuable. Building co-expression networks for each tissue will allow the discovery of tissue-specific modules, which combined with genotypic information could uncover further aspects of tissue-specific regulatory control. The topologies of the networks resulting from such approaches are however highly dependent on the methods and parameters used. Therefore, cross-validating the network predictions with the reciprocal co-twin will ensure that only genetically-relevant gene expression modules are compared.

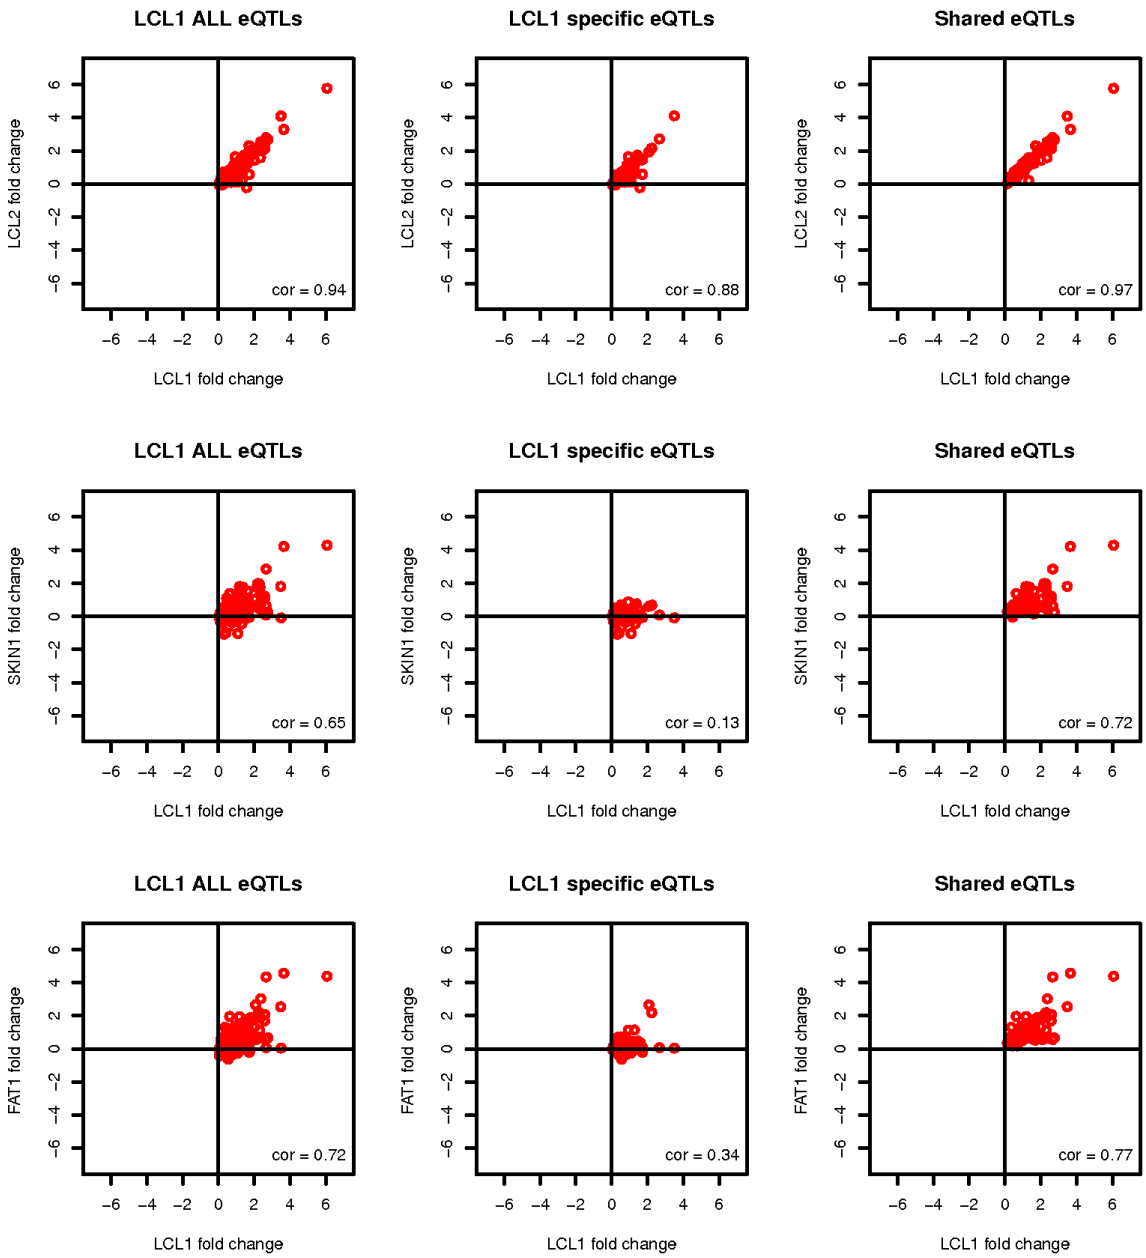

Figure 4.6. Fold change within twins and across tissues for LCL eQTLs (10^{-3} PT, SRC) discovered in Twin 1. The plotted fold change on the X and Y-axes was calculated as the difference in mean expression of homozygous genotypic classes. For each pairwise tissue comparison, the Pearson's correlation coefficient between fold changes is shown.

4.5 Conclusions

While there have been studies exploring regulatory variation in one or more tissues, the complexity of tissue-specificity in multiple primary tissues is not yet well understood. In this chapter, I explored in depth the role of regulatory variation in three human tissues: LCL, skin, and fat. The samples (156 LCL, 160 skin, 166 fat) were derived simultaneously from a subset of well-phenotyped healthy female twins of the MuTHER resource. An abundance of eQTLs in each tissue was discovered, similar to previous estimates (858 or 4.7% of genes). In addition, factor analysis (FA) was applied by removing effects of latent variables, increasing the power by at least 2-fold (1822 eQTL genes). The unique study design (Matched Co-Twin Analysis – MCTA) permits immediate replication of eQTLs with co-twins (93-98%) and validation of the considerable gain in eQTL discovery after FA correction. It was observed that the majority (>90%) of genes are regulated by single independent eQTLs with shared direction of effect across different tissues and their spatial distribution around basic gene structure landmarks was described. I highlight the challenges of comparing eQTLs between tissues and after verifying previous significance threshold-based estimates of extensive tissue-specificity, I show their limitations given their dependency on statistical power. Instead, I propose that continuous estimates of statistical significance and direct comparison of the magnitude of effect on the fold change in expression are essential properties that jointly provide a biologically realistic view of tissue-specificity. Under this framework, this study shows that 30% of eQTLs are shared among tissues, while another 29% are likely exclusively tissue-specific. However, even among the shared eQTLs a substantial proportion (10-20%) have significant differences in the magnitude of fold change between homozygote classes across tissues. These results underline the need to account for the complexity of eQTL tissue-specificity in an effort to assess consequences of such variants for complex traits.