Cell Genomics

Technology

FastGxC: Fast and powerful context-specific eQTL mapping in bulk and single-cell data

Graphical abstract

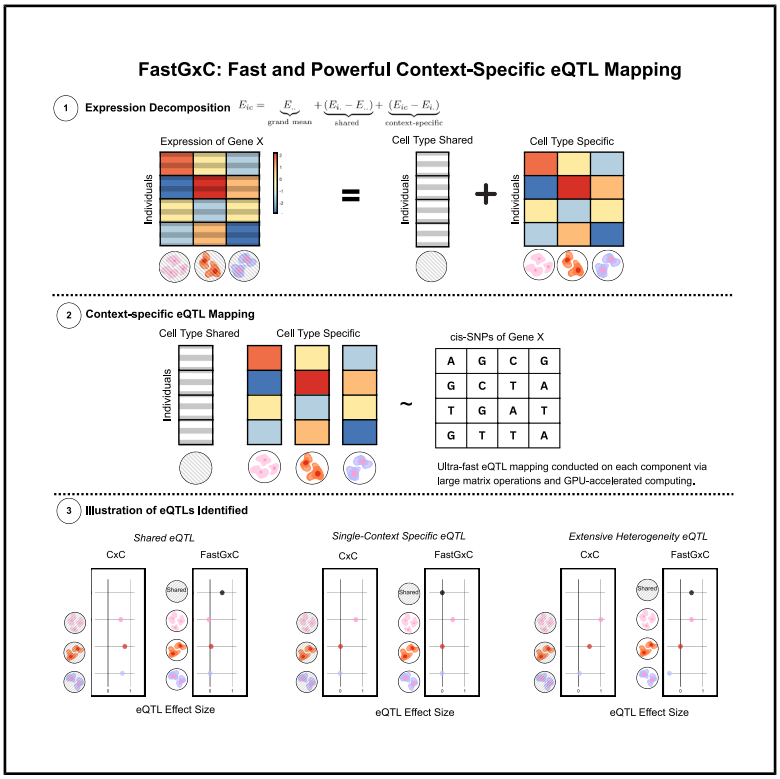

Highlights

- FastGxC is 9 times more powerful and 10^6 times faster than existing approaches
- Context-specific eQTLs are pervasive and driven by effect-size heterogeneity
- FastGxC triples precision in linking GWAS variants to relevant tissues and cell types
- FastGxC expands the number of candidate genes by 25% in cell types and by 6% in tissues

Authors

Lena Krockenberger, Andrew Lu, Mike Thompson, ..., Chun Jimmie Ye, Noah Zaitlen, Brunilda Balliu

Correspondence

bballiu@ucla.edu

In brief

Krockenberger, Lu, et al. introduce FastGxC, a method that identifies how genetic variants regulate gene activity in specific tissues and cell types more efficiently than existing approaches. Applied to large human datasets, FastGxC uncovers context-specific regulatory patterns that shed light on the genetic underpinnings of complex diseases.

Krockenberger et al., 2026, Cell Genomics 6, 101250
June 10, 2026
https://doi.org/10.1016/j.xgen.2026.101250

CellPress

Cell Genomics

CellPress
OPEN ACCESS

Technology

FastGxC: Fast and powerful context-specific eQTL mapping in bulk and single-cell data

Lena Krockenberger,[1,2,16] Andrew Lu,[3,16] Mike Thompson,[4] Lise A. Tucker,[2] Cuining Liu,[1,5] M. Grace Gordon,[6] Amanda Ramste,[7,8] Ivan Carcamo-Orive,[9,10] Joshua W. Knowles,[8] Andy Dahl,[11] Chun Jimmie Ye,[12] Noah Zaitlen,[5,13,14] and Brunilda Balliu[2,13,15,17,*]

[1]Bioinformatics Interdepartmental Graduate Program, University of California, Los Angeles, Los Angeles, CA 90095, USA
[2]Department of Pathology and Laboratory Medicine, University of California, Los Angeles, Los Angeles, CA 90095, USA
[3]UCLA-Caltech Medical Scientist Training Program, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA 90095, USA
[4]Systems and Synthetic Biology, Centre for Genomic Regulation, The Barcelona Institute for Science and Technology (BIST), 08003 Barcelona, Spain
[5]Department of Human Genetics, University of California, Los Angeles, Los Angeles, CA 90095, USA
[6]Biological and Medical Informatics Graduate Program, University of California, San Francisco, San Francisco, CA 94143, USA
[7]Institute for Molecular Medicine Finland (FIMM), HiLIFE, University of Helsinki, 00290 Helsinki, Finland
[8]Division of Cardiovascular Medicine, Cardiovascular Institute, Diabetes Research Center, Stanford University School of Medicine, Stanford, CA 94305, USA
[9]Department of Endocrinology, Metabolism, Nutrition, and Kidney Disease, Biobizkaia Health Research Institute, 48903 Barakaldo, Spain
[10]IKERBASQUE, Basque Foundation for Science, 48009 Bilbao, Spain
[11]Section of Genetic Medicine, University of Chicago, Chicago, IL 60637, USA
[12]Division of Rheumatology, Department of Medicine, University of California, San Francisco, San Francisco, CA 94143, USA
[13]Department of Computational Medicine, University of California, Los Angeles, Los Angeles, CA 90095, USA
[14]Department of Neurology, University of California, Los Angeles, Los Angeles, CA 90095, USA
[15]Department of Biostatistics, University of California, Los Angeles, Los Angeles, CA 90095, USA
[16]These authors contributed equally
[17]Lead contact
*Correspondence: bballiu@ucla.edu
https://doi.org/10.1016/j.xgen.2026.101250

SUMMARY

Context-specific expression quantitative trait loci (eQTLs) mediate genetic risk for complex diseases, but current methods limit their characterization and interpretation. We introduce FastGxC, a method for efficiently mapping context-specific eQTLs by leveraging correlation structure in multi-tissue bulk and single-cell RNA sequencing studies. In simulations, FastGxC is nine times more powerful and 10^6 times faster than existing approaches, reducing computation time from years to minutes. We applied FastGxC to bulk multi-tissue (n = 698) and peripheral blood mononuclear cell (PBMC) single-cell RNA sequencing datasets (n = 1,218), generating comprehensive tissue- and cell-type-specific eQTL maps. These eQTLs showed 4-fold enrichment in context-matched open chromatin and were twice as enriched as standard context-specific eQTLs. FastGxC improved precision in identifying relevant trait contexts by 3-fold and expanded candidate causal genes by 25% in cell types and by 6% in tissues. FastGxC provides a powerful framework for mapping context-specific eQTLs, advancing our understanding of gene regulatory mechanisms underlying complex traits.

INTRODUCTION

Genome-wide association studies (GWASs) have identified hundreds of thousands of genetic variants linked to complex traits and diseases.[1] A majority of these variants reside in non-coding regions, often overlapping DNA regulatory elements,[2] which suggests that their functional effects are mediated through transcriptional regulation.[2–5] This observation has driven significant efforts to identify expression quantitative trait loci (eQTLs)—genetic variants associated with gene expression changes—and use them to link GWAS variants to their regulatory targets.[6–15] Despite these efforts, only 21% of GWAS variants, on average per trait, overlap with known *cis*-eQTLs from bulk tissues,[12] underscoring a persistent gap between genetic associations and regulatory function.[16–18]

A key factor in this missing regulation is the context-specific nature of many disease-relevant eQTLs,[16,18] which often show stronger effects in specific tissues,[12] cell types,[10,14,15] or environmental conditions,[19–23] making them difficult to detect. While we use "context-specific" in alignment with common genomics

terminology,[24,25] the more precise term would be context-dependent eQTLs, as these show genotype-by-context (GxC) interaction effects rather than being entirely absent in other contexts. In contrast, broadly shared eQTLs, while easier to detect, are less enriched for GWAS variants,[12,14] most likely due to negative selection.[16,18] Another major factor is that many bulk[8,12] and all single-cell RNA sequencing (RNA-seq) studies rely on repeated sampling, where the same donor provides samples across multiple contexts. While this design minimizes experimental variability, it introduces intra-individual correlation, which, if unaccounted for, inflates the type I error rate in identifying eQTLs and reduces the power to test whether the eQTL is context specific.

Several methods have been developed to identify context-specific eQTLs in studies with repeated sampling (see Table S1). These methods fall into two broad categories. The first comprises approaches that jointly analyze data across contexts and test for context-specific eQTLs by incorporating a GxC interaction term. This includes (generalized) linear mixed model (LMM)-based methods,[24–29] which model the GxC effect linearly, and methods that capture non-linear GxC effects.[30] To account for repeated measurements, these methods include a random effect for the individual or cell. While powerful, their mixed model framework makes them computationally intensive for large eQTL studies. This challenge is particularly exacerbated when modeling all cell types jointly. Additionally, some of these methods[24,30] infer latent cellular contexts, further increasing computational costs.

The other category includes methods that follow a two-step process: first, they map eQTLs separately in each context (context by context [CxC]), then they define context specificity by post hoc examination of eQTL summary statistics across contexts.[10,12,15,31–33] While CxC approaches are fast, particularly those developed for (pseudo)-bulk data,[10,12,15,33] they have major limitations. First, CxC approaches can be significantly underpowered because they do not fully leverage all available data.[34] Second, many rely on ad hoc definitions of context-specificity based on subjective thresholds for effect-size differences between contexts[33,34] or the significance of an eQTL in a single context.[10,12,15] These definitions can lead to both false positive context specificity (e.g., when effects in certain contexts fail to reach significance due to chance or uneven power across contexts) and false negative context specificity (e.g., when an eQTL reaches significance in multiple contexts, leading it to be classified as shared despite having meaningful differences in effect size). Taken together, these limitations constrain the interpretation of disease-associated variants, as current methods fail to fully capture the complexity of context-specific regulatory variation.

To address these challenges, we introduce FastGxC, a statistical method that efficiently maps context-specific eQTLs while accounting for repeated sampling. In brief, FastGxC decomposes gene expression into context-shared and context-specific components and estimates genetic effects on these components using linear regression. We show, both analytically and empirically, that FastGxC's eQTL effect estimates can be viewed as computationally efficient reparametrizations of those obtained using CxC and LMM-GxC approaches. FastGxC has several key advantages over previous methods. First, it directly maps specific eQTLs without the need for post hoc analyses or arbitrary thresholds. Second, by accounting for intra-individual correlation, it adjusts for background noise and confounding factors unrelated to the context of interest, e.g., sex, age, population stratification, or sequencing batch,[35–37] maximizing power to detect context-specific eQTLs (Figure S1). Third, FastGxC leverages ultra-fast implementations of linear regression models[38,39] and GPU acceleration,[40] similar to those used in CxC eQTL mapping approaches, thereby reducing computational time from years to minutes. FastGxC can work on any continuous molecular phenotype, and its output integrates naturally with methods developed to improve the statistical power of eQTL mapping, such as mash.[34]

We first show in simulations that FastGxC is as powerful as LMM-GxC approaches but orders of magnitude faster. Both approaches significantly outperform CxC-based methods to map context-specific eQTLs. We then applied FastGxC to multi-tissue bulk RNA-seq data from the GTEx Consortium[12] (n = 698 individuals) and single-cell peripheral blood mononuclear cell (PBMC) RNA-seq data from the CLUES[14] (n = 237) and OneK1K[15] (n = 981) cohorts, which we meta-analyze, to produce comprehensive tissue- and cell-type-specific eQTL maps across 49 tissues and 8 PBMC types. FastGxC context-specific eQTLs show up to 4-fold enrichment in open chromatin regions from matched contexts and are twice as enriched as standard context-specific eQTLs, highlighting their biological relevance. We further examine their relationship with complex traits and diseases, showing that FastGxC eQTLs improve precision in identifying relevant GWAS contexts by 3-fold and expand candidate causal genes by 25% in cell types and by 6% in tissues compared to standard eQTLs. In addition, FastGxC context-specific eQTLs show a 1.2-fold increase in colocalization with complex traits compared to shared eQTLs, providing evidence that context-specific regulation helps explain regulatory mechanisms of complex diseases that remain unaccounted for by shared eQTLs. In summary, FastGxC provides a powerful framework for constructing context-specific eQTL maps, offering key insights into the gene regulatory mechanisms underlying complex human diseases.

DESIGN

FastGxC method overview

We illustrate the FastGxC method using tissues as contexts (Figure 1A), but it can be applied to any set of discrete contexts, for example, cell types[10,14,15] or environmental stimuli, sampled across overlapping individuals. FastGxC works in three steps. First, for each individual i (i = 1, …, N) and context c (c = 1, …, C), FastGxC decomposes the expression of each gene (E_{ic}) into two components: a shared component $\left(E_i^{sh}\right)$, representing the average expression across contexts, and a context-specific component $\left(E_{ic}^{sp}\right)$, representing the residual expression in the contexts after subtracting the shared component (Figure 1A, decomposition step), i.e., $E_{ic} = E_i^{sh} + E_{ic}^{sp}$. This decomposition, analogous to repeated-measures ANOVA, removes shared eQTL effects and shared noise from the context-specific components, thereby increasing power to detect context-specific eQTLs (Figure S1).[41]

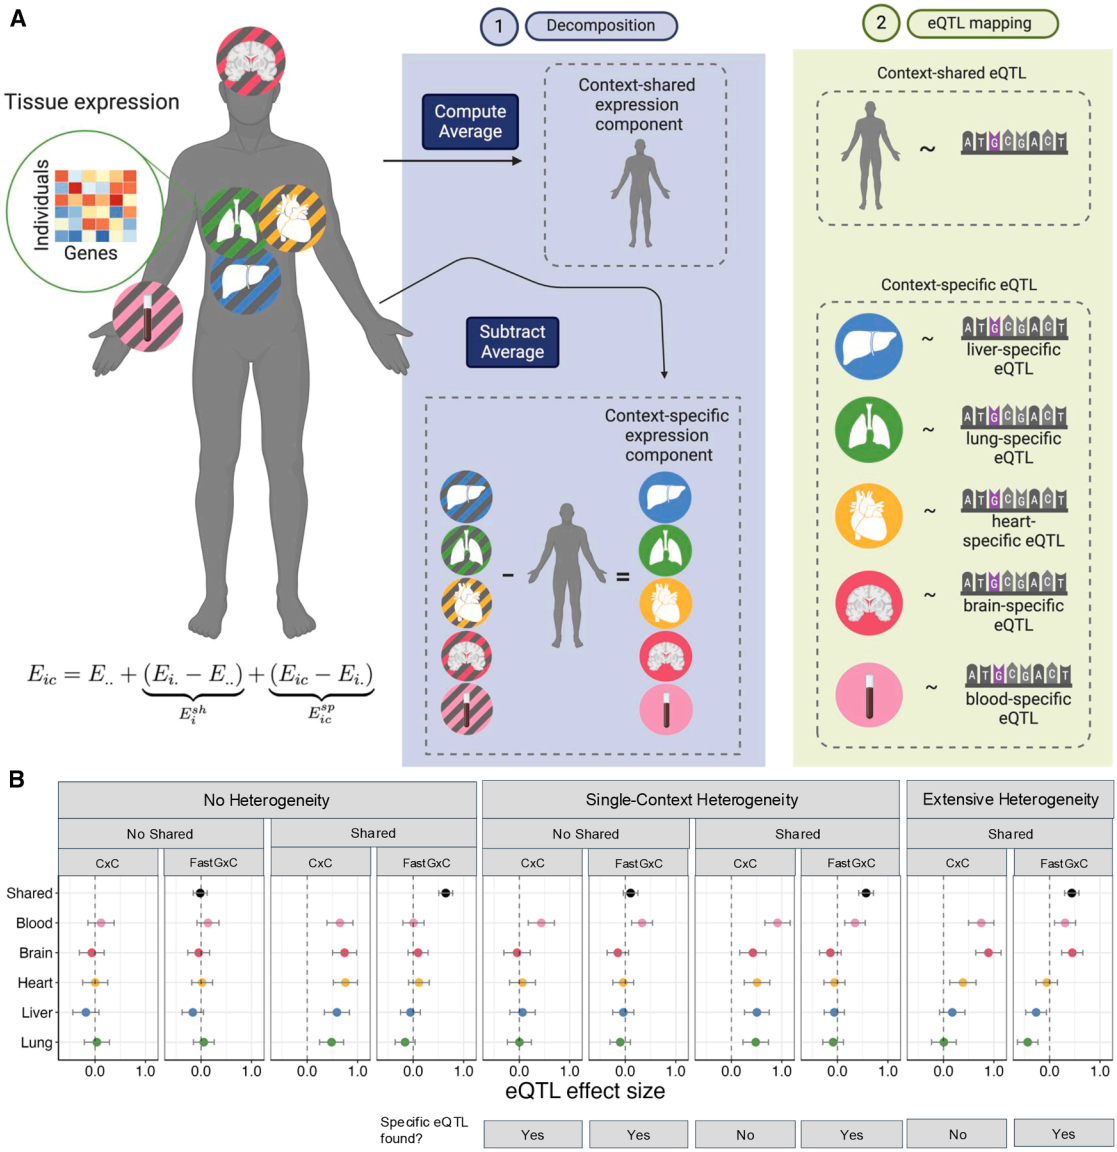

Figure 1. Overview of the FastGxC method and illustrative examples of eQTL effect patterns

(A) FastGxC decomposes gene expression for each individual into a context-shared component and context-specific components (step 1). It then estimates both the shared eQTL effect across contexts and the context-specific eQTL effect within each context by regressing these components on genotypes (step 2). See also Figure S2 for multiple testing adjustment.

(B) Illustrative examples of shared and context-specific eQTL effect patterns. The *x* axis represents the eQTL effect size, and the *y* axis and color indicate the context. Note that negative (positive) context-specific effects indicate effects weaker (stronger) than the cross-context average, not negative (positive) eQTL effects. The first example represents a scenario with no eQTL effects in any context and, thus, no shared or specific eQTLs. The second example represents a scenario with equal eQTL effects across all contexts, corresponding to a shared eQTL with no context-specific effects. The remaining examples illustrate scenarios where one or multiple contexts drive effect-size heterogeneity, resulting in context-specific eQTLs.

Next, for each gene-*cis*-SNP pair, FastGxC estimates a shared eQTL effect (β^{sh}) and *C*-specific eQTL effects $(\beta_1^{sp}, \ldots, \beta_C^{sp})$ by regressing the shared expression component and each of the *C* context-specific components on the genotype at the *cis*-SNP (Figure 1A, eQTL mapping step). This step employs ultra-fast implementations of linear regression models optimized for eQTL mapping[38] and GPU acceleration,[40] enabling computational efficiency comparable to standard eQTL mapping methods. The shared and context-specific effects represent a reparametrization of the eQTL effects obtained from conventional CxC eQTL mapping (β_1, …, β_C) (Figures 1B and S16). The total eQTL effect in each context, as estimated by CxC eQTL mapping, can be recovered by summing the shared and specific effects $\beta_c = \beta^{sh} + \beta_c^{sp}$. While point estimates are mathematically identical to CxC (supplemental methods), FastGxC provides improved

precision for context-specific effects by removing shared variance. Specifically, the shared effect corresponds to the mean eQTL effect size across contexts, i.e., $\beta^{sh} = \frac{1}{C}\sum_{c=1}^{C}\beta_c$, while the context-specific effects capture the residual eQTL effects in each context after accounting for the shared effect, i.e., $\beta_c^{sp} = \beta_c - \beta^{sh}, \forall c \in \{1,\ldots,C\}$. This decomposition separates the pleiotropic (shared) effect of an eQTL across all contexts from the context-specific effects, enabling clearer interpretation of context-specific genetic effects. Because CxC itself can be viewed as a reparametrization of the LMM-GxC framework, FastGxC provides a computationally efficient reparametrization of LMM-GxC. Full details of the analytical derivation are provided in the STAR Methods and supplemental methods.

Finally, to account for multiple testing across genes, SNPs, and contexts, FastGxC employs the hierarchical false discovery rate (FDR)-controlling procedure implemented in Peterson et al.[42] (STAR Methods; Figure S2). FastGxC defines a gene-SNP pair as an eQTL if the SNP has a significant effect on the shared or any of the specific components of gene expression, i.e., if the global null hypothesis $H_0 : \beta^{sh} = \beta_1^{sp} = \beta_2^{sp} = \cdots = \beta_C^{sp} = 0$ is rejected. If an eQTL is detected, FastGxC defines a context-specific eQTL as an SNP with a significant effect on at least one of the specific components of expression of the gene, i.e., if the global null hypothesis $H_0^{sp} : \beta_1^{sp} = \beta_2^{sp} = \cdots = \beta_C^{sp} = 0$ is rejected. This global test directly identifies context-specific eQTLs, eliminating the need for post hoc analyses or arbitrary thresholds. Finally, if significant eQTL effect-size heterogeneity is detected, FastGxC conducts C marginal tests to determine the specific context(s) driving the observed heterogeneity, i.e., $H_0^c : \beta_c^{sp} = 0 \quad \forall c \in \{1,\ldots,C\}$. Note that these tests do not specifically flag the contexts with non-zero eQTL effects; rather, they detect contexts whose effect sizes deviate significantly from the shared effect.

To illustrate how FastGxC identifies different patterns of context specificity, Figure 1B shows toy examples of eQTLs with varying patterns of effects across contexts. The first two images (“no heterogeneity”) depict scenarios under the null hypothesis of no eQTL effect-size heterogeneity across contexts, with either no shared eQTL effect (“no shared”) or a shared eQTL effect (“shared”). FastGxC does not classify either of these eQTLs as context specific, as there is no significant heterogeneity of their effects across contexts, but it would classify the second scenario as a (shared) eQTL. The remaining images illustrate scenarios under the alternative hypothesis of eQTL effect-size heterogeneity. These include heterogeneity driven by a single context (“single-context heterogeneity”) or heterogeneity spanning all contexts (“extensive heterogeneity”). FastGxC identifies all these scenarios as context-specific eQTLs, irrespective of the presence of a shared eQTL effect. By contrast, the commonly used CxC approach defines context-specific eQTLs as variants with significant eQTL effects in only a single context. As a result, this approach would classify only the first alternative scenario (“single-context heterogeneity - no shared”) as a context-specific eQTL and would overlook more complex patterns of heterogeneity, such as cases where heterogeneity exists alongside a shared effect or where heterogeneity is distributed across multiple contexts, highlighting its limitations compared to FastGxC.

RESULTS

FastGxC outperforms existing methods in simulation studies

We used a series of simulated scenarios to evaluate the performance of FastGxC in detecting an eQTL and determine if the eQTL effect is context specific as a function of intra-individual residual correlation (see STAR Methods and Table S2). In each scenario, we varied the number of individuals and contexts and the proportion of missing expression data. Some scenarios were designed to reflect those in GTEx[12] and the OneK1K cohort,[15] two of the largest bulk and single-cell RNA-seq studies, while others used complete data. Figure 2 below displays simulation results with parameters reflecting GTEx (n = 698 individuals, 49 contexts, and ~50% missing data). Results for all other parameter combinations are shown in Figures S3–S17.

The performance of FastGxC was systematically compared to three commonly used approaches: (1) the CxC approach, which performs CxC eQTL mapping and defines a context-specific eQTL as a variant with a significant effect in a single context; (2) MetaTissue,[43] a multi-tissue eQTL mapping method that combines mixed models and meta-analysis and defines a context-specific eQTL as a variant with a posterior probability greater than 0.9 that has an effect present in exactly one context; and (3) the LMM (LMM-GxC) approach, which includes a random intercept for individuals to account for intra-individual residual correlation and defines a context-specific eQTL based on the significance of the GxC interaction term (see STAR Methods). To illustrate the impact of ignoring intra-individual correlation on the identification of context-specific eQTLs, we also include the performance of a linear model with a GxC interaction term (LM-GxC) but no random intercept. Note that due to the large computational burden of MetaTissue, we did not obtain results for all scenarios with a larger sample size (N = 698).

We first evaluated the global type I error rates of each method for detecting an eQTL (Figures 2A and S3) and for testing whether an eQTL is context specific (Figures 2A and S4). FastGxC is well calibrated across all tested scenarios and in both tests. LMM-GxC is generally calibrated but becomes inflated in settings with small sample sizes and high missing-data rates. As expected, both the CxC and LM-GxC approaches, which do not account for intra-individual correlation, are miscalibrated (Figure S3). As intra-individual correlation increases, LM-GxC becomes increasingly inflated for eQTL detection and increasingly conservative for testing context specificity. The CxC approach, by contrast, remains mostly calibrated when testing for the presence of an eQTL. However, depending on sample size, the missing-data rate, and the presence or absence of a shared eQTL, CxC becomes either increasingly conservative (Figure 2A) or anti-conservative (Figure S4) when testing for (single-)context specificity. Finally, MetaTissue is consistently conservative in scenarios with a small sample size (Figure S3) when testing for the presence of an eQTL. However, when evaluating context specificity, MetaTissue becomes anti-conservative or conservative, depending on whether a shared eQTL effect is present (Figure S4).

Next, we evaluated the global power of each method to identify an eQTL. Among the calibrated methods, FastGxC and

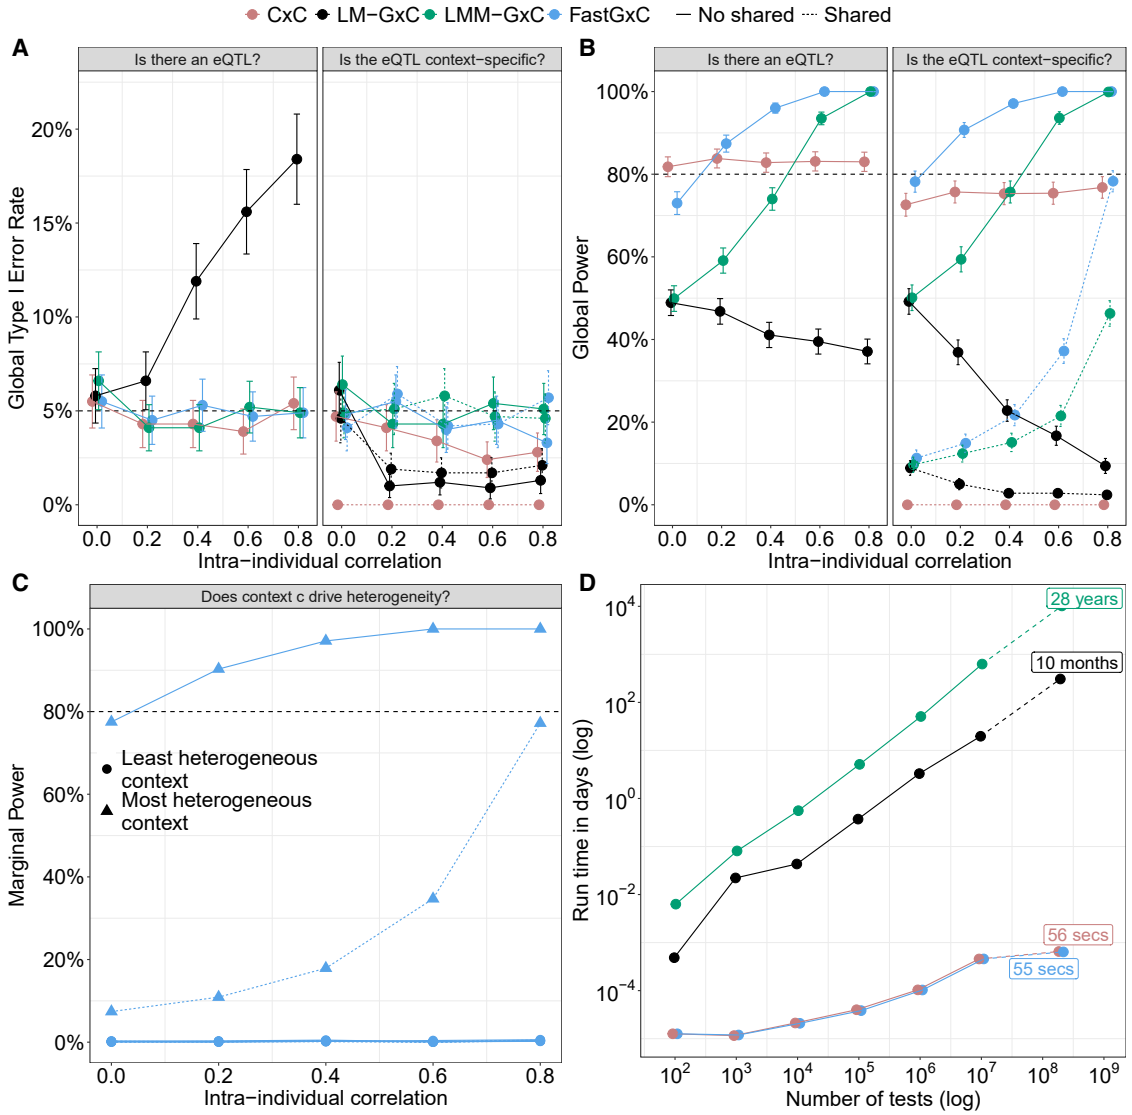

Figure 2. FastGxC outperforms existing methods in simulated data

(A–C) Global type I error rate (A) and global and marginal power (B and C) for detecting an eQTL (left images in A and B), testing for context specificity of its effect (right images in A and B), and identifying which context drives the heterogeneity (C) under the no-heterogeneity (A) and single-context-heterogeneity (B and C) scenarios (Figure 1B) across different levels of intra-individual correlation (rows). (A) and (B) include error bars representing 95% confidence intervals. (A)–(C) show results from simulations with 698 individuals and 49 contexts and GTEx missing-data patterns (63%). For effect sizes in each scenario, see Table S2. For results under the two-context and extensive-heterogeneity scenarios, with less or no missing data, lower sample sizes, and fewer contexts, see Figures S3–S14. (D) Runtime for all methods across a varying number of tests performed in a sample size of 250 individuals (average sample size across tissues in GTEx). The last points reflect the projected runtime for the entire GTEx dataset: 50 contexts, 25 K × 3 M tests, and 250 samples per context. Analyses were run on 8 cores on a 2.70 GHz Intel Xeon Gold Processor on the UCLA Hoffman2 Computing Cluster. See Figure S17 for results with a sample size of 1,000 individuals.

LMM-GxC exhibit complementary strengths. FastGxC is generally more powerful when eQTL effect-size heterogeneity is strong or driven by a few contexts (Figures 2B and S6–S8), as it leverages the Simes' method to combine *p* values. LMM-GxC is more powerful when heterogeneity is weak but spread across many or all contexts (Figures S9 and S10), due to its reliance on the likelihood ratio test (LRT). The CxC and MetaTissue approaches are less powerful than FastGxC in identifying an eQTL across all scenarios with non-zero intra-individual correlation, with the power advantage of FastGxC increasing as correlation rises. The LM-GxC method is miscalibrated for eQTL detection; therefore, we do not report or discuss its power to map an eQTL.

For all methods, power to detect an eQTL also depends on whether variability in expression is explained by shared or context-specific effects. When the shared eQTL effect explains all (no heterogeneity, Figure S5) or most ("single-context heterogeneity - shared," Figure S6) of the expression variability, power to identify an eQTL declines as intra-individual correlation increases, whereas the opposite occurs when the specific eQTL effect explains all or most variability (single-context heterogeneity - no

shared, Figure S6). In scenarios with intermediate shared and specific effects, power to identify an eQTL follows a U-shaped relationship with intra-individual correlation (extensive heterogeneity, Figure S9).

We next examined power to assess eQTL context specificity. For both FastGxC and LMM-GxC, power increases as intra-individual correlation increases, regardless of whether a shared eQTL effect is present. Notably, FastGxC is more powerful than the CxC and MetaTissue approaches even in the single-context heterogeneity scenario without a shared effect—the scenario in which these approaches are specifically used. As expected from its performance under the null, the LM-GxC method loses power to test for eQTL context specificity as intra-individual correlation increases (Figures 2B and S6–S10).

Next, to evaluate the ability of FastGxC to identify specific contexts driving effect-size heterogeneity, i.e., contexts most different from the shared effect, we examined the marginal type I error rate and power per context. Under the null hypothesis (no heterogeneity), FastGxC is calibrated in each context (FDR $\leq$ 5%, Figure S11). Under the alternative hypothesis, power was highest for contexts with effect sizes farthest from the shared effect and increased with intra-individual correlation (Figures 2C and S12–S14). For example, in single-context heterogeneity scenarios, FastGxC accurately identifies the context with the non-zero (no-shared scenario) or single strongest (shared scenario) eQTL effect (Figures 2C and S12).

In addition, we assessed FastGxC's parameter estimation accuracy by evaluating its ability to estimate the shared (i.e., β^{sh}) and specific (i.e., $\beta_c^{sp} = \beta_c - \beta^{sh}$) eQTL effect sizes, as well as the overall eQTL effect sizes in each context (i.e., $\beta_c = \beta^{sh} + \beta_c^{sp}$; Figures S15 and S16). FastGxC provided unbiased estimates for the shared, specific, and overall eQTL effect sizes under conditions with no missing data or with missing-data levels typical of single-cell RNA-seq studies such as OneK1K and CLUES ($\sim$5%). When the proportion of missing data was high (mean of 63% and up to 84% in some contexts), FastGxC estimates remained largely unbiased, with only slight deviations from the true effect in contexts with a high missing rate.

Finally, we benchmarked the computational costs of FastGxC against other approaches. To obtain practical runtimes, we used study parameters from GTEx, i.e., $\sim$50 contexts and an average of 250 individuals per context, while varying the number of tests performed (Figure 2D). When extrapolated to the entire GTEx dataset, which involves 200 million tests for 25,000 genes and 3 million SNPs, we estimated that LMM-GxC and LM-GxC would require $\sim$30 years and $\sim$10 months, respectively, to complete. In contrast, CxC and FastGxC completed the same task in under 1 min on average (based on 100 iterations). Even at a larger scale with 1,000 individuals, FastGxC remained computationally efficient, completing all tests in $\sim$5 min, whereas LMM-GxC was estimated to take over 500 years (Figure S17). MetaTissue was not included in the runtime analysis due to its substantial computational burden, which exceeds that of LMM-GxC.

Context specificity of eQTLs is widespread across tissues and PBMC types

We applied FastGxC to two datasets: multi-tissue bulk RNA-seq data from the GTEx consortium (n = 698 individuals, 49 tissues)[12] and single-cell PBMC data from the CLUES[14] and OneK1K[15] cohorts (n = 237 and n = 981 individuals, respectively, across 8 cell types). For PBMCs, we performed meta-analysis across cohorts to map *cis*-eQTLs and assess cell-type specificity (STAR Methods). Before quantifying *cis* regulation, we confirmed that FastGxC reduces background noise: the top principal components (PCs) of the decomposed expression data showed minimal correlation with technical and biological covariates compared to the original GTEx data (Figure S1).

We then assessed the extent of *cis* regulation and how context specific these effects are across tissues and cell types (Figure 3). We identified a total of 24,196 eGenes across tissues (70.21% of tested genes) and 4,564 eGenes across cell types (29.05% of tested genes), defined as genes with at least one eQTL in any context (hierarchical FDR [hFDR] $\leq$ 5%, Table S4). The majority of FastGxC eGenes (86.5% and 82.7% across tissues and cell types) had at least one shared eQTL (Figure 3A), aligning with previous observations of widespread *cis* regulation and eQTL sharing.[12,34] Despite extensive sharing, effect sizes varied substantially between contexts, with 72.1% of tissue eGenes and 63.9% of PBMC-type eGenes harboring at least one context-specific eQTL (Figure 3A). Notably, most of these context-specific eQTLs overlapped with shared eQTLs (81.3% and 73.1% across tissues and cell types), suggesting that context specificity often arises from effect-size heterogeneity rather than the presence of an eQTL in a single context (Figure 3A).

We next aimed to determine how many and which context(s) drive the effect-size heterogeneity for eGenes with context-specific eQTLs (Figures 3B and 3C; Table S3). In both bulk and single-cell data, we observed that the majority of specific eQTLs are identified in only a few contexts (Figure 3C). In tissues, much of the heterogeneity is driven by the testis (6,124 eGenes), followed by whole blood (5,219 eGenes), consistent with findings from studies mapping eQTLs specific to a single tissue.[12] Testis, which is biologically distinct from other GTEx tissues, also contributes the largest proportion (16%) of single-context-specific eQTLs—i.e., eQTLs unique to a single tissue—highlighting FastGxC's ability to detect biologically meaningful context-specific regulation (Figure 3C). For PBMC types, CD4 cells (1,907 eGenes) and classical monocytes (980 eGenes) account for the majority of heterogeneity and also contain the highest number of eQTLs unique to a single cell type. Both the number of specific eGenes and single-context specific eGenes per context are strongly correlated with the number of samples per tissue and the number of cells per cell type (Figure S18A), indicating that we may not yet have reached saturation in identifying these context-specific regulatory effects.

We next compared the eGenes identified by FastGxC to those identified by the CxC approach. Consistent with our simulation results, FastGxC identified substantially more eGenes than CxC, detecting an additional 2,159 eGenes in bulk tissues and 679 eGenes in PBMC types (Figure S18B). Broadly, most (96.6% in tissues and 62.8% in PBMCs) FastGxC shared eQTLs that overlapped with eQTLs detected in multiple contexts by CxC (Figure 4A). Importantly, FastGxC single-context-specific eQTLs mapped almost exclusively to single-context-specific eQTLs detected by CxC, demonstrating strong concordance in these cases. Furthermore, FastGxC total effect sizes (shared +

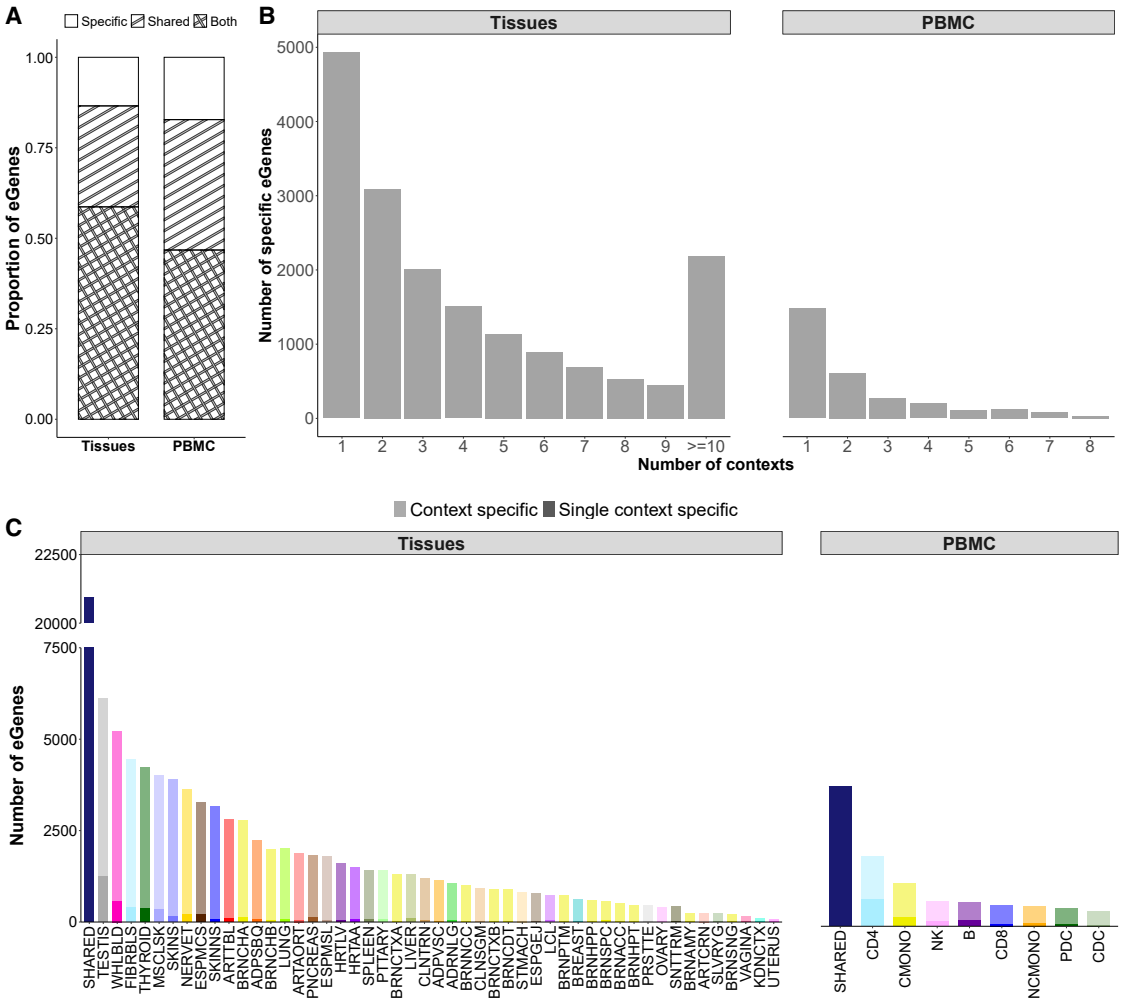

Figure 3. Context-specific eQTL mapping in bulk tissues and single-cell PBMC types

(A) Percentage of eGenes with shared-only ("shared"), specific-only ("specific"), or both shared and specific ("both") eQTLs across all tissues (left) and PBMC types (right).

(B) Number of contexts that drive effect-size heterogeneity for eGenes with context-specific eQTLs across tissues (left) and PBMC types (right).

(C) Number of eGenes with shared and context-specific eQTLs per context. For eGenes with context-specific eQTLs, color opacity indicates whether specific eQTL effects are shared with other contexts (lightest opacity) or unique to that context (darkest opacity). Tissue- and cell-type abbreviations are explained in Table S3.

See also Figure S18 and Table S4.

specific) are nearly identical to CxC effect sizes (Pearson r = 0.98 in tissues, r = 1.0 in PBMCs; Figure S19), empirically confirming their mathematical equivalence (supplemental methods). However, a substantial fraction (48.1% in tissues and 43.1% in PBMC types) of CxC single-context-specific eQTLs corresponded to FastGxC shared-only eQTLs. This discrepancy reflects the false-positive-specific effects that the CxC approach tends to identify—an issue we also observed in simulation results (Figure S4). Moreover, the number of context-specific eQTLs detected by CxC showed a stronger correlation with sample size than those detected by FastGxC (Figure S18A), further highlighting its sensitivity to power differences across contexts. These results underscore the limitations of CxC approaches that define context specificity solely by the presence of significant eQTLs in isolated contexts rather than accounting for heterogeneity in effect sizes.

Finally, we show that context-specific eQTL effect sizes are correlated within groups of biologically related tissues and cell types. For example, we see that context-specific eQTL effects are correlated among 13 brain tissues, two heart (left ventricular and atrial appendage), two artery (tibial and aorta), two esophagus (muscularis and gastro-esophageal junction), three adipose (visceral, subcutaneous, and breast), and two intestine tissues (Figure 4B, right triangle). In addition, context-specific eQTL effects are most strongly correlated across CD4 and CD8 cells, natural killer (NK) cells, and B cells between plasmacytoid and conventional dendritic cells, as well as between classical and non-classical monocytes. Furthermore, while FastGxC

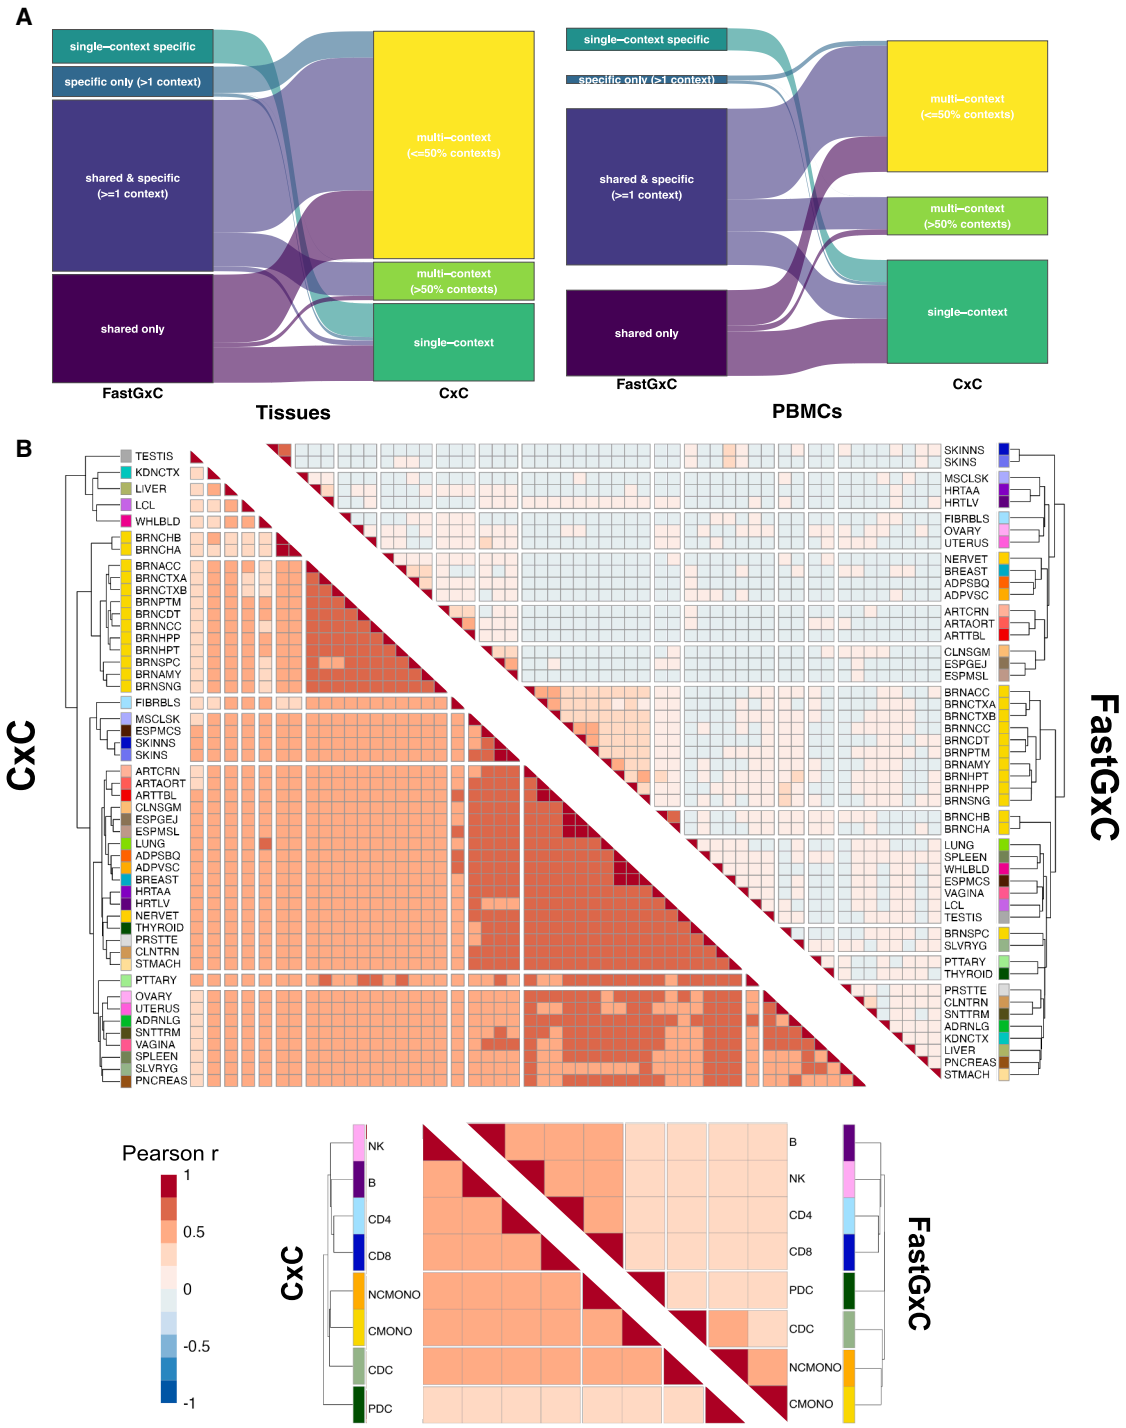

Figure 4. FastGxC context-specific eQTLs reveal distinct regulatory patterns across tissues and PBMC types

(A) Sankey diagrams showing how FastGxC eQTL categories map to CxC eQTL categories in tissues (left) and PBMCs (right). FastGxC categories are mutually exclusive: shared only (significant shared effect but no specific effects), shared and specific (significant shared and specific effects in more than one context), specific only (no shared effect, specific effects in more than one context), and single-context specific (no shared effect, specific effect in exactly one context). CxC categories are mutually exclusive: multi-context (>50%) (eQTLs detected in >50% of contexts), multi-context (≤50%) (eQTLs detected in ≤50% of contexts), and single context (eQTLs detected in exactly one context).

(B) Heatmap of Pearson's correlation of CxC eQTL effect sizes (left) and FastGxC context-specific eQTL effect sizes (right) across tissues (top) and PBMC types (bottom). Tissue- and cell-type abbreviations are explained in Table S3.

See also Figures S18 and S19 and Table S4.

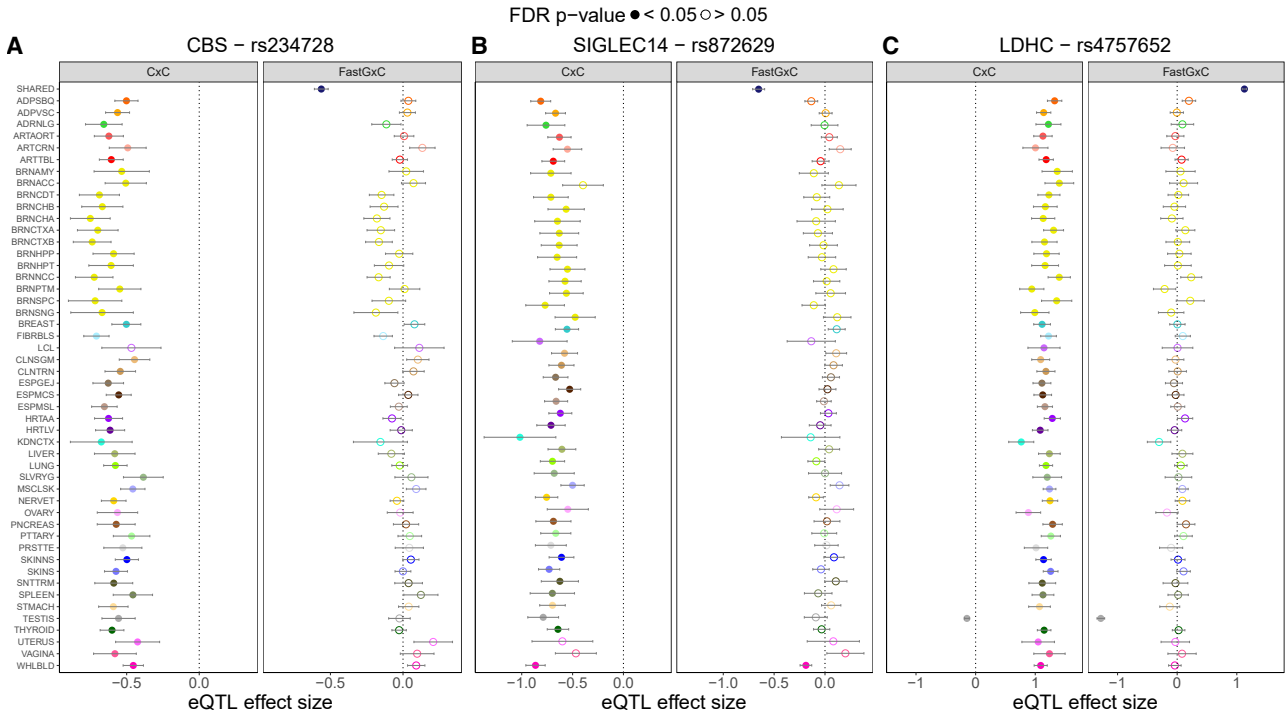

Figure 5. Examples of eQTLs with shared and tissue-specific effects

Each image shows results for a different gene: *CBS* (A), *SIGLEC14* (B), and *LDHC* (C). Each dot shows the effect-size estimates from CxC (left) and FastGxC (right) for a single tissue (color). The top entry in the FastGxC panels ("shared") represents the shared effect across all tissues. Solid dots indicate a significant effect at a 5% FDR. Error bars represent 95% confidence intervals. Tissue abbreviations are explained in Table S3. Additional examples are shown in Figure S20. See also Table S4.

context-specific eQTL effect sizes show little to no correlation outside groups of biologically related tissues and cell types, CxC effect sizes show widespread correlation across all tissues and cell types regardless of biological relationships (Figure 4B, left triangle). This again demonstrates that FastGxC is able to disentangle tissue- and cell-type-specific effects from shared effects.

To illustrate these differences, we discuss a few individual examples (Figure 5). First, we examine *CBS*, a gene that encodes cystathionine beta-synthase, an enzyme that catalyzes the rate-limiting step of the transsulfuration pathway.[44,45] This pathway acts ubiquitously across many cell types to perform diverse biological functions such as protein synthesis and methylation.[46] Indeed, CxC identifies significant eQTLs in 48 individual tissues, suggesting a universal mechanism of genetic regulation. FastGxC clarifies this by identifying a single shared eQTL and no specific eQTLs, correctly capturing the shared regulatory mechanism (Figure 5A). Second, we examine *SIGLEC14*, an immune cell surface receptor involved in the innate immune response.[47] Similar to *CBS*, CxC shows significant effects across many GTEx tissues, which could lead one to conclude that this genetic effect is invariant across the body. However, FastGxC reveals both a shared eQTL and a specific eQTL in whole blood, indicating that while *SIGLEC14* is under universal tissue-shared genetic regulation, there is also a blood-specific regulatory mechanism consistent with its known role in immunity (Figure 5B). Finally, we examine *LDHC*, which encodes lactate dehydrogenase C, a testis-specific enzyme.[48] CxC shows positive eQTL effects in most tissues but an effect in the opposite direction in the testis, suggesting distinct regulatory mechanisms in the tissue where this gene is primarily functional. FastGxC makes this testis-specific regulation immediately apparent by isolating it as a strong specific effect (Figure 5C). Additional examples of widespread sharing obscuring a tissue-specific effect are shown in Figure S20.

Context-specific eQTLs are enriched in functional genomic features from their matched context

To investigate functional differences between shared and context-specific eQTL variants, we performed enrichment analysis of regulatory genomic elements. Promoters typically drive constitutive gene expression across contexts, whereas enhancers are known to confer context specificity.[49] If FastGxC correctly distinguishes shared from context-specific effects, we would expect their variants to localize accordingly. We compared variants with only shared or only context-specific effects to minor-allele frequency (MAF)-matched non-eQTL variants (Figure 6A, right). In bulk tissues, variants with context-specific effects were enriched within enhancers (odds ratio [*OR*] = 1.06, $p = 1.16 \times 10^{-5}$, FDR $\leq$ 5%), while those with shared effects were depleted (*OR* = 0.98, $p = 2.87 \times 10^{-2}$), consistent with previous observations.[18,50] Both sets of variants were

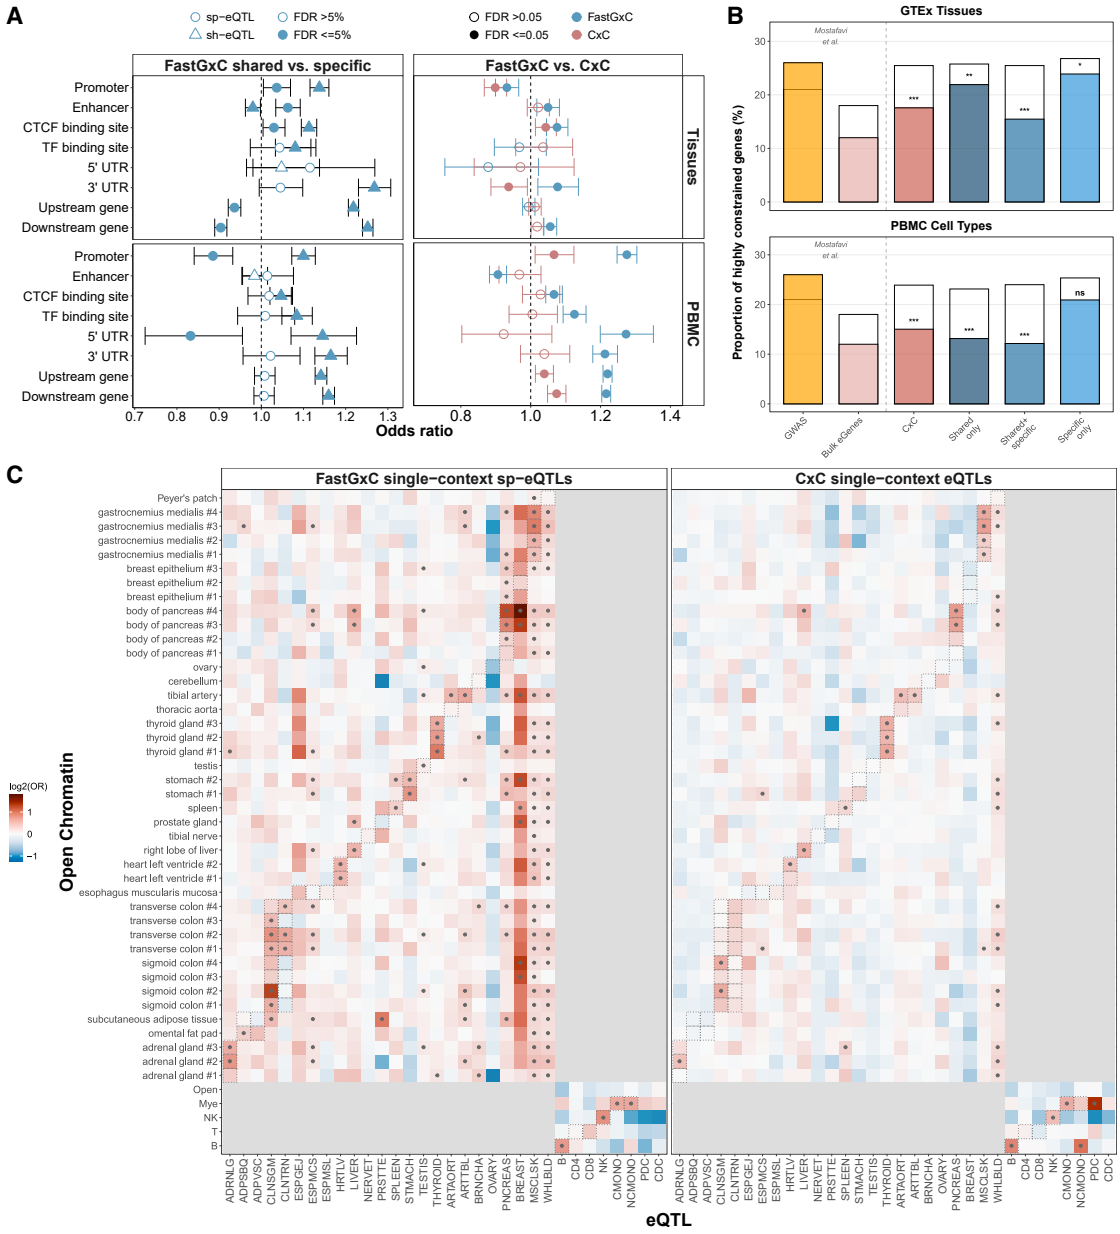

Figure 6. Context-specific eQTL variants are enriched in functional genomic features from their respective contexts

(A) Enrichment of variants with FastGxC shared or context-specific effects only (left) and variants discovered exclusively by FastGxC or CxC (right) across tissues (top) and PBMC types (bottom) in genomic elements with known regulatory effects. Shape indicates different sets of variants. Color indicates different methods. Shape fill indicates the significance of enrichment at FDR $\leq$ 5%. Error bars represent 95% confidence intervals.

(B) Proportion of highly constrained genes (pLI > 0.9, where pLI is the probability of loss-of-function intolerance) among each FastGxC category and CxC (colored bars) vs. matched background genes (white bars; matched on gene length, number of *cis*-SNPs, and expression level) for GTEx (top) and PBMC (bottom). Reference values from Mostafavi et al.[18] for GWAS genes (yellow) and bulk eGenes (pink) with their respective backgrounds are shown on the left. Significance was assessed by Fisher's exact test with Benjamini-Hochberg (BH) correction: ***$p < 0.001$, **$p < 0.01$, *$p < 0.05$, and ns, not significant.

(C) Enrichment of variants with FastGxC context-specific (left) or CxC (right) eQTL effects unique to a single context in regions of open chromatin across multiple tissues and cell types. Tissue- and PBMC-type open chromatin regions were obtained from ENCODE and Calderon et al.,[21] respectively. Boxes indicate manual matching between chromatin and expression context. Color indicates the strength of enrichment/depletion on the $\log_2$ scale. Dots indicate significant enrichment at FDR $\leq$ 5%.

See also Table S4.

enriched within promoters, but the enrichment was stronger for variants with shared effects (OR = 1.14, p = 1.39 × 10^{-37}) compared to those with specific effects only (OR = 1.04, p = 3.14 × 10^{-2}). In single-cell PBMC data, we observed a similar trend for enhancers (OR_{shared} = 0.98 and $OR_{specific}$ = 1.02), but the enrichment is not significant after multiple-testing adjustment, most likely because the number of variants with only shared or specific effects is much smaller for single-cell than bulk data. In addition, variants with shared effects only were enriched within promoters (OR = 1.10, p = 5.34 × 10^{-13}), while those with specific effects were depleted (OR = 0.88, p = 8.26 × 10^{-6}).

To understand how variants with eQTL effects mapped by the CxC approach differ functionally from those identified by FastGxC, we performed another enrichment analysis of genomic elements using sets of variants discovered only by CxC or FastGxC (Figure 6A, right). Compared to CxC-only variants, FastGxC-only variants are enriched (FDR ≤ 5%) in more genomic features (50% vs. 12.5% of annotations in tissues and 87.5% vs. 37.5% in PBMCs) and show stronger enrichment in key genomic elements, such as CTCF binding sites ($OR_{FastGxC}$ = 1.08 and OR_{CxC} = 1.04 in tissues and $OR_{FastGxC}$ = 1.07 and OR_{CxC} = 1.03 in PBMCs). Additionally, in tissues, FastGxC-only variants are significantly enriched in enhancers (OR = 1.05, p = 2.1 × 10^{-3}), while CxC-only variants are not (OR = 1.02, p = 1.8 × 10^{-1}). In PBMCs, FastGxC-only variants are depleted in enhancers but are enriched in every other genomic feature that we tested.

Recent studies have shown that genes near GWAS variants are enriched for highly constrained genes (i.e., genes with a probability of loss-of-function intolerance [pLI] score ≥ 0.90; 26% vs. 21% background), while bulk tissue eGenes in GTEx mapped using CxC are depleted (12% vs. 18% background).[18] To assess whether FastGxC recovers regulatory effects at constrained, disease-relevant genes that CxC approaches miss, we evaluated selective constraint among CxC eGenes and distinct categories of FastGxC eGenes in bulk tissues and PBMC types (Figure 6B). As expected, CxC eGenes show similar significant depletion to bulk tissue eGenes reported in Mostafavi et al.[18] (17.6% vs. 25.4% background in GTEx, 15.0% vs. 23.9% background in PBMCs). In contrast, FastGxC-specific-only eGenes show weaker depletion of constrained genes in GTEx (23.9% vs. 26.8% background) and PBMCs (20.9% vs. 25.3% background), approaching GWAS constraint levels. Shared-only and shared + specific eGenes show stronger depletion than specific-only eGenes in both bulk tissues and PBMCs. Depletion is stronger in PBMCs than in GTEx, most likely reflecting the lower statistical power of single-cell vs. bulk RNA-seq for detecting weak eQTL effects typical of constrained genes.[18,51] These results suggest that context-specific eQTLs capture regulatory effects at constrained genes that standard eQTL mapping approaches miss.

Chromatin is strongly context specific[52] and therefore provides a natural framework for validating FastGxC-mapped context-specific eQTLs and quantifying the functional differences between FastGxC- and CxC-mapped eQTLs. To this end, we tested for enrichment of variants with FastGxC or CxC single-context-specific eQTL effects in regions of open chromatin from matching tissues and cell types. Among bulk tissues, FastGxC variants were more often enriched in open chromatin from the corresponding tissues compared to CxC variants, with enrichment observed in 54% (29/54) vs. 30% (16/54) of cases (one-sided McNemar test, p = 1.95 × 10^{-3}; Figure 6C). Additionally, we observed widespread enrichment in open chromatin for both FastGxC and CxC variants in tissues with broadly distributed cell types, such as whole blood.[53,54] In PBMCs, four out of six (66.6%) cell types with matching chromatin data demonstrated significant enrichment in open chromatin regions from corresponding cell types vs. 50% of cell types for CxC eQTLs (one-sided McNemar test, p = 1). CD4 and CD8 cells lacked significant enrichment (FDR ≤ 5%) within the broader T cell group, most likely due to the aggregation of more granular subtypes, thereby reducing specificity (see STAR Methods), but the enrichment trend is similar (OR_{CD4} = 1.03 and OR_{CD8} = 1.34).

Together, these results highlight the functional relevance of FastGxC context-specific eQTLs, showing greater enrichment in functional genomic elements and improved capture of context-specific chromatin accessibility in matched contexts compared to CxC eQTLs. Additionally, context-specific eQTLs identified exclusively by FastGxC are more likely to reside in functional regions.

Context-specific eQTLs identify putatively causal contexts and genes of complex traits

Mapping eQTLs is crucial for identifying the regulatory targets and context of action of disease-associated non-coding variation. To evaluate whether FastGxC eQTLs improve our understanding of the context that mediates complex disease risk, we analyzed trait-associated variants from 539 traits in the NHGRI-EBI GWAS catalog.[55] Specifically, we tested for enrichment of variants with specific and shared eQTL effects identified by FastGxC in trait-associated variant sets, comparing them to an equal-sized set of MAF-matched non-eQTL variants (Table S5). Following the GTEx consortium protocol, we used expert curation to assign the most relevant tissues for each trait (Table S5)[12] and assessed precision and recall rates to identify the tissue labeled as relevant for each trait. These results were compared with those obtained from CxC eQTLs in individual contexts. PBMC types were excluded from this analysis due to uncertainty regarding the exact cell type relevant for each trait.

At the same recall rate, FastGxC eQTLs achieved a 3-fold increase in precision for identifying disease-relevant tissues and a 2-fold improvement in their ranking compared to CxC eQTLs (Figure 7A). While CxC eQTLs typically prioritized a median of 10 out of 49 tissues per trait, most likely due to widespread tissue sharing (Figure 4B), FastGxC prioritized a median of two tissues. This suggests that modeling the extensive sharing of eQTL effects across tissues can better localize GWAS associations to a smaller, more relevant subset of tissues.

Overall, FastGxC enrichment patterns aligned well with known trait-tissue associations (Figure 7B, FDR ≤ 5%). In cancer traits, where the relevant tissue is typically well-defined, FastGxC demonstrated superior tissue localization compared to CxC. For instance, in breast carcinoma, FastGxC showed the strongest enrichment in breast mammary tissue (OR = 5.0, p = 3.2 × 10^{-4}), while CxC prioritized Epstein-Barr virus

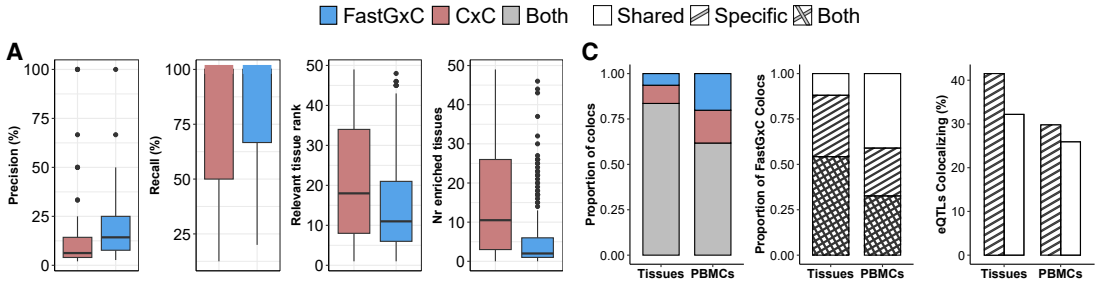

B

GWAS Trait	Method	Relevant Tissue(s) Enriched (Rank)	Nr Other Tissues Enriched	Other Enriched Tissue(s)
breast carcinoma	FastGxC	BREAST (1)	16	BRNHPT, BRNCDT, BRNCTXA, BRNPTM, BRNCTXB, BRNCHA, BRNNCC, FIBRBLS, WHLBLD, BRNCHB, ADPSBQ, NERVET, THYROID, MSCLSK, SKINS, SHARED
	CxC	BREAST (25)	38	LCL, SPLEEN, BRNNCC, CLNTRN, LIVER, BRNSPC, HRTLV, CLNSGM, PRSTTE, OVARY, ADRNLG, VAGINA, UTERUS, BRNCTXB, ADPSBQ, BRNACC, ESPMCS, WHLBLD, LUNG, SLVRYG, PNCREAS, BRNHPT, BRNCDT, ESPMSL, ARTAORT, MSCLSK, SKINS, TESTIS, ADPVSC, STMACH, THYROID, PTTARY, NERVET, ESPGEJ, HRTAA, FIBRBLS, ARTTBL, SKINNS
lung adenocarcinoma	FastGxC	LUNG (2)	1	NERVET
	CxC	LUNG (18)	21	UTERUS, ARTAORT, OVARY, ADPVSC, HRTLV, SNTTRM, ESPMCS, SKINS, PNCREAS, HRTAA, ARTTBL, BRNCHA, THYROID, FIBRBLS, BREAST, SKINNS, NERVET, ADPSBQ, WHLBLD, ESPGEJ, ESPMSL
melanoma	FastGxC	SKINNS (1), SKINS (2)	2	WHLBLD, SHARED
	CxC	SKINS (1), SKINNS (5)	6	BREAST, NERVET, WHLBLD, ADPSBQ, ESPMCS, FIBRBLS
prostate carcinoma	FastGxC	PRSTTE (1)	1	SHARED
	CxC	No enrichments	0	No enrichments
testicular germ cell tumor	FastGxC	TESTIS (1)	1	SHARED
	CxC	TESTIS (10)	11	BRNCHA, PNCREAS, ESPMCS, PTTARY, FIBRBLS, ADPVSC, ARTAORT, CLNTRN, ESPMSL, MSCLSK, SKINS
cancer	FastGxC	SHARED (1)	1	WHLBLD
	CxC	No enrichments	4	WHLBLD, ADPSBQ, BREAST, ESPMCS
coronary artery disease	FastGxC	ARTCRN (1), ARTAORT (2), HRTLV (5), HRTAA (8)	13	SKINNS, ARTTBL, SHARED, MSCLSK, SPLEEN, LIVER, BRNCHA, THYROID, ESPMCS, WHLBLD, NERVET, TESTIS, FIBRBLS
	CxC	ARTAORT (6), ARTCRN (11), HRTAA (23), HRTLV (32)	39	STMACH, CLNTRN, ARTTBL, ADPSBQ, ESPMSL, BRNACC, NERVET, SKINNS, LUNG, FIBRBLS, BREAST, OVARY, UTERUS, LIVER, BRNHPP, PNCREAS, SNTTRM, VAGINA, ESPMCS, SKINS, MSCLSK, SLVRYG, LCL, BRNCTXA, BRNCHB, CLNSGM, PRSTTE, WHLBLD, BRNNCC, ADPVSC, THYROID, BRNCHA, ADRNLG, BRNHPT, PTTARY, BRNPTM, TESTIS, ESPGEJ, SPLEEN
atrial fibrillation	FastGxC	HRTAA (1), HRTLV (2)	4	PTTARY, FIBRBLS, WHLBLD, SHARED
	CxC	HRTLV (1), HRTAA (2)	11	BRNHPP, PRSTTE, ESPMSL, ESPGEJ, THYROID, PNCREAS, BREAST, TESTIS, MSCLSK, ADPVSC, ARTTBL

Figure 7. FastGxC identifies context-relevant mechanisms and increases colocalization with complex traits

(A) Ability of FastGxC and CxC eQTLs to prioritize the most relevant tissue(s) across 539 complex traits with a strong prior indication for the likely relevant tissue(s). The number of enriched tissues for each method was computed only for traits with at least one significant enrichment in either method.

(B) Tissues prioritized by FastGxC and CxC eQTLs and the ranking of known disease-relevant tissues for specific complex traits.

(C) Colocalization of FastGxC and CxC eQTLs with GWAS summary statistics across 63 human traits. Left: proportion of colocalizations found uniquely by FastGxC or CxC and by both methods. Middle: proportion of FastGxC-identified colocalizations that are shared eQTLs only, specific eQTLs only, or both. Right: percentage of FastGxC shared and specific eQTLs that colocalized over the total number of shared and specific eQTLs tested for colocalization for tissues and PBMC types.

See also Figure S21 and Tables S5 and S6.

(EBV)-transformed lymphocytes, with breast mammary tissue ranking 25th ($OR = 2.24$, $p = 7.5 \times 10^{-4}$). In lung adenocarcinoma, CxC identified significant associations in 22 tissues, many unrelated to lung physiology (lung OR = 2.83, ranked 18th, p = 1.6×10^{-3}), whereas FastGxC found associations only in lung (OR = 5.67, $p = 2.60 \times 10^{-3}$) and the tibial nerve (OR = 20, p = 2.1×10^{-5}). For traits not specific to a single tissue, such as the "any cancers" trait, FastGxC showed the strongest enrichment for shared eQTLs, consistent with processes common across tissues. This improved tissue resolution was also evident in non-cancer traits. For example, in coronary artery disease, FastGxC identified significant associations in 17 tissues, compared to 43 for CxC, with the top tissues being cardiovascular-relevant, such as coronary (OR = 13.0, ranked 1st) and aortic (OR = 2.96, ranked 2nd) arteries, left ventricle of the heart (OR = 2.82, ranked 5th), and atrial appendage (OR = 2.37, ranked 8th).

To evaluate the ability of FastGxC eQTLs to identify the regulatory targets of trait-associated variants, we performed a colocalization analysis integrating GWAS summary statistics for 63 complex traits and diseases with FastGxC shared and specific eQTLs in bulk tissues and single-cell PBMC types (Figure 7C; Table S6). We compared these results to colocalizations based on CxC eQTLs mapped separately in each tissue and PBMC type. Across all traits and methods, we prioritized candidate causal genes for 5,726 (47.12% of tested) GWAS loci at a colocalization posterior probability (CLPP) threshold of 50%. The majority of the colocalizations (83.56% in tissues and 61.75% in PBMC types) were identified by both methods, while 6.40% and 20.18% were unique to FastGxC in tissues and PBMC types, respectively. This represents a 6.84% and 25.28% increase in significant colocalizations for tissues and PBMCs, respectively (Figure 7C), with the percentage increase remaining relatively consistent across CLPP thresholds for tissues and reaching up to 50% for PBMCs (Figure S21).

Previous studies suggest that context-specific eQTLs are more enriched for disease associations than shared eQTLs.[12,14]

To test this hypothesis, we compared the colocalization rates of FastGxC shared and specific eQTLs. In tissues, most colocalizations (54.25%) involved eQTLs with both shared and context-specific effects, while 33.80% were specific-only eQTLs. In PBMCs, colocalizations were highest for shared-only eQTLs (40.93%), followed by those with both shared and specific effects (32.62%, Figure 7C). However, after normalizing by the number of shared and specific eQTLs tested for colocalization, specific eQTLs showed higher colocalization rates than shared eQTLs in both tissues (41.52% vs. 32.19%, one-sided binomial proportion test $p = 1.25 \times 10^{-91}$) and PBMCs (29.78% vs. 25.93%, $p = 0.108$) (Figure 7C). This represents a 1.2-fold enrichment in the ability of specific eQTLs to identify candidate causal genes for trait-associated variants, reinforcing their disease relevance.

Taken together, we demonstrate that FastGxC-specific eQTLs enhance the resolution of context-trait associations, increase the number of candidate causal genes for human traits, and are more disease relevant than shared eQTLs.

DISCUSSION

We developed FastGxC, a statistical method for efficiently and powerfully mapping context-specific eQTLs by leveraging the correlation structure of functional genomic studies with repeated sampling. Through simulations, we demonstrated that FastGxC is well calibrated for both identifying eQTLs and assessing their context specificity. Furthermore, FastGxC provides unbiased estimates of overall eQTL effect sizes in each context, with only slight bias in cases of extensive missing data (over 63% of data missing). FastGxC matches the power of LMM-GxC—the only other properly calibrated method for context-specific eQTL mapping—while being orders of magnitude faster.

We applied FastGxC to bulk multi-tissue and single-cell RNA-seq datasets and identified 17,447 tissue-specific and 2,920 cell-type-specific eGenes. The majority of context-specific effects appeared in loci that exhibited context-shared effects, highlighting the importance of defining context specificity by effect-size heterogeneity rather than the presence or absence of significant eQTL effects in each context. In addition, we found that context-specific eQTLs are shared mostly between groups of biologically related contexts and are more frequently enriched in genomic elements that confer context specificity to gene expression, e.g., enhancers and context-specific regions of open chromatin, providing further evidence of their validity. Furthermore, our analysis of *LDHC* revealed testis-specific genetic regulation in addition to the known testis-specific expression of this gene, suggesting that tissue specificity can be established at multiple biological levels—a pattern that may be more widespread than currently appreciated but difficult to detect without methods such as FastGxC that explicitly model context specificity. Finally, we found that context-specific eQTLs provide increased precision for identifying disease-relevant contexts compared to CxC eQTLs, and FastGxC-specific eQTLs provided a 1.2-fold increase over shared eQTLs in identifying putative causal genes that drive human traits, confirming their utility in understanding the regulatory mechanisms underlying complex human diseases.

Additional extensions of FastGxC have the potential to further improve the power and scalability of the method, but we leave these directions to future work. The current implementation models a single shared component across all contexts, which performs well in many datasets. However, this formulation does not identify which specific contexts contribute to the shared signal and may fail to capture finer subgroup structures, such as sets of closely related tissues (e.g., brain regions in GTEx). We have previously shown that incorporating hierarchical decompositions can refine estimates of context-group-specific and context-specific eQTL effects.[56] Moreover, FastGxC defines specificity as deviation from the average effect across contexts, but some studies, such as time-course or environmental perturbation experiments, may require comparisons against a baseline context. Adjusting the decomposition step to accommodate these cases is straightforward and could expand FastGxC's applicability.

Furthermore, FastGxC assumes normally distributed expression residuals after rank-based inverse normal transformation. Extending FastGxC to handle non-normal phenotypes using generalized linear models[31,33,57] is straightforward but could be computationally costly. Similarly, FastGxC can be extended to capture non-linear genetic effects[30] but at a considerable computational cost and likely limited yield at current single-cell sample sizes. Further improvements could come from integrating methods that model shared effect patterns across contexts[34] and incorporating fine-mapping approaches, such as that of Wang et al.,[58] to refine candidate causal variants within significant loci, both of which are compatible with FastGxC.

In conclusion, we show that accounting for intra-individual correlation and extensive sharing of eQTLs across contexts reveals context-specific eQTLs that can aid downstream interpretation of disease-associated variants. Furthermore, we highlight the importance of defining context specificity based on effect-size heterogeneity rather than relying on heuristic definitions and miscalibrated tests. We anticipate that applying FastGxC to the growing number of multi-context bulk and single-cell RNA-seq studies will significantly expand our understanding of the context-specific gene regulatory mechanisms underlying complex human diseases.

Limitations

While FastGxC presents a robust framework for mapping context-specific eQTLs, several limitations merit consideration. For single-cell RNA-seq data, FastGxC operates on pseudo-bulked data, aggregating expression profiles across cells within the same context. While this approach may lead to power loss in cases of substantial cell-to-cell heterogeneity within cell types, prior studies suggest that pseudo-bulk methods perform comparably to single-cell approaches.[31,59] Additionally, FastGxC relies on predefined contexts, which can be challenging in single-cell data due to the lack of a unified framework for defining and classifying cell types.[60] Latent factor-based methods[24,30] bypass the need for predefined annotations but carry their own interpretability trade-offs for downstream disease applications. We are actively working on hierarchical decompositions that would allow FastGxC to identify context-group-specific effects and better capture finer subgroup

structures, such as sets of closely related tissues. Furthermore, while FastGxC's marginal tests are well calibrated, their interpretation requires care in settings with extensive effect-size heterogeneity across many contexts, where identifying the specific context driving specificity can be ambiguous. Finally, as our analyses were conducted on specific cohorts, findings may not fully generalize across all populations or tissue types. Nevertheless, FastGxC performs well in real-world data, as evidenced by the strong enrichment of its context-specific eQTLs in functional genomic annotations and disease associations.

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Brunilda Balliu (bballiu@ucla.edu).

Materials availability

This study did not generate new unique reagents or materials.

Data and code availability

All single-cell RNA-seq data analyzed in this study are publicly available (see below). The map of shared and context-specific eGenes, GWAS enrichment, and colocalization results are available as Tables S4, S5, and S6. These data are publicly available as of the date of publication. The FastGxC R package is available at https://github.com/BalliuLab/FastGxC. All original code, including code to reproduce manuscript figures, has been deposited at GitHub (https://github.com/BalliuLab/FastGxC_Manuscript) and Zenodo (10.5281/zenodo.18857301) and is publicly available as of the date of publication. Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

ACKNOWLEDGMENTS

We thank Anthony Zhao and Charlie Wang for their work on the FastGxC R package. B.B. is supported by NIH grants U01HG012079, R01MH125252, R01DK132775, and R01DK137889. M.T. is supported by EMBO Fellowship ALTF 266-2023. L.K. is supported by NIH T32 781460-VA-32187-05. I.C.-O. is supported by an Ikerbasque Research Fellowship, funded by the EU (H2020-MSCA-COFUND-2020-101034228-WOLFRAM2); the Department of Industry, Energy Transition and Sustainability, Basque Government, ELKARTEK program (KK-2025/00084); and grant PID2023-148986OB-I00, funded by MCIU/AEI/FEDER-EU. N.Z. is supported by NIH grants R01HG006399, R01CA227237, R01CA227-466, R01ES029929, R01MH122688, U01HG 009080, R01HG011345, R35GM133531, R01HL155024, and R01MH125252; DOD grant W81XWH-16-2-0018; and the Chan Zuckerberg Science Initiative. C.J.Y. is supported by NIH grants R01AR071522, R01AI136972, and R01HG011239 and the Chan Zuckerberg Science Initiative and is an investigator at the Chan Zuckerberg Biohub and a member of the Parker Institute for Cancer Immunotherapy. J.W.K. is supported by NIH grants P30 DK116074 (to the Stanford Diabetes Research Center), R01 DK116750, R01 DK120565, R01 DK106236, and R01 DK139112. The GTEx Project was supported by the Common Fund of the Office of the Director of the NIH and by NCI, NHGRI, NHLBI, NIDA, NIMH, and NINDS.

AUTHOR CONTRIBUTIONS

B.B. conceived the project and developed the statistical methods. L.K., A.L., C.L., and B.B. implemented simulations. A.L., L.K., B.B., M.T., A.R., and M.G.G. performed the analyses of GTEx, OneK1K, and CLUES data. A.R. and L.K. performed the colocalization analyses with input from I.C.-O. and J.W.K. L.K. and L.A.T. performed enrichment analyses. B.B., A.L., and L.K. implemented the software and prepared code and data resources. A.L., L.K., and B.B. wrote the manuscript with input from N.Z., C.J.Y., A.D., M.G.G., and M.T. All authors read and approved the manuscript.

DECLARATION OF INTERESTS

C.J.Y. is a scientific advisory board member for and holds equity in Related Sciences and ImmunAI, is a consultant for and holds equity in Maze Therapeutics, and is a consultant for TReX Bio. C.J.Y. has received research support from the Chan Zuckerberg Initiative, Chan Zuckerberg Biohub, and Genentech.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
- METHOD DETAILS
  - Overview of FastGxC method
  - Relationship between FastGxC, CxC, and LM(M)-GxC parameters
  - Simulation study
  - GTEx data quality control
  - OneK1K genotype QC and imputation
  - CLUES genotype QC and imputation
  - OneK1K and CLUES gene expression QC
  - Expression principal component analysis and correlation with covariates
  - FastGxC and CxC eQTL mapping
  - Correlations of eQTL effect sizes across context
  - Building sets of background variants and genes for enrichment analyses
  - eQTL enrichment in genomic features
  - eGene enrichment in high pLI genes
  - Enrichment in regions of open chromatin
  - Enrichment of bulk tissue eQTLs in GWAS loci
  - Colocalization of eQTLs with GWAS variants
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.xgen.2026.101250.

Received: September 12, 2025
Revised: February 9, 2026
Accepted: April 22, 2026
Published: May 18, 2026

REFERENCES

1. Sollis, E., Mosaku, A., Abid, A., Buniello, A., Cerezo, M., Gil, L., Groza, T., Güneş, O., Hall, P., Hayhurst, J., et al. (2023). The NHGRI-EBI GWAS Catalog: knowledgebase and deposition resource. Nucleic Acids Res. *51*, D977–D985. https://doi.org/10.1093/nar/gkac1010.
2. Maurano, M.T., Humbert, R., Rynes, E., Thurman, R.E., Haugen, E., Wang, H., Reynolds, A.P., Sandstrom, R., Qu, H., Brody, J., et al. (2012). Systematic Localization of Common Disease-Associated Variation in Regulatory DNA. Science (New York, N.Y.) *337*, 1190–1195. https://doi.org/10.1126/science.1222794.
3. Nicolae, D.L., Gamazon, E., Zhang, W., Duan, S., Dolan, M.E., and Cox, N.J. (2010). Trait-Associated SNPs Are More Likely to Be eQTLs: Annotation to Enhance Discovery from GWAS. PLoS Genet. *6*, e1000888. https://doi.org/10.1371/journal.pgen.1000888.
4. Gusev, A., Lee, S.H., Trynka, G., Finucane, H., Vilhjálmsson, B.J., Xu, H., Zang, C., Ripke, S., Bulik-Sullivan, B., Stahl, E., et al. (2014). Partitioning Heritability of Regulatory and Cell-Type-Specific Variants across 11 Common Diseases. Am. J. Hum. Genet. *95*, 535–552. https://doi.org/10.1016/j.ajhg.2014.10.004.

5. Gamazon, E.R., Segrè, A.V., van de Bunt, M., Wen, X., Xi, H.S., Hormozdiari, F., Ongen, H., Konkashbaev, A., Derks, E.M., Aguet, F., et al. (2018). Using an atlas of gene regulation across 44 human tissues to inform complex disease- and trait-associated variation. Nat. Genet. *50*, 956–967. https://doi.org/10.1038/s41588-018-0154-4.

6. Civelek, M., Wu, Y., Pan, C., Raulerson, C.K., Ko, A., He, A., Tilford, C., Saleem, N.K., Stančáková, A., Scott, L.J., et al. (2017). Genetic Regulation of Adipose Gene Expression and Cardio-Metabolic Traits. Am. J. Hum. Genet. *100*, 428–443. https://doi.org/10.1016/j.ajhg.2017.01.027.

7. Small, K.S., Hedman, A.K., Grundberg, E., Nica, A.C., Thorleifsson, G., Kong, A., Thorsteindottir, U., Shin, S.-Y., and Richards, H.B. (2011). Identification of an imprinted master trans regulator at the KLF14 locus related to multiple metabolic phenotypes. Nat. Genet. *43*, 561–564. https://doi.org/10.1038/ng.833.

8. Koplev, S., Seldin, M., Sukhavasi, K., Ermel, R., Pang, S., Zeng, L., Bankier, S., Di Narzo, A., Cheng, H., Meda, V., et al. (2022). A mechanistic framework for cardiometabolic and coronary artery diseases. Nat. Cardiovasc. Res. *1*, 85–100. https://doi.org/10.1038/s44161-021-00009-1.

9. De Jager, P.L., Ma, Y., McCabe, C., Xu, J., Vardarajan, B.N., Felsky, D., Klein, H.-U., White, C.C., Peters, M.A., Lodgson, B., et al. (2018). A multi-omic atlas of the human frontal cortex for aging and Alzheimer's disease research. Sci. Data *5*, 180142. https://doi.org/10.1038/sdata.2018.142.

10. Fujita, M., Gao, Z., Zeng, L., McCabe, C., White, C.C., Ng, B., Green, G.S., Rozenblatt-Rosen, O., Phillips, D., Amir-Zilberstein, L., et al. (2024). Cell subtype-specific effects of genetic variation in the Alzheimer's disease brain. Nat. Genet. *56*, 605–614. https://doi.org/10.1038/s41588-024-01685-y.

11. Orchard, P., Blackwell, T.W., Kachuri, L., Castaldi, P.J., Cho, M.H., Christenson, S.A., Durda, P., Gabriel, S., Hersh, C.P., Huntsman, S., et al. (2025). Cross-cohort analysis of expression and splicing quantitative trait loci in TOPMed. Preprint at medRxiv. https://doi.org/10.1101/2025.02.19.25322561.

12. Consortium, T.G. (2020). The GTEx Consortium atlas of genetic regulatory effects across human tissues. Science *369*, 1318–1330. https://doi.org/10.1126/science.aaz1776.

13. Gloudemans, M.J., Balliu, B., Nachun, D., Schnurr, T.M., Durrant, M.G., Ingelsson, E., Wabitsch, M., Quertermous, T., Montgomery, S.B., Knowles, J.W., and Carcamo-Orive, I. (2022). Integration of genetic colocalizations with physiological and pharmacological perturbations identifies cardiometabolic disease genes. Genome Med. *14*, 31. https://doi.org/10.1186/s13073-022-01036-8.

14. Perez, R.K., Gordon, M.G., Subramaniam, M., Kim, M.C., Hartoularos, G.C., Targ, S., Sun, Y., Ogorodnikov, A., Bueno, R., Lu, A., et al. (2022). Single-cell RNA-seq reveals cell type–specific molecular and genetic associations to lupus. Science *376*, eabf1970. https://doi.org/10.1126/science.abf1970.

15. Yazar, S., Alquicira-Hernandez, J., Wing, K., Senabouth, A., Gordon, M.G., Andersen, S., Lu, Q., Rowson, A., Taylor, T.R.P., Clarke, L., et al. (2022). Single-cell eQTL mapping identifies cell type–specific genetic control of autoimmune disease. Science *376*, eabf3041. https://doi.org/10.1126/science.abf3041.

16. Umans, B.D., Battle, A., and Gilad, Y. (2021). Where Are the Disease-Associated eQTLs? Trends Genet. *37*, 109–124. https://doi.org/10.1016/j.tig.2020.08.009.

17. Connally, N.J., Nazeen, S., Lee, D., Shi, H., Stamatoyannopoulos, J., Chun, S., Cotsapas, C., Cassa, C.A., and Sunyaev, S.R. (2022). The missing link between genetic association and regulatory function. eLife *11*, e74970. https://doi.org/10.7554/elife.74970.

18. Mostafavi, H., Spence, J.P., Naqvi, S., and Pritchard, J.K. (2023). Systematic differences in discovery of genetic effects on gene expression and complex traits. Nat. Genet. *55*, 1866–1875. https://doi.org/10.1038/s41588-023-01529-1.

19. Fairfax, B.P., Humburg, P., Makino, S., Naranbhai, V., Wong, D., Lau, E., Jostins, L., Plant, K., Andrews, R., McGee, C., and Knight, J.C. (2014). Innate Immune Activity Conditions the Effect of Regulatory Variants upon Monocyte Gene Expression. Science *343*, 1246949. https://doi.org/10.1126/science.1246949.

20. Ye, C.J., Chen, J., Villani, A.-C., Gate, R.E., Subramaniam, M., Bhangale, T., Lee, M.N., Raj, T., Raychowdhury, R., Li, W., et al. (2018). Genetic analysis of isoform usage in the human anti-viral response reveals influenza-specific regulation of ERAP2 transcripts under balancing selection. Genome Res. *28*, 1812–1825. https://doi.org/10.1101/gr.240390.118.

21. Calderon, D., Nguyen, M.L.T., Mezger, A., Kathiria, A., Müller, F., Nguyen, V., Lescano, N., Wu, B., Trombetta, J., Ribado, J.V., et al. (2019). Landscape of stimulation-responsive chromatin across diverse human immune cells. Nat. Genet. *51*, 1494–1505. https://doi.org/10.1038/s41588-019-0505-9.

22. Gutierrez-Arcelus, M., Baglaenko, Y., Arora, J., Hannes, S., Luo, Y., Amariuta, T., Teslovich, N., Rao, D.A., Ermann, J., Jonsson, A.H., et al. (2020). Allele-specific expression changes dynamically during T cell activation in HLA and other autoimmune loci. Nat. Genet. *52*, 247–253. https://doi.org/10.1038/s41588-020-0579-4.

23. Matoba, N., Le, B.D., Valone, J.M., Wolter, J.M., Mory, J.T., Liang, D., Aygün, N., Broadaway, K.A., Bond, M.L., Mohlke, K.L., et al. (2024). Stimulating Wnt signaling reveals context-dependent genetic effects on gene regulation in primary human neural progenitors. Nat. Neurosci. *27*, 2430–2442. https://doi.org/10.1038/s41593-024-01773-6.

24. Strober, B.J., Tayeb, K., Popp, J., Qi, G., Gordon, M.G., Perez, R., Ye, C.J., and Battle, A. (2024). SURGE: uncovering context-specific genetic-regulation of gene expression from single-cell RNA sequencing using latent-factor models. Genome Biol. *25*, 28. https://doi.org/10.1186/s13059-023-03152-z.

25. Cuomo, A.S.E., Heinen, T., Vagiaki, D., Horta, D., Marioni, J.C., and Stegle, O. (2022). CellRegMap: a statistical framework for mapping context-specific regulatory variants using scRNA-seq. Mol. Syst. Biol. *18*, e10663. https://doi.org/10.15252/msb.202110663.

26. Mu, W., Sarkar, H., Srivastava, A., Choi, K., Patro, R., and Love, M.I. (2022). Airpart: interpretable statistical models for analyzing allelic imbalance in single-cell datasets. Bioinformatics *38*, 2773–2780. https://doi.org/10.1093/bioinformatics/btac212.

27. Heinen, T., Secchia, S., Reddington, J.P., Zhao, B., Furlong, E.E.M., and Stegle, O. (2022). scDALI: modeling allelic heterogeneity in single cells reveals context-specific genetic regulation. Genome Biol. *23*, 8. https://doi.org/10.1186/s13059-021-02593-8.

28. Qi, G., Strober, B.J., Popp, J.M., Keener, R., Ji, H., and Battle, A. (2023). Single-cell allele-specific expression analysis reveals dynamic and cell-type-specific regulatory effects. Nat. Commun. *14*, 6317. https://doi.org/10.1038/s41467-023-42016-9.

29. Funk, M.W., Wang, Y., and Wang, L. (2025). Airqtl dissects cell state-specific causal gene regulatory networks with efficient single-cell eQTL mapping. Nat. Commun. *16*, 11403. https://doi.org/10.1038/s41467-025-66214-9.

30. Kumasaka, N., Rostom, R., Huang, N., Polanski, K., Meyer, K.B., Patel, S., Boyd, R., Gomez, C., Barnett, S.N., Panousis, N.I., et al. (2023). Mapping interindividual dynamics of innate immune response at single-cell resolution. Nat. Genet. *55*, 1066–1075. https://doi.org/10.1038/s41588-023-01421-y.

31. Zhou, W., Cuomo, A.S.E., Xue, A., Kanai, M., Chau, G., Krishna, C., Xavier, R.J., MacArthur, D.G., Powell, J.E., Daly, M.J., and Neale, B.M. (2024). Efficient and accurate mixed model association tool for single-cell eQTL analysis. Preprint at medRxiv. https://doi.org/10.1101/2024.05.15.24307317.

32. Schmiedel, B.J., Singh, D., Madrigal, A., Valdovino-Gonzalez, A.G., White, B.M., Zapardiel-Gonzalo, J., Ha, B., Altay, G., Greenbaum, J.A., McVicker, G., et al. (2018). Impact of Genetic Polymorphisms on Human Immune Cell Gene Expression. Cell *175*, 1701–1715.e16. https://doi.org/10.1016/j.cell.2018.10.022.

33. Zhang, Z.E., Kim, A., Suboc, N., Mancuso, N., and Gazal, S. (2025). Efficient count-based models improve power and robustness for large-scale

single-cell eQTL mapping. Preprint at medRxiv. https://doi.org/10.1101/2025.01.18.25320755.

34. Urbut, S.M., Wang, G., Carbonetto, P., and Stephens, M. (2019). Flexible statistical methods for estimating and testing effects in genomic studies with multiple conditions. Nat. Genet. *51*, 187–195. https://doi.org/10.1038/s41588-018-0268-8.

35. Kukurba, K.R., Parsana, P., Balliu, B., Smith, K.S., Zappala, Z., Knowles, D.A., Favé, M.-J., Davis, J.R., Li, X., Zhu, X., et al. (2016). Impact of the X Chromosome and sex on regulatory variation. Genome Res. *26*, 768–777. https://doi.org/10.1101/gr.197897.115.

36. Balliu, B., Durrant, M., Goede, O.d., Abell, N., Li, X., Liu, B., Gloudemans, M.J., Cook, N.L., Smith, K.S., Knowles, D.A., et al. (2019). Genetic regulation of gene expression and splicing during a 10-year period of human aging. Genome Biol. *20*, 230. https://doi.org/10.1186/s13059-019-1840-y.

37. Oliva, M., Muñoz-Aguirre, M., Kim-Hellmuth, S., Wucher, V., Gewirtz, A.D.H., Cotter, D.J., Parsana, P., Kasela, S., Balliu, B., Viñuela, A., et al. (2020). The impact of sex on gene expression across human tissues. Science *369*, eaba3066. https://doi.org/10.1126/science.aba3066.

38. Shabalin, A.A. (2012). Matrix eQTL: ultra fast eQTL analysis via large matrix operations. Bioinformatics (Oxford, England) *28*, 1353–1358. https://doi.org/10.1093/bioinformatics/bts163.

39. Delaneau, O., Ongen, H., Brown, A.A., Fort, A., Panousis, N.I., and Dermitzakis, E.T. (2017). A complete tool set for molecular QTL discovery and analysis. Nat. Commun. *8*, 15452. https://doi.org/10.1038/ncomms15452.

40. Taylor-Weiner, A., Aguet, F., Haradhvala, N.J., Gosai, S., Anand, S., Kim, J., Ardlie, K., Van Allen, E.M., and Getz, G. (2019). Scaling computational genomics to millions of individuals with GPUs. Genome Biol. *20*, 228. https://doi.org/10.1186/s13059-019-1836-7.

41. Potvin, P.J., and Schutz, R.W. (2000). Statistical power for the two-factor repeated measures ANOVA. Behav. Res. Methods Instrum. Comput. *32*, 347–356. https://doi.org/10.3758/bf03207805.

42. Peterson, C.B., Bogomolov, M., Benjamini, Y., and Sabatti, C. (2016). TreeQTL: hierarchical error control for eQTL findings. Bioinformatics (Oxford, England) *32*, 2556–2558. https://doi.org/10.1093/bioinformatics/btw198.

43. Sul, J.H., Han, B., Ye, C., Choi, T., and Eskin, E. (2013). Effectively Identifying eQTLs from Multiple Tissues by Combining Mixed Model and Meta-analytic Approaches. PLoS Genet. *9*, e1003491. https://doi.org/10.1371/journal.pgen.1003491.

44. Banerjee, R., and Zou, C.-G. (2005). Redox regulation and reaction mechanism of human cystathionine-beta-synthase: a PLP-dependent hemesensor protein. Arch. Biochem. Biophys. *433*, 144–156. https://doi.org/10.1016/j.abb.2004.08.037.

45. Kraus, J.P., Oliveriusová, J., Sokolová, J., Kraus, E., Vlcek, C., de Franchis, R., Maclean, K.N., Bao, L., Bukovsk, n., Patterson, D., et al. (1998). The human cystathionine beta-synthase (CBS) gene: complete sequence, alternative splicing, and polymorphisms. Genomics *52*, 312–324. https://doi.org/10.1006/geno.1998.5437.

46. Miles, E.W., and Kraus, J.P. (2004). Cystathionine beta-synthase: structure, function, regulation, and location of homocystinuria-causing mutations. J. Biol. Chem. *279*, 29871–29874. https://doi.org/10.1074/jbc.R400005200.

47. T, A., T, H., M, Y., A, V., and M, N. (2006). Discovery of Siglec-14, a novel sialic acid receptor undergoing concerted evolution with Siglec-5 in primates. FASEB (Fed. Am. Soc. Exp. Biol.) J. *20*, 1964–1973. https://doi.org/10.1096/fj.06-5800com.

48. Goldberg, E., Eddy, E.M., Duan, C., and Odet, F. (2010). LDHC THE ULTIMATE TESTIS SPECIFIC GENE. J. Androl. *31*, 86–94. https://doi.org/10.2164/jandrol.109.008367.

49. Heintzman, N.D., Hon, G.C., Hawkins, R.D., Kheradpour, P., Stark, A., Harp, L.F., Ye, Z., Lee, L.K., Stuart, R.K., Ching, C.W., et al. (2009). Histone modifications at human enhancers reflect global cell-type-specific gene expression. Nature *459*, 108–112. https://doi.org/10.1038/nature07829.

50. Arthur, T.D., Nguyen, J.P., Henson, B.A., D'Antonio-Chronowska, A., Jaureguy, J., Silva, N., Panopoulos, A.D., D'Antonio, M., and D'Antonio, M. (2025). Multiomic QTL mapping reveals phenotypic complexity of GWAS loci and prioritizes putative causal variants. Cell Genom. *5*, 100775. https://doi.org/10.1016/j.xgen.2025.100775.

51. van der Wijst, M., de Vries, D.H., Groot, H.E., Trynka, G., Hon, C.C., Bonder, M.J., Stegle, O., Nawijn, M.C., Idaghdour, Y., van der Harst, P., et al. (2020). The single-cell eQTLGen consortium. eLife *9*, e52155. https://doi.org/10.7554/eLife.52155.

52. Gross, D.S., and Garrard, W.T. (1988). Nuclease hypersensitive sites in chromatin. Annu. Rev. Biochem. *57*, 159–197. https://doi.org/10.1146/annurev.bi.57.070188.001111.

53. Mandric, I., Rotman, J., Yang, H.T., Strauli, N., Montoya, D.J., Van Der Wey, W., Ronas, J.R., Statz, B., Yao, D., Petrova, V., et al. (2020). Profiling immunoglobulin repertoires across multiple human tissues using RNA sequencing. Nat. Commun. *11*, 3126. https://doi.org/10.1038/s41467-020-16857-7.

54. Mangul, S., Yang, H.T., Strauli, N., Gruhl, F., Porath, H.T., Hsieh, K., Chen, L., Daley, T., Christenson, S., Wesolowska-Andersen, A., et al. (2018). ROP: dumpster diving in RNA-sequencing to find the source of 1 trillion reads across diverse adult human tissues. Genome Biol. *19*, 36. https://doi.org/10.1186/s13059-018-1403-7.

55. Buniello, A., MacArthur, J.A.L., Cerezo, M., Harris, L.W., Hayhurst, J., Malangone, C., McMahon, A., Morales, J., Mountjoy, E., Sollis, E., et al. (2019). The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. Nucleic Acids Res. *47*, D1005–D1012. https://doi.org/10.1093/nar/gky1120.

56. Thompson, M., Gordon, M.G., Lu, A., Tandon, A., Halperin, E., Gusev, A., Ye, C.J., Balliu, B., and Zaitlen, N. (2022). Multi-context genetic modeling of transcriptional regulation resolves novel disease loci. Nat. Commun. *13*, 5704. https://doi.org/10.1038/s41467-022-33212-0.

57. Hu, Y., Xi, X., Yang, Q., and Zhang, X. (2020). SCeQTL: an R package for identifying eQTL from single-cell parallel sequencing data. BMC Bioinf. *21*, 184. https://doi.org/10.1186/s12859-020-3534-6.

58. Wang, G., Sarkar, A., Carbonetto, P., and Stephens, M. (2020). A Simple New Approach to Variable Selection in Regression, with Application to Genetic Fine Mapping. J. R. Stat. Soc. Series B Stat. Methodol. *82*, 1273–1300. https://doi.org/10.1111/rssb.12388.

59. Cuomo, A.S.E., Alvari, G., Azodi, C.B., single-cell eQTLGen consortium; McCarthy, D.J., and Bonder, M.J. (2021). Optimizing expression quantitative trait locus mapping workflows for single-cell studies. Genome Biol. *22*, 188. https://doi.org/10.1186/s13059-021-02407-x.

60. Domcke, S., and Shendure, J. (2023). A reference cell tree will serve science better than a reference cell atlas. Cell *186*, 1103–1114. https://doi.org/10.1016/j.cell.2023.02.016.

61. Karczewski, K.J., Francioli, L.C., Tiao, G., Cummings, B.B., Alföldi, J., Wang, Q., Collins, R.L., Laricchia, K.M., Ganna, A., Birnbaum, D.P., et al. (2020). The mutational constraint spectrum quantified from variation in 141,456 humans. Nature *581*, 434–443. https://doi.org/10.1038/s41586-020-2308-7.

62. Auton, A., Abecasis, G.R., Altshuler, D.M., Durbin, R.M., Abecasis, G.R., Bentley, D.R., Chakravarti, A., Clark, A.G., Donnelly, P., Eichler, E.E., et al. (2015). A global reference for human genetic variation. Nature *526*, 68–74. https://doi.org/10.1038/nature15393.

63. ENCODE Project Consortium (2012). An integrated encyclopedia of DNA elements in the human genome. Nature *489*, 57–74. https://doi.org/10.1038/nature11247.

64. Rubin-Delanchy, P., and Heard, N.A. (2014). A test for dependence between two point processes on the real line. Preprint at arXiv. https://doi.org/10.48550/ARXIV.1408.3845.

65. Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. J. Stat. Software *67*, 1–48. https://doi.org/10.18637/jss.v067.i01.

66. Zhou, H.J., Li, L., Li, Y., Li, W., and Li, J.J. (2022). PCA outperforms popular hidden variable inference methods for molecular QTL mapping. Genome Biol. *23*, 210. https://doi.org/10.1186/s13059-022-02761-4.

67. McCaw, Z.R., Lane, J.M., Saxena, R., Redline, S., and Lin, X. (2020). Operating characteristics of the rank-based inverse normal transformation for quantitative trait analysis in genome-wide association studies. Biometrics *76*, 1262–1272. https://doi.org/10.1111/biom.13214.

68. Ho, D.E., Imai, K., King, G., and Stuart, E.A. (2007). Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference. Polit. anal. *15*, 199–236. https://doi.org/10.1093/pan/mpl013.

69. Benner, C., Spencer, C.C.A., Havulinna, A.S., Salomaa, V., Ripatti, S., and Pirinen, M. (2016). FINEMAP: efficient variable selection using summary data from genome-wide association studies. Bioinformatics *32*, 1493–1501. https://doi.org/10.1093/bioinformatics/btw018.

70. Hormozdiari, F., van de Bunt, M., Segrè, A.V., Li, X., Joo, J.W.J., Bilow, M., Sul, J.H., Sankararaman, S., Pasaniuc, B., and Eskin, E. (2016). Colocalization of GWAS and eQTL Signals Detects Target Genes. Am. J. Hum. Genet. *99*, 1245–1260. https://doi.org/10.1016/j.ajhg.2016.10.003.

71. Chang, C.C., Chow, C.C., Tellier, L.C., Vattikuti, S., Purcell, S.M., and Lee, J.J. (2015). Second-generation PLINK: rising to the challenge of larger and richer datasets. GigaScience *4*, 7. https://doi.org/10.1186/s13742-015-0047-8.

72. Li, H., Handsaker, B., Wysoker, A., Fennell, T., Ruan, J., Homer, N., Marth, G., Abecasis, G., and Durbin, R.; 1000 Genome Project Data Processing Subgroup (2009). The Sequence Alignment/Map format and SAMtools. Bioinformatics *25*, 2078–2079. https://doi.org/10.1093/bioinformatics/btp352.

73. Das, S., Forer, L., Schönherr, S., Sidore, C., Locke, A.E., Kwong, A., Vrieze, S.I., Chew, E.Y., Levy, S., McGue, M., et al. (2016). Next-generation genotype imputation service and methods. Nat. Genet. *48*, 1284–1287. https://doi.org/10.1038/ng.3656.

74. Loh, P.-R., Danecek, P., Palamara, P.F., Fuchsberger, C., A Reshef, Y., K Finucane, H., Schoenherr, S., Forer, L., McCarthy, S., Abecasis, G.R., et al. (2016). Reference-based phasing using the Haplotype Reference Consortium panel. Nat. Genet. *48*, 1443–1448. https://doi.org/10.1038/ng.3679.

75. Cronbach, L.J., and Webb, N. (1975). Between-class and within-class effects in a reported aptitude * treatment interaction: Reanalysis of a study by G. L. Anderson. J. Educ. Psychol. *67*, 717–724. https://doi.org/10.1037/0022-0663.67.6.717.

76. Price, A.L., Helgason, A., Thorleifsson, G., McCarroll, S.A., Kong, A., and Stefansson, K. (2011). Single-Tissue and Cross-Tissue Heritability of Gene Expression Via Identity-by-Descent in Related or Unrelated Individuals. PLoS Genet. *7*, e1001317. https://doi.org/10.1371/journal.pgen.1001317.

77. Yang, S., Liu, Y., Jiang, N., Chen, J., Leach, L., Luo, Z., and Wang, M. (2014). Genome-wide eQTLs and heritability for gene expression traits in unrelated individuals. BMC Genom. *15*, 13. https://doi.org/10.1186/1471-2164-15-13.

78. Simes, R.J. (1986). An improved Bonferroni procedure for multiple tests of significance. Biometrika *73*, 751–754. https://doi.org/10.1093/biomet/73.3.751.

79. Benjamini, Y., and Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. J. R. Stat. Soc., Ser. B Stat. Methodol. *57*, 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x.

80. Sherry, S.T., Ward, M.H., Kholodov, M., Baker, J., Phan, L., Smigielski, E.M., and Sirotkin, K. (2001). dbSNP: the NCBI database of genetic variation. Nucleic Acids Res. *29*, 308–311. https://doi.org/10.1093/nar/29.1.308.

81. Manichaikul, A., Mychaleckyj, J.C., Rich, S.S., Daly, K., Sale, M., and Chen, W.-M. (2010). Robust relationship inference in genome-wide association studies. Bioinformatics *26*, 2867–2873. https://doi.org/10.1093/bioinformatics/btq559.

82. Hoffman, G.E., and Schadt, E.E. (2016). variancePartition: Interpreting drivers of variation in complex gene expression studies. BMC Bioinf. *17*, 483. https://doi.org/10.1186/s12859-016-1323-z.

83. Han, B., and Eskin, E. (2011). Random-Effects Model Aimed at Discovering Associations in Meta-Analysis of Genome-wide Association Studies. Am. J. Hum. Genet. *88*, 586–598. https://doi.org/10.1016/j.ajhg.2011.04.014.

84. Barbeira, A.N., Bonazzola, R., Gamazon, E.R., Liang, Y., Park, Y., Kim-Hellmuth, S., Wang, G., Jiang, Z., Zhou, D., Hormozdiari, F., et al. (2021). Exploiting the GTEx resources to decipher the mechanisms at GWAS loci. Genome Biol. *22*, 49. https://doi.org/10.1186/s13059-020-02252-4.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
GTEx v8 expression + genotype	GTEx Consortium[12]	dbGaP: phs000424.v8.p2
OneK1K expression + genotype	Yazar et al.[15]	GEO: GSE196830
CLUES genotype data	Perez et al.[14]	dbGaP: phs002812.v1.p1
CLUES expression data	Perez et al.[14]	GEO: GSE174188
gnomAD pLI scores (v4.1/v2.1)	Karczewski et al.[61]	https://gnomad.broadinstitute.org/downloads
NHGRI-EBI GWAS Catalog	Sollis et al.[55]	https://www.ebi.ac.uk/gwas/
GWAS summary stats (63 traits)	Various; Table S5	Table S5
1000 Genomes Phase III panel	1000 Genomes[62]	https://www.internationalgenome.org
Tissue ATAC-seq peaks	Roadmap/ENCODE[63]	See Method Details
Calderon PBMC ATAC-seq peaks	Calderon et al.[21]	https://web.stanford.edu/group/pritchardlab/dataArchive/immune_atlas_web/index.html
Software and algorithms		
FastGxC R package	This paper	https://github.com/BalliuLab/FastGxC; https://doi.org/10.5281/zenodo.18857290
Reproduction code	This paper	https://github.com/BalliuLab/FastGxC_Manuscript; https://doi.org/10.5281/zenodo.18857302
MatrixEQTL	Shabalin et al.[38]	CRAN
TreeQTL	Peterson et al.[42]	CRAN
mppa	Delanchy et al.[64]	CRAN
lme4	Bates et al.[65]	CRAN
PCAForQTL	Zhou et al.[66]	CRAN
RNOmni	McCaw et al.[67]	CRAN
MatchIt	Ho et al.[68]	CRAN
FINEMAP	Benner et al.[69]	http://www.christianbenner.com
eCAVIAR	Hormozdiari et al.[70]	https://github.com/fhormoz/caviar
PLINK v1.90/v2.0	Chang et al.[71]	https://www.cog-genomics.org/plink/
bcftools v1.12/v1.18	Li et al.[72]	https://samtools.github.io/bcftools/
Minimac4	Das et al.[73]	https://genome.sph.umich.edu/wiki/Minimac4
Eagle v2.4	Loh et al.[74]	https://alkesgroup.broadinstitute.org/Eagle/
R	R Core Team	https://www.r-project.org

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

This study used publicly available, de-identified human genomic and transcriptomic data from three cohorts. Multi-tissue bulk RNA-seq and whole-genome sequencing data were obtained from the GTEx Consortium v8 (dbGaP: phs000424.v8.p2; $N = 698$ individuals, 49 tissues). Single-cell PBMC RNA-seq and genotype data were obtained from the OneK1K cohort (GEO: GSE196830; $N = 981$ individuals, 8 cell types) and the CLUES cohort (dbGaP: phs002812.v1.p1, GEO: GSE174188; $N = 237$ individuals, 8 cell types). All cohorts included adult participants of both sexes. Detailed sample characteristics, quality control procedures, and ancestry composition are described in the original publications and in the Method Details section. No new human samples were collected for this study, and all data were used in accordance with the terms of their respective data access agreements.

METHOD DETAILS

Overview of FastGxC method

Let E_{ic} be the observed expression of a gene for individual i ($i = 1, \ldots, I$) in context c ($c = 1, \ldots, C$). FastGxC first decomposes E_{ic} into an offset term, a context-shared component, and a context-specific component,[75] i.e.,

$$E_{ic} = E.. + \underbrace{(E_{i.} - E..)}_{E_i^{sh}} + \underbrace{(E_{ic} - Ei.)}_{E_{ic}^{sp}} \quad \text{(Equation 1)}$$

where $E_{..} = \left(\sum_{i=1}^{I}\sum_{c=1}^{C} E_{ic}\right)/(I \times C)$ is the average expression of the gene, computed over all I individuals and all C contexts, and $E_i = \left(\sum_{c=1}^{C} E_{ic}\right)/C$ is the average expression of the gene for individual i, computed over all contexts. In (1), $E..$ is a term that is constant across individuals and contexts for each gene, E_i^{sh} is the context-shared expression component for individual i and is constant across contexts for each gene and individual, and E_{ic}^{sp} is the context-c-specific expression component for individual i.

Next, FastGxC estimates one shared and C context-specific cis genetic effects by regressing the genotypes on each component using ultra-fast implementations of fixed-effect linear regression models[38] and GPU acceleration,[40] i.e.,

$$\begin{aligned} E_i^{sh} &= \alpha^{sh} + \beta^{sh} G_i + \varepsilon_i^{sh}, \\ E_{i1}^{sp} &= \alpha_1^{sp} + \beta_1^{sp} G_i + \varepsilon_{i1}^{sp}, \\ &\vdots \\ E_{iC}^{sp} &= \alpha_C^{sp} + \beta_C^{sp} G_i + \varepsilon_{iC}^{sp}, \end{aligned}$$

where $\alpha^{sh}, \alpha_1^{sp}, \ldots, \alpha_C^{sp}$ are intercepts. $G_i \in \{0, 1, 2\}$ is the genotype of individual i, coded as number of minor alleles, and $\beta^{sh}, \beta_1^{sp}, \ldots, \beta_C^{sp}$ are the genetic effects on the shared and each of the context-specific expression components. Finally, $\varepsilon_{i1}^{sp}, \varepsilon_{i1}^{sp}, \ldots, \varepsilon_{iC}^{sp}$ are each normally distributed residual errors with mean zero and variances $\sigma_{sh}^2, \sigma_{sp,1}^2, \ldots, \sigma_{sp,C}^2$.

Finally, to account for multiple testing across genes, SNPs, and contexts, FastGxC employs the hierarchical FDR-controlling procedure implemented in[42] (Figure S2). We define a gene-SNP pair as an eQTL if the SNP has a significant effect on the shared or any of the specific components of gene expression, i.e., if the global null hypothesis $H_0 : \beta^{sh} = \beta_1^{sp} = \beta_2^{sp} = \cdots = \beta_C^{sp} = 0$ is rejected. If an eQTL exists, we define a *shared-eQTL* as a variant with a statistically significant effect on the shared expression component, i.e., if H_0: $\beta^{sh} = 0$ is rejected, and a *context-specific eQTL* as a variant with a statistically significant genetic effect on at least one context-specific expression components, i.e., if the global null hypothesis $H_0 : \beta_1^{sp} = \beta_2^{sp} = \cdots = \beta_C^{sp} = 0$ is rejected (Figure S2). In addition, we define a *specific-eQTL in context c* as a variant with a statistically significant genetic effect on the context-c-specific expression component, i.e., if the marginal null hypothesis $H_0 : \beta_c^{sp} = 0$ is rejected.

Relationship between FastGxC, CxC, and LM(M)-GxC parameters

FastGxC's eQTL effect estimates can be viewed as a computationally efficient reparametrization of those obtained through CxC and LMM-GxC approaches. Specifically, let β_c represent the eQTL effect in context c, estimated by fitting a linear regression model for each context, i.e., $E_{ic} = \alpha_c + \beta_c G_i + \epsilon_{ic}$. Then, the CxC eQTL effect in context c is equal to the sum of the shared and context-c-specific eQTL effects from FastGxC, i.e., $\beta_c = \beta^{sh} + \beta_c^{cs}$. By construction, FastGxC specific effects $\left(\beta_c^{sp}\right)$ represent deviations from the mean effect across contexts. Positive values indicate stronger-than-average effects; negative values indicate weaker-than average-effects. In addition, let β_1 be the eQTL effect in an arbitrarily defined reference context and δ_c be the interaction eQTL effects for the non-reference context c from an L(M)M model with a genotype-by-context interaction term, i.e., $E_{ic} = (u_i) + \alpha + \beta_1 G_i + \sum_{c=2}^{C} \gamma_c K_{ic} + \sum_{c=2}^{C} \delta_c G_i \times K_{ic} + \varepsilon_{ic}$. Then, $\beta_1 = \beta^{sh} + \beta_1^{cs}$ and $\delta_c = \beta_c - \beta_1 = \beta^{sh} + \beta_c^{cs} - \beta^{sh} - \beta_1^{cs} = \beta_c^{cs} - \beta_1^{cs}$ for $c \neq 1$. Full details of the analytical derivation are provided in the Supplementary Text.

Simulation study

Genotypes were simulated using a binomial distribution with a minor allele frequency of 0.2. Gene expression data were generated under 35 scenarios, varying intra-individual correlation from 0 (independent contexts) to 0.8 and the *cis*-variant effect in each context (Table S2). Under the null hypothesis of no context-specific eQTLs (No heterogeneity), the eQTL effect was either absent across all contexts (No shared eQTL) or identical across contexts (Shared eQTL), with effect sizes explaining 5% of gene expression variability, consistent with prior estimates of *cis*-genetic contributions to gene expression heritability.[76,77]

Under the alternative hypothesis of eQTL effect size heterogeneity, we simulated three scenarios: (i) Single-context heterogeneity, where the eQTL explained 5% of variability in one context and 0% in others (No shared) or 10% in one context and 5% in others (Shared); (ii) Two-context heterogeneity, using similar effect size patterns; and (iii) Extensive heterogeneity, where effect sizes varied across all contexts, ranging from 0% to 10%. For each scenario, we simulated 1,000 datasets. To assess the impact of sample size, number of contexts, and missing data, we varied the number of individuals (100 or 698), contexts (8 or 49), and the proportion of missing expression data (approx. 63% and 7% across individuals and contexts), reflecting patterns observed in the GTEx[12] and OneK1K[15] data.

We obtained global estimates of type I error rates and power to identify an eQTL and test whether the eQTL was context-specific as follows. For the CxC-based approach, we used the MatrixEQTL R package[38] to fit linear regression models for the effect of the eQTL on expression in each context c, i.e., $E_{ic} = \alpha_c + \beta_c G_i + \epsilon_{ic}$, and obtained t test p-values for the null hypothesis of no eQTL effect in context c, H_0: $\beta_c = 0$. Following the hierarchical FDR-controlling procedure implemented in,[42] we then tested the global null

hypothesis of no eQTL effect across contexts, H_0: $\beta_1 = \ldots = \beta_c = 0$, using Simes's method,[78] as implemented in the mppa R package,[64] to combine the *t* test *p*-values. Global Type I error rate and power to identify an eQTL were computed as the proportion of datasets in which the eQTL Simes' *p*-value was significant at the $\alpha = 5\%$ level. For the MetaTissue approach, we followed the procedure implemented in[43] and obtained the RE2 *p*-values which assume no heterogeneity under the null to test the null hypothesis of no eQTL effect. Global Type I error rate and power to identify a single-context-specific eQTL were computed as the proportion of datasets in which the *t* test *p*-value was significant in only one context at *FDR* < 5% for CxC and M-value >0.9 in exactly one context for MetaTissue.

We used a similar strategy for FastGxC. Specifically, we fitted linear regression models for the effect of the eQTL on the shared and each of the *C* specific components of expression *c*, i.e., $E_i^{sh} = \alpha^{sh} + \beta^{sh} G_i + \varepsilon_i^{sh}$ and $E_{ic}^{sp} = \alpha_c^{sp} + \beta_c^{sp} G_i + \varepsilon_{ic}^{sp}$, and obtained *t* test *p*-values for the null hypothesis of no shared or context-*c*-specific eQTL effect, i.e., H_0: $\beta^{sh} = 0$ and H_0: $\beta_c = 0$. We then tested the global null hypothesis of no eQTL effect across contexts, $H_0 : \beta^{sh} = \beta_1^{sp} = \ldots = \beta_C^{sp} = 0$, and the global null hypothesis of no context-specific eQTL effect in any context, $H_0 : \beta_1^{sp} = \ldots = \beta_C^{sp} = 0$, using Simes' method to combine the corresponding *p*-values (Figure S2). We computed the global Type I error rate and power as the proportion of datasets in which the eQTL Simes' *p*-value was significant at the $\alpha = 5\%$ level.

Finally, for the LM-GxC approach, we fitted one linear model with a genotype-by-context interaction term $E_{ic} = \alpha + \beta G_i + \sum_{c=2}^{C} \gamma_c K_{ic} + \sum_{c=2}^{C} \delta_c G_i \times K_{ic} + \varepsilon_{ic}$ and tested the null hypothesis of no eQTL (H_0: $\beta = \delta_2 = \ldots = \delta_C = 0$) as well as the null hypothesis of no context-specific eQTL (H_0: $\delta_2 = \ldots = \delta_C = 0$) using likelihood ratio tests (LRT). For the LMM-GxC approach, we fitted one linear random effects model with a genotype-by-context interaction term $E_{ic} = u_i + \alpha + \beta G_i + \sum_{c=2}^{C} \gamma_c K_{ic} + \sum_{c=2}^{C} \delta_c G_i \times K_{ic} + \varepsilon_{ic}, u_i \sim N(0, \sigma_i^2)$ using the lme4 R package[65] and tested the same null hypotheses as the LM-GxC model.

To assess the ability of FastGxC to identify the heterogeneous context(s), we also obtain marginal estimates of type I error rate and power within each context by testing *C* null hypotheses of no context-specific eQTL in each context $\left(H_0^c : \beta_c^{sp} = 0\right)$ using the *t* test implemented in the MatrixEQTL R package[38] and adjust for multiple testing across contexts using the Benjamini-Hochberg procedure.[79] To compare FastGxC estimates with true simulated effect sizes, we report the mean effect size across 1,000 simulated datasets for each scenario with a 95% confidence interval computed as $\bar{\hat{\beta}} \pm Z * \frac{\sigma}{\sqrt{n}}$ where *Z* is the *Z* score at $\alpha = 0.05$.

GTEx data quality control

Fully processed, filtered, and normalized gene expression matrices (in BED format) as well as meta data including genotype PCs, PEER factors, etc (see GTEx paper[12] for more details) for each tissue across 698 individuals were downloaded through the GTEx portal (https://www.gtexportal.org/home/datasets) on March 11, 2020. Only genes expressed (as defined by the GTEx consortium[12]) in at least two tissues were considered for downstream analyses. Prior to eQTL mapping, gene expression matrices were residualized for major sources of expression variability, including PEER factors, as per the GTEx Consortium[12]).

WGS genotype VCF data were downloaded from dbGap (dbGaP Accession phs000424.v8.p2), and only individuals with both genotype and gene expression data were retained (*N* = 698). The VCF files were processed using vcftools (v0.1.16) to retain only bi-allelic SNPs, with variants filtered to include only those with minor allele frequencies greater than 5% in the tissue of interest. Genotype files were annotated with rs IDs using bcftools (v1.12),[72] and Plink (v1.90)[71] was used to transpose and convert the VCF files into a sample-by-genotype matrix, which served as input for eQTL mapping.

OneK1K genotype QC and imputation

Array genotype data for OneK1K was obtained via the Gene Expression Omnibus (GSE196830) and included genotypes for 1,104 individuals and 759,993 markers on the Illumina Infinium Global Screening Array. For downstream analyses, only individuals with gene expression data were considered (*N* = 981). Bcftools version 1.18 was used to map SNPs to the GRCh37.13 build 151 of dbSNP[80] and adjust allele strand orientation for mismatches. PLINK version 1.90[71] was used to filter SNPs and individuals with a call rate less than 0.95 (zero individuals and 11,317 SNPs), SNPs with a Hardy-Weinberg equilibrium test *p*-value less than 10^{-6} (1,857 SNPs), minor allele frequency (MAF) below 0.01 (211,894 SNPs), and individuals with ambiguous sex labeling (one individual).

To identify ancestry outlier samples, we performed principal component analysis (PCA) jointly on the OneK1K and 1000 Genome Phase I samples. 1000 Genome Phase I data was downloaded from the EMBL-EBI public endpoint (http://ftp.ebi.ac.uk/1000g/ftp/). PCA was conducted on the merged data using PLINK version 2.0, and 3 individuals with non-European ancestry, defined as being within three standard deviations from the European mean of genetic principal components 1 and 2, were excluded. Excess autosomal heterozygosity, defined as being within three standard deviations from the mean, was computed with PLINK version 1.90 and 7 individuals were removed. A genetic relationship matrix from all autosomal SNPs was generated using KING version 2.3.1.[81] No individuals were excluded based on a 0.125 relatedness threshold (second-degree relatives).

After quality control, 499,909 autosomal SNPs from 970 individuals were retained for imputation. Imputation was performed using the Michigan Imputation Server[73] with the 1000G phase III V5 reference panel[62] and was run using Minimac4 and Eagle v2.4.[74] For subsequent *cis*-eQTL analyses, 5,849,361 SNPs with imputation quality R-squared of at least 0.8 and a minor allele frequency above 0.05 were retained.

CLUES genotype QC and imputation

Genotype data for 237 individuals (188 labeled as SLE and 49 labeled as Immvar) from CLUES was obtained from dbGap (accession number phs002812.v1.p1). Genotypes from the SLE and Immvar CLUES cohorts were processed and imputed separately, as in,[14] resulting in 22,159,030 variants for the SLE cohort and 9,797,072 variants for the Immvar cohort. After imputation, variants with an imputation quality R-squared greater than 0.8 were retained, and the datasets were merged using PLINK v2.0, producing a combined total of 16,616,859 variants.

Based on genetically determined ancestry, individuals were then split into European ($N = 140$) and Asian ($N = 97$) cohorts, following the approach described in Perez et al.[14] Within each ancestry group, PLINK v2.0 was used to filter for variants with a minor allele frequency above 5%, resulting in 4,995,061 variants in the Asian cohort and 5,292,554 variants in the European cohort for further analysis. Genotype principal components were computed within each cohort using PLINK v1.90 to identify outlier individuals, defined as those within three standard deviations of the mean for genetic principal components 1 and 2. This filtering excluded 5 individuals from the Asian cohort and 8 from the European cohort.

OneK1K and CLUES gene expression QC

Gene expression data for the CLUES cohort was obtained from GEO (accession number GSE174188). For downstream analysis, we retained cells from the eight cell types used for *cis*-eQTL analysis in[14]: B cells, conventional and plasmacytoid dendritic cells (cDC and pDC), classical and non-classical monocytes (cMono and ncMono), NK cells, CD4 T cells, and CD8 T cells. As with genotype data, the CLUES cohort was divided into European and Asian ancestry groups. Only individuals with both expression and genotypes were retained, i.e., $N = 140$ and $N = 97$ European and Asian ancestry individuals.

Gene expression data for the OneK1K cohort was obtained via GEO (accession number GSE196830). From the 29 pre-defined cell type clusters, we consolidated labels for similar cell types. Specifically, memory B cells, naive B cells, and transitional B cells were grouped as B cells; natural killer cells and CD16-negative, CD56-bright natural killer cells were combined as NK cells; T alpha-beta cytotoxic, CD4-positive cells, T alpha-beta, CD4-positive cells, and T regulatory cells were grouped as CD4 T cells; and T alpha-beta, CD8-positive cells, T gamma-delta cells, and T mucosal invariant cells were grouped as CD8 T cells. For downstream analyses, only cells from the eight cell types present in CLUES were retained. In addition, only individuals with both expression and genotypes were retained, i.e., $N = 981$ individuals.

For each individual in the CLUES European, CLUES Asian, and OneK1K cohorts, a pseudo-bulk expression profile was generated by averaging counts across cells for each cell type and gene. The following steps were then applied separately within each cohort and cell type. First, expression values for each gene were adjusted for library size factors and normalized using counts per million (CPM) and transcripts per million (TPM). Genes with normalized expression greater than zero in at least 10% of samples were retained. Gene expression values were then transformed to approximate normality using the RankNorm function in the RNOmni R package.[67] Within each cohort, genes expressed in fewer than three cell types were excluded. Next, principal component analysis (PCA) was applied to identify and remove outlier samples defined as those falling outside three standard deviations from the mean of expression principal components 1 and 2. After quality control, 17,199, 16,225, and 14,701 genes and 132, 92, and 970 individuals in the CLUES European, CLUES Asian, and OneK1K cohorts, respectively, were retained for downstream analyses.

To capture major sources of expression variation, the PCA implementation in the PCAForQTL R package[66] was used with the runBE function to determine the number of principal components that explain a significant portion of variation in each cell type. Gene expression data for each cell type and cohort was then residualized for six genotype principal components, selected gene expression principal components, sex, age, batch, and, in CLUES cohorts only, SLE status. Together, these normalization and residualization steps ensure that gene expression values are on a comparable scale across contexts and reduce technical artifacts and batch effects, enabling a more direct comparison of eQTL effect sizes.

Expression principal component analysis and correlation with covariates

To evaluate the effectiveness of FastGxC in removing gene expression background noise, we first applied PCA separately to the original gene expression data and the decomposed shared and context-specific expression data, using the prcomp function in the stats R package. Next, we correlated technical and biological covariates with the first ten principal components (PCs) from each data. The correlation between expression PCs and covariates was computed using the canCorPairs function from the variancePartition R package ([82]). In short, when comparing two continuous variables (e.g., gPC1 or weight), Pearson correlation was used. In order to accommodate the correlation between a continuous and a categorical variable (e.g., cohort) canonical correlation analysis (CCA) was used. Note that CCA returns correlations values between 0 and 1.

FastGxC and CxC eQTL mapping

Residualized gene expression (see methods above) for each was mean-centered across all individuals and contexts, then decomposed into 49 tissue-specific components for GTEx and 8 -cell type-specific components for CLUES and OneK1K, along with one shared expression component. Cis genetic effects were estimated on shared gene expression levels (FastGxC), context-specific gene expression levels (FastGxC), and gene expression levels within each context (CxC) using ultra-fast linear regression models

in the MatrixEQTL R package,[38] with model = modelLINEAR and a 1 Mb window for *cis*-eQTL calls. The CLUES European, CLUES Asian, and OneK1K cohorts were then meta-analyzed within each cell type using METASOFT (v2.0)[83] random effect model (RE2) with default parameters.

Multiple testing correction was applied separately for CxC and FastGxC and for bulk and meta-analyzed single-cell results using the hierarchical FDR procedures in the TreeQTL R package[42] with an alpha level of 5% in each level. For the CxC approach, a three-level hierarchy with genes, gene-SNP pairs, and gene-SNP-context triplets was used. For FastGxC, multiple testing correction was applied separately for the shared and specific components using a two-level hierarchy for shared eQTLs (genes and gene-SNP pairs) and a three-level hierarchy for context-specific eQTLs (genes, gene-SNP pairs, and gene-SNP-context triplets) (Figure S2B).

Correlations of eQTL effect sizes across context

Pearson correlations were computed across all tissues in GTEx using effect sizes of only the significant eQTLs in each tissue for the CxC approach or tissue-specific eQTLs for FastGxC. Missing effect size values, due to an eQTL not being tested or not reaching significance in some tissues, were set to zero, and pairwise complete correlations were calculated using the cor function in R. For single PBMC cell types, Pearson correlations of eQTL effect sizes across cell types were computed using all tested, meta-analyzed METASOFT RE2 effect sizes.

Building sets of background variants and genes for enrichment analyses

The matchit function from the MatchIt R package was used to create a set of background SNP-gene pairs for each variant set of interest, matched by minor allele frequency (MAF) using the nearest neighbor matching method with a 1:1 ratio.[68] Only variants used for eQTL mapping were included in building these background sets. For eQTL sets containing more than 5,000 variants, the sets were randomly split into chunks to expedite computation. This analysis was done separately in bulk-tissues and single-cell PBMCs and for FastGxC and CxC eQTLs. The same matchit function and R package were used to construct background gene sets for each eGene category analyzed in the selective constraint enrichment analysis. Genes were matched on gene length, number of *cis*-SNPs, and quantified expression levels. Background sets were restricted to genes included in the eQTL mapping but excluded any genes identified as significant eGenes by FastGxC or CxC within their respective sets. This analysis was conducted separately for bulk tissues and PBMCs for both FastGxC and CxC eQTLs.

eQTL enrichment in genomic features

To test for enrichment of various eQTL types within genomic annotations, we first created variant sets specific to each type of eQTL. This analysis was conducted separately for bulk tissues and single-cell PBMC cell types. Specific-eQTL-only and shared-eQTL-only variant sets were derived by taking the set difference of specific-eQTL and shared-eQTL variants. The FastGxC eQTL variant set was constructed by combining shared- and specific-eQTL variants across contexts, and the CxC eQTL variant set was created by taking the union of eQTL variants across contexts. We then calculated the set difference between the FastGxC and CxC eQTL variant sets to obtain the final FastGxC-only and CxC-only eQTL variant sets. Variants from each set were annotated using the Ensembl Variant Effect Predictor (VEP) tool, which identifies variant effects, such as potential impacts on protein sequences or positioning within genomic regulatory elements. Enrichment of each eQTL set within VEP-annotated categories was tested using a one-sided Fisher's exact test from the stats R package, followed by BH correction; significance was defined as a BH-adjusted p-value less than 0.05.

eGene enrichment in high pLI genes

We annotated the constraint of all tested genes using pLI scores downloaded from gnomad.v4.1 and gnomad.v2.1 by gene data https://gnomad.broadinstitute.org/downloads#v4-constraint.[61] To test for enrichment of eGene groups within high pLI genes, we first created unique eGene sets corresponding to different FastGxC categories and CxC. We conducted this analysis separately for bulk tissues and single-cell PBMCs. FastGxC-specific-only and FastGxC-shared-only eGene sets were derived by taking the set difference of specific-eGenes and shared-eGenes. The FastGxC shared + specific eGene set was constructed by taking a set difference between all FastGxC eGenes and FastGxC-specific-only and FastGxC-shared-only eGenes, and the CxC eGene set was created by taking the set difference of all CxC eGenes and all FastGxC eGenes. Genes from each set were annotated with pLI scores, and we defined high pLI genes as genes with a pLI score >0.9. Enrichment of each eGene set within high pLI genes was tested using a one-sided Fisher's exact test from the stats R package, followed by BH correction; significance was defined as a BH-adjusted p-value less than 0.05.

Enrichment in regions of open chromatin

For bulk tissues, all available tissue ATAC-seq data in the 'not perturbed,' GRCh38, and bigBed narrowPeak categories were downloaded from www.encodeproject.org in November 2020.[63] The downloaded bigBed files were converted to bed format using the UCSC bigbedtobed tool for downstream analysis. For single-cell PBMC cell types, cell-type-specific ATAC-seq peaks were downloaded from https://web.stanford.edu/group/pritchardlab/dataArchive/immune_atlas_web/index.html[21] and grouped into major cell types (B, T, NK, Myeloid, Open), following the approach of Perez et al.[14] Bed files were then sorted using the command bedtools sort -k1,1 -k2,2n to enable memory-efficient processing for subsequent intersections.

Enrichment analysis of FastGxC and CxC single-context eQTL variants from both bulk tissues and single-cell PBMC cell types in regions of open chromatin was performed by intersecting each eQTL variant set of interest with each pre-sorted bed file, representing ATAC-seq peaks from a tissue, cell type, or sample (when multiple samples were available for each context), using the bedtools intersectBed command. A one-sided Fisher's exact test was used to obtain the statistical significance of each enrichment, followed by BH multiple testing adjustment. Significance was called for BH-adjusted *p*-values less than 0.05.

Enrichment of bulk tissue eQTLs in GWAS loci

Genome-wide association study (GWAS) data (gwas_catalog_v1.0.2-associations_e100_r2020-06-17) including 1,563 unique traits was downloaded from the NHGRI-EBI GWAS Catalog in August 2020[55] and processed for downstream analysis. Matching of variants with and without eQTL effects was performed as previously described.

Enrichment analysis of FastGxC and CxC eQTL variants was performed by intersecting each eQTL variant set of interest with trait-associated variants from the GWAS catalog based on rs IDs. Statistical significance was assessed using a one-sided Fisher's exact test for each enrichment. Only mapped traits within the GWAS catalog that contained at least 10 significant loci were included in our downstream analysis resulting in 539 traits with complete enrichment results. Multiple testing correction was applied using the hierarchical FDR procedures in the TreeQTL R package,[42] with tissues at level one and tissue-trait pairs at level two, maintaining an FDR of 5% at each level.

The most likely causal tissue(s) for the 539 GWAS traits were annotated manually, following the approach used in.[12] These annotations were used to calculate precision and recall rates. Specifically, for each trait, a contingency table was constructed to capture the frequency with which a trait of interest is both enriched in a tissue's eQTLs and assigned as the likely relevant tissue. This yielded true positive, false positive, true negative, and false negative counts (TP, FP, TN, FN). Precision was calculated as TP/(TP + FP), and recall as TP/(TP + FN).

Colocalization of eQTLs with GWAS variants

We obtained complete GWAS summary statistics for 63 unique complex traits. The list of traits analyzed, along with their references, is available in Table S5. Colocalization analysis of GWAS variants with eQTLs from bulk tissues and single-cell PBMC cell types was conducted using a custom integration of FINEMAP[69] and eCAVIAR[70] with an LD-modified colocalization posterior probability ($CLPP_{mod}$) following the method outlined in Gloudemans et al.[13] Significant colocalization was defined as $CLPP_{mod}$ above 0.5, as in,[84] and impact of this threshold on results is studied in Figure S21. To assess the colocalization contribution from FastGxC specific versus shared eQTLs in PBMCs, we limited our analysis to the 10 immune related traits out of the original 63.

QUANTIFICATION AND STATISTICAL ANALYSIS

All statistical analyses were performed in R. eQTL mapping was conducted using linear regression via the MatrixEQTL R package,[38] with significance determined by hierarchical FDR control (hFDR $\leq$5%) using the TreeQTL R package.[42] Global eQTL significance across contexts was assessed using Simes' method as implemented in the mppa R package.[64] For the LMM-GxC comparison, linear mixed models were fit using the lme4 R package[65] and significance was assessed via likelihood ratio tests. Enrichment of eQTL variants in GWAS loci was tested using Fisher's exact test with Benjamini-Hochberg correction (Method Details: Enrichment of bulk tissue eQTLs in GWAS loci). Enrichment of eQTL variants in open chromatin regions was compared between FastGxC and CxC using a one-sided McNemar test (Method Details: Enrichment in regions of open chromatin). Enrichment of eGene categories in highly constrained genes (pLI >0.9) was tested using one-sided Fisher's exact tests with Benjamini-Hochberg correction; significance was defined as BH-adjusted $p < 0.05$ (Method Details: eGene enrichment in high pLI genes). Colocalization of eQTLs with GWAS variants was performed using a custom integration of FINEMAP[69] and eCAVIAR[70] with an LD-modified colocalization posterior probability ($CLPP_{mod}$); significant colocalization was defined as $CLPP_{mod} > 0.5$ (Method Details: Colocalization of eQTLs with GWAS variants). Colocalization rates between specific and shared eQTLs were compared using one-sided binomial proportion tests. In simulation studies, 1,000 replicates were generated per scenario; type I error was assessed at the nominal 5% level. Genotypes were simulated from a binomial distribution. Sample sizes, number of contexts, and missingness rates for each simulation scenario are detailed in Table S2. Statistical test details, exact *n* values, and definitions of significance for each analysis are reported in the corresponding Method Details subsections and figure legends.