1. Supplemental figures.

Figure S1

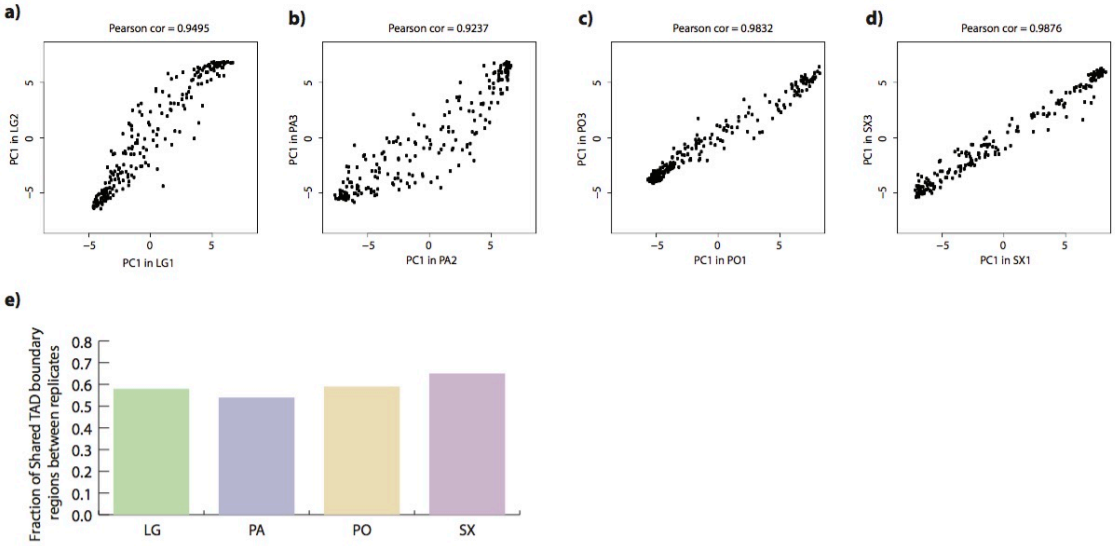

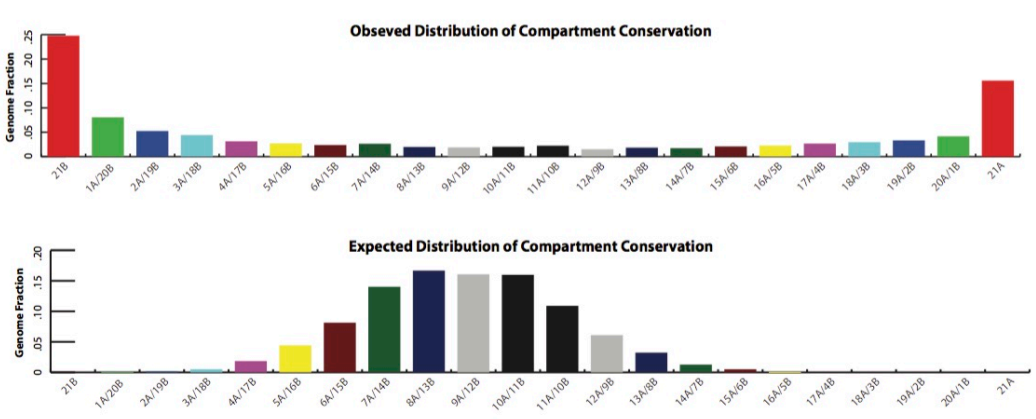

Compartment	Observed	Expected
21B	682	0
1A+20B	222	1
2A+19B	145	4
3A+18B	122	14
4A+17B	87	51
5A+16B	75	122
6A+15B	66	224
7A+14B	73	386
8A+13B	55	459
9A+12B	52	442
10A+11B	56	440
11A+10B	62	300
12A+9B	42	168
13A+8B	51	89
14A+7B	48	35
15A+6B	58	14
16A+5B	63	2
17A+4B	74	0
18A+3B	82	0
19A+2B	92	0
20A+B	115	0
21A	429	0

Chi-square test p-value < 2.2e-16
Binomial test p-value: P(# of 7B>=682) < 2.2e-16
Binomial test p-value: P(# of 7A>=429) < 2.2e-16

g)

Sample	# of TAD boundaries
0.05	2,534
0.1	2,511
0.2	2,517
0.4	2,509
0.6	2,519
0.8	2,508
All	2,516

h)

Sample	# of overlap with TAD boundary regions	% of TAD boundary regions
0.05	2,497	92.34%
0.1	2,479	91.68%
0.2	2,478	91.64%
0.4	2,479	91.68%
0.6	2,484	91.86%
0.8	2,474	91.49%
All	2,479	91.68%

i)

Sample-specificity	# of TAD boundary regions	% of TAD boundary regions
Only in one sample	128	4.73%
Shared by 2 samples	58	2.14%
Shared by 3 samples	41	1.52%
Shared by 4 samples	48	1.78%
Shared by 5 samples	51	1.89%
Shared by 6 samples	90	3.33%
Shared by all 7 samples	2,288	84.62%

Figure S1. Hi-C data reproducibility and compartment A/B conservation, related to Figure 1.

A) Scatter plots from replicates of LG, showing the genome-wide of the PC1 values used for the Compartment A/B analysis. The plot title contains the Pearson correlation coefficient of all 1Mb bin-pairs. The x- and y-axes are labeled according to their tissue type and donor. For example, LG2 corresponds to Lung tissue from donor 2.
B) Same as Panel A, except analysis of biological replicates for PO.
C) Same as Panel A, except analysis of biological replicates for PA.
D) Same as Panel A, except analysis of biological replicates for SX.
E) Bar plots showing the statistically significant fraction of overlapping TAD boundaries in LG, PO, PA, and SX (chi squarest test p value < 2.2e-16).
F) Bar plots showing the observed (top) and expected (bottom) distributions of compartment A/B conservation. Labels on the x-axis indicate the number of cell lines and compartment label for which there is conservation and the Y-axis indicates the total genome fraction that corresponds to that compartment label. For example, 16A/5B indicates the total number of 1Mb bins for which 16 human cell lines or tissues had an A compartment label and 5 samples had a B compartment label. In the bottom table, the 'Compartment' column indicates the how many cell lines are shared for each compartment label, while the 'Expected' and 'Observed' columns indicate how many 1Mb bins fall into 'Compartment' category. Statistical analysis comparing the observed and expected distributions are done with Chi-square test, and statistical analysis of having complete conservation across all samples (i.e. 21A or 21B) was done with a binomial test.
G) Table showing the total number of topological domain boundaries detected using the insulation square method (Crane et al, Nature, 2015) applied to down-sampled Hi-C from H1 cells. The left column indicates what fraction of the full H1 dataset was obtained from down-sampling, and the right column indicates the total number of TAD boundaries detected.
H) Table showing the absolute number of TAD boundary regions overlapping all putative boundaries identified across all down-sampling samples (middle column). The right column indicates the corresponding fraction out of all putative boundaries identified across all down-sampling samples. The left column indicates what fraction of the full H1 dataset was obtained from down-sampling.
I) Table showing the percentages of TAD boundaries that were unique to subsets of the down-sampled H1 datasets. The left column indicates how many of the 7 degrees of sampling share a particular TAD boundary region. The middle column indicates how many TAD boundaries regions were common to a particular subset denoted in the left column. The right column is the corresponding fraction the common TAD boundary regions are of the total putative boundaries in down-sampled H1.

Figure S2

a)

	Data Source				
Bias Factor	*Raw*	*ICE*	*VC*	*HiCNorm*	*HiCNormCis*
Effective Fragment Length	0.47	-0.39	-0.30	0.25	0.01
GC Content	0.15	0.60	0.60	0.08	0.05
Mappability	0.18	0.02	-0.03	0.10	-0.01

b)

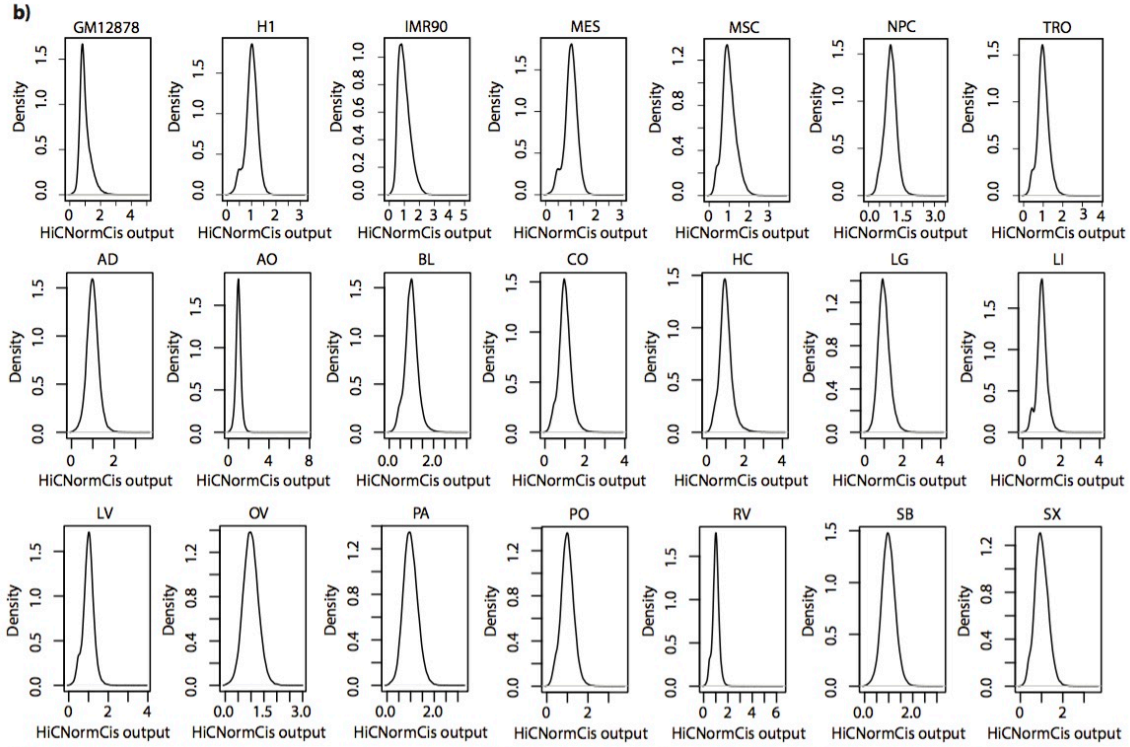

c)

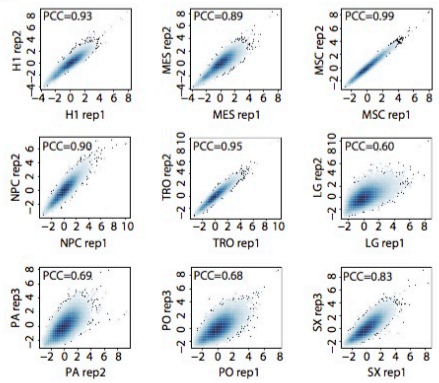

d)

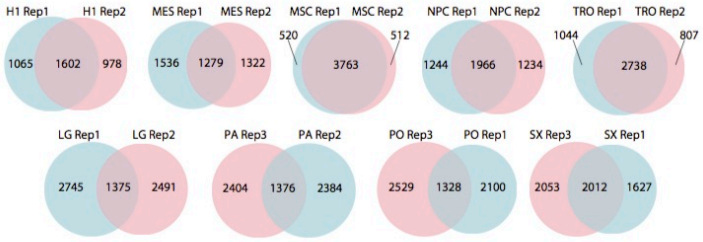

e)

Sample	cis>15kb
80%	153,016,660
60%	114,766,901
40%	76,509,297
20%	38,255,359
10%	19,130,174

Sample	80%	60%	40%	20%	10%
80%	1.00	0.97	0.95	0.94	0.94
60%	0.97	1.00	0.98	0.98	0.97
40%	0.95	0.98	1.00	0.99	0.99
20%	0.94	0.98	0.99	1.00	0.99
10%	0.94	0.97	0.99	0.99	1.00

f)

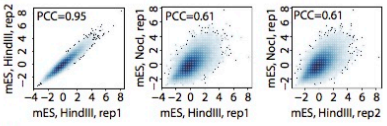

g)

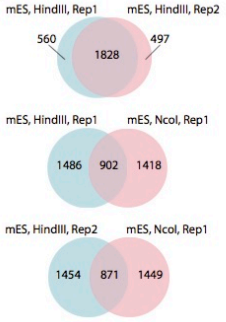

h)

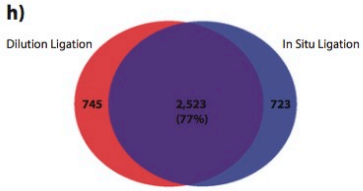

i)

Sample	before HiCNormCis			After HiCNormCis		
	F	GC	M	F	GC	M
GM12878	0.540	-0.162	0.069	0.049	-0.023	0.000
H1	0.768	-0.441	0.175	0.175	-0.117	0.000
IMR90	0.528	0.015	0.170	0.066	-0.021	-0.004
MES	0.774	-0.372	0.171	0.169	-0.113	0.000
MSC	0.696	-0.299	0.163	0.105	-0.065	-0.003
NPC	0.700	-0.451	0.189	0.105	-0.056	0.003
TRO	0.708	-0.230	0.179	0.126	-0.070	-0.003
CO	0.675	-0.257	0.214	0.113	-0.059	-0.009
HC	0.661	-0.291	0.221	0.100	-0.053	-0.010
OV	0.729	-0.552	0.177	0.149	-0.100	0.001
SB	0.757	-0.407	0.188	0.150	-0.091	-0.002
LG	0.524	0.181	0.206	0.061	0.015	-0.010
PO	0.682	-0.151	0.188	0.114	-0.054	-0.010
PA	0.627	0.121	0.181	0.105	0.000	-0.013
SX	0.211	0.509	0.136	-0.031	0.128	-0.007
LI	0.720	-0.396	0.115	0.108	-0.069	-0.001
LV	0.770	-0.474	0.099	0.152	-0.095	0.001
RV	0.758	-0.432	0.100	0.142	-0.091	0.000
AD	0.771	-0.406	0.185	0.154	-0.099	-0.004
AO	0.799	-0.538	0.121	0.177	-0.118	0.004
BL	0.721	-0.297	0.135	0.124	-0.072	-0.003

j)

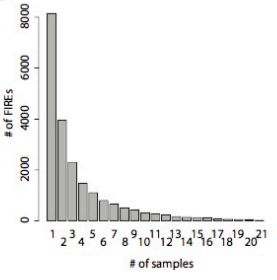

Figure S2. FIRE calling methodology, related to Figure 2.

A) Box plot showing the distribution of the total raw 15-200kb cis interactions per bin in each sample. Box plots for various degrees of down-sampling of H1 rep2 (Dixon et al, Nature, 2015) are shown in green, cell lines are shown in blue, and primary tissue Hi-C data is shown in yellow.
B) Table showing the Pearson correlation coefficients (PCC) between local contact summation of each bin with their respective effective restriction fragment length, GC content, and mappability (as rows). The normalization method (or lack thereof for raw matrix) to prepare the Hi-C contact data is listed as column headers. PCC values are rounded to the nearest hundredth.
C) Density plots showing the distribution of HiCNormCis outputs for each sample. The y-axes show the density and the x-axes are the HiCNormCis output values. The sample name is indicated in the title of each plot.
D) Scatterplots showing the genome-wide pairwise correlation of FIRE score between two biological replicates for H1, MES, MSC, NPC, TRO, LG, PA, PO, and SX. Inset is the Pearson correlation coefficients.
E) Pie charts showing the overlapping FIRE calls in 9 pairs of biological replicates from cell lines or tissues. Same 9 samples as Panel D. (Chi-square test p value < 2.2e-16).
F) Left, table showing the number of long-range cis interactions in a down-sampled replicate of H1 (H1 rep2 from Dixon et al, 2015) Hi-C data. The 'Sample' column indicates what fraction of the full dataset was extracted during down-sampling, and 'cis>15kb' is the total number of long-range cis interactions from the down-sampled data. To the right, a table showing the Pearson correlation coefficient (PCC) of the genome-wide FIRE scores for down-sampled H1 data. Each row/column corresponds to what down-sampled fraction of the Hi-C data was used for the correlation analysis. Each table entry is the PCC.
G) Scatter plots showing the genome-wide Pearson correlation coefficient (PCC) between 3 different samples, including two biological replicates of mES cells prepared using HindIII and 1 sample of mES cells prepared using NcoI (data from Dixon et al, 2012). Inset is the genome-wide PCC value.
H) Pie chart showing the significant FIRE bin overlap between two biological replicates of mES cells prepared with HindIII (left), or mES HindIII rep1 and mES NcoI (middle), or mES HindIII rep2 and mES NcoI (right). (Chi-square test p value < 2.2e-16).
I) Pie charts showing the significant FIRE bin overlap between samples either prepared using the in situ ligation procedure (right) or the "dilution" ligation procedure (left). (Chi-square test p value < 2.2e-16).
J) Table showing the Pearson correlation coefficients (PCC) for total cis interactions counts (within 15-200kb distance) and fragment length of a given bin (column 'F'), GC content (column 'GC'), and mappability (column 'M'), either before (group 'Before HiCNormCis'), or after normalization (group 'After HiCNormCis'), and for each sample (rows).

Figure S3

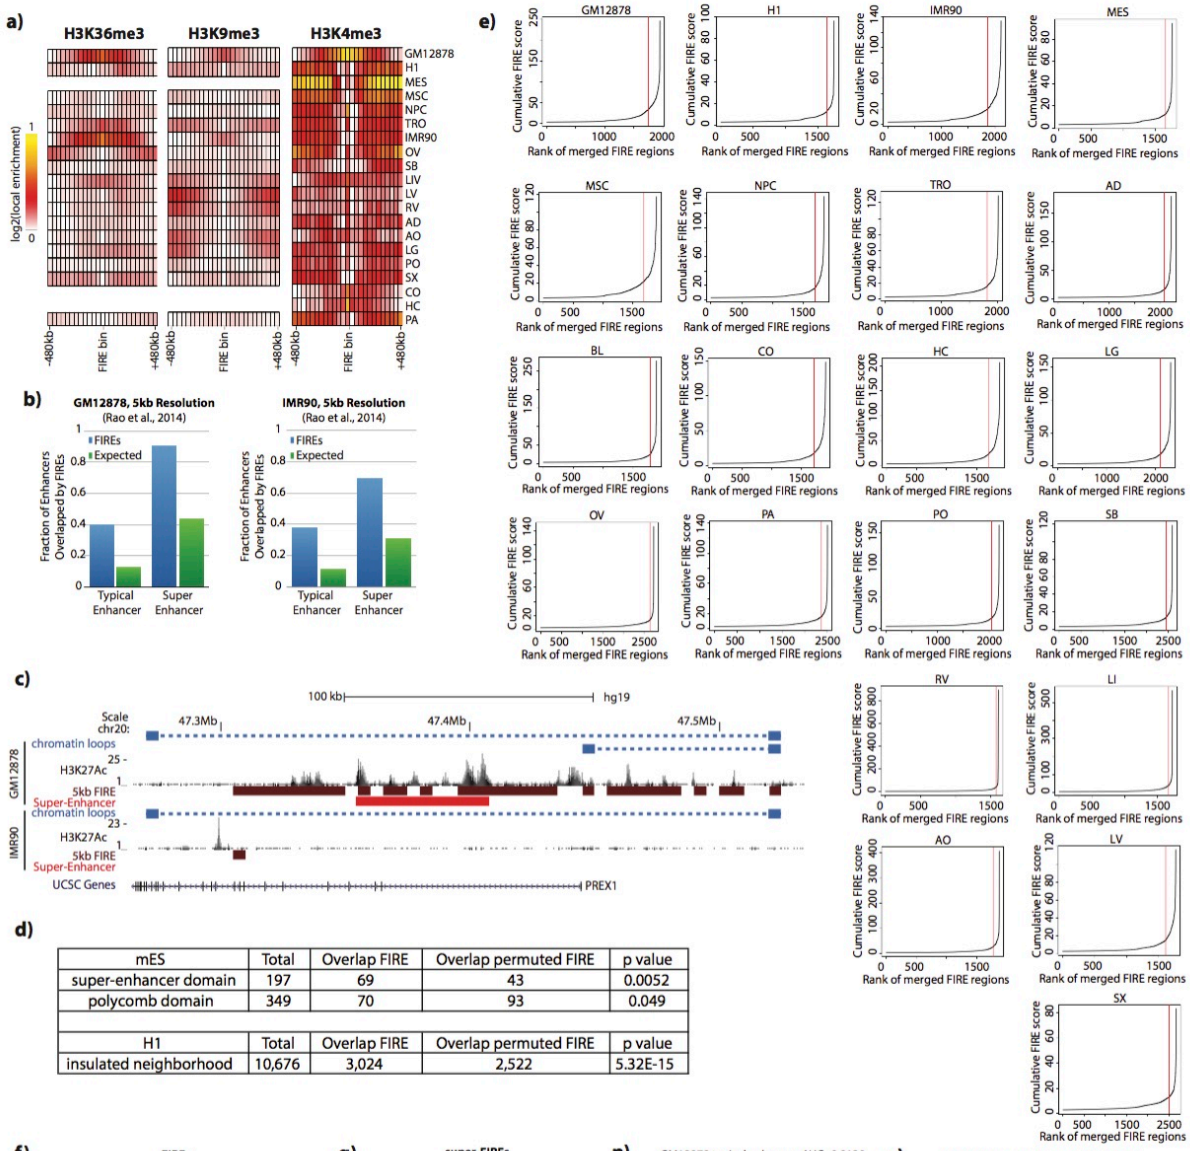

mES	Total	Overlap FIRE	Overlap permuted FIRE	p value
super-enhancer domain	197	69	43	0.0052
polycomb domain	349	70	93	0.049
H1	Total	Overlap FIRE	Overlap permuted FIRE	p value
insulated neighborhood	10,676	3,024	2,522	5.32E-15

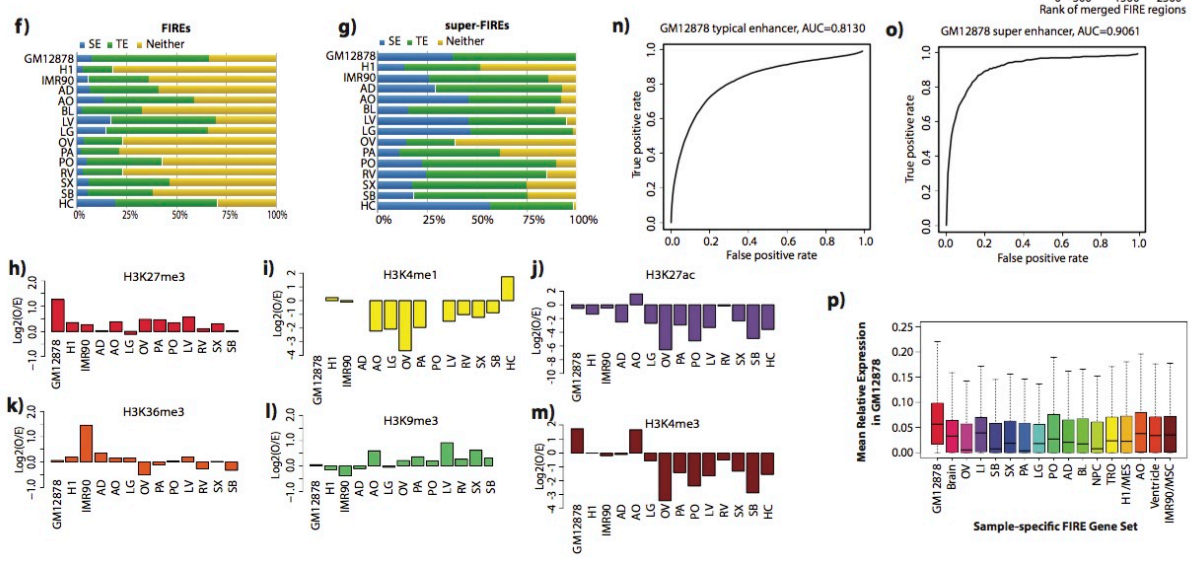

Figure S3. Analysis of chromatin biochemical features at FIREs and super-FIREs, related to Figure 4.

A) Heatmaps showing the local enrichment (see Supplemental Methods) of H3K36me3 (left), H3K9me3 (middle), and H3K4me3 (right), centered on FIRE bins for each cell line or adult tissue. Local enrichment is calculated relative to the peaks per bin for H3K4me3, and RPKM values for H3K36me3 and H3K9me3. H3K36me3 and H3K9me3 data was not available for CO or HC.
B) Bar plots showing the fraction of typical or super-enhancers overlapped by observed FIRE calls (blue bars) in GM12878 (left plot) and IMR90 (right plot) at 5kb resolution (Rao et al., 2014), or size-matched randomly permuted FIRE calls (green bars). Within each plot, analysis of typical enhancers is on the left, analysis of super-enhancers is on the right.
C) Genome browser snapshot of the PREX1 locus (chr20:47,263,536-47,534,527) in GM12878 (top set of tracks) and IMR90 (bottom set of tracks). Shown for each cell line are previously annotated (Rao et al., 2014) chromatin loops (blue; square is loop anchor, dash to loop), H3K27Ac signal (black), FIREs defined at 5kb resolution (brown), and previously annotated (Hnisz et al., 2013) super-enhancers (red). The bottom of the snapshot shows the positioning of USCS genes at this locus.
D) Table showing the overlap between FIREs, super-enhancer domains, polycomb domains in mESCs (Dowen et al., 2014) (top section) and insulated neighborhoods in H1 cells (Ji et al., 2016) (bottom section). Tabulated are the Total number of domains or insulated neighborhoods, how many are overlapped by a FIRE, and how many are expected to overlap based on random permutation of FIRE positioning in that respective cell type. The Chi-square test p-value is reported in the right column.
E) Line plots showing the cumulative FIRE scores (y-axis) of ranked stitched FIRE bins (x-axis) from the FIREs with the lowest cumulative FIRE scores (left side) to the highest FIRE scores (right side). The red vertical line indicates the inflection point, whereby stitched FIRE bins to the right of this line are called as super-FIREs.
F) Stacked bar plots showing the fraction of FIREs containing at least 1 super-enhancer (SE, blue bars), typical enhancer (TE, green bars), or no SE or TE (yellow bars). Each row is the analysis of a different cell or tissue type.
G) Same as Panel F, except analysis of super-FIREs.
H) Bar plots showing the enrichment (y-axis) of H3K27me3 at super-FIREs that do not contain any annotated typical enhancer or super-enhancers. Each bar represents the analysis of a different tissue, which has been previously annotated for super-enhancers (Hnisz et al., 2013). Hippocampus (HC) tissue is not shown because there is no H3K27me3 ChIP-seq data in HC.
I) Same as Panel H, except analysis of H3K4me1.
J) Same as Panel H, except analysis of H3K27ac.
K) Same as Panel H, except analysis of H3K36me3. No ChIP-seq data available for HC.
L) Same as Panel H, except analysis of H3K9me3. No ChIP-seq data available for HC.
M) Same as Panel H, except analysis of H3K4me3.
N) Line plot showing the relationship between the True Positive rate, defined as the fraction of FIRE bins overlapping typical enhancers (Hnisz et al., 2013), and the False Positive rate, defined as the fraction of FIRE bins not overlapping a typical enhancer, as a function of the significance threshold using to define FIREs in GM12878 cells. (AUC=0.813).
O) Same as Panel N, except for super-enhancers (Hnisz et al., 2013). (AUC=0.9061).
P) Genome-wide analysis showing the relative gene expression levels for genes within 200kb of GM12878-specific FIREs. Genes within 200kb of GM12878-specific FIREs were collected, and then for each sample, the relative gene expression levels are calculated. Shown are the box plots of the distribution of relative gene expression levels for each sample indicating that GM12878 relative gene expression levels are higher than any other sample (Two-sample t-test p-value < 2.2e-16 compared to brain, OV, LI, SB, SX, PA, LG, AD, NPC, ventricle, and IMR90/MSC; p-value < 5.66e-7 compared to PO; p-value < 2.93e-8 compared to BL; p-value < 1.04e-9 compared to TRO; p-value < 4.84e-10 compared to H1/MES; p-value < 9.26e-6 compared to AO). Boxplots show the median (black line) and interquartile range.

Figure S4

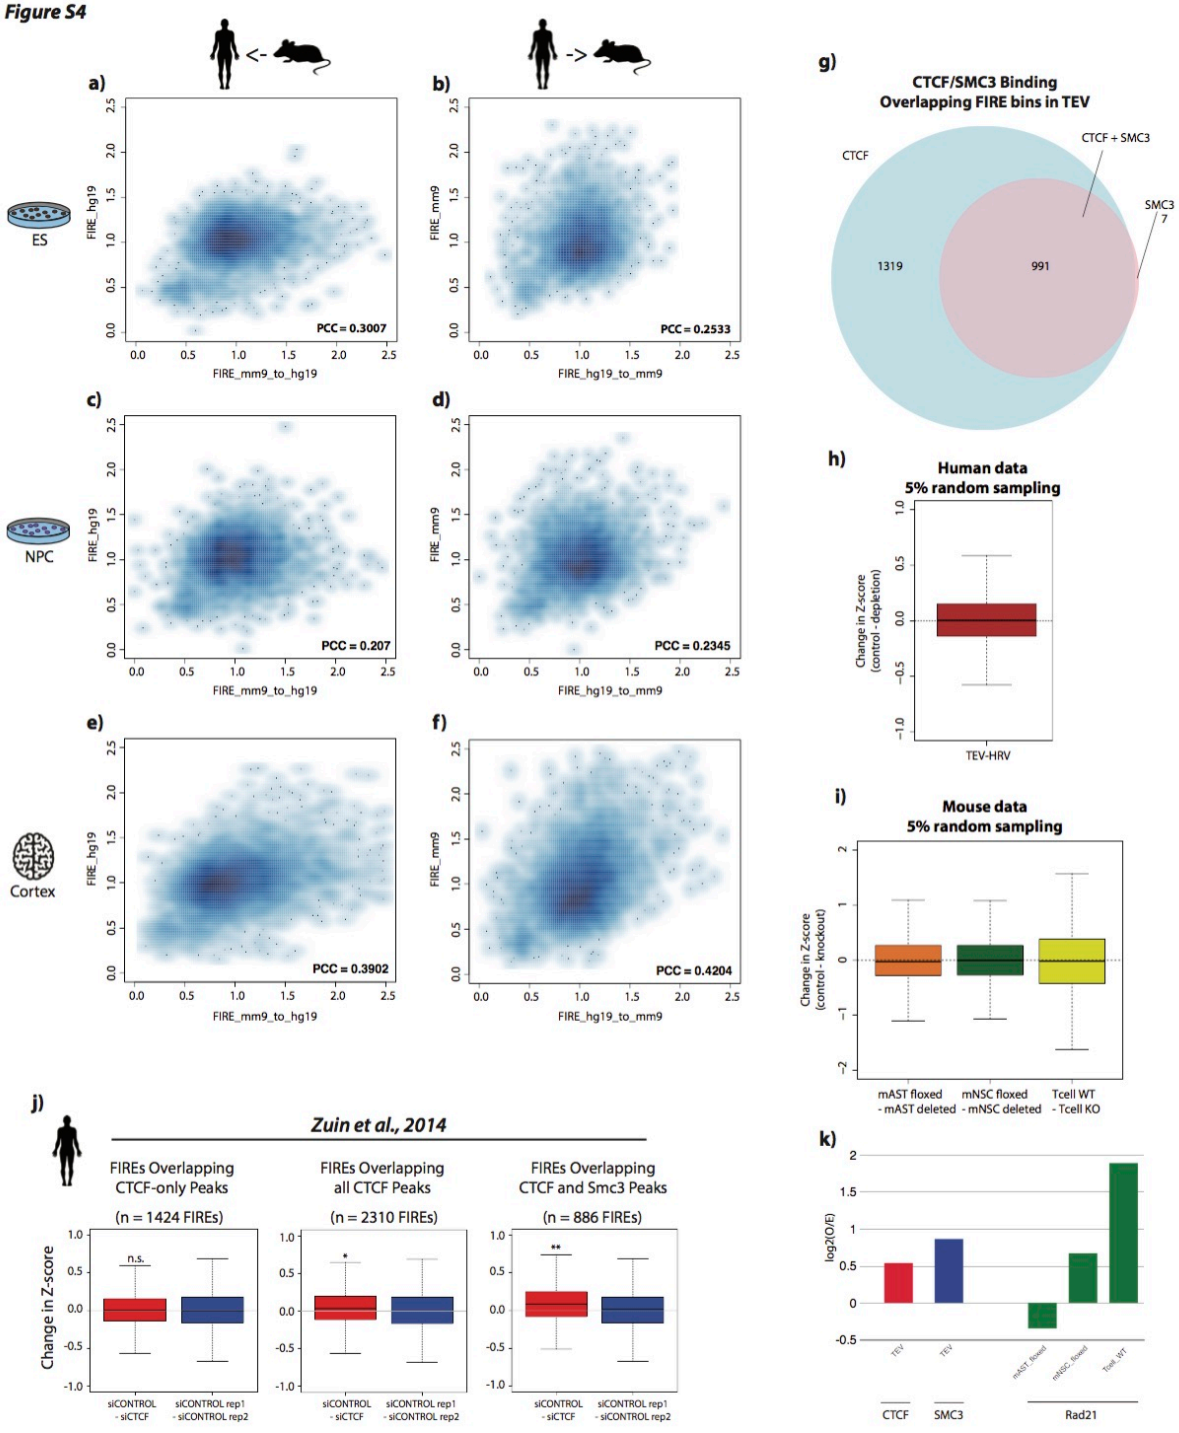
a)
b)
c)
d)
e)
f)
ES
NPC
Cortex
FIRE_hg19
FIRE_mm9
FIRE_mm9_to_hg19
FIRE_hg19_to_mm9
PCC = 0.3007
PCC = 0.2533
PCC = 0.207
PCC = 0.2345
PCC = 0.3902
PCC = 0.4204
g)
CTCF/SMC3 Binding
Overlapping FIRE bins in TEV
CTCF
CTCF + SMC3
SMC3
7
1319
991
h)
Human data
5% random sampling
Change in Z-score
(control - depletion)
TEV-HRV
i)
Mouse data
5% random sampling
Change in Z-score
(control - knockout)
mAST floxed
- mAST deleted
mNSC floxed
- mNSC deleted
Tcell WT
- Tcell KO
j)
Zuin et al., 2014
FIREs Overlapping
CTCF-only Peaks
(n = 1424 FIREs)
FIREs Overlapping
all CTCF Peaks
(n = 2310 FIREs)
FIREs Overlapping
CTCF and Smc3 Peaks
(n = 886 FIREs)
Change in Z-score
n.s.
*
**
siCONTROL
- siCTCF
siCONTROL rep1
- siCONTROL rep2
k)
log2(O/E)
TEV
TEV
mAST_floxed
mNSC_floxed
Tcell_WT
CTCF
SMC3
Rad21

Figure S4. FIRE score species conservation and reduction upon loss of Cohesin, related to Figure 5.

A) Scatterplot showing the correlation between randomly selected non-FIRE bins in mouse ES cells that liftover to the hg19 reference genome. Shown on the x-axis are the FIRE scores from the randomly selected mouse bins that can be liftover to hg19. Shown on the y-axis are the FIRE scores in the corresponding human bins. The PCC value is shown in the bottom right corner.
B) Scatterplot showing the correlation between randomly selected non-FIRE bins in human ES cells that liftover to the mm9 reference genome. Shown on the x-axis are the FIRE scores from the randomly selected human bins that can be liftover to mm9. Shown on the y-axis are the FIRE scores in the corresponding mouse bins. The PCC value is shown in the bottom right corner.
C) Same as Panel A, expect using NPC cell data.
D) Same as Panel B, except using NPC cell data.
E) Same as Panel A and C, except using cortex tissue data.
F) Same as Panel B and D, except using cortex tissue data.
G) Pie charts showing the overlap between FIRE bins called in the TEV sample and bins bound by CTCF only (blue shading, left), SMC3-only (pink shading, right), or co-bound peaks (blue+pink overlap, center).
H) Box plots depicting the change in Z-score in a random sampling of 5% of bins in TEV and HRV cells. There is no significant change in FIRE score in either comparison. Change in z-score is used for comparison, rather than change in FIRE score (-ln(p-value)), since z-score has approximate Gaussian distribution.
I) Same as Panel H, except for comparing mAST (floxed – deleted, left boxplot), mNSC (floxed-deleted, middle boxplot), and T-cells (WT-Knockout). In all cases, there is not significant change in FIRE score at a random sampling of FIRE bins. Change in z-score is used for comparison, rather than change in FIRE score (-ln(p-value)), since z-score has approximate Gaussian distribution.

J) Box plots showing the change in Z-score at FIREs overlapping bins bound by CTCF but not SMC3 “CTCF-only” (left column), all CTCF peaks (middle column), and CTCF and SMC3 co-binding (right column) for the comparison of siCONTROL and siCTCF samples. The red boxes show distributions of FIRE score change at FIRE bins called in wild type cells minus the mutant cells, while the blue boxes are distributions for FIRE score change at FIRE bins called in wild type cells but between biological replicates of wild type cells. These comparisons show the significant reduction of FIRE score at all CTCF peaks, and especially at CTCF SMC3 co-bound peaks overlapping FIRE bins (*p=4.88e-5, **p=3.89e-9; two sample t-test). Change in z-score is used for comparison, rather than change in FIRE score (-ln(p-value)), since z-score has approximate Gaussian distribution.

K) Bar plots showing the significant enrichment of CTCF, SMC3, or Rad21 in FIREs from control samples in 3 different studies (From left to right - One-sample t-test p value < 1.11e-15, < 6.54e-14, < 1.71e-10, < 1.33e-13, and < 2.2e-16). The sample name is indicated across the x-axis, and the log2(O/E) values are plotted on the y-axis.

Figure S5

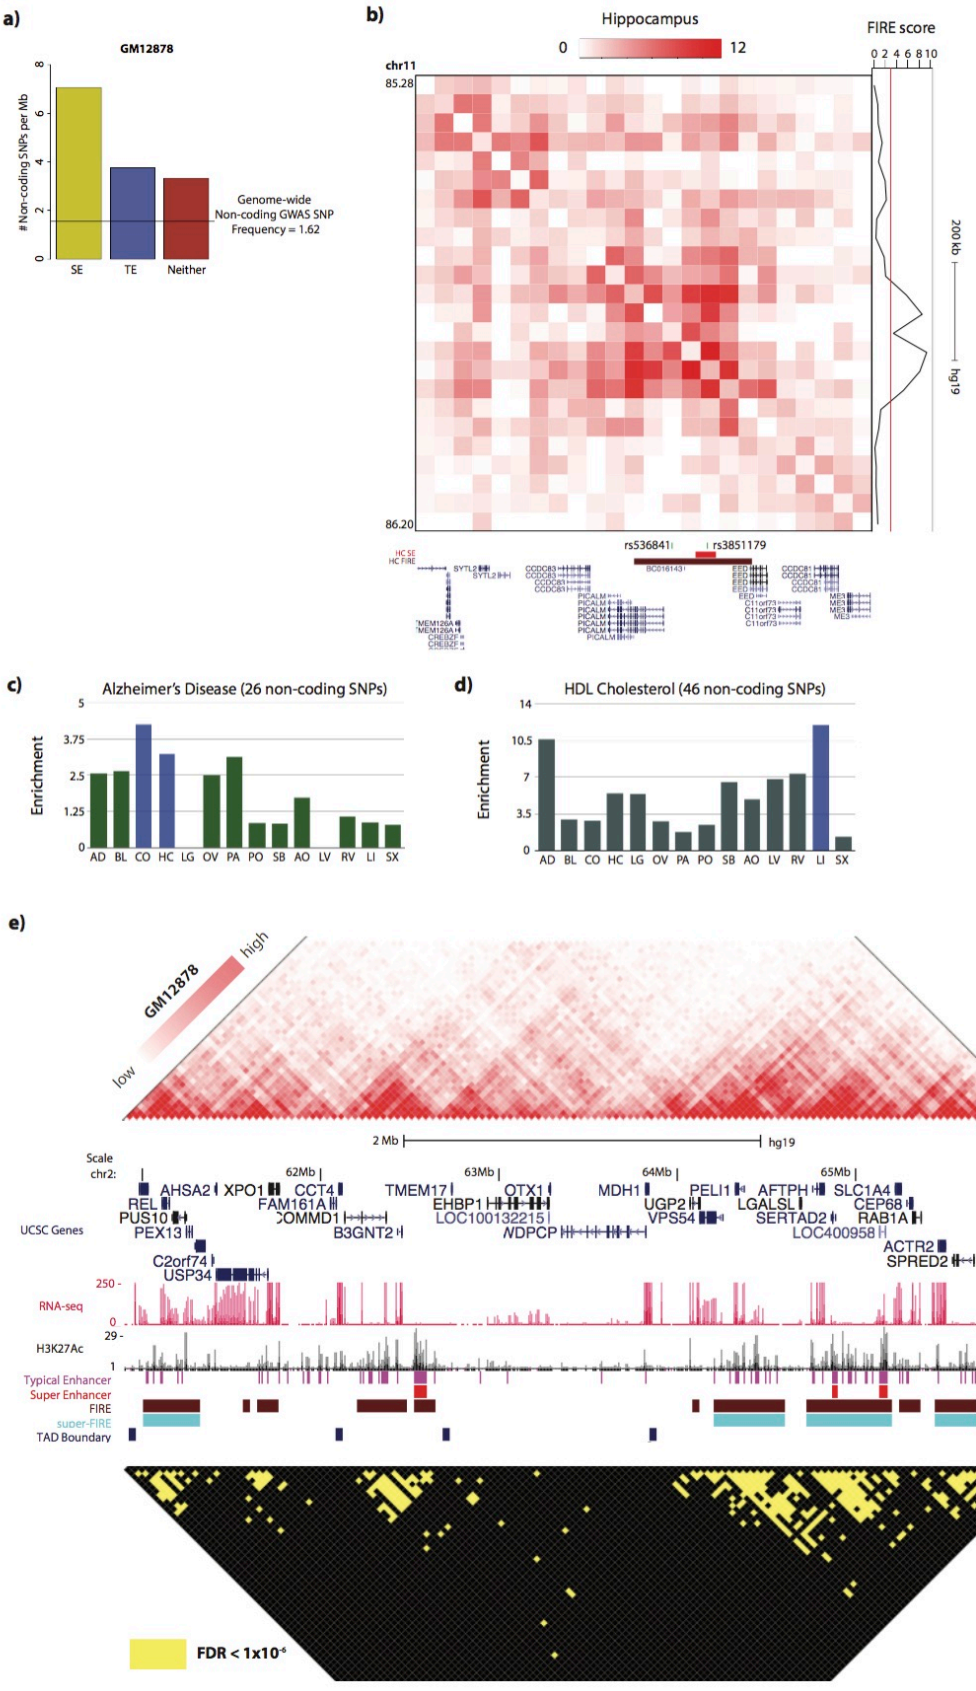
a)
GM12878
Non-coding SNPs per Mb
Genome-wide
Non-coding GWAS SNP
Frequency = 1.62
SE
TE
Neither
b)
Hippocampus
0
12
FIRE score
chr11
85.28
86.20
200 kb
hg19
rs536841
rs3851179
c)
Alzheimer's Disease (26 non-coding SNPs)
Enrichment
AD BL CO HC LG OV PA PO SB AO LV RV LI SX
d)
HDL Cholesterol (46 non-coding SNPs)
Enrichment
AD BL CO HC LG OV PA PO SB AO LV RV LI SX
e)
GM12878
high
low
2 Mb
hg19
Scale
chr2:
62Mb
63Mb
64Mb
65Mb
UCSC Genes
AHSA2
XPO1
CCT4
TMEM17
OTX1
MDH1
PELI1
AFTPH
SLC1A4
REL
FAM161A
EHBP1
UGP2
LGALSL
CEP68
PUS10
COMMD1
LOC100132215
VPS54
SERTAD2
RAB1A
PEX13
B3GNT2
NDPCP
LOC400958
ACTR2
C2orf74
SPRED2
USP34
RNA-seq
250
0
H3K27Ac
29
1
Typical Enhancer
Super Enhancer
FIRE
super-FIRE
TAD Boundary
FDR < 1x10⁻⁶

Figure S5. Analysis of non-coding disease-associated SNPs in FIREs and FIRE-FIRE contacts, related to Figure 6.

A) Bar plot showing the number of non-coding GWAS SNPs per megabase in FIRE overlapping super-enhancers (SE), FIREs overlapping typical enhancers (TE), and FIREs not overlapping either TE or SE. The horizontal line indicates the genome-wide SNP frequency. All analysis was done using GM12878 FIRE data.
B) Normalized Hi-C contact matrix of a 920kb locus (chr11:85,280,000-86,200,000) in human hippocampus tissue (HC). The tracks below show the presence of two Alzheimer's disease associated SNPs (rs536841 and rs3851179) located within a broad FIRE region (brown, chr11:85,840,000-85,880,000). One SNP resides within a HC super-enhancer (red) and the other SNP resides outside of the super-enhancer but within the FIRE region. Both SNPs reside in close proximity to gene PICALM, as shown in the bottom UCSC gene track. Right of the Hi-C contact matrix is the continuous FIRE score across this locus.
C) Enrichment of Alzheimer's disease-associated SNPs across 14 primary adult tissue FIRE annotations, showing the highest enrichment in FIREs from both brain tissues (CO and HC).
D) Enrichment of SNPs associated with quantitative HDL cholesterol metrics across 14 primary adult tissue FIRE annotations, showing the highest enrichment in liver FIREs.
E) Normalized Hi-C contact matrix (top) in GM12878 for a 5.14Mb locus (chr2:60900000-66040000) illustrating the extent of statistically significant FIRE-FIRE interactions. Hi-C color scale ranges from low to high, corresponding to the 15^{th} and 99^{th} percentile contact frequencies within this locus. The reflected matrix shows the statistically significant (FDR<1e-6) Hi-C contacts within 2Mb genomic distance across the locus. Only bin-pairs with FDR<1e-6 are yellow, and the rest are black. Between the matrices are a UCSC gene annotations (blue, top), RNA-seq data (red), H3K27Ac data (black), typical enhancer annotations (Hnisz et al., 2013) (purple), super-enhancer annotations (Hnisz et al., 2013) (red), FIRE annotations (brown), super-FIRE annotations (cyan), and TAD boundary calls (blue).

2. Supplemental tables.

Table S1. Hi-C Data Manifest and Quality Metrics, Related to Figure 1

Table S2. Compartment A/B Patterns and PC1 values, Related to Figure 1, 2

Table S3. TAD boundary annotations, Related to Figure 1, 2

Table S4. Fit-Hi-C peak calling summary, and related analyses, Related to Figure 6, 7

Table S5. Genome-wide FIRE scores, Related to Figure 2

Table S6. FIRE calls and sample-specific FIRE calls in the primary cohort, Related to Figure 2, 3

Table S7. Observed and Expected Values of FIREs in Compartment A and B, Related to Figure 2

Table S8. Gene Ontology (GO) analysis of genes near sample-specific FIREs; top biological process terms, Related to Figure 3

Table S9. Gene Ontology (GO) analysis of genes near sample-specific FIREs; top disease ontologies, related to Figure 3

3. Supplemental experimental protocols.

Tissue Collection

For all human tissues except for dorsolateral prefrontal cortex (CO) and hippocampus (HC), samples were collected as previous described as part of the Epigenome Roadmap Consortium collection (The Roadmap Epigenomics Consortium, 2015). Human dorsolateral prefrontal cortex (CO) and hippocampus (HC) tissue were obtained from the National Institute of Child Health and Human Development (NICHD) Brain Bank for Developmental Disorders. Ethics approval was obtained from the University Health Network and The Hospital for Sick Children for use of the tissues. The two specimens used here were from a single male donor, age 31, who was classified as healthy.

4. Computational methods.

Histone ChIP-Seq data processing and peak-calling

Published single- or paired-end ChIP-Seq raw data were downloaded for H3K4me1, H3K4me3, H3K27ac, H3K9me3, H3K27me3, and H3K36me3 from GEO database under accession number GSE16256 and from SRA database under accession umber SRP000941 (Roadmap Epigenomics Consortium et al, 2015). The raw data were aligned to hg19 human reference genome using BWA-mem. Unmapped, non-uniquely mapped, and low quality (less than 10 quality score) reads were removed. We also removed PCR duplicate reads with PicardTools. ChIP-seq peaks were identified using MACS2 with the following parameters (--format=BAM -g mm -m 5 50 -p 1e-5) with corresponding input ChIP-Seq data as a background model. We also calculated input normalized RPKM values for H3K9me3, H3K27me3, and H3K36me3 in each 40kb bin.

RNA-Seq data processing

Published RNA-Seq raw sequencing data were downloaded from GEO database under accession number GSE16256 SRP000941 (Roadmap Epigenomics Consortium et al, 2015). RNA-Seq raw reads were aligned to hg19 human reference genome using BWA-mem. Unmapped and non-uniquely mapped reads were removed. Transcription levels were obtained based on GENCODE annotation v19 and normalized to FPKM values using Cufflinks. FPKM values from multiple replicates or multiple donors were combined together and the mean FPKM value was calculated for each gene.

Hi-C data processing

Unpublished Hi-C libraries described in this manuscript were sequenced on either Illumina HiSeq2000 or HiSeq2500 instrument. All other published Hi-C data were downloaded from SRA and converted to paired-end FASTQ files. Paired-end reads were then aligned independently to either the hg19 human reference genome or mm9 mouse reference genome using BWA-mem. As BWA-mem retains multiple alignments for a single read-end if it maps in two locations (i.e. a chimeric read), we kept only the 5' alignments for each read-end. Read-pairs in which both read-ends had mapping quality greater than 10 were paired using in-house scripts and converted into BAM files using Samtools. PCR duplicates were then removed using PicardTools. If down-sampling was performed, we then used PicardTools 'DownsampleSam' function to down-sample this final processed BAM file. Then, raw contact matrices were constructed using in-house scripts, and then further processed using HiCNormCis (described below) for the FIRE analysis. For all other Hi-C analyses not pertaining to FIRE scores, Hi-C data were normalized using HiCNorm (Hu et al., 2012), Vanilla Coverage (Rao et al., 2014), or ICE (Imakaev et al., 2012), where indicated. For all datasets of similar nature [such as the main cell lines in this study (GM12878, IMR90, H1, H1-derived) or the primary tissue collection, or the samples from each respective publications], we performed quantile normalization on HiCNorm matrices to normalize for differences in sequencing depth between samples within each group. This was done prior to any downstream comparative analyses.

Compartment A/B Calling

Compartment A/B analysis was performed at 1Mb resolution as previously described (Lieberman-Aiden et al., 2009). First of all, we calculated the average read count for each 1Mb bin in each sample. For cell line data, we removed 1Mb bins with average read count <=100. For tissue data, we removed 1Mb bins with average read count <=10. We used different thresholds for cell line data and tissue data, since tissue data have generally lower sequencing depth than the cell line datasets. Such filtering step has removed around 10% low coverage regions in the entire genome. Due to varying sequencing depths, the filtered regions are slightly different in each sample, and only bins which had a numeric value across all samples were used for downstream compartment analysis. After generating the first three principle components using the 'prcomp' function in R on the Pearson correlation matrix, we visually examined the first principle component (PC1) in each of 7 cell lines and 14 tissues, and found that for a few tissues the PC1 vectors of chr3 and chrX correspond to two chromosome arms, instead of A/B compartment. In specific, these outliers are PC1 vector of chr3 in bladder (BL), dorsolateral prefrontal cortex (CO), hippocampus (HC), lung (LG), psoas muscle (PO), aorta (AO), left ventricle (LV), right ventricle (RV), and PC1 vector of chrX in adrenal gland (AD), dorsolateral prefrontal cortex (CO), hippocampus (HC), pancreas (PA), psoas muscle (PO), left ventricle (LV), right ventricle (RV). For those outlier, the second principle component (PC2) was used to call A/B compartment. Visual examination of those PC2 vectors confirmed they match to the plaid-pattern observed in the normalized Hi-C contact matrices, instead of two chromosome arms.

Compartment A/B Conservation Analysis

To estimate the degree of compartment label conservation (related to Figure 1b, c; Figure S1f), we first scanned every 1Mb bin across the genome and counted the number of cell lines or tissue types that shared the same compartment label, and recorded which label was shared. By performing this at genome-wide scale, we obtained an observed distribution of A/B compartment conservation (Figure 1c, Figure S1f). To statistically determine if this distribution deviates from expectation, or to statistically test the significance of ubiquitous conservation (same label in all cell lines and tissue types), we first created an expected distribution of compartment conservation. First, for each cell line or tissue type, we randomly permuted the compartment label for each bin, while preserving the total number of A or B compartments on each chromosome. We then conducted the same conservation enumeration described for the observed data, and obtained an expected distribution of conservation (Figure S1f). This distribution was compared to the observed distribution using a Chi-square test. Testing the significance of observing the same compartment label ("ubiquitous conservation") across all cell lines or tissue types was done by comparing to the expected values using a binomial test.

TAD Boundary Reproducibility and Conservation Analyses

To estimate the degree of TAD boundary region conservation across samples in the primary cohort (related to Figure 1d, e; Figure S1e), we first identified TAD boundaries at 40Kb bin resolution for each sample independently, and then concatenated unique boundary bins across all samples into a single putative boundary region reference file. Consecutive TAD boundaries within 200Kb distance were also merged into a TAD boundary "region". Merging of adjacent boundary bins was performed because often times larger TAD boundaries (up to 400Kb) may result in slightly shifted (by a few bins) boundary calls between samples, and though they do not directly overlap, then both are a bin within the same boundary region. Moreover, in previous reports, TAD boundaries have been defined as 40-400Kb (Dixon et al., 2012) while regions >400kb are characterized as regions of "disorganized chromatin". Given this, and after defining boundary "regions" using our approach, the final list of unique TAD boundary regions ranged in size from 40-400Kb, consistent with previous definitions (Dixon et al., 2012). Using the cumulative list of TAD boundary regions, we evaluated the fraction of the total number of cell lines and tissues that had a boundary bin overlap with the given boundary region. To evaluate the overlap of TAD boundaries between tissue Hi-C biological replicates (LG, PA, PO, SX), boundaries within 80kb of each other were considered overlapping, which may underestimate the true boundary overlap since TAD boundaries have been previously defined as up to 400kb, and large boundaries regions are subject to technical variation in TAD calling at 40kb resolution. A chi-square test was used to evaluate statistical significance of TAD boundary overlap between replicates.

TAD Boundary Reproducibility and Conservation Analyses

To understand if our TAD identification method is robust across the sequencing depths used in this manuscript, we down-sampled H1 rep2 Hi-C data (Dixon et al., 2015) as described above, and constructed HiCNorm contact maps. We then applied the insulation square method (Crane et al., 2015) to identify TAD boundaries. To determine what fraction of TAD boundaries within a given down-sampled dataset overlap other putative TAD boundaries in H1 down-sampled data, we first collected all putative TAD boundary regions from each of the 7 samples and made a reference putative boundary file (approx. 2,700 putative TAD boundary regions). For each down-sampled dataset, we then asked what fraction of TAD boundary regions overlaps the boundaries in the reference putative boundary list (related to Figure S2h). To understand what fraction of TAD boundary regions are shared across all down-sampled datasets we calculated the percentage of TAD boundaries that were unique to subsets of the down-sampled files, including TAD boundaries that were shared across all down-sampling datasets (related to Figure S2i).

Comparison of FIREs and chromatin loops and insulated neighborhoods

To explore the relationship between FIREs and chromatin loops, we called FIREs using the methods described in this manuscript, except at 5kb resolution using in situ Hi-C data in GM12878 and IMR90 (Rao et al., 2014). To compute the enrichment of chromatin loops in FIREs, we first assigned each chromatin loop anchor to a 5kb bin using the previously published loop annotations. We then computed the observed overlap between 5kb FIREs and 5kb loop anchors, and the expected overlap by permuting the FIRE positioning. Statistical significance was computed using Chi-square test. Conversely, to analyze the enrichment for FIREs at chromatin loop anchors, we conducted the same type of analysis, except asking what fraction of loop anchors are overlapped by a FIRE.

To explore the relationship between FIREs and insulated neighborhoods, super-enhancer domains and polycomb domains, we computed the enrichment (observed overlap / expected overlap) of 40kb FIREs at insulated neighborhoods defined in H1 cells (Ji et al., 2016), and the enrichment of 40kb FIREs at super-enhancer domains and polycomb domains in mESCs (Dowen et al., 2014). Statistical significance was computed using Chi-square test.

Identifying super-FIREs

To identify super-FIREs, we used a similar approach of that used to identify super-enhancers (Hnisz et al., 2013). First we merged all book-ended FIRE bins into large continuous FIRE regions. We then ranked the merged FIRE regions by their cumulative Z-score, and plotted the ranked FIRE regions as a function of their cumulative Z-score (related to Figure S3c). We then found the inflection point of the line plot, and defined the FIRE regions to the right of the inflection point as super-FIREs. The same procedure can be done for 5kb bin resolution FIREs, but by stitching FIRE bins within 15kb of one another.

Enrichment of FIRE in compartment A or compartment B

Using the compartment A/B calls at 1Mb resolution for each sample, observed FIRE bins were categorized into either compartment A or compartment B, depending on which compartment the FIRE bin resided. For all observed FIRE calls, the total compartment A overlap and compartment B overlap were enumerated ($O_{FIRE(A)}$ or $O_{FIRE(B)}$). To generate expected values, FIRE bins were randomly permuted while preserving the total number of FIREs per sample and per chromosome, and then re-categorized into either compartment A or compartment B ($E_{FIRE(A)}$ or $E_{FIRE(B)}$). Enrichment for compartment A or compartment B was calculated as either $\log2(O_{FIRE(A)}/ E_{FIRE(A)})$ and $\log2(O_{FIRE(B)}/ E_{FIRE(B)})$, respectively. To statistically evaluate the significance of enrichment of FIREs in compartment A or compartment B, for we created a two by two table using total compartment A overlap and compartment B overlap in observed FIRE calls ($O_{FIRE(A)}$ or $O_{FIRE(B)}$) and expected FIRE calls ($E_{FIRE(A)}$ or $E_{FIRE(B)}$), respectively. Chi-square test was performed to access the statistical significance (related to Table S7) and the process was performed independently for each sample.

FIRE positioning relative to TAD

For each sample and each FIRE bin, we found the TAD for which the FIRE bin resides using TAD calls for that given sample (related to Figure 2e, f). For each FIRE bin within a given TAD, we set the center position of the TAD to 0.5 relative distance units, corresponding to ‘halfway’ between each adjacent TAD boundary. We then computed

the distance from the TAD center to the boundary (D_{center}), as well as the distance of the FIRE bin to the nearest boundary (D_{FIRE}). Selecting the nearest boundary ensures the D_{FIRE} will always be less than or equal to D_{Center}. The relative distance units of the FIRE within a TAD are then computed as (D_{FIRE}/D_{Center})/2.

FIRE clustering analysis

We performed hierarchical clustering analysis using all samples in our primary cohort. Specifically, we first used the normalized total cis interaction (HiCNormCis) value for each 40Kb bin, and calculated the Euclidean distance of two genome-wide FIRE score vectors between any two samples, using the R function "dist". We then used the R function "hclust" with option "single linkage" to perform the hierarchical clustering analysis (related to Figure 3a). Next, we selected 40Kb bins which are cell line or tissue specific FIREs, and visualized their HiCNormCis scores using software JAVA TreeView (Saldanha, 2004).

Genomic Regions Enrichment of Annotations Tool (GREAT) analysis

We performed the GREAT analysis (McLean et al., 2010) to investigate the biological processes and disease ontologies for genes in the neighborhood of cell line or tissue specific FIRE bins (related to Figure 3d, e; Table S8-9). Specifically, we input our list of cell- or tissue-specific FIRE bins for each sample into the GREAT software (http://bejerano.stanford.edu/great/public/html/), and allowed the software to test neighboring genes for biological process and disease ontology enrichment. GREAT then evaluates the statistical significance of enrichment for each biological process, compared to the whole genome background. A Bonferroni-corrected Binomial test was used to obtain the p-value. Reported are the top fifteen biological processes ranked by the most significant p-values, in GM12878-specific FIREs and brain-specific FIREs, respectively (related to Figure 3e, f) and top terms for all samples as well as top disease ontologies are found in Tables S8-9.

Histone Local Enrichment Analysis

For each 40Kb FIRE bin in each sample, we calculated either the number of peaks per bin (for narrow peaks H3K27ac, H3K4me1 and H3K4me3) or the RPKM values per bin (for broad peaks H3K27me3, H3K9me3 and H3K36me3) and then calculated these values for each of the 12 bins upstream and 12 bins downstream of the FIRE bin, creating a vector of 25 values, centered on the FIRE bin (related to Figure 4b; Figure S3a). Those 25 values represent the histone mark profile in 1Mb region centered at each FIRE bin. As a control, to generate an expected histone mark profile, we randomly permuted the location of FIRE bins ten times within each sample, and calculated the averaged peak count or RPKM value at each position across ten random permutations. To calculate the local enrichment, we first calculated the ratio between observed value and expected value for each of the 25 positions around a FIRE bin, creating an enrichment score profile. Then, to assess the magnitude of local enrichment, we normalized each enrichment score relative to the local minima, by taking the log2 of the position enrichment divided by the minimum local enrichment. This converts the data to have a local enrichment of 0 at the local minima and specifically allows one to appreciate the enrichment of FIRE bins relative to the local neighboring bins, rather than relative to genome-wide levels.

Mean-rank Gene Set Test

To determine if genes near sample-specific FIREs tend to be expressed predominantly in the same tissue, we adapted the Mean-rank Gene Set Test concept, originally described in the 'Limma' R package (Ritchie et al., 2015) (https://bioconductor.org/packages/release/bioc/html/limma.html). Conceptually, the mean-rank gene set test evaluates whether a particular subset of genes is highly ranked relative to other genes in terms of a given statistic. Then using a Wilcoxon test, evaluates the null hypothesis that the mean rank of one's subset of genes is not different than the expected mean ranking. A 'p-value' is generated by using the 'WilcoxGST' function in the Limma R package whereby the statistic parameter is a ranked list of relative gene expression values (with 1 being the gene with the highest relative expression, defined more below), and the index parameter is the positional indices of the genes within 200kb of a sample-specific FIRE set. However, the Wilcoxon test only evaluates if the mean rank of the test genes are different from the expected ranking, therefore not specifically addressing whether the mean rank is more towards 1 compared to the expected ranking. Therefore, we present the results as the difference between the

expected rank and actual mean rank, whereby a positive value indicates that the mean ranking is closer to 1 than the expected ranking.

In more detail, for each cell line or tissue, we first collected genes whose transcription start site (TSS) is within 200kb of a sample-specific FIRE. The collection of these genes within 200kb of sample-specific FIREs make up the sample-specific FIRE gene set, termed "FIRE genes". To prepare the Relative Expression rank file for each cell line or tissue, we used RNA-Seq data to first filter out genes with zero FPKM in all 21 samples, and then transformed the expression values into Log2(FPKM+1) values. Next, we divided each gene expression value by its cumulative gene expression sum across all 21 samples, to create the relative gene expression value (related to Figure 4f). For each sample, we then sorted all genes by their relative gene expression to assign each gene an expression rank, with 1 being the gene with the highest relative gene expression in that sample. Using these ranks for each sample, we calculated the mean expression rank for genes from a sample-specific FIRE gene set (related to Figure 4h), and then across all sample-specific FIRE gene sets (related to Figure 4g). A gene set enriched for sample-specific expression is expected to have a lower numeric mean rank (towards 1). By random chance, the mean rank will be approximately half of the total number of expressed genes. Therefore, we defined the enrichment score as the expected mean rank – observed mean rank. A large positive enrichment score indicates that genes within 200kb of sample-specific FIREs are primarily expressed in that sample relative to others, whereas a large negative enrichment score indicates that genes within 200kb of sample-specific FIREs are lowly expressed in that sample relative to other samples.

FIRE bin conservation

To investigate the degree of conservation of FIRE bins between human and mouse in three difference cell types (related to Figure 5a, b), we first identified FIRE bins using our HiCNormCis approach in the human and mouse samples. Next we identified breakpoints of major genomic rearrangements between human and mouse based on UCSC "net" alignments (Chiaromonte et al., 2001; Kent et al., 2003; Schwartz et al., 2003). To identify breakpoints in hg19, we used the alignment where hg19 is the target genome and mm9 is the query genome (http://hgdownload.soe.ucsc.edu/goldenPath/hg19/vsMm9/hg19.mm9.net.gz). To identify breakpoints in mm9, we used the alignment where mm9 is the target and hg19 is the query (http://hgdownload.soe.ucsc.edu/goldenPath/mm9/vsHg19/mm9.hg19.net.gz). From each alignment, we calculated the genomic coordinates of the boundaries of all "fill" and "gap" blocks of size >50kb. We sorted these coordinates and then recursively merged those that are separated within 25kb into a single genomic interval. The resulting set of merged intervals defined our breakpoints. Any FIRE bins containing human<->mouse synteny breakpoints as defined above were removed from downstream analyses. UCSC liftover tool was then used to convert the genomic location of FIREs between hg19 human reference genome and mm9 mouse reference genome. Since in many cases the a 40kb bin in one species lifts over to a region that is part of 2 40kb bins in the other species, we considered a "conserved FIRE" if 1 of the 2 bins was a FIRE call. As a control, we also lifted over the genomic location of randomly permuted FIREs (that don't contain a breakpoint) between human and mouse, and calculated the number of FIREs that are conserved. For each of the six comparisons in Figure 6a, we also obtained the expected level of conservation. A Chi-square test was used to evaluate the statistical significance of FIRE conservation between human and mouse.

FIRE score conservation

To estimate the FIRE score conservation between human and mouse across a range of FIRE scores (related to Figure S4a-f), we randomly selected 4,000 40Kb bins, and used UCSC liftover tool to convert the genomic location of the randomly selected 40Kb bins between hg19 human reference genome and mm9 mouse reference genome. Since in many cases the a 40kb in one species lifts over to a region that is part of 2 40kb bins in the other species, we took the average FIRE score of the 2 40kb bins when conducting the correlation analysis. We then made a scatter plot of FIRE scores between the paired human and mouse datasets at the syntenic 40Kb bins, and calculated the Pearson correlation coefficients.

Change in FIRE score upon loss of Cohesin or CTCF

To investigate the impact of Cohesin loss on local interaction frequency (i.e. on FIRE tendency), we evaluated the change in local interaction frequency (as 'Change in Z-score') upon loss of Cohesin (related to Figure 5c-e) or CTCF (related to Figure S4j). In these analysis, we used the Z-score for each FIRE bin, instead of negative -ln(p-value), since Z-scores has approximate Gaussian distribution. For comparison of Z-score change between "control cells" (defined within each experiment as the condition without Cohesin manipulation or CTCF knockdown) and experimental cells (defined within each experiment as the condition with Cohesin depletion or knockout, or CTCF knockdown), we first identified the most confident FIRE bins in control cells, defined as FIRE bins in both control biological replicates. Next, we calculated the change of Z-score between control and experimental, at those selected most confident FIRE bins. As an analysis control, we also calculated the change of Z-score between two control biological replicates at the same set of high confidence FIRE bins. A two sample t-test was used to evaluate the statistical significance of the difference in Z-scores between control vs. experimental, as well as between two biological replicates of control samples. Since two WT biological replicates are symmetric, we took the absolute value of the difference in Z-score between the biological replicates. Therefore, the Z-score difference between two control biological replicates is always positive, and is a fair comparison to the Z-score difference between control and experimental.

CTCF and SMC3/Rad21 Enrichment Analysis

To determine if FIREs are enriched for CTCF or SMC3 (in TEV sample) or Rad21 (in mAST_floxed mNSC_floxed or Tcell_WT samples), we calculated how many CTCF or Cohesin subunit peaks are present in FIREs. We also permuted FIRE positioning 10 times, and asked the same question to obtain a distribution of expected values. To determine statistical significance, we compared this observed value to the expected distribution using a one-sample t-test.

FIRE and disease-associated SNP analyses

We collected the 4,378 non-coding disease associated GWAS SNPs (referred to hereafter as "SNPs") used in a previous study (Hnisz et al., 2013), and converted each SNP ID to its genomic location in hg19 human reference genome, using NCBI dbSNP online tool (http://www.ncbi.nlm.nih.gov/projects/SNP/dbSNP.cgi?list=rslist), resulting in 4,327 SNPs. Next, we mapped each SNP to FIRE bins identified from each of 7 cell lines and 14 tissues, and calculated the SNP density, defined as the number of mapped SNPs per 1Mb of FIRE bins. We further divided FIRE bins based on their overlap with typical enhancers and super-enhancers, and calculated the SNP density within each sub FIRE groups. Additionally, we performed disease-based FIRE SNP overlap analysis. For each of 456 diseases, we defined the enrichment score as the ratio between the proportion of SNPs overlapped with FIRE bins and the proportion of FIRE bins in the genome. Higher enrichment score indicates stronger overlap between SNPs and FIRE bins.

Calling Significant Interaction Pairs in Hi-C data

Statistically significant contacts in Hi-C data were identified using Fit-Hi-C, as previously described (Ay et al., 2014). First, Fit-Hi-C assumes that the expected contact frequency is a function of genomic distance. Fit-Hi-C also assumes the observed contact counts follow a Poisson model for non-peak Hi-C bin-pairs, (i.e. $O_{ij} \sim Poisson(\lambda(d_{ij}))$), and assumes an observed contact count is significantly higher than this Poisson variable for a peak bin-pairs (i.e. a statistically significant Hi-C contact). Fit-Hi-C conducts fitting and removing outliers iteratively. Fit-Hi-C requires the user to specify the range of genomic distance to assess for statistical significance. Based on this genomic distance input and for each iteration, Fit-Hi-C first bins the specific genomic distance into B bins (by default B=100), then estimates the mean observed contact count of currently labeled non-peak bin-pairs from each bin and then fits a spline curve $\lambda(d_{ij})$ based on average observed count at each distance determined by B and the user-input distance cutoff. For example, if one were to input B=50 and 2Mb genomic distance, then the spline curve will fit the mean contact count across 50 distance data points. Then, Fit-Hi-C tests each observed count O_{ij} against the calibrated Poisson distribution $Poisson(\lambda(d_{ij}))$. Fit-Hi-C rejects the null hypothesis when p value is small and labels this observation as a significant bin-pair "peak" (a significant Hi-C contact). In the next iteration, Fit-Hi-C conducts the same processes of calibrating the background distribution and significance testing. After converting our Hi-C contact matrix into the correct input format for Fit-Hi-C, we used the default Fit-Hi-C code to

calculate a p value and q value (a false discovery rate, FDR) for each bin-pair within 2Mb genomic distance. The generic example code for Fit-Hi-C can be found here: (https://noble.gs.washington.edu/proj/Fit-Hi-C/). For all analyses in this study (except where noted) we used a conservative peak-calling threshold of FDR<1e-6. This is based on the observation that more relaxed peak calls (FDR<0.05, the Fit-Hi-C default parameter) seemed to overcall peaks, and, FDR<1e-6 corresponds to ~1 million total peaks in IMR90, very similar to previous reports (Jin et al., 2013).

eQTL Enrichment Analyses

Statistically significant SNP-gene pairs were downloaded from the GTex Portal (http://www.gtexportal.org/home/), using Version 6 (filed called GTEx_Analysis_V6_eQTLs.zip). Since only a subset of our tissue types can be found in the GTex dataset, we extracted 6 GTex datasets corresponding to 6 of our higher depth tissue Hi-C datasets. The following files were used from the GTex datasets: Adrenal_Gland_Analysis.snpgenes, Liver_Analysis.snpgenes, Brain_Frontal_Cortex_BA9_Analysis.snpgenes, Artery_Aorta_Analysis.snpgenes, Heart_Left_Ventricle_Analysis.snpgenes, Heart_Left_Ventricle_Analysis.snpgenes.

To evaluate whether statistically significant contacts emanating from FIRE bins are enriched for SNP-gene pairs, and also to address whether the most significant Hi-C peaks are further enriched for SNP-gene pairs compared to less significant Hi-C peaks, we first used Fit-Hi-C to generate q values (i.e. FDRs) for all bin-pairs within 2Mb genomic distance for each tissue type and sub-selected higher depth tissue datasets in which we also obtained GTex information (i.e. 6 tissues listed above). For the analysis of each sample, we first ranked significant bin-pairs by their FDR, from most significant pairwise contact to contacts with FDR approaching 0.05 (default Fit-Hi-C significance cutoff). This generates a genome-wide ranked list of significant pairwise contacts. We then divided significant bin-pairs into two groups depending on whether the anchor bin is a FIRE bin or non-FIRE bin, creating two groups termed “FIRE bin peaks” and “non-FIRE bin peaks”. In order to evaluate whether there is a difference in the presence of known SNP-gene pairs emanating from FIRE bins compared to non-FIRE bins, we selected the top 1K-20K significant FIRE peaks at 1K step size. As a control, we randomly selected a size-matched statistically significant bin-pairs emanating from non-FIRE bins. To evaluate whether FIRE peaks contained more SNP-gene pairs than non-FIRE bin peaks, we tested whether the average number of SNP-gene pairs captured by the top set of FIRE peaks is significantly higher than the size-matched control set (from non-FIRE bin peaks), using a one-side two-sample t test. Due to the random nature of selecting the size-matched control set, we generated 10 control datasets for each comparison (i.e. 1k, 2k…20k). To assess if the most significant FIRE bin peaks are more enriched for SNP-gene pairs than less significant FIRE bin peaks, we have plotted the log2(O/E) values for the top 1k,2k,3,4k,5k,10k,15k FDR groups (related to Figure g-j). Using a p value here is not entirely appropriate to address this analysis since p values for two-sample t tests are sensitive to sample size.

FIRE peak analyses

To evaluate whether FIREs have more local peaks than non-FIREs, we used Fit-Hi-C peak-calling results at stringent statistical significance (FDR<1e-6) to obtain distributions of the number of peaks emanating from FIRE bins or size-matched randomly permuted non-FIRE bins. To determine if the observed number of peaks from FIREs is greater than non-FIREs, we used a two-sample t-test.

To determine if FIREs self-interact at higher frequency than FIREs with non-FIREs or non-FIREs with non-FIREs, we first collected all FIRE bins, and then for each distance (d) from 40kb to 2Mb, we calculated the mean interaction frequency in which a FIRE bin was contacting another FIRE bin. Therefore, for each distance increment, we obtain a mean FIRE-FIRE interaction frequency. We then repeated the same procedure, but this time calculating the interaction frequency of FIREs with non-FIREs at each distance increment. Lastly, we randomly permuted FIRE bin locations to obtain a set of random non-FIRE bins and then calculated the interaction frequency with other non-FIRE bins for each distance increment. Then, for each genomic distance increment, we compared the FIRE-FIRE frequency with either the FIRE-nonFIRE or nonFIRE-nonFIRE using a two-sample t-test (related to Figure 7e; Table S4). This process was done independently for each sample.

To evaluate if FIREs are often the significant contact target of other FIREs we first collected all significant (FDR<1e-6) FIRE target bins determined by Fit-Hi-C, as well as all FIRE bins. We then intersected the FIRE target bins and FIRE bin annotations, creating three groups: FIRE targets that are non FIREs, FIRE targets that are FIREs,

and FIRE bins that are not targets of other FIREs (related to Figure 7g, h). The statistical significance of whether a FIRE bin is more likely a target of another FIRE bin was evaluated using a chi-square test.

Supplemental References:

Ay, F., Bailey, T.L., and Noble, W.S. (2014). Statistical confidence estimation for Hi-C data reveals regulatory chromatin contacts. Genome Res 24, 999–1011.

Chiaromonte, F., Yap, V., and Miller, W. (2001). Scoring pairwise genomic sequence alignments. Pacific Symp.

Dixon, J.R., Selvaraj, S., Yue, F., Kim, A., Li, Y., Shen, Y., Hu, M., Liu, J.S., and Ren, B. (2012). Topological domains in mammalian genomes identified by analysis of chromatin interactions. Nature 485, 376–380.

Dixon, J.R., Jung, I., Selvaraj, S., Shen, Y., Antosiewicz-bourget, J.E., Lee, A.Y., Ye, Z., Kim, A., Rajagopal, N., Xie, W., et al. (2015). Chromatin architecture reorganization during stem cell differentiation. Nature 518, 331–336.

Dowen, J.M., Fan, Z.P., Hnisz, D., Ren, G., Abraham, B.J., Zhang, L.N., Weintraub, A.S., Schuijers, J., Lee, T.I., Zhao, K., et al. (2014). Control of cell identity genes occurs in insulated neighborhoods in mammalian chromosomes. Cell 159, 374–387.

Hnisz, D., Abraham, B.J., Lee, T.I., Lau, A., Saint-André, V., Sigova, A.A., Hoke, H.A., and Young, R.A. (2013). Super-enhancers in the control of cell identity and disease. Cell 155, 934–947.

Hu, M., Deng, K., Selvaraj, S., Qin, Z., Ren, B., and Liu, J.S. (2012). HiCNorm: removing biases in Hi-C data via Poisson regression. Bioinformatics 28, 3131–3133.

Imakaev, M., Fudenberg, G., McCord, R.P., Naumova, N., Goloborodko, A., Lajoie, B.R., Dekker, J., and Mirny, L.A. (2012). Iterative correction of Hi-C data reveals hallmarks of chromosome organization. Nat. Methods 9, 999–1003.

Ji, X., Dadon, D.B., Powell, B.E., Fan, Z.P., Borges-Rivera, D., Shachar, S., Weintraub, A.S., Hnisz, D., Pegoraro, G., Lee, T.I., et al. (2016). 3D Chromosome Regulatory Landscape of Human Pluripotent Cells. Cell Stem Cell 18, 262–275.

Jin, F., Li, Y., Dixon, J.R., Selvaraj, S., Ye, Z., Lee, A.Y., Yen, C.A., Schmitt, A.D., Espinoza, C.A., and Ren, B. (2013). A high-resolution map of the three-dimensional chromatin interactome in human cells. Nature 503, 290–294.

Kent, W.J., Baertsch, R., Hinrichs, A., Miller, W., and Haussler, D. (2003). Evolution's cauldron: duplication, deletion, and rearrangement in the mouse and human genomes. Proc. Natl. Acad. Sci. U. S. A. 100, 11484–11489.

Lieberman-Aiden, E., van Berkum, N.L., Williams, L., Imakaev, M., Ragoczy, T., Telling, A., Amit, I., Lajoie, B.R., Sabo, P.J., Dorschner, M.O., et al. (2009). Comprehensive mapping of long-range interactions reveals folding principles of the human genome. Science (80-.). 326, 289–293.

McLean, C.Y., Bristor, D., Hiller, M., Clarke, S.L., Schaar, B.T., Lowe, C.B., Wenger, A.M., and Bejerano, G. (2010). GREAT improves functional interpretation of cis-regulatory regions. Nat. Biotechnol. 28, 495–501.

Rao, S.S.P., Huntley, M.H., Durand, N.C., and Stamenova, E.K. (2014). A 3D Map of the Human Genome at Kilobase Resolution Reveals Principles of Chromatin Looping. Cell 1–16.

Ritchie, M.E., Phipson, B., Wu, D., Hu, Y., Law, C.W., Shi, W., and Smyth, G.K. (2015). limma powers differential expression analyses for RNA-sequencing and microarray studies. Nucleic Acids Res. 43, e47.

Saldanha, A.J. (2004). Java Treeview--extensible visualization of microarray data. Bioinformatics 20, 3246–3248.

Schwartz, S., Kent, W.J., Smit, A., Zhang, Z., Baertsch, R., Hardison, R.C., Haussler, D., and Miller, W. (2003). Human-mouse alignments with BLASTZ. Genome Res 13, 103–107.

The Roadmap Epigenomics Consortium (2015). Integrative analysis of 111 reference human epigenomes.