

Transcriptional immune suppression and upregulation of double stranded DNA damage and repair repertoires in ecDNA-containing tumors

Reviewed Preprint

Published from the original preprint after peer review and assessment by eLife.

[About eLife's process](#)

Reviewed preprint posted

August 10, 2023 (this version)

Sent for peer review

June 15, 2023

Posted to bioRxiv

April 24, 2023

Miin S. Lin, Se-Young Jo, Jens Luebeck, Howard Y. Chang, Sihon Wu, Paul S. Mischel , Vineet Bafna 

Bioinformatics and Systems Biology Graduate Program, University of California at San Diego, La Jolla, CA, USA • Department of Biomedical Systems Informatics and Brain Korea 21 PLUS Project for Medical Science, Yonsei University College of Medicine, Seoul, South Korea • Department of Computer Science and Engineering, University of California at San Diego, La Jolla, CA, USA • Center for Personal Dynamic Regulomes, Stanford University, Stanford, CA, USA • Department of Genetics, Stanford University, Stanford, CA, USA • Howard Hughes Medical Institute, Stanford University, Stanford, CA, USA • Children's Medical Center Research Institute, University of Texas Southwestern Medical Center, Dallas, TX, USA • Sarafan Chemistry, Engineering, and Medicine for Human Health (Sarafan ChEM-H), Stanford University, Stanford, CA, USA • Department of Pathology, Stanford University School of Medicine, Stanford, CA, USA • Halicioğlu Data Science Institute, University of California at San Diego, La Jolla, CA, USA

 https://en.wikipedia.org/wiki/Open_access

 <https://creativecommons.org/licenses/by/4.0/>

Abstract

Extrachromosomal DNA is a common cause of oncogene amplification in cancer. The non-chromosomal inheritance of ecDNA enables tumors to rapidly evolve, contributing to treatment resistance and poor outcome for patients. The transcriptional context in which ecDNAs arise and progress, including chromosomally-driven transcription, is incompletely understood. We examined gene expression patterns of 870 tumors of varied histological types, to identify transcriptional correlates of ecDNA. Here we show that ecDNA containing tumors impact four major biological processes. Specifically, ecDNA containing tumors upregulate DNA damage and repair, cell cycle control, and mitotic processes, but downregulate global immune regulation pathways. Taken together, these results suggest profound alterations in gene regulation in ecDNA containing tumors, shedding light on molecular processes that give rise to their development and progression.

eLife assessment

This study of extrachromosomal DNA (ecDNA) aims to identify genes that distinguish ecDNA+ and ecDNA- tumors. This timely study is **important** in addressing the genes responding to the amplification of the ecDNA. The data presented are for the most part **solid**, there were concerns regarding the clarity in the description of the analysis methods and whether the evidence for specific genes required to maintain the ecDNA+ state was entirely conclusive.

Introduction

Extrachromosomal DNA (ecDNA) are large, functional, circular double-stranded DNA molecules that are enriched for oncogenes, highly amplified, and frequently observed in a wide variety of cancer types^{1,2}. ecDNAs lack centromeres and are asymmetrically segregated into daughter cells during cell division, driving intratumoral genetic heterogeneity, accelerated evolution, and rapid treatment resistance^{3,4}. Further, recent studies demonstrate strong positive selection for ecDNA during tumor progression⁵. ecDNAs also exhibit highly accessible chromatin and altered cis- and trans-regulation, including cooperative intramolecular interactions⁶, promoting elevated expression of oncogenic transcriptional programs^{7–9}, further contributing to poor outcome for patients².

The recent development of computational tools that enable detection of ecDNA from whole genome sequencing data, has facilitated analyses of well-curated, publicly available datasets, including The Cancer Genome Atlas (TCGA), thereby providing an important opportunity to identify transcriptional repertoires that are preferentially detected in bona fide, clinical ecDNA-containing tumors. To shed new light on the gene expression patterns that may enhance ecDNA development and progression, we examined global transcriptional analysis of ecDNA-containing tumors.

Results

A recent analysis utilized the tools AmpliconArchitect and AmpliconClassifier on 1,921 tumors from The Cancer Genome Atlas (TCGA) to suggest that ecDNA prevalence ranges from 0% to 59.6% across multiple tumor tissue subtypes². Using AmpliconClassifier (AC), the analysis classified tumor samples into five subtypes: ecDNA(+), Breakage Fusion Bridge (BFB), complex non-cyclic, linear, and no-amplification. However, due to limitations imposed by short-read sequencing, AC may classify some ecDNA(+) structures as complex non-cyclic when breakpoints are missed. Secondly, BFB cycles can give rise to ecDNA formation, making discernment of the two modes of amplification difficult. To limit false-negative ecDNA classifications in the ecDNA(-) set, we treated samples with only a linear or no-amplification status as ecDNA(-), removing complex non-cyclic and BFB(+) samples from the analysis. In order to understand the transcriptional programs active in maintaining ecDNA, we selected 870 samples from 14 tumor types with at least three ecDNA(+) samples each, and compared the gene expression data of the resulting 234 ecDNA(+) and 636 ecDNA(-) samples (Table S1).

Machine learning identifies candidate genes for ecDNA maintenance

In lieu of identifying genes that are highly differentially expressed between ecDNA(+) and ecDNA(-) samples but driven by a small subset of cases (e.g. gene A in Figure S1a), we sought to identify genes (e.g. gene B) whose expression level was predictive of ecDNA presence. We assumed that genes that were persistently over-expressed or under-expressed in ecDNA(+) samples relative to ecDNA(-) samples were more likely to be involved in ecDNA biogenesis or maintenance, or in mediating the cellular response to the presence of ecDNA. For exposition purposes, we refer to these functions as “maintenance” activities below.

To identify a minimal set of genes whose expression values were consistently predictive of ecDNA presence, we used Boruta,¹⁰ an automated feature selection algorithm (Fig. 1a and Methods). Given the unequal representation of ecDNA(+) and ecDNA(-) samples within each of the 14 tumor

types, we performed Boruta on 200 datasets, each consisting of a random selection of 80% of the 870 samples (**Fig. 1a**), and chose the criterion of a gene being labeled as a Boruta gene in at least 10 of the 200 runs to be selected for downstream analysis. The Boruta analysis identified a set of 408 genes with persistent differential expression, hereafter denoted as the Core gene set.

Extending the Core set with co-expressed genes

We note that the Core gene set is not a comprehensive list of discriminatory genes, using a toy example. Consider gene “B”, a member of the core gene set, and another gene, “C”, whose expression values across all samples are nearly identical to the expression values of core gene B. The Boruta analysis would not need to assign gene C to the core set in addition to gene B, because adding both genes incurs the same predictive power as adding one. However, either, or both genes may play an important role in ecDNA maintenance. To correct this, we ran `pvcust` to cluster all gene expression values, and to identify stable clusters using multiscale bootstrap resampling (**Fig. 1b**; Methods). We used an approximately unbiased (AU) confidence value of 0.95 to select the most highly co-expressed gene clusters. An AU confidence value of 0.95 represents the rejection of the null hypothesis that a group of genes fail to form a stable cluster at a significance level of 0.05. Recomputing the number of Boruta runs that members of a cluster were selected in, we selected clusters that appeared in at least 10 of the 200 Boruta runs (Methods). This resulted in the selection of 354 recurring clusters (**Table S2**).

Notably, among the 354 clusters, only 2 clusters (with 14 total genes) did not contain any Core genes. As most genes do not have completely identical expression patterns, we would expect one gene to be consistently picked as a Boruta gene over another co-expressed gene. Consistent with this hypothesis, most (344/354) clusters contained only 1 or 2 Core genes (**Fig. 1c**). When selecting clusters that contained at least 1 Core and 1 co-expressed gene, 53 of 71 clusters contained 1 to 3 Core genes (**Fig. S1b**), confirming that a few genes per co-expressed cluster provide sufficient predictive value, but other co-expressed genes might still play an important role in ecDNA maintenance. This is true for clusters of various sizes, including the 2-member cluster #74 and the 21-member cluster #3. In cluster #74, *CSTF1* had similar expression values to the Core gene *RAE1*, which is a mitotic checkpoint regulator implicated in tumor progression¹² (**Fig. S1c**; **Table S3**). While not necessarily increasing the predictive value, *CSTF1* is also a proto-oncogene involved in aberrant alternative splicing events¹³. In cluster #3, 12 genes were highly co-expressed with 9 Core genes (**Fig. S1d**; **Table S3**), and were enriched in cell-cycle related biological processes (Methods). Importantly, the total number of genes per cluster was also small (**Fig. 1d**), with only 7 of 354 clusters carrying more than 10 genes. This suggests that genes involved in ecDNA maintenance have specific roles that cannot be accomplished by multiple other genes.

Summarizing, the 354 clusters contained 643 genes, which included 408 Core genes and 235 additional genes (**Fig. 1b**). Together, we define these genes as the CorEx (Core+co-expressed) genes (**Table 1**). The remaining manuscript investigates the functional properties of these genes.

CorEx gene expression correlates strongly with ecDNA occurrence

We validated the relevance of CorEx genes in ecDNA maintenance by running cross-validation experiments (**Fig. 2a**; Methods) to test the predictive power of CorEx gene expression in determining the ecDNA status of the sample. For exposition, only the maximum recall was plotted for experiments with similar precision (**Fig. 2a,b**), with all data points shown in **Fig. S2a,b**. Expectedly, the predictive performance did not change when switching between Core genes and CorEx genes, because each of the non-core gene in the CorEx list had an expression pattern similar to at least one Core gene (**Fig. 2b**, **S2a**). While the precision and recall numbers were not high enough for the gene expression to be used directly in predicting the ecDNA status, they were significantly better than a random subset of genes. For precision values at least 0.7, the maximum




Figure 1.

Genes predictive of ecDNA status

(a) The feature selection algorithm, Boruta, was applied to 200 datasets of randomly selected subsets consisting of 80% of all samples. Genes selected by Boruta in at least 10 of the 200 runs were identified as the Core set of genes (408) involved in ecDNA maintenance. (b) Identification of highly co-expressed and stable gene clusters using pvclust expanded the Core set by an additional 235 genes to the final list of 643 CorEx genes. (c) Out of 354 clusters, the majority (344) of clusters contained 1 or 2 Core genes. (d) Most clusters were small, with only 7 clusters containing more than 10 genes.

(median) recall using CorEx genes was 0.65 (0.48) (**Fig. 2b**, **S2b**). In contrast, the maximum (median) recall for randomly selected genes at the same precision level was 0.37 (0.24) (p -value 1.8×10^{-35} ; Mann-Whitney U test).

To test the persistence of CorEx genes across tumor types, we re-computed Cliff's delta values¹⁴ for each of the 11 TCGA tumor types that had at least 10 ecDNA(+) and at least 10 ecDNA(-) samples. The directionality of gene expression patterns was significantly similar to TCGA in each tissue type, with one exception (**Fig. 2c**, **Table S4**). The sole exception was the tumor type of Sarcoma (SARC). It is notable that the TCGA-SARC samples included many liposarcomas. In addition to containing ecDNA, liposarcoma samples are known to have extensively rearranged structures indicative of chromothripsis and neo-chromosome formation¹⁶. For other tissue types, the p -values against a null hypothesis of no match to the pan-cancer prediction ranged from 4.2×10^{-12} to 3.5×10^{-85} (Fisher's exact test) for the significant associations (Methods). The results were similar if we tested using only Core genes (**Fig. S2c**). Summarizing, the 643 CorEx genes are differentially expressed across a multitude of tumor types, and have consistently higher or lower expression in ecDNA(+) samples relative to ecDNA(-) samples. These results are consistent with a pan-cancer role of CorEx genes in ecDNA maintenance.

CorEx genes are very different from the most differentially expressed genes

We asked if the CorEx gene list was significantly different from a list of genes with the highest change in expression between the two classes. We performed a differential expression analysis using DESeq2¹⁷ and picked the 643 most significantly differentially expressed genes in terms of the absolute value of their shrunken log-fold change estimate (LFC; Methods). Using the sign of the LFC value as the determinant for directionality, 240 of these genes were up-regulated, while 403 were down-regulated. Notably, only 86 of these top-|LFC| genes overlapped with the CorEx gene set (**Table S5**; **Fig. S2d**). As discussed above, these highly differentially expressed genes could be driven by a small subset of samples, and therefore, were not as predictive of ecDNA status as the CorEx genes. For precision values of at least 0.7, the maximum recall using top-|LFC| genes was 0.52 (median: 0.39), significantly less than the recall of 0.65 (median: 0.48) using CorEx genes (Mann-Whitney U, p -value 4.8×10^{-21} ; **Fig. 2b**, **S2b**).

The top-|LFC| genes were also different from the CorEx genes by other metrics. Not surprisingly, the log-fold change (LFC) values of the top-|LFC| genes were higher than the LFC values of the CorEx genes (**Fig. 2d**, MWU p -value 1.83×10^{-158}). However, much of the LFC change was due to the very low expression of the top-|LFC| genes in either ecDNA(+), or ecDNA(-) samples. In fact, the CorEx genes had higher expression in both ecDNA(+) and ecDNA(-) samples compared to the Differentially Expressed (DE) genes (**Fig. 2e**, MWU p -value $< 2 \times 10^{-308}$). While the absolute log fold-change in expression of CorEx genes between ecDNA(+) and ecDNA(-) samples was not that high (median: 0.30, mean: 0.41), it was persistent across all samples (variance: 0.14, standard deviation: 0.37).

For example, the genes *ITLN1* and *PNMT* had the second and eighth-highest absolute LFC values of 3.92 and 2.89 in the top-|LFC| list. However, their normalized expression values in most ecDNA(+) samples were low. *ITLN1* had a normalized RSEM expression value ≤ 8 (21st percentile) in 210/234 ecDNA(+) samples. Similarly, the normalized RSEM expression value of *PNMT* in 223/234 ecDNA(+) samples was less than 8.5 (rank percentile: 41.1%). For *PNMT*, the differential expression was mediated by 11 ecDNA(+) samples having an expression value ≥ 11 , and 5 of the 11 samples contained *PNMT* on an ecDNA amplicon (**Fig. S2e**). Similarly, 3 samples with high RSEM contained *ITLN1* on an amplicon (**Fig. S2f**), partly accounting for the high |LFC| value. In contrast, the CorEx gene, *RAE1*, had a high normalized expression value in both ecDNA(+) and ecDNA(-) samples (average 9.72, rank percentile 74.3%), with a small but persistent LFC value of 0.33.




Figure 2.

Validation of CorEx genes

(a,b) Cross-validation experiments validating the predictive value of CorEx genes. Precision denotes the fraction of predicted samples that were truly ecDNA(+). Recall refers to the fraction of ecDNA(+) samples that were predicted correctly. For multiple points with a similar precision, the maximum recall is plotted. **(a)** The curves for CorEx and Core genes overlap, suggesting similar predictive power. **(b)** CorEx genes have higher predictive rates compared to 643 randomly selected genes and the top 643 differentially expressed genes based on logarithmic fold changes from a DESeq2 analysis (Top-|LFC| genes). **(c)** CorEx genes were consistently up- or down-regulated in ecDNA(+) samples across tumor types, with the exception of SARC. **(d)** Of the 643 Top-|LFC| genes, 240 were up-regulated while 403 were down-regulated in ecDNA(+) samples. Of the CorEx genes, 325 were up-regulated while 318 were down-regulated. The absolute LFC values of the Top-|LFC| gene set was significantly greater than that of the CorEx genes (p -value $1.83e-158$). **(e)** The normalized gene expression values of the CorEx genes were significantly higher than that of the Top-|LFC| gene set (p -value $< 2e-308$). *** p -value < 0.001 .

The results confirm our intuition that differential expression can arise due to multiple reasons, including low expression of the gene in a majority of samples, or the copy number amplification of a gene in a few samples. In contrast, the CorEx genes were selected based on persistent over- or under-expression in ecDNA(+) samples.

CorEx genes primarily up-regulate three biological processes: Cell Cycle, Cell division, and DNA Damage Response

To identify enriched biological processes specific to either up-regulated or down-regulated genes in ecDNA(+) samples, we combined two metrics of effect size, Cliff's delta¹⁴, ¹⁵, and log fold change¹⁷ to determine the directionality of CorEx genes (**Table S6**; Methods). The two effect size metrics were mostly in agreement in terms of directionality. Of the 7,288 genes that passed the negligible effect size thresholds in both metrics, only 14 were not in concordance. This more stringent approach, in comparison to a simple directionality based on the sign of a single effect size value, was applied given that gene set enrichment is dependent on not only the number of up- or down-regulated genes but also the degree of overlap with genes under a specific biological process term (Methods). Using this approach, among the 643 CorEx genes, 262 genes were found to be up-regulated in ecDNA(+), while 271 were found to be down-regulated (**Table 1**; Methods). 110 genes did not make the effect size cut-off. The numbers were similar for the 408 Core genes, with 190 up-regulated, 196 down-regulated, and 22 genes not making the cut-off.

We performed a gene set enrichment analysis to identify the Gene Ontology (GO) biological processes that are enriched in CorEx genes (Methods). Briefly, we applied a one-sided Fisher's exact test using gene sets from MSigDB^{18–20}, using a false discovery rate of 5% (Benjamini-Hochberg procedure). The UP-regulated genes were enriched for 187 Biological processes (**Table S7**). Note that the GO-biological process (BP) terms are not independent, because of their hierarchical organization, and sharing of genes across different GO terms. Therefore, we used an approach similar to that used in DAVID²¹ to cluster the biological processes enriched by the UP-regulated genes into 11 broad categories (**Table S8**; **Fig S3**; Methods). The 11 categories were assigned a name using manual inspection of the constituent GO terms, or called "Other." The 11 categories (including "Other") are shown in a waterfall plot to explain the contribution of each gene to a category (**Fig. 3a**).

The 10 categories included expected participation of biological processes involved in (a) cell-cycle regulation (Mitotic/Meiotic Cell Cycle, G1/S, G2/M) (b) cell-division (Spindle Organization, Cell Division, Chromosome Condensation, Chromosome Segregation), (c) DNA Damage response (DNA Repair), and (d) the HOX Gene cluster. Indeed, one of the largest clusters, cluster #3, containing 9 Core genes and 12 co-expressed genes, was enriched in GO-BP terms related to the cell cycle (**Table S9**). Notably, the enriched categories also included a role for the *HOX* genes with 17 members of the *HOX* family up-regulated in ecDNA(+) cancers (**Table S8**). Many recent reports have associated *HOX* genes with cancer, including an association with their phenotypic "hallmarks"²². Genes involved in angiogenesis (*HOXA2*, *HOXC5*), genome instability (*HOXC5*, *HOXC11*), deregulating cellular energetics (*HOXA4*, *HOXC5*), and metastasis (*HOXA2*) were all up-regulated in ecDNA(+) cancers.

While Meiotic cell-cycle was also enriched, only 7 genes were allocated specifically to the group of enriched terms: *SEPP1*, *SNRPA1*, *TMEM203*, *RNF114*, *TAF4*, *TNFAIP6*, and *PTX3* (**Table S10**). Though these genes have roles in the meiotic cell cycle, they have also been implicated in cancer and other inflammatory diseases. The spliceosomal protein *SNRPA1* is a pro-metastatic splicing enhancer²³. *TNFAIP6*, together with *PTX3*, activates the Wnt/ β -catenin pathway to promote gastric carcinoma cell invasion²⁴. *TMEM203* is a STING-centered signaling regulator implicated in inflammatory diseases²⁵. *RNF114* is a zinc-binding protein whose over-expression is an indicator of epithelial inflammation and implicated in various tumors²⁶. *TAF4*, a transcription initiation factor, when overexpressed, is implicated in ovarian cancer by playing a role in




Figure 3.

Up-regulated CorEx genes

(a) GO biological processes enriched in up-regulated genes were clustered into 11 broad categories. **(b)** Genes up- or down-regulated in processes involved in major double-strand break (DSB) damage repair pathways. Many critical genes in the c-NHEJ pathway were down-regulated in ecDNA(+) samples relative to ecDNA(-) samples.

dedifferentiation that promotes metastasis and chemoresistance²⁷. Finally, in looking at the genes in the “Other” category, *SHCBP1* was the only gene unique to it. A member of the neural precursor cell proliferation process, *SHCBP1* is reported to promote tumor cell signaling and proliferation²⁸.

Taken together, the 11 biological process categories explain 165 of the 262 up-regulated CorEx genes, and suggest that Mitotic Cell-Division, Cell cycle regulation, and DNA Damage response are the three broad categories of biological processes up-regulated in ecDNA maintenance, along with an up-regulation of genes in the *HOX* cluster.

CorEx genes upregulate specific Double-strand break repair pathways

The 16 enriched DNA damage response GO terms contained terms “double-strand break repair” and “recombination”, but none contained the terms “single-strand”, “nucleotide-excision”, or “mismatch-repair” (**Table S8**), suggesting that the CorEx genes are largely composed of genes involved in multiple double-strand break (DSB) repair pathways, which include classical non-homologous end-joining (c-NHEJ), Alternative end-joining (Alt-EJ), single-strand annealing (SSA), or homology directed repair (HDR)²⁹.

The choice of these varied DSB repair mechanisms for ecDNA maintenance is not well understood. We compiled and hand-curated a list of 129 genes involved in DSB repair and marked them for their role in one or more of these 4 pathways (**Table S11**). Of these genes, a high number (67) were up-regulated in ecDNA(+), while a smaller number (15) were down-regulated, relative to ecDNA(-) samples. This breakdown of 129 DDR genes contrasts with an analysis using all genes where a near identical number of genes (5,256, and 5,251) were up- and down-regulated in ecDNA(+) samples, confirming that DDR genes are significantly up-regulated relative to all differentially expressed genes (p -value < 0.0001; Fisher’s exact test). When broken down to the roles of genes in individual DSB repair pathways, we found that Alt-EJ with 11 up-regulated and 1 down-regulated genes (p -value 0.0063), SSA (11 up, 1 down (p -value 0.0063)), and HR (46 up, 8 down; p -value < 0.00001) were all up-regulated. However, classical NHEJ (14 up, 7 down; p -value: 0.19) was not significantly up-regulated in ecDNA(+) samples relative to ecDNA(-) samples (Methods, **Table S12**, **Table S13**).

The expression of key genes in these pathways raises the possibility of an increased role of non-classical-NHEJ processes in ecDNA development or progression, relative to c-NHEJ (**Fig. 3b**; **Table S11**). A number of genes involved in c-NHEJ were downregulated in ecDNA-containing tumors relative to non-ecDNA tumors. These included *XLF/NHEJ1* (MWU p -value 2.05e-03), which is a key member of the ligase complex required for c-NHEJ; *LIG4*, another member of the ligase complex (MWU p -value 0.03), *PNKP*, which generates 5-phosphate/3-hydroxyl DNA termini required for ligation (MWU p -value 3.90e-06); and also, DNA polymerases λ (*POLL*; MWU p -value 1.09e-21) and μ (*POLM*; MWU p -value 0.01), which promotes the ligation of terminally compatible overhangs requiring fill-in synthesis and promotes the ligation of incompatible 3’ overhangs³⁰ in a template independent manner, respectively. This does not imply a defect in these repair processes, but rather, potentially additional or preferential utilization of alternative DSB repair pathways in ecDNA-containing tumors.

TP53BP1 is key to blocking resection and promoting the c-NHEJ pathway choice, but is displaced by BRCA1 and the MRN complex to initiate resection in the broken strands^{31, 32}. *BRCA1* was significantly up-regulated in ecDNA(+) samples (**Fig. 3b**), while *TP53BP1* was significantly down-regulated (MWU p -value 0.016), although with negligible effect size (**Table S11**). Supporting the role of alternative pathway choice for DDR, key genes in the Alt-EJ pathway, including *PARP-1*, DNA polymerase θ (*POLQ*), *LIG1*, *LIG3*, *FEN1* were all significantly up-regulated in ecDNA(+) samples. Homology directed repair is the preferred pathway when a sister chromatid is available

to act as a template. HDR is initiated by additional and extensive resection. The genes *BLM*, *EXO1*, *RPA1*, *RPA3* which promote additional resection, as well as *BRCA1*, *BRCA2*, *RAD51*, and others that support HDR were all found to be significantly up-regulated. We can conclude that the specific pathway choice for DSB repair in ecDNA(+) samples is dominated by alt-EJ and homology directed repair pathways, while c-NHEJ is not a preferred choice.

CorEx genes primarily down-regulate immune system processes

Using methodology similar to the analysis of the up-regulated genes, the down-regulated genes enriched 73 GO terms (Table S14), and could be clustered into seven broad categories, including “Other” (Fig. 4a; Table S15; Fig. S4). Surprisingly, all categories were immunomodulatory. The most enriched broad category contained 75 CorEx genes relating to the Lymphocyte activation pathway. It included genes enriching “T-cell activation” (28 CorEx genes; p -value 2.58e-05), and “Positive regulation of cell-cell adhesion” (16 CorEx genes; p -value 4.49e-03). Other down-regulated pathways included Cytokine activation, especially for genes in the IL-12 pathway (6 CorEx genes, p -value 7.17e-03), TNF super-family (12 CorEx genes, p -value 7.24e-03), and Inflammation, including, for example, down-regulation of Toll-like receptor 2 signaling (4 CorEx genes, p -value 4.49e-03). Finally, the broad category of Leukocyte chemotaxis was also enriched among the down-regulated genes. The chemotaxis genes include many chemokines and their receptors involved in trafficking of T cells to the site of the tumor. The remaining down-regulated genes included four fucosyltransferases, and the category marked “Other.” *FUT2* silencing is associated with reduced adhesion and increased metastatic potential³³. Notably, the category marked “Other” was dominated by genes in NF- κ B pathway regulation (14 CorEx genes, p -value 2.28e-02).

NF- κ B signaling represents a prototypical, proinflammatory pathway³⁴ with multiple roles, including apoptosis. Specifically, 8 of the 14 down-regulated genes involved caspase activation (Table S16), representing the pro-apoptotic arm of NF- κ B signaling. A parallel pathway for sensing endogenous ligands secreted in cell death and cancer is mediated by Toll-like receptor (TLR) proteins³⁵. Remarkably, all ten TLRs were significantly down-regulated in ecDNA(+) tumors. They included TLRs expressed on the cell membrane that bind lipids and proteins as well as TLRs expressed on endosomal membranes that bind DNA. The CorEx down-regulated genes also included many involved in TLR signaling, such as *TLR3*, *CYBA*, *LYN*, and *TIRAP*.

Four of the seven broad categories mapped to facets of the cancer immune cycle³⁶ (Fig. 4b). We tested if CorEx genes in these categories were more likely to be down-regulated rather than up-regulated, when compared to the non-CorEx differentially expressed genes. The Inflammation category, which mapped to the “Cancer antigen presentation” facet, showed 18 up-regulated and 76 down-regulated CorEx genes (p -value 0.005, Fisher exact test, Table S15). Similarly, the CorEx genes related to the “Trafficking of T cells” facet (“Leukocyte migration and chemotaxis” category, 13 up, 49 down-regulated; p -value 0.03) and “Infiltration and recognition of tumor cells by cytotoxic T cells” facet (“Lymphocyte activation” category, 23 up, 75 down-regulated; p -value 0.0075) were also significantly down-regulated. However, down-regulation in the “Priming and activation” facet (“Cytokine production” category, 6 up, 28 down-regulated) was not significant at the 5% level.

As the RNA data were bulk-sequenced, transcripts were sampled from tumor cells and cells from the tumor microenvironment. Thorsson *et al.*³⁷ mined immune cell expression signatures to identify six immune subtypes: wound healing (C1), IFN- γ dominant (C2), inflammatory (C3), lymphocyte depleted (C4), immunologically quiet (C5), and TGF- β dominant (C6). A recent study analyzing the tumor microenvironment (TME) of ecDNA(+) vs. ecDNA(-) samples in seven tumor subtypes revealed an association of ecDNA presence with immune evasion³⁸. Our results (Fig. S5), which used an updated version of the classification method for these ecDNA(+) samples, were broadly consistent with those from the Wu *et al.*³⁸ study. Our results suggested an increase in C1 and C2 subtypes and a depletion of C3 and C6 between ecDNA(+) and ecDNA(-) categories (p -value 3.96e-03, Chi-squared test). Notably, the C3 (inflammatory) subtype is associated with lower levels




Figure 4.

Down-regulated CorEx genes

(a) GO biological processes enriched in down-regulated genes were clustered into 7 broad categories. (b) Four of these categories map to steps in the cancer-immunity cycle. CorEx genes in three of the four categories were significantly down-regulated compared to all genes (Fisher's exact test).

of somatic copy number alterations, and C6 with high lymphocyte infiltration, while C1 is associated with elevated levels of angiogenic genes. These are consistent with our findings of increased somatic copy numbers, increased expression of angiogenic genes on ecDNA(+) samples, and reduced lymphocyte infiltration.

ecDNA(+) samples carry a higher mutational burden relative to ecDNA(-) samples

In order to understand if the change in transcriptional program was driven by mutations to the genes, we checked if ecDNA(+) samples have differential levels of mutation relative to ecDNA(-). Intriguingly, we found that the total mutation burden was significantly higher in ecDNA(+) samples relative to ecDNA(-) samples (**Fig. 5a**). The result was significant also when mutations were limited to deleterious substitutions as measured by SIFT or PolyPhen2, and high-impact insertions and deletions (**Fig. S6a**). However, when controlling for cancer type, only glioblastoma (GBM; lower mutations in ecDNA(+)), low-grade gliomas (LGG; higher mutations in ecDNA(+)), and uterine corpus endometrial carcinoma (UCEC; lower mutations in ecDNA(+)) continued to show differential total mutational burden (**Fig. S6b**). Next, we tested if specific genes were differentially mutated between the two classes (**Fig. 5b**). For deleterious/high-impact mutations, *TP53* was the only gene whose mutational patterns were significantly higher in ecDNA(+) compared to ecDNA(-) (OR 2.67, Bonferroni adjusted *p*-value 4.22e-07). BRAF mutations, however, were more common in ecDNA(-) samples and were significant to an adjusted *p*-value < 0.1 (OR 0.27). The excess of *TP53* mutations in ecDNA(+) samples provides additional support to the hypothesis that mutations in DNA damage response or cell cycle checkpoints are important for ecDNA(+) formation and maintenance. Other genes that are differentially mutated with nominal significance (unadjusted *p*-value < 0.005) are shown in **Table S17**.

We also tested if a collection of gene mutations could predict ecDNA status using XGBoost³⁹, which uses an adaptive boosting of “weak classifiers” to predict class. Here, each mutated gene was treated as a weak classifier of ecDNA status. However, the two classes could not be separated with high accuracy (**Fig. S6c**). An unsupervised principal component analysis did not separate the two classes either. Only the first principal component explained a significant proportion (14%) of the total variance (**Fig. S6d**) and did not separate the bulk of the samples. Finally, we recapitulated earlier findings that ecDNA(+) samples enrich for APOBEC activity through the presence of the mutation signatures SBS2 and SBS13^{40, 41} (**Fig. S7**). The enrichment in *TP53* mutations was also consistent with previous findings⁵. On balance, however, collections of gene mutations did not distinguish ecDNA(+) samples from ecDNA(-) samples, at least at a pan-cancer level, in contrast to the gene expression data.

Persistently occurring genes in ecDNA(+) samples represent potential vulnerabilities

To rank CorEx genes by importance, we computed harmonic mean rank values based on three categories: a) the average GINI importance statistic from the trained random forest models; b) the number of Boruta runs that a gene was selected in; and c) the number of Boruta runs (out of 200) that a gene was selected in when counting by cluster (Methods). 65 genes that were up-regulated (47 genes) or down-regulated (18 genes) had a harmonic rank lower than 3 (**Table 1**). The next highest ranked gene had harmonic rank exceeding 17. These 65 genes represent the most persistent differentially expressed CorEx genes. Expectedly, the high-ranked up-regulated genes impacted cell division (16 genes), cell cycle regulation (10 genes), DNA damage response (16 genes), with only 12 of the 47 genes not specifically enriching any known GO biological process. Many of these genes were from small CorEx clusters with less than 3 members, but we also found 6 genes from the HOX gene cluster (cluster #17), and another cluster of 21 genes (cluster #3). Members of cluster #3 appeared in all 200 Boruta runs; however, there were three genes all involved in cell-division (*TPX2*, *KIF2C*, and *AURKA*), each of which appeared in at least 180 Boruta




Figure 5.

Mutational characteristics of ecDNA-containing tumors

(a) Total mutation burden of ecDNA(+) and ecDNA(-) samples. ecDNA(+) samples have significantly higher mutation burden than the ecDNA(-) samples (p -value < 0.0001, Mann-Whitney test). **(b)** Odds ratios of differentially mutated genes in ecDNA(+) and ecDNA(-) (p -value < 0.005). The size of the dot indicates whether the corresponding gene belongs to the Cancer Gene Census (CGC) or not (Non-CGC). Only *TP53* and *BRAF* showed significance at the level of FDR < 0.1 (Benjamini-Hochberg).

runs. High expression among these three genes is associated with poor prognosis⁴², and due to the highly persistent nature of their differential expression across ecDNA(+) samples, they represent a possible widespread vulnerability for ecDNA(+) samples.

Intriguingly, 14 of the 18 down-regulated genes with low harmonic rank came from a single cluster (#2; **Table S2**), and 13 of the 18 genes did not specifically enrich any specific BP ontology. Six of the down-regulated genes appeared in 180 or more Boruta runs (*CHMP7*, *XPO7*, *INTS9*, *TACR1*, *KIAA1967*, and *PCM1*). Some of these genes (*CHMP7*, *XPO7*, *KIAA1967*) are reported to be tumor suppressor genes^{43–45}. However, the exact functional role of down-regulating these genes in ecDNA(+) samples remains to be elucidated.

Discussion

ecDNA is increasingly recognized as a major cause of oncogene amplification, intratumoral genetic heterogeneity, accelerated evolution, and treatment resistance, but many of the underlying processes involved in its formation, function, and progression are not fully understood. The ability to conduct multi-omic studies of well-curated, bona fide clinical tumor samples, such as the TCGA, presents an opportunity to learn about differentially regulated gene expression programs that may be involved in ecDNA biogenesis or maintenance, and in worse outcomes for patients^{7–9}. Using a relatively intuitive set of principles, we have developed a machine learning approach that identifies differentially expressed, co-regulated genes in ecDNA-containing tumors, highlighting four main biological processes: non-c-NHEJ DSB repair, cell cycle, proliferation control, and immune regulation.

The GO analysis revealed three core biological processes that were up-regulated and only the immune system processes as being down-regulated. These observations strengthen the case for targeting proteins involved in mitotic cell-division⁴⁶, cell-cycle regulation, and DNA damage response in ecDNA(+) cancers, but also reveal roles for the *HOX* cluster of genes. Also, in this paper, we did not extensively study the role of ncRNA in ecDNA maintenance. We do note that *HOTAIR*, encoded in the *HOXC* locus, is independently associated with metastasis and poor outcomes⁴⁷. Further experiments are needed to provide a mechanistic basis for the role of *HOX* cluster genes in ecDNA maintenance, as also for involvement of ncRNA.

The DNA damage genes are broadly up-regulated in ecDNA(+) samples, especially in double-strand break repair. Within this broad category of mechanisms, our analysis suggests that alternative DSB repair pathways such as Alt-EJ are preferred relative to classical NHEJ. This is consistent with previous observations of small microhomologies at breakpoint junctions^{2, 48}, and has important implications in therapeutic selection that will need to be validated in future experimental studies. We note, however, the microhomology analyses typically study breakpoint junctions, and might ignore double-strand breaks in non-junctional sequences which could be observed, for example at replication-transcription junctions.

The down-regulated genes were primarily immunomodulatory in nature, in addition to a few persistently down-regulated tumor suppressor genes. Lowered expression of immunomodulatory genes in ecDNA(+) samples has been previously reported^{2, 38}, but not mechanistically explained. Remarkably, the down-regulated immunomodulatory genes encompassed most aspects of the cancer immune cycle, suggesting impaired recognition of tumor DNA and proteins as foreign in ecDNA(+) tumors. Sensing of foreign DNA, including tumor DNA, is often mediated by the cGAS/STING pathway^{49, 50}. Intriguingly, cGAS was significantly up-regulated in ecDNA(+) samples, while STING was significantly down-regulated, suggesting a role for STING agonists in intervention. Finally, in addition to the down-regulation of genes in the toll-like receptor family,

we observed a down-regulation of genes involved in regulating TLR signaling pathways that were part of the CorEx list. Understanding the mechanisms of broad down-regulation of TLRs could provide insight into vulnerabilities of ecDNA(+) tumors.

Mutation data alone does not provide as clear a picture of the genes involved in ecDNA maintenance. We did observe that the total mutation burden (TMB) was higher in ecDNA(+) samples. However, that relationship is much less clear after controlling for cancer type. High TMB has been positively correlated with sensitivity to immunotherapy⁵¹, and better patient outcomes; however, the gene expression patterns suggest that immunomodulatory genes are down-regulated in ecDNA(+) samples, and patients with ecDNA(+) tumors have worse outcomes². Notably, other results have suggested that the correlation between TMB and response to immunotherapy is not uniform, and it can vary across different tumor subtypes⁵². In general, no collection of gene mutations was predictive of ecDNA status, although mutations in *TP53* were more likely in ecDNA(+) samples, and perhaps are an important driver for ecDNA formation⁵.

These results suggest that cancer cells that contain ecDNA have profound alterations in their global transcriptional patterns. Importantly, these transcriptional differences do not arise solely from genes on the ecDNAs themselves, but rather suggest that fundamental global processes involved in DSB repair, cell cycle control, and immune regulation contribute to ecDNA formation and pathogenesis.

Methods

TCGA sample ecDNA status classification

Amplicon Classifier (version 0.4.9, <https://github.com/jluebeck/AmpliconClassifier>) classified amplicons detected in 1,921 TCGA samples into five sub-types: ecDNA, BFB, complex non-cyclic, linear, and no-amplification. When classifying a sample with multiple amplicons, the order of preference is as follows: ecDNA, BFB, complex non-cyclic, linear, and no-amplification. Given the challenges of detecting ecDNA from short read data, and to avoid possible false-negative ecDNA classifications, samples with a BFB or complex non-cyclic status which were not called ecDNA(+), were removed from the analysis. We treated samples with the linear amplification and no-amplification classifications as ecDNA(-). Of the 1,921 samples, 1,535 samples classified as ecDNA(+) and ecDNA(-) had RNA-seq data, including 1,406 primary solid tumor samples, 95 tumor metastasis samples, and 34 primary blood derived cancer – peripheral blood samples. Removing metastases results in a total of 1,440 samples, including 243 ecDNA(+) and 1,197 ecDNA(-) samples. While the set of 1,440 samples represented 24 tumor types, ten of these tumor types had insufficient numbers of ecDNA(+) samples, including four tumor types with no ecDNA(+) samples. To prevent the 561 ecDNA(-) samples representing these tumors from skewing the analysis, we removed 570 samples representing tumor types with less than three ecDNA(+) samples. This resulted in a total of 870 samples representing 14 tumor types, of which 234 were classified as ecDNA(+) and 636 were classified as ecDNA(-).

Gene expression datasets

Gene expression data for 32 studies part of the TCGA Pan-cancer Atlas was downloaded from cBioPortal (01.05.2021) (<https://www.cbioportal.org/>). The cBioPortal “data_RNA_Seq_v2_expression_median.txt” data is sourced from the file “EB++AdjustPANCAN_IlluminaHiSeq_RNASeqV2.geneExp.tsv” (synapse id: syn4976363). Briefly, the matrices contain batch corrected values of the upper-quartile (UQ) normalized RSEM estimated counts data from Broad firehose (tumor.uncv2.mRNAseq_RSEM_all.txt). Missing values due to the batch effect correction process were imputed using *K*-nearest neighbors (KNN). For each tumor type, values were imputed based on gene vectors under the assumption that genes are

similarly expressed between samples of the same tumor type. For genes with less than 60% of samples with missing values, values were imputed using the logarithmic (base 2) of the gene expression value plus one, and subsequently back-transformed when writing the imputed matrices to file. The resulting gene expression matrix used for the Boruta analysis described below consisted of 870 TCGA samples and 16,309 protein-coding genes (based on “hgnc_complete_set.txt” downloaded from HGNC on 7.24.2018). To generate a RSEM raw counts matrix for the DESeq2 analysis described below, mRNAseq_Preprocess.Level_3 data was downloaded from Broad Firehose (tumor.uncv2.mRNAseq_raw_counts.txt).

Boruta analysis

To identify a minimal set of genes whose expression values were predictive of the sample being ecDNA(+), we used an automated feature selection algorithm, Boruta¹⁰. Given a gene expression matrix, M , of dimension $r \times c$, where r is the number of samples and c is the number of genes, Boruta generates c shadow feature vectors by random shuffling of gene vectors in matrix M , generating a new matrix M' of dimension $r \times 2c$. A random forest algorithm is then used to quantify the importance of each gene in separating ecDNA(+) from ecDNA(-) samples. In a given Boruta run, for each iteration $i \in n$, where $i = \{1, 2, 3, \dots, n\}$, there are two possible outcomes for a gene: 1) if the gene scores higher than the best scoring shadow feature, the gene is considered a “hit”, and 2) if the gene scores lower than the best scoring shadow feature, the gene is considered a “non-hit”. The number of hits, k_g , for each gene, g , after i iterations is stored in a list, L_i . Similar to a coin toss event of n trials, the probability that a gene is a “hit” or “non-hit” in each iteration is 0.5, and follows the binomial distribution.

After i iterations, the probability that a gene has k or more hits, $P(X > x)$, where $x = k - 1$, can be computed using the binomial distribution survival function. This probability is computed for each gene, g , with k_g number of hits and stored in list P_i . As Boruta involves testing multiple genes per iteration and testing the same genes repeatedly over i iterations, two multiple hypothesis testing corrections are applied to list P_i . First, the Benjamini-Hochberg procedure is applied to control the False Discovery Rate (FDR) at 5%, and second, the Bonferroni correction is applied to control for the family-wise error rate at $\alpha = 0.05$. If the probability of a gene satisfies both corrections, the gene is accepted as a Boruta gene. Separately, the probability that a gene, g , has k_g or lower hits, $P(X \leq k_g)$, can be computed using the binomial cumulative distribution function. For genes that are not accepted, if the probability satisfies both corrections, the gene is rejected as a Boruta gene.

The Boruta script used for this analysis is a revised version of the BorutaPy python package (6.21.2021; https://github.com/scikit-learn-contrib/boruta_py) with the addition of early stop parameters (≤ 50 tentative feature counts stagnant for 5 iterations) and output of hits per iteration. We randomly split the 870 samples into training (80%) and testing (20%) data to enable the evaluation of selected features. Given the unequal representation of ecDNA(+) and ecDNA(-) samples within each of the tumor types, we opted to perform the following procedure 200 times. For each of the 200 runs, Boruta is performed on the training data. In short, the Boruta method takes as input the original expression matrix, M , appends shadow features that are randomized versions of gene features in M , performs random forest classification on the data (class_weight=balanced_subsample, max_depth=7), selects the highest feature importance score, s , of shadow features as a cut off, and returns features with a higher score than s . Of the 941 genes identified in the Boruta analysis, 408 genes were selected in at least 10 of the 200 Boruta runs, and subsequently defined as the Core set of genes in downstream analyses.

Highly co-expressed genes

To identify genes co-expressed with the core set of Boruta genes, hierarchical clustering of the 16,309 genes was performed using the R package pvclust¹¹ (ver. 2.2-0; dist.method=correlation, method=ward.D2, nboot=1000). A total of 843 significant clusters (AU > 0.95) with at least 1 Boruta gene were selected, consisting of 1,375 genes. To obtain the final list of CorEx genes, we apply a

minimal count of 10 runs for the gene or 10 runs for the cluster of genes seen in 200 Boruta runs. A cluster is determined to be seen in a Boruta run if at least one of its members is selected in the run. This results in 354 clusters, with a total number of 643 genes, of which 408 are Core genes.

Evaluation of CorEx genes

To evaluate the 408 Core genes as predictive of ecDNA presence in tumor samples, a random forest classifier was trained on the 80% data sets and tested on the 20% data sets defined previously. Hyper-parameters were tuned using RandomizedSearchCV followed by GridSearchCV ($k_{\text{fold}}=5$, $f1$ score). Furthermore, we performed a similar evaluation using 408 randomly selected genes instead of the Core genes for each of the 200 training and testing data sets. A similar evaluation was performed using 643 randomly selected genes instead of the CorEx genes. We further evaluated a set of 643 most differentially expressed genes based on the absolute log-fold change estimates from a conventional DE analysis using DESeq2¹⁷ as described below.

Precision, recall curve

To select data points (i.e., (precision, recall) values for a training/testing data pair described above) representative of the predictive power of a gene list (i.e., CorEx genes, Core genes, Top-|LFC| genes, or randomly selected genes), for each consecutive window of precision, we selected the data point with the highest recall (sensitivity) value. Specifically, using 16 consecutive windows of precision with a step-size of 0.0625 from 0.0 to 1.0, we selected 8 representative data points for Core genes, 8 for CorEx genes, 8 for Top-|LFC| genes, and 11 for randomly selected genes. The representative points were then used to interpolate 500 new points using the piecewise cubic hermite interpolating polynomial (pchip) function in the Scipy python package. For the figures generated, we only showed data points with a precision value greater or equal to 0.55.

Default DE analysis

We performed a default DESeq2¹⁷ (R package, ver. 1.36.0) analysis to obtain shrunken maximum *a posteriori* (MAP) log-fold change estimates for effect size (i.e., LFC). Specifically, to obtain i) LFC effect size values per gene for integration with its Cliff's delta effect size value when determining if a gene is up- or down-regulated in ecDNA(+) samples, and ii) a list of n top-ranked genes by absolute value of the LFC (with application of an adjusted p -value <0.05 cutoff and LFC threshold of $\log_2(1.1) = 0.13$) for use in comparison against genes selected as important in the prediction of ecDNA in samples. For comparisons against Core genes, n is set to 408, and for comparisons against CorEx genes, n , is set to 643.

To obtain the LFC effect size metric between ecDNA(+) vs. ecDNA(-) samples for each gene's expression, we fed as input to DESeq2 a matrix of raw RSEM estimated counts. To take into account batch effects, we included the center and platform information of samples, downloaded from synapse id syn4976363 (EB++GeneExpAnnotation.tsv), in the design of the DESeq object:

```
DESeq_object <- DESeqDataSetFromMatrix(countData =
AC_0_4_9_TCGA_matrix, colData = coldata, design =
~batch+condition)
```

To compute results, the lfcThreshold was set to $\log_2(1.1)$ for an accurate computation of p -values and the contrast set to $c(\text{"condition", "ecDNA(+)", "ecDNA(-)"})$ to obtain the logarithmic fold change of the form $\left(\frac{\text{ecDNA}(+)}{\text{ecDNA}(-)}\right)$. By setting the lfcThreshold , the null hypothesis tested is that $|LFC| \leq \theta$, where $\theta = (1.1)$, and the alternative hypothesis is that $|LFC| > \theta$. A $\log_2(1.1)$ value is chosen as the

minimal value/negligible effect size threshold as it represents a 10% fold-change, and anything below this fold-change would likely not be of biological interest⁵³. The specific commands run are as follows:

```
DESeq_object$condition <- relevel(DESeq_object$condition, ref =
"Non_ecDNA")
DESeq_object <- DESeq(DESeq_object)
results <- results(DESeq_object, lfcThreshold= log2(1.1),
contrast=c("condition","ecDNA","Non_ecDNA"))
```

To obtain the shrunken MAP log-fold change estimates, we used the `lfcShrink` function provided in DESeq2, using the default `apeglm` method for the empirical Bayes shrinkage procedure⁵⁴:

```
lfcShrink(DESeq_object, lfcThreshold= log2(1.1),
coef="condition_ecDNA_vs_Non_ecDNA", type="apeglm")
```

Up- or Down-regulated genes in ecDNA(+) samples

To categorize genes as “up-” or “down-” regulated in ecDNA(+) samples, we integrated two effect size metrics, Cliff’s delta (d)^{14, 15} and the DESeq2 shrunken MAP log-fold change estimate (LFC)¹⁷. Effect size is a measure of the magnitude of deviation from the null hypothesis, and unlike p -values, has the advantage of not being impacted by sample size⁵⁵. This property is especially useful when comparing effect size values of a gene between tests where sample sizes differ. Comparing p -values between such tests would be invalid.

Cliff’s delta, d , is a non-parametric measure of the separation between two distributions and ranges from -1 to 1 . Given two distributions, $X = \{i_1, i_2, \dots, i_m\}$ and $Y = \{j_1, j_2, \dots, j_n\}$, comparisons are made between each of m values in X and n values in Y . Cliff’s delta is computed as $d = \frac{\#(i > j) - \#(i < j)}{mn}$, where $\#(i > j)$ is the number of times a member of X is greater than a member of Y , and $\#(i < j)$ is the number of times a member of X is less than a member of Y ¹⁴. A negative d indicates that values in Y tend to be higher than X , while a positive d indicates that values in X tend to be higher than Y . The magnitude of the effect size of Cliff’s delta can be separated into four levels: $|d| < 0.147$ for negligible effects, $0.147 \leq |d| < 0.33$ for small effects, $0.33 \leq |d| < 0.474$ for medium effects, and $|d| \geq 0.474$ for large effects⁵⁵. The python package used to compute Cliff’s delta values can be accessed at <https://github.com/neilernst/cliffsDelta>. The input values used to compute Cliff’s delta are log-transformed normalized gene expression values plus one as described in the “Gene expression datasets” section of methods.

The DESeq2 LFC is computed as described above. To allow integration of the Cliff’s delta effect size with the DESeq2 LFC, we also separated the LFC values into four levels: $|LFC| < \log_2(1.1)$ for negligible effects, $\log_2(1.1) \leq |LFC| < \log_2(1.5)$ for small effects, $\log_2(1.5) \leq |LFC| < \log_2(2)$ for medium effects, and $|LFC| \geq \log_2(2)$ for large effects.

For each gene, g , whether the gene is up-regulated or down-regulated in ecDNA(+) samples is determined by the signage and magnitude of its effect sizes d_g and LFC_g . The initial criteria for d_g and LFC_g to be used as a determinant in the direction of a gene is for it to have a magnitude larger than that of a negligible effect. If the signage of d_g and LFC_g are both positive, gene g is considered up-regulated in ecDNA(+). If only a single value has a magnitude larger than the negligible effect threshold (e.g. $d_g > 0$ and $LFC_g = -0.1$), gene g is considered up-regulated in ecDNA(+). In the case of conflicting signages between the two values, the effect size with a larger magnitude takes precedence. For example, if $d_g = -0.2$ and $LFC_g = 0.847$, given that d_g has a small negative effect and LFC_g has a large positive effect, gene g is considered up-regulated in ecDNA(+) samples.

Tumor heatmap

A Cliff’s delta effect size matrix representing 643 CorEx genes was generated to compare TCGA with tumor expression patterns. For each of the 11 tumor types with at least 10 ecDNA(+) and at least 10 ecDNA(-) samples, we re-computed Cliff’s delta. Using a Fisher’s exact test (`fisher_exact`

function from the `scipy.stats` python package; alternative hypothesis: two-sided), we tested the null-hypothesis of whether the up- and down-directionality of CorEx genes in TCGA vs. each tumor were independent of each other. The contingency table is as below. The directionality of a gene (up or down) is based solely on the signage of the gene's Cliff's delta effect size value.

	Tumor	
	UP	DOWN
TCGA	a	b
	c	d

Gene set enrichment analysis

To identify gene ontology biological processes that are enriched in our CorEx set of genes, we applied a hypergeometric test (one-sided Fisher's exact test; `fisher_exact` function from the `scipy.stats` python package) using gene set to gene mappings from `msigdb` (`c5.go.bp.v7.5.1.entrez`) and a universe of $N=16k$ for the number of genes measured in the RNAseq data. The false discovery rate was controlled at 5% and adjusted p -values were computed using the Benjamini-Hochberg procedure (`fdr` correction from `python statsmodels` package). A final set of GOBP terms with adjusted p -value < 0.05 was used for downstream analysis. The contingency table is as follows, where N is the total number of genes in the universe, n is the number of DE genes (either up- or down-regulated in `ecDNA(+)`), m is the number of genes in the gene set, and k is the number of DE genes in the gene set.

	DE	Non-DE
Inside gene set	k	m-k
Outside gene set	n-k	N+k-n-m

Clustering gene sets

To cluster enriched gene sets into categories for visualization purposes, Cohen's kappa coefficient (`python sklearn cohen_kappa_score`) was used to determine term-term "connectivity" (agreement of term-term pairs) – an approach described in the DAVID paper²¹. Given a $r \times c$ binary matrix, where enriched GOBP terms are rows, CorEx genes are columns, and values are 1 if a CorEx gene is part of the GOBP term or 0 otherwise, kappa scores were computed between each pair of terms, where term $t_i \in t$ and $i = \{1, 2, 3, \dots, n\}$: $Kappa_Score(t_x, t_y)$, where $x \in i$ and $y \in i$.

Each term, t_i , formed an initial seeding group, g_i , where a term t_x is part of g_i if $Kappa_Score(t_i, t_x) \geq \text{Score threshold}$. If at least 50% of term-term pairs in g_i have a $Kappa_Score \geq \text{score threshold}$, the initial seeding group g_i is retained for the next step. The second criteria ensures that terms within the same seeding group have strong interconnectivity. An iterative merging of seeding groups then follows: groups sharing $p\%$ or more members are merged. The representative term for each group was determined as the member with the highest interconnectivity score with other members of the group. After the automatic grouping process, a manual inspection leads to the merging of outliers or smaller groups into representative groups. We used a kappa score threshold of 0.5, and condition of $p \geq 25\%$ of shared members when merging for the down-regulated genes, and a kappa score threshold of 0.6 and condition of $p \geq 50\%$ of shared members when merging for the up-regulated genes.

DDR pathway genes

We hand-curated 88 genes for double-stranded break DNA damage repair pathways (a-EJ, HR, c-NHEJ, SSA) via an extensive literature search, and added an additional 51 genes from the following MSigDB (c5.go.bp.v2022.1.Hs.entrez.gmt) GO biological process terms: GO:0097680 (double strand break repair via classical nonhomologous end joining), GO:0097681 (double strand break repair via alternative nonhomologous end joining), GO:1905168 (positive regulation of double strand break repair via homologous recombination), and GO:0045002 (double strand break repair via single strand annealing). This resulted in a final list of 129 genes. The directionality of the genes, as either up- or down-regulated in ecDNA(+) samples, is based on the full set of 1,440 samples, consisting of 243 ecDNA(+) and 1,197 ecDNA(-) samples representing 24 tumor types (**Table S12**; **Table S13**).

To test whether the number of genes passing our effect size thresholds for all genes 5,256 (UP) and 5,251 (DOWN) was significantly different from the up-/down-regulated genes implicated in each of the 4 pathways, we performed a Fisher's exact test (fisher_exact function from the scipy.stats python package) on the contingency table below, where *c* and *d* are the up- and down-regulated genes for each of the pathways tested.

Contingency table:

	UP	DOWN
All genes	5,256	5,251
Pathway	<i>c</i>	<i>d</i>

Pathway values:

	<i>c</i>	<i>d</i>
c-NHEJ	14	7
Alt-EJ	11	1
HR	46	8
SSA	11	1

Physical presence on amplicons

To determine physical presence of a gene on an amplicon, gene coordinates listed in the gene annotations file GRCh37/human_hg19_september_2011/Genes_July_2010_hg19.gff, downloaded from the AA repo on 3/21/2022, were mapped to amplicon genomic intervals (bed files). A gene is determined to be physically present on an amplicon if its genomic coordinates are fully encompassed within the amplicon genomic intervals.

Mutational analysis

We pulled the list of mutations from the open-access version of MC3 dataset (<https://ellrottlab.org/project/mc3/>), then investigated differences between ecDNA(+) and ecDNA(-) samples. Synonymous mutations were excluded when calculating the mutation burden.

For the differentially mutated gene analysis, only damaging mutations were selected by using snpEff⁵⁶ annotation which is originally included in the MC3 dataset. First, mutations annotated as HIGH in the IMPACT column were selected to obtain frameshift INDELs and stop gain SNVs. Next, mutations predicted to be damaging by SIFT⁵⁷ and PolyPhen2⁵⁸, were selected to obtain damaging missense mutations. Finally, we generated 2-by-2 contingency tables for each gene with cells a , b , c , and d representing the number of individuals with and without damaging mutations in ecDNA(+) and (-) tumors. The odds ratios were computed as $OR = \frac{ad}{bc}$, where a is the number of individuals in ecDNA(+) with the mutation, b is the number of individuals in ecDNA(+) without the mutation, c is the number of individuals in ecDNA(-) with the mutation, and d is the number of individuals in ecDNA(-) without the mutation. To determine if a gene contained mutations that were implicated in cancer, we checked genes against the Cancer Gene Census (CGC) database (v97)⁵⁹, marking genes in the database as CGC and those that were not as non-CGC.

Classification with mutational status

First, we created a binary matrix representing whether a gene is damaged or not, from the MC3 damaging mutation set described as above. Then we divided the whole matrix into 80% of the training set and 20% of the test set. Hyperopt was applied on the training set to select the best parameters for the XGBoost³⁹ model. The optimal parameters estimated by Hyperopt (eta = 0.1, max_depth = 6, min_child_weight = 3.0, scale_pos_weight = 4.9) were input to the XGBoost model, and the ecDNA status of each sample were also input as an answer set. The performance of the model was checked by inputting the 20% test set to the model and comparing the output result with the answer. Principal component analysis was also performed with the same mutational matrix as above, using the scikit.learn package.

Ranking of CorEx genes

To rank CorEx genes by importance, we computed harmonic mean rank values based on three categories: a) the average Gini importance statistic or mean decrease impurity (MDI) MDI (feature importance) values extracted from the trained random forest models on 200 training sets during the evaluation method described above, b) the number of Boruta runs (out of 200) that a gene is selected in, rounded to 2-digits, and c) the number of Boruta runs (out of 200) that a gene is selected in when counting by cluster, rounded to 2-digits. The ranks for each category are adjusted separately so that genes with the same value share the same rank value. For example, if using the Boruta run count as the rank value:

	Run count	round(run_count/10.0)	Rank	Adjusted rank
Gene A	200	20	1	1
Gene B	200	20	2	1
Gene C	200	20	3	1
Gene D	187	19	4	4

The harmonic mean rank is defined as:

$$\text{harmonic mean rank} = \frac{3}{\frac{1}{a} + \frac{1}{b} + \frac{1}{c}}$$

Tumor immune subtype

We classified ecDNA(+) and ecDNA(-) samples into the 6 immune subtype categories (Thorsson et al., 2018³⁷) provided in Table S6 from Bagaev et al., 2021⁶⁰.

Author Contributions

M.S.L., P.S.M, and V.B. conceived and designed the experiments and core narrative of the manuscript. All analyses were performed by M.S.L., apart from the following ones. Mutational analysis was performed by S.J. The amplicon classifications were performed by J.L. S.W, V.B, H.Y.C. and P.S.M. provided guidance on the analysis and interpretation of results. M.S.L., S.J., J.L., and V.B. wrote the manuscript with feedback from all authors.

Acknowledgements

This project was supported by Cancer Grand Challenges CGCSDF-2021\100007 with support from Cancer Research UK and the National Cancer Institute (V.B, H.Y.C., P.S.M.) M.S.L., J.L., and V.B. were supported in part by grants U24CA264379, R01GM114362, and OT2CA278635 from the NIH and CGCATF-2021/100025 from CRUK. S.J. was supported by a grant of the Korea Health Technology R&D Project through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea (grant number: HI19C1330).

Competing Interests

V.B. is a co-founder, paid consultant, SAB member and has equity interest in Boundless Bio, inc. and Abterra, Inc. H.Y.C. is a co-founder of Accent Therapeutics, Boundless Bio, Cartography Biosciences, Orbital Therapeutics, and an advisor of 10x Genomics, Arsenal Biosciences, Chroma Medicine, and Spring Discovery. P.S.M. is a co-founder and advisor of Boundless Bio. J.L. receives compensation as a consultant for Boundless Bio. The remaining authors declare no competing interests.

References

1. Turner K. M., et al. (2017) **Extrachromosomal oncogene amplification drives tumour evolution and genetic heterogeneity** *Nature* **543**:122–125
2. Kim H., et al. (2020) **Extrachromosomal DNA is associated with oncogene amplification and poor outcome across multiple cancers** *Nat. Genet* **52**:891–897
3. Nathanson D. A., et al. (2014) **Targeted therapy resistance mediated by dynamic regulation of extrachromosomal mutant EGFR DNA** *Science* **343**:72–76
4. Lange J. T., et al. (2022) **The evolutionary dynamics of extrachromosomal DNA in human cancers** *Nat. Genet* **54**:1527–1533
5. Luebeck J., et al. (2023) **Extrachromosomal DNA in the cancerous transformation of Barrett's oesophagus** *Nature* <https://doi.org/10.1038/s41586-023-05937-5>
6. Hung K. L., et al. (2021) **ecDNA hubs drive cooperative intermolecular oncogene expression** *Nature* **600**:731–736
7. Wu S., et al. (2019) **Circular ecDNA promotes accessible chromatin and high oncogene expression** *Nature* **575**:699–703
8. Morton A. R., et al. (2019) **Functional Enhancers Shape Extrachromosomal Oncogene Amplifications** *Cell* **179**:1330–1341
9. van Leen E., Brückner L., Henssen A. G. (2022) **The genomic and spatial mobility of extrachromosomal DNA and its implications for cancer therapy** *Nat. Genet* **54**:107–114
9. Kursa M. B., Rudnicki W. R (2010) **Feature Selection with the Boruta Package** *J. Stat. Softw* **36**
10. Suzuki R., Shimodaira H (2006) **Pvclust: an R package for assessing the uncertainty in hierarchical clustering** *Bioinformatics* **22**:1540–1542
11. Kobayashi Y., et al. (2021) **Mitotic checkpoint regulator RAE1 promotes tumor growth in colorectal cancer** *Cancer Sci* **112**:3173–3189
12. Wang Q., et al. (2019) **Prognostic Potential of Alternative Splicing Markers in Endometrial Cancer** *Mol. Ther. - Nucleic Acids* **18**:1039–1048
13. Cliff N (1993) **Dominance statistics: Ordinal analyses to answer ordinal questions** *Psychol. Bull* **114**:494–509
14. Cliff N (1996) **Answering Ordinal Questions with Ordinal Data Using Ordinal Statistics** *Multivar. Behav. Res* **31**:331–350
15. Garsed D. W., et al. (2014) **The Architecture and Evolution of Cancer Neochromosomes** *Cancer Cell* **26**:653–667
16. Love M. I., Huber W., Anders S (2014) **Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2** *Genome Biol* **15**

17. Subramanian A., et al. (2005) **Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles** *Proc. Natl. Acad. Sci* **102**:15545–15550
18. Liberzon A., et al. (2011) **Molecular signatures database (MSigDB) 3.0** *Bioinformatics* **27**:1739–1740
19. Liberzon A., et al. (2015) **The Molecular Signatures Database Hallmark Gene Set Collection** *Cell Syst* **1**:417–425
20. Huang D., et al. (2007) **The DAVID Gene Functional Classification Tool: a novel biological module-centric algorithm to functionally analyze large gene lists** *Genome Biol* **8**
21. Feng Y., et al. (2021) **Homeobox Genes in Cancers: From Carcinogenesis to Recent Therapeutic Intervention** *Front. Oncol* **11**
22. Fish L., et al. (2021) **A prometastatic splicing program regulated by SNRPA1 interactions with structured RNA elements** *Science* **372**
23. Cui H., et al. (2022) **TNFAIP6 Promotes Gastric Carcinoma Cell Invasion via Upregulating PTX3 and Activating the Wnt/ β -Catenin Signaling Pathway** *Contrast Media Mol. Imaging* **2022**:1–9
24. Li Y., et al. (2019) **TMEM203 is a binding partner and regulator of STING-mediated inflammatory signaling in macrophages** *Proc. Natl. Acad. Sci* **116**:16479–16488
25. Feng Z., et al. (2022) **RNF114 Silencing Inhibits the Proliferation and Metastasis of Gastric Cancer** *J. Cancer* **13**:565–578
26. Ribeiro J. R., Lovasco L. A., Vanderhyden B. C., Freiman R. N (2014) **Targeting TBP-Associated Factors in Ovarian Cancer** *Front. Oncol* **4**
27. Xu N., et al. (2020) **SHCBP1 promotes tumor cell proliferation, migration, and invasion, and is associated with poor prostate cancer prognosis** *J. Cancer Res. Clin. Oncol* **146**:1953–1969
28. Chang H. H. Y., Pannunzio N. R., Adachi N., Lieber M. R (2017) **Non-homologous DNA end joining and alternative pathways to double-strand break repair** *Nat. Rev. Mol. Cell Biol* **18**:495–506
29. Ghosh D., Raghavan S. C (2021) **20 years of DNA Polymerase μ , the polymerase that still surprises** *FEBS J* **288**:7230–7242
30. Daley J. M., Sung P (2014) **53BP1, BRCA1, and the Choice between Recombination and End Joining at DNA Double-Strand Breaks** *Mol. Cell. Biol* **34**:1380–1388
31. Fouquin A., et al. (2017) **PARP2 controls double-strand break repair pathway choice by limiting 53BP1 accumulation at DNA damage sites and promoting end-resection** *Nucleic Acids Res* **45**:12325–12339
32. He L., Guo Z., Wang W., Tian S., Lin R (2023) **FUT2 inhibits the EMT and metastasis of colorectal cancer by increasing LRP1 fucosylation** *Cell Commun. Signal. CCS* **21**
33. Lawrence T (2009) **The Nuclear Factor NF- κ B Pathway in Inflammation** *Cold Spring Harb. Perspect. Biol* **1**:a001651–a001651

34. Urban-Wojciuk Z., et al. (2019) **The Role of TLRs in Anti-cancer Immunity and Tumor Rejection** *Front. Immunol* **10**
35. Chen D. S., Mellman I (2013) **Oncology Meets Immunology: The Cancer-Immunity Cycle** *Immunity* **39**:1–10
36. Thorsson V., et al. (2018) **The Immune Landscape of Cancer** *Immunity* **48**:812–830
37. Wu T., et al. (2022) **Extrachromosomal DNA formation enables tumor immune escape potentially through regulating antigen presentation gene expression** *Sci. Rep* **12**
38. Chen T., Guestrin C. (2016) **Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 785–794 (Association for Computing Machinery, 2016). doi:10.1145/2939672.2939785. XGBoost: A Scalable Tree Boosting System. in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining :785–794 <https://doi.org/10.1145/2939672.2939785>**
39. Bergstrom E. N., et al. (2022) **Mapping clustered mutations in cancer reveals APOBEC3 mutagenesis of ecDNA** *Nature* **602**:510–517
40. Hadi K., et al. (2020) **Distinct Classes of Complex Structural Variation Uncovered across Thousands of Cancer Genome Graphs** *Cell* **183**:197–210
41. Luca M. D., Lavia P., Guarguaglini G (2006) **A Functional Interplay Between Aurora-A, Plk1 and TPX2 at Spindle Poles: Plk1 Controls Centrosomal Localization of Aurora-A and TPX2 Spindle Association** *Cell Cycle* **5**:296–303
42. Guo Y., Shi J., Zhao Z., Wang M (2021) **Multidimensional Analysis of the Role of Charged Multivesicular Body Protein 7 in Pan-Cancer** *Int. J. Gen. Med* **14**:7907–7923
43. Innes A. J., et al. (2021) **XPO7 is a tumor suppressor regulating p21 CIP1-dependent senescence** *Genes Dev* **35**:379–391
44. Qin B., et al. (2015) **DBC1 functions as a tumor suppressor by regulating p53 stability** *Cell Rep* **10**:1324–1334
45. Von Hoff D. D., et al. (1992) **Elimination of extrachromosomally amplified MYC genes from human tumor cells reduces their tumorigenicity** *Proc. Natl. Acad. Sci* **89**:8165–8169
46. Gupta R. A., et al. (2010) **Long noncoding RNA HOTAIR reprograms chromatin state to promote cancer metastasis** *Nature* **464**:1071–1076
47. Sanborn J. Z., et al. (2013) **Double Minute Chromosomes in Glioblastoma Multiforme Are Revealed by Precise Reconstruction of Oncogenic Amplicons** *Cancer Res* **73**:6036–6045
48. Samson N., Ablasser A. (2022) **The cGAS–STING pathway and cancer** *Nat. Cancer* <https://doi.org/10.1038/s43018-022-00468-w>
49. Sun L., Wu J., Du F., Chen X., Chen Z. J (2013) **Cyclic GMP-AMP Synthase Is a Cytosolic DNA Sensor That Activates the Type I Interferon Pathway** *Science* **339**:786–791
50. Rizvi N. A., et al. (2015) **Mutational landscape determines sensitivity to PD-1 blockade in non-small cell lung cancer** *Science* **348**:124–128

51. McGrail D. J., et al. (2021) **High tumor mutation burden fails to predict immune checkpoint blockade response across all cancer types** *Ann. Oncol* **32**:661–672
52. McCarthy D. J., Smyth G. K (2009) **Testing significance relative to a fold-change threshold is a TREAT** *Bioinformatics* **25**:765–771
53. Zhu A., Ibrahim J. G., Love M. I (2019) **Heavy-tailed prior distributions for sequence count data: removing the noise and preserving large differences** *Bioinformatics* **35**:2084–2092
54. Romano Jeanine, Kromrey Jeffrey D., Coraggio Jesse, Skowronek Jeff (2006) **Appropriate statistics for ordinal level data: Should we really be using t-test and Cohen'sd for evaluating group differences on the NSSE and other surveys?** *Annu. Meet. Fla. Assoc. Institutional Res* **177**
55. Cingolani P., et al. (2012) **A program for annotating and predicting the effects of single nucleotide polymorphisms, SnpEff: SNPs in the genome of Drosophila melanogaster strain w1118; iso-2; iso-3** *Fly (Austin)* **6**:80–92
56. Ng P. C., Henikoff S (2003) **SIFT: Predicting amino acid changes that affect protein function** *Nucleic Acids Res* **31**:3812–3814
57. Adzhubei I., Jordan D. M., Sunyaev S. R (2013) **Predicting functional effect of human missense mutations using PolyPhen-2** *Curr. Protoc. Hum. Genet. Chapter 7*
58. Sondka Z., et al. (2018) **The COSMIC Cancer Gene Census: describing genetic dysfunction across all human cancers** *Nat. Rev. Cancer* **18**:696–705
59. Bagaev A., et al. (2021) **Conserved pan-cancer microenvironment subtypes predict response to immunotherapy** *Cancer Cell* **39**:845–865

Author information

Miin S. Lin

Bioinformatics and Systems Biology Graduate Program, University of California at San Diego, La Jolla, CA, USA

Se-Young Jo

Department of Biomedical Systems Informatics and Brain Korea 21 PLUS Project for Medical Science, Yonsei University College of Medicine, Seoul, South Korea

Jens Luebeck

Department of Computer Science and Engineering, University of California at San Diego, La Jolla, CA, USA

Howard Y. Chang

Center for Personal Dynamic Regulomes, Stanford University, Stanford, CA, USA, Department of Genetics, Stanford University, Stanford, CA, USA, Howard Hughes Medical Institute, Stanford University, Stanford, CA, USA

Sihan Wu

Children's Medical Center Research Institute, University of Texas Southwestern Medical Center, Dallas, TX, USA

Paul S. Mischel

Sarafan Chemistry, Engineering, and Medicine for Human Health (Sarafan ChEM-H), Stanford University, Stanford, CA, USA, Department of Pathology, Stanford University School of Medicine, Stanford, CA, USA

For correspondence: pmischel@stanford.edu

Vineet Bafna

Department of Computer Science and Engineering, University of California at San Diego, La Jolla, CA, USA, Halicioğlu Data Science Institute, University of California at San Diego, La Jolla, CA, USA

For correspondence: vbafna@ucsd.edu

Editors

Reviewing Editor

Thomas Gingeras

Cold Spring Harbor Laboratory, United States of America

Senior Editor

Richard White

Ludwig Institute for Cancer Research, University of Oxford, United Kingdom

Reviewer #1 (Public Review):

Recently discovered extrachromosomal DNA (ecDNA) provides an alternative non-chromosomal means for oncogene amplification and a potent substrate for selective evolution of tumors. The current work aims to identify key genes whose expression distinguishes ecDNA+ and ecDNA- tumors and the associated processes to shed light on the biological mechanisms underlying ecDNA genesis and their oncogenic effects. While this is clearly an important question, the analysis and the evidence supporting the claims are weak. The specific machine learning approach seems unnecessarily convoluted, insufficiently justified and explained, and the language used by the authors conflates correlation with causality. This work points to specific GO processes associated (up and down) with ecDNA+ tumors, many of which are expected but some seem intriguing, such as association with DSB pathways. My specific comments are listed below.

A. The claim of identifying genes required to 'maintain' ecDNA+ status is not justified - predictive features are not necessarily causal.

B. The methods and procedures to identify the key genes is hyperparameterized and convoluted and casts doubt on the robustness of the findings given the size and heterogeneity of the data.

a. In the first two paragraphs of Boruta Analysis Methods section, authors describe an iterative procedure where in each iteration, a binomial p-value is computed for each gene based on number of iterations thus far in which the gene was selected (higher GINI index than max of shadow features). But then in the third paragraph they simply perform Random Forest in 200 random 80% of samples and pick a gene if it is selected in at least 10/200. It is ultimately not clear what was done. Why 10/200? Also "the probability that a gene is a "hit" or "non-hit" in each iteration is 0.5" is unclear. That probability is of a gene achieving GINI index higher than the max of shadow features. How can it be 0.5?

b. The approach of combining genes with clusters is arbitrary. Why not start with clusters and evaluate each cluster (using some gene set summary score) for their ability to discriminate? Ultimately, one needs additional information to disambiguate correlated genes (i.e. in a co-

expression cluster) in terms of causality.

c. The cross-validation procedure is not clear at all. There is a mention of 80-20 split but exactly how/if the evaluation is done on the 20% is muddled. The way precision-recall procedure is also a bit convoluted - why not simply use the area under the PR curve?

d. The claim is that Boruta genes are different from differentially expressed genes but the differential expression seems to be estimated without regards to cancer type, which would certainly be highly biased and misleading. Why not do a simple regression of gene expression by ecDNA status, cancer type and select the genes that show significant coefficient for ecDNA status?

C. After identifying key features (which the authors inappropriately imply to be causal) they perform a series of enrichment/correlative analysis.

a. It is known that ecDNA status associates with poor survival, and so are cell cycle related signal. Then the association between Boruta genes and those processes is entirely expected. Is it not? The same goes for downregulation of immune processes.

b. The association with DSB specifically is interesting. Further analysis or discussion of why this should be would strengthen the work.

c. On page 15, second paragraph, when providing the up versus down CorEx genes, please also provide up versus down for non-CorEx genes as well to get a sense of magnitude.

d. The finding that Boruta genes are associated with high mutation burden is intriguing because in general mutation burden is associated with better survival and immunotherapy response. This counter-intuitive result should be scrutinized more to strengthen the work.

e. On page 17 "12 of the 47 genes not specifically enriching any known GO biological Process" is confusing. How can individual gene enrich for a GO process?

Reviewer #2 (Public Review):

In their manuscript entitled "Transcriptional immune suppression and upregulation of double stranded DNA damage and repair repertoires in ecDNA-containing tumors" Lin et al. describe an important study on the transcriptional programs associated with the presence of extrachromosomal DNA in a cohort of 870 cancers of different origin. The authors find that compared to cancers lacking such amplifications, ecDNA+ cancers express higher levels of DNA damage repair-associated genes, but lower levels of immune-related gene programs.

This work is very timely and its findings have the potential to be very impactful, as the transcriptional context differences between ecDNA+ and ecDNA- cancers are currently largely unknown. The observation that immune programs are downregulated in ecDNA+ cancers may initiate new preclinical and translational studies that impact the way ecDNA+ cancers are treated in the future. Thus, this study has important theoretical implications that have the potential to substantially advance our understanding of ecDNA+ cancers.

Strengths

The authors provide compelling evidence for their conclusions based on large patient datasets. The methods they used and analyses are rigorous.

Weaknesses

The biological interpretation of the data remains observational. The direct implication of these genes in ecDNA(+) tumors is not tested experimentally.

Reviewer #3 (Public Review):

Summary:

Using a combination of approaches, including automated feature selection and hierarchical clustering, the author identified a set of genes persistently associated with extrachromosomal DNA (ecDNA) presence across cancer types. The authors further validated the gene set identified using gene ontology enrichment analysis and identified that upregulated genes in

extrachromosomal DNA-containing tumors are enriched in biological processes like DNA damage and cell proliferation, whereas downregulated genes are enriched in immune response processes.

Major comments:

1. The authors presented a solid comparative analysis of ecDNA-containing and ecDNA-free tumors. An established automated feature selection approach, Boruta, was used to select differentially expressed genes (DEG) in ecDNA(+) and ecDNA(-) TCGA tumor samples, and the iterative selection process and two-tier multiple hypothesis testing ensured the selection of reliable DEGs. The author showed that the DEG selected using Boruta has stronger predictive power than genes with top log-fold changes.
2. The author performed a thorough interpretation of the findings with GO enrichment analysis of biological processes enriched in the identified DEG set, and presented interesting findings, including the enrichment in DNA damage process among the genes upregulated in ecDNA(+) tumors.
3. Overall, the authors achieved their aims with solid data mining and analysis approaches applied to public data tumor data sets.
4. While it may not be the scope of this study, it will be interesting to at least have some justification for choosing Boruta over other feature selection methods, such as Recursive Feature Elimination (RFE) and backward stepwise selection.
5. The authors showed that DESEQ-selected DEGs with top log-fold changes have less strong predictive power and speculated that this may be due to the fact that genes with top log-fold changes (LFC) are confined only to a small subset of samples. It will be interesting to select DEGs with top log-fold changes after first partitioning the tumor samples. For example, randomly partition the tumor samples, identify the DEGs with top LFC, combine the DEGs identified from each partition, then evaluate the predictive power of these DEGs against the Boruta-selected DEGs.
6. While the authors showed that the presence of mutations was not able to classify ecDNA(+) and (-) tumor samples, it will be interesting to see if variant allele frequencies of the genes containing these mutations have predictive power.