frontiers | Frontiers in Pharmacology

TYPE Original Research
PUBLISHED 03 February 2025
DOI 10.3389/fphar.2024.1516621

OPEN ACCESS

EDITED BY
James Costello,
University of Colorado Anschutz Medical Campus, United States

REVIEWED BY
Shouhai Zhu,
Mayo Clinic, United States
Joshua Black,
University of Colorado Anschutz Medical Campus, United States

*CORRESPONDENCE
Elizabeth Brunk,
elizabeth_brunk@med.unc.edu

†These authors have contributed equally to this work and share first authorship

RECEIVED 24 October 2024
ACCEPTED 12 December 2024
PUBLISHED 03 February 2025

CITATION
Goble K, Mehta A, Guilbaud D, Fessler J, Chen J, Nenad W, Ford CG, Cope O, Cheng D, Dennis W, Gurumurthy N, Wang Y, Shukla K and Brunk E (2025) Leveraging AI to automate detection and quantification of extrachromosomal DNA to decode drug responses.
Front. Pharmacol. 15:1516621.
doi: 10.3389/fphar.2024.1516621

Leveraging AI to automate detection and quantification of extrachromosomal DNA to decode drug responses

Kohen Goble[1†], Aarav Mehta[2†], Damien Guilbaud[3], Jacob Fessler[2], Jingting Chen[4], William Nenad[5,6], Christina G. Ford[3], Oliver Cope[3], Darby Cheng[1], William Dennis[3], Nithya Gurumurthy[3], Yue Wang[7], Kriti Shukla[1] and Elizabeth Brunk[1,3,7,5,8*†]

[1]Department of Chemistry, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [2]Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [3]Integrative Program for Biological and Genome Sciences (IBGS), University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [4]Department of Biochemistry and Biophysics, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [5]Computational Medicine Program, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [6]Curriculum in Bioinformatics and Computational Biology, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [7]Department of Pharmacology, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States, [8]Lineberger Comprehensive Cancer Center, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States

Introduction: Traditional drug discovery efforts primarily target rapid, reversible protein-mediated adaptations to counteract cancer cell resistance. However, cancer cells also utilize DNA-based strategies, often perceived as slow, irreversible changes like point mutations or drug-resistant clone selection. Extrachromosomal DNA (ecDNA), in contrast, represents a rapid, reversible, and predictable DNA alteration critical for cancer's adaptive response.

Methods: In this study, we developed a novel post-processing pipeline for automated detection and quantification of ecDNA in metaphase Fluorescence in situ Hybridization (FISH) images, leveraging the Microscopy Image Analyzer (MIA) tool. This pipeline is tailored to monitor ecDNA dynamics during drug treatment.

Results: Our approach effectively quantified ecDNA changes, providing a robust framework for analyzing the adaptive responses of cancer cells under therapeutic pressure.

Discussion: The pipeline not only serves as a valuable resource for automating ecDNA detection in metaphase FISH images but also highlights the role of ecDNA in facilitating swift and reversible adaptation to epigenetic remodeling agents such as JQ1.

KEYWORDS

cytogenetics, extrachromosomal DNA, ecDNA, double minute chromosomes, machine learning, computer vision, fluorescence *in situ* hybridization, deep neural networks

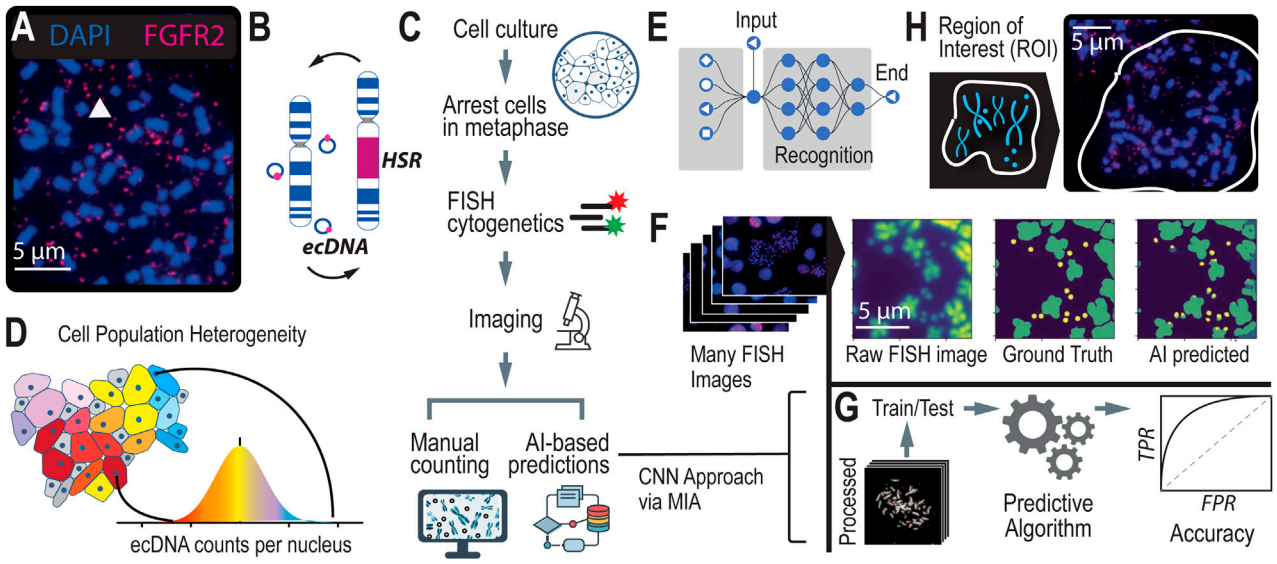

FIGURE 1
Scaling Annotations of ecDNA in Metaphase FISH images using AI **(A)**. Microscopy methods, such as scanning electron microscopy (SEM) and DNA Fluorescence *in situ* Hybridization (FISH) are gold standard approaches for visualizing ecDNA. DNA FISH uses fluorescent probes that bind to specific DNA sequences that indicate which regions of the genome are amplified by ecDNA. Multiple species can be observed depending on which fluorescent probe is detected. **(B)** Under stress, ecDNAs can re-enter chromosomes to form homogeneous staining regions (HSRs). **(C)** FISH is the gold standard method for detecting and analyzing ecDNA in cell nuclei. It consists of several steps which include cell cycle synchronization, fixation, and hybridization. With advances in AI, time-intensive manual labeling and counting of ecDNA is accelerated and can be scaled up. **(D)** Asymmetric division of ecDNA molecules into daughter cells during replication and division leads to a heterogeneous population of cells with various ecDNA counts and species. **(E)** Convolutional Neural Networks is a computer vision approach that is well suited to image segmentation. **(F)** A computer vision approach applied to detecting and counting ecDNAs in metaphase FISH images. **(G)** Using ground truth (i.e., manually labeled metaphase FISH images), we can assess the accuracy of the computer vision method at detecting and counting ecDNAs per cell nucleus. **(H)** Region of Interest (ROI), defined as the human-annotated boundary representing the nuclear spread of a single cell. This is the area within the image where ecDNA is predicted and counted.

Introduction

Cancer cells are masters of adaptation, often using irreversible genomic changes to survive treatment (Labrie et al., 2022). While traditional drug discovery focuses on blocking temporary, protein-level responses, it overlooks the fast-acting genomic responses on extrachromosomal DNA (ecDNA, Figure 1A), which are harder to detect and even harder to target. This blind spot leaves 15% of cancer patients without effective therapies (Kim et al., 2020). EcDNA is a distinct form of genetic alteration that amplifies oncogenes outside the constraints of chromosomal DNA, providing cancer cells with unique evolutionary advantages. Unlike chromosomal mutations, deletions, or stable amplifications integrated within homogeneous staining regions (HSRs), ecDNA is a dynamic structure capable of rapid replication and uneven segregation during cell division. In tumors, ecDNA is believed to drive tumor heterogeneity by creating cell-to-cell genetic and molecular differences that influence drug responses. This variability allows certain cells to adapt and persist during therapy, contributing to the development of drug resistance.

Next-generation sequencing has revolutionized cancer biology, enabling unprecedented insights into the genetic landscape of cancer. Large-scale consortia such as The Cancer Genome Atlas (TCGA), the Dependency Map (DepMap) (Barretina et al., 2012; Tsherniak et al., 2017), and the Cancer Cell Line Encyclopedia (CCLE) have generated extensive multi-omics datasets, characterizing thousands of cancer genomes and tumor samples. Among these datasets, whole-genome sequencing (WGS), whole-exome sequencing (WES), and ATAC-seq are the most comprehensively covered across samples. Leveraging this data, computational algorithms (Yang et al., 2023; Deshpande et al., 2019) have been developed to detect unique structural variants, including ecDNA and double minute chromosomes (DMs), which frequently appear in cancer cell nuclei. These advancements have fundamentally shifted our understanding of cancer genome structure and karyotype, opening new avenues for research and potential therapeutic interventions.

However, despite their transformative potential, sequencing-based approaches have significant limitations in practice. While they are effective at identifying amplified genes and reconstructing the circular structures of the most dominant ecDNA species, their prediction accuracies typically range between 60%–70% (Fessler et al., 2024). This level of accuracy is often insufficient for precise detection and characterization, particularly when detecting ecDNA in cells with fewer copy number counts. Additionally, a major limitation of sequence-based prediction algorithms is that they are not able to distinguish between ecDNA and other amplification types, such as HSRs, which are chromosomal regions that can also exhibit gene amplification with "copy-paste-amplify" patterns (Figure 1B).

In this way, sequencing alone cannot differentiate between freely mobile ecDNA, which can independently replicate and segregate during cell division, and HSRs, which are stable and integrated into chromosomes. However, it is important to note that ecDNA and HSRs are not mutually exclusive entities and can interconvert under

certain conditions. This dynamic interconversion further complicates the interpretation of sequencing-based methods, which may miss critical aspects of cancer biology, including the role of ecDNA in driving tumor heterogeneity and drug resistance. Therefore, to accurately detect and characterize ecDNA within cancer cell nuclei, it is essential to integrate additional approaches, such as high-resolution nuclear imaging or functional assays, alongside sequencing data. These complementary methods provide the necessary context and precision, enabling a more comprehensive understanding of ecDNA's role in cancer, including its potential to transition between ecDNA and HSR forms.

Currently, the only methods capable of accurately detecting and quantifying ecDNA, as well as determining its precise location within cell nuclei, are gold-standard cytogenetic imaging techniques. Techniques like Fluorescence *in situ* Hybridization (FISH) and G-banding karyotyping use fluorescence-based methods to visualize ecDNA within the nucleus (Figure 1C). In FISH, cells are typically arrested in metaphase when chromosomes are condensed and spread out, allowing for better visualization of ecDNA. The cell nuclei are stained with fluorescent dyes, such as DAPI, to enhance visibility. When the genetic sequence of ecDNA is known, fluorescent probes can be designed to hybridize with specific regions on the ecDNA, providing unambiguous localization within the nucleus. A standard practice in the field involves generating 20–100 images of individual cell nuclei to capture the heterogeneity in ploidy and ecDNA counts across a cell population (Figure 1D). However, the subsequent manual annotation and analysis of these images are time-intensive and low-throughput, limiting the scalability of these image-based approaches. This laborious process underscores the need for more efficient methods that can achieve the same level of accuracy without the associated bottlenecks.

Leveraging AI to automate the detection of ecDNA and HSRs in metaphase FISH images represents a transformative shift from traditional, labor-intensive karyotyping methods. By training computer vision models to identify ecDNA in metaphase FISH images (Figure 1E), we can overcome the historical challenges of scalability, enhancing both accuracy and precision in detection. This approach complements sequence-based prediction methods by providing a robust alternative that directly visualizes ecDNA. Several efforts (Turner et al., 2017; Rajkumar et al., 2019) have been made to automate image-based detection of ecDNA using various approaches, such as thresholding and Convolutional Neural Networks (CNNs, Figure 1E). These models employ supervised learning, where they are trained on manually labeled imaging data (i.e., "ground truth" data), with each ecDNA precisely identified and annotated by experts. These annotations supply the model with the exact coordinates of ecDNA in each image, enabling the algorithm to learn the features necessary to predict the presence of ecDNA in non-annotated images (Figure 1F). Accuracy of the model is determined by comparing the predicted counts or locations of ecDNA within metaphase FISH images to the ground truth data (Figure 1G).

Autodetection of ecDNA in metaphase FISH images remains a significant challenge due to several factors. Metaphase FISH images are often noisy, with considerable cell-to-cell variation that complicates traditional image processing techniques. Key issues that reduce prediction accuracy include high noise ratios in pixel intensities, variations in the size and morphology of chromosomes versus ecDNA, and a severe class imbalance where the majority of the image is background or chromosomes. Existing algorithms (Turner et al., 2017; Rajkumar et al., 2019) struggle to generalize across different image types, particularly when transitioning from low-resolution, black-and-white, tiled images to high-resolution images. These models, trained on lower-quality images, are heavily dependent on the specific microscope and image processing conditions, making them non-transferable. As a result, they detect only 30%–40% of the ecDNA in our images and suffer from a high false negative rate, especially when ecDNA is close to chromosomes. Moreover, these models are outdated and not publicly available for retraining. There is a critical need for updated, publicly available CNN-based models that can be retrained with higher-resolution images to enhance the accurate detection of ecDNA in metaphase FISH images.

Recently, the Microscopy Image Analyzer (MIA) was developed as an end-to-end interface designed to help researchers scale image analysis across large datasets. MIA is unique in its ability to integrate image processing, analysis, and interpretation within a single platform, providing a streamlined approach that significantly reduces the time and effort required for manual annotation. We applied MIA to our specific challenge of auto-detecting ecDNA in metaphase FISH images, leveraging its capabilities to create a high-throughput post-processing pipeline (Figure 1F). This pipeline dramatically increased the accuracy of ecDNA detection by refining image segmentation, reducing noise, and improving feature recognition. Currently, the field lacks standardized methods for processing and analyzing metaphase FISH images, making our contribution particularly valuable for researchers interested in automating ecDNA detection. Additionally, we have applied this approach to pharmacological studies to monitor changes in ecDNA during drug treatment. These experiments require the analysis of numerous images across many conditions, making the ability to scale and autodetect essential for reducing the time and labor involved in such studies.

Results and discussion

Image processing pipeline for MIA

The quality of images and annotations plays a crucial role in the success of computer vision models, particularly in tasks such as detecting ecDNA in non-annotated images. High-quality images with precise annotations are essential for the computer to accurately learn and identify the critical features necessary for reliable ecDNA detection. Poor image quality or inaccurate annotations can lead to misrepresentations of features, ultimately undermining the model's performance and predictive accuracy. In this section, we discuss the comprehensive steps we took to post-process our image database, focusing on rigorous image ranking and quality control (QC) measures. By ensuring that only the highest quality images were used for training and validation, we aimed to optimize the performance of our MIA-based model for ecDNA detection.

Pre-processing of metaphase FISH images

Accessing image quality

Image quality can vary due to several technical and biological factors, including the type of cell line used, cell-to-cell variations, post-imaging processing, sample handling, probe potency, and the characteristics of the microscope. A raw image must undergo post-imaging processing to qualify for counting. It is common for any human to over- or under-process an image and incorrectly annotate ecDNAs or HSRs, which generates false negatives and/or positives. Additionally, the biology of the nuclei itself introduces variability. For instance, each image may contain one or more cell nuclei that have burst open, exposing chromosomes and ecDNA. However, the way these nuclei burst on the microscope slide is non-uniform and can occur on slightly different planes, resulting in some nuclei being more in focus than others. Cell-to-cell karyotypic differences after bursting, such as the number and clustering of chromosomes, also greatly fluctuate visibility of ecDNA. This variability makes counting ecDNA within a single nucleus subjective and can lead to inconsistencies in manual annotation. Moreover, the use of DNA probes involves heating samples to accelerate probe hybridization, which can degrade the samples and make them more difficult to image accurately.

It is important to create a database of images from which an AI algorithm can train on that are consistent and generalizable. Any overlapping nuclei, debris clouds, or low resolution images decrease the confidence of attribution and negatively impacts the ability of the algorithm to learn correct features for future predictions. While we found that it can be rare to find extremely well resolved and separated bursted nuclei that generate data with high confidence, we tagged and excluded images with lower confidence from training using a systematic grade-based approach.

We ranked images on a scale from 0 to 4 based on several key quality parameters:

1. Resolution and Focus: This parameter assesses the sharpness and clarity of the chromosomes in the image. Images where chromosomes are blurry and indistinguishable, such that the chromosome arms cannot be clearly seen, receive a lower score (0–1). Conversely, images with sharp, well-defined chromosomes receive a higher score (3–4).
2. Uniformity: Uniformity refers to the evenness and consistency of the image. This is evaluated by observing the spatial distribution of the bursted nuclei. If the nuclei are unevenly spread, with overlapping or indistinct regions, the image is marked lower. If the nuclei are evenly distributed with clear boundaries and minimal overlap, the image receives a higher score.
3. Spatial Distinction: This parameter assesses how well the bursted nuclei and chromosomes are separated from surrounding artifacts or debris. If a nucleus is too clustered or overlaps with its surroundings, making it difficult to distinguish individual chromosomes, the image is ranked lower (0–1). Images with well-separated, distinct nuclei that have a clear "splash zone" with no overlap are ranked higher (3–4).
4. Debris and Artifacts: The amount of debris or artifacts in the image is also considered. Images with minimal debris and clean backgrounds are ranked higher, while those with significant debris (Supplementary Figure 1) or distracting artifacts are ranked lower.

Images marked as "0" represent the lowest quality, characterized by poor resolution, lack of uniformity, and high debris, while images marked as "4" are of the highest quality, with clear, well-focused chromosomes, uniform distribution, and minimal artifacts. This ranking system ensures that only the best quality images are used for further analysis and model training.

Assessing annotation quality

Annotation quality in metaphase FISH images for counting ecDNA is a complex and time-consuming task, often leading to inconsistencies due to its subjective nature. Different researchers may label ecDNA differently, resulting in variations in the data. Typically, metaphase FISH images use DAPI dye to localize ecDNA, which stains all nucleotides but does not differentiate between chromosomal DNA and ecDNA. While using DNA probes that hybridize to specific regions of ecDNA could reduce ambiguity, the exact sequences of ecDNA are not always known. Additionally, these probes are expensive, and there can be heterogeneity in ecDNA sequences within a single sample, leading to incomplete labeling if some ecDNA contain the targeted genes and others do not.

Given these challenges, it is more practical to develop algorithms that predict the location of ecDNA in metaphase FISH images using DAPI dye alone. However, different researchers may employ various methods to annotate ecDNA using DAPI, such as adjusting contrast to differentiate between DNA and debris or comparing the general shape and size of potential ecDNA across images. These subjective approaches can lead to significant variation in manual annotations, as there are no clear, standardized criteria for what constitutes an ecDNA. Furthermore, the fact that ecDNA can change shape and size from cell to cell complicates the establishment of consistent annotation rules, making it difficult to generalize across samples.

The annotation of the 3,000 images used in this study was carried out by multiple students over several years. Consequently, the annotation styles, approaches, and quality varied, creating a challenging dataset for machine learning. To address this, our post-processing image ranking system was essential in selecting the most consistent and representative images for training the algorithm. While the model is designed to learn the patterns of ecDNA from the majority of correct annotations, discrepancies in the data—such as errors in labeling of ecDNA that is not visible in the DAPI channel due to its proximity to or overlap with chromosomes—can negatively impact the model's measured accuracy, making it appear less effective than it actually is.

The images used to train this model were first ranked on a scale from 0 to 4, with 0 indicating images that are highly subjective and difficult to count due to unclear focus, and 4 indicating images that are clearly well-counted and in focus. After an initial ranking, the annotations were reviewed more closely, especially in cases where multiple fluorescent images were merged. Particular attention was given to the DAPI channel, as this is crucial for accurate ecDNA detection. If many annotation markers were found to be skewed, missing, or incorrect, the image received a lower rank. Additionally, notes were taken on potential improvements to enhance image quality for future use.

From the 2,312 images that were reviewed and ranked from previous experiments, about 1,242 images were ranked a 3 or 4 and were chosen to be used for highest quality training data. The table below displays descriptions for each rank and examples in their FISH probe merged with DAPI and grayscale DAPI only with annotations. Images that had little annotation issue were flagged to be perfected and upgraded in ranking.

Generating regions of interest

In the context of counting ecDNAs within cell nuclei, generating regions of interest (ROIs, Figure 1H) in metaphase FISH images is a critical initial step. Metaphase FISH images typically contain multiple nuclei, some of which may be intact while others are bursted, revealing the underlying chromosomes and ecDNAs. These images can be complex, with overlapping structures and varying degrees of focus, making it challenging to accurately count ecDNAs without precise guidance. By defining ROIs, researchers can direct the algorithm to focus specifically on the areas of interest—namely, the individual nuclei or their remnants—thereby filtering out irrelevant parts of the image. This targeted approach significantly enhances the accuracy of ecDNA counts, ensuring that the algorithm is analyzing only the relevant nuclear material and not extraneous background or overlapping cells. By clearly delineating ROIs, the post-processing steps are streamlined, leading to more reliable and reproducible results in the study of ecDNA within cancer cells.

With the selected set of high quality images, it was essential to address the presence of debris (Supplementary Figure 1), as well as chromosomes and ecDNA that had drifted from other nuclei, particularly at the edges of the images. These extraneous elements were not originally annotated and needed to be removed to focus on the relevant data. To achieve this, a lasso tool was employed to draw a region of interest (ROI) that encapsulated all annotated ecDNA while minimizing the inclusion of non-annotated material. The resulting ROIs were saved as binary masks. During this process, image contrast was significantly enhanced to differentiate and either include or exclude faintly visible features. For nuclei that were close together but clearly separated, each was carefully reviewed, with distinct cells isolated into their own ROIs when possible or otherwise removed from the dataset. This approach resulted in a variety of ROI shapes, ranging from tight circles to elongated ovals and irregular polygons, effectively tracing around unburst nuclei and isolated debris clouds. Debris and debris clouds were differentiated from ecDNA based on their size and their distance to the centroid of the bursted nuclei. For example, the size of ecDNA from NCI-H2170 cells was estimated to be approximately 0.25 μm (Supplementary Figure 2) from a prior study (Madren et al., 2024), using Scanning Electron Microscopy (SEM) and Correlative Light Electron Microscopy (CLEM) (Madren et al., 2024).

As previously mentioned, subjectivity can arise when interpreting the boundaries of a cell's bursted nuclei radius and distinguishing between circularized debris and ecDNA, leading to inconsistencies in annotations between different individuals. Defining ROIs not only provides the algorithm with a focused area for counting and segmentation but also reduces the influence of these subjective interpretations. By constraining the analysis to clearly defined areas, we minimize the need for multiple revisions of the boundary during post-prediction review. Additionally, this approach facilitates a more consistent and reliable evaluation of ground truth image annotations, ensuring that they can be accurately reviewed and audited within a well-defined context.

Expanding image annotations

Our images were annotated by placing a dot on a single pixel to mark the location of ecDNA. However, this dot may not always align perfectly with the pixel of highest intensity for that ecDNA, which can impact the quality of annotations and subsequently affect the performance of machine learning algorithms. Recognizing the importance of precise annotations in training an algorithm, we developed a method to improve these annotations. We created a mask around each annotated pixel, shaped like a diamond, to include nearby pixels surrounding the ecDNA mark. This approach provided a more comprehensive representation of the ecDNA, allowing the algorithm to better identify and classify ecDNA by capturing a broader context of the signal. This enhancement was crucial in improving the accuracy and robustness of our machine learning model.

Evaluation datasets: human error, ceiling performance, and full model benchmarking

Human error

To better understand the amount of human error involved in annotating metaphase FISH images, we conducted an experiment where we selected a subset of metaphase FISH images. Three different individuals were asked to annotate this subset independently, and one individual annotated the same subset in three separate replicates. Our findings revealed that across the three individuals, there was a 8% ± 1% error in the annotations, indicating the presence of human error. This analysis suggests that any algorithm trained on such data is inherently limited by this level of technical error, implying that the maximum achievable accuracy for the algorithm is constrained by the variability in human annotations. This establishes a benchmark for understanding the limitations of our model's performance and provides a realistic expectation for its accuracy (Figure 2A).

Ceiling performance

To determine the highest achievable accuracy of MIA, we focused on a subset of hand-picked, top-quality metaphase FISH images that were most amenable to automated AI learning. We selected 144 metaphase FISH images that ranked highest according to our stringent image quality criteria, ensuring they were clear, well-focused, and exhibited minimal noise or overlap. By training MIA on this cherry-picked dataset, we aimed to create a "ceiling model" that represents the maximum potential accuracy of ecDNA detection under ideal conditions. While we anticipate that this model will demonstrate superior performance, we also recognize that it may not be generalizable to more diverse or lower-quality images. Nonetheless, this approach provides valuable insight into the upper limits of algorithmic accuracy, establishing a benchmark for what can be achieved in optimal scenarios.

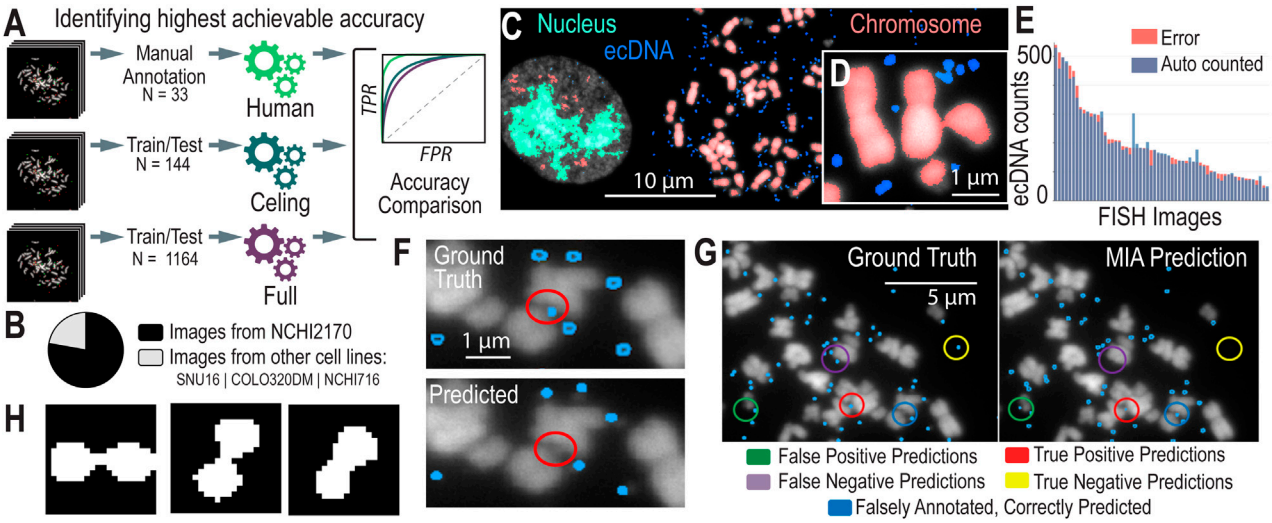

FIGURE 2
AI Models and Assessment of their Accuracies **(A)**. We created different models to test the highest possible accuracy that we could expect our model to reach. **(B)** The breakdown of images used to train the Convolutional Neural Networks (CNNs) by cell line model system. **(C)** CNN-predicted chromosomes, ecDNA, and in-tact nuclei from a metaphase FISH image. **(D)** A zoomed-in image of ecDNA near chromosomes, highlighting the challenges in detecting individual ecDNAs when they cluster close together. **(E)** Accuracy metrics for the full (n = 1,164) model, showing the distribution of error (the difference in ecDNAs predicted versus manually counted) across images. **(F)** Existing algorithms struggle to predict ecDNAs in close proximity to chromosomes. **(G)** Annotated results from MIA, showing instances of true positives/negatives and false positives/negatives as well as cases where MIA finds ecDNA that were not correctly annotated. **(H)** Contour detection of conjoined segmentations. The figure illustrates examples of closely clustered ecDNA entities that are identified as a single object by the algorithm. Contour detection enables the separation of these conjoined objects into distinct segments, improving the accuracy of ecDNA boundary identification and differentiation.

Full model

To develop a robust and generalizable model, we created a full model dataset by selecting images with a ranking of 3 or higher, ensuring a higher standard of image quality for training. This dataset comprises 1,164 images, with 78% derived from our primary model system, NCI-H2170 cells, while the remaining 22% come from a diverse set of cell line model systems that harbor ecDNA, including NCIH716, SNU16, and COLO320 (Figure 2B). By incorporating a variety of cell lines, this dataset aims to strike a balance between achieving high accuracy and maintaining broad applicability across different cellular contexts. The full model represents our comprehensive effort to train an algorithm capable of accurately detecting ecDNA across a range of image qualities and biological conditions.

Segmentation

Segmentation in computer vision is the process of dividing an image into distinct segments or regions, where each segment corresponds to different objects or parts within the image. The primary objective of segmentation is to simplify or transform the representation of an image into something more meaningful and easier to analyze, making it a critical step for tasks like ecDNA recognition.

Semantic segmentation is particularly effective for identifying and analyzing ecDNA within cell nuclei in FISH (Fluorescence *in situ* Hybridization) images, as demonstrated in previous approaches (Rajkumar et al., 2019). By classifying each pixel in an image into predefined categories—such as "ecDNA," "chromosome," "unbursted nuclei," or "background"—semantic segmentation treats all objects of the same category as a single class (Figure 2C). While it does n't distinguish between different instances of the same object, it excels in detecting and segmenting ecDNA within the complex cellular environment (Figure 2D). This approach overcomes challenges posed by overlapping structures and varying signal intensities in metaphase FISH images, enabling more accurate analysis of ecDNA.

Microscopy Image Analyzer (MIA) builds on these principles by incorporating various image processing techniques, including semantic segmentation, depending on the specific model or task. MIA is designed for flexibility and adaptability across different types of microscopy images, including those used for ecDNA detection in metaphase FISH images. While it can employ semantic segmentation similar to previous approaches, MIA primarily focuses on streamlining the analysis of large-scale image datasets through object detection, classification, and region-of-interest (ROI) selection (Figure 1H). By applying techniques like thresholding, edge detection, and ROI selection, MIA enables researchers to identify and quantify features of interest—such as ecDNA—in a high-throughput, efficient manner.

Accuracy and performance

We began by testing the accuracy and performance of our ceiling model. Accuracy was assessed by comparing the total ecDNA count in each image within the validation set against the annotated ground truth. This metric guided the initial iterations of model improvement. The first training sets consisted of approximately 144 ultra-high-confidence, hand-picked images. When evaluated against a validation set of 45 selected images, the model exhibited an absolute difference in total counts of around 8% compared to the ground truth. While this level of accuracy was unexpectedly high, it

TABLE 1 Accuracy metrics across all models.

Model	Pred count	Actual count	Accuracy	Precision	Recall	Matthew	F1
A. MIA model metrics							
Ceiling	203.016	193.778	0.83	0.842	0.877	0.653	0.857
2,170 (3+ images)	212.184	215.868	0.796	0.835	0.817	0.587	0.824
Full model	157.728	214.927	0.637	0.724	0.561	0.305	0.621
CV full model	132.668	180.703	0.621	0.709	0.538	0.279	0.597
B. EcSeg model metrics							
2,170 (3+ images)	191.197	215.868	0.656	0.712	0.641	0.324	0.667
Full model	79.31	214.927	0.424	0.441	0.199	−0.058	0.278
CV full model	78.81	180.703	0.438	0.435	0.237	−0.05	0.315
C. Drug response predictions							
2,170 control (MIA)	284.532	326.675	0.719	0.796	0.696	0.444	0.74
2170 JQ1 (MIA)	100.758	165.379	0.603	0.777	0.47	0.281	0.579
2,170 control (ecSeg)	97.286	326.675	0.413	0.505	0.164	−0.048	0.242
2170 JQ1 (ecSeg)	40.121	165.379	0.375	0.399	0.101	−0.147	0.166
D. Drug response predictions in other cell lines							
COLO320 JQ1 (MIA)	35.828	45.969	0.544	0.62	0.481	0.149	0.535
COLO320 JQ1 (ecSeg)	10.328	45.969	0.347	0.338	0.078	−0.19	0.184
SNU16 JQ1 (MIA)	142.08	260.3	0.658	0.738	0.522	0.382	0.571
SNU16 JQ1 (ecSeg)	18.8	260.3	0.335	0.119	0.013	−0.213	0.039

In Section A, MIA is the CNN-based model and we assess accuracy for the Ceiling, 2,170 (with 3+ or higher ranked images from only NCI-H2170 cells), and Full models. The Full model was split twice to evaluate the differences between training on different images in the dataset. In Section B, ecSeg is the CNN-based model used to assess all but the Ceiling model. In Section C, we use the Full MIA model to predict ecDNA counts in FISH data from control NCI-H2170 cells and cells that have been treated with JQ1. We also assess the Full ecSeg model to predict ecDNA counts in the same set of images. In Section D, we use the Full MIA and Full ecSeg models to predict ecDNA counts in other cell lines treated with JQ1 to evaluate the generalizability of these models to other cell lines.

likely represents the upper limit of achievable performance. In subsequent iterations, models were trained on a larger dataset comprising 1,164 images, which were refined by selecting ROIs and filtered to include only images with ranks of 3 and 4. After incorporating contour detection (discussed below) to address conjoined segmentations, the error rate for the resulting models, evaluated on a new set of 55 higher quality images (33 ranked 3+), was 7.62% (Figure 2E).

After achieving accuracy levels comparable to human annotations, we introduced an additional precision metric to further evaluate model performance. This metric involved comparing the predicted segmentation masks generated by the model to the ground truth masks (created using a pixel-patched approach to address class imbalance). Specifically, we assessed whether each unique object in the ground truth had at least one pixel overlap with a predicted mask, considering these overlapping objects as correctly identified. However, we refined the process to ensure that each object in the ground truth was exclusively matched to a single predicted object, preventing multiple predictions from being counted as overlaps with the same ground truth object. This approach helped us identify and balance false positives and negatives that might otherwise have been missed, thereby improving the overall quality of the training set.

Comparison to existing approaches

Previous approaches to automating ecDNA detection, such as ecDetect and ecSeg (Turner et al., 2017; Rajkumar et al., 2019), have made significant strides in addressing the challenges of accurately counting and localizing ecDNA within cell nuclei. ecDetect primarily relies on size thresholding to distinguish ecDNA from other nuclear components, applying a filter to isolate the region of interest (ROI) before analysis. However, this method can struggle with noise and closely clustered ecDNA, leading to potential inaccuracies in annotation and resolution (Figure 2F). ecSeg, on the other hand, utilizes a U-Net architecture with a ResNet50 backbone and focuses on segmenting images into patches for analysis. While effective, ecSeg was trained on a relatively small dataset of 483 annotated images, which may limit its generalizability and resolution at the chromosomal level.

Using the Microscopy Image Analyzer (MIA) tool builds on these previous methods by incorporating advanced image post-processing techniques that enhance annotation quality and improve resolution. Trained on nearly 1,300 high-quality images, MIA benefits from a larger and more diverse training dataset and introduces a refined segmentation process that accurately identifies and separates closely clustered ecDNA, resolving them even when

colocalized around or between chromosomes (Figure 2G). Unlike ecDetect, MIA's method does not rely heavily on size thresholding, allowing it to maintain high accuracy without filtering out subtle ecDNA signals.

Moreover, MIA's ability to classify chromosomes within the segmented ROIs provides an additional layer of precision, which is crucial for distinguishing between ecDNA and chromosomal artifacts. This capability directly impacts the overall accuracy of ecDNA detection, as it reduces false positives and negatives that arise from misclassification in complex cellular environments. By comparing the performance of MIA with ecDetect and ecSeg (Table 1, Sections A and B; Supplementary Table 1), we demonstrate significant improvements in both the accuracy of ecDNA identification and the ability to handle unannotated images.

Post-processing analyses

Contour detection

Contour detection is a technique used in image processing and computer vision to identify and delineate the boundaries or edges of objects within an image. A contour represents a curve that connects all the continuous points along a boundary that share the same color or intensity. In the context of image segmentation, contour detection is crucial for accurately separating and defining the shapes and structures of different objects within an image, such as nuclei, chromosomes, or ecDNA.

When dealing with conjoined segmentations—instances where multiple objects (e.g., closely clustered ecDNA) are touching or overlapping (Figure 2H), contour detection helps to identify the precise boundaries of each individual object. This allows for the separation of these conjoined objects into distinct segments, thereby improving the accuracy of the model in distinguishing between different ecDNA entities or other features in the image.

To enhance the accuracy of ecDNA detection, we applied contour detection as a critical post-processing step. By accurately delineating the boundaries of individual ecDNA within these complex regions, we were able to significantly improve the model's ability to distinguish between different ecDNA entities. The improvement after using contour detection underscores the importance of advanced post-processing techniques in refining the performance of computer vision models for complex biological image analysis.

Merging data from other fluorescent channels

Once MIA was executed to predict ecDNA in unannotated images, our next step was to identify the specific genes located on these predicted ecDNA, enabling us to count and differentiate between different ecDNA species and assess oncogene amplification levels. In our case, we labeled two oncogenes known to localize to ecDNA in each cell line: ERBB2 and MYC for NCI-H2170 cells, and FGFR2 and MYC for SNU16 and NCIH716 cells. Beyond using DAPI fluorescence, which provides a general overview of nucleic acids in our images, we employed FISH DNA probes that bind directly to genomic regions within each of these genes. To accurately map the predicted ecDNA locations with the actual locations identified through other channels (e.g., FITC and Texas Red), we developed a post-processing script. This script enabled the alignment of MIA-predicted ecDNA with the fluorescent signals from the gene-specific probes. When comparing the results of this script to human annotations, we found that the accuracy ranged within 97% ± 2% (Supplementary Figures 3–5; Supplementary Table 2), demonstrating the effectiveness of our approach (See Supplementary Methods for more details).

Applications in systems pharmacology: changes in ecDNA during drug treatment

Approximately 15% of tumors harbor ecDNA (Kim et al., 2020), yet research into the critical response mechanisms involving ecDNA remains limited, and only a few drugs have been thoroughly explored. Identifying treatments that specifically target cancer cells with ecDNA presents a significant challenge. To address this challenge, we developed CytoCellDB (Fessler et al., 2024), a comprehensive global data resource designed to streamline drug response evaluations. This platform allowed us to systematically assess drug efficacy across 139 cell lines harboring ecDNA and over 400 cell lines without ecDNA, providing a robust framework for efficiently identifying impactful treatments. Through this approach, we identified a cell line whose ecDNA had not yet been reported, that co-amplifies ERBB2 and MYC on the same ecDNA amplicon. The discovery of new ecDNA model systems highlight the potential of CytoCellDB to uncover new insights in ecDNA and cancer research.

Co-amplification of oncogenes on ecDNAs presents unique opportunities for diversification of gene and gene dosage combinations across cells in a population. As demonstrated by Figure 3A, there are three identified "species" of unique ecDNA and one unidentified species. These different ecDNA species localize different genes or gene combinations on the same fragment of ecDNA, "mix and matching" genes that are not commonly that close to one another in genomic space. For example, the gene MYC is encoded in chromosome 8 and the gene ERBB2 is encoded in chromosome 17. Suddenly, on ecDNA they become neighbors in the ecDNA species that co-localize these genes on the same amplicon. In NCI-H2170's case, we see ecDNA that localizes MYC and ERBB2 separately, and those that co-amplify these genes. We also see an "unknown" entity that amplifies genes other than MYC or ERBB2. These different species generate far more population genetic diversity than having one species alone. Increasing genetic population heterogeneity may be a strategy of these cancer cells to diversify their populations to increase the likelihood of generating advantageous traits to overcome drug resistance or changing environments, effectively "hedging their bets" (Figure 3B).

Currently, we lack a systematic method to quantify changes in ecDNA during treatment. Previous studies that have studied the effects of drugs on ecDNA have been fragmented and inconsistent, with significant variations in cell lines, drugs, dosages, and time points. This lack of standardization has led to an incomplete understanding of the dynamic processes involved in ecDNA drug response. Most research to date has primarily focused on the effects of chemotherapeutic agents (Haque et al., 2001; Storlazzi et al., 2010; Kitajima et al., 2001; van Leen et al., 2022), such as methotrexate, hydroxyurea, cisplatin, etoposide, doxorubicin, and vincristine, on DM reintegration. Interestingly, tyrosine kinase inhibitors (TKIs) have been shown to induce reversible reintegration of DMs (Nathanson et al., 2014), highlighting the complexity of the interactions between drugs, ecDNA and cellular adaptation strategies.

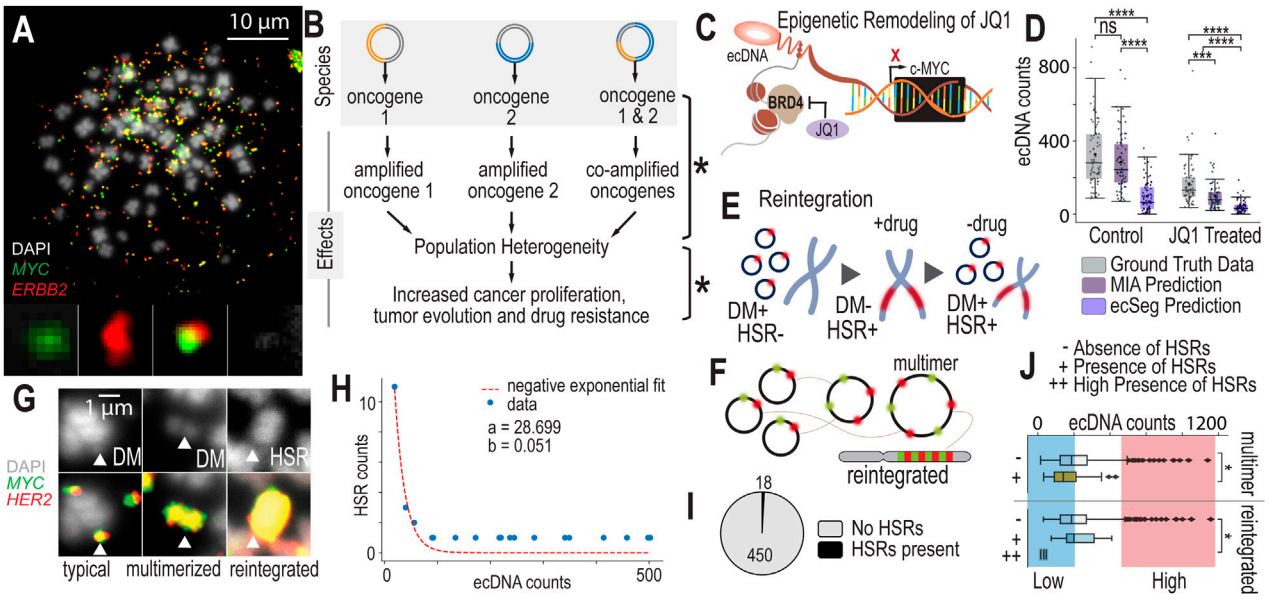

FIGURE 3
Scaling Data Analytics to Probe ecDNA-Mediated Drug Response Mechanisms **(A)**. Metaphase FISH image with dual probe labeling of NCI-H2170 cells that have four or more ecDNA species present in most of their cell nuclei. **(B)** Having multiple species of ecDNA present across cell nuclei can generate widespread cell-to-cell heterogeneity. As ecDNA segregates unevenly into daughter cells, each of these species will segregate unevenly. This generates extreme population genetic heterogeneity in terms of copy number variation differences across cells. This heterogeneity could lead to accelerated adaptation and drug resistance. **(C)** JQ1 is a molecule that was recently discovered to impact ecDNA higher-order clustering. It inhibits BRD4, impacting MYC transcriptional activity. Because MYC is often amplified on ecDNA, JQ1 may globally influence cells that harbor ecDNA-based MYC amplications. **(D)** Experimental results and AI-predicted ecDNA counts for MIA and ecSeg comparing ecDNA counts in control (untreated) and drug-treated (JQ1) NCI-H2170 cells. Significance was determined using a Mann-Whitney-Wilcoxon test with p-values < 0.0001 (****) and < 0.001 (***). **(E)** Under treatment, cells eliminate their ecDNA rapidly and reintegrate them into chromosomes, forming homogeneous staining regions (HSRs). **(F)** Drug treatment increases the number of multimerized ecDNAs that are seen in metaphase FISH images, which may be a precursor to chromosomal reintegration. **(G)** FISH data showing different structural views of ecDNA during drug treatment, including double minutes, multimerized ecDNA or "hubs" and HSRs. **(H)** A model predicting ecDNA reintegration rate during drug treatment. **(I)** The fraction of cells that were observed to have HSRs during metaphase FISH image analysis. **(J)** Cells with lower ecDNA counts are more likely to be observed to have HSRs compared to cells with high ecDNA levels. Significance was determined with a Chi-Square test with p-value < 0.05 (*). Similarly, cells with lower ecDNA levels are more likely to have multimerized structures compared to cells with high levels of ecDNA.

Recent studies suggest that epigenetic remodeling agents, such as JQ1, exert unique effects on cell lines harboring ecDNA. JQ1 and other epigenetic modulators alter chromatin accessibility, effectively modifying the "openness" or "closed-ness" of DNA regions, including those amplified on ecDNA. JQ1 specifically targets BRD4, a BET bromodomain protein known to enhance MYC transcription. BRD4 may also tether double minutes (DMs, a type of ecDNA) in transcriptional "hubs" (Zhou et al., 2020; Wu et al., 2019; Hung et al., 2021). By inhibiting BRD4, JQ1 not only disrupts ecDNA-specific transcription (Wu et al., 2019) but also impacts the broader transcriptional landscape, reducing the activity of oncogenes amplified on ecDNA, such as MYC and PVT1 (Figure3C). This dual impact of JQ1 on both transcription and chromatin organization highlights its potential as a therapeutic agent in ecDNA-driven cancers.

In addition to these transcriptional effects, JQ1 may disrupt ecDNA "hubs," the spatially organized groups of ecDNA molecules that cluster together. Such hubs are believed to be a critical adaptive mechanism in cancer cells that harbor ecDNA, enabling enhanced transcriptional output and rapid responses to selective pressures. The disruption of these hubs by JQ1 could undermine a key survival strategy in cancers harboring ecDNA, offering an exciting new avenue for drug development targeting these aggressive tumor phenotypes. Furthermore, JQ1's influence on chromatin accessibility may indirectly impact the segregation and stability of ecDNA during cell division, adding another layer of complexity to its effects.

As a pharmacology application, we leverage our AI-accelerated image annotation platform, using the Microscopy Image Analyzer (MIA), to study changes in ecDNA counts before and after drug treatment across three distinct cancer cell lines: a lung cancer cell line (NCI-H2170), a gastric cancer cell line (SNU16), and a colorectal cancer cell line (COLO320DM). Our goal is to apply our image processing pipeline and MIA to predict ecDNA counts in images from untreated, control (DMSO-only) and drug-treated samples, with a focus on assessing the impact of JQ1 treatment on ecDNA counts after 24 h. This approach aims to provide deeper insights into the effects of JQ1 on ecDNA copy number changes and its potential role in disrupting critical adaptive mechanisms in cancer cells.

Model performance for predicting ecDNA counts in JQ1-treated samples

The MIA model and the ecSeg model were evaluated for their ability to predict ecDNA counts in both untreated and JQ1-treated

samples. Our objective was to determine how accurately each model could quantify the reduction in ecDNA levels induced by JQ1 treatment, a BET inhibitor known to impact chromatin accessibility and gene expression.

The MIA model achieved reasonable accuracy in predicting ecDNA counts across both untreated and treated conditions. For untreated samples, the model obtained an accuracy of 0.719, with precision and recall values of 0.796 and 0.696, respectively, resulting in an F1 score of 0.74. However, when applied to JQ1-treated samples, the model's performance declined, reflecting the biological complexity of drug-induced changes. In treated samples, the MIA model reported an accuracy of 0.603, with precision at 0.777 and recall at 0.47, yielding an F1 score of 0.579. The drop in recall suggests that the model undercounted ecDNA, possibly due to the altered ecDNA patterns post-treatment.

The ecSeg model exhibited lower overall performance compared to the MIA model. For untreated samples, it achieved an accuracy of 0.413 with a precision of 0.505 and a recall of 0.164, indicating difficulty in detecting the full ecDNA population. In JQ1-treated samples, the ecSeg model's performance further declined, with an accuracy of 0.375 and a precision of 0.399, while recall dropped to 0.101, resulting in a negative Matthew's correlation coefficient (MCC) of −0.147. These metrics highlight the ecSeg model's limitations in adapting to ecDNA changes following drug treatment.

Model comparison and implications

The performance of both models underscores the challenges in predicting ecDNA counts under drug treatment conditions. While the MIA model demonstrated moderate generalizability, achieving reasonable performance across untreated and treated samples, the ecSeg model struggled with both untreated and JQ1-treated datasets. The higher precision but lower recall in JQ1-treated samples for both models suggests that while some ecDNA features were accurately identified, a significant portion of ecDNA was missed, possibly reflecting altered ecDNA morphology or reintegration patterns in response to the drug.

These findings indicate that while our models, especially MIA, offer potential for predicting ecDNA behavior, further model refinement is needed to improve recall and robustness under varying biological conditions. The results also demonstrate the need to incorporate additional training data reflecting post-treatment changes to better capture the dynamics of ecDNA in response to drugs like JQ1 (Table 1, Sections A and B).

Cells rapidly eliminate ecDNA after 24 h of JQ1 treatment

Comparing ecDNA counts before and after 24 h of JQ1 treatment, we observe that cells rapidly eliminate their ecDNAs. In NCI-H2170 cells, the mean ecDNA count in untreated wild-type cells is 330 ± 167, while after JQ1 treatment, the mean ecDNA count drops to 220 ± 132 (Figure 3D). Similarly, in COLO320 cells, the mean ecDNA count decreases from 50 ± 37 in wild-type cells to 45 ± 30 post-treatment. In SNU16 cells, the mean ecDNA count shifts from 370 ± 203 in wild-type cells to 298 ± 156 following JQ1 exposure. Using a Wilcoxon rank-sum test, we find that the difference in ecDNA counts between untreated and JQ1-treated populations is highly significant across all 3 cell lines. We see a similar trend between DMSO (vehicle-only control) and JQ1-treated (Supplementary Figure 6). This suggests that JQ1 induces a similar drug response mechanism as other chemotherapeutic agents, where cells rapidly eliminate their ecDNA, possibly through the selection of less sensitive clones, cells with fewer ecDNAs, or active elimination processes. Remarkably, the distributions of all ecDNA species—whether containing only MYC, only ERBB2, or both MYC and ERBB2—show significant differences before and after drug treatment (Supplementary Figures 7–10).

Cells reintegrate ecDNAs into chromosomal homogeneous staining regions (HSRs)

During drug treatment, cells reintegrate ecDNA into chromosomes, forming HSRs, or regions of highly amplified genes within chromosomes (Figure 3E). In this way, ecDNA, which is extrachromosomal and independent of chromosomes, shapeshifts into chromosomal DNA. As a field, we still do not definitively understand how this phenomena increases cell fitness or survival during drug treatment. The overarching opinion is that cells "hide" and "store" their ecDNA in HSRs when it will serve them better in future conditions. Furthermore, ecDNAs form higher-order structures in which it appears that they multimerize and combine into larger ecDNA entities (Figure 3F). Multimerization may occur prior to reintegration; however, we currently lack sufficient evidence to confirm this process definitively (Figure 3G).

Comparing the number of HSR counts in NCI-H2170 cells before and after JQ1 treatment, we find that cells reintegrate ecDNAs at a rate of approximately 5% during drug exposure (Figure 3H). In this cell line, no HSRs are observed in untreated cells that contain the same set of genes as the ecDNAs, indicating that any HSR detected in the drug-treated samples containing MYC or ERBB2 represents a reintegrated ecDNA. After 24 h of drug treatment, we observe 18 reintegration events across 450 cells (Figure 3I), confirming that ecDNA reintegration occurs as part of the cellular response to the drug.

ecDNA reintegration mostly occurs in cells with fewer ecDNA counts

Reintegration events are more likely to occur in cells that have lower numbers of ecDNA. Analysis of over 600 metaphase FISH images from NCI-H2170 cells showed that DMs reintegrate into HSRs more frequently in cells with fewer DMs (p = 0.027, chi-squared test), consistent with findings from previous studies (Haque et al., 2001) (Figure 3J). One possible explanation for this observation is that cells with fewer ecDNAs may face lower cellular allocation costs, allowing them to invest more resources into processes such as reintegration. Alternatively, these cells may have different cell cycle distributions, potentially spending more time in specific phases that favor reintegration events. This could enable a higher frequency of HSR formation in cells with fewer ecDNAs, offering them a survival advantage under drug treatment conditions.

This analysis provides a powerful framework for quantifying how JQ1 influences ecDNA elimination and reintegration, opening new avenues to explore whether its effects are mediated through direct disruption of ecDNA hubs, transcriptional regulation, or broader processes like cell cycle dynamics. By systematically measuring changes in ecDNA, this work lays the foundation for

future studies to unravel the complex mechanisms underlying JQ1's impact on ecDNA behavior.

Model limitations

While the Microscopy Image Analyzer (MIA) offers significant advantages in automating the detection of ecDNA in metaphase FISH images, it has several limitations that users should consider.

Algorithm-specific limitations

One key limitation is that the algorithm has been trained exclusively on metaphase FISH images. As such, applying MIA to other types of images (e.g., interphase cells or tissue sections) would likely require retraining the model on an appropriate dataset. Additionally, the detection of ecDNA is dependent on size thresholds inherent to the algorithm's training. While it is difficult to provide an exact size threshold in terms of spatial resolution, the ecDNA analyzed in the training dataset typically ranged from 1 to 4 Mb in DNA content or approximately 0.25 μm in size (Supplementary Figure 2). Further research is needed to clarify the minimum detectable size for ecDNA and how this correlates with biological and technical resolution limits. Finally, MIA was trained on metaphase FISH images captured at ×60 magnification, and its performance is optimal with images of similar resolution. Deviations in magnification (e.g., ×20, ×40, or ×100) or image resolution may reduce accuracy or limit the ability to analyze full metaphase spreads efficiently.

Biological considerations

The necessity for metaphase spreads limits the application of this algorithm to cell line models and samples where such spreads can be reliably generated. This may preclude its immediate use in primary patient samples or tissues that cannot be processed into metaphase spreads. Additionally, the interpretation of ecDNA distributions under JQ1 treatment warrants careful consideration. While our data suggest that JQ1 influences the population of ecDNA-positive cells entering metaphase, it is possible that JQ1 does not directly alter ecDNA itself but rather affects the subset of cells capable of entering mitosis or the chromosomal context of ecDNA and HSRs. This underscores the need for further studies to disentangle direct effects on ecDNA from secondary effects mediated through cell cycle regulation.

Technical limitations

From a software perspective, MIA is restricted to a graphical user interface (GUI) and lacks command-line functionality, preventing it from running in a "headless" mode. This limits its integration into larger, script-based workflows for high-throughput analyses. The software also requires substantial memory resources, which can be a constraint when analyzing large datasets or using less powerful machines. Additionally, MIA does not allow direct processing from master folders, requiring manual file management and input, which can be cumbersome for large-scale analyses.

A detailed guide on preparing, ranking, and analyzing metaphase FISH images with MIA is provided in the Supplementary Material. This documentation includes step-by-step instructions on image preparation, ranking criteria, and running the software for optimal results. These limitations highlight opportunities for future improvement in both the biological scope and technical functionality of MIA to enhance its applicability in ecDNA research.

Conclusion

In this study, we successfully automated the detection of extrachromosomal DNA (ecDNA) using the Microscopy Image Analyzer (MIA), providing a robust and accurate framework for analyzing metaphase FISH images. Through the development and application of our custom pipeline, we enhanced MIA's capabilities specifically for ecDNA detection and quantification, establishing it as a powerful tool for investigating the effects of pharmacological treatments on cancer cells.

Our approach has yielded new insights into cellular responses during JQ1 treatment, revealing significant shifts in ecDNA counts, patterns of ecDNA reintegration into chromosomes, and the tendency of certain cells within the population to favor ecDNA reintegration. These findings underscore the potential of MIA-based analysis to advance our understanding of how cancer cells adapt to therapeutic pressures, offering valuable perspectives for drug discovery and cancer treatment strategies.

Additionally, we observed promising signs of generalizability. Models trained primarily on images from 1 cell line could predict ecDNA patterns in another cell line with moderate success, indicating that the framework holds potential for cross-cell-line applications. Similarly, models trained on images of untreated cells were able to predict ecDNA behavior in drug-treated cells with reasonable accuracy, demonstrating that the training process is sufficiently robust to generalize across treatment conditions. These findings suggest that the system is flexible and capable of adapting to varying biological contexts, though further validation will be needed to fully assess the extent of this generalizability.

Methods

Primary cell culture

NCI-H2170, SNU16, and COLO320DM and NCIH716 cells were purchased from ATCC and were grown in RPMI media (Gibco) supplemented with 10% heat-inactivated fetal bovine serum (Gibco) in a humidified incubator with 5% CO_2. SUM159PT cells were given as a gift from Gary Johnson's laboratory (UNC Chapel Hill) and were STR verified and checked for *mycoplasma*. SUM159PT cells were grown in Ham's F-12 media supplemented with 5% heat inactivated fetal bovine serum along with 10 mM HEPES, 1ug/mL Hydrocortisone, and 5ug/mL Insulin. Cells were harvested with 0.25% Trypsin in DPBS (Gibco). Viable cells were counted using a Countess 3 (Invitrogen) counter and trypan blue (Invitrogen, T10282). Cells were collected for metaphase and karyotype (G-banding) experiments within 3 passages.

Drug treatments

For NCI-H2170 cells, an xCELLigence Real-Time Cell Analyzer (Agilent) was used to monitor cell activity over time in the presence and absence of JQ1, recording sweeps every 15 min using RTCA Software Pro 2.6.0. The Agilent SP station was placed in a humidified CO2 incubator. Background readings were recorded using 80 µL of media alone in a 96-well PET E-plate. NCI-H2170 cells were seeded at a concentration of 20,000 cells per well in 100 µL of media. After a 24-h attachment period, cells were treated with JQ1 at six different doses ranging from 0 µM to 100 µM, with DMSO as the vehicle control. Impedance-based measurements were collected at 15-min intervals and converted into Cell Index (CI) values, which reflect changes in cell viability, adherence, and morphology. For each JQ1 concentration, CI values were normalized to the initial pre-treatment measurements, and a dose-response curve was generated by plotting normalized CI values against JQ1 concentrations at the 24-h endpoint. The IC50 dose, calculated to be 36.48 µM, was determined by fitting the data to a nonlinear regression model {log [(JQ1)] vs. response} using GraphPad Prism, identifying the concentration of JQ1 that reduced the Cell Index by 50%. Cells were treated at this concentration and collected at 24 h for metaphase FISH experiments.

For semi-adherent and suspension cell lines (SNU16 and COLO320DM), we determined IC50 values using CellTiter-Glo (Promega). SNU16 cells were treated with six dilutions of JQ1 (0–10 µM) with DMSO as the vehicle control, and growth was measured using luminescent signal intensities proportional to ATP levels. For COLO320DM cells, 10 dilutions of JQ1 (0–50 µM) were tested. For each cell line, a dose-response curve was generated by plotting relative luminescence against the JQ1 concentration. The IC50 values were determined by fitting the curve to a nonlinear regression model [log (inhibitor) vs. response] using GraphPad Prism, identifying the concentration of JQ1 at which 50% of cell viability was inhibited. The IC50 doses were calculated to be 14.67 µM and 32.80 µM for SNU16 and COLO320DM cells, respectively.

Metaphase sample preparation and imaging

Generating condensed chromatin during metaphase allows for optimal imaging of ecDNA. Before karyotyping, cells underwent a four-stage preparation: arrest at metaphase, incubation in a hypotonic solution, cell fixation, and staining. Samples were prepared using cells cultured by the Brunk Lab.

Cells were arrested at metaphase by treating with colcemid at 0.1 µg/mL (10 µg/mL Colcemid Solution, FUJIFILM Irvine Scientific) in cell culture media when cells reached ~70% confluency. Colcemid arrests cellular division during mitosis by binding to tubulin, preventing spindle formation and cytokinesis. Cells were incubated with colcemid for 12–20 h before being harvested following standard cell culture procedures. For adherent or semi-adherent cells, trypsinization was used to detach cells, and this was quenched with a mixture of cold colcemid-spiked media and PBS wash to maximize yield. Cells were resuspended in 1 mL of 1x PBS by pipetting and transferred to 1.5 mL microcentrifuge tubes for centrifugation at 5,000 rpm for 2 min.

The cells were incubated with 600 µL of pre-warmed 37°C 0.075M KCl (Gibco), added dropwise with gentle agitation to resuspend cells. After 15 min at 37°C, the cells became swollen and fragile due to osmotic pressure, making them ready for fixation. Freshly prepared modified Carnoy's fixative (3:1 methanol:glacial acetic acid) was added dropwise to each sample. Tubes were immediately centrifuged at 5,000 rpm for 2 min. After leaving ~150 µL of supernatant, pellets were gently agitated and resuspended. Another 600 µL of fixative was added dropwise, followed by agitation and centrifugation for 2 min at 5,000 rpm. This fixation step was repeated three times, with the final addition of fixative adjusted to achieve ~6 million cells/mL (0–1 mL), ensuring optimal density for single-cell imaging.

Microscope slides were prepared using **Superfrost™ Microscope Slides** (Fisherbrand, Cat. No: 12550123), which are uncharged. Slides were humidified using water vapor immediately before a drop (10 µL) of the prepared cell suspension was dropped from a height of ~60–70 cm onto the slide. The slides were left to air dry for an hour. Dried slides were equilibrated in ×2 saline sodium citrate (SSC) (Ultrapure™ 20X SSC buffer, Invitrogen, Cat. No: 15557-036) and dehydrated through an ascending ethanol series (70%, 85%, 100%) for 2 min each. 5 µL of fluorescent DNA probes (Empire Genomics) were applied to the center of a metaphase spread. A 22 × 22 mm cover glass (Electron Microscopy Sciences, Cat. No: 72210-10) was placed over the slide. Slides were denatured at 70°C for 2 min and incubated at 37°C overnight (16–20 h) in a humidified slide moat.

Post-hybridization, slides were washed sequentially in 0.4× SSC and 2× SSC + 0.05% Tween-20 for 2 min each, followed by a final dip in 2× SSC. After removing residual washing buffer, **SlowFade™ Diamond Antifade Mountant with DAPI** (Invitrogen, Cat. No: S36964) was added to the center of the slide. A 24 × 60 mm microscope cover glass (Fisherbrand, Cat. No: 12541036) was placed on top and sealed with nail polish.

All images were captured using an **Echo Revolution Microscope** (Echo, San Diego, CA) at **×60 magnification**. Images were taken from the same slide or occasionally from two slides prepared from the same metaphase spread to ensure consistency in experimental analysis. While uncharged slides were used for all metaphase spreads, potential differences in ecDNA adherence between charged and uncharged slides were not specifically tested.

Image ranking

Images were evaluated on a 0 to 4 scale, with 0 indicating the lowest confidence in image quality and 4 representing the highest. The ranking process considered factors such as image resolution, clarity, and human oversight. The initial assessment focused on the content and visual quality. Resolution was reviewed to identify blurriness, out-of-focus chromosomes, or artifacts. This preliminary review set a ceiling for the final score; for instance, if the resolution was insufficient to clearly resolve individual chromosomes or ecDNA, the image could receive a maximum score between 0 and 2. In contrast, images with sharp features that clearly resolved chromosomes and ecDNA could achieve a higher ranking of 3 or 4.

Another factor considered was the distribution of material around nuclei, referred to here as the nuclear splash pattern. In some cases, nuclear material spreads unevenly, influenced by factors like nearby artifacts or popped nuclei, which disrupt the typical elliptical shape of a nucleus. The rank was reduced only if these artifacts caused significant overlap between regions, making it difficult to confidently distinguish boundaries. Images with highly crowded fields or poorly defined ROI boundaries were also downgraded, as such conditions made consistent annotation across different users more challenging.

After evaluating image quality, the annotations were reviewed for accuracy. Merged images with multiple probe layers were separated to isolate the DAPI stain, and the original annotations were re-imported. To enhance visibility of small ecDNA, the contrast was increased. Annotations were then verified based on their proximity to the original markings. If most annotations aligned within a few pixels of the original, the image was assigned a rank of 4. However, the rank was lowered if significant errors were detected. False positives were noted when annotations overlapped with chromosomes, debris (Supplementary Figure 1) or debris clouds containing probe molecules, or when ecDNA was duplicated due to layer misalignment or indicator splitting. False negatives occurred when dim ecDNA was missed, when ecDNA was too close to chromosomes or artifacts, or when clusters of ecDNA were undercounted. If such errors were prevalent, the image was downgraded to a rank of 1. In total, we had 183 images marked as 0, 290 images marked as 1, 700 images marked as 2, 970 images marked as 3, and 441 images marked as 4.

Finally, each image was assigned a final rank, and detailed notes were recorded. These notes included suggestions for improving annotations to increase the image's rank, as well as any observations—such as uneven brightness or scaling—that could explain deviations in accuracy across images but were not directly captured by the ranking criteria.

ROI segmentation approach

Images were isolated to the single-nucleus level by cytogeneticists, technicians, or scientists who defined the perimeter of the region of interest. This was accomplished using the lasso tool in the FIJI software (Schindelin J. et al., 2012; Schneider C. et al., 2012) to create a binary image representing the region of interest (ROI) for each nucleus. Masks were then generated for each image to tightly constrain the focus to the annotated regions and chromosomes. This step aimed to minimize the inclusion of unannotated data that could interfere with model training. However, in some cases, artifacts such as unpopped nuclei could not be fully excluded when defining the ROI. These artifacts were allowed to remain in the final ROI to provide counterexamples for the model, which could enhance its robustness.

Expanding pixels around annotations

During the original annotation phase, the FIJI counting tool was used by multiple different operators/technicians/scientists to document the count. This tool will give a total count of the annotations made to the image and tag a single pixel the location. Each image was 2048 by 2448 pixels and could have up to a thousand annotations. Due to the large class imbalance each annotation was expanded using a python script to a 7 pixel tall 7 pixel wide diamond patch to decrease the class imbalance and build in padding for each annotation not being exactly in the center of each ecDNA. The pixel patch was transferred as a perimeter into a single.npz file that would translate to a segmentation by MIA.

Performing analysis with MIA

Within MIA, each image was dissected into 224 × 224 pixel patches and randomly augmented with any affine transformation. The patches were then fed into a U-Net with a Res-Net-50 encoder neural network. The model used a focal conservative loss with a class weighing of 0.82. The batch size (number of images per model update) was set at 8, and the learning rate was set 0.001 with a 0.1 reduction factor for loss plateaus over three epochs until 1.00 * (10^{-6}). Contours were not predicted during the training due to memory limitations, but were optimized to match model training data ecDNA counts in post-processing. These hyperparameters were tuned based on the minimum loss and computational resources over 50 epochs of a reduced dataset.

For training of large models, the model was trained for between 65 and 75 epochs based on the final plateau of loss. Contours were separated in the masks to segregate overlapping ecDNA based on ecDNA count. The segmentation masks were then exported from the.npz coordinates into a.tif image.

A detailed guide on how to prepare, rank, and analyze images using the MIA with metaphase FISH images is provided in the Supplementary Material. This documentation includes step-by-step instructions on image preparation, ranking criteria, and running MIA for optimal results.

Accessing accuracy of predictions

To assess the model's performance, we used a standard accuracy metric defined as the difference between the total predicted count and the total ground truth count, normalized by the ground truth total. This metric provides a measure of the model's deviation from the true count. The majority of the observed error resulted from the model underestimating the true count, indicating a tendency toward undercounting.

Error metric and model evaluation

To evaluate the model's performance, we calculated the Mean Absolute Error (MAE) between the ground truth ecDNA count and the predicted ecDNA count. The MAE is defined as:

$$MAE = (1/n) \sum |Predicted\,Count_i - Ground\,Truth\,Count_i|$$

This metric provides an average of the absolute differences between predicted and actual counts, offering a straightforward measure of the model's accuracy. We selected the MAE because it reflects the magnitude of the prediction error without being influenced by the direction of the error (over- or underestimation).

Precision

Two measures of precision were calculated: location-based precision and count-based precision.

Location-based precision

This metric evaluated precision at the pixel level by determining whether individual pixels from the model's segmentation output overlapped with the ground truth map. Each pixel present in both the ground truth pixel patch and MIA's segmentation output was counted as a true positive. Conversely, pixels present in the model's output but not in the ground truth were counted as false positives. In this method, each pixel was treated as an individual label, providing a highly granular precision measure.

Count-based precision

This measure compared the overlap between predicted segmentation labels and the corresponding ground truth pixel patches. If at least one pixel of a predicted segment overlapped with the corresponding ground truth patch, the segment was considered a true positive. To avoid overestimation, any predicted segment overlapping multiple ground truth patches was counted only once as a true positive. The total number of overlapping segments was divided by the total ground truth segment count to yield the precision score.

Any difference between the total number of predicted segments and the overlapping true positives was recorded as the false positive count.

Object-oriented model performance evaluation

One approach to assess model performance was object-oriented evaluation. In this method, the true positives (TPs) were calculated by counting the number of ground truth (GT) objects that overlapped with at least one predicted object, while subtracting the number of predicted objects that overlapped with more than 1 GT object. This ensures that each true positive consists of a unique pairing between one predicted object and 1 GT object, avoiding overcounting.

The false positives (FPs) were determined by subtracting the total number of true positives from the total number of predicted objects. Similarly, the false negatives (FNs) were computed by subtracting the number of true positives from the total number of GT objects.

The true negatives (TNs) were identified by counting the pixels labeled as zero in both the prediction masks and the GT labels. To adjust for the disproportionate number of background pixels compared to ecDNA pixels, we scaled the number of true negatives by dividing it by the ratio of background pixels to ecDNA pixels in the GT for each image. This result was further divided by 25, representing the size of each GT object, to ensure that the number of true negatives was comparable to the number of true positives.

This scaling ensures that the accuracy metric reflects the importance of true positives and is not disproportionately influenced by the abundance of true negatives in the data.

High-confidence image set and model performance

To estimate the highest achievable accuracy, a curated set of 194 high-confidence images was used. These images were stained exclusively with a DAPI probe, and their annotations were carefully reviewed and corrected. The regions of interest (ROIs) were meticulously defined to minimize noise and extraneous content while ensuring that all annotated ecDNA remained in view. Special attention was given to the ROI edges to avoid abrupt cuts into the bright fields surrounding objects. Any images with artifacts or blur within the ROI were excluded from the dataset.

The dataset was split into a training set of 131 images and a test set of 63 images, selected randomly. The model was trained using the same procedures and parameters as the main model (ensure that this is referenced if detailed elsewhere), for a total of 60 epochs.

Model performance on the validation set resulted in a mean absolute error (MAE) of 6.84%, while the training set predictions achieved a MAE of 5.14%.

Estimating user-to-user variability

Given that the data inherently reflects biases introduced by the human operators (technicians/scientists), we aimed to estimate the extent of user-to-user variability. To do this, we conducted a test with 11 images, where multiple operators independently annotated and counted the same images in both the DAPI-only layer and the merged images.

The analysis revealed an average error of approximately 8.4% for the DAPI-only images, while the merged images exhibited a higher variance, with an average error of 12.0% between users.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found in the article/Supplementary Material.

Ethics statement

Ethical approval was not required for the studies on humans in accordance with the local legislation and institutional requirements because only commercially available established cell lines were used. Ethical approval was not required for the studies on animals in accordance with the local legislation and institutional requirements because only commercially available established cell lines were used.

Author contributions

KG: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing–original draft. AM: Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Writing–original draft. DG: Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Writing–original draft. JF: Investigation, Validation, Writing–review and editing. JC: Data curation, Formal Analysis, Visualization, Writing–review and editing. WN: Formal Analysis, Software, Writing–review and editing. CF: Data curation, Methodology, Writing review and editing. OC: Data curation, Formal Analysis, Writing–review and editing. DC: Formal Analysis, Writing–review and editing. WD: Data curation, Formal Analysis, Writing–review and editing. NG: Data curation, Formal Analysis, Writing–review and editing. YW: Data curation, Writing–review and editing. KS: Formal Analysis, Writing–review and editing. EB: Conceptualization, Data curation, Formal Analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing–original draft, Writing–review and editing.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. The authors gratefully acknowledge the funding support provided by the IBM Junior Faculty Development Award from the University of North Carolina at Chapel Hill (to EB), the Computational Medicine Pilot Award (to EB) and the NIGMS Grant 2R25GM089569-14 (to DG).

Acknowledgments

We extend our appreciation to the Renaissance Computing Institute (RENCI) at the University of North Carolina at Chapel Hill, with special thanks to Dr. Matt Suttusky, Prof. Ashok Krishnamurthy for their valuable insights and time spent discussing key concepts and theories. We also appreciate the discussions with Prof. Marc Niethammer and Qin Lu on computer vision.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that Generative AI was used in the creation of this manuscript. In accordance with NIH guidelines, Generative AI was used solely as a copy-editing tool in the preparation of this manuscript.

Publisher's note

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fphar.2024.1516621/full#supplementary-material

References

Barretina, J., Caponigro, G., Stransky, N., Venkatesan, K., Margolin, A. A., Kim, S., et al. (2012). The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature* 483, 603–607. doi:10.1038/nature11003

Deshpande, V., Luebeck, J., Nguyen, N.-P. D., Bakhtiari, M., Turner, K. M., Schwab, R., et al. (2019). Exploring the landscape of focal amplifications in cancer using AmpliconArchitect. *Nat. Commun.* 10, 392–414. doi:10.1038/s41467-018-08200-y

Fessler, J., Ting, S., Yi, H., Haase, S., Chen, J., Gulec, S., et al. (2024). CytoCellDB: a comprehensive resource for exploring extrachromosomal DNA in cancer cell lines. *Nar. Cancer* 6, zcae035. doi:10.1093/narcan/zcae035

Haque, M. M., Hirano, T., Nakamura, H., and Utiyama, H. (2001). Granulocytic differentiation of HL-60 cells, both spontaneous and drug-induced, might require loss of extrachromosomal DNA encoding a gene(s) not c-MYC. *Biochem. Biophys. Res. Commun.* 288, 586–591. doi:10.1006/bbrc.2001.5798

Hung, K. L., Yost, K. E., Xie, L., Shi, Q., Helmsauer, K., Luebeck, J., et al. (2021). ecDNA hubs drive cooperative intermolecular oncogene expression. *Nature* 600, 731–736. doi:10.1038/s41586-021-04116-8

Kim, H., Nguyen, N.-P., Turner, K., Wu, S., Gujar, A. D., Luebeck, J., et al. (2020). Extrachromosomal DNA is associated with oncogene amplification and poor outcome across multiple cancers. *Nat. Genet.* 52, 891–897. doi:10.1038/s41588-020-0678-2

Kitajima, K., Haque, M., Nakamura, H., Hirano, T., and Utiyama, H. (2001). Loss of irreversibility of granulocytic differentiation induced by dimethyl sulfoxide in HL-60 sublines with a homogeneously staining region. *Biochem. Biophys. Res. Commun.* 288, 1182–1187. doi:10.1006/bbrc.2001.5892

Labrie, M., Brugge, J. S., Mills, G. B., and Zervantonakis, I. K. (2022). Therapy resistance: opportunities created by adaptive responses to targeted therapies in cancer. *Nat. Rev. Cancer* 22, 323–339. doi:10.1038/s41568-022-00454-5

Madren, J. A., Chen, J., Dennis, W., Ford, C., K White, K., and Brunk, E. (2024). A standardized protocol for sample preparation for scanning electron microscopy to visualize extrachromosomal DNA. *Biotechniques* 76, 311–321. doi:10.1080/07366205.2024.2346042

Nathanson, D. A., Gini, B., Mottahedeh, J., Visnyei, K., Koga, T., Gomez, G., et al. (2014). Targeted therapy resistance mediated by dynamic regulation of extrachromosomal mutant EGFR DNA. *Science* 343, 72–76. doi:10.1126/science.1241328

Rajkumar, U., Turner, K., Luebeck, J., Deshpande, V., Chandraker, M., Mischel, P., et al. (2019). EcSeg: semantic segmentation of metaphase images containing extrachromosomal DNA. *iScience* 21, 428–435. doi:10.1016/j.isci.2019.10.035

Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., et al. (2012). Fiji: an open-source platform for biological-image analysis. *Nat. Methods* 9, 676–682. doi:10.1038/nmeth.2019

Schneider, C., Rasband, W., and Eliceiri, K. (2012). NIH Image to ImageJ: 25 years of image analysis. *Nat. Methods* 9, 671–675. doi:10.1038/nmeth.2089

Storlazzi, C. T., Lonoce, A., Guastadisegni, M. C., Trombetta, D., D'Addabbo, P., Daniele, G., et al. (2010). Gene amplification as double minutes or homogeneously

staining regions in solid tumors: origin and structure. *Genome Res.* 20, 1198–1206. doi:10.1101/gr.106252.110

Tsherniak, A., Vazquez, F., Montgomery, P. G., Weir, B. A., Kryukov, G., Cowley, G. S., et al. (2017). Defining a cancer dependency map. *Cell* 170, 564–576.e16. doi:10.1016/j.cell.2017.06.010

Turner, K. M., Deshpande, V., Beyter, D., Koga, T., Rusert, J., Lee, C., et al. (2017). Extrachromosomal oncogene amplification drives tumour evolution and genetic heterogeneity. *Nature* 543, 122–125. doi:10.1038/nature21356

van Leen, E., Brückner, L., and Henssen, A. G. (2022). The genomic and spatial mobility of extrachromosomal DNA and its implications for cancer therapy. *Nat. Genet.* 54, 107–114. doi:10.1038/s41588-021-01000-z

Wu, S., Turner, K. M., Nguyen, N., Raviram, R., Erb, M., Santini, J., et al. (2019). Circular ecDNA promotes accessible chromatin and high oncogene expression. *Nature* 575, 699–703. doi:10.1038/s41586-019-1763-5

Yang, M., Zhang, S., Jiang, R., Chen, S., and Huang, M. (2023). Circlehunter: a tool to identify extrachromosomal circular DNA from ATAC-Seq data. *Oncogenesis* 12, 28. doi:10.1038/s41389-023-00476-0

Zhou, S., Zhang, S., Wang, L., Huang, S., Yuan, Y., Yang, J., et al. (2020). BET protein inhibitor JQ1 downregulates chromatin accessibility and suppresses metastasis of gastric cancer via inactivating RUNX2/NID1 signaling. *Oncogenesis* 9, 33–14. doi:10.1038/s41389-020-0218-z