



Cite this: DOI: 10.1039/d5lc00738k

MaGIC-OT: an AI-guided optical tweezers platform for autonomous single-cell isolation in microfluidic devices

 Jan-Philipp Cieslik, ^{ac} Xiaoye Xia^b and Ali Salehi-Reyhani ^{*ab}

Automating the isolation of rare cells such as circulating tumour cells (CTCs) within crowded microfluidic environments remains a bottleneck in liquid biopsy workflows. Optical tweezers offer contact-free, selective manipulation but traditionally rely on expert operators. We present MaGIC-OT (machine-guided isolation of cells using optical tweezers), a platform that integrates classical path planning and deep reinforcement learning (DRL) to automate single-cell manipulation inside a microfluidic chip. We built a high-fidelity simulation to train and benchmark control policies and show that cooperative, human-in-the-loop training improves DRL performance. Trained agents outperform expert users in speed and isolation success *in silico*, and we demonstrate proof-of-concept isolation of a cancer cell from a spiked blood sample on-chip. MaGIC-OT provides a flexible framework for intelligent optical manipulation, aligning microfluidic device design with autonomous control strategies and offering a pathway toward high-purity, label-free single-cell workflows.

 Received 25th July 2025,
 Accepted 3rd January 2026

DOI: 10.1039/d5lc00738k

rsc.li/loc

Introduction

Liquid biopsies have transformed how we monitor tumour evolution and treatment response by enabling minimally invasive, longitudinal sampling of circulating tumour cells (CTCs). Yet, the extreme rarity of CTCs in peripheral blood – only ~20% of patients with primary breast cancer present ≥ 1 CTC per 7.5 mL, with a median of ~3 CTCs per 7.5 mL in metastatic disease – remains a major bottleneck for high-throughput, information-rich single-cell studies.^{1,2} This scarcity demands instrumentation that can both sensitively detect and selectively isolate individual viable cells from dense backgrounds of blood cells and debris.³ Size and deformability-based platforms, including deterministic lateral displacement (DLD),⁴ inertial/Dean-flow focusing,⁵ viscoelastic migration,⁶ and microfiltration/constrictions,⁷ offer label-free enrichment at moderate to high throughput, but select primarily on biophysical properties yet typically result in high residual leukocyte carry-over, limiting single-cell purity without multiple enrichment stages. Hydrodynamic traps,⁸ microwells,^{9,10} and pneumatically valved arrays¹¹ deterministically capture single cells for analysis, yet require prior enrichment to achieve purity. Droplet microfluidics excels at partitioning and downstream

assays but is not, by itself, a selective isolation step.¹² Field-based methods exploit intrinsic physical properties of cells. Bulk and surface-acoustic-wave acoustophoresis separate by compressibility/density with excellent viability but struggle with phenotypic specificity.¹³ Dielectrophoresis (DEP) discriminates by complex polarisability and can capture and route single cells from mixtures deposited on microelectrode arrays.¹⁴ Practical constraints include the need for low-conductivity media (~10–100 mS m⁻¹) to limit Joule heating and can struggle with live cell isolation when operating in physiological buffers (>1 S m⁻¹).¹⁵ For cell isolation, the microelectrode grid array ‘quantises’ motion by limiting the translation of cells in orthogonal directions along the rectangular pixel array; this reduces operational loading densities and throughput is ultimately constrained by imaging and route planning and the need to avoid cell contact during transfers. Optoelectronic/optically induced DEP (OET/ODEP) uses a photoconductive substrate and light-defined virtual electrodes, obviating fixed metal microelectrode grids and enabling parallel, light-addressed traps.¹⁶ However, ODEP share the same buffer constraints as classical DEP. Current commercial solutions span microelectric field-based systems (*e.g.* DEPArray) and micromanipulation platforms (*e.g.* CellCelector), which can achieve high purity but often at the cost of manual intervention, protracted run times, and limited scalability.^{17,18}

Optical trapping (optical tweezers) offers a non-contact, label-free and spatially precise modality for manipulating microscopic objects, from single molecules to entire cells.¹⁹ In cellular systems, single-beam tweezers can position and

^a Department of Surgery & Cancer, Imperial College London, London, W12 0HS, UK. E-mail: ali.salehi-reyhani@imperial.ac.uk

^b Institute for Molecular Science and Engineering, Imperial College London, SW7 2AZ, UK

^c Dept. Obstetrics & Gynaecology, University Hospital and Medical Faculty of the Heinrich-Heine University Düsseldorf, Düsseldorf, Germany



sort individual cells with sub-micron accuracy and preserve viability.^{20–25} Optical tweezers (OT) exert force *via* a tightly focussed near-infrared beam, requiring no electrodes, photoconductors or low conductivity media. OT operates robustly in physiological buffers and supports continuous, free-space trajectories through crowded environments, making them intrinsically attractive for isolating rare targets such as CTCs. Despite their precision, optical tweezers remain predominantly manual tools, requiring significant operator expertise that constrains throughput, hinders standardisation, and limits broader translational adoption.

To improve throughput and usability, optical tweezers have been integrated with microfluidic architectures that simplify optical-based sorting. Transiting cells can be directed into desired microfluidic channel by calculating a trapping trajectory to displace cells.²⁶ In micromanipulation and micro-electric fields the transport of a cell to its destination is relatively straightforward. Either no obstacles must be avoided, or isolation is conducted on microelectrode grids which drastically simplify path planning. Optical trapping requires no on-chip functionalisation and no patterning of electrodes but does pose challenges to automation. During optical trapping, the operator must actively circumnavigate debris while only observing a narrow field of view covering a small fraction of the route. A classic approach would be to plan trajectories using path finding algorithms (such as A*) and control the instrument to traverse such paths.^{27,28} In optical trapping the field of view is limited by the need for high numerical aperture (NA) objectives. Acquiring tiles of images allows the coverage of areas larger than a single field of view to establish the state of large environments; however, this can take time. Fields of cells are not static since non-diffusional transport of fluid can occur and acts to disturb the field. Thus, a scan of a microscopic slide can become outdated quickly and repeatedly updating the state of the entire environment after each isolation event is not an efficient strategy. Reducing particle density mitigates some of these challenges, but at a substantial cost to sample processing efficiency.

In practice, the human operator relies on a limited set of cues to steer the trap and avoid obstacles, such as awareness of the general direction to the target, partial memory of the environment and live observation within the local field of view. The challenge of developing autonomous systems to mimic this process presents a problem space remarkably similar to autonomous driving and AI gameplay in 2D spatial environments. One avenue to improve performance beyond the limitations of static algorithms is through machine learning (ML).²⁹ Two widely known approaches are supervised learning (SL), where an ML agent learns by reviewing pre-recorded data, and reinforcement learning (RL), where an ML-agent explores a simulated environment and learns by optimising a reward function.

AI systems have surpassed expert-level human performance in strategic games such as Go,^{30,31} Atari-based platformers^{32,33} and real-time strategy games like StarCraft II.³⁴ Although these

games benefit from deterministic and repeatable simulation environments, real-world robotics and manipulation tasks require agents to generalise under physical uncertainty. Autonomous racing, in particular, has emerged as a compelling challenge for RL due to its need for high-speed planning and adaptation. Racing games have advanced to a degree where they can serve as realistic physics simulators and models have been trained that even outperform professional human drivers.³⁵

With this in mind, we developed MaGIC-OT (machine guided isolation of cells using optical tweezers), a digital simulation environment designed to support model training and evaluation for autonomous optical trapping (Fig. 1). Previous work has demonstrated virtual and augmented optical trapping environments in low-complexity fields populated with microspheres, including applications in outreach and basic proof-of-concept demonstrations.³⁶ While these represent key steps toward AI integration in optical systems, these environments lacked biological realism and were not tailored to the clinical and technical challenges associated with rare-cell isolation. In contrast, MaGIC-OT was designed to emulate the real-world constraints of single-cell trapping in microfluidic devices, including high-density environments, the presence of debris and dynamically evolving local occlusions. It supports both classic pathfinding

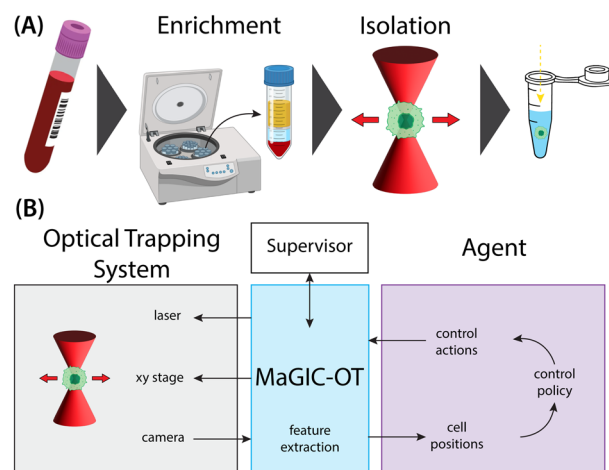


Fig. 1 Processing liquid biopsy samples. a) Workflow for isolating circulating tumour cells (CTCs) from cancer patient blood samples. Enrichment steps typically use microfluidic approaches to deplete red and white blood cells while concentrating CTCs. In this work we use density centrifugation to enrich for CTCs. Cells from the enriched product can be sorted and isolated using optical tweezers. Sorting is based on visual identification, enabling 100% purity. Isolated CTCs remain intact and viable for downstream analysis. b) Schematic representation of the digital simulation environment (MaGIC-OT) for model training and evaluation. The system captures and/or modifies the states of the laser, XY translation stage, and microscope camera data to determine the real-time configuration of the physical optical trapping system. MaGIC-OT extracts key features (e.g. cell type, number, and positions) and provides this information to the agent, which computes control actions such as moving a cell to a target location. The agent can be either a human or a digital model. A supervisory layer oversees MaGIC-OT, ensuring optimal system performance and decision-making.



and deep reinforcement learning models, enabling rigorous benchmarking in simulation prior to deployment.

The MaGIC-OT platform supports dense, biologically relevant simulation environments that reflect key bottlenecks in rare-cell isolation. We demonstrate that MaGIC-OT can be used to evaluate both classic and ML-based models for optical trapping. We show that cooperative learning approaches, involving alternating control between humans and agents, can significantly improve agent performance. Benchmarking against skilled human operators, we show that trained models can outperform expert users in success rate and efficiency, underscoring the translational potential of machine-guided optical trapping for biomedical applications.

Experimental methods

Virtual environment

The environment supports the instantiation of randomised particle fields within bounded arenas that represent microfluidic channels and analysis chambers using a JSON file (Fig. 2). Walls are implemented as impenetrable geometric primitives and all particles are assigned properties including type, radius, position, and mobility. Particle type includes target cells such as CTCs and non-target objects such as other cells (*e.g.* white blood cells, erythrocytes) and debris. Target cells are marked using metadata accessible to the agent only through spatial coordinates, while the visual channel (used by the agent's convolutional encoder) renders the environment at a fixed resolution. A single-beam optical trap is simulated at the centre of the agent's field of view, representing the laser focus of our optical tweezer system. The trap can be programmatically toggled on or off and translated in discrete steps along the X and Y axes by the control agent. When the active laser intersects a particle (*i.e.*, a cell or debris object), a trapping force is applied, calculated as a function of the radial displacement between the laser centroid and the particle's centre of mass.³⁷ This force model captures the effective restoring behaviour of the optical potential, simulating realistic laser-mediated translation dynamics (Fig. 2B). In addition to the rendered visual scene, the simulation provides spatial metadata to the agent or downstream learning model. These include: (i) the current X/Y coordinates of the optical trap within the simulated field; (ii) a soft-labelled region estimating the location of the nearest target cell based on a proximity heuristic; and (iii) the nearest region within the target isolation zone relative to the target cell's current location. This spatial information is appended as a structured numerical vector alongside visual input during model training and inference.

The primary objective within the MaGIC-OT simulation environment is to transport a designated target cell into a predefined target zone. Upon successful delivery, both the cell and its associated target zone are removed from the simulation, and a new target cell–zone pair is introduced, thereby supporting uninterrupted training across sequential episodes. To ensure reproducibility of initial particle distributions,




Fig. 2 Virtual learning environment. (A) Annotated view of a prototypical MaGIC-OT simulation, replicating a microfluidic cell isolation platform. In this configuration, the enrichment product flows through a primary channel and target cells are directed into designated analysis chambers. Centrally positioned within the field of view is the optical trapping laser, which can be activated, deactivated, and repositioned by the agent. The environment is populated with target CTCs (green) alongside non-target cells and debris (grey), including contaminating red and white blood cells. (B) The agent's behaviour is contingent on whether the target CTC is untrapped (*i.e.* not influenced by the laser) or trapped (*i.e.* under active laser influence). Yellow chevrons indicate the direction of the target zone, where CTCs must be deposited to achieve isolation. (C) Classical digital agents employ path-finding algorithms to compute an optimal trajectory from the CTC's current position to the target zone. (D) Human and deep learning-based agents, which emulate human operators, rely solely on real-time visual input within the current field of view. These agents must actively explore and dynamically adapt to environmental changes to successfully navigate and isolate the target CTCs.

MaGIC-OT allows for deterministic seeding: a user-defined random seed set at the start of each run generates an identical configuration of cells and debris, which can be cached for rapid reuse in future sessions.

To facilitate effective training and debugging of artificial agents, the simulator includes a suite of configurable parameters. An action limit can be defined such that the simulation terminates automatically after a maximum number of discrete actions, preventing agents from stalling in unproductive or infinite loops. Additionally, an optional “start helper” can be enabled to initialise the optical trap and field of view at a fixed distance and random angle from the nearest target cell, ensuring standardised initialisation without hardcoding start positions; if omitted, default positions are loaded from predefined level files.



MaGIC-OT also supports fully automatic recording of gameplay sessions. These recordings capture both the rendered simulation and time-resolved spatial metadata, including trap positions, target locations and selected actions. Visual overlays can annotate each frame with the agent's reward at that timestep and arrows indicating the direction of the chosen action. Output data can be exported as image sequences or structured numerical arrays, with accompanying metadata files to facilitate *post hoc* analysis. This infrastructure supports reproducibility, quantitative benchmarking, and detailed behavioural audits of agent training and performance.

Computing environment

Training pipelines were implemented in Python using PyTorch with GPU acceleration enabled. Frame-wise logging of image data, spatial metadata, and user actions was handled *via* the MaGIC-OT CellLogger system. Statistical analysis, including comparison of isolation success rates and path planning performance, was performed in R using standard statistical packages. Significance testing included Student's *t*-test, with *p*-values < 0.05 considered statistically significant.

Classic algorithm

MaGIC-OT can execute classic path-finding algorithms and display the resulting path as an image, to guide the human operator, or as an additional input for machine learning tasks. The path can be broken down into waypoints and used inside the simulation to integrate traditional dynamic programming approaches. The environment can be exported as an image (or 2D matrix of pixel values) which is dynamically cropped to extend slightly beyond the area of interest. This view aims to simulate a scan of a microscope slide and can either be returned as a normal grayscale image or as a binary mask describing traversable and non-traversable areas of the microfluidic chip. Further, the binary mask can be annotated with a cost function to discourage paths near obstacles. This matrix can then be passed to a pathfinding algorithm. Here, we employ the classic A* algorithm for planning collision-free trajectories from a target cell's starting position to an empty analysis chamber. The planning space is represented as a probabilistic roadmap constructed over the free space of the microfluidic geometry.³⁸ Each node corresponds to a candidate waypoint, and connectivity is determined by local visibility in Euclidean space, filtered through a collision-checking function. Obstacle inflation is applied to account for the physical dimensions of the trapped cell.

Machine learning

Neural network architecture. The neural network architecture underlying MaGIC-OT is designed to process multimodal input derived from the simulation environment. It accepts two distinct input streams: (i) a grayscale image of the

current field of view, represented as a two-dimensional matrix, and (ii) a vector of spatial information, comprising position-related metadata as described in the previous section. The image input is processed through a sequence of three convolutional layers. Each convolutional layer is followed by batch normalisation and a leaky rectified linear unit (leaky ReLU) activation function. The resulting feature maps are then flattened into a one-dimensional vector. The spatial information is concatenated to this vector, yielding a unified representation of visual and positional features. This composite vector is subsequently passed through three fully connected (linear) layers, each employing a standard ReLU activation function. The final output is a vector with dimensionality equal to the number of discrete actions available to the agent; this output is interpreted as the unnormalised *Q*-values associated with each possible action. Weight initialisation for the convolutional layers was performed using the Kaiming He method, which is optimised for ReLU-type activations.³⁹ To mitigate overfitting and encourage regularisation, dropout layers were inserted between each convolutional layer and between the fully connected layers.

Hyperparameter selection for the network was conducted using a Bayesian optimisation sweep. The objective of the sweep was to maximise classification accuracy on a held-out test dataset during supervised learning. Parameters optimised during the sweep included convolutional kernel size, stride and the number of channels per layer, as well as the dimensionality of the fully connected layers. Additional optimisation was performed for training parameters, including learning rate, batch size and dropout probability. The loss function employed during supervised learning was categorical cross-entropy and model performance was evaluated based on the proportion of correctly classified actions within the test dataset. To address class imbalance within the action space, the loss function can be optionally weighted by the inverse frequency of each action in the training dataset. This weighting increases the penalty associated with misclassification of infrequent actions, thereby improving representation of rare but critical behaviours. Following supervised training, a second Bayesian sweep was conducted to tune the parameters governing the reinforcement learning training regime. This included optimisation of reward function parameters and the ϵ -greedy exploration schedule.

Data augmentation was explored as a strategy to improve sample efficiency and generalisability. Rotation of training examples was implemented, though this required coordinated transformation of the input image, spatial information and action labels. Two rotation modes were implemented in MaGIC-OT. In rotation mode 1, all coordinate-based metadata are rotated about the laser's position, which remains fixed at the centre of the image. This preserves the laser's coordinates while transforming all other spatial references. In rotation mode 2, the entire environment is rotated about its global centre, resulting in updated coordinates for the laser and all other spatial elements. These procedures enable the reuse of image-state-action triplets while maintaining geometric consistency, a



technique widely used in vision-based machine learning systems.

Reinforcement learning. The RL agent was trained using PyTorch (1.10, Cuda 11.3) and a deep Q-learning (DQN) algorithm with an ϵ -greedy exploration strategy. Episodes were conducted within the MaGIC-OT simulator using randomly seeded environments with variable cell density, obstacle distributions and goal locations. The reward function was shaped to include positive rewards for progress toward the goal, penalties for deviation or collision and a terminal bonus for successful isolation.

Supervised and cooperative learning. For supervised learning, models were trained using human gameplay trajectories recorded within the MaGIC-OT environment. The dataset was annotated using a frame-by-frame logging system (MaGIC-OT CellLogger; see below) that stores both image-state pairs and corresponding human actions. Data augmentation included random rotations and flips of chip environments to enhance generalisability. Weighted loss functions were trialled to upweight rare or context-sensitive actions. To improve policy robustness, a cooperative learning framework was implemented. Here, the human and agent alternate control of the optical trap during a training episode, allowing the agent to encounter challenging or unstable states that arise from suboptimal human decisions. This approach significantly enriched the replay buffer and enabled the agent to learn recovery strategies otherwise absent in clean demonstrations.

Data logging and replay

To enable supervised learning and performance benchmarking, a logging suite (MaGIC-OT CellLogger) was developed. Within the simulation environment, the system records real-time optical trap coordinates, user actions, system state, and environmental metadata (e.g. cell positions, goal locations) with each camera frame. All simulation data were stored in a structured format with parallel JSON metadata to ensure compatibility with downstream machine learning frameworks. Replay functionality was implemented to allow step-by-step reconstruction of both human and agent isolation attempts. Logged simulation episodes could be loaded into the MaGIC-OT simulator, enabling identical re-execution of actions and visual verification of performance under controlled conditions. The logging infrastructure supports selective filtering (e.g. success-only, high-density episodes) and export to PyTorch-compatible datasets. The CellLogger module also supports in-line annotation of key events such as trap loss, collision, or recovery, which were used during cooperative training experiments to identify salient transitions for experience replay augmentation.

For experimental validation, MaGIC-OT can be operated by a human user interacting with the physical microscope, stage and laser. In this context, interaction with the microscope, stage and laser must be captured through a data logger. CellLogger acts as the intermediary between control inputs and the microscope stage while capturing a live image from the microscope camera.

The application supports a multitude of physical input devices (game controller, joysticks, etc.), which are relayed to the microscope stage controller through a RS232 connection. Additionally, CellLogger tracks the position of the stage. The gathered data can then be used to train models. Once trained, CellLogger can utilise a trained model to move the stage, while a ratio of human and AI interaction can be set, to switch between operators. Automated and manual switching between human and AI operator was implemented to generate richer training datasets and allow human operators to “rescue” the AI from situations not present in a recording purely operated by a human.

Cell isolation

Optical trap platform. A single-beam optical trap was constructed around an inverted microscope (Ti2; Nikon, Japan) by integrating a continuous-wave Ytterbium fibre laser (1070 nm; IPG Photonics, UK). The laser beam was expanded using a telescope system to slightly overfill the back aperture of a high-NA oil-immersion objective (60 \times , NA = 1.4; Nikon, Japan). Beam alignment was achieved using two steering mirrors. The laser was focused into a microfluidic chamber containing the sample, enabling trapping of individual cells. A brightfield illumination source provided real-time imaging, with a camera (Andor; Oxford Instruments, Ireland) used for tracking. The laser power at the sample plane was adjusted *via* a half-wave plate and polarising beam splitter to optimise trapping forces while minimising photodamage. MaGIC-OT could software control the motorised stage (ProScan III; Prior Scientific Instruments, UK) using manufacturer supplied drivers (Prior SDK v.1.9.2). With a 60 \times objective, 140 \times 10 tiles are required to scan the single cell analysis chip (Fig. 3 and 6), resulting in an overall acquisition time of 16–17 min; this is the minimum time cost to update the state of the entire environment for path planning. Faster acquisition times (25 \times 3 tiles, 1 min) can be achieved using a lower power 10 \times objective at the cost of resolution. Continuous switching between dry and immersion objectives is often problematic due to the propensity for fluid residue (water or oil) to remain on the sample; this obscures imaging with the dry objective and is not a reliable strategy in this setup.

Chip fabrication and use. Microfluidic devices were fabricated using standard soft lithography techniques. Briefly, photomasks were used to pattern SU-8 photoresist (MicroChem) on silicon wafers through conventional UV photolithography, producing a negative relief mould with a channel height of 35 μ m. Polydimethylsiloxane (PDMS; Sylgard 184, Dow Corning) was mixed at a 10:1 (base:curing agent) ratio, degassed under vacuum and poured onto the master mould. After curing at 70 $^{\circ}$ C for 3 hours, individual chips were cut from the PDMS slab and fluidic inlets/outlets were made using a desktop drill. The PDMS devices were irreversibly bonded to glass coverslips (No. 1.5, 170 μ m thickness) following surface activation *via* oxygen plasma treatment. Following bonding, channels were vacuum-filled



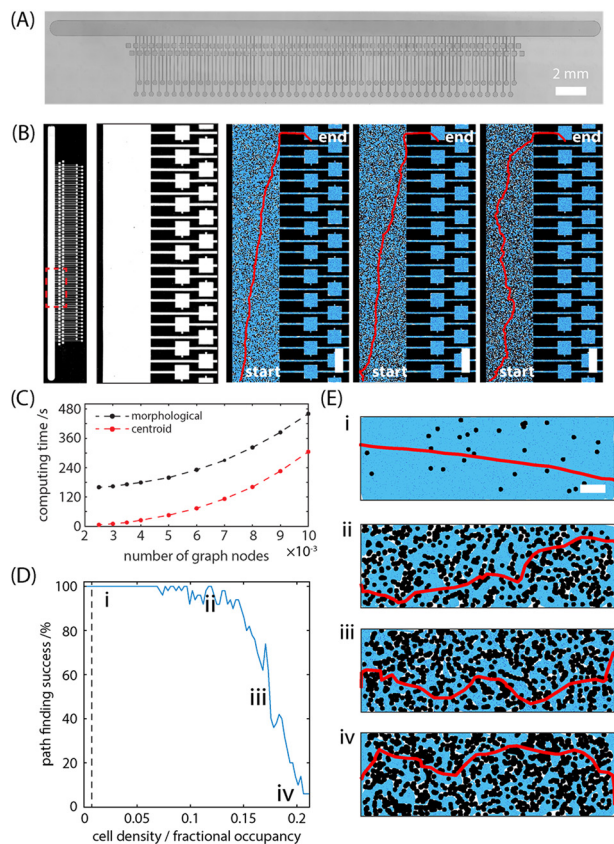


Fig. 3 Path planning for single-cell isolation using the A* algorithm. (A) Full-field brightfield image of the microfluidic device used for cell isolation, comprising a central transport channel for enriched cell suspensions and lateral microchambers designed as isolated analysis wells. (B) Binarised representation of the same field, where traversable regions (white) and non-traversable structures (black) are delineated. The area outlined in red (dashed box) indicates the subregion used for planning experiments. A probabilistic roadmap is generated over the free space and isolation paths are computed from the designated target cell origin to an assigned analysis chamber. Simulations are conducted under varying cell densities to evaluate long-range path planning performance. Scale bars: 500 μm . (C) As cellular crowding increases, a higher density of roadmap nodes is required to capture feasible paths, resulting in non-linear scaling of computational cost. Incorporating obstacle inflation – expanding object centroids (e.g., erythrocytes, WBCs, CTCs) by a fixed radius equivalent to the trapped cell – substantially improves computational efficiency. Depending on environmental density, this approach yields 1.5- to 25-fold reductions in planning time. (D) Isolation success rate as a function of cell density, measured over $n = 50$ trials per condition. The vertical dashed line denotes the operational cell loading density of the DEPArray system, a benchmark CTC isolation platform. (E) Graphs (i–iv) show representative simulation snapshots at selected density values, illustrating the effect of environmental complexity on isolation success. Scale bar: 50 μm .

with 4% (w/v) PBSA (phosphate-buffered saline with 4% bovine serum albumin) and incubated to passivate channel surfaces and prevent nonspecific cell adhesion. Chip architecture was based on previously validated designs for single-cell isolation and analysis.^{40–42}

Cell sample preparation. MCF7 (ATCC) were cultured using low glucose Dulbecco's modified Eagles medium (DMEM;

Thermo Fisher Scientific, UK) supplemented with 10% (v/v) fetal bovine serum (FBS; ThermoFisher Scientific, UK) and 1% penicillin–streptomycin (Thermo Fisher Scientific, UK) in polystyrene flasks in a 5% CO_2 37 °C cell incubator. For spiked blood experiments, 1000 cells were spiked into 10 mL blood draws from healthy volunteers. Human samples used in this research project were obtained from the Imperial College Healthcare Tissue Bank (ICHTB). ICHTB is supported by the National Institute for Health Research (NIHR) Biomedical Research Centre based at Imperial College Healthcare NHS Trust and Imperial College London. ICHTB is approved by Wales REC3 to release human material for research (17/WA/0161). Approvals were provided by the ICHTB and informed consent was obtained from all human participants of this study. Spiked cells were enriched by density centrifugation (OncoQuick; Greiner Bio-One, Austria).⁴³ The fraction containing 'CTCs' was carefully extracted by pipette and washed with 4% PBSA by centrifugation before being resuspended in 0.5 mL 4% PBSA. The recovery rate of spiked cells was 25% on average (range 5–50%, $n = 3$), which was evaluated in separate spiking experiments. The enriched CTC solution was introduced stepwise into the microfluidic chip using a syringe pump since the load volume exceeded the channel volume; approximately 50 μL of solution was processed. Density centrifugation is a facile enrichment step, not a final isolation step. Its value in our workflow is to quickly reduce non-target cell load while retaining enough target cells so MaGIC-OT can be evaluated in performing the end-point, single-cell isolation step. Steps that improve cell recovery and/or purity post density centrifugation have been reported.⁴⁴

Statistical analysis

All the experimental data in this work were measured at least three times and recorded as means \pm SD. Shapiro–Wilk test student's *t*-test was performed on data. Details can be found in the figure legends including sample size (n), and probability (P) value. $P < 0.05$ was considered to be statistically significant. Statistical analyses were conducted using python or MATLAB.

Results

MaGIC-OT environment

MaGIC-OT was designed to provide a high-fidelity digital simulation of real-world optical tweezers experiments, while supporting seamless integration with machine learning workflows (Fig. 2). Built on a 2D game engine (PyGame), it renders a microscope's field of view in real time as seen by an operator through the eyepiece or camera (Fig. 2A). The simulation can track tens of thousands of particles and compute realistic optical trapping forces on each (Fig. 2B). We quantified force and displacement behaviours *in silico* using an analytical force model for spherical particles and aligned them qualitatively with user observations of the physical system; however, full quantitative force calibration *in vitro* may be necessary in future.



MaGIC-OT is implemented as a modular Python application to maximize flexibility and portability. It can be embedded into custom Python scripts, and a lightweight wrapper abstraction allows the underlying engine to be swapped or upgraded without altering the core code. Users can easily modify environment parameters *via* configuration files (JSON), enabling changes to channel geometry, spawn locations, and particle properties without recompilation. Every aspect of the virtual setup is dynamic: walls can be repositioned, cell spawn areas can be redefined, and new obstacle layouts can be loaded on the fly. In summary, MaGIC-OT achieves its goal of faithfully emulating optical trapping of rare cells while providing a versatile platform for integration with automated control algorithms and machine learning models (Fig. 2C and D). This foundation allows rigorous *in silico* experimentation, bridging the gap between physical optical tweezer setups and computational approaches for single-cell handling.

Classic approach with A*

MaGIC-OT enables *in silico* experiments and comparisons between pathfinding algorithms. Using the MaGIC-OT simulator, we first evaluated a classical path-planning strategy for cell isolation (Fig. 3). The A* algorithm, a best-first graph search method was applied to plan an optimal path for moving a target cell to an isolated chamber.²⁷ To facilitate this, we generated a probabilistic roadmap (PRM) of the traversable space within a microfluidic channel (Fig. 3B).³⁸ Roadmaps are determined by the number of nodes and the maximum connection distance between nodes. The roadmap nodes occupy the “free” space inside channel walls not occupied by any cells (either the target or bystander cells), based on a chip design for single-cell analysis. Edges connect nodes that are within a defined distance, without imposing a grid, taking full advantage of the continuous freedom of optical trap movement. A* then efficiently searches this roadmap for the shortest path from the target cell’s initial location to the pre-defined isolation zone, using the sum of travelled distance g and a heuristic estimate h (straight-line distance to the goal) to guide its exploration. An environment increasingly congested with cells requires a higher density of nodes with a shorter connection distance to maximise the probability of calculating a feasible path through a random cell field with a cost in computation time.

A key step when determining the traversable collision-free space of the environment is to recognise that the path must consider the diameter of the cell being translated by the OT. Padding, or inflating, all non-traversable objects and areas can be computationally demanding performed in a generalist way across large pixel-dense environments. We achieved significant improvements in computation time when inflating the centroid of objects within the environment (erythrocytes, CTCs, WBCs, *etc.*) with a set inflation radius *i.e.* that of the currently trapped cell (Fig. 3C).

This inevitably sacrifices nodes in the roadmap for improved computation time. In practice, when isolating rare cells, this

limits the density at which cells can be isolated with high probability. For environments where all objects have been identified as cells, this is not at issue. However, objects which significantly differ in shape from single cells, such as debris or cell aggregates, pose challenges to this approach. Ultimately users must decide which approach is suitable for their application and the trade-offs this makes. To mitigate drift, we experimented with incremental re-planning of local segments, but full-field rescans remained the time-limiting step (≥ 16 min at 60 \times) and were therefore not practical for rapid serial isolations. Importantly, our implementation of A* assumes a static snapshot of the environment during the planning; any cell movements after the path is calculated are not accounted for at this stage (we address this limitation below).

We next examined how the rate of finding successful paths changed as the complexity of the environment increased. Both human and algorithmic operators began to fail in isolating the target cell when the field became highly congested with other cells. Fig. 3D shows the relationship between cell density and isolation success. Up to 0.15 fractional occupancy (1.74×10^{-3} cells μm^{-2}) (Fig. 3E ii), path finding success was $\geq 90\%$, after which a decrease to 50% at 0.18 fractional occupancy (2.05×10^{-3} cells per μm^2) (Fig. 3E iii) and 6% at 0.21 fractional occupancy (2.50×10^{-3} cells per μm^2) (Fig. 3E iv) is observed. For comparison, the vertical dashed line indicates the operational loading density (0.007 fractional occupancy, $\sim 0.08 \times 10^{-3}$ cells per μm^2) of the DEPArray, a benchmark dielectrophoresis-based single cell isolation platform used in conjunction with the CellSearch CTC isolation system (Fig. 3E i).¹⁸

We explored how MaGIC-OT Classic would perform against human operators. In this scenario, the state of the microfluidic environment was determined and paths were planned for execution (Fig. 4). Skilled human operators were shown a binarised top-down view of the simulated chip and asked to manually draw a path to deliver the target cell to the goal (isolation chamber), avoiding other cells. MaGIC-OT Classic, in turn, computed its path on the same initial layout using A* with inflation optimisation enabled. Only successful trials, where a complete path to the goal existed, were analysed; this ensured that path length and timing comparisons were not biased by failed attempts. It is, perhaps, unsurprising that MaGIC-OT Classic was able to identify collision free paths significantly faster than its human counterparts (Fig. 4A). We measured the length of the paths and found them to be similar (Fig. 4B) with the machine-identified paths being on average slightly shorter. This is an encouraging result for both operators. On the one hand, MaGIC-OT Classic is able to outperform skilled human operators to reduce overall isolation path lengths and therefore minimise trapping times. While suggesting that experienced optical tweezer users intuitively plan near-optimal trajectories, although the automated planner can still eke out modest improvements in efficiency. Shorter paths directly translate to reduced trapping time for the target cell, which is beneficial for cell viability. All human-drawn paths






Fig. 4 Comparative performance of MaGIC-OT classic and expert human operators. (A) Computation time required by MaGIC-OT classic to generate collision-free isolation trajectories using the A* pathfinding algorithm is benchmarked against manual path planning by expert optical tweezers operators across environments of increasing cellular density. (B) Total path length from target cell to isolation chamber is quantified for both machine-generated and human-planned trajectories. The algorithm consistently identifies shorter or equivalent paths with reduced computation time, particularly in complex, high-density microenvironments, highlighting the efficiency of algorithmic routing in constrained biological systems. Error bars: standard deviation ($n = 10$).

were validated in the simulator to be collision-free before comparison, to ensure a fair baseline.

Overall, the classic A*-based planner demonstrated the ability to match or outperform skilled humans in both speed and path quality for single-cell isolation in static scenes. This confirms that even a relatively simple algorithm can automate the key task of routing a cell through a crowded microfluidic channel with precision. However, a clear limitation of the static path planning approach is the need for a stable environment during execution: if cells move appreciably after the initial plan, the pre-computed path may no longer be viable. In a real experiment, updating the plan would require pausing to re-image the entire field and then recompute a new path, incurring a significant time cost.

Deep reinforcement learning

We sought an approach that could handle dynamic environments without constant rescanning. To achieve this, we developed a deep reinforcement learning (DRL) framework that enables adaptive control of the optical trap in real time.^{45,46} In a microfluidic device, cells such as CTCs are not fixed in place: they can drift due to flow instabilities or Brownian motion, meaning a path that was initially collision-free may become obstructed moments later. As a result, deterministic path-planning approaches become increasingly ineffective over time unless continuously re-initialised with updated spatial information. While computing a new path is relatively fast, obtaining the updated positions of all cells requires performing another microscope scan of the device, which is considerably slower. In our setup, the time to perform a full brightfield scan of the microfluidic chip (to capture all cell positions) can take on the order of several tens of minutes, depending on the objective used. Identifying specific target cells requires additional fluorescence imaging in multiple channels (*e.g.* for CTCs: nuclear stain, EpCAM+,

CD45-).³ This creates a fundamental bottleneck: the global state of the system cannot be refreshed frequently without pausing the procedure, during which time the cells may continue to move. Of course, specialised optical microscopes may be built or configured to simultaneously capture all fields, but this does not obviate the need to rescan. Without global path planning, environmental information is only partially known and is limited to the extent of the field of view. Human operators overcome this by using an egocentric strategy, making decisions based only on the local field of view and continuously adjusting as they go, rather than relying on a perfect global map. Inspired by this, we trained a deep reinforcement learning agent to perform cell isolation in a similar way, using only local, real-time visual inputs and reactive decision-making (Fig. 5).

A deep Q-network (DQN) agent was deployed as the central decision-making architecture within the MaGIC-OT control framework, enabling data-driven inference of optimal actions from high-dimensional visual and spatial inputs (Fig. 5A). The agent received two modalities of input: a pixel-level rendering of the microscope field and spatial coordinates indicating the positions of the target cell and destination chamber. The action space comprised discrete translational operations for the optical trap. A custom reward function was defined to incentivise biologically relevant behaviours, including: i) minimising the distance between the trap and the target cell, ii) translating the trapped cell toward the goal, iii) successful deposition of the cell within the analysis zone, and iv) penalising events such as loss of the target cell through collision with a wall or another particle.

Learning was facilitated using an ϵ -greedy exploration policy and iterative updates to the Q-function *via* stochastic gradient descent. Through successive training episodes in synthetic environments, the agent converged on robust policies for obstacle avoidance and target delivery. Following training on a curated set of simplified microfluidic geometries, the DQN agent exhibited competent navigational capabilities, reliably executing isolations while dynamically circumventing mobile obstructions (Fig. 5B). Notably, the agent demonstrated consistent performance across diverse layout configurations. These results demonstrate that a convolutional neural network can overcome these challenges to learn successful control policies from raw video data in complex RL environments.

To enhance generalisation and sample efficiency, a hybrid learning approach in cooperation with a human operator was introduced. DRL with human feedback has been applied in various domains, including robotic control, game playing and natural language processing. It enables agents to learn complex tasks more efficiently by leveraging human feedback as a strong signal during the learning process. Initial policy weights were obtained *via* supervised learning (SL) on recorded human operator trajectories, leveraging historical isolation attempts. During recordings, a smart guidance tool helps the human operator to understand the surroundings by displaying the same positional data that machine learning



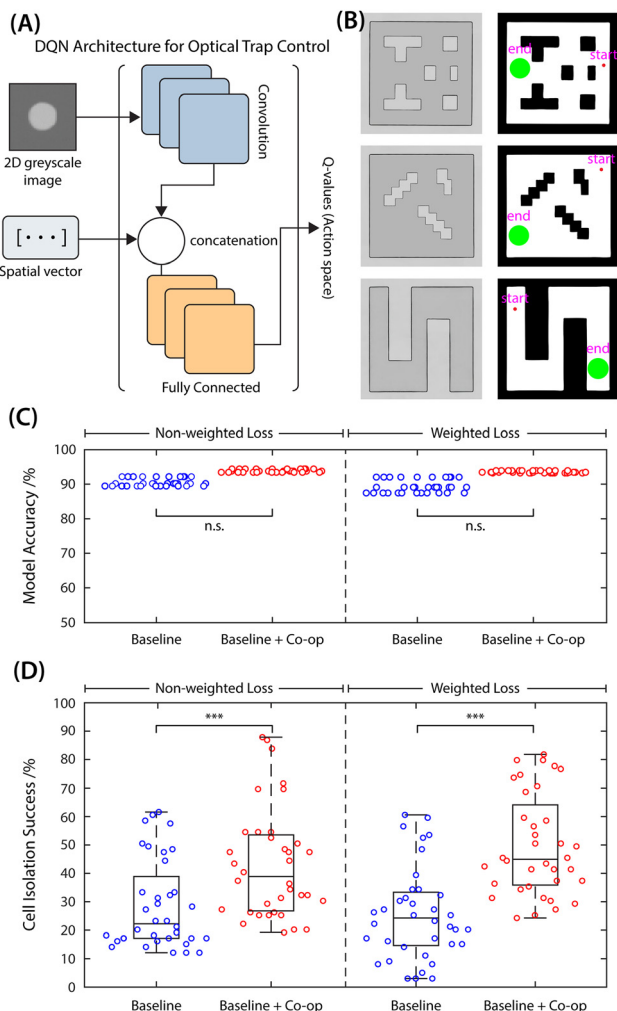


Fig. 5 Architecture and performance of the deep reinforcement learning agent. (A) Schematic representation of the DQN model architecture used for optical trap control. The network receives dual inputs: a 2D image array representing the local microscope field of view and a vector encoding spatial metadata (coordinates of the target cell and destination zone). The image input is processed through three convolutional layers, each followed by leaky ReLU activation and dropout regularisation to prevent overfitting. The resulting tensor is flattened and concatenated with the spatial vector, producing a unified representation that is passed through three fully connected layers with ReLU activations and dropout. The output is a vector corresponding to the discrete action space, and the action with the highest predicted Q-value is selected *via* an argmax operation. (B) Representative examples of training environments used during agent development. The target cell (small red circle) and the goal zone (large green circle) are placed randomly within predefined spawn regions. The spatial configurations and obstacle placements vary between levels to enhance policy generalisation. (C) and (D) Quantitative evaluation of model performance under different training paradigms. “Baseline” models are trained solely *via* supervised learning on human demonstration data. The “Baseline + Coop” variant incorporates cooperative training sessions between a human operator and the agent, exposing the model to a broader distribution of state-action trajectories. Cooperative training substantially improves isolation success rates in simulation. In contrast, applying class-weighted loss during supervised learning has no statistically significant impact on test accuracy or simulation performance. These results highlight the critical role of experiential diversity, rather than loss function rebalancing in developing robust, generalisable cell manipulation policies. Students t-test: n.s., no significant difference at $P < 0.05$; ***, significant difference at $P < 0.001$.

agents receive. The pre-recorded footage was split into training and testing datasets (9:1 ratio). The neural networks of the SL and DQN have the same layout, allowing the weights and biases to be transferred between them. Cooperative training sessions alternated control between a human operator and the AI agent. This strategy exposed the agent to complex or adverse conditions unlikely to be encountered through naive exploration, allowing refinement of sub-optimal behaviours and enriching the training with critical recovery behaviours when encountering dead ends.

To assess the impact of model training methodology, we implemented a series of ablation experiments using baseline supervised learning models trained exclusively on recorded human gameplay data. In all scenarios, success was defined as delivery of the target cell to the isolation chamber without incurring collision or loss. These models achieved high accuracy on held-out test datasets, typically exceeding 90% irrespective of whether weighted or non-weighted loss functions were employed (Fig. 5C). The inclusion of a weighted loss, which increases the contribution of infrequent actions during training, did not significantly alter test performance ($p > 0.05$), albeit reducing variance, suggesting that the baseline datasets were sufficiently balanced or that rare-action sparsity did not limit predictive fidelity.

However, while test accuracy remained consistently high, static evaluation on curated datasets did not fully capture real-time operational performance. When deployed in dynamic simulations within MaGIC-OT, supervised learning models achieved a mean success rate of approximately 22% (non-weighted) and 24% (weighted) across 20 test environments (Fig. 5D). By contrast, models trained with cooperative learning achieved significantly improved performance, with a mean success rate of 40% (non-weighted) and 45% (weighted), and peak success rates approaching 80–90% in certain environments. These results demonstrate the practical gains afforded by hybrid training strategies and contextual learning.

The discrepancy between static test accuracy and live simulation success underscores the importance of evaluating model performance in contextually realistic, temporally extended scenarios. Successful isolation requires sequencing correct decisions across long action horizons, with error accumulation posing a challenge for purely supervised policies. Nevertheless, the encouraging performance of the hybrid-trained agent in simulation, where it completed isolation tasks in complex, obstacle-rich environments, highlights its real-world applicability.

Overall, the DRL-enhanced MaGIC-OT framework demonstrates strong capacity for automated cell manipulation in dynamically evolving microenvironments. By eliminating the need for exhaustive whole-field rescanning and enabling on-the-fly control adjustments, the agent can successfully isolate target cells while reducing latency and photochemical/thermal stress. These attributes are critical for ensuring sample integrity and viability, particularly in workflows requiring the isolation of rare and sensitive cells for downstream molecular analysis.



Isolation of cancer cells

We conducted proof-of-concept isolations to demonstrate the real-world applicability of the MaGIC-OT platform using the single cell isolation microfluidic device used to train MaGIC-OT (Fig. 6). A spiked blood sample mimicking a patient liquid biopsy is processed to enrich for CTCs and the product introduced into the chip. While cancer cells are enriched, the product is not pure and contains a significant proportion of contaminating cells from other blood factors *i.e.* WBCs, erythrocytes and platelets that can confound downstream assays without further processing. In our current implementation, image analysis generates a candidate list (segmented objects with basic features), and an operator confirms the target before initiating autonomous manipulation. This human-in-the-loop designation reflects how many clinical workflows are run today and deliberately decouples recognition from control, allowing the control policy to be benchmarked independently of the classifier. The optical trap was activated and centred on a target cell, which can then navigate the captured cell toward the nearest empty analysis chamber on the chip. A collision-free path is achieved that avoids obstacles, such as neighbouring cells and microfluidic walls. The main-channel loading was 0.13 ± 0.02 fractional occupancy ($n = 5$ fields). Approximately 250 target cells were recovered and resuspended in 0.5 mL (target-

cell concentration $\sim 5 \times 10^2$ cells per mL). The overall suspension, however, remained dominated by residual blood components (RBCs, WBCs, platelets and debris). A fractional occupancy of 0.13 corresponds to an effective total object concentration on the order of $1\text{--}2 \times 10^7$ cells per mL, demonstrating that optical tweezer-based isolation can be performed in a highly crowded regime despite a low target-cell concentration. The precise trajectory taken by the trapped cell is illustrated in Fig. 6A, where successive fields of view are shown in sequence. The trajectory highlights how the laser-guided cell moves out of the main channel and into the side chamber (Fig. 6B). Throughout this transfer, no other cells were pushed or co-transported as the trapped cell circumnavigated any nearby cells, thereby preventing accidental capture of unwanted cells or debris. After transport, the target cell was released into the designated chamber by deactivating the optical trap, achieving a physically isolated single-cell sample. The result was an isolation event in which no additional cells were observed in the destination chamber, indicating visually confirmed single-cell capture without detectable co-transport. Fig. 6B iii demonstrates this, showing a candidate CTC (green arrowhead) that has been successfully moved from the mixed population into an empty analysis chamber. The starting location of the target cell in the enriched sample is surrounded by contaminating cells and debris (white arrowheads), whereas its final location is a purified chamber free of any other particles. In total, 5 operator-designated target cells were presented to MaGIC-OT and 5 were deposited into single-cell chambers. The isolation path length was $889 \pm 59 \mu\text{m}$ (average \pm standard deviation) and time to isolate was 71 ± 5 s, yielding a mean transport speed of $15.0 \pm 1.7 \mu\text{m s}^{-1}$. Under these conditions, this corresponds to ~ 51 cells per h, which is comparable to DEPArray recovery rates.^{18,44}

This demonstrates that the platform can extract a single viable cell from a complex sample and deposit it into a separate compartment with no visible cross-contamination. This highlights MaGIC-OT's potential for real-world liquid biopsy applications. The ability to target a single rare tumour cell from a blood sample and isolate it opens opportunities for comprehensive single-cell analysis of CTCs in a clinical context. For instance, an isolated viable CTC can be analysed for molecular markers, cultured to test drug responses, or examined for metastasis-related traits. The high-purity isolation achieved here is particularly valuable because it ensures that downstream analyses are not confounded by background blood cells.

Conclusions

MaGIC-OT delivers three substantive advances. (i) A centroid-inflated A/PRM* planner that inflates each obstacle by the trapped-cell radius at its centroid, preserving clearance yet achieving order-of-magnitude faster planning in crowded channels. (ii) Cooperative human-in-the-loop training, where a human operator rescues the policy from rare failure states are incorporated into learning, measurably improving




Fig. 6 Optical trapping-based isolation of candidate CTCs. (A) The dashed yellow line traces the path of the optical trap. (B) Isolation of a target cell (green arrowhead) from an enriched spiked blood sample. The cell is initially surrounded by contaminating cells and debris (white arrowheads). The cell moves from a crowded field (i and ii) into a clean, physically separated analysis chamber (iii). The final isolated chamber contains only the target cell, confirming high-purity capture. Scale bars in subpanels: 50 μm .



isolation success and robustness beyond imitation or self-play alone. (iii) Deterministic single-cell delivery in heterogeneous, cell-laden suspensions that mirror post-enrichment CTC workflows, rather than on beads of sparse fields. MaGIC-OT demonstrates that single-beam optical tweezers can be endowed with closed-loop autonomy for single-cell isolation in microfluidic devices, using either deterministic (A^*) or learned (DRL) control policies. In dense environments and under realistic drift, DRL trained entirely *in silico* surpassed expert performance for isolation time and success, while the classical planner produced near-shortest trajectories when the environment was quasi-static. This delineates when each paradigm is most effective. Graph-based planning excels at minimising path length but requires whole-field imaging to refresh its map; as soon as cells drift appreciably, the advantage erodes. Conversely, the egocentric DRL policy is reactive to local obstacles without needing global state reconstruction, making it tolerant to Brownian and advective motion. A practical implementation could therefore combine the two: a global A^* plan to set a coarse route, with local deviations handed off to the DRL controller for real-time obstacle negotiation.⁴⁷ Such a hybrid scheduler would couple the efficiency of graph search with the adaptability of learned policies and could be generalised to multi-trap or 3-D trapping architectures.

Several limitations should be acknowledged. First, although we were successful in training models for use with the single-cell chip, further work is needed to quantify performance across diverse chip geometries and experimental conditions. Second, global throughput is still constrained by the time required for fluorescence-based target identification and by stage-scanning speeds; we expect hardware acceleration will be required for high-volume clinical use. Third, optical exposure was kept within viability limits for single transfers, but systematic phototoxicity studies and long-term cell-function assays are needed. Finally, while we demonstrated a spiked-blood proof-of-concept, extension to true patient samples at clinically representative CTC numbers will require integrated enrichment strategies and rigorous recovery/purity benchmarking. Addressing these points will position AI-guided optical tweezers as a versatile component of next-generation single-cell diagnostic platforms and as a general framework for autonomous manipulation in microfluidics.

Author contributions

ASR conceived and designed the research and secured funding. JPC and ASR wrote the manuscript. JPC developed code, performed *in silico* experiments and analysed the simulation data. ASR contributed to code development and conducted the classical path-planning simulations, with additional contributions from XX. ASR performed the cell isolation experiments.

Conflicts of interest

There are no conflicts to declare.

Data availability

All data are available in the main text. All relevant data are available from the authors upon reasonable request for academic, non-commercial use. The MaGIC-OT training environment is provided under the Creative Commons Attribution Non-Commercial 4.0 International License (CC BY-NC 4.0) on request.

Acknowledgements

This work was supported by a Community of Analytical Measurement Science Lectureship award and an Engineering and Physical Science Research Council (EPSRC) Innovation fellowship awarded to ASR. All authors are grateful to the NIHR Biomedical Facility at Imperial College London for infrastructure support.

References

- 1 W. J. Janni, B. Rack, L. W. Terstappen, J. Y. Pierga, F. A. Taran, T. Fehm, C. Hall, M. R. de Groot, F. C. Bidard, T. W. Friedl, P. A. Fasching, S. Y. Brucker, K. Pantel and A. Lucci, *Clin. Cancer Res.*, 2016, **22**, 2583–2593.
- 2 F. C. Bidard, D. J. Peeters, T. Fehm, F. Nole, R. Gisbert-Criado, D. Mavroudis, S. Grisanti, D. Generali, J. A. Garcia-Saenz, J. Stebbing, C. Caldas, P. Gazzaniga, L. Manso, R. Zamarchi, A. F. de Lascoiti, L. De Mattos-Arruda, M. Ignatiadis, R. Lebofsky, S. J. van Laere, F. Meier-Stiegen, M. T. Sandri, J. Vidal-Martinez, E. Politaki, F. Consoli, A. Bottini, E. Diaz-Rubio, J. Krell, S. J. Dawson, C. Raimondi, A. Rutten, W. Janni, E. Munzone, V. Caranana, S. Agelaki, C. Almici, L. Dirix, E. F. Solomayer, L. Zorzino, H. Johannes, J. S. Reis-Filho, K. Pantel, J. Y. Pierga and S. Michiels, *Lancet Oncol.*, 2014, **15**, 406–414.
- 3 A. J. Rushton, G. Nteliopoulos, J. A. Shaw and R. C. Coombes, *Cancers*, 2021, **13**, 5.
- 4 A. Hochstetter, R. Vernekar, R. H. Austin, H. Becker, J. P. Beech, D. A. Fedosov, G. Gompper, S.-C. Kim, J. T. Smith, G. Stolovitzky, J. O. Tegenfeldt, B. H. Wunsch, K. K. Zeming, T. Krüger and D. W. Inglis, *ACS Nano*, 2020, **14**, 10784–10795.
- 5 M. E. Warkiani, B. L. Khoo, L. Wu, A. K. P. Tay, A. A. S. Bhagat, J. Han and C. T. Lim, *Nat. Protoc.*, 2016, **11**, 134–148.
- 6 J. Zhou and I. Papautsky, *Microsyst. Nanoeng.*, 2020, **6**, 113.
- 7 S. Ribeiro-Samy, M. I. Oliveira, T. Pereira-Veiga, L. Muineloromay, S. Carvalho, J. Gaspar, P. P. Freitas, R. López-López, C. Costa and L. Diéguez, *Sci. Rep.*, 2019, **9**, 8032.
- 8 D. Jin, B. Deng, J. X. Li, W. Cai, L. Tu, J. Chen, Q. Wu and W. H. Wang, *Biomicrofluidics*, 2015, **9**, 014101.
- 9 T. M. Gierahn, M. H. Wadsworth, T. K. Hughes, B. D. Bryson, A. Butler, R. Satija, S. Fortune, J. C. Love and A. K. Shalek, *Nat. Methods*, 2017, **14**, 395–398.



- 10 J. Yuan and P. A. Sims, *Sci. Rep.*, 2016, **6**, 33883.
- 11 A. K. White, M. VanInsberghe, O. I. Petriv, M. Hamidi, D. Sikorski, M. A. Marra, J. Piret, S. Aparicio and C. L. Hansen, *Proc. Natl. Acad. Sci. U. S. A.*, 2011, **108**, 13999–14004.
- 12 K. Matula, F. Rivello and W. T. S. Huck, *Adv. Biosyst.*, 2020, **4**, e1900188.
- 13 Y. Fan, X. Wang, J. Ren, F. Lin and J. Wu, *Microsyst. Nanoeng.*, 2022, **8**, 94.
- 14 B. Sarno, D. Heineck, M. J. Heller and S. D. Ibsen, *Electrophoresis*, 2021, **42**, 539–564.
- 15 P. G. Bonacci, G. Caruso, G. Scandura, C. Pandino, A. Romano, G. I. Russo, R. Pethig, M. Camarda and N. Musso, *Transl. Oncol.*, 2023, **28**, 101599.
- 16 M. C. Wu, *Nat. Photonics*, 2011, **5**, 322–324.
- 17 C. Nelep and J. Eberhardt, *Cytometry, Part A*, 2018, **93**, 1267–1270.
- 18 M. Di Trapani, N. Manaresi and G. Medoro, *Cytometry, Part A*, 2018, **93**, 1260–1266.
- 19 P. H. Jones, O. M. Maragò and G. Volpe, *Optical Tweezers: Principles and Applications*, Cambridge University Press, Cambridge, 2015.
- 20 T. Avsievich, R. Zhu, A. Popov, A. Bykov and I. Meglinski, *Rev. Phys.*, 2020, **5**, 100043.
- 21 T. Imasaka, Y. Kawabata, T. Kaneta and Y. Ishidzu, *Anal. Chem.*, 2002, **67**, 1763–1765.
- 22 E. Eriksson, K. Sott, F. Lundqvist, M. Sveningsson, J. Scrimgeour, D. Hanstorp, M. Goksor and A. Graneli, *Lab Chip*, 2010, **10**, 617–625.
- 23 K. O. Greulich, G. Pilarczyk, A. Hoffmann, G. Meyer Zu Horste, B. Schafer, V. Uhl and S. Monajembashi, *J. Microsc.*, 2000, **198**, 182–187.
- 24 N. Neve, S. S. Kohles, S. R. Winn and D. C. Tretheway, *Cell. Mol. Bioeng.*, 2010, **3**, 213–228.
- 25 H. Zhang and K. K. Liu, *J. R. Soc., Interface*, 2008, **5**, 671–690.
- 26 B. Landenberger, H. Höfemann, S. Wadle and A. Rohrbach, *Lab Chip*, 2012, **12**, 3177–3183.
- 27 P. Hart, N. Nilsson and B. Raphael, *IEEE Trans. Syst. Sci. Cybern.*, 1968, **4**, 100–107.
- 28 Y. H. Wu, D. Sun, W. H. Huang and N. Xi, *IEEE-ASME Trans. Mechatron.*, 2013, **18**, 706–713.
- 29 V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg and D. Hassabis, *Nature*, 2015, **518**, 529–533.
- 30 D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel and D. Hassabis, *Nature*, 2016, **529**, 484–489.
- 31 D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. T. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel and D. Hassabis, *Nature*, 2017, **550**, 354–359.
- 32 M. G. Bellemare, Y. Naddaf, J. Veness and M. Bowling, arXiv, 2012, preprint, arXiv:1207.4708, DOI: [10.48550/arXiv.1207.4708](https://doi.org/10.48550/arXiv.1207.4708).
- 33 V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra and M. Riedmiller, arXiv, 2013, preprint, arXiv:1312.5602, DOI: [10.48550/arXiv.1312.5602](https://doi.org/10.48550/arXiv.1312.5602).
- 34 O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Y. Wang, T. Pfaff, Y. H. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps and D. Silver, *Nature*, 2019, **575**, 350+.
- 35 P. R. Wurman, S. Barrett, K. Kawamoto, J. MacGlashan, K. Subramanian, T. J. Walsh, R. Capobianco, A. Devlic, F. Eckert, F. Fuchs, L. Gilpin, P. Khandelwal, V. Kompella, H. C. Lin, P. MacAlpine, D. Oller, T. Seno, C. Sherstan, M. D. Thomure, H. Aghabozorgi, L. Barrett, R. Douglas, D. Whitehead, P. Dürr, P. Stone, M. Spranger and H. Kitano, *Nature*, 2022, **602**, 223–228.
- 36 M. Praeger, Y. Xie, J. A. Grant-Jacob, R. W. Eason and B. Mills, *Mach. Learn. Sci. Technol.*, 2021, **2**, 035024.
- 37 A. A. R. Neves and C. L. Cesar, *J. Opt. Soc. Am. B*, 2019, **36**, 1525–1537.
- 38 L. E. Kavraki, P. Svestka, J. C. Latombe and M. H. Overmars, *IEEE Trans. Robot. Autom.*, 1996, **12**, 566–580.
- 39 J. J. Liu, Y. F. Liu and Q. C. Zhang, *Neurocomputing*, 2022, **483**, 171–182.
- 40 A. Salehi-Reyhani, J. Kaplinsky, E. Burgin, M. Novakova, A. J. deMello, R. H. Templer, P. Parker, M. A. Neil, O. Ces, P. French, K. R. Willison and D. Klug, *Lab Chip*, 2011, **11**, 1256–1261.
- 41 S. Chatzimichail, P. Supramaniam and A. Salehi-Reyhani, *Anal. Chem.*, 2021, **93**, 6656–6664.
- 42 P. Supramaniam, Z. Wang, S. Chatzimichail, C. Parperis, A. Kumar, V. Ho, O. Ces and A. Salehi-Reyhani, *ACS Synth. Biol.*, 2023, **12**, 1227–1238.
- 43 R. Rosenberg, R. Gertler, J. Friederichs, K. Fuehrer, M. Dahm, R. Phelps, S. Thorban, H. Nekarda and J. R. Siewert, *Cytometry*, 2002, **49**, 150–158.
- 44 Z. Guo and W. Xia, *Med-X*, 2024, **2**, 28.
- 45 T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver and D. Wierstra, arXiv, 2015, preprint, arXiv:1509.02971, DOI: [10.48550/arXiv.1509.02971](https://doi.org/10.48550/arXiv.1509.02971).
- 46 S. Levine, C. Finn, T. Darrell and P. Abbeel, arXiv, 2015, preprint, arXiv:1504.00702, DOI: [10.48550/arXiv.1504.00702](https://doi.org/10.48550/arXiv.1504.00702).
- 47 M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel and W. Zaremba, arXiv, 2017, preprint, arXiv:1707.01495, DOI: [10.48550/arXiv.1707.01495](https://doi.org/10.48550/arXiv.1707.01495).

