PureSentry:

Can real-time contamination detection transform cell therapy manufacture?

cambridgeconsultants
Part of Capgemini Invent

Contents

© Cambridge Consultants 2022

Executive summary

PureSentry represents a significant breakthrough for contamination detection in cell and gene therapy (CGT) and bioprocessing monitoring. As a prototype device for on-line, real-time detection, it acts as a proof of concept for event-based cameras in rapid microbial detection and a case study for how AI can be applied to solve problems in industry 4.0.

The team here at Cambridge Consultants is proud of PureSentry's development and excited by its potential to help transform the way industry analyses liquid biologics. There is certainly a pressing imperative for change. Cell therapy is acknowledged as the future of medicine, but advances are hampered by the slow and laborious sterility testing that makes biomanufacturing so hard to scale.

PureSentry throws a light on future possibilities in automated testing. It is sensitive to single microbes and cost-effective enough to check batches 24/7. On-line monitoring via a closed loop system ensures continuous detection of batches that are never compromised. Time is saved, costs and risks are reduced, and therapies are delivered to patients sooner. The concept is both immensely powerful and broadly applicable.

This whitepaper has been written as an informative and comprehensive guide to PureSentry's development. It tracks the key challenges faced and solutions deployed during the journey, and explores specific milestones, from addressing user needs and engineering requirements to the creation of detection algorithms for specific types of bioreactors.

The paper captures a number of notable results, including the ability to distinguish sterile cell culture medium, from medium contaminated with live microbes, in a label-free manner. These results are applicable to static and mixing bioreactors, the latter being a more challenging detection problem as human cells are present in the data as well as the contaminant.

The PureSentry project reflects the multidisciplinary capabilities of Cambridge Consultants as a whole, and the particular expertise of our bioinnovation team, which excels at the intersection of engineering, biology and advanced computation. The project group drew on broad experience, listened closely to clients working in the space and followed industry trends before ultimately proposing a solution to an industry-wide problem using a combination of novel imaging technology and advanced AI. That solution represents many facets of our product development expertise, as well as our ability to integrate fluidics, electronics, optics, software and UX development and AI with microbial and human cell culture.

We hope that the paper will help to provoke a wider debate and sharing of ideas as the industry faces up to the vital challenge of advancing manufacturing processes for CGT. We believe that the immediate future is full of opportunity. Real-time monitoring of biological systems, as exemplified by PureSentry, has applications beyond cell therapy manufacture – from food production and cellular agriculture to water monitoring and pathogen surveillance.

© Cambridge Consultants 2022

1 Cell therapy is the future of medicine

Many manufactured products, from food to cosmetics and importantly medicine, require thorough testing for microbial growth. Freedom from contamination is critical when it comes to cell therapy because the product will be infused into a patient and therefore requires the highest standard of quality assurance.

"Freedom from contamination is crucial when it comes to cell therapy."

2 Biomanufacturing is hard to scale because sterility testing is slow and not fully automated

In recent years, increasing numbers of CGT manufacturers have become our clients. This is thanks to our experience of developing medical devices and breakthrough manufacturing technology, but also because the industry faces a difficult challenge. Standard sterility testing methods derive from growth-based tests which aren't scalable and add up to two weeks delay in declaring the therapy safe.

One leading cell therapy costs around $500,000 per dose, meaning not everyone who needs it will receive it. The high price tag reflects how challenging and expensive it is to manufacture. A review of autologous cell therapy manufacture in Japan found that out of 30,000 batches, 0.18% failed due to contamination (Reference 1), so around 1 per 1000 patients treated may have a delay in receiving their therapy or not get one at all. Even more concerning is that a recent study by Glaxo Smith Kline (GSK) found that 59% of the cost of manufacture is Current Good Manufacturing Practice (CGMP) cell handling and analytics combined. Whilst contamination itself stops a small number of patients getting timely treatment, the cost of manufacture could prevent many more patients from even having the option.

We have worked with cell therapy manufacturers to identify the technologies that will transform their practices and make advanced therapies scalable. Crucially, we identified a lack of options when it comes to on-line, continuous monitoring for signs of contamination.

Growth-based tests require highly skilled technicians in expensive cleanrooms to obtain and process samples. The modern rapid sterility testing equipment options we identified have a wide range of costs from $10,000 up to in excess of $500,000 and take a fixed number of samples at a time. It is clear that current solutions increase overheads and hinder cost-efficient scale up. Growth-based tests and rapid microbial tests are typically destructive and require a technician to take the sample, restricting monitoring to individual timepoints.

"It is clear that current solutions increase overheads and hinder cost-efficient scale up."

© Cambridge Consultants 2022

3 What is the best way to automate contamination detection?

We set out to find another way: PureSentry. Our prototype device is sensitive to single microbes and cost-effective enough that every batch can be monitored 24 hours a day, seven days a week. PureSentry reduces labour costs and eliminates sampling, reducing contamination risks. Cell therapy manufacture must be made cheaper and more automated for it to reach its life-saving potential, and PureSentry demonstrates a way to achieve that.

The on-line monitoring approach has many advantages. Monitoring via a closed loop system means the batch is never compromised and cell therapy technicians can work on other tasks. PureSentry monitors continuously and the longer it monitors, the more it will see. This is more comprehensive than sampling at a fixed number of checkpoints. How would continuous monitoring for contamination impact existing processes of manufacturing a cell therapy and delivering it to a patient?

4 How could PureSentry positively disrupt cell therapy manufacturing?

4.1 Changing timelines

PureSentry can begin monitoring a cell therapy batch as soon as it arrives in the manufacturing facility. This means it can flag potential contamination issues faster than off-line growth-based tests can report a result. In the short term, we envisage manufacturers using on-line monitoring as a complementary tool alongside their standard sterility testing protocols. Once the manufacturer gains confidence in on-line monitoring for signs of contamination they will be able to take actions such as:

"PureSentry can flag potential contamination issues faster than off-line growth-based tests can report a result."

- Notifying healthcare providers that the batch may not pass QC and may need to start again
- Conducting additional observations or sterility tests and reviewing records to identify potential sources of contamination
- Following regulator approval, therapies may be able to be released sooner, particularly in urgent cases

In some situations discarding the batch may have benefits. For example:

- Saving expensive reagents and lowering the overall cost of cell therapy manufacture
- Freeing up throughput for another patient
- Restarting manufacture with a back up sample from a patient whose first dose has been compromised

© Cambridge Consultants 2022

Either way the costs associated with the further CGMP cell handling, complex analytical measurements and fill/finish processes can be saved. Finally, as with various rapid microbial testing technologies, on-line monitoring with PureSentry could one day be evaluated against the standard sterility methods. If found to be equivalent to USP 71 (Reference 2), it could replace manual sampling and off-line sterility tests altogether. This would enable manufacturers to eliminate the long delay where a therapy is cryopreserved, ready for the patient but not yet deemed safe for release. This proposed process is shown in Figure 1.

4.2 Scalable contamination detection

A key requirement for PureSentry is that the system must be affordable enough to be deployed on millions of cell therapy batches at once across the industry. We addressed this aim by using a set of low-cost off-the-shelf consumables suitable for single use (microfluidic chip, tubing and connectors) and selected the Prophesee event-based camera sensor, partly as it can be manufactured with economies of scale. We are confident that monitoring every batch 24/7 with PureSentry is as economically feasible as it is imperative for patient safety.

"We are confident that monitoring every batch 24/7 with PureSentry is as economically feasible as it is imperative for patient safety."

Outline CAR-T Manufacturing Process

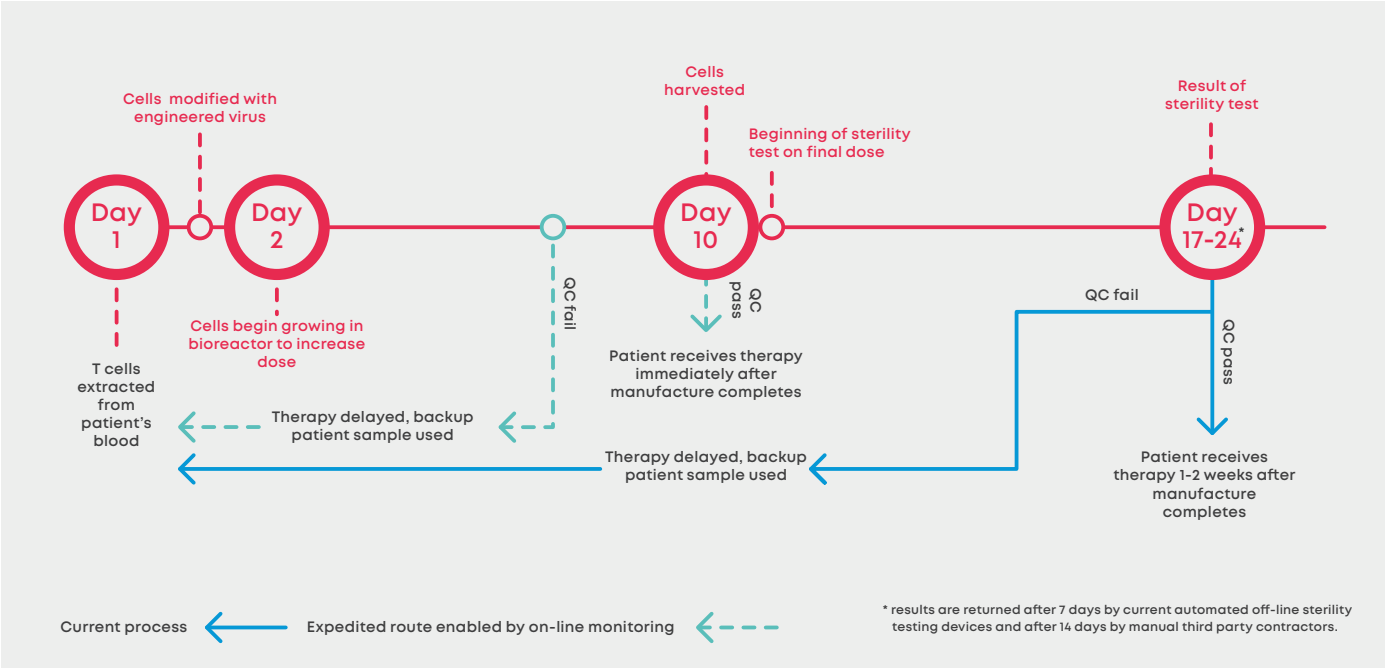

Figure 1: On-line monitoring could enable QC fails to be predicted, and one day with sufficient evidence and regulatory approval could replace manual sampling and off-line sterility testing

© Cambridge Consultants 2022

5 Our Solution: PureSentry

To kick off the project we gathered internal experts from different disciplines who had worked with our cell therapy manufacturing clients. We generated a list of user needs and validated our assumptions with several representatives from the Cell and Gene Therapy Catapult, a UK organisation working to accelerate the development and lower the cost of cell and gene therapy manufacturing. User needs were converted into formal requirements that either specify a type of component or a property that the system must have (see examples in Figure 2).

The requirements drove the selection of the correct components and ensured we built a working system in a short space of time. We were able to meet the requirements for the first prototype using a standard lab microscope instead of custom optics, though our optics specialists have identified components suitable for a low-cost integrated optical set-up in future iterations.

PureSentry takes seconds to set up, thanks to a flip-top pump. Having a touchscreen-ready UI reduces the number of parts and ensures that the device is easy to clean. The fluidic loop connects directly to an industry standard culture vessel, as shown in Figure 3.

5.1 PureSentry performance

Cell therapy manufacturers use a range of specialist equipment including a variety of bioreactors. In order to make PureSentry broadly applicable, we developed detection algorithms for two major types – static bioreactors and mixing bioreactors. These present distinct challenges because the static bioreactor permits sampling of the medium 'headspace' without disturbing the T cells. Therefore, our first algorithm needed to detect only microbes in cell culture medium, whereas our second algorithm for mixing bioreactors needed to detect microbes within a pool of T cells. The following section is in two parts, dealing first with the algorithm for static bioreactors.

"Having a touchscreen-ready UI reduces the number of parts and ensures that the device is easy to clean."

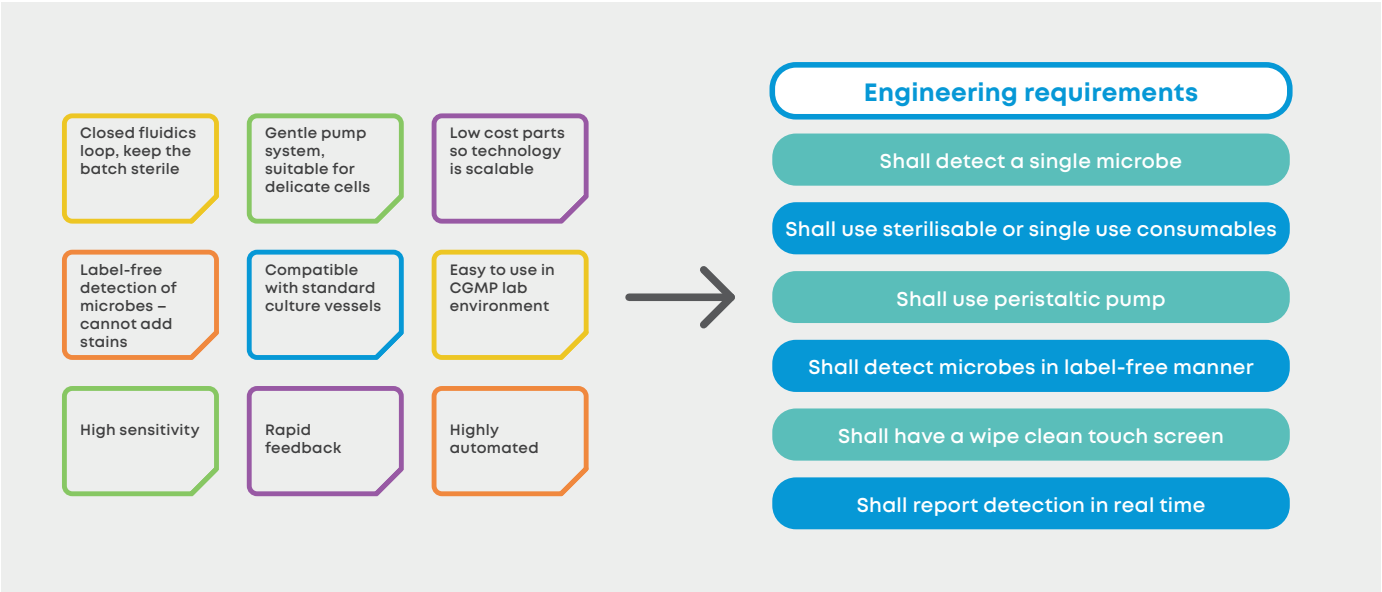

Figure 2: Example of user needs of a cell therapy manufacturing technician considered and translated into engineering requirements

© Cambridge Consultants 2022

Figure 3: Overview of PureSentry prototype and key features pictured with a static bioreactor

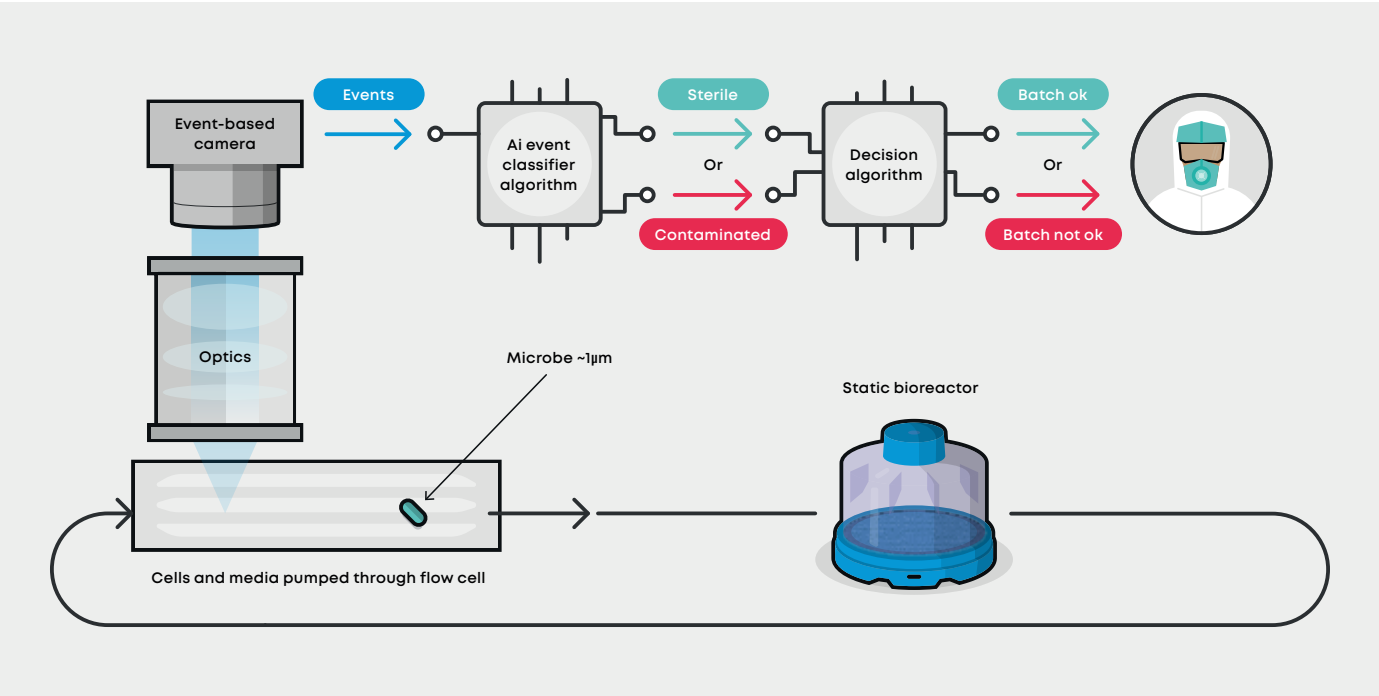

Figure 4: Schematic of PureSentry operating on a static bioreactor, sampling medium containing potential microbes without disturbing the T cells

© Cambridge Consultants 2022

5.1.1 Static bioreactors

PureSentry continuously samples medium from a static bioreactor without disturbing the T cells which grow at the base of the flask. The medium passes through the closed fluidic loop, briefly illuminated by a standard white microscope lamp (Figure 4). The optics relay the image of the sample onto the sensor of the Prophesee event-based camera. PureSentry software converts events from pixels that detect changes into a 'pseudoframe' by combining all the events in a 20 millisecond timespan. An algorithm analyses each frame and classifies it as 'sterile' or 'contaminated' based on training data imaged from contrived samples known to be sterile or containing microbes. The false positive rate is determined empirically and set as a parameter in the code that indicates whether the batch is safe or not. When the rate of contaminated frames exceeds the known false positive rate, the batch is classified as contaminated.

In order to test the PureSentry design, we first characterised the growth of the model contaminant, *E. coli*, in an industry standard culture vessel. In just a few hours, with no agitation, we saw an increase from <100 Colony Forming Units (CFUs) per millilitre to >4000CFUs/ml in around eight hours (Figure 5). This confirmed that our model contaminant grows rapidly in a culture vessel and in mammalian cell culture medium. On-line monitoring can therefore detect earlier than off-line monitoring and is likely to be a complementary tool that makes cell therapies safer in the long run.

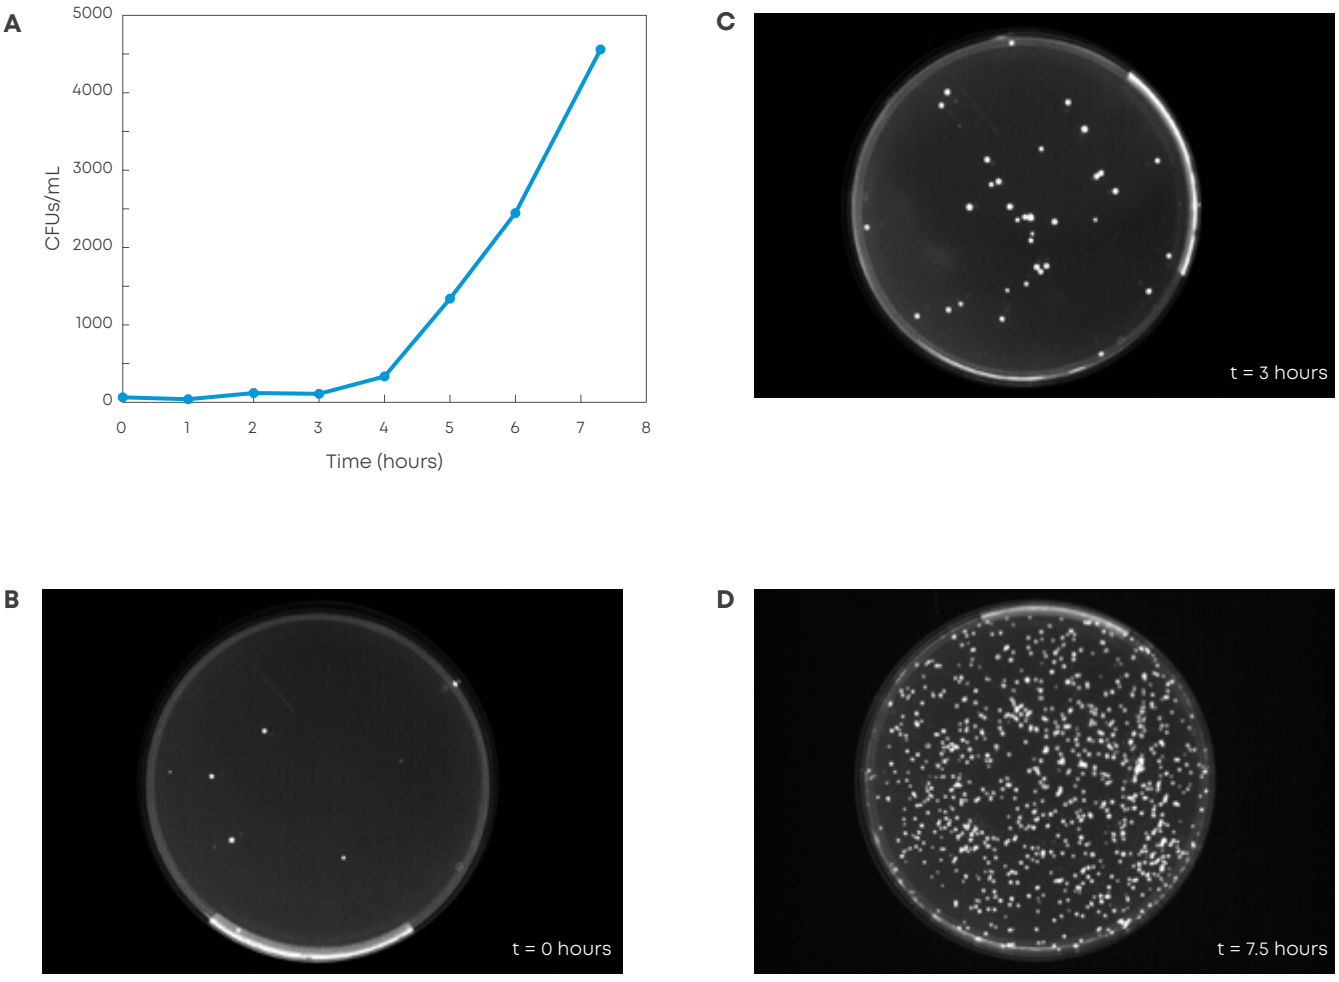

Figure 5: Exponential growth of *E. coli* in cell culture medium in a static bioreactor flask (A). CFUs/mL increases over time from t = 0 hours (B), t = 3 hours (C), t = 7.5 hours (D)

© Cambridge Consultants 2022

We next observed sterile and contaminated medium using PureSentry. Apart from small amounts of pixel noise, the sterile medium produced blank grey images, meaning nothing is changing in the scene and the camera observes no contaminants flowing past. In contrast, the contaminated medium produced pseudoframes with *E. coli* sized objects (approximately one micron diameter) usually with a black leading edge (indicating the *E. coli* are blocking some of the light as they pass a pixel) and a white 'tail' (indicating that the light incident on the pixel increases once the *E. coli* have passed). See Figure 6 for examples of sterile and contaminated pseudoframes captured by PureSentry. We verified these detected objects were indeed *E. coli* cells by comparing a strain of bacteria engineered to express a green fluorescent protein to a non-engineered bacterial strain (MG1655) via the event-based camera. Both strains produced identical signals on the event-based camera, and the fluorescent strain enabled us to confirm spots imaged under brightfield were also visible under fluorescence (data not shown).

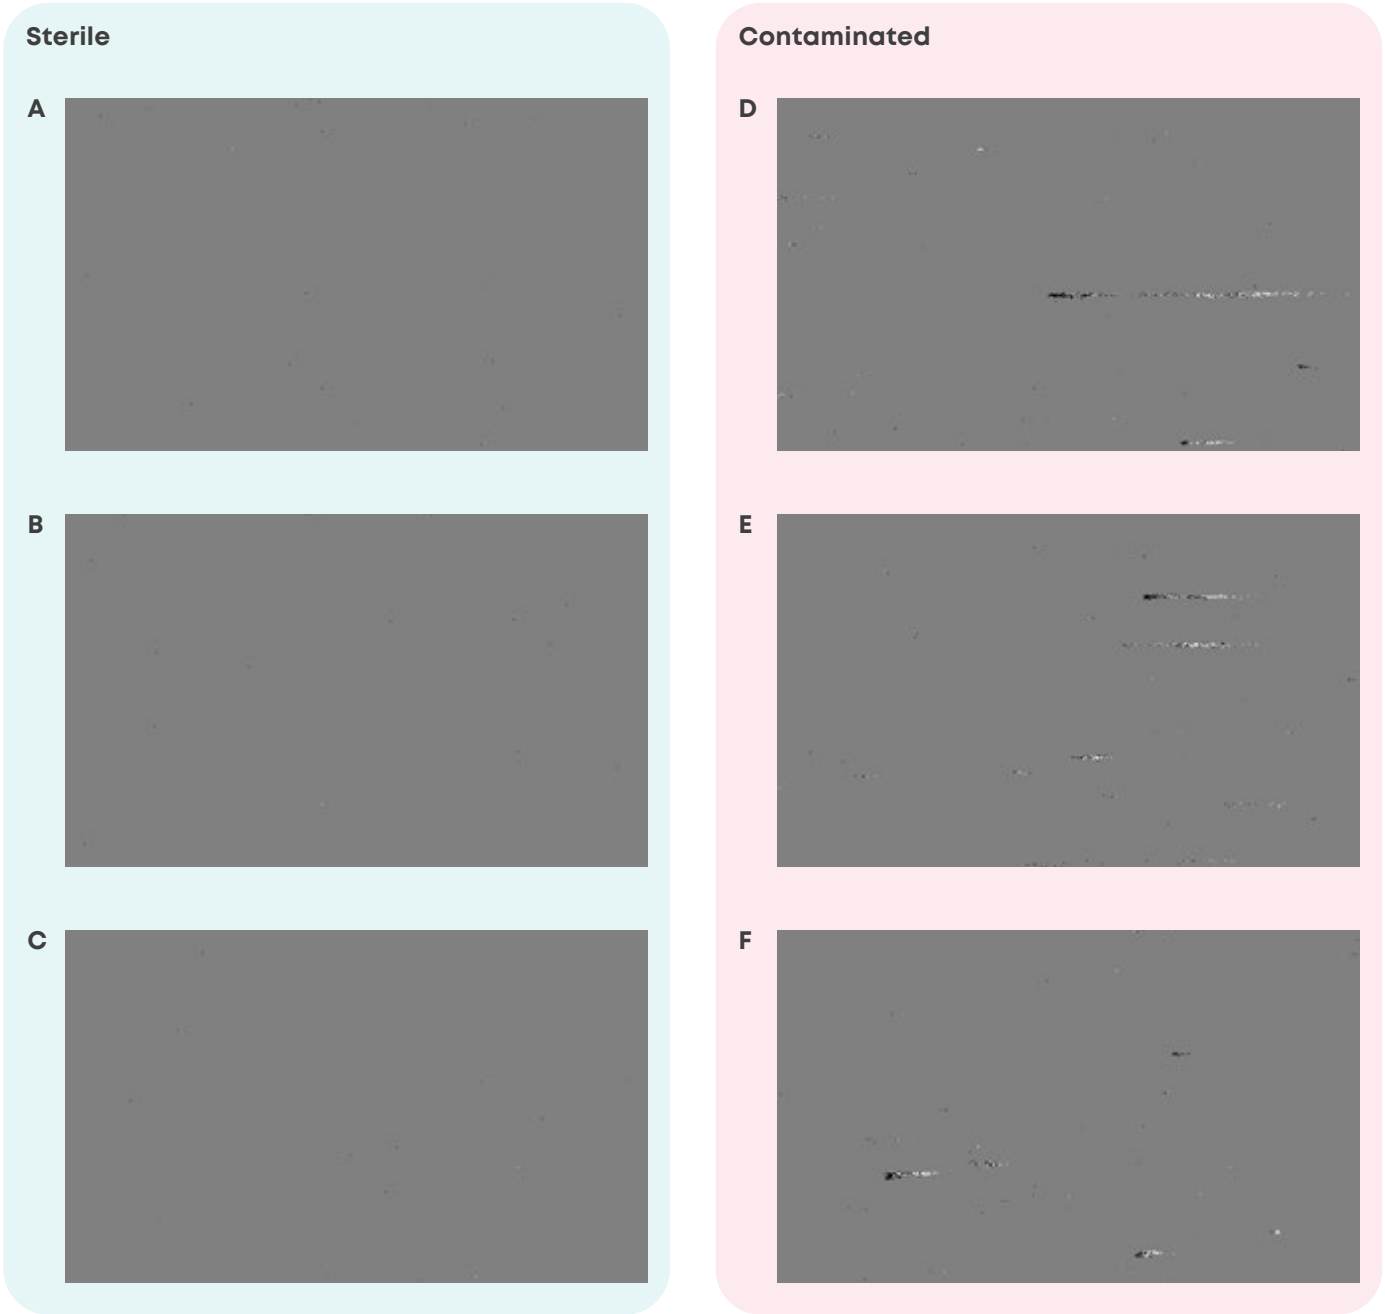

Figure 6: The Prophesee event-based camera detects Gram-negative bacteria in a label-free non-destructive imaging set-up with standard brightfield illumination

© Cambridge Consultants 2022

The first detection algorithm development began with a comparison of several different neural networks well suited to classifying images. We trained a convolutional neural network (CNN), a recurrent neural network (RNN) and a spiking neural network (SNN) on a set of 19,000 images labelled manually as either containing evidence of contaminating bacteria or appearing sterile. The SNN was ruled out as it learned at a slower rate than the other options and did not complete training in time to be compared. Both the RNN and CNN could complete the task, but the CNN was faster to train and quicker at analysing frames. Satisfactory results were achieved with the CNN after 24 hours of training and it achieved high accuracy with a classification latency under 20ms. This speed enabled real-time monitoring with no accumulation of a backlog of data to process. Figure 7 provides an overview of the development and performance of the PureSentry static bioreactor algorithm.

The trained detection algorithm was tested on sterile and contaminated cell culture medium and detected a >100-fold higher signal in contaminated versus sterile samples within just 15 minutes. The cumulative number of contaminated frames rises sharply in a contaminated sample whilst staying relatively flat in a sterile sample. The rate of contaminated frames in the sterile sample indicates the false positive rate and can be used to calculate the accuracy of the CNN (Figure 8B).

"We trained a convolutional neural network (CNN), a recurrent neural network (RNN) and a spiking neural network (SNN) on a set of 19,000 images."

PureSentry static bioreactor algorithm

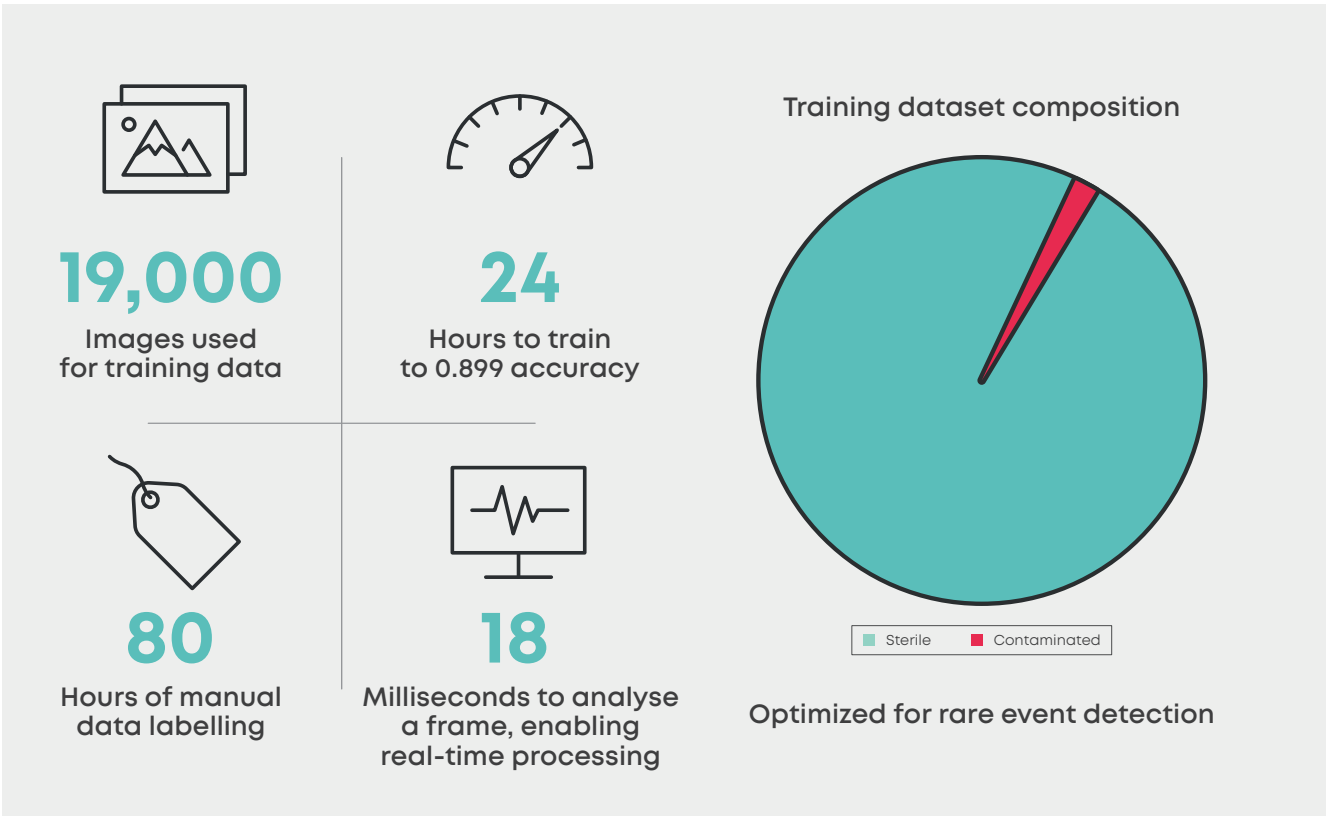

Figure 7: Summary of development and performance of the PureSentry static bioreactor algorithm

© Cambridge Consultants 2022

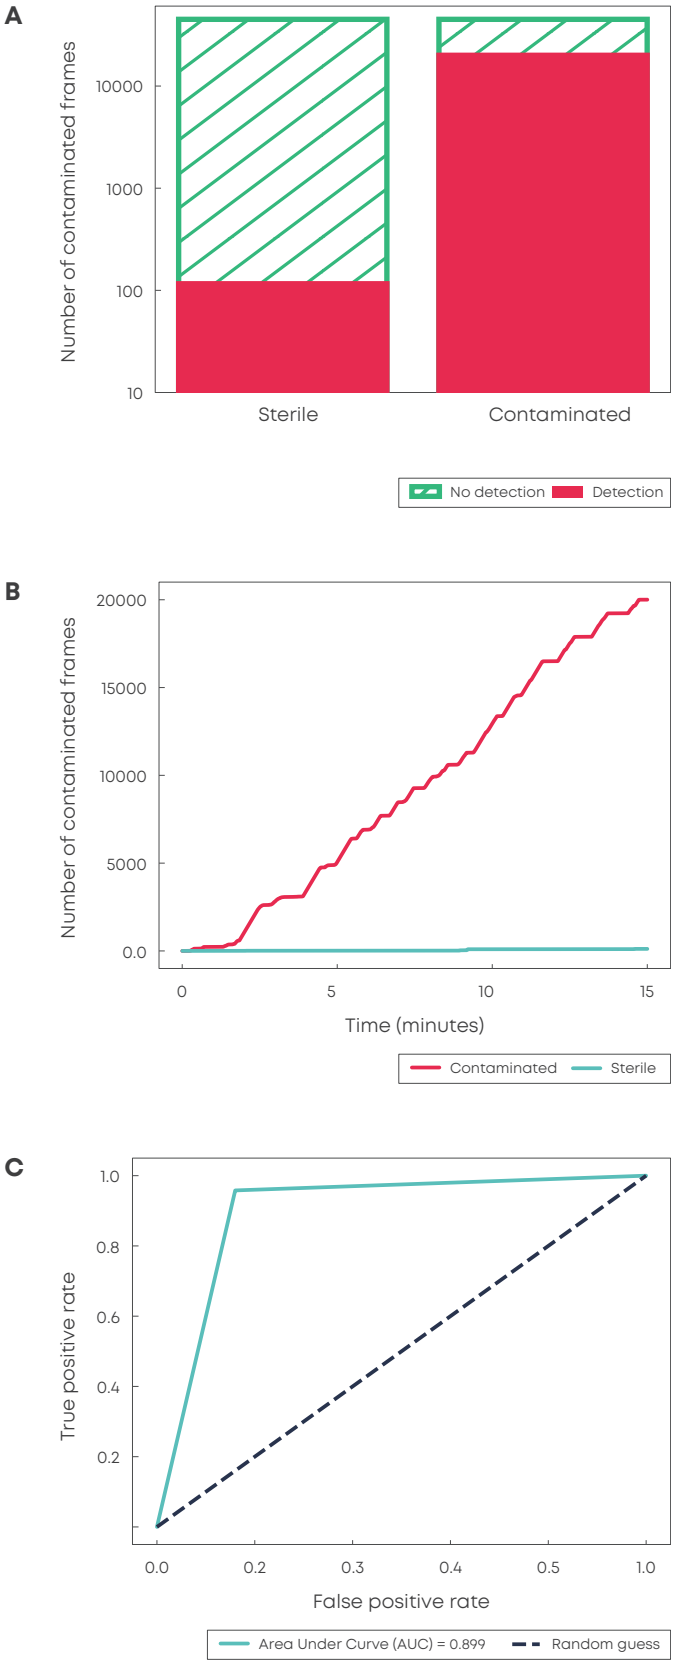

The Receiver Operating Characteristic (ROC) curve is a typical plot for visualising the performance of a machine learning algorithm. The ROC curve is shown in Figure 8C with the true positive rate plotted against the false positive rate for a range of inputs. An ideal algorithm has the maximum true positive rate and minimum false positive rate for all inputs (an L-shape rotated 90 degrees clockwise). Strongly performing algorithms have an inflection point close to x = 0, y = 1, and are very far from the line of y = x, which indicates an algorithm no better than random chance (dashed line). The accuracy of the algorithm is calculated by summing the number of true positives and the number of true negatives and dividing by the total number of frames. We initially achieved an accuracy of 0.966 with a classification latency slightly greater than 20ms. After tuning the algorithm to analyse each frame in under 20ms the final accuracy was 0.899, which demonstrates excellent performance.

In order to relate this accuracy level into biologically meaningful terms, we compared the number of detected events and frames categorised as contaminated in sterile and contaminated runs with the number of bacterial colonies grown on agar plates from the same media.

As seen in earlier experiments, the sterile batch produced very little signal in the vast majority of frames (Figure 9A), but once again small objects were detected in the contaminated batch, representing *E. coli* (Figure 9D). The plots of events and contaminated frames classified over time show very different profiles. The sterile run is almost completely comprised of background noise events and has a handful of frames that PureSentry classified as potentially contaminated in the entire run of 25 minutes (Figure 9B). In contrast, the contaminated run (Figure 9E) shows frequent event spikes which are correctly classified.

Figure 8: PureSentry analyses 45,000 frames in 15 minutes, detecting >100 times more contaminated frames in a contaminated batch than in a sterile batch (A). The cumulative signal is higher relative to the background later in the run (B). Receiver operating characteristic curve for the trained CNN shows it has a high accuracy as the True Positive Rate is much greater than the False Positive Rate across a wide range of false positive rates (C)

© Cambridge Consultants 2022

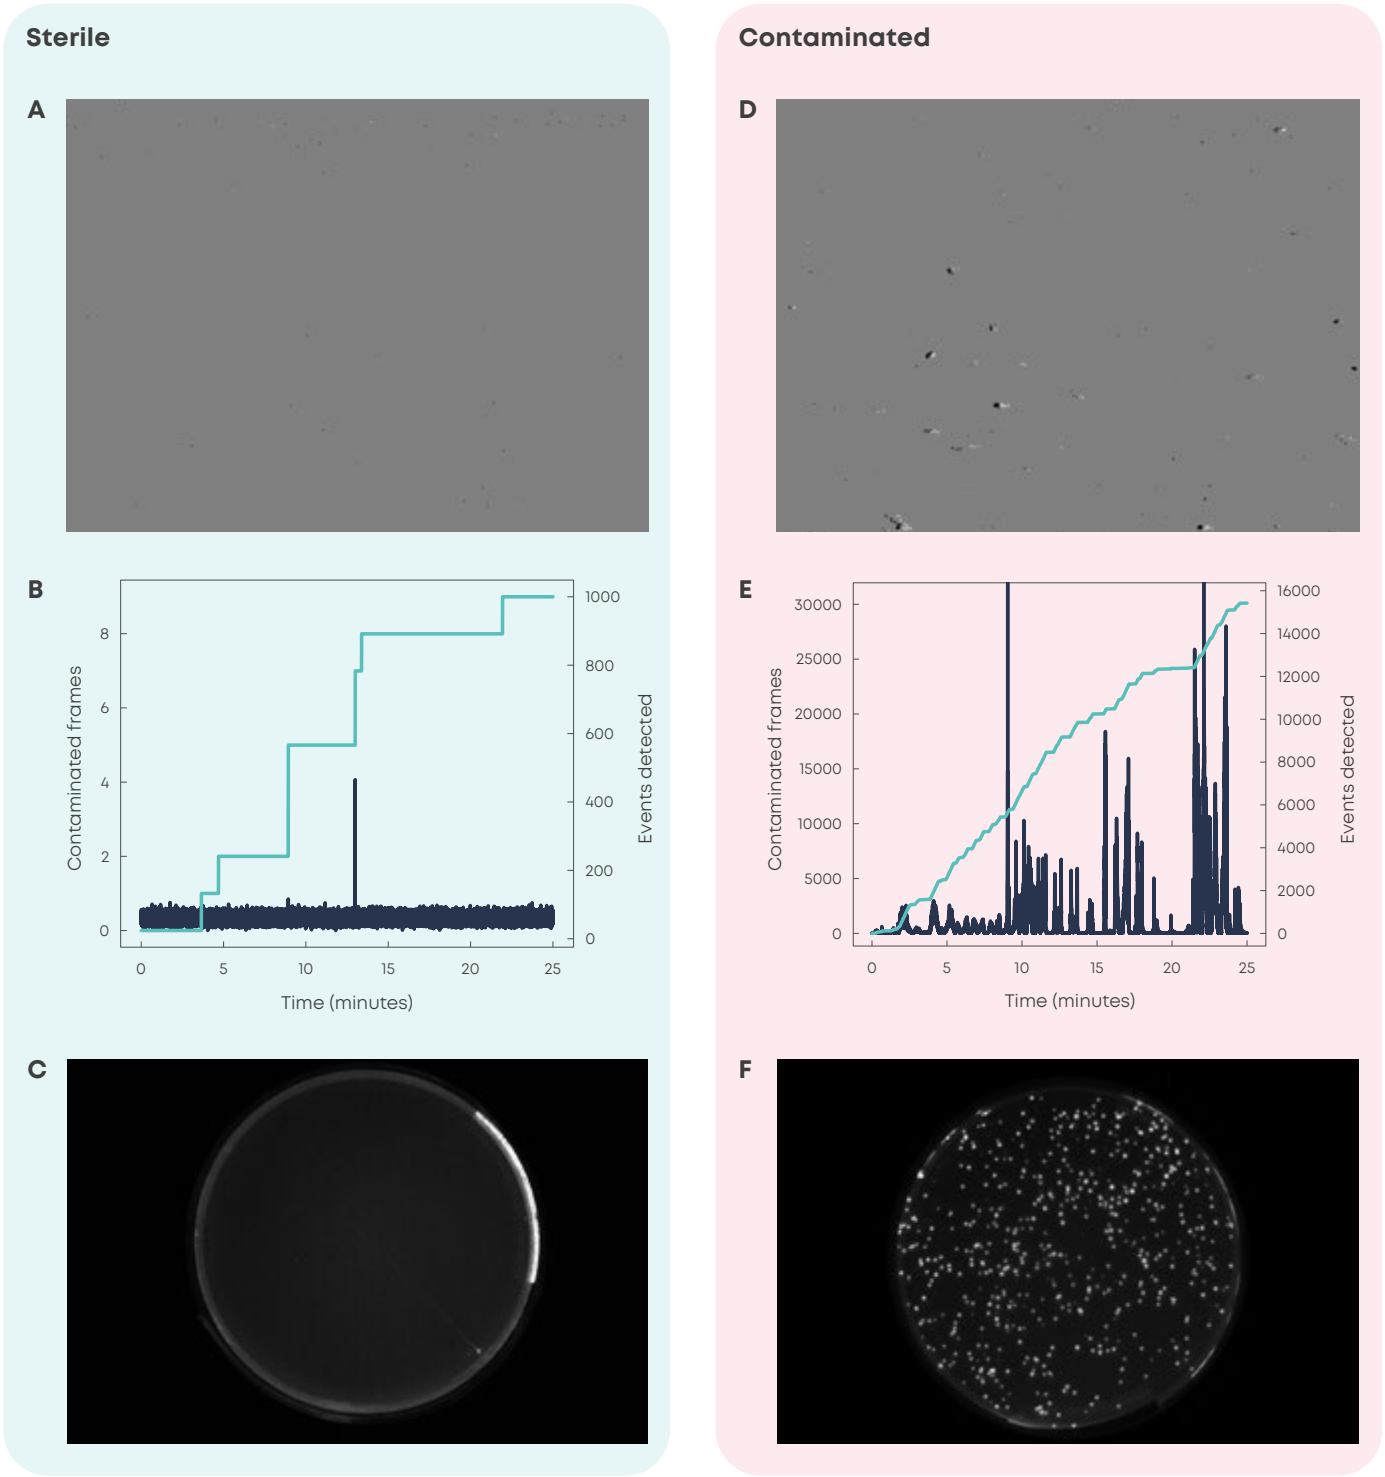

Figure 9: Sterile frame (A), contaminated frame (D), events over time (blue line) and contaminated frames (green line) for sterile medium (B) and contaminated medium (E). Agar plate showing no colony growth for sterile (C) and 750CFUs/mL for contaminated (F)

© Cambridge Consultants 2022

Figure 9E shows several thousand events for many of the frames but each pixel is counted as an event, so this may indicate only a handful of bacterial cells along with their 'tails' of fired pixels as imaged by the event-based camera. Finally, the agar plates confirm the negative control samples were indeed sterile, with no colony growth at all (Figure 9C). For the contaminated batch we observed a mean of 75 colonies (n=3) grown from a 100µL sample (Figire 9F), giving 750CFUs/mL as the starting inoculation level. Both the sterile and contaminated samples were incubated overnight before measurement with PureSentry to simulate the early stage of the cell therapy manufacturing process. The European Medicines Agency (EMA) defines the acceptable limit for parenteral (injectable) drugs such as cell therapies as not more than 10CFUs/mL (see Reference 3) and whilst the PureSentry static bioreactor algorithm is yet to be proven sensitive down to this limit, our experiments with G-Rex bioreactors show that a sample containing 90CFUs/mL will be detectable at >1000CFUs/mL after less than five hours of growth at 37°C. This means the benefits of continuous monitoring are clear even using this early prototype. Furthermore, the effective limit of detection for PureSentry is much lower than 750CFUs/mL on a multiday process, such as CAR-T therapy manufacturing, which typically includes around ten days' worth of cell expansion time.

In order to assess whether the PureSentry static bioreactor algorithm could be applicable for other contaminants of interest, we built a model based on Poisson statistics to determine the time to reliably detect contamination. The model is based on doubling times of a contaminating microbe, a culture volume of 100mL and the volume of medium imaged in our microfluidic chip. Our initial results for the doubling time of *E. coli* (20 minutes, see Reference 4) agreed with our empirical finding that *E. coli* is detectable by PureSentry in around five hours from a starting CFUs/mL of <100. We reviewed the literature on ten contaminating microorganisms that were highlighted by the Cell and Gene Therapy Catapult as being of concern to manufacturers. The organisms were grouped into fast, medium, slow and 'worst case' based on published values for their doubling times, as shown in Figure 10.

Group	Organisms	Doubling time	Time to detect in days		
			1CFUs/mL	10CFUs/mL	100CFUs/mL
Fast	*E. coli*, Staph aureus	30min	0.2	0.1	0.1
Medium	C. sporogenes, B. subtilis, C. albicans, P. aeruginosa, A. brasiliensis	120min	0.9	0.7	0.4
Slow	P. acnes, K. rhizophila	360min	2.9	2.1	1.3
Worst case	Mycoplasma spp.	720min	5.5	4.0	2.5

Figure 10: A Poisson statistics based model of PureSentry's continuous frame classification monitoring approach suggests that even the slowest reproducing contaminants could be detected by PureSentry within the ten day expansion phase of CAR-T manufacture

© Cambridge Consultants 2022

All organisms of concern reach a detectable limit well within ten days of cell culture, meaning this approach is likely to be applicable across a wide range of relevant contaminants and will outperform a traditional approach of taking a small number of samples for off-line growth-based sterility testing. This is a compelling case for cell therapy manufacturing facilities to look at continuous, on-line monitoring as a natural evolution towards smarter manufacturing.

In this section we have introduced a novel, scalable, on-line approach to contamination detection and demonstrated that it operates within relevant sensitivities and timeframes for the cell therapy industry.

5.1.2 Mixing bioreactors

In contrast to static bioreactors, mixing bioreactors such as rocking and stirred tank bioreactors employ agitation to ensure adequate gas exchange occurs to support the respiration of the cells. Sampling from a well-mixed cell culture means that T cells will inevitably be captured in many of the images of the culture medium. Therefore, any detection algorithm must be able to detect contaminants whilst ignoring the cells that are also present. In order to make PureSentry applicable to a wider range of bioreactors, we tested the Prophesee event-based camera's performance detecting T cells label-free and selected a new AI approach, better suited to this scenario, as shown in Figure 11.

"This is a compelling case for cell therapy manufacturing facilities to look at continuous, on-line monitoring as a natural evolution towards smarter manufacturing."

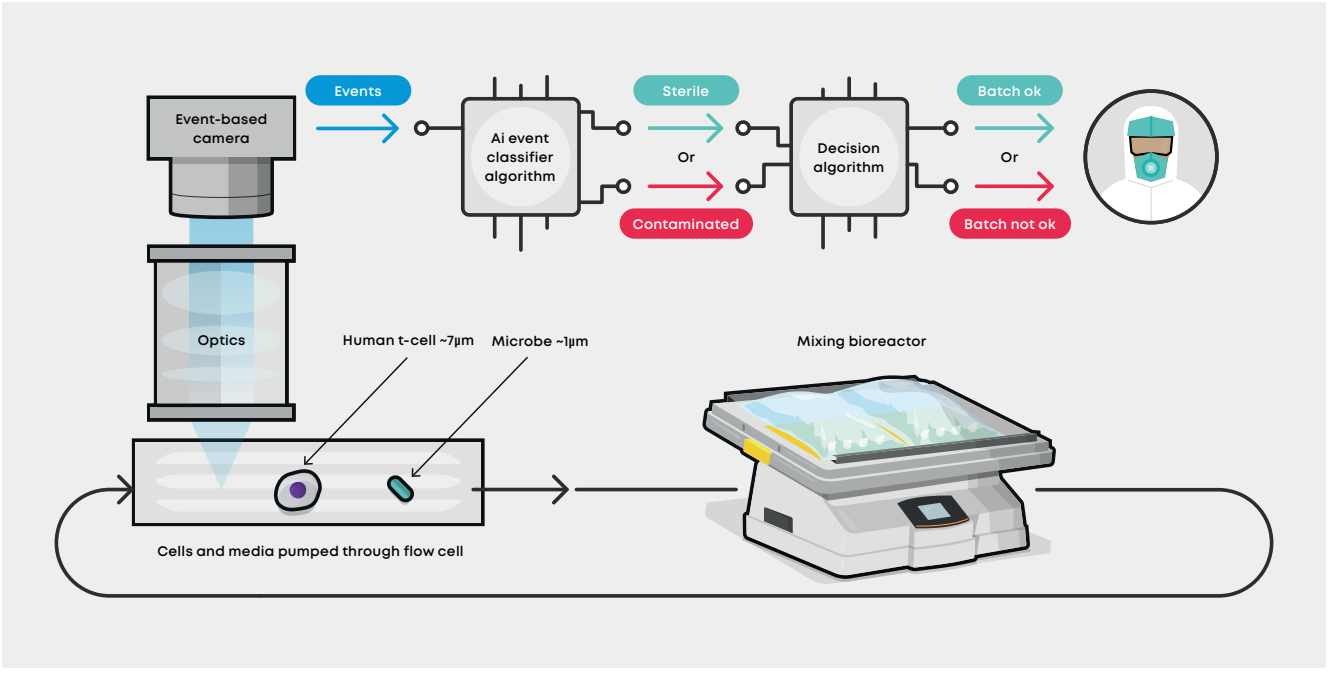

Figure 11: Schematic of PureSentry operating on a mixing bioreactor, where contaminants and T cells frequently occupy the same field of view

© Cambridge Consultants 2022

Firstly, as found with *E. coli* the event-based camera was able to detect T cells reliably without any sample preparation or added labels. T cells detected by the camera are distinct from the model contaminant in that they are larger (approximately seven micrometres diameter versus one micrometre for *E. coli*). The Prophesee camera is able to detect T cells and *E. coli* simultaneously as seen by comparing Figure 12A, Figure 12 B and Figure 12 C (sterile T cell culture) with Figure 12 D, Figure 12 E and Figure 12 F (T cell culture contaminated with *E. coli*).

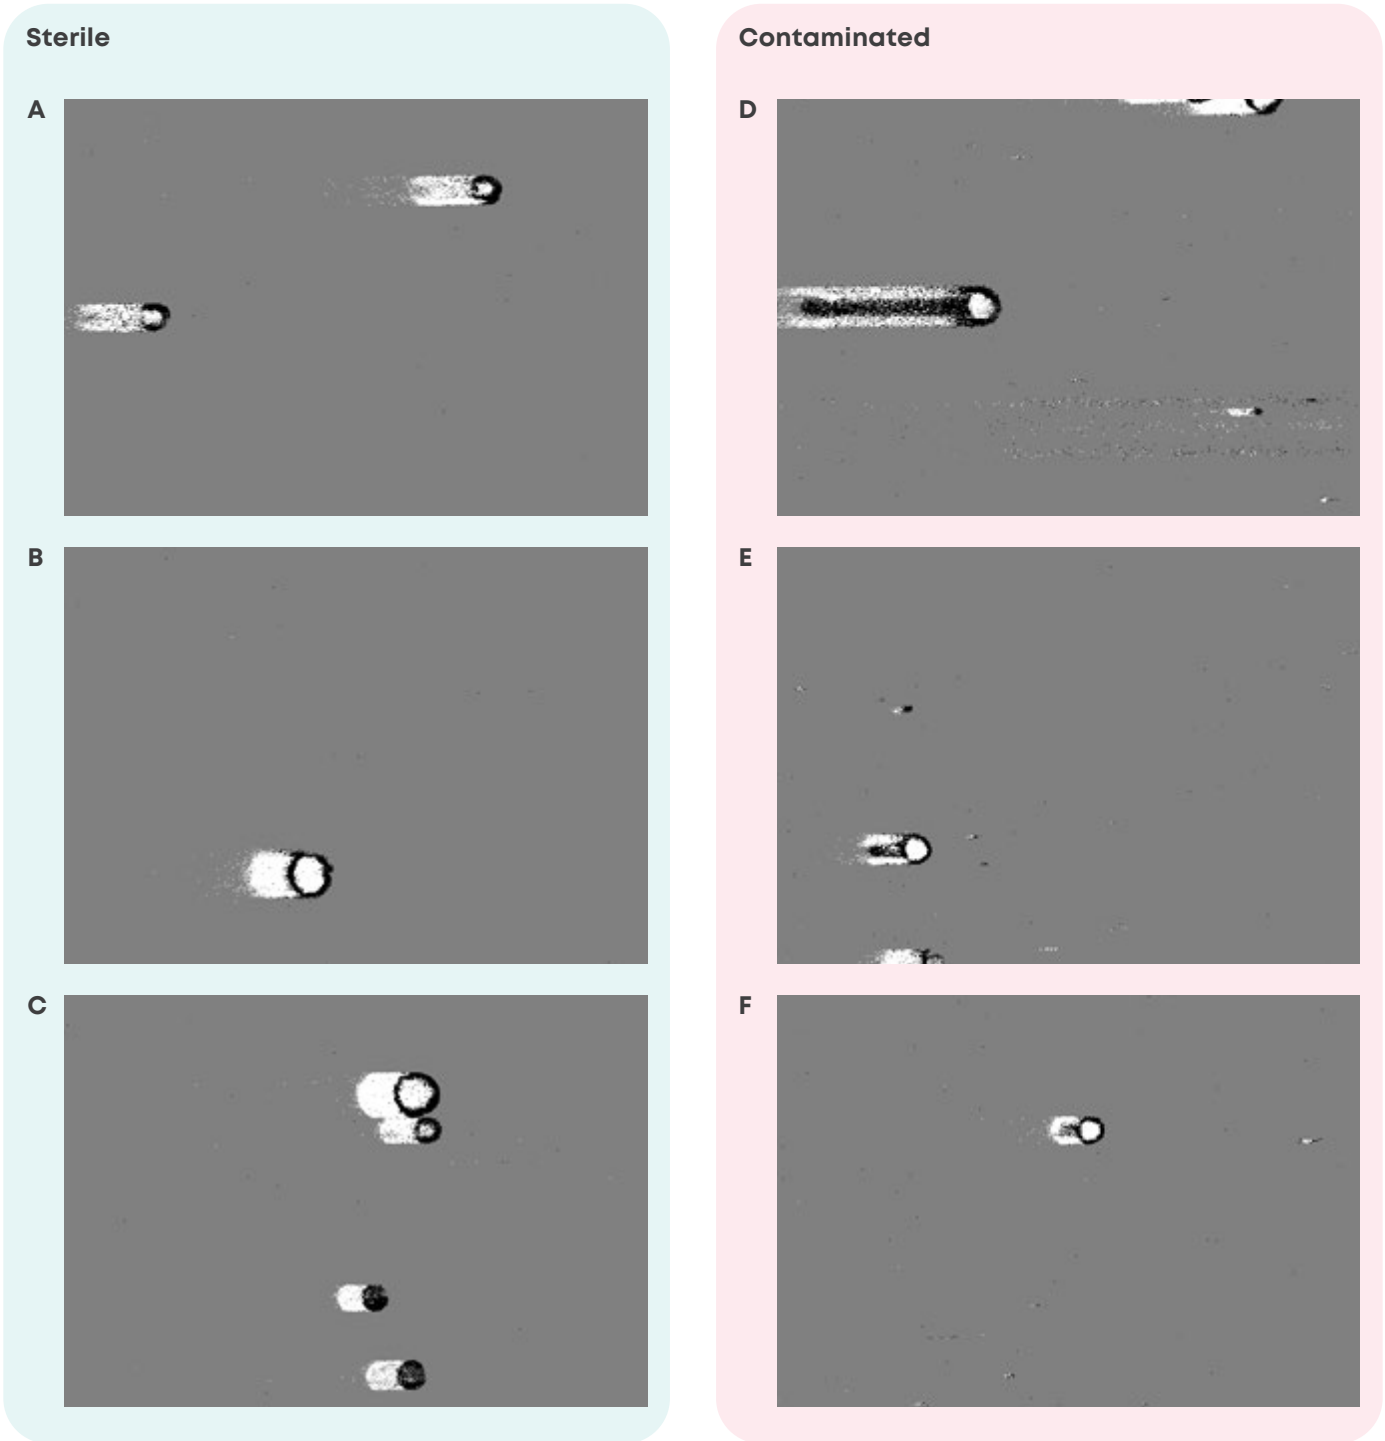

Figure 12: PureSentry detects T cells and *E. coli* simultaneously in a label-free manner

© Cambridge Consultants 2022

As PureSentry is focused on rare-event detection such as the appearance of a small number of contaminating cells within a ten day process, our training data acquisition experiments needed to be many hours at a minimum. Reviewing all of the frames manually is not practical or efficient so we developed a computer vision system to filter out background noise events and segment, measure and count the remaining objects. This system is able to distinguish T cells and *E. coli* based on size as seen in Figure 13 where the sterile T cell culture has a monomodal distribution (except for some noise events) with a modal object size between 10 and 25 micrometres.

The contaminated culture is bimodal, with one similar peak to the sterile culture dataset and an additional peak of small objects representing the additional bacterial cells. This gave us confidence that our data has significant differences that could be addressed by an AI-based detection system. Although this tool in itself provides the user with valuable information about the culture – and could be sufficient to detect contamination in some circumstances – we believe an adaptive algorithm is the best solution. Adaptive algorithms enable the user to train the system based on their cell type of interest (offering applicability to other cell therapy types) and lay the groundwork for detecting multiple species of contaminant in the future.

"This gave us confidence that our data has significant differences that could be addressed by an AI-based detection system."

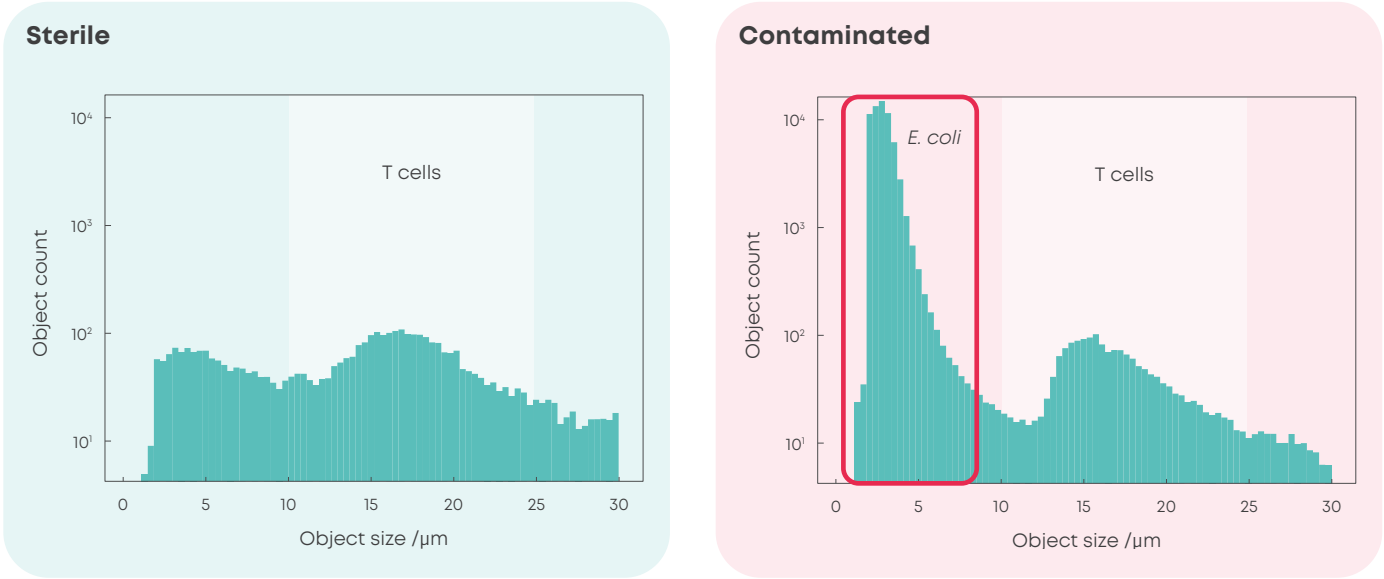

Figure 13: Computer vision detects a distinct bimodal distribution of objects in contaminated T cell cultures

© Cambridge Consultants 2022

The challenge of detecting contaminants in a mixing bioreactor containing millions to billions of human T cells is a harder detection problem than detection in a static bioreactor. To help overcome this challenge we selected a different machine-learning algorithm that we believed would be better able to cope with the added T cell background signal. This method is known as a 'long short-term memory' (LSTM) network and a schematic in Figure 14 provides an overview of how the data is processed. The model receives the same type of event-based camera data, grouped into 20ms 'pseudoframes', though this time T cells are present in both sterile and contaminated datasets. In the contaminated scenario, bacteria are also present. The images are labelled as either belonging to a sterile or a contaminated sample, but the individual frames are not manually labelled to indicate presence of contaminants. This essentially removes the labour-intensive labelling step which was necessary for the static bioreactor detection algorithm.

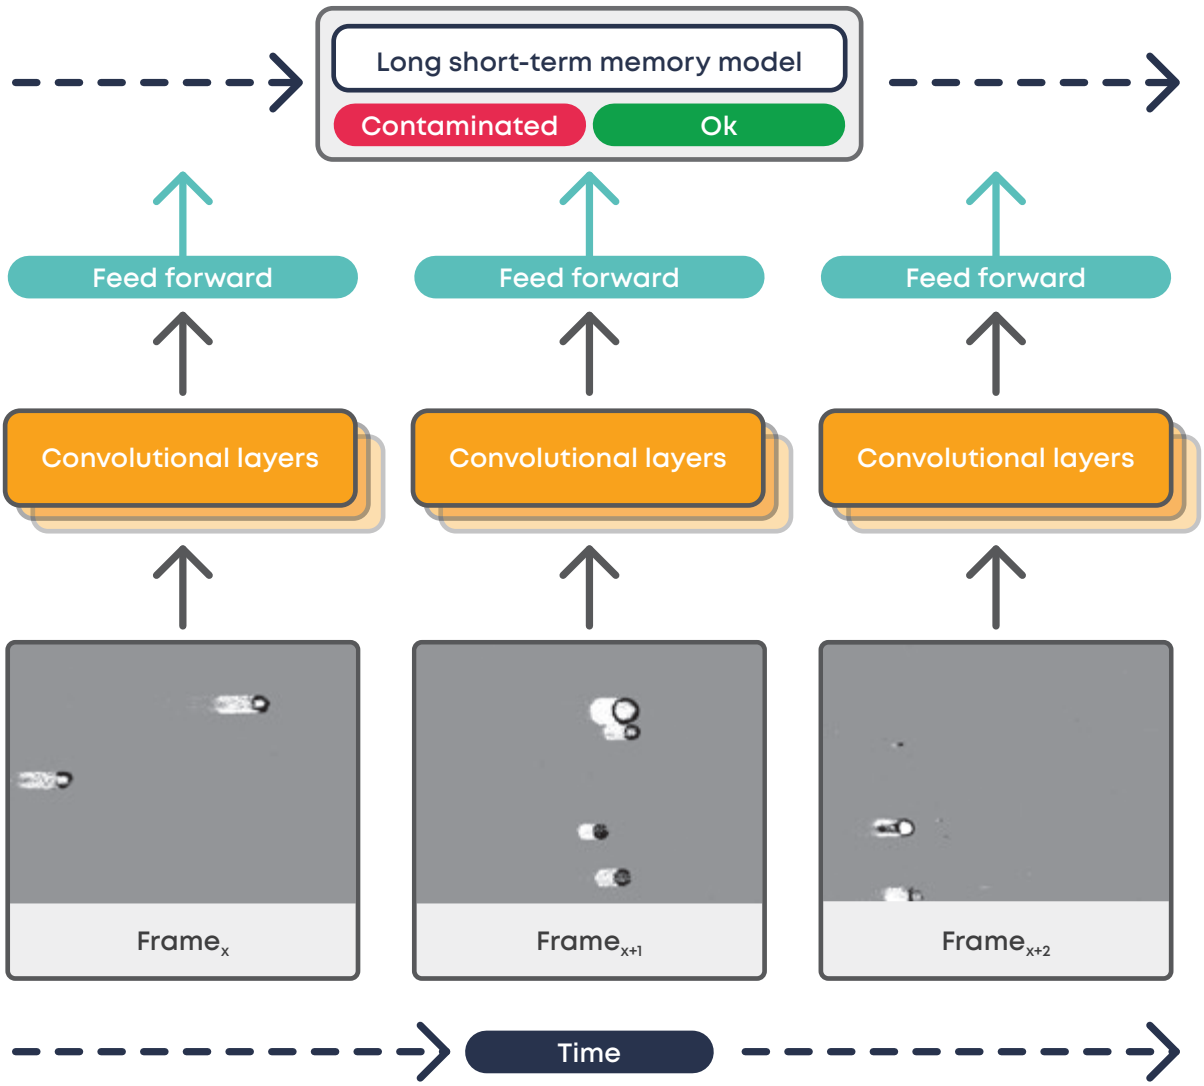

Figure 14: Schematic of the detection algorithm for mixing bioreactors

© Cambridge Consultants 2022

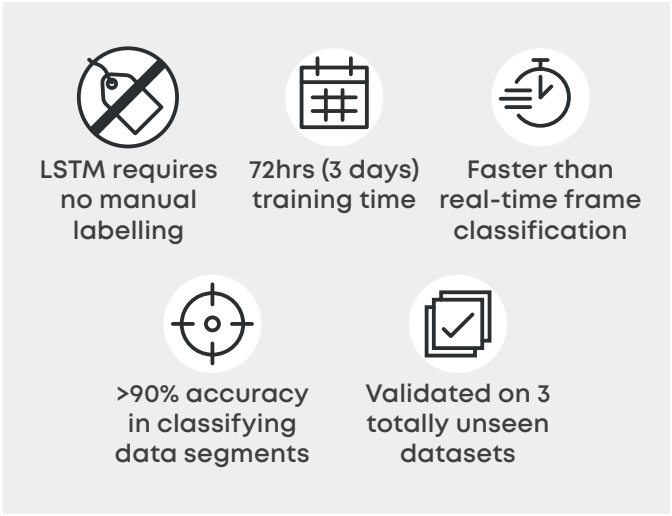

One similarity between the two algorithms is that the data is fed into several convolutional layers in the LSTM model, and as with the CNN, these are mathematical operations that represent the information from the original images in different ways that are not human readable but facilitate the machine learning algorithm in telling apart the two classes of data, sterile and contaminated. The key difference with the LSTM model is that it has an internal memory, meaning that the model updates its own decision-making algorithm with each new frame of data that arrives. This enables the model to understand long term trends and characteristics of the data. In this case, the data was divided into ten second 'movie clips' of the sample that PureSentry analysed. The data as a whole was divided into 90% training data, 5% test and 5% validation data, with each grouping being composed of 50% sterile and 50% contaminated data. After 100 epochs and three days of training, the LSTM model could distinguish sequences of frames from contaminated and sterile T cell cultures with over 90% accuracy.

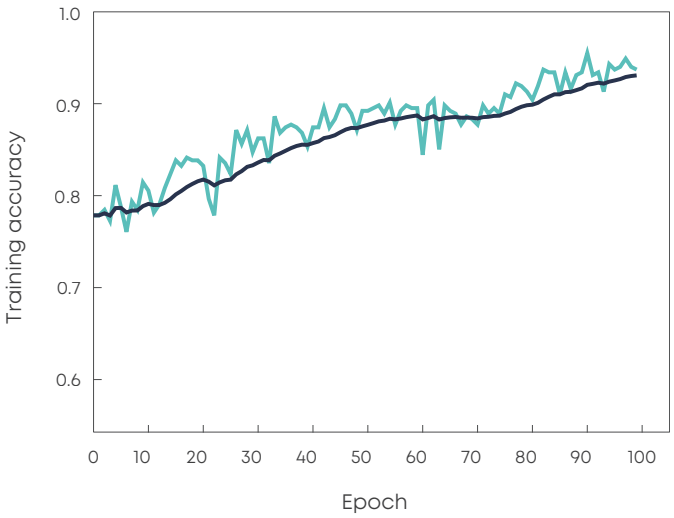

After this initial successful result, we devised a test to ensure the algorithm was able to cope with data outside of its own training set. We used the same PureSentry set-up but performed three biological replicates with new T cell batches and fresh cultures of contaminating microbes. The LSTM model exceeded our expectations, maintaining over 90% accuracy on three entirely new datasets, as summarised in Figure 15. This provides significant confidence in our approach; not only can PureSentry adapt to different bioreactor set-ups, we can also introduce new cell types while still confidently determining if a batch is contaminated. Furthermore, the algorithm requires essentially zero data labelling besides a few carefully prepared training samples.

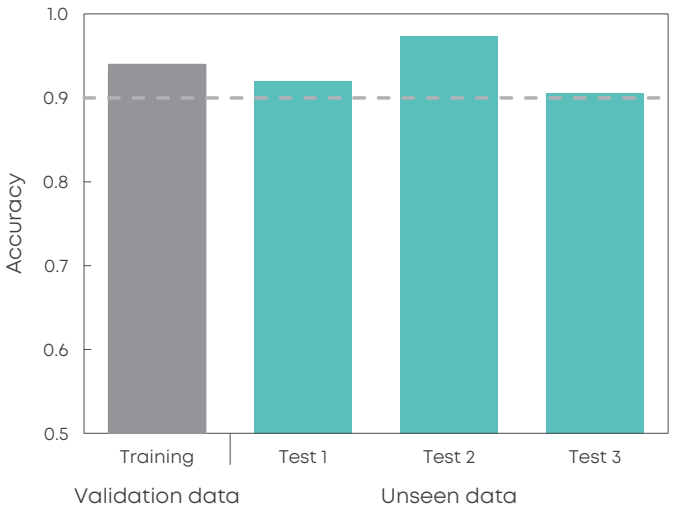

"Not only can PureSentry adapt to different bioreactor set-ups, we can also introduce new cell types while still confidently determining if a batch is contaminated."

Figure 15: Summary of training data and performance for PureSentry mixing bioreactor algorithm

© Cambridge Consultants 2022

6 Be part of the future of cell therapy manufacturing

In conclusion, cell therapy is already delivering astonishing clinical results in the treatment of a range of diseases, but to realise its full potential this treatment needs to be scaled. Before we can do that, manufacture needs to be more automated and more cost efficient. For these reasons, Cambridge Consultants has pushed the latest camera technology into the realm of bioimaging and developed real-time custom detection AI algorithms. This system demonstrates that contaminants which cause cell therapy batches to fail can be detected in a fully automated process at a cost several orders of magnitude below the current cost of a cell therapy treatment. Adopting an on-line monitoring approach in cell therapy could save manufacturers money and help deliver therapies to patients sooner. Moreover, it could ensure that therapies are made at such a scale that anyone who needs one can have access.

If you are focused on solving the pressing problem of automating cell therapy analytics, interested in helping to develop PureSentry further, or have a problem that requires biological, engineering or machine learning expertise to solve, please reach out to us at bioinnovation@cambridgeconsultants.com

© Cambridge Consultants 2022

References

1 Manabu Mizutani, Hazuki Samejima, Hiroshi Terunuma, Masahiro Kino-oka, Experience of contamination during autologous cell manufacturing in cell processing facility under the Japanese Medical Practitioners Act and the Medical Care Act. Regenerative Therapy 2016 Dec;5:25-30. doi: 10.1016/j.reth.2016.06.004.

2 United States Pharmacopeia (USP): <71> Sterility Tests. Pharmacopeial Forum 2009;34(6). Interim Revision Announcement No. 6, December 1, 2008, official on May 1, 2009. https://www.usp.org/sites/default/files/usp/document/harmonization/gen-method/q11_pf_ira_34_6_2008.pdf

3 Yang H, Li N, Chang S. A Risk-based Approach to Setting Sterile Filtration Bioburden Limits. PDA J Pharm Sci Technol. 2013 Nov-Dec;67(6):601-9. doi: 10.5731/pdajpst.2013.00942. PMID: 24265301.

4 Gibson B, Wilson DJ, Feil E, Eyre-Walker A. The distribution of bacterial doubling times in the wild. Proc Biol Sci. 2018;285(1880):20180789. doi:10.1098/rspb.2018.0789.

Authors

Matthew Armean-Jones, Project Manager
Josh Gibson, Lead Engineer

© Cambridge Consultants 2022

About Cambridge Consultants

Cambridge Consultants, part of Capgemini Invent, has an exceptional combination of people, processes, facilities and track record. Brought together, this enables innovative product and services development and insightful technology consulting. We work with companies globally to help them manage the business impact of the changing technology landscape. We're not content to deliver business strategy based on target specifications, published reports or hype. We pride ourselves on creating real value for clients by combining commercial insight with engineering rigour.

We work with some of the world's largest blue-chip companies as well as with innovative start-ups that want to change the status quo fast. With a team of around 800 staff in Cambridge (UK), Boston, San Francisco and Seattle (USA), Singapore and Tokyo, we have all the inhouse skills and lab facilities needed to help you. We are currently collaborating with ambitious pharma and biotech companies to deliver the promise of cell therapies, gene therapies and regenerative medicine to patients. Our expertise covers equipment, process, QC and delivery – and the ability to address vital issues such as cost, scale up and accessibility.

For further information or to discuss your requirements, please contact:

James Hallinan, Head of Business Development, Bioinnovation
james.hallinan@cambridgeconsultants.com

cambridge consultants
Part of Capgemini Invent

UK — USA — SINGAPORE — JAPAN

www.cambridgeconsultants.com

Cambridge Consultants is part of Capgemini Invent, the innovation, consulting and transformation brand of the Capgemini Group. www.capgemini.com

Lit-WHI-056 v0.7 '2022_PureSentry_Whitepaper'