

METHOD FOR DETECTING THE PRESENCE OF A CELL OF INTEREST OR A CELLULAR
FRAGMENT OF INTEREST IN A SAMPLE OF ORGANIC FLUID

5 Cross-Reference to Related Applications

This Patent Application claims priority from Italian Patent Application No. 102023000019245 filed on September 19, 2023, the entire disclosure of which is incorporated herein by reference.

10 Technical Field of the Invention

The present invention relates to a method for detecting the presence of a cell of interest and/or a cellular fragment of interest in a sample of organic fluid. The invention further relates to a control unit for performing the method
15 and a corresponding computer program.

Prior Art

The presence of specific cells or specific cellular fragments in a sample of organic fluid of a person may be indicative of a disorder or a condition in the person and
20 associated prognosis. For example, the presence of Circulating Tumor Cells (CTCs) and tumor cell fragments, e.g. tumor-derived extracellular vesicles (tdEV), in a blood sample of an individual is strongly associated with cancer prognosis.

25 However, CTCs are rare cells whose density in the blood of a patient may be very low, for example of few units, in the blood volume investigated which typically ranges from 1 to 10 milliliter of blood, but also applies to the complete blood volume (~5 liters).

30 Therefore, identification and enumeration of CTCs and tdEV are challenging.

A known technique that allows to identify CTCs and tdEVs in a blood sample is based on processing the blood sample with the CellSearch® system (Menarini Silicon Biosystems), which performs immuno-magnetic enrichment of the blood sample targeting EpCAM; staining of the enriched sample with DAPI, CD45-APC and CK-PE; and selective fluorescent imaging of the stained cell suspension.

Identification and enumeration of CTCs on said fluorescent images is performed by manual review. Only for research purposes, the identification of CTCs and tdEVs can be partially automated through the open-source imaging program ACCEPT, as described in "Circulating tumor cells, tumor-derived extracellular vesicles and plasma cytokeratins in castration-resistant prostate cancer patients" by A. Nanou et al., published on Oncotarget 2018 Apr 10;9(27):19283-19293, DOI: 10.18632/oncotarget.25019.

Currently, identification of CTCs and tdEVs is performed by trained operators through visual assessment, which is a time-consuming procedure that could be affected by subjective interpretations. In addition, agreement between different trained operators is often poor. About 40% of CTCs can be identified with high confidence, but 60% remains uncertain, as described in "How to Agree on a CTC: Evaluating the Consensus in Circulating Tumor Cell Scoring" by L.L. Zeune et al. Cytometry part A, December 2018; 93(12): 1202-1206. A CTC definition is challenging due to the high dimensional decision boundary which has fluent transitions in several dimensions. Furthermore, the decision on inclusion criteria is a tradeoff between precision and recall. Adding more events as CTC can increase the total number of true positives in a sample, but will also increase

the number of false positives.

In the document "Deep learning of circulating tumor cells" by L. Zeunet et al., published on 10 February 2020 on Nature Machine Intelligence
5 (<https://doi.org/10.1038/s42256-020-0153-x>), hereinafter referred as REF1, an improved automated workflow has been proposed.

The method of REF1 proposes the use of a single machine learning algorithm based on classification, namely a deep
10 learning convolutional neural network, for the identification of CTCs and tdEVs.

The neural network of REF1 is trained to identify CTCs and tdEVs in the fluorescence images acquired with the CellSearch® system. The neural network of REF1 is trained by
15 using training images that have been labeled by human reviewers.

The identification of CTCs based on the neural network of REF1 has a better prognostic value than a manual identification of CTCs.

20 However, the Applicant has verified that the detection performance of the method disclosed in REF1 may be improved.

Applying machine learning to classify whether a cell is e.g. a CTC or not and enumerate such cells in a patient sample can represent a useful approach, but the specific
25 machine learning algorithm and the respective training can make the difference to achieve a successful result.

Summary of the Invention

It is therefore an object of the present invention to overcome the drawbacks of the prior art.

30 This object is achieved by a method for detecting the presence of a cell of interest, in particular a rare cell,

and/or a cellular fragment of interest, a respective computer program and a control unit for implementing the method, as defined in the appended set of claims.

Brief Description of the Drawings

5 Fig. 1 shows a flowchart of the method for detecting the presence of cells of interest and/or cellular fragments of interest according to the invention.

 Fig. 2 shows a block diagram of a control unit according to the invention.

10 Fig. 3 shows a detailed flowchart of a classification step of the method of figure 1.

 Fig. 4 shows an example of an image of a marked sample of organic fluid obtained from an individual.

15 Fig. 5 shows examples of thumbnails extracted from the image of figure 4.

 Fig. 6 shows a flowchart of a method for training a first classifier of the control unit of figure 2, according to an embodiment of the invention.

20 Fig. 7 shows a flowchart of a method for training the first classifier of the control unit of figure 2, according to a different embodiment of the invention.

 Fig. 8 shows a detailed flowchart of a step of the method of Fig. 7, according to an embodiment of the invention.

25 Fig. 9 shows a flowchart of a method for training a second classifier of the control unit of figure 2, according to the invention.

 Fig. 10 shows an exemplificative distribution, in a feature space, of objects to be labeled.

30 Fig. 11 shows an exemplificative distribution of metastatic-dominated regions in the feature space of Fig.

10.

Fig. 12 and Fig. 13 show results of a classification according to the invention performed on illustrative images.

Definitions

5 Unless defined otherwise, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art to which this invention pertains. Although many methods and materials similar or equivalent to those described herein may be used
10 in the practice or testing of the present invention, preferred methods and materials are described below. Unless mentioned otherwise, the techniques described herein for use with the invention are standard methodologies well known to persons of ordinary skill in the art.

15 The expression "cells of interest" is intended to indicate cells whose presence and number in an organic fluid sample of a patient may be indicative of the presence of a disorder or a specific condition in the patient.

In a preferred embodiment, cells of interest are rare
20 cells. The expression "rare cells" is intended to indicate cells whose presence and number in an organic fluid sample of a patient may be indicative of the presence of a disorder or a specific condition in the patient and whose density in the organic fluid of the patient is low.

25 For example, the rare cells may be circulating tumor cells (CTCs), whose presence in the blood of a patient for which cancer has been diagnosed may be indicative of the presence of metastasis, or whose presence in the blood of a patient for which cancer is suspected may be indicative of
30 an elevated risk of metastatic disease. The density of CTCs in the blood may be below 200 / ml whole blood, in particular

of few units, or even smaller than one unit, per milliliter of blood. For example, in establishing overall survival for patients with castration-resistant prostate cancer, patients are divided into those with 5 or more, or less than 5, CTCs
5 per 7.5 ml of blood.

For example, the rare cells may be circulating endothelial cells (CEC) which may be indicative of cardiovascular disease, infectious disease, or cancer. Density of CECs in blood ranges from 0 to 200 / ml whole
10 blood, in particular from 0 to 50 / ml whole blood.

For example, the rare cells may be circulating multiple myeloma cells (CMMC). The presence of CMMC and their number in the blood of a patient may be indicative of their presence in the bone marrow at the time of diagnosis. Furthermore,
15 their number may be monitored to evaluate myeloma, as well as the efficacy of different therapies for this tumor. Density of CMMCs in blood ranges from 0 to 20000 / ml whole blood, in particular 0 to 200 / ml whole blood.

Other kinds of rare cells may be fetal cells (such as
20 erythroblast or trophoblasts), tumor-associated fibroblasts, stromal cells, respiratory virus cells, and/or cells collected from cerebrospinal fluid taps.

The expression "cellular fragments of interest" and the like is intended to indicate phospholipid membrane-enclosed
25 structures, in particular cellular particles including e.g. exosomes, microvesicles, apoptotic bodies, whose presence and number in an organic fluid sample of a patient may be indicative of the presence of a disorder, or a specific condition, in the patient.

30 According to a preferred embodiment, the cellular fragments of interest may be rare cellular fragments of

interest, that is cellular fragments whose presence and number in an organic fluid sample of a patient may be indicative of the presence of a disorder or a specific condition in the patient and whose density in the organic
5 fluid of the patient is low.

For example, when the expression "cells of interest" is used to indicate CTCs, then the expression "cellular fragments of interest" may indicate tdEVs, which may also be indicative of metastasizing tumors, in particular carcinoma.

10 For example, when the expression "cells of interest" is used to indicate CECs, then the expression "cellular fragments of interest" may indicate circulating cell fragments, hereinafter indicated as endothelium-derived extracellular vesicles, edEV, which may also be indicative
15 of cardiovascular disease, infectious disease, or cancer.

For example, when the expression "cells of interest" is used to indicate CMMCs, then the expression "cellular fragments of interest" may indicate circulating cell fragments, hereinafter indicated as multiple myeloma-derived
20 extracellular vesicles, mmEV, which may also be used to evaluate myeloma, as well as the efficacy of different therapies for this tumor.

The rare cells may be intended to indicate also clusters of rare cells and/or cellular fragments of interest, for
25 example clusters of CTCs.

By metastatic cancer it is intended a cancer that has developed, or is in the process of developing, secondary malignant growths at a distance from the primary site. By non-malignant cancer or benign cancer, it is intended a
30 cancer that is localized and is not capable of metastasizing.

With benign samples, it is intended samples from

patients who have been either: 1) diagnosed with non-malignant cancer; 2) suspected of cancer but diagnosed with a non-cancer (e.g. infection, cyst, etc.); or 3) assumed to be healthy (e.g. part of a healthy donor program that includes only donors without any suspicion of cancer).

Detailed Description of the Invention

The method according to the invention is for detecting the presence of cells of interest, and/or cellular fragments of interest in a sample of organic fluid, in particular a blood sample, of a patient (individual).

In detail, the following description will be focused on a method for detecting the presence of rare cells, in particular Circulating Tumor Cells (CTC), and rare cellular fragments, in particular tumor-derived Extracellular Vesicles (tdEV) in a sample of organic fluid of a patient. For the sake of simplicity and without loss of generality, in the following description, the expressions "rare cells and/or fragments thereof", "rare cells and/or fragments", "rare cells and/or cellular fragments" and the like are also used to indicate the expression "rare cells and/or rare cellular fragments of interest", wherein the expressions "rare cells" and "rare cellular fragments of interest" have the meaning defined in the above "Definitions" section.

However, the present method may be applied also to detect other kinds of rare cells, such as CEC, CMMC, etc., and/or other kind of cellular fragments of interest, such as edeV, mmEV, etc., as defined in the above "Definitions" section.

Figure 1 shows a flowchart of the method according to an embodiment of the invention.

At a step S01, a sample of organic fluid of a patient,

in particular a blood sample, is provided. Then, the sample of organic fluid is processed so as to enable selective imaging of the sample and detection of the presence of rare cells and/or fragments in the sample.

5 In detail, at a step S02A, the sample of organic fluid may be enriched. For the detection of CTCs and tdEVs, the blood sample may be subject to immunomagnetic enrichment targeting EpCAM.

 At a step S02B, the sample of organic fluid is marked
10 with at least one marker specific for the rare cell, here in particular a CTC, and/or the fragment, here in particular a tdEV, to be detected, and/or optionally with at least one marker specific for cells and/or fragments other than the rare cell and/or the fragment to be detected.

15 The markers are identifiable by selective imaging.

 According to a preferred embodiment, the markers are fluorescent markers that can be imaged by fluorescence imaging. For example, for detection of CTCs and/or fragments, the enriched sample may be stained with one or more
20 fluorescent markers, in particular DAPI, CD45-APC and CK-PE.

 However, processing of the sample may be different, according to the specific application, in order to allow detection of the presence or absence of the markers of interest, for example using scatter markers or temporal
25 behavior in response to a stimulus, or any other known method.

 At a step S03, images of the marked sample are captured, in a per se known way.

 In detail, in this embodiment, the images are
30 fluorescent images.

 For example, the marked sample may be arranged within

a cartridge and digital images of the surface for each of the fluorophores may be recorded.

According to a preferred embodiment, one or more of steps S02A, S02B and S03, in particular all steps S02A, S02B
5 and S03, may be performed by the CellSearch System®. The step S03 could be performed by other imaging systems such as the DEPArray™.

Then, at a step S04, the captured images are processed by a control unit 20, of which a block diagram is shown in
10 figure 2, in order to determine the presence of CTCs and/or tdEVs in said images and, therefore, in the sample of organic fluid of the patient.

In the embodiment illustrated in figure 1, the method further comprises a further step S05 at which the control
15 unit 20 provides, based on the output of the image classification of step S04, a rare-cells counter indicative of the number of rare cells, e.g. the number of CTCs, and/or the number of cellular fragments, e.g. the number of tdEVs, that have been detected in the organic fluid sample of the
20 patient.

The control unit 20 comprises a first segmentation module 22, a first classifier 24, a second segmentation module 25, a feature extraction module 26 and a second classifier 28, operatively coupled together and described in
25 detail hereinafter.

In the embodiment illustrated, the control unit 20 further comprises an interface unit 29 for coupling the control unit 20 with external devices, which may be useful for performing the method according to the invention.

30 In the embodiment illustrated, the control unit 20 further comprises also labeling units 30M and 30A, which are

units to enable manual labeling (30M) or automated labeling (30A), and which may be used for training the classifiers 24 or 28, as discussed in detail hereinafter.

Step S04 is described in detail hereinafter with reference to the flowchart of figure 3. For sake of simplicity and without loss of generality, step S04 is described with reference to one of the images that are captured from the sample at step S03. It is clear that step S04 may be performed on the whole set of images that are captured from the sample and that represent the entire sample.

In detail, at a step S10, the control unit 20 receives the images that are captured at step S03, for example through the interface unit 29.

The images may be loaded in the control unit 20 either automatically or by an operator, e.g. medical staff.

Figure 4 shows, for illustrative purpose only, an example of an image IMG that is received by the control unit 20.

The image IMG is a digital image, for example scanning an area of $700\text{ }\mu\text{m} \times 900\text{ }\mu\text{m}$ of the sample imaged at step S03.

The image IMG may be a single-color image or a multicolor image, for example RGB-coded, depending on the specific method and equipment used for acquiring the image IMG and the specific post-processing steps performed on the image IMG.

In the example of figure 4, for sake of simplicity, the image IMG is a multicolor, RGB-coded image.

However, the image IMG may be a multi-dimensional image, i.e., formed by a plurality of sub-images, with at least as many sub-images as the number of different markers to be

identified.

In detail, in this embodiment, where the images are fluorescent images, the image IMG may be formed by a sub-image for each fluorescent channel, e.g. four sub-images
5 corresponding each to a respective marker channel of interest. However, the number of channels may be higher or lower, depending on the specific application.

The image IMG comprises one or more objects or events
32 (of which zoomed-in examples are shown in figure 5 and indicated individually as OBJ₁, OBJ₂, OBJ₃) to be classified,
10 i.e. to be assessed whether they represent a CTC, a tdEV or a different non relevant object.

Each object OBJ_i corresponds to a respective area of the image IMG representing a potential candidate for being
15 classified as a CTC or a tdEV. For example, with reference to the zoomed-in images of figure 5, the objects OBJ₁, OBJ₂, OBJ₃ correspond each to the portion of the image within a respective segmentation outline.

The objects OBJ are portions of the image IMG comprising
20 one or more signals from the markers used in step S02 for marking the sample.

In practice, the objects OBJ are marker-stained objects.

Then, at a step S12, the first segmentation module
25 receives the image IMG and runs a segmentation algorithm on the image IMG.

The segmentation module 22 identifies, on the image IMG, the objects OBJ to be classified.

In detail, the segmentation module 22 identifies the
30 outlines (or contours) of the objects OBJ to be classified.

Then, a plurality of thumbnails TBN[1:N] can be

extracted from the image IMG by the control unit 20.

In detail, at step S13, a thumbnail TBN_i is gathered for each object OBJ_i that has been identified at step S12. Each thumbnail TBN_i is a portion of the image IMG that
5 contains one object OBJ_i to be classified. Additionally, each thumbnail TBN_i may contain one or more objects that should not be classified.

By way of example only, three thumbnails TBN_1 , TBN_2 , TBN_3 are shown in figure 5.

10 Each thumbnail TBN_i is an image, in particular a portion of the image IMG, containing one of the objects OBJ to be classified.

For example, each thumbnail TBN_i may have a size of 80 pixels x 80 pixels; however, the size of each thumbnail TBN_i
15 may be chosen depending on the specific application.

For example, each thumbnail TBN_i may be centered around the respective object OBJ_i to be classified.

For example, each thumbnail TBN_i may have a size that encompasses the object OBJ_i with some margin of zero or more
20 pixels around the object OBJ_i .

The segmentation module 22 is configured to implement a per se known segmentation algorithm. For example, the segmentation algorithm may be a machine learning algorithm, in particular an artificial neural network such as a
25 convolutional neural network or a semantic segmentation network, or a different type of algorithm, for example an active contour method, a thresholding method, or a localized contrast method.

For example, the segmentation may be performed using
30 the ACCEPT toolbox developed by Leonie Zeune et al and available at <https://github.com/LeonieZ/ACCEPT>.

According to a preferred embodiment, the segmentation algorithm may be based on the StarDist algorithm available at <https://github.com/stardist/stardist>. A detailed implementation example of the StarDist algorithm may be found, for example, in the document by Michiel Stevens et al, "StarDist Image Segmentation Improves Circulating Tumor Cell Detection", published on 13 June 2022 (<https://doi.org/10.3390/cancers14122916>).

Then, step S14, the first classifier 24 receives the thumbnails $TBN[1:N]$ and, for each thumbnail TBN_i , runs a first classification algorithm on the thumbnail TBN_i in order to classify the respective object OBJ_i contained therein.

In practice, the first classifier 24 runs the first classification algorithm starting from an image containing an object to be classified, in particular here from the thumbnail TBN_i containing the object OBJ_i to be classified.

The first machine learning algorithm is trained to identify if the marker-stained object OBJ_i is a CTC or a tdEV. In other words, the first classification algorithm is a machine learning algorithm trained to identify the presence of a CTC and/or a tdEV in an image. In practice, the first classifier 24 may have been trained to identify if a marker-stained object in an image is a rare cell (e.g., CTC), or trained to identify if a marker-stained object in an image is a cellular fragment (e.g., tdEV), or trained to identify if a marker-stained object in an image is a rare cell, a cellular fragment or a different object.

In detail, in this embodiment, the first classification algorithm is trained with labelled training data (in particular thumbnails with object outline), by using a training model (or approach) wherein labelling of the

training images is performed, at least in part, manually, i.e. provided by an expert operator, e.g. medical staff. In other words, labeling of the training images of the first classifier 24 is operator-assisted.

5 Different embodiments of the training method of the first classifier 24 will be described in detail hereinafter with reference to figures 6, 7 and 8.

According to an embodiment, the first classifier 24 is configured to implement a deep learning algorithm, such as
10 an artificial neural network, in particular a convolutional neural network.

The first classifier 24 may be a binary classifier, i.e. trained to classify the object OBJ_i of the input thumbnail TBN_i in one of two classes, for example 'CTC' and
15 'not-CTC', or a multiclass classifier, i.e. trained to classify the input thumbnail TBN_i in one of multiple classes.

According to a preferred embodiment, the first classifier 24 is configured to classify the input thumbnail TBN_i among five different classes, namely:

- 20 - 'CTC', which indicates that the object OBJ_i is a CTC;
- 'tdEV', which indicates that the object OBJ_i is a tdEV;
- 'WBC', which indicates that the object OBJ_i is a white blood cell;
25 - 'bare nucleus', which indicates that the object OBJ_i is only the nucleus of a cell; and
- 'other object', which indicates that the object OBJ_i is an object that does not fall under any of the other four classes.

30 An example of the architecture of the neural network implemented by the first classifier 24 may be found in the

document, hereinafter referred to as document REF1, by Leonie L. Zeune et al, "Deep learning of circulating tumor cells", published on 10 February 2020 on Nature Machine Intelligence (<https://doi.org/10.1038/s42256-020-0153-x>).

5 For example, the first classifier 24 may be the neural network identified as "standard CNN" in document REF1.

 For example, the architecture of the neural network implemented by the first classifier 24 may be formed by four convolutional layers each followed by max-pooling and then
10 a fully connected network for classification.

 The first classifier 24 outputs, for each thumbnail TBN_i , the class that is assigned to the respective object OBJ_i .

 In detail, the first classifier 24 may output, for each
15 thumbnail TBN_i a plurality of probability values, one for each class. In this case, the class having the highest probability value is the one that is assigned to the object OBJ_i . Alternatively, a class may be assigned only if one of the probability values exceeds a threshold.

20 Then, step S16, the control unit 20 checks whether the output of the first classifier 24 has classified the object OBJ_i as a CTC or a tdEV.

 In the embodiment of figure 3, if the first classifier 24 has not classified the object OBJ_i as a CTC or as a tdEV,
25 branch N from step S16, then (step S17) a counter for the corresponding class is incremented (for example if the first classifier 24 has classified the object OBJ_i as a white blood cell, then a white blood cell counter is incremented). In other words, in this embodiment, at branch N from step S16,
30 the classification of the object OBJ_i by the first classifier 24 is assumed to be correct and is not subject to a further

validation by the second classifier 28.

At step S17, the control unit 20 may also provide an output signal indicating that the object OBJ_i in thumbnail TBN_i is not a CTC or a tdEV, so as to warn an operator.

5 After step S17, the method returns to step S13 and the control unit 20 repeats step S14 for the next thumbnail TBN_{i+1} of the image IMG.

In case the first classifier 24 classifies the object OBJ_i , branch Y from step S16, as a CTC or a tdEV, then the
10 output of the first classifier 24 is confirmed by running a second classification algorithm starting from the thumbnail TBN_i , by the second classifier 28.

The second classification algorithm is a machine learning algorithm trained to identify if the marker-stained
15 object OBJ_i in an image is a CTC or a tdEV. In other words, the second classification algorithm is a machine learning algorithm trained to identify the presence of a CTC and/or a tdEV in an image. In practice, the second classifier 28 may have been trained to identify if a marker-stained object
20 in an image is a rare cell (e.g., CTC), or trained to identify if a marker-stained object in an image is a cellular fragment (e.g., tdEV), or trained to identify if a marker-stained object in an image is a rare cell, a cellular fragment or a different object.

25 The second classification algorithm has been trained with training data (e.g. objects in an image) that has been labelled using a different approach with respect to the first classification algorithm.

In detail, in this embodiment, the second classifier 28
30 is trained using training objects whose labelling is provided in an automated manner, in particular based on a result of

a nearest neighbors analysis on the training objects.

The training method of the second classification algorithm will be described in detail hereinafter with reference to figure 9.

5 In detail, in this embodiment, the confirmation of the output of the first classifier 24 comprises performing a new segmentation (step S18), feature extraction (step S20) and running a second classification algorithm (step S22).

Step S18 is optional.

10 Step S18 is performed by the second segmentation module 25.

According to an embodiment, at step S18, the segmentation module 25 applies a thresholding method to find the outline of the signal in each channel (for example the
15 fluorescent signal in each sub-image of the image IMG) which is located within the outline of the objects to be classified OBJ_i that has been identified at step S12 by the first segmentation module 22.

For example, the applied threshold may be equal to the
20 thumbnail background level plus the maximum of the following three values: $\underline{a} \cdot I_p$, wherein I_p is the peak intensity relative to the background, with $\underline{a}$ being, for example, equal to 0.25; $\underline{b} \cdot B_v$, wherein B_v is an estimation of the background variance, with $\underline{b}$ being, for example, equal to 2.4; and $\underline{C}$, with for
25 example $C=2$, if the first two values are too small (e.g. smaller than $\underline{C}$) so as to suppress noise segmentation.

Step S18 may be useful in case the first segmentation (step S12) enforces a star-convex shape, even though not all the objects OBJ to be identified may have a star-convex
30 shape. Moreover, the segmentation of step S18 may be useful to increase the accuracy of the outlines of the objects OBJ

to be identified, since some of the features extracted by module 26 may depend on the outline shape.

At step S20, the feature extraction module 26 extracts, for the object OBJ_i , a feature vector FTR_i representing the
5 object OBJ_i . The feature vector FTR_i may be extracted from the thumbnail TBN_i and/or be based on the class probabilities that are output by the first classifier 24.

The feature vector FTR_i comprises data extracted from the thumbnail TBN_i that are indicative of the presence of a
10 CTC or a tdEV in the thumbnail TBN_i .

Each object OBJ_i has a corresponding outline (or contour) that delimits the object OBJ_i to be classified in thumbnail TBN_i . In practice, each thumbnail TBN_i is formed by three portions: the area of the thumbnail within the
15 outline, which represents the object OBJ_i to be classified, the area of the thumbnail outside the outline, but within other outlines to be classified, and the area of the thumbnail outside all outlines which represents the background of the thumbnail.

20 The feature vector FTR_i may comprise any one, or a combination, of the following data extracted for object OBJ_i : pixel intensity, sharpness, shape and size of the object OBJ_i to be classified, uniqueness, overlap and similarity between the channels of the thumbnail TBN_i , class probabilities
25 output from step S14, and data indicative of the background of the thumbnail TBN_i .

In detail, the features indicative of the pixel intensity of object OBJ_i may comprise any one or more of: the minimum intensity; the median intensity; the mean
30 intensity; 90th percentile and maximum intensity of the pixels inside the channel outline, inside the event outline,

and outside the event outline, wherein the channel outline is the outline determined for a single channel at step S18 and the event outline is the outline for the whole event on which the classification of step S14 has been made.

5 According to an embodiment, all intensity values are relative to the intensity of the background of the thumbnail TBN_i ; this can increase the usefulness of the intensity measures, since background can vary more than the difference between positive and negative staining.

10 The features indicative of the sharpness of the object OBJ_i to be classified may comprise the mean L2 norm of the Sobel filters in lateral directions along the channel segmentation contour.

15 The features indicative of the shape and size of the object OBJ_i to be classified may comprise any one or more of: length of the object OBJ_i along the major axis and/or the minor axis of the object OBJ_i ; area and/or convexity of the object OBJ_i ; perimeter of the outline of the object OBJ_i ; perimeter of convex hull; and various ratios thereof.

20 The features indicative of the uniqueness of the object OBJ_i to be classified may comprise the portion of area, for example a percentage value, of the thumbnail TBN_i indicating how many pixels within the channel segmentation outlines lie inside the event outline. For example, if the segmentation
25 outline of one channel lies completely within the event outline, then the uniqueness of thumbnail TBN_i may be 100%. Conversely, if none of the channel segmentation outlines overlap with the event outline, then the uniqueness of thumbnail TBN_i may be 0%.

30 The features indicative of overlap and similarity between the channels of the object OBJ_i to be classified may

comprise the intersection of two channel segmentations in pixels and as fraction of total area, intersection over union of two channels, and channel similarity measures such as the root mean square error of two channels each scaled from min-max to [0,1].

The features indicative of the probabilities of the object OBJ_i may comprise the individual class probabilities output by the first classifier 24 at step S14; the next highest probability, i.e. the second highest class probability output by the first classifier 24; and entropy of all class probabilities p_m defined as the sum over all classes of $-p_m \cdot \ln(p_m)$.

The features indicative of the background of the thumbnail TBN_i may comprise one or more of: the 20th, 50th, and 60th percentile for each channel of all pixels of the thumbnail TBN_i that are not inside any channel segmentation of the first segmentation at step S12.

At step S22, the second classifier 28 receives the feature vector FTR_i of the object OBJ_i and runs the second classification algorithm starting from the feature vector FTR_i .

The second classifier 28 is trained to provide an output indicating whether the object OBJ_i is a rare cell and/or whether the object OBJ_i is a cellular fragment. In other words, the output of the second classifier 28 is used to validate (confirm), i.e., maintain or revert, the output of the first classifier 24.

In detail, in this embodiment, the second classifier 28 comprises two classifiers, namely a CTC classifier 28A and a tdEV classifier 28B.

The CTC classifier 28A is a binary classifier that is

trained to associate the object OBJ_i to either a 'CTC' class or a 'non-CTC' class, depending on whether the CTC classifier 28A identifies the object O_i as a CTC.

5 The tdEV classifier 28B is a binary classifier that is trained to associate the object OBJ_i to either a 'tdEV' class or a 'non-tdEV' class, depending on whether the tdEV classifier 28B identifies the object OBJ_i as a tdEV.

10 At step S22, the object OBJ_i is processed by the CTC classifier 28A if the first classifier 24 has determined, at step S14, that object OBJ_i is a CTC. On the other hand, the object OBJ_i is processed by the tdEV classifier 28B if the first classifier 24 has determined, at step S14, that the object OBJ_i is a tdEV.

15 In this embodiment, both the CTC classifier 28A and the tdEV classifier 28B are configured to implement each a respective classification algorithm, namely here each a respective random forest algorithm.

20 Use of random forests for the second classifier 28 may be preferred because they are simpler than some state-of-the-art machine learning models, thereby requiring less computational resources. At the same time, the Applicant has verified that performance in detection of rare cells and/or fragments is primarily limited by labelling errors in the training data rather than in the complexity of the machine learning model that is used.

25 However, the CTC classifier 28A and the tdEV classifier 28B may be configured to implement a machine learning algorithm based on classification that is different from the random forest.

30 For example, the CTC classifier 28A and/or the tdEV classifier 28B may be configured to implement an artificial

neural network, a decision tree or a support vector machine.

In another embodiment, the classifiers 28 are an ensemble of classifiers, possibly implementing several classification approaches and deciding the final output by majority voting.

The control unit 20, step S24, checks the output of the second classifier 28.

If, branch N from step S24, the output of the second classifier 28 indicates that the object OBJ_i is not a rare cell and/or a fragment, then the method returns to step S17. In other words, at branch N from step S24, the second classifier 28 has reverted the output of the first classifier 24.

If, branch Y from step S24, the output of the second classifier 28 indicates that the object OBJ_i contains a rare cell and/or a fragment, then the method proceeds to step S26 and the control unit 20 increases the value of a respective counter.

In detail, in this embodiment, if the CTC-classifier 28A classifies the thumbnail TBN_i as a CTC, then the control unit 20 increases a CTC-counting value CTC-CNT. If the tdEV-classifier 28B classifies the thumbnail TBN_i as a tdEV, then the control unit 20 increases a tdEV-counting value tdEV-CNT.

In practice, at branch Y from step S24, the control unit 20 has determined the presence of a rare cell and/or fragment in the sample of organic fluid of the patient, as a result of the classification of the object OBJ_i .

By increasing the value of the rare-cell counter, the control unit 20 may warn an external operator that a rare cell and/or a fragment has been detected in the organic fluid

sample of the patient.

For example, the current value of the counter may be provided to the external operator.

After step S26, the method returns to step S13 and
5 repeats step S14 on the next object OBJ_{i+1} .

Steps S14 and subsequent steps may be repeated for each thumbnail TBN_i that are extracted from the image IMG.

Step S04 may be repeated for all the images that are captured from the sample of organic fluid of the patient.

10 After step S04 has been repeated for all the images captured from the sample of organic fluid of the patient, the value of the CTC counter CTC-CNT and the tdEV counter tdEV-CNT are indicative of the total number of CTCs and the total number of tdEVs that have been detected in the organic
15 fluid sample of the patient.

With reference to step S05 (figure 1), the control unit
20 may provide the CTC counter CTC-CNT and/or the value of the tdEV counter tdEV-CNT, for example a visual warning through the interface unit 29, so as to warn an operator
20 about the total number of total number of CTCs and the total number of tdEVs that have been detected in the organic fluid sample of the patient.

Additionally, at the end of step S04, the control unit
20 may also provide, through the interface unit 29, also the
25 value of the other counters described with reference to step S17.

An embodiment of a training method 50 of the first classifier 24 is described hereinafter with reference to figure 6.

30 The control unit 20 receives, step S50, a set of training images $T1-IMG[1:K]$.

The training images $T1-IMG[1:K]$ are obtained from organic fluid samples of patients in the same way described with reference to steps S01 to S03 of figure 1.

5 The set of training images $T1-IMG[1:K]$ comprises images that are obtained from organic fluid samples of patients for which a type of cancer has been diagnosed, in particular patients with various types of metastatic cancers including breast, prostate, colorectal, pancreas and lung cancers.

10 Selection of objects for training may be needed to ensure sufficient positive training examples are present in the dataset.

According to an embodiment, the set of training images $T1-IMG[1:K]$ comprises only images obtained from organic fluid samples of patients for which a type of cancer has
15 been diagnosed, in particular patients with various types of metastatic cancers including breast, prostate, colorectal, pancreas and lung cancers.

According to a different embodiment, the training images $T1-IMG[1:K]$ may be obtained from blood or diagnostic
20 leukapheresis samples of patients with various types of cancers including primary or metastatic breast cancer, prostate cancer, colorectal cancer, lung cancer and pancreatic cancer as well as samples from healthy controls, patients with (non-)malignant tumors or other non-cancer
25 diseases.

Then, step S51, the first segmentation module 22 segments each training image $T1-IMG_j$ of the plurality of training images $T1-IMG[1:K]$.

A plurality of training thumbnails $T1-TBN$ are extracted
30 and gathered, as discussed with reference to steps S12 and S13 of figure 3. Each training thumbnail $T1-TBN_i$ is a portion

of the training image T1-IMG that contains a segmented marker-stained object T1-OBJ_i to be labelled.

At step S51, the segmentation module 22 may run a known segmentation algorithm, for example the same segmentation
5 algorithm as the one described with reference to step S12 of figure 3.

Then, step S52, each of the training thumbnails T1-TBN_i is labelled by assigning a specific class thereto indicating if the respective training marker-stained object T1-OBJ_i is
10 a rare cell or a fragment, in particular here a CTC or a tdEV. In other words, each of the training thumbnails T1-TBN_i is labelled by assigning thereto one of the output classes of the first classifier 24.

The labelling of step S52 is, at least in part, manual,
15 i.e. performed by one or more operators or reviewers, e.g. experienced medical staff.

The reviewers look at the training thumbnails T1-TBN and decide, based on experience, to which class each of the training thumbnail T1-TBN_i belongs.

20 For example, the reviewers may assign to each training thumbnail T1-TBN_i any one of the following classes: 'CTC', 'tdEV', 'WBC', 'bare nucleus', and 'other object'.

According to an embodiment, the operators may be instructed to label a training thumbnail T1-TBN_i as a rare
25 cell or fragment even if the operators are not confident that the training thumbnail T1-TBN_i actually contains a rare cell or a fragment. By doing so, it is possible to avoid missing potential thumbnails containing a rare cell or a fragment for training the first classifier 24.

30 The step of labeling S52 may be performed completely by the operators. Labeling can be performed by 3-5 persons at

a rate of approximately 1,000 - 3,000 objects per day, depending on the difficulty of classification. The inventors have verified that classification accuracy improves going from 5,000 to 10,000 to 25,000 to 50,000 labeled objects, with the increase suggesting that further gains are possible if more data is available. An order of magnitude or more larger training dataset is expected to yield further improvements in classifier performance and robustness but may not be attainable with human labeling alone.

10 Alternatively, the step of labeling S52 may be partially assisted by a specific automatic algorithm and partially assisted by a reviewer.

For example, the training thumbnails T1-TBN may be first processed by means of a specific processing module, for example implemented by the control unit 20 and configured to implement one or more previously trained algorithms, for example through the ACCEPT toolbox (<https://github.com/LeonieZ/ACCEPT>), and, subsequently, the reviewers manually score if the extracted candidate really belonged to the automatically assigned class. A detailed implementation of such approach may be found in document REF1 under the section "Methods - Labelling of cells".

20 Then, step S53, the first classifier 24 runs a training algorithm by using the training thumbnails T1-TBN that have been labelled at step S52.

As discussed above, the first classifier 24 is configured to implement an artificial neural network.

According to an embodiment, the first classifier 24 may be trained as discussed in document REF1, under the section "Methods - training parameters and set up".

30 For example, the first classifier 24 may be trained by

minimizing a cross-entropy loss function as discussed in document REF1. Alternatively, the first classifier 24 may be trained by minimizing a focal loss function.

According to an embodiment, the training method 50 may
 5 comprise a further step of processing the training thumbnails T1-TBN, before step S53. In this case, the same processing step may be performed on the thumbnails TBN before step S14 of figure 3.

For example, each thumbnail may be processed as
 10 discussed in the section "Methods - Image processing" of document REF1.

According to an embodiment, the thumbnail processing may comprise performing the product of each training thumbnail $T1-TBN_i$ and the respective segmentation.
 15 Therefore, by way of illustrative example only, if the original thumbnail had a size of 80x80x3, the resulting processed thumbnail has a size of 80x80x6. The resulting processed thumbnail is the one that is used at step S53 for training the first classifier 24. This allows to provide
 20 information to the first classifier 24 on which event should be classified when multiple events are visible in the training thumbnail.

According to an embodiment, the thumbnail processing may comprise scaling, for each fluorescent channel, the
 25 intensity of each pixel of the training thumbnail as described in the following equation:

$$I_{pixel,s} = \frac{I_{pixel} - I_m}{\max(I_{max} - I_m, \delta)} , \text{ wherein } I_m \text{ is a measure of the}$$

intensity background and δ is a design parameter that may be chosen depending on the specific application. I_m may be the
 30 minimum value of all pixels, a percentile determined on

pixels not belonging to any event, or a percentile of all pixels, or another method to quantify the background.

According to an embodiment, I_m may be the median intensity of the unsegmented areas in the thumbnail. The Applicant has verified that said scaling may help in improve the issue of empty thumbnails being misinterpreted as a cell.

According to an embodiment, the training method 50 may comprise a further step of pre-selection of the thumbnails to be labeled and therefore used for training the first classifier 24 at step S53. In fact, of all the training thumbnails that are extracted from the training images, typically only a small percentage of thumbnails contains a rare cell. Accordingly, the pre-selection step may comprise uniformly spaced sampling of 100 events per cartridge from a t-stochastic nearest neighbor (tSNE) representation of all events, supplemented by up to 200 events whose StarDist segmentation was inside a CellSearch segmentation. If more than 200 events were available, the CellSearch segmentations that were selected as CTC were given priority. This strategy led to a reasonably frequent representation of CTC and tdEV in the test and training initialization set, but possibly underrepresented the 10-20% of CTC that were missed by CellSearch segmentation.

The pre-selection may also comprise a step of active learning, which allows to supplement the training data with events near the decision boundary.

A different embodiment of the training method of the first classifier 24 is described in detail with reference to figures 7 and 8 and indicated by 60.

At step S61, images obtained from samples of patients with metastatic disease are retrieved. The images may be

stored in archives in the control unit 20. All the samples have been processed as described in steps S01-S03 of figure 1.

Then, step S62, from all the samples of the patients, the samples that will be used to prove the performance of the first classifier 24 are excluded. The remaining samples form an available pool of samples.

At step S63, N samples $S_1, \dots, S_N$ are randomly selected from the available pool of samples.

At step S64, a number of training marker-stained objects L1-OBJ to be labeled are selected in the images obtained from samples $S_1, \dots, S_N$. Each object L1-OBJ _{i} is represented by the respective thumbnail that contains it and may be further represented by the position of the object L1-OBJ _{i} in the thumbnail and/or auxiliary information useful for representing or guiding the labeling of object L1-OBJ _{i} in the subsequent step.

The selection step S64 may be performed in accordance with a specific selection procedure that may be chosen depending on the specific application.

For example, the selection procedure may be based on active learning, in particular using uncertainty sampling wherein utility measures any of "least confident", "minimum margin" and "maximum entropy", or a combination thereof.

In alternative or in addition, the selection procedure may comprise implementing a multi-view solution to the cold-start problem. In this solution, events of classes of interest are overweighted in inverse proportion to their frequency, selecting from successive populations or views, wherein each succession is less likely to contain events of interest and has more events than the previous. Selection

may continue until a maximum number of events is reached so that if the first view is numerous, the procedure moves on to the next sample.

A detailed embodiment of a possible selection procedure
5 is described later with reference to figure 8.

The N samples $S_1, \dots, S_N$ are removed from the available pool of samples (step S65).

Then, step S66, the thumbnails containing the selected objects L1-OBJ to be labeled are labeled by operator(s) or
10 expert reviewer(s). In other words, the labeling of step S66 is human labeling, based on experience.

Each thumbnail is labelled by assigning a specific class thereto indicating if the respective object L1-OBJ_i is a rare cell or a fragment, in particular here a CTC or a tdEV. In
15 other words, each thumbnail T1-TBN_i is labelled by assigning thereto one of the output classes of the first classifier 24.

For example, the reviewers may assign to each thumbnail any one of the following classes: 'CTC', 'tdEV', 'WBC', 'bare
20 nucleus', and 'other object'.

According to an embodiment, the operators may be instructed to label each thumbnail as a rare cell or fragment even if the operators are not confident that the training thumbnail actually contains a rare cell or a fragment. By
25 doing so, it is possible to avoid missing potential thumbnails containing a rare cell or a fragment for training the first classifier 24.

Additionally, step S67, a confidence parameter may be established, which is indicative of the human confidence or
30 reliability of the labeling step. For example, the confidence parameter may be indicative of a degree of agreement among

the reviewers.

At step S68, the selected thumbnails with the respective operator-assigned labels form the training data of the first classifier 24.

5 Then, the first classifier 24 is trained by using the training data, by using a per se known training algorithm.

At step S70, it is assessed if the confidence parameter has an acceptable value, e.g. if compared to a confidence or reliability threshold that can be determined by the reviewers
10 depending on the specific application.

If the confidence or reliability is not acceptable, branch N from step S70, then the training of the first classifier 24 is deemed complete and thus terminated (step S71). The first classifier 24 is thus ready to be used within
15 the classification method of figure 3 for assessing new samples from patients.

If the confidence or reliability is still acceptable, branch Y from step S70, then the training of the first classifier 24 is continued by using other samples from the
20 available pool of samples. In other words, the training method 60 goes back to step S63 and is repeated until the confidence or reliability is not deemed acceptable anymore (branch N from step S70).

As the number of iterations increases, first the
25 confidence of reviewers increases. However, after a number of iterations, the confidence of reviewers decreases as the number of iterations increases. The initial increase is due to practice, and the later decrease is because the first classifier 24 has become better and the uncertainty sampling
30 selects more events that are truly difficult. Therefore, the selection of objects starts to yield many events for which

labeling of the human reviewers is inconsistent. Accordingly, adding more training data with many labeling errors may not improve performance anymore.

5 The Applicant has verified that the training method 60 is particularly suited for the correct classification of rare events, such as the correct classification of CTCs and/or tdEVs.

10 An embodiment of the selection procedure of step S64 of figure 7 is described hereinbelow with reference to figure 8.

The selection procedure starts, S81, with gathering the images obtained from the samples $S_1, \dots, S_N$, that have been randomly selected from the available pool of samples in step S63 (figure 7).

15 At step S82, it is checked whether training data for the first classifier 24 already exists.

20 If, branch N from step S82, the training data database for training the first classifier 24 is empty, for example at a first training iteration of the first classifier 24, then (step S83) the first segmentation module 22 segments all the images belonging to a sample S_i of the randomly selected samples $S_1, \dots, S_N$, thereby identifying a number of segmented objects S-OBJ.

25 For each segmented object S-OBJ_k, the control unit 20 gathers, step S84, a thumbnail T_k containing the object S-OBJ_k.

30 Additionally, for each segmented object S-OBJ_k, also the following information is gathered: a set of features F_k extracted from the object S-OBJ_k, a prior labeling pL_k and prior thumbnails pT_k associated to the object S-OBJ_k. The "prior" corresponds to segmentations and human operator

labels within the CellSearch system. A prior thumbnail has the presence of two markers required for the positive identification of CTC. A prior label reflects the presence of at least one CTC in a prior thumbnail, irrespective of the number of objects inside that prior thumbnail. The majority of prior thumbnails contain one object, but the largest prior thumbnail is $> 1000 \times 1000$ pixels containing 100's of objects. So, a prior thumbnail with prior label CTC has a high likelihood of being a CTC, and otherwise is likely near a CTC. A prior thumbnail with prior label "not CTC" has low likelihood of being a CTC, but is probably harder for an automated system to classify than randomly selected events.

Then, step S85, up to M_1 objects are randomly selected from a group G_A of the segmented objects S-OBJ, thereby yielding a number $Q_A \leq M_1$ of objects. Group G_A is formed by those segmented objects that are inside prior thumbnails with the prior thumbnail being priorly labeled as containing at least one CTC.

Then, step S86, up to $M_1 - Q_A$ objects are randomly selected from a group G_B of the segmented objects S-OBJ, thereby yielding a number Q_B of objects. Group G_B is formed by those segmented objects having a prior thumbnail pT being priorly labeled as not containing a CTC and that are not contained in group G_A .

Then, step S87, up to $M_1 - Q_A - Q_B$ objects are randomly selected from a group G_C of the segmented objects S-OBJ, thereby yielding a number Q_C of objects. Group G_C is formed by those segmented objects having a set of features that may indicate the object being a CTC and that are not contained neither in group G_A nor G_B .

Then, step S88, up to M_2 objects are randomly selected

from a group G_D of the segmented objects S-OBJ, thereby yielding a number Q_D of objects. Group G_D is formed by those segmented objects having a set of features that may indicate the object being a tdEV and that are not contained neither
5 in group G_A nor G_B nor G_C .

Then, step S89, up to M_3 objects are selected from a manifold representation of the set of features extracted from the objects belonging to a group G_E of the segmented objects S-OBJ, thereby yielding a number Q_E of objects. Group
10 G_E is formed by the objects S-OBJ that are not contained in any of the preceding groups, i.e. in neither of groups G_A , G_B , G_C , G_D .

In practice, steps S83 to S89 forms an embodiment of a multi-view solution to the cold-start problem. The
15 embodiment may allow the training data not to become dominated by the few samples with many, e.g. 1000's, CTC (which may tend to artificially inflate classifier accuracy). Also a good balancing may be achieved; for example, as the data from the multi-view solution contains
20 15% CTC and 11% tdEV, where 20% each would be balanced, and the input is ~0.01% CTC and ~0.1% tdEV.

Then, the index i that is associated to the sample S_i that has been analyzed in steps S83-S89 is increased ($i=i+1$) and the control unit 20 checks, step S90, whether the next
25 sample S_{i+1} still belongs to the randomly selected N samples, e.g., if $i+1>N$.

In the negative case, branch N from step S90, the selection procedure returns to step S82 and checks if training data already exists.

30 In case training data for the first classifier 24 already exists, branch Y from step S82, the following steps

are performed.

First, step 91, as discussed for step S83, all images obtained from sample S_i are segmented, thereby providing a number of segmented objects S-OBJ.

5 Then, step S92, the first classifier 24 (as currently trained based on the existing training data collected so far) is run on the segmented objects S-OBJ. The first classifier that is used in step S92 is a preliminary version of the first classifier 24 that is described with reference
10 to figure 3, since training thereof has not been completed yet. However, for the sake of simplicity, the first classifier used in step S92 is still indicated by 24.

For each segmented object $S\text{-OBJ}_k$, the following data is extracted: a respective thumbnail T_k , a respective set of
15 features F_k and respective class probabilities p_k as output by the first classifier 24.

Then, step S93, up to M_1 objects are selected from a group G_A of the segmented objects S-OBJ, thereby yielding a number $Q_A \leq M_1$ of objects. Group G_A is formed by those
20 segmented objects that, in step S92, have been classified by the first classifier 24 in a respective class of interest (i.e., CTC or tdEV) with a high entropy on the class probabilities p_k . For example, selection could be performed from events with a Shannon entropy > 1.25 .

25 Then, step S94, up to M_2 objects are selected from a group G_B of the segmented objects S-OBJ, thereby yielding a number Q_B of objects. Group G_B is formed by those segmented objects that, in step S92, have been classified by the first classifier 24 in a respective class of interest (i.e., CTC
30 or tdEV) with a small prediction margin and are not contained in group G_A . Events with a small prediction margin have a

low difference between the highest probability $p_{\max}$ and the second highest probability p_{next} output by the first classifier 24. For example, selection could be performed from events with a difference in probabilities $p_{\max} - p_{\text{next}} <$
 5 0.2.

Then, step S95, up to M_3 objects are selected from a group G_C of the segmented objects S-OBJ, thereby yielding a number Q_C of objects. Group G_C is formed by those segmented objects with low confidence, and are not contained neither
 10 in group G_A nor G_B . Low confidence means for which the highest probability $p_{\max}$ output by the first classifier 24 in step S92 has a small value, for example selection could be performed from events with a maximum class probability $<$
 0.5.

15 Any of steps 93-95 could be applied only to events with a maximum probability class of CTC or tdEV, or to events of any maximum probability class.

Then, step S96, up to M_3 objects are selected from a manifold representation of the set of features extracted
 20 from the objects belonging to a group G_E of the segmented objects S-OBJ, thereby yielding a number Q_E of objects. Group G_E is formed by the objects S-OBJ that are not contained in any of the preceding groups, i.e. in neither of groups G_A , G_B , G_C .

25 The selection of objects described above with reference to steps S93 to S96 may be quasi-random. If there are a lot of objects in G_A , G_B or G_C , a random selection may be performed from the respective group. The selection for group G_E may be essentially random.

30 In practice, steps S91 to S96 forms an embodiment of an active learning procedure.

Then, step S90 is repeated.

After all the samples $S_1, \dots, S_N$ have been analyzed (i.e. either through steps S83 to S89 or through steps S91 to S96), the selection procedure of step S64 proceeds to step S97.

5 At step S97, for each sample $S_1, \dots, S_N$, the respective selected objects QA to QE are prepared for the subsequent human labeling. For example, for all the selected objects a respective thumbnail is generated. Moreover, information about the position of the object in the thumbnail may be
10 provided to the expert reviewer. Additionally, also selected variables from the set of features extracted from the objects may be provided to the expert reviewer, so as to aid labeling.

 The Applicant has verified that the specific selection
15 procedure described with reference to Figure 8 may help optimizing the training of the first classifier 24.

 A training method 100 of the second classifier 28 is described hereinafter with reference to figure 9.

 The training method 100 may be performed for each one
20 of the classifiers that form the second classifier 28, namely here the CTC classifier 28A and the tdEV classifier 28B.

 In the following, for sake of simplicity, the training method 100 is described with reference to the CTC classifier 28A only.

25 A set of training samples is provided, which are obtained from organic fluid samples of patients in the same way described with reference to steps S01 to S03 of figure 1.

 The control unit 20 retrieves, step S101, image archives
30 T2-IMG that are captured from the set of training samples.

 Training data for the second classifier 28A is extracted

from the images T2-IMG of the training samples.

The set of training samples is formed by a metastatic subset of samples M_I , which are obtained from organic fluid samples of patients for which metastatic cancer has been diagnosed (metastatic patients), and a non-malignant subset of samples B_J , which are obtained from organic fluid samples of patients for which cancer has not been diagnosed (benign tumors or non-cancerous diseases). Patients with non-metastatic malignant disease are not included in either group because they have an elevated risk of being misdiagnosed.

In particular, the metastatic subset of samples M_I is obtained from organic fluid samples of patients in which the presence of metastatic cancer is established.

The number of organic fluid samples M_I , B_J of metastatic patients and non-malignant patients may be chosen depending on the specific application. For example, the ratio thereof may be approximately 1:1. However, in many applications, the number of objects from metastatic samples will be far greater than the number of objects from non-malignant samples. It is possible to balance the number of objects from such samples.

Moreover, the metastatic samples M_I and the non-malignant samples B_J may have been selected from a wider set of samples, similarly to what described with reference to steps S61-S64 of figure 7.

Then, step S102, the segmentation module 22 applies a segmentation algorithm on the images T2-IMG that are obtained from the metastatic subset of samples M_I and the non-malignant subset of samples B_J , so as to identify a plurality of objects T2-OBJ that may be used for the subsequent training.

The segmentation may be the same as the one described with reference to step S12 of figure 3.

Then, step S103, thumbnails T2-TBN are gathered for the segmented objects T2-OBJ.

5 Then, step S104, the first classifier 24 is run on the objects T2-OBJ, as described with reference to step S14 of figure 3.

10 Objects T2-OBJ may undergo a segmentation refinement, step S106, as described with reference to step S18 of figure 3.

A feature vector $T-FTR_i$ is extracted for each object OBJ_i , step S107, as described with reference to step S20 of figure 3 and therefore not described in detail hereinafter.

15 Of the objects T2-TBN, the labeling unit 30A gathers the objects DL-OBJ and corresponding feature vectors DL-FTR_i that the first classifier 24 has classified in the 'CTC' class.

20 The selected objects DL-OBJ comprises both metastatic objects M-OBJ, i.e. the portion of gathered objects DL-OBJ that belongs to the metastatic subset of samples M_I , and benign (non-malignant) objects B-OBJ, i.e. the portion of gathered objects DL-OBJ that belongs to the benign (non-malignant) subset of samples B_J .

25 Then, step S108, the labeling unit 30A performs a dimensionality reduction of the gathered objects DL-OBJ, based on the feature vectors T-FTR extracted for all the gathered objects DL-OBJ.

The dimensionality reduction groups the objects DL-OBJ with similar features together.

30 The dimensionality reduction algorithm may be, for example, a principal component analysis (PCA), a t-

distributed stochastic nearest neighbors (t-SNE), or a uniform manifold approximation and projection (UMAP).

The dimensionality reduction algorithm provides at output a reduced feature vector for each object DL-OBJ_i starting from the corresponding original feature vector extracted at step S107. For example, at step S108, the reduced feature vectors may form a 2D feature space.

Step S108 may be useful for visualizing the distribution of the selected thumbnails in a 2D space, or for avoiding the curse of dimensionality.

At step S109, the labeling unit 30A adjusts the balance between the metastatic objects M-OBJ and the non-malignant objects B-OBJ.

In detail, the labeling unit 30A may sample from all gathered objects DL-OBJ in the overrepresented sample type (metastatic or non-malignant) so as to reduce the number of objects. The reduction may be performed randomly or with some sampling bias to ensure all samples are represented and avoid the overrepresentation of a few samples.

The reduction may help to correct the difference between the number of metastatic objects M-OBJ and the number of non-malignant objects B-OBJ. In fact, the number of non-malignant objects B-OBJ that is gathered at step S104 may be much lower than the number of metastatic objects M-OBJ that is gathered at step S104. This may result in a class imbalance at input of the subsequent analysis. The reduction of step S109 may help adjusting precision and recall of the CTC classifier 28A. If the method in S110 allows this, an alternative for this sampling could be a weighted loss function, bias initialization, or some other method for addressing class imbalance.

In practice, after step S109, the gathered objects DL-OBJ comprise a number M' of metastatic objects M' -OBJ and a number B' of non-malignant objects B' -OBJ. Each gathered object has the respective reduced feature vector T-FTR output
5 from step S108.

The feature vector T-FTR defines a feature space having a dimension N , wherein N is the number of features forming the feature vector T-FTR, e.g. a 30-dimensional feature vector is reduced to a 2 dimensions after the dimensionality
10 reduction of step S108.

Then, step S110, a classification algorithm based on proximity, for example a k nearest neighbors (kNN) classification, a spectral clustering, or a mixture of gaussians, is run on the gathered objects M' -OBJ, B' -OBJ in
15 order to identify one or more regions of the feature space that are dominated by metastatic objects, i.e. objects coming from the metastatic subset of samples M_T .

For example, the regions of the feature space that are dominated by metastatic objects are regions of the feature
20 space wherein the density of metastatic objects is greater than the density of non-malignant objects. In practice, density may be expressed as the number in a small region around each point in the feature space.

Then, step S111, the labeling unit 30A labels the
25 objects contained within the metastatic dominated regions as belonging to the 'CTC' class, and all the objects that are outside the metastatic dominated regions as belonging to the 'NOT-CTC' class.

Figure 10 shows an exemplificative distribution of the
30 metastatic objects M' -OBJ (upper plot) and of the non-malignant objects B' -OBJ (lower plot) in the feature space.

The feature space is here a two-dimensional space formed by a first feature tsne1 and a second feature tsne2 as obtained at step S108.

In the example of figure 10, the proximity
5 classification algorithm may identify four metastatic-dominated regions indicated by REG₁, REG₂, REG₃, REG₄.

Figure 11 shows an example of the feature space of figure 10, wherein the metastatic-dominated regions MD-R are identified by the lighter gray areas of the feature space.

10 Then, the second classifier 28A is trained by using, as labelled training examples, the objects M'-OBJ and B'-OBJ that have been labelled at step S111.

For example, when the second classifier 28 is a random forest or decision tree, the second classifier 28 may be
15 trained by using the feature vectors T-FTR of the objects M'-OBJ and B'-OBJ, together with the respective labels as assigned thereto at step S111.

In detail, in this embodiment, training the second classifier 28A may comprise training multiple variations of
20 the second classifier 28A.

In fact, due to random sampling and the non-deterministic nature of the training process, using the same dataset will produce slightly different models having different performance.

25 Then, step S113 the multiple variations of the second classifier 28A may be evaluated based on prognostic performance of the multiple variations of the second classifier 28A when used for prediction of other patients of other studies. The prognostic performance may be evaluated
30 in a known way, for example by performing a Cox regression on the log of the number of CTC per patient for samples taken

at comparable times in treatment for patients with a comparable disease. Comparable times are, for example, just before initiation of a new therapy, or 4-6 weeks after initiation of a new therapy. Comparable disease is, for example, a group of patients with castration resistant prostate cancer, or a group of patients with stage IV breast cancer.

Based on said prognostic performance, the variation of the second classifier 28A having the best performance is selected (step S114) to be used for prediction of new samples (figure 3). This terminates training of the second classifier 28A.

The training method described with reference to figure 9 has been described with reference to the training of the CTC classifier 28A. However, the person skilled in the art would understand that the same steps may be applied, *mutatis mutandis*, for training the tdEV classifier 28B. For example, for training the tdEV classifier 28B, at step S105, the objects that have been classified in the 'tdEV' class are gathered and the labeling of step S111 is modified in 'tdEV' and 'NOT-tdEV'.

Figure 12 shows, by way of examples only, a series of thumbnails extracted from an image, for example the image of figure 4, for which both the outputs of the first classifier 24 and the output of the CTC classifier 28A are reported.

All the thumbnails of figure 12 have been classified as CTC by the first classifier 24, but only some of them have been classified as CTC also by the CTC classifier 28A.

In detail, the thumbnails of rows R1, R3 and R4 have been classified as CTC by the first classifier 24 and as not-CTC by the CTC classifier 28A. The thumbnails of rows

R2, R5 and R6 have been classified as CTC both by the first classifier 24 and the CTC classifier 28A.

Figure 13 shows, by way of examples only, a series of thumbnails extracted from an image, for example the image of figure 4, for which both the outputs of the first classifier 24 and the output of the tdEV classifier 28B are reported.

All the thumbnails of figure 10 have been classified as tdEV by the first classifier 24, but only some of them have been classified as tdEV also by the tdEV classifier 28B.

In detail, thumbnails of rows R1 and R2 have been classified as tdEV by the first classifier 24 and as not-tdEV by the tdEV classifier 28B. The thumbnails of rows R3 and R4 have been classified as tdEV both by the first classifier 24 and the tdEV classifier 28B.

The fact that the first classifier 24 and the second classifier 28 are trained with two different training methods or approaches allows to optimize the final output of the classification.

In detail, the Applicant has verified that the method for detecting the presence of rare cells and/or fragments according to the invention allows to obtain an improved estimation of the prognosis of overall survival (Hazard Ratio) of a patient, if compared with the known methods for detecting rare cells. For example, CTC identified with the proposed method in two different datasets (metastatic and non-metastatic breast cancer) improved the Hazard Ratio (HR) for Overall Survival to 2.6 and 2.1 respectively, compared with 1.2 and 0.8 obtained by standard CTC selection using the ACCEPT software. On the same two datasets, tdEV identification by the proposed method increased HR to 1.6 and 2.9 respectively, compared to 1.5 and 1.0 provided by

ACCEPT.

Therefore, the double-step classification described above, wherein the second classifier 28 is used to validate the output of the first classifier 24, improves the accuracy
5 of detection of rare cells and/or fragments in organic fluid samples of a patient.

In particular, the use of an automated labelling procedure in training the second classifier 28 has proven to maximize the contrast between rare cells (and/or fragments)
10 that are detected by the first classifier 24 in metastatic and non-malignant patients, thereby improving the present detection method with respect to the known methods.

The use of a proximity classification algorithm for labeling is less demanding in terms of computational
15 resources than other known supervised learning approaches, in particular with respect to the human effort required for labeling when the labeling is operator-assisted. Therefore, the training method 60 of the second classifier 28 may be applied on a large dataset without an excessive effort.

Moreover, the reduction described with reference to
20 step S109 (figure 9) allows to adjust precision and recall of the second classifier 28, depending on the specific application.

Finally, it is clear that modifications and variations
25 may be made to what has been described and illustrated herein, without thereby departing from the scope of the present invention, as defined in the annexed claims. For example, the different embodiments described can be combined with each other to provide further solutions.

For example, the first classifier 24 may be configured
30 to implement a machine learning algorithm based on

classification that is not a deep learning algorithm.

For example, the second classifier 28 may be used to validate the output of the first classifier 24, even when the first classifier 24 classifies the thumbnail TBN_i as non-CTC or non-tdEV. In other words, with reference to figure 3, the branch N from step S16 may also be followed by steps S18 to S22, before discarding the thumbnail TBN_i .

For example, step S108 of dimensionality reduction may be optional. In this case, the proximity classification algorithm may be run on the original feature vectors as extracted at step S107.

In addition or in alternative, the second classifier 28 may be trained (step 112) by using either the reduced feature vectors of the labelled examples that are extracted at step S108 or the original feature vector of the labelled examples that are as extracted at step S107. Using the original feature vectors before dimensionality reduction may improve training of the second classifier 28.

For example, the second classifier 28 may be a single classifier trained to recognize both CTCs and tdEV. In other words, the second classifier 28 may be a non-binary classifier.

For example, the method according to the invention, in particular the steps described with reference to figure 3, may be used to detect the presence of rare cells only, e.g. CTC, not fragments, e.g. tdEV. Alternatively, the method may be used to detect presence of fragments only, e.g. tdEV, and not a full rare cell, e.g. CTC.

For example, the location of tdEV, CTC, and other cell types relative to each other may be used to automatically identify and/or enumerate the number of clusters of various

types, e.g. a cluster of multiple CTC, cluster of CTC with white blood cell. Clusters of multiple CTC are expected to have a lower false positive rate than single CTC, CTC with white blood cells are potentially useful to monitor efficacy
5 of immunotherapies.

For example, the first classifier 24 and the second classifier 28 may be trained to identify, in an image, the presence of rare cells different from the CTCs, for example CEC, CMMC, fetal cells (such as erythroblast or
10 trophoblasts), tumor-associated fibroblasts, stromal cells, etc., or other kinds of cells of interest. In addition or in alternative, the first classifier 24 and the second classifier 28 may be trained to identify, in an image, the presence of cellular fragments of interest other than tdEV,
15 for example edEV, mmEV, and other kinds of fragments. In this case, the training methods 50, 60 of the first classifier 24 and, respectively, the second classifier 28 may be modified accordingly. For example, with reference to the description of the training method of the second
20 classifier, the metastatic subset of samples may indicate a malignant subset of samples obtained from individuals for which the specific disorder associated to the respective cell of interest and/or cellular fragment of interest has been diagnosed, and the metastatic-dominated regions may
25 indicate a malignant-dominated region of the feature space that is dominated by objects coming from the malignant subset of samples.

For example, the training methods 50, 60 and 100 may be performed, in full or in part, by a control unit external to
30 the control unit 20 of figure 2. For example, the labeling unit 30M and/or 30A may be external to the control unit 20.

For example, the labeling unit 30M may be used to assist in training of the first classifier 24 (e.g., gathering, visualizing and selecting the thumbnails to be labeled, running the training algorithm of the first classifier, etc.).

Moreover, the training methods 50, 60 and 100 may further comprise per se known test steps for testing the performance of the first and/or second classifiers 24, 28.

For example, the output of the second classifier 28 may be used as a labelled training example for training the first classifier 24. This allows to substantially increase the size of the dataset used for training the first classifier 24.

According to an embodiment, the method of the invention may further comprise using the classification output of the second classifier 28 (28A and/or 28B) to label training examples (i.e. objects in an image as discussed for the first classifier 24) of a third classifier and train the third classifier using said labeled training examples, so that the third classifier is trained to identify if a marker-stained object in an image is a rare cell or a fragment (as defined for the first classifier 24). In addition, the step of classifying the marker-stained object comprises running the third classifier on the marker-stained object, so that the classification result of the third classifier only may be used to determine whether the marker-stained object in the image is a rare cell or a fragment. The third classifier may be, for example, a variation on MobileNet, Inception, NASNet, another deep convolutional neural network, a Siamese neural network, or some other classifier developed for the classification of image data. According to an embodiment, an

operator may review and modify the labeling of the classification made by the second classifier in providing a label for training examples of the third classifier. The operator decisions may be added to the training data for the
5 third classifier, and the classifier weights are regularly updated to include the new training data.

For example, the control unit 20 may be configured to perform one or more of the steps described with reference to steps S02 and S03 of figure 1. For example, the control unit
10 20 may also comprise the equipment for processing the organic fluid sample of the patient and/or the image acquisition tool that is used for capturing the images of the processed organic fluid sample.

For example, the steps described with reference to
15 figure 3 may be performed in a different order. For example, step S14 may be performed on all the thumbnails extracted from the image and, subsequently, steps S16-S22 may be performed in a batch mode. This may increase the speed of the classification step S04.

20 The segmentation module 22 and the first classifier 24 have been described as independent blocks with reference to figure 2; this allows to optimize both the segmentation and the classification steps, in particular in cases where the events of interest to be classified are rare, as in the
25 detection of rare cells. However, the segmentation module 22 may be a sub-module of the first classifier 24 or may be absent, depending on the specific implementation; in this case, for example, the first classifier 24 may be trained to be executed directly on the image IMG.

30 In alternative, the segmentation module 22 may be external to the control unit 20; in this case, for example,

the control unit 20 may acquire directly a thumbnail TBN as image containing the marker-stained object to be classified by the first classifier 24.

The segmentation module 25, the feature extraction module 26 and the second classifier 28 have been described as independent blocks with reference to figure 2; this allows to optimize both the segmentation and the classification. However, the segmentation module 25 and the feature extraction module 26 may be part of the second classifier 28 or may be absent, depending on the specific implementation.

All the blocks described with reference to the control unit 20 of figure 2 may be hardware modules, software modules or a mix of hardware and software modules, depending on the specific application.

The control unit 20 may be single processing unit or may be a distributed computing system.

The method of figure 3 may be performed completely or partially in cloud; for example, partially on a local computer of medical staff and partially on an external server.

In another embodiment, in addition or in alternative to a counter that keeps track of the classified CTC and tdEV (step S05 of figure 1), the intensity of marker signal in a fluorescent channel not used for classification may be provided. Such marker signals may be of clinical interest, for example used for selection of therapies which target specific markers. The extraction of this data for each event could be performed by feature extraction module 26. The form of data provided could be: 1) the raw data for each event; 2) the number of events that are negative, dimly positive, strongly positive; 3) the proportion of events that are

negative, dimly positive, strongly positive; or 4) some summary statistic that expresses the marker expression for the rare events in the sample overall.

The method of the invention may also comprise a sorting step, after the classification step (e.g., step S04 of figure 1), which allows to isolate the cells of the organic fluid sample that have been identified as cells of interest (e.g., rare cells, in particular CTCs or fetal cells) or as cellular fragments of interest (e.g., circulating cellular fragments, in particular tdEV) for performing downstream processing, including molecular analysis of the cells of interest (e.g. whole genome amplification, low-pass sequencing, next-generation sequencing, Copy Number Variation, FISH, RT-PCR, qPCR, dPCR, STR analysis, Detection of Mutations, Gene Expression Profiling, etc.) or for culture and proliferation of the cells of interest. The sorting step may be performed by the DEPArray™ or by other known image-based cell sorters.

CLAIMS

1. A method for detecting the presence of a cell of interest and/or a cellular fragment of interest in a sample of organic fluid comprising the steps of:

5 - providing (S01) the sample of organic fluid obtained from an individual;

 - processing the sample (S02A, S02B), wherein processing the sample comprises marking the sample with at least one marker specific for the cell of interest and/or the cellular
10 fragment of interest and/or optionally at least one marker specific for cells and/or fragments other than the cell of interest and/or the cellular fragment of interest, the marker/s being identifiable by selective imaging,

 the method further comprising, by a control unit (20):

15 - acquiring (S03, S10-S13) at least one image (IMG, TBN_i) of the marked sample, the image comprising at least one marker-stained object (32, OBJ_i) to be classified;

 - classifying (S04) the marker-stained object; and

 - determining whether the marker-stained object is a cell
20 of interest or a cellular fragment of interest based on a result of the classification of the marker-stained object,

 wherein classifying the marker-stained object comprises:

 - running, by a first classifier (24) of the control unit, starting from the image, a first machine learning
25 algorithm based on classification and trained to identify if a marker-stained object in an image is a cell of interest or a cellular fragment of interest; and

 - validating an output of the first classifier by running, by a second classifier (28, 28A, 28B) of the control
30 unit, starting from the image, a second machine learning algorithm based on classification and trained to identify if

a marker-stained object in an image is a cell of interest or a cellular fragment of interest,

wherein the first classifier has been trained by using a first training model and the second classifier has been trained by using a second training model that is different from the first training model.

2. The method according to claim 1, wherein the first and the second classifiers are trained using labelled training data, the first training model being based, at least in part, on manual labeling of a first set of training data, for example one or more images (T1-IMG[1:K], T1-TBN) containing one or more marker-stained training objects (T1-OBJ, L1-OBJ), by an operator, the second training model being based on automated labeling of a second set of training data, for example one or more images (T2-IMG[1:M], T2-TBN) containing one or more marker-stained training objects (T2-OBJ).

3. The method according to the preceding claim, further comprising training (100) the second classifier (28, 28A, 28B) by using the second training model, wherein the second training model comprises:

- acquiring (S101-S103) the second set of training data (T2-IMG, T2-TBN, T2-OBJ) from a malignant subset (M_1) and a non-malignant subset (B_j) of samples of organic fluid, the malignant subset being obtained from individuals for which a disorder that is associable to the cell of interest and/or cellular fragment of interest has been diagnosed, the non-malignant subset being obtained from individuals for which a disorder that is associable to the cell of interest and/or cellular fragment of interest has not been diagnosed, the second set of training data comprising marker-stained

training objects (M-OBJ, B-OBJ) to be labeled; and

performing an automated labeling of the second set of training data, comprising:

- running (S104), by the first classifier (24), the
5 first machine learning algorithm on the second set of training data, thereby providing a classification of the marker-stained training objects to be labeled;
- selecting (S105, S109), by a labeling unit (30A), a portion (DL-OBJ, M'-OBJ, B'-OBJ) of the second set of
10 training data based on an output of the first classifier;
- obtaining (S107, S108) a feature vector (T-FTR_i) for each selected marker-stained training object (DL-OBJ, M'-OBJ, B'-OBJ), the feature vectors defining a feature space;
- 15 - running (S110), by the labeling unit, a proximity classification algorithm on the feature vectors of the selected marker-stained training objects (DL-OBJ, M'-OBJ, B'-OBJ), the proximity classification algorithm being configured to identify one or more malignant-dominated
20 regions (REG₁-REG₄, MD-R) of the feature space dominated by marker-stained training objects (M-OBJ, M'-OBJ) belonging to the malignant subset,
- labeling (S111), by the labeling unit, the marker-stained training objects of the selected portion (DL-OBJ,
25 M'-OBJ, B'-OBJ) based on an output of the proximity classification algorithm.

4. The method according to the preceding claim, wherein training the second classifier further comprises:

- training (S112) multiple variations of the second
30 machine learning algorithm;
- evaluating (S113) the prognostic performance of each

variation of the second machine learning algorithm; and

- selecting (S114) one of the variations of the second machine learning algorithm based on a result of the prognostic performance evaluation.

5 5. The method according to claim 3 or 4, wherein the marker-stained training objects contained in the one or more malignant-dominated regions (REG₁-REG₄, MD-R) are labelled as being a cell of interest or cellular fragment of interest.

10 6. The method according to any of claims 3-5, wherein training the second classifier further comprises performing, by the control unit, a dimensionality reduction algorithm (S108) on the feature vectors of the selected marker-stained training objects, before running the proximity classification algorithm.

15 7. The method according to any of claims 3-6, wherein selecting the marker-stained training objects comprises selecting only the marker-stained objects that the first classifier has identified as a cell of interest or a cellular fragment of interest.

20 8. The method according to any of claims 3-7, wherein selecting the marker-stained training objects comprises sampling (S109) from the selected marker-stained training objects so as to reduce the number of selected marker-stained training objects and so as that the selected marker-stained
25 training objects have a predetermined balance between marker-stained training objects belonging to the malignant subset and marker-stained training objects belonging to the non-malignant subset, before running the proximity classification algorithm.

30 9. The method according to any of claims 2-8, wherein the first set of training data is obtained from, in

particular only from, organic fluid samples of patients for which a disorder that is associable to the cell of interest and/or cellular fragment of interest has been diagnosed.

10. The method according to the preceding claim, wherein
5 training the first classifier comprises performing (S64) a selection on the first training data, so as to obtain selected marker-stained training objects (L1-OBJ) to be labeled.

11. The method according to claim 9 or 10, wherein
10 training the first classifier comprises performing multiple training iterations (S63-S70) of the first classifier, each iteration being performed on a different subset of the first training data, each training iteration comprising:

- assessing a confidence parameter indicating the
15 labeling confidence of the operators; and
- compare the confidence parameter with a confidence threshold; and
- ending the training of the first classifier (S71) or performing a next training iteration of the first classifier,
20 based on a result of the comparison of the confidence parameter with the confidence threshold.

12. The method according to any of the preceding claims, wherein the first machine learning algorithm is a deep learning algorithm, in particular a convolutional neural
25 network.

13. The method according to any of the preceding claims, wherein validating an output of the first classifier comprises extracting (S20) a feature vector (FTR) from the image representing the marker-stained object (OBJ_i) to be
30 classified and running the second machine learning algorithm on the feature vector.

14. The method according to the preceding claim, wherein the feature vector comprises data indicative of one or more of: pixel intensity, sharpness, shape and size of the marker-stained object; and/or one or more of: the outputs of the first classifier, similarity of the marker-stained object between different channels of the image, and overlap of the marker-stained object between different channels of the image.

15. The method according to any of the preceding claims, wherein validating an output of the first classifier (24) comprises confirming, for example maintaining or reverting, the output of the first classifier based on an output of the second classifier (28).

16. The method according to any of the preceding claims, wherein the marker-stained object is determined to be a cell of interest or a cellular fragment of interest based on the result of the classification of the marker-stained object by the second classifier (28), in particular in response to the second classifier identifying the marker-stained object as a cell of interest or a cellular fragment of interest.

17. The method according to any of the preceding claims, wherein the step of validating an output of the first classifier is performed in response to the first classifier identifying the marker-stained object as a cell of interest and/or a cellular fragment of interest,

the marker-stained object being determined to be a cell of interest or a cellular fragment of interest in response to the second classifier identifying the marker-stained object as a cell of interest or a cellular fragment of interest.

18. The method according to any of the preceding claims,

wherein classifying the marker-stained object further comprises, before running the first machine learning algorithm, segmenting the image (S12), by a first segmentation module (22) of the control unit, to identify,
5 in the image, the marker-stained object to be classified.

19. The method according to any of the preceding claims, wherein the second classifier (28) is a binary classifier that is trained to identify if the marker-stained object is either a cell of interest or not a cell of interest, and/or
10 is either a cellular fragment of interest or not a cellular fragment of interest.

20. The method according to any of the preceding claims, wherein the second machine learning algorithm is a random forest, a decision tree, a support vector machine, or an
15 artificial neural network.

21. The method according to any of the preceding claims, wherein classifying (S04) the marker-stained object further comprises, before running the second machine learning algorithm, segmenting the image (S18), by a second
20 segmentation module (25) of the control unit, to identify, in the image, the marker-stained object to be classified by the second classifier.

22. The method according to any of claims 2-21, wherein the first training data comprises training examples that are
25 labeled, at least in part, based on an output of the second classifier.

23. The method according to any of the preceding claims, further comprising using the output of the second classifier to label training examples of a third classifier; and
30 training the third classifier using said labeled training examples, so that the third classifier is trained to identify

if a marker-stained object in an image is a cell of interest or a cellular fragment of interest, wherein classifying the marker-stained object comprises running the third classifier on the marker-stained object, so that the marker-stained
5 object in the image is determined to be a cell of interest or a cellular fragment of interest based on an output of the third classifier.

24. The method according to claim 22 or 23, wherein labeling comprising performing a review, by an operator, of
10 classifications made by the second classifier.

25. The method according to any of the preceding claims, further comprising:

- counting (S26) a number of cells of interest and/or a number of cellular fragments of interest in the sample by
15 increasing a cell-of-interest counter (CTC-CNT) in response to each marker-stained object being determined to be a cell of interest, and/or by increasing a cellular-fragment-of-interest counter (tdEV-CNT) in response to each marker-stained object being determined to be a cellular fragment of
20 interest; and

- providing (S06) the cell-of-interest counter and/or the cellular-fragment-of-interest counter.

26. The method according to any of the preceding claims, wherein the method is for detecting the presence of a cell
25 of interest and a cellular fragment of interest.

27. The method according to any of the preceding claims, wherein the cell of interest is a rare cell and the cellular fragment of interest is a rare cellular fragment of interest.

28. The method according to the preceding claim, wherein
30 the rare cell is one of the following: CTC, CEC, CMMC, fetal cells, tumor-associated fibroblasts, stromal cells,

respiratory virus cells; and the rare cellular fragment of interest is any of the following: tdEV, edEV, mmEV.

29. The method according to the preceding claim, wherein the rare cell is a CTC and the cellular fragment of interest
5 is a tdEV.

30. The method according to any of the preceding claims, further comprising extracting, from the marker-stained object in the image, features of one or more markers, for example intensity of a marker in the image, that are not
10 used for identification of cells of interest or cellular fragments of interest and whose presence or absence on a cell of interest or cellular fragment of interest is of clinical interest.

31. The method according to any of the preceding claims, further comprising, by an image-based cell sorter, sorting
15 at least one of the cells that has been determined to be a cell of interest or a cellular fragment of interest.

32. The method according to the preceding claim, wherein the step of sorting is for downstream molecular analysis or
20 for culture and proliferation of the sorted cells.

33. A control unit (20) configured to:
- acquire at least one image (IMG, TBN_i) of a marked sample, wherein the marked sample is a sample of organic fluid obtained from an individual and the sample has been marked
25 with at least one marker specific for a cell of interest and/or a cellular fragment of interest and/or optionally at least one marker specific for cells and/or fragments other than the cell of interest and/or the cellular fragment of interest, the marker/s being identifiable by selective
30 imaging, wherein the image comprises at least one marker-stained object (32, OBJ_i) to be classified;

- classify the marker-stained object; and
- determine whether the marker-stained object is a cell of interest or a cellular fragment of interest based on a result of the classification of the marker-stained object,

5 wherein, for classifying the marker-stained object, the control unit comprises:

- a first classifier (24) configured to run, starting from the image, a first machine learning algorithm based on classification and trained to identify if a marker-stained
10 object in an image is a cell of interest or a cellular fragment of interest; and

- a second classifier (28) configured to run, based on an output of the first classifier, starting from the image, a second machine learning algorithm based on classification
15 and trained to identify if a marker-stained object in an image is a cell of interest or a cellular fragment of interest,

 wherein the first classifier has been trained by using a first training model and the second classifier has been
20 trained by using a second training model that is different from the first training model.

34. A computer program comprising instructions that, when executed by a control unit, causes the control unit to:

- acquire at least one image (IMG, TBN_i) of a marked
25 sample, wherein the marked sample is a sample of organic fluid obtained from an individual and the sample has been marked with at least one marker specific for a cell of interest and/or a cellular fragment of interest and/or optionally at least one marker specific for cells and/or
30 fragments other than the cell of interest and/or the cellular fragment of interest, the marker/s being identifiable by

selective imaging, wherein the image comprises at least one marker-stained object (32, OBJ_i) to be classified;

- classify the marker-stained object; and

- determine whether the marker-stained object is a cell
5 of interest or a cellular fragment of interest based on a result of the classification of the marker-stained object,

wherein, for classifying the marker-stained object, the control unit comprises:

- a first classifier (24) configured to run, starting
10 from the image, a first machine learning algorithm based on classification and trained to identify if a marker-stained object in an image is a cell of interest or a cellular fragment of interest; and

- a second classifier (28) configured to run, based on
15 an output of the first classifier, starting from the image, a second machine learning algorithm based on classification and trained to identify if a marker-stained object in an image is a cell of interest or a cellular fragment of interest,

20 wherein the first classifier has been trained by using a first training model and the second classifier has been trained by using a second training model that is different from the first training model.