

SYSTEM AND METHOD FOR NON-DESTRUCTIVE RAPID FOOD PROFILING USING ARTIFICIAL INTELLIGENCE

TECHNICAL FIELD

The present application relates to a system and method using artificial intelligence, and more specifically to a system and method using machine learning for non-destructive rapid food analysis at a molecular level.

BACKGROUND

[0001] Consumers and regulatory authorities expect a consistent and reliable food product from manufacturers/suppliers. However, due to the nature of agricultural and farm products, such consistency is difficult to achieve as the products have intrinsic variations as a result of different cultivars, terroir, altitudes, preparation methods, etc. Even products grown at a single location may exhibit differences in properties due to seasonality, climate changes, and quality of soil. These difficulties may be compounded by the blending of multiple cultivars, which is often performed to try to compensate for such variations. There have been attempts by manufacturers to improve the efficiency of measuring characteristics of food products, for example, by using gas chromatography, electronic tongues which operate based on electro-potentiometric differences or by using mass spectrometry. Unfortunately, these approaches require destroying the sample being tested and long processing lead time. In practice, food manufacturers/suppliers commonly end up relying on the human experts to visually inspect, physically touch, and taste the products.

[0002] SUMMARY

[0003] According to one aspect, there is provided a system for non-destructive profiling of a tea, the system including: a portable apparatus including: a receptacle configured to move a sample in a volumetric sampling space, in which the sample includes the tea in a non-homogenized form; a source configured to direct light towards the volumetric sampling space; an optical device having an input port and an output port, the input port being configured to sense reflectance from at least a part of the sample in the volumetric sampling space, the optical device being configured to output a component of the reflectance through the output port; and

[0004] a detector coupled to the output port, the detector being configured to convert the component of the reflectance into captured data, the captured data being characterized by an overtone spectrum; and a computing device coupled to the detector, the computing device being configured to: apply at least one first machine learning model to the captured data to predict at least one facet corresponding to at least one selected wavelength from the overtone spectrum; and provide a signature data characteristic of the tea using a plurality of the at least one facet.

[0005] According to another aspect, there is a system for non-destructive profiling of a sample of a food, the system including: a receptacle configured to move the sample in a path so as to intersect a volumetric sampling space, the sample being in a non-homogenized form; a sensor configured to sense reflectance from at least a part of the sample in the volumetric sampling space, the sensor being configured to output a component of the reflectance as captured data, the captured data being characterised by an overtone spectrum over a range of wavelengths; and a computing device configured to apply at least one first machine learning model to the captured data to: predictively determine selected wavelengths from the range of wavelengths; predict at least one facet corresponding to the selected wavelengths; and provide a signature data using the at least one facet.

[0006] In any of the above systems, the computing device may be further configured to apply at least one second machine learning model to the signature data to: predict at least one descriptor; and provide a signature characteristic of the tea using the at least one descriptor.

[0007] According to another aspect, there is provided a method of non-destructive profiling of a sample of a food, the method comprising: providing captured data to a computing device, the captured data being characterised by an overtone spectrum of a measure of a reflectance over a range of wavelengths, wherein the reflectance is sensed from at least a part of the sample in a volumetric sampling space, the sample including non-homogenised food; predicting at least one facet by applying at least one first machine learning model to the captured data, wherein the at least one facet corresponds to at least one selected wavelength predictively determined from the range of wavelengths; and using the at least one facet to provide a signature data characteristic of the sample.

[0008] The above method may further include: applying at least one second machine learning model to the signature data, wherein the at least one second machine learning model is

configured to: predict at least one descriptor; and provide the signature of the food using at least one descriptor. The method may further include: predicting a blend intended to have a taste associated with the food, wherein the blend has a composition of ingredients different from a composition of the food. The method may further include: applying a second machine learning model to the signature data, wherein the second machine learning model is at least one selected from a group consisting of: a taste profile prediction module, a variety prediction module, a blend configuration module, an adulteration detection module, and a food grade/quality control module, and a nutritional analysis module.

[0009] These and other aspects of the embodiments will be further described in following description and appended drawings.

BRIEF DESCRIPTION OF DRAWINGS

[0010] Fig. 1 is a schematic diagram of a system according to one embodiment.

[0011] Fig. 2A and Fig. 2B illustrate examples of a volumetric sampling space for data capture.

[0012] Fig. 3 illustrates an overtone spectrum with seemingly no distinguishing peak.

[0013] Fig. 4 shows a spectrum with distinct peaks.

[0014] Fig. 5 is a schematic block diagram of a machine learning system according to one embodiment.

[0015] Fig. 6 is a flow diagram of a machine learning method according to one embodiment.

[0016] Fig. 7 illustrates a machine learning method according to another embodiment.

[0017] Fig. 8 illustrates the method of Fig. 7 as applied to predicting a signature of tea.

[0018] Figs. 9A to 9C show examples of predicted descriptors.

[0019] Fig. 10 show examples of predicted signatures.

[0020] Fig. 11 illustrates examples of overtone spectra for samples not in the solid state.

DETAILED DESCRIPTION

[0021] It will be readily understood that parts of the embodiments, as generally described and illustrated in the figures herein, may be arranged and designed in a wide variety of different configurations in addition to the described example embodiments. Thus, the following more detailed description of the example embodiments, as represented in conjunction with the figures, is not intended to limit the scope of the embodiments, as claimed, but is merely representative of example embodiments.

[0022] Reference throughout this specification to “one embodiment”, “another embodiment” or “an embodiment” (or the like) means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment. Thus, the appearance of the phrases “in one embodiment” or “in an embodiment” or the like in various places throughout this specification are not necessarily all referring to the same embodiment. As used herein, the singular “a” and “an” may be construed as including the plural “one or more” unless clearly indicated otherwise.

[0023] Furthermore, the described features, structures, or characteristics may be combined in any suitable manner in one or more embodiments. In the following description, numerous specific details are provided to give a thorough understanding of embodiments. One skilled in the relevant art will recognize, that the various embodiments be practiced without one or more of the specific details, or with other methods, constituents, materials, etc. In other instances, some or all known structures, materials, or operations may not be shown or described in detail to avoid obfuscation.

[0024] A system and method using machine learning for material analysis at a molecular level will be described with reference to Fig. 1 to illustrate non-destructive rapid taste profiling using tea leaves as an example.

[0025] The system 100 includes an apparatus 102 in which is provided a source 110 of light (electromagnetic radiation) and a sensor 120. The sensor may be one suitable for sensing light at wavelengths commonly referred to as visible-to-near infra-red (VIS-NIR) light. Alternatively, the sensor may be one suitable for sensing light at a wider range of wavelengths. Alternatively, the sensor may be one suitable for use with the selected source. The source is configurable to direct light towards a receptacle 130. The source, the receptacle, and the sensor may be disposed in the apparatus with a housing suitable for reducing extraneous light from being picked up by the sensor. In contrast to a typical set-up for nuclear magnetic resonance,

the housing for a system of Fig. 1 may be sized and configured to facilitate portability or movement from one place to another. The rotatable receptacle, the source, and the sensor may form part of a portable apparatus. Thus, embodiments of the present disclosure are suitable for use “on-site” if so desired. The sensor may be coupled to a computing device 140. The computing device may be provided as part of the apparatus 102 or as a separate device configured to receive data/signals output from the sensor (as shown). The computing device may be coupled to a memory/storage provided as part of the computing device or provided over a network connection.

[0026] The receptacle 130 may be disposed in the apparatus 102 with the receptacle in turn configured to receive or support a sample 200 for testing (also referred to as a test sample). In this example, the receptacle is an essentially flat disk without a raised edge around its circumference or side 132. Such a receptacle is suitable for directly disposing a non-fluid sample on a surface 134 of the receptacle. At the same time, disposing a sample on the receptacle can include providing the sample in a cup 210 and disposing the cup on the receptacle. The term “cup” is used to refer generally to a vessel, which may come in a range of different shapes and sizes, suitable for holding and/or transporting the sample. Using a cup to hold the sample may be useful when the sample is at least partially fluid, or when the sample is in a powder form. In an example where the sample is an amount of dried tea leaves, the sample can be directly disposed on the receptacle without a cup. Alternatively, the dried tea leaves can be disposed in a cup, and the cup can be placed on or otherwise received by the receptacle.

[0027] The system of Fig. 1 is further configured such that the receptacle is rotatable about an axis of rotation 136. As the receptacle is rotated about the axis, the sample will be displaced relative to the source and/or sensor. The relative movement between the sample and the source and/or the sensor provide different surfaces of the sample to light coming from the source. Light would be expected to fall on different surfaces of the sample at different angles. Light reflected off the sample and sensed by the sensor would also be expected to vary with the relative movement between the sample and the source and/or the sensor. The variations can be amplified by positioning the cup in an offset position so that the cup is not substantially coaxial with the axis of rotation of the receptacle. Counter-intuitively, it is not required to precisely position the sample relative to the source and/or sensor in order to achieve a high degree of repeatability. To illustrate, the cup may be positioned at an offset relative to the receptacle, and

the offset may be of a different amount from one test to another. Figs. 2A and 2B illustrate schematically that when the sample is in relative motion to the source and/or sensor, the incident light on the sample traces a path on the sample such that the data capture is from a region, rather than a precisely targeted spot. The data capture in this case includes reflectance contributed by molecules beneath the “surface” of the sample. That is to say, the region from which data is captured includes a depth dimension. The data capturing process according to embodiments of the present disclosure can thus be described as a volumetric data capture, in contrast to a spot data capture or a line data capture. The receptacle is configured to move relative to the source and/or the optical device/sensor to define a volumetric sampling space. The source, receptacle, and the sensor may be adapted so that at least part of the sample is disposed in the volumetric sampling space when there is relative motion between at least two of these elements. In other words, the sample may be moved relative to the volumetric sampling space. The system of Fig. 1 is an example of providing captured data from a volumetric sampling space 300 (Figs. 2A, 2B), described to aid understanding and not intended to be limiting. For the sake of brevity, each set of movement bringing the incident light and/or sample through a designated path without retracing a previously traversed path is referred to as one scan. In one example, one scan will be completed in one complete revolution of the receptacle. It has been found that as few as one scan can be enough to provide the captured data for predicting a signature of the food. In some embodiments, one scan can be completed in as little as 3 seconds. A signature (e.g. taste signature) characteristic of a food (e.g., tea) can be output by the system in about 5 seconds from initiating rotation of the receptacle.

[0028] In one example, as shown in Fig. 1. the sensor 120 includes an optical device 122 and a detector 128. The optical device is configured to receive light reflected (reflectance) from the sample 200 and to output light to the detector. The detector is configured to indirectly receive light reflected by the sample. The optical device is configured to receive scattered, divergent, or diffused light reflected by the sample. The optical device may be described as a device configured to provide further diffuse reflection of the light received. The optical device may be described as a device configured to provide multiple diffuse reflection of the light received. The optical device may be described as a device configured to amplify the light received. The optical device is configured to output a component of the light received to the detector. The detector is configured to convert light output from the optical device into data/electrical energy receivable by the computing device.

[0029] The optical device may be described as having an input port 124 and an output port 126. The input port is orientated towards the receptacle 130 so as to sense reflectance or receive light reflected by what is disposed on the receptacle. The optical device, or more specifically the input port, is configured to sense reflectance or receive light that is reflected from the sample disposed on the receptacle. As illustrated schematically by Fig. 1, the system 100 is configured to provide a source 110, a receptacle 130, and the optical device 120 in a physical arrangement suitable for light to travel a direct path from the source to the sample 200 (or the receptacle where a sample may be disposed) before being reflected and received into the optical device through the input port 124. The optical device 120 is configured to remove spatial information from the light received through the input port 124, and to output a component of the light through the output port 126. The optical device may be described as a device configured to remove spatial information from its input such that its output is without spatial information. The optical device may be described as a device configured to extract non-spatial information from an input such that its output is a component or a measure of the input, or in other words, a linear component or a magnitude of the input. The optical device may include, e.g., an integrating sphere, a goniophotometer, a focusing sphere, a Coblentz sphere, or similar devices to receive scattered, divergent, or diffused light, and to output a linear component of the received light. Conceivably, the optical device may alternatively be in the form of metasurfaces, or as a device integrated with the detector, etc.

[0030] Using dried tea leaves as a sample helps to illustrate the utility of the system and method in analysing non-homogenized (also referred to as non-homogenous) articles in the solid state, even when the articles are randomly disposed in various orientations in three-dimensional space. In the case of articles such as tea leaves that are ready for use in brewing, non-homogeneity presents challenges. Each tea leaf is intrinsically unique, and when dried, the tea leaves can take on different shapes and configurations, but often without visible characteristics for distinguishing the taste of one tea from another. This is unlike the case where non-homogeneity of the appearance of plant parts is used to identify a stage of growth of a plant part. For example, fresh plant parts at different stages of growth may be recognized by identifying visible characteristics (such as the presence or absence of a bud or a leaf). Using conventional visual inspection, this means that the fresh plant parts must be laid out with no relevant parts concealed in order to facilitate identification.

[0031] The spectrum characteristic of data obtainable for a non-homogenized sample can be described as an overtone spectrum. An overtone spectrum is characterized by a lack of distinct or sharp peaks. An overtone spectrum is characterized by gradual changes in intensity (in this case, the reflectance) over a range of wavelengths. Fig. 3 illustrates an example of an overtone spectrum 360 to aid understanding. According to the embodiment of Fig. 1, the system is configured to provide irradiation (by the source) on multiple points (surfaces) of the sample. This would be expected to generate a “noisier” spectrum or a spectrum with a larger degree of overtone, and typically deemed undesirable for repeatability as well as rapid and accurate determination of a molecular signature of a sample.

[0032] The resulting spectrum would be expected to be a seemingly “noisy” and “characterless” spectrum or a spectrum with a larger degree of overtone, as compared to a spectrum obtained by FTIR (Fourier Transform Infrared) spectroscopy, mass spectrometry or NMR (nuclear magnetic resonance). “Noise” may include fluctuations of the data/signals which occurs even when the sample is of the same molecular composition. The overtone spectrum obtainable may include fluctuations in the distribution of the data/signals over wavelength, as well as fluctuations in the absolute light intensity measured, at least partially as a result of the non-homogenized nature of the physical aspects of the sample, and/or the diffuse nature of the light reflected by the sample. The molecular signature of the sample cannot be easily extracted from an overtone spectrum using a simple filter to filter out noise.

[0033] Thus, in conventional FTIR (Fourier transform Infrared-Red) spectroscopy, pre-processing to provide a homogenized sample is generally required. Homogenized samples are sometimes made by powderising, followed by filtering the powder to a specific size and then forming compressed pellets of a standard weight from the powder. A spectrum 460 obtainable using conventional FTIR spectroscopy (such as illustrated in Fig. 4 for example) uses a homogenized sample so as to obtain distinct or sharp peaks 462, which can then be used for accurate identification of the sample. In conventional methods, pre-processing of the sample for the purpose of improving homogeneity or repeatability is usually deemed an essential procedure. Pre-processing or homogenisation may take the form of brewing the tea, or dissolving or extracting substances from the tea leaves. Pre-processing or homogenisation may take the form of pulverising the tea leaves to produce a powder having particles of similar size and shape. Such pre-processing or homogenisation processes are unfortunately destructive.

[0034] The system and method described in accordance with embodiments of the present disclosure are found to be suitable for non-destructive analysis of samples, even non-homogenized samples, such as tea leaves. The term “non-destructive” as used in this document includes analysing the sample “as is”. In this example, this means that the tea leaves need not be pulverized or otherwise made into a powder form for the purpose of undergoing analysis. In this example, “non-destructive” can be understood to mean that the tea leaves need not be used or brewed to form a liquid for the purpose of undergoing analysis. In this example, it can also mean that the tea leaves need not have other substances (which may contaminate the sample) added as a substrate for the purpose of undergoing analysis. In this example, “non-destructive” can also mean that the tea leaves need not undergo pre-processing such as homogenisation, extraction, etc., before being received by the receptacle. The sample can thus be described as “non-homogenized” or “heterogeneous” in the sense that it need not undergo pre-processing to ensure or increase homogeneity throughout the whole of the sample. Understandably, these lists of examples of “non-destructive” and “non-homogenized” are not exhaustive, and they may vary when the sample includes different types of food products. The term “food” as used in this document includes food and beverage, and/or related products, whether directly or indirectly available from nature. Tea leaves and tea beverages are just a few examples of food. Understandably, such a system and method can come in particularly useful in situations where it is desirable to maintain the sample in the same form before and after the analysis. Without limiting applications, such a system and method can be useful when the sample is expensive, or when the sample is to be stored for future sale/consumption after carrying out the analysis. Without the need for pre-processing for the sake of performing the analysis, the overall time required for carrying out the analysis can be significantly reduced.

[0035] The system is thus configured to be “smart” enough to overcome challenges presented by non-homogeneity, and to use an overtone spectrum for accurate analysis. The system and method are also capable of reproducing the same results for the same sample in different rounds of analysis. It will be appreciated that the term “same results” is to be understood in a statistical sense. Repeatability or reproducibility can be useful for authentication, certification, or such validation purposes.

[0036] It is found that reproducibility is also achievable by embodiments of the present disclosure when different amounts of samples are used. The sample, and/or the articles constituting the sample, can thus be disposed randomly, that is, without specific geometric or

spatial requirements relative to the receptacle. The dried tea leaves can overlap one another partially or fully, in whichever orientation in three-dimensional space, to various depths, various thickness and/or various number of layers. This can come in useful when it is preferred not to transfer the sample out of its original container. It is not necessary to arrange the articles in a single layer in which none of the articles overlap or contact each other. This can come in useful when it is preferred not to subject the articles constituting the sample to excessive handling.

[0037] As shown in Fig. 1, the sensor is configured to output the captured data/signals (that is, the captured data based on the sample reflectance over a range of wavelengths detected) to the computing device 140 coupled to the sensor 120. The computing device may be coupled to a memory configured to store at least one machine learning algorithm executable by the computing device. Referring to a method 500 illustrated by the schematic block diagram of Fig. 5, the output of data capture 510 is fed to at least one first machine learning model 520. The output of the at least one first machine learning model may be input to at least one second machine learning model 530. The at least one second machine learning model 530 may be configured as one or more artificial intelligence modules 530', such as, a taste profile prediction module, a variety prediction module, a blend configuration module, an adulteration detection module, and a food grade/quality control module, a nutritional analysis module, etc. Training and/or trained data in a database 540 (in a memory/storage accessible by the computing device) may be applied by the at least one first machine learning model and/or the at least one second machine learning model. Various examples will be described below.

[0038] Referring now to Fig 6 for a method 600 for predicting a signature of a food, captured data 141 (Fig. 1) from a three-dimensional or volumetric sampling space 300 (Figs. 2A, 2B) is provided to at least one first machine learning model 520. The computing device 140 is configured to apply the at least one first machine learning model 520 to predict a signature data 610 based on the captured data 141. The captured data may be captured from one or more samples 200. One example of the captured data is sample reflectance. The captured data may be captured from one or more scans of the sample. The at least one first machine learning model may include normalization 522 of the captured data and dimensionality reduction. The dimensionality reduction technique may include principal component analysis, noise elimination, selection of variables, and/or application of weightages. The computing device may be configured to apply at least one first machine learning model selected from a plurality

of first machine learning models to predict at least one facet of the signature data. The computing device may be configured to apply a plurality of selected first machine learning models to predict respective facets of the signature data. The at least one first machine learning model may be described as an unsupervised learning algorithm. The resulting predicted signature data 610 may be provided as an output of the system, stored in a database, and/or be fed to at least one second machine learning model 530. As another example, the first machine learning model may be configured as a supervised learning model, or a partial supervised learning model.

[0039] The at least one second machine learning model (also referred to here as an artificial intelligence (AI) module) 530 may be configured as a supervised machine learning model. The AI module may be one of a plurality of AI modules, and may be configured to apply logistic regression, naive Bayes, support vector machine, neural network, random forest, etc. The second machine learning model is configured to predict a signature 620 of the food based on the predicted signature data 610.

[0040] The computing device may be configured to apply the first machine learning model and/or the second machine learning model with training/trained data stored in a memory or a storage device, in the computing device or accessible via a network, such as in cloud storage. The training/trained data may be part of a database configured to store one or more supervised machine learning models and training/trained data. The database may store a plurality of signatures of a food, such as the taste signatures of the food.

[0041] In one embodiment, as shown in Fig. 7, the method 600' may be performed using captured data from a known sample. The second machine learning model 530 may be a supervised machine learning algorithm configured to predict 710 a signature, and to compare 720 the prediction against known or training data from the database. Training feedback 730 is made so as to improve the second machine learning model.

[0042] Referring still to Fig. 7, the method 600' may also be performed using captured data from an unknown sample. The second machine learning model 530 may be a supervised machine learning algorithm configured to predict 710 and to output a signature 620.

[0043] Fig. 8 shows an example where an embodiment of the present disclosure is applied to predicting a sensory experience associated with a tea, and more specifically to the taste of the tea. In some examples, the database is configured to include one or more training models, taste

signatures, taste signature data, cultivars of tea, origins of tea varieties, tea classifications, etc. In some examples, the database is configured to store ProfilePrint™ data, a proprietary example of signature data. Examples of different types of signatures that can then be predicted using the signature data include: taste profile, cultivar, origin, classification, etc.

[0044] An example 800 of performing a signature prediction based on a sample of dried tea leaves will be described to aid understanding, without intention to be limiting. For example, signature data 610 of a known tea sample is obtained and together with a corresponding tea taster input, act as input into training one or more second machine learning models 530'. Continuing with the example of predicting the signature of tea leaves, a signature prediction module for tea may be built by collecting input (training data) from a tea taster, for example.

[0045] As an example, the second machine learning model is a signature prediction module 530'. In this example, the signature prediction uses vocabulary of a master tea taster to describe a tea brewed from dried tea leaves taken from the sample. It has been found that the system and method as described can predict a taste or flavour, similar to a human master tea taster. It can be appreciated that the signature prediction is a complex skill when performed by a human. Contrary to its name, "taste signature" is not simply about the flavour of the brewed tea as sensed by the taste buds on the tongue. A master tea taster would take into account the appearance, aroma, flavour, and mouth feel, of the tea. It therefore takes multiple descriptors to characterize a signature characteristic of a specific tea sample – for example, a taste signature may include descriptor "bright" and descriptor "smoky". It is said that there are at least 3,000 varieties of tea in the world. The Indian and Ceylon tea industry uses a vocabulary of about 35 descriptors, such as "bright", "pungent", "smoky", etc., to compose different taste signatures. In another example, a signature can be a combination of different degrees of descriptors, such as "sourness", "bitterness", "astringency", "sweetness", "richness", "Umami", "saltiness", "aftertaste-A", "aftertaste-B", etc. The same plant in a plantation may produce tea with different signature data from harvest to harvest. Therefore, simply matching the signature data would not produce consistently accurate taste predictions, and accurate taste predictions are required for correct prediction of the tea variety.

[0046] The second machine learning model may include models such as logistic regression, naive Bayes, support vector machines, artificial neural networks, and random forests. In this example, a predictive model is implemented to predict the presence or absence of each descriptor. Each descriptor may be predicted by a separate binary logistic regression model.

Alternatively, a plurality of the descriptors may be predicted by a multivariate regression model. In binary logistic regression for example, one model may be implemented to predict the presence or absence of the “bright” descriptor. Another model may be implemented to predict the presence or absence of the “pungent” descriptor. In another example, the predictive model is configured to predict a value for each descriptor, in which the value is one out of a plurality of possible values.

[0047] The second machine learning model 530' may include multiple predictive models trained on training data, in which signatures are independent variables, and the human tea taster-defined taste characteristics are dependent variables. Referring to Figs. 9A to 9C, the nine plots 901, 902, 903, 904, 905, 906, 907, 908, 909 show examples of predicted descriptors. A different predictive model may be applied to predict each descriptor. The second machine learning model in this example is configured to apply multiple predictive models to predict or arrive at a signature (combination of descriptors) for a tea sample. Examples of the predicted signatures of 15 different samples of tea are shown in Fig. 10. The signature prediction may be fed back to the supervised machine learning model to improve future predictions. The database is therefore configured to update the training data with new data, for subsequent use in the predictive model. Fig. 9A, Fig. 9B, and Fig. 9C additionally illustrate correlation between the predicted descriptor and the observed data. As the predictive models are fed with more data over time, the correlation between a predicted descriptor and the observed data is expected to improve.

[0048] Commonly recognized categories of tea include: black, dark, oolong, yellow green and white. As shown below in Table 1, tests to identify variants were performed on four different categories of tea using dried tea leaves using a system and method as described above.

Tea Category	Captured Data	Signature Data
Black	47 +/- 6%	86 +/- 7%
Green	89 +/- 8%	94 +/- 8%
Dark	73 +/- 16%	81 +/- 12%
Oolong	78 +/- 13%	85 +/- 8%

Table 1. Accuracy in identifying tea variants within each category

[0049] In each category of tea, there can be over 300 variants, having differences resulting from different cultivars, seasons, origins, etc. After machine learning has been applied to the

captured data, the resulting signature data demonstrated an improvement in the accuracy of categorisation/classification by as much as 40% in the case of black tea. Therefore, in addition to predicting a sensory experience such as taste, embodiments of the present disclosure are also useful for categorisation or classification.

[0050] Further, the system and method according to embodiments of the present disclosure are found to be sufficiently accurate at the level of resolution required for predicting the different tea varieties within a category of tea. Table 2 below shows the results from conducting 100 blind tests per sample of Ceylon tea. The experiment used black tea (same category) of different varieties (Dimbula, Kandy, Ruhuna, Uva) from different regions of Sri Lanka. The signatures within each tea variety can exhibit a range of variations in the signature. For example, the Dimbula region produces teas with taste ranging from “full bodied” to “light and delicate”. All 400 hundred tests performed correctly predicted the tea variety.

Predictions	Number of Tests Performed on Each Tea Variety			
	Dimbula	Kandy	Ruhuna	Uva
Dimbula	100	0	0	0
Kandy	0	100	0	0
Ruhuna	0	0	100	0
Uva	0	0	0	100

Table 2. Tea Variety Predictions Based on Dried Tea Leaves

Predictions	Grading of Dimbula Tea by Human Experts		
	Grade 1	Grade 2	Grade 3
Grade 1	100	0	0
Grade 2	10	85	5
Grade 3	4	10	86

Table 3. Predictions of Grades of Tea Based on Dried Tea Leaves

[0051] The system and method according to embodiments of the present disclosure are found to be sufficiently accurate at a level of still finer resolution required for predicting the different grades of tea within a tea variety. Table 3 below shows the results from conducting 100 blind tests per sample of the Dimbula variety. Grading is conventionally performed by human experts

or master tea tasters. The results are compared to how master tea tasters graded the same samples. Of the 100 tests performed by the system on a sample graded by human experts as “grade 1”, the system accurately predicted in 100% of the tests that the sample is of “grade 1”. Of the 100 tests performed by the system on a sample graded by human experts as “grade 2”, the system accurately predicted in 85% of the tests that the sample is of “grade 2”. Of the 100 tests performed by the system on a sample graded by human experts as “grade 3”, the system accurately predicted in 86% of the tests that the sample is of “grade 3”.

[0052] Thus, it can be appreciated that the embodiment of the present disclosure can provide richer information, such as cultivars, origins, provenances or tastes. Such levels of details were previously not obtainable using NIR and other micro-spectroscopy-based methods.

[0053] Many food products, e.g., coffee, tea, rice, spices, whiskies, etc., on the market are actually blends or mixtures so as to deliver a consistent sensory experience for the consumer. For instance, a specific champagne product sold under the same brand name is actually a different mixture of different ingredients from year to year, otherwise, the taste of the product will vary from batch to batch in part owing to the intrinsic uniqueness of nature. The conventional method of blending involves having a human expert taster to select the ingredients to use in the blend, and to determine the proportion of each ingredient in the blend.

[0054] Embodiments of the present disclosure provides a non-destructive rapid system and method of predicting a blend composition. In this example, it is desired to replicate or mimic a taste of a known tea (target taste) by creating a blend using tea selected from a plurality of ingredients. The ingredients may differ from one another in terms of harvests, geographical origins, suppliers, batches, sources, varieties, and/or grades, etc.

[0055] A database of signatures is stored in a storage (e.g., local server, cloud storage, etc) accessible by the computing device 140, in which the computing device is configured to predict blend composition according to an embodiment of the present disclosure. Part of the database may have been previously developed by human master tea tasters using industry-standard vocabulary descriptors to describe the taste of tea. Additionally, at least part of the database may be developed using an artificial intelligence-based system including a first machine learning model and a second machine learning model. The database may include a plurality of signatures, in which each signature is a predicted combination of a plurality of descriptors, and in which each descriptor is associated with a predicted binary variable. Each of such signatures

may be referred to as a TasteMapTM of a taste of a particular tea variety or blend. If the target signature is not available in the database of signatures, a signature may be predicted for the target taste and stored in the database, using the system and method described above. If the signature of any ingredient is not found in the database of signatures, a signature for that ingredient may be predicted and stored in the database.

[0056] For example, using a system of Fig. 1, captured data from at least one sample of a first ingredient can be obtained. The computing device can be configured to input the captured data into the first machine learning model. The first machine learning model is configured to normalize the captured data. The first machine learning model is configured to subject the normalized captured data to dimensionality reduction. The first machine learning model is configured to predict at least one facet of a signature data, in which the at least one facet corresponds to at least one predictively selected wavelength from the overtone spectrum. In one example, dimensionality reduction includes principal component analysis to determine discrete wavelength values within the VIS-NIR wavelength range. The determined discrete wavelength values can be described as an overtone spectrum with inadequately distinctive features when compared to other spectrometer readings. The first machine learning model is configured to use the determined discrete wavelength values to predict a signature data (also referred to as a ProfilePrintTM data) for the first ingredient.

[0057] In other words, the detector is coupled to the output port to receive a component of the light. The detector is configured to convert the component of the light into captured data. The captured data is characterized by an overtone spectrum at least partially in a visible-to-near infra-red wavelength range. The computing device is coupled to the detector to receive the captured data. The computing device is configured to apply a plurality of first machine learning models to the captured data to predict a plurality of signature data, and to apply a plurality of second machine learning models in which the plurality of signature data is used to predict a plurality of descriptors. The plurality of descriptors can be used to form the signature of the tea. Each of the plurality of the second machine learning models may be configured to predict a respective one of the plurality of descriptors.

[0058] This process may be iteratively performed to predict a signature for each of the ingredient. The signature predicted can be stored in a database forming part of training data for one or more machine learning models. This process may therefore be iteratively performed over time to improve a signature data prediction for the ingredient. The machine learning

models may be trained on the signatures as independent variables. The machine learning models may be trained on human tea taster-defined descriptors (or characteristics) as dependent variables. The descriptors used may be in the form of words or symbols understood by a user.

[0059] In one example, multiple machine learning models are provided to predict one signature. In one example, the second machine learning model is configured to predict the presence or the absence of a single descriptor. In one example, the second machine learning model is configured to predict the presence (e.g., “bright”) or the absence (e.g., “not bright”) of a descriptor in a signature. In one example, the second machine learning model is configured to apply different machine learning models for each of a plurality of descriptors.

	Composition (%)					% Difference from Target	Cost Savings
	Tea 1	Tea 2	Tea 3	Tea 4	Tea 5		
Blend	\$\$\$	\$\$\$\$\$	\$	\$\$\$\$	\$\$		
A	0.00%	0.0%	61.3%	0.0%	38.7%	10.00%	5.34%
B	0.00%	0.0%	44.9%	0.0%	55.1%	5.00%	4.37%
C	0.00%	0.0%	36.7%	0.0%	63.3%	2.50%	3.90%
Target	20.0%	20.0%	20.0%	20.0%	20.0%	0%	0.00%

Table 4. Predicted Blends

[0060] A signature (e.g., TasteMap™) may be used in predicting one or more blends. Table 4 is a table including at least one possible substitute blend for a Target Blend. The second machine learning model has predicted three possible blends (Blend A, Blend B, and Blend C) that may be perceived by consumers as having a similar taste experience as the Target Blend. The second machine learning model is configured to apply a machine learning algorithm to predict a blend that may produce an overall taste experience similar to that of the target taste experience. The degree of similarity improves over iterative training, and as the database of signatures develops. The second machine learning model is configured to predict one or more ingredients to be selected.

[0061] In this case, the Target Blend is a mixture of five ingredients in equal proportions, that is, 20% Tea 1, 20% Tea 2, 20% Tea 3, 20% Tea 4, and 20% Tea 5. A signature (target signature) is first predicted for the Target Blend. The five ingredients (Tea 1, Tea 2, Tea 3, Tea 4, and Tea

5) may differ from one another only in terms of the grade of tea leaves. The five ingredients may differ from one another in terms of the source or origin.

[0062] It can be understood that having five ingredients is just an example to aid understanding; more than five or fewer than five ingredients may be involved in other examples. For example, the results of the predictions show surprisingly that as few as two ingredients may be used to reproduce the target taste of the Target Blend, although the Target Blend has traditionally been a blend of five ingredients. In this example, the second machine learning model predicted Blend A as consisting of 61.3% of Tea 3 mixed (or blended) with 38.7% of Tea 5.

[0063] The second machine learning model may be configured to predict more than one blend. For example, the same machine learning model may be configured to also predict a second predicted blend B which consists of 44.9% of Tea 3 and 55.1% of Tea 5, as well as a third predicted blend C. The predicted blend C consists of 36.7% of Tea 3 and 63.3% of Tea 5. Thus, the computing device can be configured to apply the second machine learning model to a signature to predict a blend of at least two ingredients, where the blend is intended to have a taste associated with the tea.

[0064] The second machine learning model may be configured to provide a numerical indicator of the percentage difference between each predicted blend and the target blend. According to the numerical indicators in Table 4, of the three predicted blends, Blend B and Blend C are more likely to match the taste of the Target Blend than Blend A. When human tasters (among whom included a tea taster, a master tea taster, and senior management in the tea industry) were brought in to verify the predicted results, Blend B and Blend C were found by all to match the taste of the Target Blend. This suggests that, in the event that one or more of the ingredients Tea 1, Tea 2 or Tea 3 are not available, the food product manufacturer has the option of using the available ingredients to provide the consumer with a consistent taste. The second machine learning model can also be a tool to help the food product manufacturer segment the markets, for example, into a general consumer market segment (Blend B) and another market segment catering to the tea aficionado (Blend C).

[0065] The second machine learning model may be further configured to take into consideration, in addition to the target signature, one or more user-defined parameters in predicting the ingredients to be selected. Examples of user-defined parameters include the amount or the proportion of a chosen ingredient to be used in a blend, the price or cost-related

data of the ingredients. In Table 4, the relative cost of the ingredients is symbolically represented simply for illustration purposes. Based on the cost of the ingredients and the predicted blends, the cost savings of switching blends can be calculated by the computing device. The second machine learning model may be configured to take into consideration such user-defined parameters as part of the predictive learning to enable more complex decision making.

[0066] To illustrate that embodiments of the present disclosure are applicable to other food products besides tea, an example involving edible bird's nest samples will be described. Four samples (EBN A, EBN F, EBN K, and EBN B) were provided. Each sample includes solid articles of various sizes, ranging in size from dust particle size to pieces larger than 10 millimetres (mm) in length or width. Each sample was scanned five times using a system of Fig. 1. The contents were gently disturbed between successive scans. Human visual inspection of the captured data (before it is input to the first machine learning model) failed to identify any distinctive patterns as the captured data seems similar and consistent across all the four samples. The captured data was fed to the first machine learning model to predict signature data (ProfilePrint™ data) for each of the samples. The signature data is subject to a predictive algorithm, and the results are shown below in Table 5. Based on the predicted signature data, pure samples of authentic edible bird's nest (EBN A) can be differentiated from the adulterated samples, even if the adulteration is not visible to the naked eye, or not visible from the captured data. EBN F and EBN K were adulterated with substance X and substance Y respectively. EBN B were associated with different predicted results, corresponding to was adulterated with a heavy dosage of substance X. This is reflected in Table 5, demonstrating that the quality of the other samples and a measure of the quality (e.g., purity) can be correctly predicted based on the predicted signature data.

Sample	EBN A	EBN F	EBN K	EBN B
EBN A	100.00%	79.93%	81.71%	58.82%
EBN F	79.93%	100.00%	92.94%	74.86%
EBN K	81.71%	92.940%	100%	80.27%
EBN B	58.82%	74.86%	80.27%	100%

Table 5. Predictions of Edible Bird's Nest

[0067] The experiment shows that the system is capable of predicting signature data from overtone spectrum of materials other than tea leaves. This also demonstrates that the system can be used for screening or adulteration detection of various substances, and specifically, of food or food-related substances.

[0068] In addition to samples in a solid state, samples in the liquid state have also been found suitable for use with embodiments of the present disclosure. Fig. 11 shows the overtone spectra (or captured data) 1000 for milk samples obtained using an embodiment of the system of Fig. 1. The data captured for a refrigerated sample of milk is spread over a region bounded by a dotted line 1110. The data captured for a milk sample after different lengths of time unrefrigerated are also shown by regions bounded by a dashed line 1120 and a solid line 1130 respectively. It can be appreciated that the overtone spectra do not have distinct peaks. It was found that the system was able to predict signature data from the captured data for prediction and/or meaningful analysis, even though the captured data at any one wavelength may spread over a range of reflectance values. In this case, it was also demonstrated that different volumes of the samples (or the depth of the milk in the cup) need not be a constant for all the scans. This greatly simplifies the process of data capture for the non-expert user, and thus increases the accessibility to such technology in regions where laboratory resources are scarce or prohibitively expensive. From the captured data for the milk samples, signature data is predicted. The first machine learning model may be trained on the signature data thus captured to provide a non-destructive rapid determination of a quality of the milk, for example, the freshness of the milk, etc.

[0069] This disclosure has been presented for purposes of illustration and description but is not intended to be exhaustive or limiting. Many modifications and variations will be apparent to those of ordinary skill in the art. The example embodiments have been chosen and described in order to explain principles and practical application, and to enable others of ordinary skill in the art to understand the disclosure for various embodiments with various modifications as are suited to the particular use contemplated.

[0070] Thus, although illustrative example embodiments have been described herein with reference to the accompanying figures, it is to be understood that this description is not limiting and that various other changes and modifications may be effected therein by one of ordinary skill in the art without departing from the scope of the disclosure.

CLAIMS

1. A system for non-destructive taste profiling of a sample of a food, the system comprising:
 - a receptacle configured to move the sample in a path so as to intersect a volumetric sampling space, the sample being in a non-homogenized form;
 - a sensor configured to sense reflectance from at least a part of the sample in the volumetric sampling space, the sensor being configured to output a component of the reflectance as captured data, the captured data being characterised by an overtone spectrum over a range of wavelengths; and
 - a computing device configured to apply at least one first machine learning model to the captured data to:
 - predictively determine selected wavelengths from the range of wavelengths;
 - predict at least one facet corresponding to the selected wavelengths; and
 - provide a signature data using the at least one facet, wherein the signature data is characteristic of a taste of the food.
2. The system of claim 1, wherein the receptacle is configured to move relative to the source and/or the optical device to define the volumetric sampling space.
3. The system of claim 2, wherein the sensor is configured to remove spatial information from the reflectance.
4. The system of claim 2, further configured to apply at least one second machine learning model to the signature data to:
 - predict at least one descriptor; and
 - provide the signature of the sample using the at least one descriptor.
5. The system of claim 4, wherein the computing device is further configured to apply the at least one second machine learning model to the signature data to predict a blend

intended to have a taste associated with the sample, wherein the blend has a composition of ingredients different from a composition of the sample.

6. The system of claim 4, wherein the computing device is configured to apply the at least one second machine learning model to the signature data to predict a variant and/or a category of the sample.
7. The system of claim 4, wherein the computing device is configured to apply the at least one second machine learning model to the signature data to predict a degree of purity of the sample.
8. A method of non-destructive taste profiling of a sample of a food, the method comprising:
 - providing captured data to a computing device, the captured data being characterised by an overtone spectrum of a measure of a reflectance over a range of wavelengths, wherein the reflectance is sensed from at least a part of the sample in a volumetric sampling space, the sample including non-homogenised food;
 - predicting at least one facet by applying at least one first machine learning model to the captured data, wherein the at least one facet corresponds to at least one selected wavelength predictively determined from the range of wavelengths; and
 - using the at least one facet to provide a signature data characteristic of the sample, wherein the signature data is characteristic of a taste of the food.
9. The method of claim 8, further comprising: applying at least one second machine learning model to the signature data, wherein the at least one second machine learning model is configured to:
 - predict at least one descriptor; and
 - provide the signature of the food using the at least one descriptor.
10. The method of claim 9, further comprising: predicting a blend intended to have a taste associated with the food, wherein the blend has a composition of ingredients different from a composition of the food.

11. The method of claim 9, wherein the second machine learning model is at least one selected from a group consisting of: a taste profile prediction module, a variety prediction module, a blend configuration module, an adulteration detection module, and a food grade/quality control module, and a nutritional analysis module.