(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
30 December 2021 (30.12.2021)

WIPO | PCT

(10) International Publication Number
WO 2021/260349 A1

(51) International Patent Classification:
A24F 40/00 (2020.01) *G16H 20/13* (2018.01)
A61M 11/00 (2006.01) *G16H 40/63* (2018.01)
A61M 15/06 (2006.01) *G16H 50/20* (2018.01)

(21) International Application Number:
PCT/GB2021/051523

(22) International Filing Date:
16 June 2021 (16.06.2021)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
2009489.2 22 June 2020 (22.06.2020) GB

(71) Applicant: **NICOVENTURES TRADING LIMITED** [GB/GB]; Globe House, 1 Water Street, London WC2R 3LA (GB).

(72) Inventors: **MOLONEY, Patrick**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **JAUREGUI, Juan Estebam Paz**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **CHAN, Justin Han Yang**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **BALAN, Catalin Mihai**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **KUENZEL, Johanna**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **TURAKULOV, Laziz**; BAT, (Holdings) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **HODGSON, Matthew**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **NANDRA, Charanjit**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB). **ROUGHLEY, Howard**; British American Tobacco, (Investments) Limited, Globe House, 1 Water Street, London WC2R 3LA (GB).

(74) Agent: **EALEY, Douglas**; D Young & Co LLP, 120 Holborn, London EC1N 2DY (GB).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,

(54) Title: USER FEEDBACK SYSTEM AND METHOD

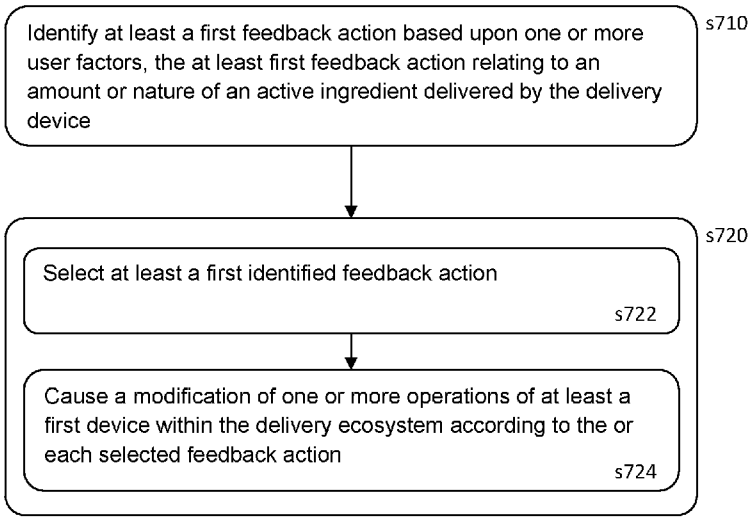

Figure 7

(57) Abstract: A user feedback system for a user of a delivery device within a delivery ecosystem comprises an estimation processor adapted to identify at least a first feedback action based upon one or more user factors, the at least first feedback action relating to an amount or nature of an active ingredient delivered by the delivery device, the feedback action being expected to alter a state of the user as indicated at least in part by the one or more user factors; and a feedback processor adapted to select at least a first identified feedback action, and to cause a modification of one or more operations of at least a first device within the delivery ecosystem, according to the or each selected feedback action.

WO 2021/260349 A1

[Continued on next page]

HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

USER FEEDBACK SYSTEM AND METHOD

Technical Field

The present invention relates to a user feedback system and method for a user of a delivery device.

Background

The "background" description provided herein is for the purpose of generally presenting the context of the disclosure. Work of the presently named inventors, to the extent it is described in this background section, as well as aspects of the description which may not otherwise qualify as prior art at the time of filing, are neither expressly or impliedly admitted as prior art against the present disclosure.

Aerosol provision systems are popular with users as they enable the delivery of active ingredients (such as nicotine) to the user in a convenient manner and on demand.

As an example of an aerosol provision system, electronic cigarettes (e-cigarettes) generally contain a reservoir of a source liquid containing a formulation, typically including nicotine, from which an aerosol is generated, e.g. through heat vaporisation. An aerosol source for an aerosol provision system may thus comprise a heater having a heating element arranged to receive source liquid from the reservoir, for example through wicking / capillary action. Other source materials may be similarly heated to create an aerosol, such as botanical matter, or a gel comprising an active ingredient and/or flavouring. Hence more generally, the e-cigarette may be thought of as comprising or receiving a payload for heat vaporisation.

While a user inhales on the device, electrical power is supplied to the heating element to vaporise the aerosol source (a portion of the payload) in the vicinity of the heating element, to generate an aerosol for inhalation by the user. Such devices are usually provided with one or more air inlet holes located away from a mouthpiece end of the system. When a user sucks on a mouthpiece connected to the mouthpiece end of the system, air is drawn in through the inlet holes and past the aerosol source. There is a flow path connecting between the aerosol source and an opening in the mouthpiece so that air drawn past the aerosol source continues along the flow path to the mouthpiece opening, carrying some of the aerosol from the aerosol source with it. The aerosol-carrying air exits the aerosol provision system through the mouthpiece opening for inhalation by the user.

Usually an electric current is supplied to the heater when a user is drawing/ puffing on the device. Typically, the electric current is supplied to the heater, e.g. resistance heating element, in response to either the activation of an airflow sensor along the flow path as the user inhales/draw/puffs or in response to the activation of a button by the user. The heat generated by the heating element is used to vaporise a formulation. The released vapour mixes with air drawn through the device by the puffing consumer and forms an aerosol. Alternatively or in addition, the heating element is used to heat but typically not burn a botanical such as tobacco, to release active ingredients thereof as a vapour / aerosol.

How the user interacts with the e-cigarette (for example the amount of vaporised / aerosolised payload consumed by the user, and/or their pattern of use), and their actual or perceived utility from the interaction, may be influenced by the user's state, which at least in part may be expressed colloquially as their mood(s) and/or subjective need(s).

Consequently it would be useful to provide a delivery mechanism that was more responsive to the user's state.

SUMMARY OF THE INVENTION

In a first aspect, a user feedback system for a user of a delivery device within a delivery ecosystem is provided in accordance with claim 1.

In another aspect, a user feedback method for a user of a delivery device within a delivery ecosystem is provided in accordance with claim 27.

Further respective aspects and features of the invention are defined in the appended claims.

It is to be understood that both the foregoing general summary of the disclosure and the following detailed description are indicative, but are not restrictive, of the disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete appreciation of the disclosure and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following detailed description when considered in connection with the accompanying drawings, wherein:

- Figure 1 is a schematic diagram of a delivery device in accordance with embodiments of the description.
- Figure 2 is a schematic diagram of a body of a delivery device in accordance with embodiments of the description.
- Figure 3 is a schematic diagram of a cartomiser of a delivery device in accordance with embodiments of the description.
- Figure 4 is a schematic diagram of a body of a delivery device in accordance with embodiments of the description.
- Figure 5 is a schematic diagram of a delivery ecosystem in accordance with embodiments of the description.
- Figure 6 is a schematic diagram of a user feedback system in accordance with embodiments of the description.
- Figure 7 is a flow diagram of a user feedback method for a user of a delivery device within a delivery ecosystem, in accordance with embodiments of the description.

DESCRIPTION OF THE EMBODIMENTS

A user feedback system and method is disclosed. In the following description, a number of specific details are presented in order to provide a thorough understanding of the embodiments of the present disclosure. It will be apparent, however, to a person skilled in the art that these specific details need not be employed to practice embodiments of the present disclosure. Conversely, specific details known to the person skilled in the art are omitted for the purposes of clarity where appropriate.

As described above, the present disclosure relates to a user feedback system. This user feedback system is for improving the responsiveness of a delivery device for a user.

The term 'delivery device' may encompass systems that deliver a least one substance to a user, and include non-combustible aerosol provision systems that release compounds from an aerosol-generating

material without combusting the aerosol-generating material, such as electronic cigarettes, tobacco heating products, and hybrid systems to generate aerosol using a combination of aerosol-generating materials; and aerosol-free delivery systems that deliver the at least one substance to a user orally, nasally, transdermally or in another way without forming an aerosol, including but not limited to, lozenges, gums, patches, articles comprising inhalable powders, and oral products such as oral tobacco which includes snus or moist snuff, wherein the at least one substance may or may not comprise nicotine.

The substance to be delivered may be an aerosol-generating material or a material that is not intended to be aerosolised. As appropriate, either material may comprise one or more active constituents, one or more flavours, one or more aerosol-former materials, and/or one or more other functional materials.

Currently, the most common example of such a delivery device is an aerosol provision system (e.g. a non-combustible aerosol provision system) or electronic vapour provision system (EVPS), such as an e-cigarette. Throughout the following description the term “e-cigarette” is sometimes used but this term may be used interchangeably with delivery device except where stated otherwise or where context indicates otherwise. Similarly the terms ‘vapour’ and ‘aerosol’ are referred to equivalently herein.

Generally, the electronic vapour / aerosol provision system may be an electronic cigarette, also known as a vaping device or electronic nicotine delivery device (END), although it is noted that the presence of nicotine in the aerosol-generating (e.g. aerosolisable) material is not a requirement. In some embodiments, a non-combustible aerosol provision system is a tobacco heating system, also known as a heat-not-burn system. An example of such a system is a tobacco heating system. In some embodiments, the non-combustible aerosol provision system is a hybrid system to generate aerosol using a combination of aerosol-generating materials, one or a plurality of which may be heated. Each of the aerosol-generating materials may be, for example, in the form of a solid, liquid or gel and may or may not contain nicotine. In some embodiments, the hybrid system comprises a liquid or gel aerosol-generating material and a solid aerosol-generating material. The solid aerosol-generating material may comprise, for example, tobacco or a non-tobacco product. Meanwhile in some embodiments, the non-combustible aerosol provision system generates a vapour / aerosol from one or more such aerosol-generating materials.

Typically, the non-combustible aerosol provision system may comprise a non-combustible aerosol provision device and an article (otherwise referred to as a consumable) for use with the non-combustible aerosol provision system. However, it is envisaged that articles which themselves comprise a means for powering an aerosol generating component (e.g. an aerosol generator such as a heater, vibrating mesh or the like) may themselves form the non-combustible aerosol provision system. In one embodiment, the non-combustible aerosol provision device may comprise a power source and a controller. The power source may be an electric power source or an exothermic power source. In one embodiment, the exothermic power source comprises a carbon substrate which may be energised so as to distribute power in the form of heat to an aerosolisable material or heat transfer material in proximity to the exothermic power source. In one embodiment, the power source, such as an exothermic power source, is provided in the article so as to form the non-combustible aerosol provision. In one embodiment, the article for use with the non-combustible aerosol provision device may comprise an aerosolisable material.

In some embodiments, the aerosol generating component is a heater capable of interacting with the aerosolisable material so as to release one or more volatiles from the aerosolisable material to form an

aerosol. In one embodiment, the aerosol generating component is capable of generating an aerosol from the aerosolisable material without heating. For example, the aerosol generating component may be capable of generating an aerosol from the aerosolisable material without applying heat thereto, for example via one or more of vibrational, mechanical, pressurisation or electrostatic means.

In some embodiments, the aerosolisable material may comprise an active material, an aerosol forming material and optionally one or more functional materials. The active material may comprise nicotine (optionally contained in tobacco or a tobacco derivative) or one or more other non-olfactory physiologically active materials. A non-olfactory physiologically active material is a material which is included in the aerosolisable material in order to achieve a physiological response other than olfactory perception. The aerosol forming material may comprise one or more of glycerine, glycerol, propylene glycol, diethylene glycol, triethylene glycol, tetraethylene glycol, 1,3-butylene glycol, erythritol, meso-Erythritol, ethyl vanillate, ethyl laurate, a diethyl suberate, triethyl citrate, triacetin, a diacetin mixture, benzyl benzoate, benzyl phenyl acetate, tributyrin, lauryl acetate, lauric acid, myristic acid, and propylene carbonate. The one or more functional materials may comprise one or more of flavours, carriers, pH regulators, stabilizers, and/or antioxidants.

In some embodiments, the article for use with the non-combustible aerosol provision device may comprise aerosolisable material or an area for receiving aerosolisable material. In one embodiment, the article for use with the non-combustible aerosol provision device may comprise a mouthpiece. The area for receiving aerosolisable material may be a storage area for storing aerosolisable material. For example, the storage area may be a reservoir. In one embodiment, the area for receiving aerosolisable material may be separate from, or combined with, an aerosol generating area.

Alternatively or in addition to aerosol provision systems, a delivery device may include any device that causes/enables the introduction of an active ingredient into the body of the user in a manner that allows the active ingredient to take effect.

Example delivery devices may thus for example include a device that disperses an aerosol into a receptacle, after which a user may take the receptacle from the device and inhale or sip the aerosol. Hence the delivery device does not necessarily have to be directly engaged with by the user at the point of consumption.

In this regard, a delivery device may alternatively or in addition provide a reminder or usage regime for a user, for example reminding a user when to use a snus pouch, or other active deliverable such as a pill. The delivery device may optionally store and dispense such consumables according to the reminder or usage regime.

Similarly, an example delivery device may be a home refill station, which mixes e-liquid components for the user and uses the mix to fill a reservoir of their e-cigarette, thereby determining the type, blend, and/or concentration and/or amount of active ingredients that the user will consume, all else being equal. Such a home refill station may be referred to as a 'dock', as may a power recharging station, or a device that combines both functions.

In this regard, a delivery device operating as a vending machine may similarly provide consumable refills or disposable devices based on mixes and/or selections of e-liquid components, either mixed on demand or equivalently selected from a range of pre-prepared mixes. Similarly, in other implementations, the vending machine may dispense oral products (such as for example snus, snuff,

gums, gels, sprays, and other delivery systems such as patches) or other consumable products containing active ingredients and/or flavourants, for example.

In each case, the delivery device is operable to influence one or more of the amount, timing, type, blend, and/or concentration and/or amount of active ingredient consumed by the user.

Hence more generally a delivery device is operable to influence a property of an active ingredient consumed by a user.

It will be appreciated that several delivery devices may operate in tandem to provide this influence. For example a home refill station, or a vending machine, may operate in conjunction with an e-cigarette to actually deliver a modification of active ingredient, or other feedback, to a user. Similarly a mobile phone may operate in parallel with an e-cigarette to provide information or analysis relevant to the modification or other feedback.

In this sense a delivery device may actually be a delivery system comprising multiple devices operating sequentially and/or in parallel to affect the desired influence / feedback. Hence references to a delivery device or delivery system herein may be considered interchangeable except where stated otherwise.

Referring now to the drawings, wherein like reference numerals designate identical or corresponding parts throughout the several views, Figure 1 is a schematic diagram of a vapour / aerosol provision system such as an e-cigarette 10 (not to scale), providing a non-limiting example of a delivery device in accordance with some embodiments of the disclosure.

The e-cigarette has a generally cylindrical shape, extending along a longitudinal axis indicated by dashed line LA, and comprises two main components, namely a body 20 and a cartomiser 30. The cartomiser includes an internal chamber containing a reservoir of a payload such as for example a liquid comprising nicotine, a vaporiser (such as a heater), and a mouthpiece 35. References to 'nicotine' hereafter will be understood to be merely an example and can be substituted with any suitable active ingredient. References to 'liquid' as a payload hereafter will be understood to be merely an example and can be substituted with any suitable payload such as botanical matter (for example tobacco that is to be heated rather than burned), or a gel comprising an active ingredient and/or flavouring. The reservoir may be a foam matrix or any other structure for retaining the liquid until such time that it is required to be delivered to the vaporiser. In the case of a liquid / flowing payload, the vaporiser is for vaporising the liquid, and the cartomiser 30 may further include a wick or similar facility to transport a small amount of liquid from the reservoir to a vaporising location on or adjacent the vaporiser. In the following, a heater is used as a specific example of a vaporiser. However, it will be appreciated that other forms of vaporiser (for example, those which utilise ultrasonic waves) could also be used and it will also be appreciated that the type of vaporiser used may also depend on the type of payload to be vaporised.

The body 20 includes a re-chargeable cell or battery to provide power to the e-cigarette 10 and a circuit board for generally controlling the e-cigarette. When the heater receives power from the battery, as controlled by the circuit board, the heater vaporises the liquid and this vapour is then inhaled by a user through the mouthpiece 35. In some specific embodiments the body is further provided with a manual activation device 265, e.g. a button, switch, or touch sensor located on the outside of the body.

The body 20 and cartomiser 30 may be detachable from one another by separating in a direction parallel to the longitudinal axis LA, as shown in Figure 1, but are joined together when the device 10 is in use by a connection, indicated schematically in Figure 1 as 25A and 25B, to provide mechanical and electrical

connectivity between the body 20 and the cartomiser 30. The electrical connector 25B on the body 20 that is used to connect to the cartomiser 30 also serves as a socket for connecting a charging device (not shown) when the body 20 is detached from the cartomiser 30. The other end of the charging device may be plugged into a USB socket to re-charge the cell in the body 20 of the e-cigarette 10. In other implementations, a cable may be provided for direct connection between the electrical connector 25B on the body 20 and a USB socket.

The e-cigarette 10 is provided with one or more holes (not shown in Figure 1) for air inlets. These holes connect to an air passage through the e-cigarette 10 to the mouthpiece 35. When a user inhales through the mouthpiece 35, air is drawn into this air passage through the one or more air inlet holes, which are suitably located on the outside of the e-cigarette. When the heater is activated to vaporise the nicotine from the cartridge, the airflow passes through, and combines with, the generated vapour, and this combination of airflow and generated vapour then passes out of the mouthpiece 35 to be inhaled by a user. Except in single-use devices, the cartomiser 30 may be detached from the body 20 and disposed of when the supply of liquid is exhausted (and replaced with another cartomiser if so desired).

It will be appreciated that the e-cigarette 10 shown in Figure 1 is presented by way of example, and various other implementations can be adopted. For example, in some embodiments, the cartomiser 30 is provided as two separable components, namely a cartridge comprising the liquid reservoir and mouthpiece (which can be replaced when the liquid from the reservoir is exhausted), and a vaporiser comprising a heater (which is generally retained). As another example, the charging facility may connect to an additional or alternative power source, such as a car cigarette lighter.

Figure 2 is a schematic (simplified) diagram of the body 20 of the e-cigarette 10 of Figure 1 in accordance with some embodiments of the disclosure. Figure 2 can generally be regarded as a cross-section in a plane through the longitudinal axis LA of the e-cigarette 10. Note that various components and details of the body, e.g. such as wiring and more complex shaping, have been omitted from Figure 2 for reasons of clarity.

The body 20 includes a battery or cell 210 for powering the e-cigarette 10 in response to a user activation of the device. Additionally, the body 20 includes a control unit (not shown in Figure 2), for example a chip such as an application specific integrated circuit (ASIC) or microcontroller, for controlling the e-cigarette 10. The microcontroller or ASIC includes a CPU or micro-processor. The operations of the CPU and other electronic components are generally controlled at least in part by software programs running on the CPU (or other component). Such software programs may be stored in non-volatile memory, such as ROM, which can be integrated into the microcontroller itself, or provided as a separate component. The CPU may access the ROM to load and execute individual software programs as and when required. The microcontroller also contains appropriate communications interfaces (and control software) for communicating as appropriate with other devices in the body 10.

The body 20 further includes a cap 225 to seal and protect the far (distal) end of the e-cigarette 10. Typically there is an air inlet hole provided in or adjacent to the cap 225 to allow air to enter the body 20 when a user inhales on the mouthpiece 35. The control unit or ASIC may be positioned alongside or at one end of the battery 210. In some embodiments, the ASIC is attached to a sensor unit 215 to detect an inhalation on mouthpiece 35 (or alternatively the sensor unit 215 may be provided on the ASIC itself). In either case, the sensor unit 215, with or without the ASIC, may be understood as an example of a sensor platform. An air path is provided from the air inlet through the e-cigarette, past the airflow

sensor 215 and the heater (in the vaporiser or cartomiser 30), to the mouthpiece 35. Thus when a user inhales on the mouthpiece of the e-cigarette, the CPU detects such inhalation based on information from the airflow sensor 215.

At the opposite end of the body 20 from the cap 225 is the connector 25B for joining the body 20 to the cartomiser 30. The connector 25B provides mechanical and electrical connectivity between the body 20 and the cartomiser 30. The connector 25B includes a body connector 240, which is metallic (silver-plated in some embodiments) to serve as one terminal for electrical connection (positive or negative) to the cartomiser 30. The connector 25B further includes an electrical contact 250 to provide a second terminal for electrical connection to the cartomiser 30 of opposite polarity to the first terminal, namely body connector 240. The electrical contact 250 is mounted on a coil spring 255. When the body20 is attached to the cartomiser 30, the connector 25A on the cartomiser 30 pushes against the electrical contact 250 in such a manner as to compress the coil spring in an axial direction, i.e. in a direction parallel to (co-aligned with) the longitudinal axis LA. In view of the resilient nature of the spring 255, this compression biases the spring 255 to expand, which has the effect of pushing the electrical contact 250 firmly against connector 25A of the cartomiser 30, thereby helping to ensure good electrical connectivity between the body 20 and the cartomiser 30. The body connector 240 and the electrical contact 250 are separated by a trestle 260, which is made of a non-conductor (such as plastic) to provide good insulation between the two electrical terminals. The trestle 260 is shaped to assist with the mutual mechanical engagement of connectors 25A and 25B.

As mentioned above, a button 265, which represents a form of manual activation device 265, may be located on the outer housing of the body 20. The button 265 may be implemented using any appropriate mechanism which is operable to be manually activated by the user – for example, as a mechanical button or switch, a capacitive or resistive touch sensor, and so on. It will also be appreciated that the manual activation device 265 may be located on the outer housing of the cartomiser 30, rather than the outer housing of the body 20, in which case, the manual activation device 265 may be attached to the ASIC via the connections 25A, 25B. The button 265 might also be located at the end of the body 20, in place of (or in addition to) cap 225.

Figure 3 is a schematic diagram of the cartomiser 30 of the e-cigarette 10 of Figure 1 in accordance with some embodiments of the disclosure. Figure 3 can generally be regarded as a cross-section in a plane through the longitudinal axis LA of the e-cigarette 10. Note that various components and details of the cartomiser 30, such as wiring and more complex shaping, have been omitted from Figure 3 for reasons of clarity.

The cartomiser 30 includes an air passage 355 extending along the central (longitudinal) axis of the cartomiser 30 from the mouthpiece 35 to the connector 25A for joining the cartomiser 30 to the body 20. A reservoir of liquid 360 is provided around the air passage 335. This reservoir 360 may be implemented, for example, by providing cotton or foam soaked in liquid. The cartomiser 30 also includes a heater 365 for heating liquid from reservoir 360 to generate vapour to flow through air passage 355 and out through mouthpiece 35 in response to a user inhaling on the e-cigarette 10. The heater 365 is powered through lines 366 and 367, which are in turn connected to opposing polarities (positive and negative, or vice versa) of the battery 210 of the main body 20 via connector 25A (the details of the wiring between the power lines 366 and 367 and connector 25A are omitted from Figure 3).

The connector 25A includes an inner electrode 375, which may be silver-plated or made of some other suitable metal or conducting material. When the cartomiser 30 is connected to the body 20, the inner electrode 375 contacts the electrical contact 250 of the body 20 to provide a first electrical path between the cartomiser 30 and the body 20. In particular, as the connectors 25A and 25B are engaged, the inner electrode 375 pushes against the electrical contact 250 so as to compress the coil spring 255, thereby helping to ensure good electrical contact between the inner electrode 375 and the electrical contact 250.

The inner electrode 375 is surrounded by an insulating ring 372, which may be made of plastic, rubber, silicone, or any other suitable material. The insulating ring is surrounded by the cartomiser connector 370, which may be silver-plated or made of some other suitable metal or conducting material. When the cartomiser 30 is connected to the body 20, the cartomiser connector 370 contacts the body connector 240 of the body 20 to provide a second electrical path between the cartomiser 30 and the body 20. In other words, the inner electrode 375 and the cartomiser connector 370 serve as positive and negative terminals (or vice versa) for supplying power from the battery 210 in the body 20 to the heater 365 in the cartomiser 30 via supply lines 366 and 367 as appropriate.

The cartomiser connector 370 is provided with two lugs or tabs 380A, 380B, which extend in opposite directions away from the longitudinal axis of the e-cigarette 10. These tabs are used to provide a bayonet fitting in conjunction with the body connector 240 for connecting the cartomiser 30 to the body 20. This bayonet fitting provides a secure and robust connection between the cartomiser 30 and the body 20, so that the cartomiser and body are held in a fixed position relative to one another, with minimal wobble or flexing, and the likelihood of any accidental disconnection is very small. At the same time, the bayonet fitting provides simple and rapid connection and disconnection by an insertion followed by a rotation for connection, and a rotation (in the reverse direction) followed by withdrawal for disconnection. It will be appreciated that other embodiments may use a different form of connection between the body 20 and the cartomiser 30, such as a snap fit or a screw connection.

Figure 4 is a schematic diagram of certain details of the connector 25B at the end of the body 20 in accordance with some embodiments of the disclosure (but omitting for clarity most of the internal structure of the connector as shown in Figure 2, such as trestle 260). In particular, Figure 4 shows the external housing 201 of the body 20, which generally has the form of a cylindrical tube. This external housing 201 may comprise, for example, an inner tube of metal with an outer covering of paper or similar. The external housing 201 may also comprise the manual activation device 265 (not shown in Figure 4) so that the manual activation device 265 is easily accessible to the user.

The body connector 240 extends from this external housing 201 of the body 20. The body connector 240 as shown in Figure 4 comprises two main portions, a shaft portion 241 in the shape of a hollow cylindrical tube, which is sized to fit just inside the external housing 201 of the body 20, and a lip portion 242 which is directed in a radially outward direction, away from the main longitudinal axis (LA) of the e-cigarette. Surrounding the shaft portion 241 of the body connector 240, where the shaft portion does not overlap with the external housing 201, is a collar or sleeve 290, which is again in a shape of a cylindrical tube. The collar 290 is retained between the lip portion 242 of the body connector 240 and the external housing 201 of the body, which together prevent movement of the collar 290 in an axial direction (i.e. parallel to axis LA). However, collar 290 is free to rotate around the shaft portion 241 (and hence also axis LA).

As mentioned above, the cap 225 is provided with an air inlet hole to allow air to flow when a user inhales on the mouthpiece 35. However, in some embodiments the majority of air that enters the device when a user inhales flows through collar 290 and body connector 240 as indicated by the two arrows in Figure 4.

Referring now to Figure 5, the e-cigarette 10 (or more generally any delivery device as described elsewhere herein) may operate within a wider delivery ecosystem 1. Within the wider delivery ecosystem, a number of devices may communicate with each other, either directly (shown with solid arrows) or indirectly (shown with dashed arrows).

In Figure 5, as an example delivery device an e-cigarette 10 may communicate directly with one or more other classes of device (for example using Bluetooth ® or Wifi Direct ®), including but not limited to a smartphone 100, a dock 200 (e.g. a home refill and/or charging station), a vending machine 300, or a wearable 400. As noted above, these devices may cooperate in any suitable configuration to form a delivery system.

Alternatively or in addition the delivery device, such as for example the e-cigarette 10, may communicate indirectly with one or more of these classes of device via a network such as the internet 500, for example using Wifi ®, near field communication, a wired link or an integral mobile data scheme. Again, as noted above, in this manner these devices may cooperate in any suitable configuration to form a delivery system.

Alternatively or in addition the delivery device, such as for example the e-cigarette 10, may communicate indirectly with a server 1000 via a network such as the internet 500, either itself for example by using Wifi, or via another device in the delivery ecosystem, for example using Bluetooth ® or Wifi Direct ® to communicate with a smartphone 100, a dock 200, a vending machine 300, or a wearable 400 that then communicates with the server to either relay the e-cigarette's communications, or report upon its communications with the e-cigarette 10. The smartphone, dock, or other device within the delivery ecosystem, such as a point of sale system / vending machine, may hence optionally act as a hub for one or more delivery devices that only have short range transmission capabilities. Such a hub may thus extend the battery life of a delivery device that does not need to maintain an ongoing WiFi® or mobile data link. It will also be appreciated that different types of data may be transmitted with different levels of priority; for example data relating to the user feedback system (such as user factor data or feedback action data, as discussed herein) may be transmitted with a higher priority than more general usage statistics, or similarly some user factor data relating to more short-term variables (such as current physiological data) may be transmitted with a higher priority than user factor data relating to longer-term variables (such as current weather, or day of the week). A non-limiting example transmission scheme allowing higher and lower priority transmission is LoRaWAN.

Meanwhile, the other classes of device in the ecosystem such as the smartphone, dock, vending machine (or any other point of sale system) and/or wearable may also communicate indirectly with the server 1000 via a network such as the internet 500, either to fulfil an aspect of their own functionality, or on behalf of the delivery system (for example as a relay or co-processing unit). These devices may also communicate with each other, either directly or indirectly.

In an embodiment of the description, to form a user feedback system as will be described later herein, the server 1000, the delivery device, such as for example the e-cigarette 10, and/or any other device within the delivery ecosystem, may utilise one or more sources of information within the delivery

ecosystem or accessible by one or more devices within it in order to be more accurately responsive to the user's state. These may include a wearable or mobile phone (or any other source, such as the dock or vending machine), or sources such as a storage system1012 of the server. The delivery device may also provide information (such as data relating to interaction with an e-cigarette) to one or more data receivers within the ecosystem, which again may comprise one or more of a wearable, mobile phone, dock, or vending machine, or the server.

To form a user feedback system as will be described later herein, a device within the delivery ecosystem, such as the delivery device 10, may utilise one or more processors to analyse or otherwise process this information, in order to estimate the user's state and/or estimate a form of feedback action determined to alter the estimated state of a user (whether a typical / default user, or a user of a similar demographic to the current user, or specifically the current user), for example by causing modification of one or more operations of the delivery device or another device in the delivery ecosystem.

It will be appreciated that the delivery ecosystem may comprise multiple delivery devices (10), for example because the user owns multiple devices (for example so as to easily switch between different active ingredients or flavourings), or because multiple users share the same delivery ecosystem, at least in part (for example cohabiting users may share a charging dock, but have their own phones or wearables). Optionally such devices may similarly communicate directly or indirectly with each other, and/or with devices within the shared delivery ecosystem and/or the server.

It will be appreciated that references to 'the user's state' encompass one of many states of the user, or equivalently one aspect of the overall state of the user. Hence for example the user's level of stress, which as a non-limiting example may be a combination of social circumstance and cortisol levels, is an example of 'the state of the user', but does not completely define the user. In other words, the state of the user is a state relevant to the potential intervention of one or more feedback actions as described elsewhere herein.

User Feedback System

Referring now to Figure 6, in an embodiment of the description, a user feedback system 2 for a user of a delivery device within a delivery ecosystem 1 comprises an obtaining processor 1010 operable to obtain one or more user factors indicative of user state, an estimation processor 1020 operable to calculate an estimation of user state based upon one or more of the obtained user factors, and a feedback processor 1030 operable to select a feedback action for at least a first device within the delivery ecosystem, responsive to the estimation of user state, in a manner expected to alter the estimated state of a user.

Figure 6 illustrates one possible embodiment of such a user feedback system as a non-limiting example.

In this embodiment, the obtaining processor 1010, estimation processor 1020, and feedback processor 1030 are located within the server 1000. However, it will be appreciated that any one or more of these processors may be located elsewhere within the ecosystem 1, or its role may be shared between two or more processors in server and/or the ecosystem. For example the obtaining processor may be located in an e-cigarette or mobile phone, or the feedback processor may be located in a vending machine or e-cigarette, or the functionality of these processes may be shared between the server and such devices. In other examples, these processors may be local to the delivery device (e.g. an e-cigarette), or to a delivery system comprising the delivery device and a mobile phone.

Obtaining Processor

The obtaining processor 1010 obtains or receives one or more user factors from one or more sources, with the user factors being in one or more classes of data.

Such user factors may have a causal and/or correlating relationship with the user's state, or some other predictable relationship with it. Whilst such a state may be associated with what is colloquially referred to as the user's 'mood', the user's subjective mood *per se* is not a primary consideration of the feedback system; rather, the feedback system relates to the correspondence between obtained user factor(s) and user states, and user states and a form of feedback action that may alter such a state of the user, typically in a predetermined manner that is beneficial to the user.

Further it will be appreciated that where there is a correspondence between user factor(s) and states, and states and feedback, there is also in principle a correspondence between the user factor(s) and the feedback, without the intervening state necessarily needing to be explicitly estimated.

The classes of data obtained by or for the obtaining processor include but are not limited to: indirect or historical data; neurological or physiological data; contextual data; environmental or deterministic data; and use-based data.

Indirect or historical data

Indirect or historical data provides background information about the user that is not necessarily related to their immediate circumstances (e.g. not their immediate environment or context), but which may nevertheless have an influence on the user's state.

Examples of indirect or historical data include but are not limited to the user's purchase history, previously input user preference data, or behavioural patterns in general. Hence more generally, user choices or actions, typically relating to the delivery device but typically not directly derived from use of the delivery device itself.

Optionally, such information (or indeed any persistent information, such as preferred user settings, or model data for user state and/or feedback action as described elsewhere herein, account details, or other stored user factor data), can be transferred between devices where a given user purchases or uses different delivery devices, so that such information does not need to be re-acquired for new or respective devices. Such information can be transferred or shared for example by direct data transfer via Bluetooth® link between old and new devices. However, since a potential reason for buying a new device is because a previous one has been lost, alternatively or in addition the information may be transferred or shared by (also) holding the information remotely in association with an account/user ID to which different delivery devices / systems of the user are then also associated. Hence a system with learnt / obtained indirect or historical data on an old device may be transferred or shared to a new device either directly between devices or via a centralised user account.

As an example of historical information, purchase history may be indicative of a user's state, being indicative of a general state of the user long term (for example in terms of significant or recurrent purchases), and/or a recent state of the user (for example in terms of recent purchases, or purchases that are likely to still influence the user).

Hence purchase history that may be indicative of the user's state includes type(s) of products bought, frequency of purchase, and the like (not necessarily limited to products directly related to the delivery device or its consumables), how they are bought (e.g., online vs shop), and volume of purchases in a

time period. The correspondence between how purchases (and the purchased product or service) affect a user's state can be initially determined on a population basis (e.g.to enable a statistically significant amount of data to be collated), or on a subset of such a population having similar demographics to the user, and/or on the basis of the individual user. Purchases may assist with this process for example, by being marked as associated with certain states, whether using human readable or machine readable markings (such as QR codes); if a consumable or other purchase comprises a machine readable mark, this may be registered as an indicator of mood. Similarly, a consumable may comprise a means for it to be recognised as indicative of mood when inserted or otherwise loaded into the delivery device; for example a microchip with a code, or another uniquely identifiable means of electronically detecting a payload type (such as a binary pattern of conductive dots on the consumable's surface that may be detected by corresponding contacts on the delivery device), may be used. Such identifiable types may vary by composition (e.g. flavours, active ingredients or concentrations of either) or default administration (e.g. two types could be identical except for indicating to the device a different heating profile that results in a different inhalation effect).

The obtaining processor may obtain indirect or historical data from a number of sources, including user profile data held in storage 1012 at the server, for example comprising previously input user preference data, and/or similarly logs of interactions and/or usage patterns; web or Internet based data 110 such as purchasing records received from vendors or other partners; information gathered with consent by a mobile phone 100 of the user, variously relating to input user preference data, on-line purchases, interaction/usage data (for example where the phone operates in tandem with an e-cigarette or other delivery device as a delivery system local to the user), user questionnaires, and the like. Similarly alternatively or in addition the obtaining processor may obtain such data from the delivery device itself.

Neurological and/or physiological data

Neurological and/or physiological data is descriptive of the physical state of the user, in terms of mind and/or body. The data can be descriptive of the user's state on various timescales, including immediate status or changes in state (such as for example heart rate), longer term status or changes in state (such as hormonal cycles), or chronic status, such as fitness levels.

Non-limiting examples of long-term data, for example in the order of multiple months to years, include indicators of the user's metabolism, body shape (e.g. ectomorph, mesomorph, endomorph) or body mass index; chronic disease; any other long term condition such as pregnancy; and activity/fitness level.

Such data may be obtained by or for the obtaining processor from one or more user questionnaires (for example either a questionnaire completed specifically to assist the user feedback system, and/or a questionnaire completed for any third-party partner, for example for a fitness wearable device or social media provider); medical or insurance records by consent; or at least in part from other devices such as a fitness wearable 400 and/or other devices in a wider ecosystem 1 such as smart scales.

Non-limiting examples of medium-to-long-term data, for example in the order of multiple weeks to months, include a user's hormonal levels or hormonal cycles for hormones such as oestrogen, testosterone, dopamine and cortisol; any acute condition or illness; and activity/fitness level.

Non-limiting examples of medium term data, for example in the order multiple days to weeks, include a user's sleep cycle; any acute condition or illness; and a user's hormonal levels or hormonal cycles for hormones such as oestrogen, testosterone, dopamine and cortisol.

Non-limiting examples of medium to short-term data, for example in the order of multiple hours to days, include the user's degree of wakefulness; their degree of activity; appetite or fullness; blood pressure; temperature; and again any acute condition or illness, and/or hormones.

Again such medium term data (whether longer or shorter) may be obtained by or for the obtaining processor from questionnaires, medical or other records, or fitness or other smart devices. Hence for example hormonal levels may be obtained or inferred from questionnaires, medical or other records, diary or calendar entries with consent, and/or fitness or other smart devices, including for example pinprick blood tests. Similarly blood pressure, temperature, degree of activity and the like can be obtained from smart devices (typically wearables) or user input.

Non-limiting examples of short term data, for example in the order of multiple minutes to hours, include the user's sweat response; galvanic skin response (phasic and/or tonic); their degree of activity; appetite or fullness; blood pressure; breathing rate; temperature; muscle tension; heart rate and/or heart rate variability; and again any acute condition or illness, and/or hormones.

In addition, neurological and/or physiological information specific to the delivery device may also be obtained by the obtaining processor, such as the cumulative amount of vapour generated within the short term (for example within a preceding period corresponding to one, two or more times the pharmacological half-life of the active ingredient in the user's body).

Non-limiting examples of immediate data, for example in the order of seconds to minutes, include the user's body position; blink rate; breathing rate; heart rate; heart rate variability; brain wave pattern; galvanic skin response (e.g. phasic); muscle tension; skin temperature; voice (e.g. qualities such as volume, pitch, breathiness); and their degree of activity.

Again short-term and immediate and data may be obtained by or for the obtaining processor typically from biometric sensing, for example using smart devices, or using any suitable approach described herein. For example galvanic skin response could be measured by electrodes on the delivery device; heart rate can be obtained by optical scanning of a blood vessel on the wrist by a wearable device, or by use of an electrocardiogram (ECG) or other dedicated strap-on device. Similarly brainwave patterns can be detected by an electroencephalogram (EEG), and muscle tension can be detected by electromyogram (EMG). Meanwhile body position, blinking and the like can be captured for example by a camera on a phone or in a vending machine.

To the extent that the same examples span different time frames in the above description, it will be appreciated for example that different hormones, hormonal cycles, fitness levels and the like can have shorter and longer term characteristics. It will also be appreciated that where an example of data is included in one list but not another, this does not preclude the data being gathered / used over a different time frame; for example blood pressure may be listed as an example of short term data, but clearly may also be part of longer term data, for example due to ongoing high blood pressure.

As with indirect or historical data, data of a plurality of these types and/or from multiple sources may be used in any suitable combination.

In addition to directly measured neurological or physiological data, any suitable analysis or data fusion may be implemented to obtain data of particular relevance to the delivery device regarding the user's state.

For example, the feedback system may be operable to estimate a current nicotine concentration (as a non-limiting example of an active ingredient), or a concentration of active or inactive compounds that break down from the consumed ingredient, within a user (and subsequently deliver nicotine / the active ingredient accordingly).

Hence in principle the feedback system (for example in a pre-processor or subsystem of the obtaining processor) may estimate the concentration of nicotine in the user based on monitoring the nicotine consumed, the time at which it is consumed, and having stored the value for the half-life of nicotine in the body (around 2 hours, although this value can be refined based on information regarding the individual, such as height, weight etc.). Such monitoring can be performed based on usage data from the delivery device. Hence for example based on the original active ingredient concentration, and a predetermined relationship between heating/aerosol generator power and aerosol mass output, an mass of active ingredient per unit volume inhaled may be estimated; from that, using predetermined absorption relationships (optionally based on analysis of depth/duration of inhalation, using airflow data), the amount of active absorbed may be determined; finally the body mass of the user, and potentially other factors such as a age, gender and the like may be used to determine the concentration of active ingredient and/or breakdown products in the user over time. Again, here nicotine is a non-limiting example of an active ingredient.

It has been found that users typically try to have a nicotine level which is between an upper and lower threshold (which may be different between users), which collectively may be regarded as defining a 'baseline' level. The feedback system can establish such a baseline (e.g., by monitoring use over time), and, as will be described in more detail later herein, the feedback system can select and optionally cause modification of one or more operations of the delivery device to deliver nicotine to match the baseline. The baseline may be a steady value or may vary, e.g. with time of day or day of week. It may be initially estimated based on a profile of the user obtained for example from a questionnaire, and/or built up or refined by information (measured and/or self-reported) from the user.

Such a modification may be expected to alter the estimated state of a user, in a positive manner, as it has been previously determined that the chance that a user will be in a positive mood increases when their nicotine levels are close to their personal baseline or thresholded range.

Where a user consumes several different active ingredients, each may have its own baseline thresholds. Optionally, the feedback system can monitor whether consumption of one active ingredient overlaps another to the extent that one active ingredient may affect the baseline of another, and if so modify these accordingly, for example based on stored pharmokinetic data relating to such overlaps.

As noted previously, in these circumstances it is likely that the user interacts with multiple delivery devices to consume the different active ingredients, and the usage from each device may be combined for the associated user. Alternatively, where a single device can switch between payloads (for example hating different gels), or has a mixed payload of actives, the currently heated payload or payload mix can be communicated to the feedback system for the purposes of tracking consumption.

Contextual data

Contextual data relates to situational factors other than environmental factors (see elsewhere herein) that may affect the user's state. Typically such situational factors affect the user's psychological state or disposition towards stress, calm, happiness, sadness, or certain patterns of behaviour, and hence may

also influence and/or have a correlation with neurological and physiological user factors such as dopamine or cortisone levels, blood pressure, heart rate and the like as described elsewhere herein.

Examples of contextual data include the user's culture, including at a broad scale where they live, their religion if they have one, and at a narrower scale their job and/or employment status, educational attainment and the like, and social economic factors that may interact with these such as gender and relationship status.

Such information may be obtained by or for the obtaining processor from user questionnaires, social media data, and the like.

Other contexts include the season (e.g. winter, spring, summer, autumn) or month, and any particular events or periods within that season or month, such as Lent, Easter, Ramadan, Christmas and the like. For example, users are more likely to see consumption at or below their personal baseline as a positive thing during Lent, or the first few weeks of January.

Such information may be obtained by or for the obtaining processor from a calendar and database of events, suitably filtered if appropriate according to other contexts such as country, religion, employment, gender and the like as described previously.

Other contexts include the user's agenda or calendar, which can indicate sources of stress or relaxation, and how busy or otherwise the user is at a given time. Hence for example a social event may be associated with a positive influence on user state, for example raising dopamine levels, whereas a medical appointment or driving test may be associated with stressors such as an increase in cortisol and heart rate. Similarly events, appointments, and/or reminders in rapid succession may indicate a negative effect on the user's state.

The user's agenda or calendar can also provide an indication of the user's likely location, which may affect either their state, or their ability to use the delivery device in a manner that may modify that state. For example, the user may have different typical states, and different abilities to use their delivery device, depending on whether they are at home, at work, in outdoor or indoor public spaces, in an urban or countryside environment, or commuting. The relationship between user state and location may at least initially be based on data from a corpus of users. Alternatively or in addition this relationship may be built up or refined based on data from the user (e.g. measured or self-reported). It will be appreciated that the user's location may also be determined from a GPS signal obtained by the delivery device or an associated device such as a smartphone, or the registered location of a vending machine or point of sale unit.

With regards to commuting or other modes of travel, the type of travel may influence the user's state. For example, walking may have a more positive effect on the user's state than driving, for example in terms of heart rate, blood pressure and the like. It will be appreciated that this context illustrates the potential for the combination of contexts to be significant, as walking in the sun versus walking in the rain may have different effects on the user's state. The type of travel may for example be inferred from GPS data from the user's phone, or the pairing of the phone or delivery device with a vehicle, or the purchase of public transport tickets, or a questionnaire indicating travel habits/times.

Such information may be obtained by or for the obtaining processor from work or personal digital calendars, for example on the user's phone. It will also be appreciated that the user's phone, or other

smart wearable, may directly provide an indication of the user's location, and/or historical patterns of location, for example corresponding to a user's home and work locations and average commuting times.

Other contexts include the weather in the user's location or the upcoming weather in the user's location or upcoming location. Depending on the user, sunny weather is likely to improve the user's mood and sociability, whilst poor weather is likely to lower the user's mood and potentially reduce their sociability or affect their ability to socialise. For example, some users are likely to behave so as to consume active ingredients to an extent that reflects their expectations of mood as suggested by the weather, optionally in conjunction with other contextual factors and further user factors as described herein.

Such information may be obtained by or for the obtaining processor from a weather app, which may be located on a smart phone 100 of the user, or accessed directly for example by the server 1000. More generally weather data may be obtained in response to GPS data (for example by the smartphone), and/or using a local weather measuring system such as a barometer.

Other contexts include the user's proximity to other people, either generally in terms of crowds or social setting, or specifically in terms of other individuals with which there is in principle a measurable correlation with user behaviour. For example, a user may have a different state depending on whether they are in proximity to their boss, their work colleagues, the friends, their partner, their children, or their parents. Hence for example the user may have a different state in a crowded or sociable environment versus when alone or with a partner or family members.

Such proximity can be inferred from the user's agenda or calendar, their mobile phone, their delivery device, or their location. The user may self-report their social status either specifically for the purpose of the user feedback system herein, or generally for example on social media; meanwhile a phone and/or delivery device for example may detect signals from other phones and/or delivery devices for more than a predetermined period of time, indicating they are remaining in each other's presence. Optionally a phone's camera may be used to detect others, but this may not be available if the phone is in a pocket or bag. The feedback system can also determine the proximity of users of the delivery device with other users of such a delivery device - e.g. any suitable delivery device whose location can be determined by the feedback system (for example directly or via an associated mobile phone), whether or not that other delivery device is part of a feedback system itself. Similarly the feedback system can determine the proximity of specific people whom the user has, with permission, identified to the feedback system; for example by providing their phone number to the feedback system, or the system associating a detected Bluetooth ® or other ID with that user.

A user may also indicate (for example via a questionnaire) their typical state in response to different social situations, groups or individuals, whether at a broad level such as 'introverted' or 'extroverted', or more specifically.

It will be appreciated that other contexts exist that may influence the user's state, such as recently consumed information; social media content, news articles, streamed video, e-books, e-magazines, photos and other similar content that may be obtained by or for the obtaining processor. Some content may be assumed to have a universally consistent effect on the state of users, such as for example news of a natural disaster, whilst other content may affect individuals differently, such as the results for a user's preferred sports team, and be assessed individually, for example based upon results of a user questionnaire.

The content of the consumed information may be assessed, for example for keywords, to generate a rating for positive or negative influence on the user's state. Optionally only the rating may be obtained by or for the obtaining processor, or any suitable digest, such as a keyword selection. More generally the obtaining processor may only receive a digest of user factors as appropriate, particularly where the source material does not itself enumerate some user factor property.

Likewise, usage of devices other than the delivery device may influence the user's state. In particular a choice of apps on the user's phone, and the interaction, type of interaction, and/or duration of interaction with them may have correlations with the user's state; for example social media or playing a gaming app may raise dopamine and/or cortisol levels, heart rate, and the like; whilst listening to a music app may reduce heart rate and/or cortisol levels. The duration of interaction may have a linear or non-linear relationship with these changes of state, or may with time indicate a different state; for example playing a game for a long time may indicate boredom.

It will be appreciated that for many user factors, not merely contextual but of other types as well, a situational response (e.g. an expected state) may at least be initially based upon data from a cohort of users (for example a prior test population of users), but alternatively or in addition may be built up or refined from information obtained from the user (whether measured, received or self-reported).

Environmental and deterministic data

Environmental and deterministic data effectively relate to long-term context data outside of the user's choice or influence. There is some overlap with longer term contextual influences such as culture; hence for example the user's upbringing, their genetics, gender, biome internally (for example they gut biome) and/or externally (for example whether they live in an arid or verdant environment), and age.

As with other data described herein, such environmental and deterministic data may be obtained by or for the obtaining processor from one or more user questionnaires (for example either a questionnaire completed specifically to assist the user feedback system, and/or a questionnaire completed for any third-party partner, for example for a fitness wearable device or social media provider). Amongst other things, such a questionnaire may ask for details such as sex/gender, height, weight, ethnicity, age, etc. Such a questionnaire may also comprise psychometric test questions to estimate a user's mental predisposition and/or history (e.g. one or more of extrovert / introvert, active/passive, optimist/pessimist, calm/anxious, independent/dependent, content/depressed, and the like). Such a questionnaire may also ask questions related to the user's culture and beliefs (e.g. one or more of: a country of own or parent's origin; religion, if any; political persuasion, if any; newspapers or news websites read, if any; other media consumption, if any; and the like). Again as with other data described herein, some such environmental and deterministic data may be obtained by or for the obtaining processor from medical or insurance records by consent; and/or may be inferred from the user's location, as appropriate.

Not all environmental and deterministic data need be long-term; hence for example the time of day, day of the week and month of the year may be considered environmental and deterministic data. Hence for example the user state may vary over the course of a day or week, for example being different during weekdays and weekends, and/or during work hours of the weekday versus evenings, and also potentially at specific times of day. Similarly there may also be overlap for example with other contextual data, such as the weather. Again there may also be synergy between different user factors; for example the time of year may affect the amount of daylight (in terms of both the length and

potentially also weather patterns). The level and/or duration of daylight, either as measured (e.g. using a light sensor / camera on a device within the delivery ecosystem) or as inferred from the date, may also have a detectable relationship with the user's state. The quality of light (e.g. colour temperature, indoor / flickering or outdoor) may also be treated as a user factor.

Use-based data

Use-based data relates to direct interactions of the user with the delivery device and/or optionally any other device within the delivery ecosystem or which can report on interactions with it to the feedback system (e.g. to the obtaining processor). These interactions may relate to vaping/consumption and/or manipulation/handling and/or setting the device.

Vaping/consumption based interactions may relate to inter-inhalation properties such as the number, frequency, and/or distribution/pattern of puffs/acts of consumption within one or more chosen periods. Such periods may include daily, hourly, as a function of location, as a function of pharmokinesis (for example the active ingredient half-life within the body for one or more delivered active ingredients), or any other period that may be relevant to the user's state, and/or chosen to increase the apparent correlation between number, frequency and/or distribution/pattern of puff/consumption and a user's state; for example the period may be equal to the average period of time taken to smoke a conventional cigarette, either for the individual user or as a general population average.

Vaping based interactions may also relate to intra-inhalation properties such as individual vaping actions or statistical descriptions of a cohort thereof (for example but not limited to a cohort within one of the above-described chosen periods), such as duration, volume, average airflow, airflow profile, active ingredient ratio, active ingredient delivery timing, heater temperature, and the like.

Data relating to vapes and vaping behaviour (or more generally consumption) as described above may be obtained by or for the obtaining processor from a delivery device itself, for example via a Wi-Fi ® connection to the server 1000, or via communication with a companion mobile phone 100 or other local computing device, paired to the delivery device 10 for example via a Bluetooth ® connection to form a delivery system. However, in principle at least some data relating to consumption may be obtained from one or more other devices within the delivery ecosystem; for example an associated mobile phone may log vaping events to collate frequency / distribution data. Similarly, a wearable sensor may determine the degree of volume of inhalation based on just movement. Hence one or more sensors relating to the determination of vaping based interactions may be located externally to the delivery device, although typically least one will be internal to the vaping device, the most frequent being an airflow sensor that is typically used to detect the onset of inhalation and activate the aerosolisation mechanism of the device (in turn, typically a heater as discussed previously herein). In any event, such internal or external sensors either singly or in combination represent examples of a sensor platform.

In any event, consequently the user feedback system comprises at least a first sensor platform (internal and/or external to the delivery device 10) comprising at least a first sensor operable to detect at least a first physical property associated with at least a first user inhalation action.

As noted previously, the or each physical property can be one or more of an intra-inhalation property or an inter-inhalation property.

The delivery device may comprise one or more airflow sensors as described previously herein to determine when the user vapes and/or how the user vapes, for example as characterised above, and

raw data relating to vaping/consumption events may be stored in the memory of the delivery device or transmitted to the companion mobile phone or any other suitable device within the delivery ecosystem. The data may then be used to determine features such as the number, frequency, and/or distribution/pattern of puffs/acts of consumption within one or more chosen periods, and/or the duration, volume, average airflow, airflow profile, average ingredient ratio, and/or heater temperature values for one or more vaping/consumption events, using a processor of the delivery device and/or the any other device within the delivery ecosystem.

Optionally at least one sensor of a sensor platform may be adapted to sense at least two of puff profile, puff frequency, puff duration, number of puffs, session length, peak puff pressure and determine the state/mood of the user from the sensed information.

Puff profile, for example, characterises the variation of inhalation strength over the duration of an inhalation (or statistically over a cohort of inhalations), and may indicate for example short sharp inhalations that are relatively shallow or relatively deep and may for example be indicative of higher stress or a feeling by the user that they have a need for more of the active ingredient, or slower and longer inhalations that may be relatively shallow or relatively deep and be indicative of lower stress. Hence for example the airflow rate of a puff may be used to characterise the puff profile, with higher airflow rates associated with short sharp inhalations being likely indicative of high stress than lower airflow rates.

Puff frequency may similarly have a correlation with stress such that in stressed conditions the puff frequency may be higher than when the user is calm.

Puff duration may be considered a subset of puff profile. In puff profile, the variation in inhalation strength (for example as indicated by a proxy measure of airflow rate) over the duration of the inhalation provides the profile and when integrated also the total inhaled volume of the puff. However to a first approximation, the duration is also indicative of the type of inhalation being taken, typically with a correlation between shorter puffs in stressful situations and longer puffs with the user is calm.

The number of puffs within a session can also be indicative of the user's state. A session can be understood to either be a fixed period of time, such as hourly intervals or intervals of in minutes, where N may be any suitable value, such as for example 1, 5, 10, 20, 30, or 45 minutes, or a session can be defined functionally as a period comprising inhalations that are separated by less than a predetermined period of time that is taken to indicate that the session is over. This period may be any suitable value such as for example again 1, 5, 10, 20, 30, or 45 minutes.

In any case, for any given session, all else being equal the number of puffs taken by a user is likely to be greater when the user is in the stressed state than when the user is calm.

Similarly, where a session is defined functionally, sessions are likely to be shorter when the user is in a stressed state than when the user is calm.

Peak puff pressure may also be considered a subset of puff profile, and is indicative of how sharply user inhales. Both the peak pressure and its relative position within the duration of the inhalation may be characteristic of &of inhalation performed by the user during the puff. A high peak, particularly if early in the inhalation, is indicative of user stress, or the user's perceived wish to ingest more of the active ingredient. Meanwhile a low peak, typically in the middle of the inhalation, is indicative of the user being less stressed and simply maintaining a rate of ingestion close to their preferred baseline level.

Alternatively or in addition to the number of puffs within the session, the frequency of puffs within a predetermined period such as 24 hours, or one or more sessions as described above, or the period of time at a given location (e.g. work/home), may follow a predictable pattern; as a nonlimiting example a user may have bursts of frequent use early in the day, as lunch break, and shortly after work, and a small increase in frequency late in the evening before bed. This frequency pattern may be learned and used to anticipate the user's state, and/or to be used as a used as a factor where the user's usage pattern deviates from the learned pattern. It will be appreciated that the frequency of puffs is only one feature of inhalation based user interaction that may be subject to pattern analysis; for example the distribution of inhalation actions within a predetermined period may have a characteristic property that then may be used to predict the user state and/or to detect deviations from habitual behaviour. Hence for example either as a function of frequency and/or distribution, if the user cannot vape during work meetings, and as a result frequency drops effectively to 0, and/or the usage distribution shows a prolonged gap compared to the learned normal distribution for the user, then this is likely to be indicative of stress.

It will be appreciated that any other measurable property described herein, such as depth of inhalation, duration of inhalation and the like, which may vary on average throughout the day, may be modelled as a pattern or distribution that may be used for prediction purposes or to identify deviations from normal behaviours or situations. Such profiles may be built up for a single notional day, or a notional working day and rest day, or notional individual days of the week.

It will be appreciated that the above measurements may be obtained using one or more sensors of a sensor platform, such as an airflow rate sensor, air speed sensor, dynamic pressure sensor, microphone, or the like, whose measurements can be related to the degree of inhalation by the user and hence used to provide intra- and inter-inhalation data such as that described above.

In any event as noted above, such information may then be packaged and sent to the obtaining processor as one or more user factors.

Manipulation/handling based interactions may relate to how the user interacts with the delivery device when not actively vaping on it; for example to characterise whether the delivery device is kept in a bag until immediately prior to use, or whether the user plays or fidgets with the delivery device in between uses.

Hence for example the delivery device or any other handheld device within the delivery ecosystem, such as the user's mobile phone may comprise a sensor for detecting handshake; that is to say, small involuntary movements (so-called micro-movements) of the user's hand, such as trembling. Such micro-movements may be indicative of a state of the user. For example the amount, frequency, or prevalence of such micro-movements, and/or the amplitude of such micro-movements, are likely to have a correlation or correspondence with one or more of user stress, user fatigue, user focus, and a user's deviation from a preferred baseline amount of active ingredient within their body.

The delivery device may comprise one or more touch by sensors or accelerometers to determine such interactions. Similarly, the device may comprise buttons and other settings for which user interactions may be logged. Interactions with buttons and other settings relating to the delivery device on a companion mobile phone may also be logged. Such interaction data may then be packaged and sent to the obtaining processor is one or more user factors.

It will be appreciated that detecting touch may be one of several functions of a sensor in a sensor platform; for example physiological data may also be obtained using such sensors, or conversely such physiological sensors may also provide a touch detection function. Hence a galvanic skin response detector and/or heart rate detector may detect a touch and other physiological properties of the user at the same time. Such a sensor may be located on a grip part of the delivery device, for example where one or more of the user's fingers and/or where the user's palm are likely to hold the device for a prolonged period of time (for example when compared to contact with the mouthpiece of the delivery device or any buttons or other user interface elements of the delivery device).

Galvanic skin response detectors typically work by measuring skin conductivity or electrodermal activity, which in turn is typically a function of user perspiration (often in minute amounts, and typically do this by applying a low constant voltage to the user's skin (for example through a grip part of the delivery device) and then measuring how skin conductance (resistance) varies. Typically there is a tonic or slow fluctuating component in the order of seconds to minutes, and a faster varying phasic component fluctuating within seconds. Either component may be indicative of the user's state and may hence be a physical property contributing to user factors for the user feedback system. Notably both positive and negative stimuli (for example joy or stress) can increase galvanic skin response; hence optionally other contextual information may be useful to disambiguate the signal. However, separately there is also a clear correlation or correspondence between galvanic skin response and the consumption of certain active ingredients such as for example nicotine.

Meanwhile heart rate detectors of the type most frequently found in wearables and which may for example be found in the delivery ecosystem, for example in a wearable or in a mobile phone or delivery device typically comprise an LED light source and sensor; the sensor detects reflections from the light source after it has passed through the user's skin and been reflected back at least in part by blood as it pulses through veins and arteries; the pulsing action results in a characteristic variance in the amount of light reflected, and this can be detected to determine the user's heart rate. It will also be appreciated that similar heart rate detectors based on electrodes are available (electrocardiogram or ECG sensors), which detect the electrical activity of the heart, or variations in electrical properties associated with blood pulses.

As noted elsewhere herein, the user's heart rate (whether instantaneous or averaged over a predetermined period of time) may be indicative of their state, and hence may be a physical property contributing to user factors of the user feedback system. Similarly, variability of the user's heart rate may be indicative of the user state, with high variability being associated with stress. It will be appreciated that a heart rate monitor can in principle generate instantaneous, average, and/or variability-based data using the same sensors.

Other sensors associated with physiological measurements may be similarly optionally included within the delivery device or any other device of the delivery ecosystem that the user is likely to interact with in a manner enabling such measurements. These include for example a muscle tension sensor, and/or a cortisol sensor.

Muscle tension can be detected using an electromyogram (EMG), which again may use surface electrodes; typically the EMG data is based on a voltage difference between a recording site and a reference site, where the reference site typically is a bony low muscle point in the body. For a handheld device such as the delivery device, therefore an appropriate site for the reference electrode may be coincident with the fold of a finger or thumb joint; such a position can be predicted based on the

moulding of the device (for example a grip portion) and the location of activation buttons or any other user interface elements.

Meanwhile cortisol can be detected using a sensor known in the art and positioned on the mouthpiece of the delivery device; cortisol can be measured in saliva, and so this may be measured from the lips of the user during an inhalation action. Cortisol is also present in sweat, and so in principle could alternatively or in addition be detected using a sensor incorporated into the body of the delivery device where it is held by the user. As noted elsewhere herein, there is a correlation between cortisol and stress levels in a user.

It will be appreciated that electrodes built into the delivery device, for example in the grip region (or any other device in the delivery ecosystem, as described elsewhere herein) may be used for two or more modes of detection such as galvanic skin conductance, heart rate, muscle tension or the like, either in parallel by respective analyses of the same raw signal data, or in a sequential cycle.

Such sensors typically require two electrodes to measure skin conductivity between them. On a relatively small delivery device, optionally the electrodes may be concentric (for example an outer circle and inner circle or disc/point) in order to provide a compact sensor that may be used for example with a fingertip.

The delivery device itself, and/or in combination with any other suitable device of the delivery ecosystem, may optionally comprise one or more of the above sensors in any combination.

Interaction with buttons or other user interface elements may also provide information relating to a state of the user during use of the delivery device. For example, in a delivery device where activation uses a button press or other UI interface, the delivery device may measure the time between such activation and inhalation occurring. This period of time is likely to have a correlation or correspondence with one or more of user stress, user fatigue, user focus, and a user's deviation from a preferred baseline amount of active ingredient within their body. Hence for example the period of time is likely to be shorter if the user is stressed that if the user is calm.

Similarly, the degree of force applied to a button or element of the user interface may be measured, for example in terms of peak force/pressure applied, and/or a force profile, may be indicative of a state of the user. Hence for example a high degree of force (for example above a predetermined threshold) and/or a short interaction with the button or other user interface element may be indicative of user stress, and hence there may be a correlation or correspondence between the degree of force or the shortness of activation and a degree of stress of the user.

As noted above, the delivery device may comprise one or more accelerometers, and/or similarly gyroscopes or other motion sensors, through which motion of the delivery device may be determined. Using telemetry from one or more such motion sensors within the delivery device, the user feedback system can detect for example incidental or subconscious manipulation of the device; for example changes in orientation whilst overall position remains within a predetermined radius and/or moves slowly or generally in a horizontal direction; such motions are indicative of the user toying with the device within their hand whilst either stationary or walking. Such toying may be indicative of a state of the user; for example it may be indicative of at least a subconscious wish to use the device, or to use the device more than is currently the case, and hence correlate with heightened stress, a lack of focus, and/or a user's deviation from a preferred baseline amount of active ingredient within their body.

Similarly such telemetry can be used to detect characteristic gestures associated with use, such as lifting the device up and into an engagement position with the user's mouth, and any subsequent disengagement motion. The speed and/or exploration of these actions may similarly correlate or have a correspondence with the user's mood, for example with more rapid movements being associated with increased stress, and slow movements being associated with the user being calm.

Likewise such telemetry can be used to detect characteristic gestures not associated with use, such as gesticulation by the user, or gross movements of the user for example when climbing the stairs or using a lift, or travelling at speeds and/or in speed profiles consistent with cycling, driving by car, travelling by bus, train or plane; these activities in turn may indicate the state of the user, either in terms of their internal state with regards to breathlessness or exhaustion (in relation to gross movement), agitation or stress (in relation to gesticulation), or in terms of their external state with regards to how easily they can use a delivery device, for example when cycling or on public transport.

Such telemetry can likewise be used to detect other motion, such as a small pendulum action associated with being in a bag, or a larger pendulum action associated with being held in the user's hand as they walk, or a pattern of motion consistent with being in a user's pocket.

In addition to physical manipulation, other interactions with the delivery device or with devices within the delivery ecosystem may optionally be similarly evaluated. For example a microphone in the delivery device or the user's mobile phone may be used to detect the user's voice (for example when speaking specifically to the device, or to other people nearby, or on a phone call, or optionally as an ongoing background activity in a manner similar to a voice activated personal digital assistant). Properties of the user's voice such as volume, word speed, timbre, tonality, pitch, and/or non-harmonic content may be analysed to determine whether the user is vocalising in a calm or a stressed manner, optionally after calibration for example to the user's neutral voice. Similarly optionally such a device in the delivery ecosystem may monitor for keywords indicative of different states of the user, whether positive and/or negative.

In a similar manner to vocal expression, alternatively or in addition optionally facial expression may be monitored. In this case, the delivery device or devices within the delivery ecosystem such as the user's mobile phone, or a vending machine, may comprise a camera. In the case of the delivery device, it may comprise one or more cameras positioned to have the user's face within its field of view during inhalation and/or during the action of lifting the device to the user's face (for example on a similar side to the mouthpiece); alternatively or in addition there may be a camera facing away from the user during inhalation, in order to capture details of the user's environment.

Data from images from such cameras can be obtained pertaining to the user's state, including for example the user's overall facial expression, which typically has a strong correlation with the user's subjective mood, but also for example muscle tension in the face, which tends to correlate with stress, strain, or pain. Meanwhile eye movements can indicate a user's degree of focus and/or the nature of some activities the user is undertaking (for example patterns of eye movement and/or blinking will be different when driving, reading, or socialising, and tend to differ when a person is alert or drowsy). Similarly, if capable of being resolved by the camera, micro-movements in the face or neck can be indicative of heart rate.

Such a camera may also be used to obtain other data, such as for example motion based on the relative movement of a scene relative to the camera or key points therein, the detection of people significant to

the user, such as a partner or children, the extent or nature of a social situation, such as the number of people in proximity to the user. Similarly such a camera may be used to determine whether the user is indoors or outdoors, based for example on the detection of sky, colour temperature, light flickering, characteristic indoor features such as windows or TV screens, or the like.

It will also be appreciated that user interaction may comprise a specific indication of user state by the user. In this case, a user interface is provided that allows the user to select a setting that is an indicator of their state. This indicator may be explicit, for example providing a selection of user states and optionally values (for example from 1 to 100 indicating the degree of a state, so that the user can directly input their subjective assessment of their own state. As noted elsewhere herein, this may be useful for the evaluation processor, and/or for evaluation model training purposes or the construction of rules or look up tables for associating user factors with user states. Alternatively the user interface may be more indirect, for example having a 'calm' mode and 'boost' mode, where the mode is a default for when the user is calm, whereas the 'boost' mode delivers more of the active ingredient per volume of aerosol inhaled and hence may have a correlation with user stress.

In a similar manner to the indications provided by use of a calm mode and boost mode, the selection of a particular consumable (for example one with a normal or calm concentration of active ingredient or one with a high or boosted concentration of active ingredient) may be indicative of the user's degree of stress or calm, typically the start of the day when such consumables are being selected; hence this may be indicative of more chronic levels of stress.

It will be appreciated that where a user has multiple delivery devices 10, usage may be aggregated across these devices, either by obtaining user factor data from each device, or already aggregated via an intermediary such as a phone app or one of the delivery devices acting as a hub for this purpose. Where different devices deliver different active ingredients (whether type or concentration), this may also be accounted for in modelling use, as a non-limiting example in relation to pharmokinesis.

Multiple data sources

As noted above, and as shown in Figure 6, the obtaining processor may receive multiple user factors of the types described herein from one or more sources, such as those in the delivery ecosystem 1, the Internet 110, and records held by the feedback system 1012, for example at the server 1000.

As noted above, these user factors may variously be classified as indirect or historical data; neurological or physiological data; contextual data; environmental or deterministic data; and/or use-based data.

In the case of use-based data, it will be appreciated that multiple sensors, and/or a sensor with multiple sensing capabilities may be used in a sensor platform to obtain some or all of such use-based data.

Obtaining processor operation

Turning again to Figure 6, the obtaining processor 1010 is typically part of a remote server 1000, and may receive user factors from diverse sources such as the server's own storage/database 1012, on-line sources 110, and devices within the user's delivery ecosystem 1, such as the delivery device 10 itself, a mobile phone 100, a fitness wearable 400, a docking unit 200, a vending machine 300, and any other suitable device that may provide information relevant to the user's state, such as a voice-activated home assistant, smart thermostat, smart doorbell or other Internet of things (IOT) device.

The obtaining processor 1010 may comprise one or more physical and/or virtual processors, and may be located within the remote server, and/or its functionality may be distributed or further distributed over multiple devices, including but not limited to the user's mobile phone 100, a docking unit 200, a vending machine 300, and the delivery device 10 itself. The obtaining processor may comprise one or more communication inputs, for example via network connections, and/or via local connections to local storage. The obtaining processor may also comprise one or more communication outputs, for example via network connections, and/or via local connections, for example to the estimation processor 1020.

The obtaining processor may comprise pre-processors or sub-processors (not shown) adapted to parse and/or convert obtained information into user factors where this information is not immediately usable as such; examples may include keyword or sentiment analysis of consumed media, for example to determine as a user factor a net positive or negative influence on an aspect of user state, or similarly keyword analysis of the user's calendar to determine locations and events, for example again to determine as a user factor a net positive or negative influence on an aspect of user state. Other inputs, such as ambient temperature or probability of rain, may similarly be converted to a scale appropriate to user factors, for example being normalised or classified according to influences on user state. Similarly noisy data may be processed to remove statistical outliers or to perform smoothing functions, or calculate averages or other statistical values, or the like. It will also be appreciated that such pre-processing or sub-processing may be performed at one or more devices within the user's delivery ecosystem on behalf of the obtaining processor.

The obtaining processor may thus be operable to generate and/or relay user factors for input to the estimation processor at varying degrees of abstraction from the original source material.

Hence optionally original source data may be enumerated, codified, classified, formatted, or otherwise processed, or simply passed through and provided as input to the estimation processor, so that there are potentially as many or more inputs as there are original sources of data. As will be appreciated from the description above, this may result in a large number of inputs.

Hence optionally one, some or all of the original source data may be any rated, codified, classified, formatted, or otherwise processed, or simply passed through as appropriate to an optional intermediate user factor generation stage of the obtaining processor; this may determine positive or negative influences from the submitted inputs on a specific subset of user factors that may be relevant to the user state but not directly or easily measurable, such as effects on dopamine and/or cortisol, heart rate, satiety, and the like.

Similarly such an intermediate user factor generation stage of the obtaining processor may combine inputs from similar classes to generate a class-level user factor for one or more of the classes of data described herein.

Hence as non-limiting examples, indirect or historical data could be summarised as how actively the user modifies or updates their device, or how receptive they are to such modifications, on a given scale. Neurological or physiological data could be summarised as how stressed the user appears to be, on a given scale, and/or their trajectory on that scale. Contextual data could be summarised as how sociably desirable use of the delivery device is currently, on a given scale. Environmental or deterministic data could be summarised by how likely the user is to want to use the delivery device in a given timeframe; and use-based data could be summarised as how frequently or deeply the user is or has recently used the delivery device.

It will be appreciated that in practice only source data from some or one of the classes may be available, and even where data from one class is available, a class-level user factor such as in the examples above may not be generated, or different kinds of class level user factor may be generated depending on the type of data received within that class (e.g. different subsets of individual user factors); similarly, class-level user factors may be generated for input to the estimation processor in parallel with individual user factors.

The contributing values and/or influences from different individual, subset and/or or class level user factors may then be presented as inputs to the estimation processor, with the selection of class, subset and/or individual user factors being chosen to give a good discrimination between different user states.

For example, galvanic skin response may provide a good indicator of a user's state, and is also responsive to nicotine as an active ingredient by reducing the response; as such it may optionally be a candidate for an individual source of data to be used as an input to the estimation processor. Other physiological measures to provide good discrimination include muscle tension (EMG), heart rate, skin temperature, brainwaves (EEG), and breathing rate. Any of these, where available, may be considered for inclusion as an individual source of data, optionally after being any rated, codified, classified, formatted or otherwise processed, alternatively or in addition combined in any combination with these or other user factors described elsewhere herein.

Similarly location, social setting, time-of-day, and hormonal levels are all good indicators of the user's state and may be candidates for use as individual sources of data as input to the estimation processor.

Hence more generally user factors may be obtained by or for the obtaining processor and provided to the estimation processor after any suitable parsing or processing, either individually and/or as combined subset or class values with one or more others (for example based on weighted contributions, statistical functions, trained machine learning outputs, look-up tables of precomputed correspondences between values of the obtained data and values of a target user factor, and the like), as for example individual, subset and/or or class level user factors.

Estimation processor

The estimation processor 1020 is operable to calculate an estimation of user state, based upon one or more of the inputs received from the obtaining processor comprising or based upon obtained user factors. The calculation of an estimation of user state can be either explicit to generate an output reflective of a user's state prior to generating a proposed feedback action (which may be thought of as a two-step process), or implicit to identify a proposed feedback action expected to alter a user's state (which may be thought of as a single step process).

Like the obtaining processor, the estimation processor may comprise one or more physical and/or virtual processors and may be located within the remote server, and/or its functionality may be located in a device of the delivery ecosystem, such as the delivery device 10, or distributed or further distributed over multiple devices, including but not limited to the user's mobile phone 100, a docking unit 200, a vending machine 300, and the delivery device 10 itself. The estimation processor may comprise one or more communication inputs, for example to receive data from the obtaining processor 1010. The estimation processor may also comprise one or more communication outputs, for example to provide a proposed feedback action to the feedback processor 1030.

Explicit state estimation

In an embodiment of the description, in a two-step process the estimation processor initially explicitly estimates a state of the user in a first step before then generating a proposed feedback action in response to the estimated state in a second step. This estimated state may itself take the form of a single value or category, or may be a multivariate description of the user's state.

As non-limiting examples of a single value state, the estimated state may describe:

i. a stress level of the user;
ii. a degree of benefit the user is expected to subjectively experience in response to a unit consumption of a proposed active ingredient; and
iii. a social flexibility score, indicative of how easily the user can currently use the delivery device and hence alter their state through modification of delivery;

As non-limiting examples of a state category, the estimated state may be:

i. one of a plurality of state classifications, all, some, or none of which may correspond to what are colloquially referred to as moods; hence for example happy, sad, low cortisol, medium cortisol, high cortisol, calm, stressed, receptive to change (for example willing to use their delivery device to alter their state), or unreceptive to change.
ii. one of a plurality of state classifications chosen to have a subsequent clear correlation with either inputs from the obtaining processor and/or an available feedback action, the classifications not necessarily fitting a notional category such as 'happy' or 'high cortisol', but having classification boundaries driven at least in part by their correspondence to either the available inputs from the obtaining processor or outputs for the feedback processor.

As non-limiting examples of a multivariate description of the user's state, the estimated state may comprise:

i. the user's stress level according to physiological indicators, and separately according to contextual indicators, together with an indication of their current social flexibility based on time-of-day, location, and/or proximity to specific individuals;
ii. an indicator of the user's physiological state based upon galvanic skin response and heart rate, together with current position in a hormonal cycle, and indicators of mental state derived from questionnaire and/or social media analysis.

These examples may be used to provide non-limiting illustrations of the operation of the estimation processor, as follows.

The estimation processor may use predetermined rules, algorithms and/or heuristics to convert input data from the obtaining processor into estimated states.

- For example, a single value state such as a stress level of the user may be derived by applying a predetermined combination to a plurality of user factors, such as a weighted sum, with the result normalised according to the number of currently available inputs contributing to the sum.
- Similarly a single value state such as the degree of benefit expected for the user may be derived by estimating the user's positive or negative emotional state based upon summing indicator values for positive or negative keywords or sentiments in on-line media recently consumed or produced, and positive or negative values associated with a classification of the user's location.
- Likewise an estimated state category may be selected by template matching user factor values to predetermined values indicative of a given category, or similarly identifying a least-mean-

squares error between user factors and a template of user factor values for each candidate category, optionally with different categories and greater error having different linear or non-linear weightings, reflecting their relative salience in identifying the category.

- Finally as an example, a multivariate state may comprise deriving individual indications of state according to any of the above examples; hence a single value stress level may be generated as discussed above for each of physiological and contextual indicators, and a social flexibility value may be determined based upon scores previously associated with different times of day, location and class of specific individual (e.g., partner versus child); or a social flexibility classification may be based matching templates to such scores and/or values for the underlying input data.

Alternatively or in addition, the estimation processor may use look up tables to convert input data from the obtaining processor into estimated states.

In one instance, these look up tables may simply provide a precomputed implementation of the above predetermined rules, algorithms and/or heuristics, to avoid repetition of these calculations either at the server, or on a device within the delivery ecosystem that has limited processing capability but is acting as the estimation processor or sharing its role, such as the delivery device 10, or a dock 200, vending machine 300, wearable device 400, or associated phone 100.

In another instance, such look up tables may provide associations between input values from the obtaining processor and output values of user states, state classifications and/or multivariate states previously derived according to any suitable mechanism, such as for example feedback from extensive user testing, or as described later herein, the output of a machine learning system; again in this latter case, a look up table may potentially provide a computationally simpler facsimile of such a machine learning system by recording pairs of inputs and outputs for common values that may be easier to implement on devices within the delivery ecosystem having a comparatively low computational power.

Alternatively or in addition, the estimation processor may model correlations between input data and estimated states of the user. Such correlations may be due to causal links between a user factor and a user state, or a tendency for the user factor to accompany a cause of the user state, hence acting as a proxy, typically with particular degree of probability. Similarly such correlations may be due to the user factor and the user state both responding to a separate cause or circumstance in a manner that is sufficiently repeatable to form a correlation. Likewise such correlations may be due to the user state giving rise to the user factor. Hence more generally correlations relate to measurably predictable correspondences between one or more user factors (whether individual, subsets or class level user factors as output by the obtaining processor) and a user state (whether a single value, a classification or multivariate), typically either due to a causal link (in either direction between user factor and state), a common cause resulting in responses in user factor and user state with a repeatable relationship at least at a statistical level, and/or a measurable correspondence regardless of whether a direct or indirect causal link is known.

Where the estimation processor models correlations, it can be trained using a data set comprising as inputs data corresponding to the above described outputs of the obtaining processor, and as target outputs descriptors of a state of the user, whether single value, a classification, or multivariate, for example based upon a direct measurement of a user's state and/or user self-reporting regarding their state.

The specific means by which such correlations may be derived include any suitable technique for estimating such correlations, including a correlation map between inputs and outputs, where presentation of an input and output at the same time (or within a predetermined time window, if a temporal factor is included) results in a reinforcement of the link between the specific inputs and outputs (for example by increment of a connective weight). Once trained on the dataset, a new input will result, by virtue of the connective weights, in the activation to a greater or lesser extent of one or more candidate states correlating with that input; the candidate state with the strongest activation may then be chosen as the user state, or such states may be ranked by activation strength. It will be appreciated that in such a system, multiple input values may be provided simultaneously, corresponding to individual, subset of class level user factors as described elsewhere herein, and the generated outputs may correspond to a single value state, a classification, or a multivariate state with a number of values representing different aspects of the user state as described elsewhere herein being output.

A specific example of a correlation map is a neural network, and any suitable form can be considered.

More generally, any suitable machine learning system capable of determining a correlation or other predictable correspondence between one or more inputs and one or more outputs may be considered.

Given the above described dataset, such machine learning systems are typically supervised and may for example be a supervised classification learning algorithm, for example if the user state is a classification; or a supervised regression learning algorithm, for example if the user state is a single value or multivariate. Other forms of machine learning are also suitable, such as reinforcement learning or adversarial learning, or semi-supervised learning. Furthermore multiple independent machine learning systems separately trained on different or partially overlapping individual, subset or class level outputs of the obtaining processor can be ensembled to improve modelling results, for example to accommodate different configurations of source data due to different patterns of ownership of devices in the delivery ecosystem of different users, and different permissions and habits affecting the availability of online sources of information. It will also be appreciated that a mixture of different machine learning systems can be used in parallel, for example to generate a multivariate state of the user, with for example one or more different elements of the multivariate description been generated by different respective machine learning systems. These respective machine learning systems can be on separate hardware (e.g. based on dedicated neural processors) but more typically may be on the same hardware (e.g. software based machine learning systems that are loaded and run as required).

Meanwhile unsupervised learning algorithms may also be considered; hence for example associative learning may determine the probability that if one input or pattern of input is present then the user will be in a given state.

Examples of the above machine learning systems, in the forms of algorithms and/or neural networks, will be known to the skilled person.

Meanwhile machine learning may optionally also be used to prepare (e.g. pre-process) data, either at the estimation processor and/or in the obtaining processor; hence for example clustering (for example k-means clustering) may be used to classify a diverse set of inputs into a class level user factor of the type described previously herein. Such an approach may be used or also used for example to derive classifications for user states, as per the second example of a state category classification described previously herein, in response to inputs from the obtaining processor or available feedback actions of the feedback processor.

Similarly as a preparatory step at the estimation processor and/or obtaining processor, dimension reduction, such as principal component analysis, may be employed to reduce the number of inputs whilst retaining information having a significant correspondence with the user state.

Hence in summary the estimation processor, if generating explicit estimates of user state, uses a repository for the correspondence between available inputs from the obtaining processor and the estimated states, where that repository for the correspondence may be embodied in algorithms, rules or heuristics, and/or in one or more look up tables, and/or in one or more trained machine learning systems.

In each case, the result is an estimation of the user state, which may take the form of a single value, a category, or a multivariate description/representation of the user state as described previously herein.

Meanwhile the operation of the estimation processor if generating implicit estimates of user state, is described later herein.

Feedback proposals from an estimated state

As noted previously herein, the estimation processor may operate in a two-step process; in the first step estimating a user state from inputs comprising one or more user factor or data derived from such user factors by the obtaining processor, as described previously herein, and in the second step generating a proposed feedback action expected to alter a user's state, as will be described below.

In principle, the second step may be implemented by the feedback processor rather than the estimation processor, or may be shared between the feedback processor and the estimation processor. Alternatively the feedback processor may simply receive the proposed feedback action. In any case, the feedback processor may then select the feedback action (either by default if only one is proposed, or selecting one or more if a plurality are proposed), or optionally act to cause one or more feedback actions proposed by the estimation processor to occur in an appropriate manner within the delivery ecosystem.

For the purposes of explanation, the second step is described herein as occurring within the estimation processor.

The two-step process may be chosen for practical reasons; for example, training sets for use in modelling the correspondence/correlation between user factors or their derivations by the obtaining processor and user state may be easier to generate or acquire than training sets for use when directly modelling the correspondence/correlation between user factor based inputs and proposed feedback actions, because the user's state may be either directly measurable, or straightforward for a user to report.

Similarly it may be easier to generate a training set determining the correspondence/correlation between a measurable and/or self-reported user state and a proposed feedback action, based for example upon user questionnaires ranking feedback actions for given states, and/or upon the subsequent effectiveness of an implemented feedback action in altering the state of the user toward a more desirable state, as measured and/or reported by the user. Typically a more desirable state is one which improves the user's subjective sense of well-being and/or moves physiological or neurological indicators of the user's state toward a preferred norm (for example reducing elevated heart rate, galvanic skin response, elevated skin temperature, and/or breathing rate, and the like).

The input for the second step will typically be an estimation of the user state, represented by a single value, a category, or a multivariate description as described previously herein, or a plurality of these if multiple states are estimated (for example with varying degrees of activation/strength of correlation in response to the inputs of the first step). Optionally, inputs to the second step may also comprise one or more user factors and/or inputs as provided by the obtaining processor; for example, as described elsewhere herein certain physiological measurements may be useful indicators/proxies for the user state, such as galvanic skin response, heart rate, breathing rate, skin temperature and the like. Hence optionally one or more of these or any other inputs to the first stage, may also be provided for the second stage in conjunction with the or each estimated state.

In any event, as with the estimation of the user state, the generation of a proposed feedback action may use any suitable mechanism that embodies a correspondence/correlation between the estimated user state and the proposed feedback action.

As noted previously, this may include predetermined rules, algorithms and or heuristics to convert estimated states into proposed feedback action.

- For example, a single value state (such as degree of stress) may drive a corresponding proposed feedback action, such as increasing the proportion of active ingredient within an inhaled unit volume of generated aerosol, which in turn may be achieved by modifying heater, air flow, reservoir and/or other payload storage settings, and the like, which as described later herein may be managed by the feedback processor. The relationship between the degree of stress and the change in active ingredient may be linear or non-linear, or may change qualitatively at different values, for example not changing at all for low levels of stress, have a linear relationship for medium levels of stress, and have an asymptotic relationship for high degrees of stress up to a maximum proportion of active ingredient, and for example at or near this maximum also modifying a behaviour of the user interface of the delivery device or other device within the ecosystem, such as issuing a warning or calming message on the user's phone.
- Meanwhile for example a single category state may have a corresponding proposed feedback action.
- Finally for example a multivariate state may result in a corresponding proposed feedback action being based on weighted or unweighted contributions from the different elements of the state description, and/or different feedback actions may be proposed based on overlapping or non-overlapping subsets of the elements of the state description. Hence for example if the state description suggests that the user is stressed and are in a work environment, then the feedback action may assume that they are implicitly stressed because they are in the work environment, but are currently unable to increase their intake of active ingredient, and therefore issue a message on a UI of the delivery device or other devices in the ecosystem, such as the user's phone, to suggest the user that they take a break. Meanwhile if the user is stressed but not in their work environment, then the feedback action may be similar to the degree of stress example above, resulting in an increase in the proportion of active ingredient delivered to the user.
- As noted above, any one of these may also be accompanied by one or more inputs to the first step.

Again like the estimation of user state, the estimation processor may alternatively or in addition use look up tables to convert state estimation data into proposed feedback actions.

Alternatively or in addition, again like the estimation of user state, the estimation processor may model correlations between estimated user states and proposed feedback actions, and the use similar techniques to do so.

Where the estimation processor models correspondences/correlations, it can be trained using a data set comprising as inputs data corresponding to the estimated user state (for example in the form of single values, classifications, or multivariate descriptions, or a combination of these) and optionally also inputs from the obtaining processor as described previously herein, and as target outputs proposed feedback actions.

The proposed feedback actions are discussed in more detail later herein, but may typically comprise at least one type of action and optionally one or more variables characterising the performance of that action. Hence for example a change in vaporisation temperature is a type of action, and an increase or decrease, or amount of increase or decrease, would represent a variable characterising the performance of that action. Similarly modifying active ingredient concentration and/or amount in the aerosol is a type of action, and an increase or decrease in concentration and/or amount, or an amount of increase or decrease, would represent a variable characterising the performance of that action.

Hence as a non-limiting example in the context of a machine learning system, different output nodes may represent different types of action, and the values of those nodes may represent either a flag indicating selection of that feedback action, or a value relating to a variable of that feedback action, depending on how the system is trained. It will also be appreciated that multiple output nodes may be associated with one or more types of action in a machine learning system, depending on the training regime.

It will be appreciated that potentially a plurality of feedback actions may be indicated in response to an estimated user state. In such circumstances, the feedback processor may subsequently determine whether to select just one feedback action, for example based upon the degree of change caused by the action as implied by its associated variable or variables, or implement multiple feedback actions in parallel or sequentially, in the latter case optionally a sequence determined by a predetermined order, or again responsive to the strength of activation of a flag output node for each feedback action, and/or the degree of change implied by each action's associated variable or variables.

It will also be appreciated that to train such a machine learning system, measured and/or reported user states could be provided as inputs, and respective proposed feedback actions could be provided as targets, with actions and values selected according to their reported efficacy during user trials for users having the corresponding user state; again in this case efficacy or effectiveness typically relates to the user's perceived improvement in state, and/or a change in neurological and/or physiological state toward a predetermined norm or preferred state.

Optionally as first training phase, simulated states and corresponding feedback actions could be used to provide initial training (for example based on questionnaire results as described previously), with a proportionally smaller cohort of real-world training data then being used to refine the model.

Optionally, feedback from the user themselves as to the efficacy and/or suitability, desirability, practicality etc., of any feedback actions could be further used to refine the model and effectively personalise it to the user. Again this feedback may be reported by the user for example via a user interface on a device within the delivery ecosystem such as the delivery device or their phone, and/or

based on measurements of neurological and/or physiological response. Where plural feedback actions are implemented or indicated, optionally the user may rank them in order of preference.

In summary, the two-step process, comprising an explicit estimation of user state as a first or interim step, may be of use where these steps better fit the available underlying empirical data sets used to model the correspondences/correlations, whether this is done by rule-based techniques or machine learning.

Objectively, the operation of the estimation processor in this mode is thus to take inputs from the obtaining processor, typically in the form of different individual, subset and/or or class level user factors, and output one or more proposed feedback actions either simply identifying the action in a manner similar to a flag, and/or identifying the degree of relevance of that action to the estimated state based on the activation level and output corresponding to the proposed feedback action, and/or indicating a change or amount of change of one or more variables that at least in part characterise the proposed feedback action.

The explicit estimation of the user state is thus typically an internal, interim step. However it will be appreciated that this estimate could be relayed to the user for their information, and optionally the user could modify the estimate, particularly where the estimate or a component of the estimate in a multivariate description relates to a subjective measure or to a proxy of a subjective measure such as the user's sense of stress. Hence for example the estimate could be displayed on a user interface of the user's mobile phone, and the user could use this information to self-assess, and make changes to the estimate as a result. The modified estimate of the user's state could then be used together with or instead of the originally generated estimate in the second step to identify/generate a proposed feedback action that may be more accurate than the proposal based on the original estimate of the user's state.

Furthermore, any changes made to the estimate of the user's state could be used to update and refine the model of the first step, and indeed for certain machine learning techniques, a lack of correction by the user may similarly be taken as a positive reinforcement of the estimate for the purposes of training.

As mentioned previously herein, if further training is not desired, then optionally the relationships between input and output values derived by the machine learning process may be captured in one or more look up tables, which may be computationally simpler to use (though may occupy more memory).

Implicit state estimation

In an embodiment of the description, rather than using the two-step process described above, the estimation processor performs a single step process that implicitly estimates a state of the user as part of the relationship between individual, subset and/or or class level user factors provided as input by the obtaining processor, and proposed feedback actions generated as an output and typically expected to alter a user's state.

Hence an estimation processor (1020) adapted to calculate an estimation of user state based upon one or more of the obtained user factors may equally be an estimation processor (1020) adapted to identify/generate a proposed feedback action state based upon one or more of the obtained user factors; in this case the user state is implicit in the relationship between the user factors and the proposed feedback action, which is expected to alter the implicitly estimated state of a user.

In a similar manner to the two-step process described previously herein, the estimation processor may use predetermined rules, algorithms and/or heuristics to convert input data from the obtaining processor into estimated states. These may for example combine the processes for the two separate steps of the explicit state estimation embodiments, and/or refine some or all of the rules, diagrams and/or heuristics in response to the single step nature of the implicit state estimation approach, or may be derived from scratch for the singe step process.

Again like the two-step process, the estimation processor may alternatively or in addition use look up tables to convert input data into proposed feedback actions. Again these may be concatenations of look up tables from the two-step approach, and/or may be further processed to provide single step look-up tables, or may be derived from scratch for the singe step process.

Again like the two-step process, the estimation processor may alternatively or in addition use machine learning. In this case for example, inputs used in the first step of explicit state estimation, and targets used in the second step of generating proposed feedback actions from the estimated steps, may be used to train a machine learning system that identifies measurable correspondences between them.

It will be appreciated that to present corresponding inputs and targets for training purposes, the training set should have captured this correspondence; as noted previously herein, it may be that datasets exist for the inputs and a user state, and user states and effective feedback actions; consequently inputs and feedback actions can be married for training purposes based upon the common user state value, class or multivariate descriptors as appropriate; clearly also where the training datasets were collected by users for whom user factors were measured and/or self-reported, user states were measured and/or self-reported, and subsequent efficacy and/or suitability, desirability, practicality etc., of feedback actions were measured and/or self-reported, then these self-consistent sets of input user factors (as provided by the obtaining processor) and target feedback actions can be used for training.

Alternatively or in addition, a two-step system with explicit state estimation, which has been trained on separate datasets, and/or uses respective rules, algorithms and/or heuristics from the two steps, and/or uses look up tables from the two steps, can be used as a data source.

For example, a single step look-up table may be created by running through the first and second steps of look-up tables or rules, algorithms and/or heuristics, and/or machine learning systems for a two-step estimation to provide look up links between inputs as provided by the obtaining processor, and proposed feedback actions identified/generated by running through the two-step process using those inputs.

Alternatively or in addition, a single step machine learning system may be trained by running through the first and second step of look up tables or rules, algorithms and/or heuristics, and/or machine learning systems for a two-stage estimation to provide inputs as provided by the obtaining processor, and provide as targets for training proposed feedback actions identified/generated by running through the two step process using those inputs.

Optionally, a single step machine learning system trained in this manner may then have its training refined using additional data, such as a combined training set as described above and/or, in a similar manner to that described previously herein for the step scheme, data received from one or more users during use of the user feedback system.

It will also be appreciated for example that a training set may be based directly on capturing the desired input and target values rather than using an amalgam of datasets or processes.

It will be appreciated that for either the two-step approach or the single step approach, training data may be gathered using one or more devices in the delivery ecosystem, for example to build a training set relating user factors to user states. Such a training set may be generated using a version of the user feedback system that does not generate a proposed feedback action, but simply gathers the user factors and user state information. Similarly a training set relating user states to proposed feedback actions may initially be based upon asking users, for whom their state is known (e.g. measured/reported) to rate proposals for feedback actions, for example via a user interface on their phone as part of a user testing scheme. Hence in this case the feedback system may propose feedback actions and select one or more of the proposed actions, but in different versions or modes may either present the selected proposed feedback action(s) to the user for evaluation (for example via a user interface) for example during a training-data gathering phase or a calibration phase (for example to characterise the user within a subgroup to which responses may be better tailored, as disclosed elsewhere herein), or may cause the selected proposed feedback action(s) to be implemented, modifying of one or more operations of at least a first device within the delivery ecosystem, responsive to the estimation of user state (whether explicitly or implicitly modelled), in a manner expected to alter the estimated state of a user. Training data relating user factors to proposed feedback actions may be obtained in a similar manner.

Hence such datasets may be obtained using a version or mode of the user feedback system that as noted above does not actually cause a modification to one or more operations of a device in the ecosystem (optionally except for eliciting a response from the user, e.g. for training data purposes).

This preceding generation of the user feedback system, or training/refinement mode of the user feedback system, could thus comprise an obtaining processor (1010) operable to obtain one or more user factors indicative of user state, and operable to obtain user state data (for example based on measurements similar to those of user factors, and/or self-reporting by the users), and or feedback action preference/efficacy data; the estimation processor would then comprise a training or development phase in which the correspondences/relationships/correlations between inputs based on the user factors as described previously and targets based on the user states (in the two-step scheme) or the proposed feedback actions (in the single-step scheme) are modelled as described previously, for example once a sufficient corpus of data had been amassed.

Alternatively or in addition, in such a preceding generation and/or in a training mode of a feedback system, the delivery device and/or other participating devices in the delivery ecosystem may consequently only upload data to the obtaining processor, but not download feedback actions (or optionally any other data) from the feedback system.

Similarly, in either such a preceding generation and/or in a training mode of a feedback system, and/or in providing improved or supplementary input for the feedback system, then as mentioned elsewhere herein user factors such as from neurological / physiological data (e.g. from biometric sensing), motion and/or location user factors (e.g. from touch, accelerometer or GPS sensors), contextual user factors, and/or any of the other user factors disclosed herein may be accompanied by direct input on the user's state as reported by the user. This may be used for the generation of a training set, as described previously, but alternatively or in addition the user's reported state may be treated as a user factor by the obtaining processor, or directly by the estimation processor. In principle the user's reported state may optionally be used in lieu of an explicit state estimation by the estimation processor, but it is

possible that at least in some cases the user's reported state will be approximate compared to what may be derivable, or estimated from some measurements (if these are available), or the user may not be informed by all the facts available to the feedback system. Furthermore, some users may normalise their state and self-report in a biased manner, particularly for pathological states such as depression. Hence optionally the user's direct input on their state may be used in conjunction with one or more other user factors from the obtaining processor as described above as input to the estimation processor, in the first (or only) step as described above. Optionally, alternatively or in addition the user's direct input on their state may be used in conjunction with an estimation of their state as input to the second step of the estimation processor, if the two-step technique is used.

Other variations in training and input may also be considered. For example it will be appreciated that as noted previously herein, different user factors operate or vary over different time scales. Consequently for either the two-step approach or the single step approach for the estimation processor as described herein, user factors that are not expected to have changed within an interval between successive operations of the estimation processor may be stored and re-used (for example in storage 1012), rather than being re-obtained.

Furthermore, some parts of the estimation model relating to these longer term factors may not need to be re-run if the outcomes for those factors are expected to remain the same. This may be straightforward for rule, algorithm and/or heuristic methods, and/or look-up tables, but for a machine learning system it may require a modified architecture; for example a two-stage ML or multi-layer system may be trained on all inputs, but subsequently run with inputs or outputs relating to long-term user factors clamped, and the previously computed intermediate results of that part of the ML system fed into the remaining part of the ML system in conjunction with newly generated intermediate results from user factors with shorter time frames.

It will also be appreciated that as noted previously herein, different users may have different combinations of devices within their delivery ecosystem, and/or different combinations of these devices may be active at any one time; similarly, different users may have greater or lesser presence on social media, or use their digital calendar to a greater or lesser extent, and the like. Consequently the user factors available to the obtaining processor and hence also the inputs available to the estimation processor may differ from user to user, and/or from time to time. Accordingly, the estimation processor may use different models (explicit or implicit, as discussed above) to propose feedback actions, depending upon the inputs available. Alternatively or in addition, where an input to a model is missing, a neutral input value may be provided so as to reduce or remove the influence of that missing input on the feedback action proposed. The number of different models provided by/for the estimation processor may therefore depend upon the number of data sources assumed for a model (with more sources or more diverse sources making the model potentially more fragile), and the robustness of the model to the replacement of inputs with placebo/neutral values where an input is currently unavailable; in this latter case it will be appreciated that some inputs may be more critical than others, so at least some individual inputs may be required for a model to run. Hence depending upon the complexity and robustness of the model, it may be that only one model is needed, or a suite of models anticipating different scenarios. Optionally, a subset of all available models is selected for a user depending upon the devices known to exist in their delivery ecosystem; meanwhile new models may be added when new devices join the delivery ecosystem, whether permanently for example in the case of the user buying a new dock 200, or temporarily for example in the case of the user interacting with a vending machine or point-of-sale device.

Estimation processor output

Whether a single step or two step process is used, and whether the estimation for any step is based on rules, algorithms, and/or heuristics, look up tables, and/or machine learning, the output of the estimation processor is a proposed feedback action.

Possible feedback actions differ qualitatively and/or quantitatively.

Hence for example they may vary qualitatively based on whether they relate to modifying the generation of aerosol for the user (whether in response to current circumstances or pre-emptively); modifying the user's interaction with the delivery device or system, either during inhalation or between inhalations; modifying the user interface of the delivery device or system; reminding the user to use or change their use of a delivery device or system; recommending an operation or selection of a delivery device or delivery device consumable; and/or recommending/activating/modifying the operation of a device that is not directly related to the delivery of active ingredient, but may nevertheless change the user's state, either directly (for example through biofeedback) or indirectly (for example by activating noise cancellation in a user's headphones).

Hence more generally feedback actions may fall into categories that are behavioural, focusing on altering actions and/or habits of the user to change their state; pharmaceutical, focusing on how one or more active ingredients delivered to the user may change their state; and non-consumption interventions, focusing on alternative first or third party options (i.e. relating to the delivery device or other devices in the delivery ecosystem or elsewhere) to change a user's state.

Meanwhile proposed feedback actions may vary qualitatively depending on the extent to which the effect of the feedback action is desired to make a positive change in the user's state; hence for example in the delivery device a change to heater temperature, payload aerosolisation, payload composition, or the like may comprise a quantitative value indicating the degree of change, or class of change, as appropriate. Similarly modifications to a user interface in the delivery device or another device of the delivery ecosystem may comprise incremental steps relating to the number of user interactions required or prompted with the delivery system, and the nature of those user interactions; for example running through five categories, with the first category having no notifications to minimise interruption of the user, a second category only having critical notifications such as for low battery or low payload, a third category corresponding to a default in which critical and non-critical notifications are provided, a fourth category further including recommendations and/or prompts to engage the user with other features of the user interface, and a fifth category additionally including an audible tone. These five categories may be selected according to the user state on a scale of how stressed they are (for example with minimal notifications for high stress), and/or how bored they are (for example with high notification for high boredom).

As noted previously, the type of feedback action, and/or the amount or class of change, if appropriate, may be identified according to rules, algorithms, and/or heuristics, look up tables, and/or machine learning as appropriate.

Similarly as noted previously, where more than one type of feedback action, and/or more than one amount or class of change is calculated/estimated to be an appropriate response to the user factors/user state, optionally multiple feedback actions may be proposed accordingly, or the top N

feedback actions may be selected based for example upon strength of activation, where N may be one or more.

Feedback processor

The feedback processor 1030 is operable to implement one or more proposed feedback actions, thereby causing modification of one or more operations of a device within the delivery ecosystem, responsive to the estimation of user state.

Hence the feedback processor may act to cause the feedback action or actions proposed by the estimation processor to occur in an appropriate manner within the delivery ecosystem.

The feedback action or actions are typically implemented in a manner expected to alter the estimated state of a user. This user may be considered a generic, average, notional user; it will be appreciated that the model or models upon which the generation of the proposed feedback action(s) is/are based are typically developed or trained using data from a corpus of users, and hence relate to changing the state of a generic, average, or notional user.

However, typically this will nevertheless similarly alter the state of the particular user of a respective delivery device, on the basis that most users are likely to respond in a similar manner to these changes.

However as described elsewhere herein, if the feedback system can receive further feedback from the individual user (for example by measurement or self-reporting) as to the efficacy of proposed feedback actions, then optionally the system can become increasingly tailored towards the particular user, for example through supplementary training and/or refinement of parameters, and hence implement feedback actions responsive to the estimation of user state in a manner expected to alter the estimated state of the particular user. Similarly, separate rules, algorithms, and/or heuristics, look up tables, or machine learning systems may be generated for different user groups, for example based on demographics and/or patterns of response to feedback actions, so that the proposed feedback actions are better tailored to a particular user falling within one of these groups, even if measured or reported assessments of feedback efficacy are not available from the particular user, or are too sparse to effectively refine the training of a machine learning system or alter the parameters of an algorithm etc., to personalise its response to them.

Like the obtaining processor and the estimation processor, the feedback processor 1030 may comprise one or more physical and/or virtual processors and may be located within the remote server 1000, and/or its functionality may be distributed or further distributed over multiple devices within the delivery ecosystem, including but not limited to the user's mobile phone 100, a docking unit 200, a vending machine 300, and the delivery device 10 itself. The feedback processor may comprise one or more communication inputs, for example to receive data from the estimation processor 1010, and one or more communication outputs, for example to communicate with the delivery device 10, and/or another device within the delivery ecosystem 1 such as those listed above, or any other device that may participate in a feedback action.

In particular, the feedback processor may optionally comprise a selection and notification sub-processor (not shown) which may be located at the server and/or at a device within the delivery ecosystem with suitable computational power, such as a vending machine, mobile phone, or indeed a suitable delivery device, to optionally select one or more feedback actions and select one or more respective devices within the ecosystem for implementing one or more feedback actions; and optionally an action

implementation sub-processor (not shown) at one or more respective devices within the ecosystem for managing the implementation of a feedback action. Optionally, the action implementation sub-processor may be considered a separate processor to the feedback processor.

Herein, references to the selection and notification sub-processor and the feedback processor, or the action implementation sub-processor and the feedback processor, may each be considered interchangeable; it will be appreciated that whilst these sub-processors may be complementary hardware to the feedback processor, and/or effectively share a role of the feedback processor, they may equally be functions of the feedback processor operating under suitable software instruction. Meanwhile as noted above, optionally at least the action implementation sub-processor may be a separate processor to the feedback processor, for example communicating with the feedback processor via the Internet.

Selection and notification

Optionally, the selection notification sub-processor may select one or more feedback actions generated by the estimation processor in a manner as described previously herein, if the estimation processor indicates more than one feedback action may be appropriate. Clearly, if only one feedback action is proposed, then as a default this would be selected.

For a selected feedback action, the selection notification sub-processor may then select which device or devices within the delivery ecosystem should implement the feedback action, and formulates a command/notification/instruction for the or each device characterising the type and/or amount of the feedback action. It will be appreciated that where a device is only capable of one feedback action, then the type can be implicit in the act of notification, and similarly where a device is only capable of one amount feedback action, then the amount can be implicit in the act of notification. Any device within the delivery ecosystem could potentially comprise a feedback means. It will be appreciated therefore that potentially the device or devices that provide user factor data to or for the obtaining processor within the delivery ecosystem are different to the device or devices implementing the or each feedback action.

Optionally, the selection notification sub-processor may poll devices within the delivery ecosystem to determine their availability for the purposes of providing a feedback action. For devices that may be accessible by the processor, for example by the Internet, then devices registered in association with the user or the user's delivery device (such as for example the delivery device 10, mobile phone 100, wearable device 400, docking device 200) may be polled directly.

For devices that may only be accessed via an intermediary device, for example via Bluetooth ® connection to an accessible device, the accessible device may be asked to poll such indirect devices. Hence for example the selection notification sub-processor may cause/request the user's mobile phone 100 to poll a delivery device 10, wearable device 400 or docking device 200, if these are only accessible via a local wired or wireless connection.

For devices that are not formally associated with the user or only intermittently associated with the user, such as a vending machine 300 or other point of sale system, the selection notification sub-processor may receive location data from a device within the delivery ecosystem associated with the user, such as their mobile phone 100 or delivery device 10, and compare this with a registered or reported location of the vending machine 300; if the locations are within a threshold distance of each

other, then the vending machine may be considered part of the delivery ecosystem whilst that condition holds true. Alternatively or in addition, the selection notification sub-processor may instruct an accessible device to either poll for any compatible vending machines, or broadcast a Bluetooth beacon identifying the accessible device, for example using a single-use ID so that the accessible device is identifiable without revealing details of the user or their associated devices; such an ID may comprise a component identifying the purpose of the ID to enable detection by the vending machine, followed by the single use component unique to the user or their associated device; a compatible vending machine in accordance with embodiments of the present invention may then optionally recognise the single use ID and relay it back to the selection notification sub-processor, thereby informing it that the user has accessible devices within local wireless range of the vending machine. It will be appreciated that whilst the above makes reference to a vending machine, this is an example for the purposes of explanation, and these techniques may apply to any device not formally associated with a user or only intermittently associated with them, such as a car or train, a WiFi® or Bluetooth ® hotspot in a shop, a smart TV, or the like.

Optionally, devices outside the user's own delivery ecosystem may be selected. For example, a delivery device and/or a phone or other device of a friend or family member associated with the user (for example following registration of these people by the user) may be used to inform that friend or family member of the user's status, so that the friend or family member can intervene. Optionally the user can set the conditions under which this occurs, and/or which friends or family members are notified. Similarly, devices within a predetermined proximity of the user may be selected. For example, if the user is in a good mood, compatible devices within a predetermined radius of the user may all synchronise a feature such as a colour of a light, to signal to these users that there is scope for an enjoyable social encounter.

Using one or more of these techniques, the selection notification sub-processor may thus determine what devices are currently available to deliver a feedback action.

Typically, a feedback action will be specific to a particular device within the delivery ecosystem or a pair of devices cooperating to fulfil a function; consequently, for a proposed feedback action or selected feedback action, optionally the selection notification sub-processor may only poll the device or devices within the delivery ecosystem that are relevant to that feedback action.

More generally however, a feedback action may be specific to a particular capability required to deliver that feedback action; hence for example a feedback action comprising a message prompting the user to perform a certain action such as breathing more slowly/calmly in between uses of the delivery device or inhaling more slowly/calmly during use of the delivery device may be implemented on any device within the delivery ecosystem capable of displaying such a message; hence for example it may be provided by one or more of the delivery device itself, if it comprises a display; the user mobile phone, or a fitness wearable, or a suitably equipped docking unit of the delivery device. In principle such a message may similarly be provided to the user by a vending machine or other point of sale device.

Similarly, it will be appreciated that certain devices within the delivery ecosystem may provide input data to the feedback system that is used to generate the proposed feedback action; consequently such input activity from devices may be logged as indicating their accessibility, and/or it may be implicit from the proposed feedback action that certain devices are currently accessible to the feedback system; in either case, a poll of the devices may not be necessary, or the receipt of input data may be treated as an effective poll result if a polling scheme is in place.

In the event that the device or devices relevant to that feedback action are not available (e.g. do not respond to the poll), then optionally the feedback processor/selection notification sub-processor may choose the next proposed feedback action in the top N feedback actions, if multiple feedback actions were proposed by the estimation processor. If no relevant device is available for a feedback action, then the feedback processor may not implement any feedback action, and/or send a notification to the user to that effect, for example via a user interface of the user's phone, or if the user's phone, as the accessible device for linking to other devices within the ecosystem, is not available, then notifying the user via a text or similar other mechanism that will reach the user once they are contactable again. Similarly, if there is no effective communication currently available between the feedback processor and the relevant device or devices, or (depending on where the feedback processor is at least in part located), there is no effective communication between the feedback processor at a relevant device or devices and the estimation processor or other parts of the feedback system, then the relevant device or devices may default to a normal otherwise default delivery or other default behaviour suitable to that device.

In the event that the device or devices relevant to a feedback action are available (i.e. do respond to the poll, or have responded to a poll within a predetermined preceding period of time during which it can be assumed the device is still accessible, or have contributed input data within a predetermined preceding period of time), then the feedback processor will transmit one or more commands to the device or devices for implementing the feedback action as proposed by the estimation processor.

As noted above, the nature of the commands may depend upon the proposed action and the target device or devices. In some cases, the simple existence of the command will be sufficient to specify the proposed action, for example to turn a device on when it is off. In other cases, the command will need to specify the type of feedback action, for example in relation to changing heater function, payload type, user interface behaviour or the like within the delivery system. In either of these cases, the command may need to specify the amount of feedback action, for example to specify the change in temperature, the concentration of active ingredient or flavouring within the payload, or selected parameters for the user interface.

As noted above, the command may be directly to an accessible device, or may be to request that an accessible device relays a command to another device within the ecosystem, or itself issues a command to such a device; for example the feedback processor may instruct the user's mobile phone 100 to issue commands to the delivery device 10. Further degrees of indirection may be envisaged, such as the user's mobile phone issuing commands to a dock 200, which in turn may modify settings of the delivery device when it is docked (for example to charge power or payload). It will similarly be appreciated that the feedback processor may issue commands of different kinds to different devices; hence for example a command may be issued directly to the mobile phone to change aspects of its user interface, and to a dock 200 (either directly if it is capable, or via the phone), causing it to change a composition of a payload to be provided to the delivery device, and also change one or more settings on the delivery device when it is docked. It will be appreciated that other permutations of commands such as these, whether direct, indirect, or a mix of the two, can be envisaged within the delivery ecosystem.

As described elsewhere herein, will be appreciated that different feedback actions may relate to behavioural, pharmaceutical and/or non-consumption aspects of the delivery ecosystem.

Behavioural feedback actions are typically focused on altering actions and habits of the user relating to operations of, and/or interactions with, a device within the delivery ecosystem other than operations

relating to an amount or nature of an active ingredient delivered by the delivery device itself, although this can occur in parallel. Examples may relate to the use of changes in flavour or flavour concentration, changing vapour mass delivery to modify inhalation behaviour; modification to scheduling schemes or reminders relating to delivery device usage or correlated with delivery device usage; changes to user interfaces, whether on the delivery device on another device in the delivery ecosystem, in terms of information provided, mode of feedback (e.g. haptic and/or visible such as coloured lights, graphical themes, and/or messages); for example providing a traffic light UI display on the delivery device, such as an LED, to alert the user to how they are using the device), and the like.

Hence for example the selection of a flavour may involve choosing a flavour that encourages a behaviour complementary to the user's current state. Hence for example a peppermint flavour may be invigorating, if the user is tired, whereas a lavender flavour may be calling or soporific if the user is stressed. The relationship between flavours and user states may be determined empirically. The choice of flavour can also affect user behaviour based on how much the user likes the flavour, with less preferred and more preferred flavours reducing or increasing consumption. A change in flavour may also act as a prompt to the user to change behaviour in a previously decided manner; for example, different flavours may be marketed with imagery corresponding to different moods, behaviours, or user states, so that when the feedback processor causes selection of a particular flavour, the user is prompted according to the associated marketing/imagery.

Switching between flavours may be provided for example by using gel patches of respective flavours, and selectively heating the appropriate patch, or similar selective heating or supply of alternative flavourings within the aerosol generation process; other techniques may comprise use of multiple reservoirs for liquid flavourings, with selective supply, and the like.

Flavour concentration may similarly modify a user's behaviour. For example, disabling flavour entirely may cause the user to reduce consumption; meanwhile patterning flavour concentration (for example over a one-hour period, or a 20 minute period, or a usage session demarcated by a previous lack of use for a predetermined period of time, start with a high flavour concentration and progressively reduce it down, so that the user gets an initial sense of intervention from the delivery device, but also a sense of diminishing returns, encouraging more rapid cessation of use within the period/session. More generally, a user will associate a stronger flavour with a stronger placebo effect; hence where the action of using the delivery device is part of the modification of the user state, then a stronger flavour may enhance the effectiveness of this action. Hence optionally flavour concentration may be used as a modifier for other feedback actions, including in particular pharmaceutical feedback actions, as discussed elsewhere herein.

Changing vapour mass delivery, independent of active ingredient delivery, has a similar effect to changing flavour, in that it gives the user the impression of a greater or lesser amount of aerosol/vapour being inhaled; increasing vapour mass delivery gives the user the impression that they have inhaled more active ingredient than they actually have, and conversely decreasing vapour mass delivery gives the user the impression that they have inhaled less.

Hence for example the user can be encouraged to decrease usage by increasing vapour mass delivery.

Changes in usage frequency or patterns can alternatively or in addition be modified by direct changes to scheduling or reminding by the delivery device or any other device in the delivery ecosystem, such as for example the docking unit or the user's mobile phone.

Other forms of feedback may be provided by the device within the delivery ecosystem, for example by changing a colour scheme of a user interface component, whether this is a single LED, a full display, or anything between, or indeed through any other user interface medium such as haptics or audio. Hence for example a traffic light scheme could be used by a single LED to prompt the user to change their behaviour for example in a pre-defined manner; for example an LED may progress from green through amber to red (or directly from green to red) in a feedback action responsive to user factors such as physiological signs of stress such as heart rate, breathing rate, galvanic skin response and the like as discussed elsewhere herein, and/or responsive to other indicators of stress such as keywords in social media or text posts by the user, or circumstances indicated by their calendar, such as a particularly stressful location.

The more capable the user interface, the more detailed and/or tailored to the individual user the feedback can be. Hence if a device within the delivery ecosystem comprises a text capable display, then specific messages can be provided to the user. As noted previously herein, examples may include prompting the user to perform a certain action such slow or more calm breathing in between uses of the delivery device, and/or slow or more calm inhalation during use of the delivery device, or advice to take longer gaps between inhalations orientation sessions, or to change one or more settings on the device (particularly if these cannot be done automatically).

In addition to advising the user on how to modify their use of the delivery device, or any other device in the delivery ecosystem, such feedback actions can advise the user to modify their behaviour more generally; for example recommending that the user takes time out from a stressful situation, or performs an invigorating exercise, or conversely a meditative activity such as yoga. Such recommendations may be selected according to previously received user preferences; hence for example yoga may not be suggested to someone who does not already attend yoga classes.

The advice or prompt may be related to those physiological or situational factors that are likely contributing to the user's current state; hence for example if the user has physical signs of stress and also the background environment has been to detected as being noisy, the advice may be to wear headphones and listen to relaxing music; hence the advice need not be directly related to the use of the delivery device or consumption of materials through it. Hence more generally the wording of any message provided to the user may be modified according to one or more user factors.

Hence optionally a device within the delivery ecosystem may prompt the user to use either another device within the delivery ecosystem, optionally in a particular way, or to use some other device which may or may not be related to vaping or equivalent activities.

Hence a device within the delivery ecosystem may prompt a user to use a certain product suited to a user's current state; for example the device may recommend the user switches to a snus pouch instead of using an e-cigarette; this may occur for example where the user appears stressed but is in an environment where use of an e-cigarette is the delivery device may not be possible (for example if the user appears to be indoors, and the user's calendar suggests they are at a restaurant).

It will be appreciated that this may also interact with other forms of feedback action described above; for example a user may have two separate delivery devices providing separate flavours (or independently of the above separate active ingredients or active ingredient concentrations), and a device within the delivery ecosystem advises the user on which is currently the best to use, based on the feedback action identified in response to the currently available user factors relating to the user.

More generally, the feedback action may provide a prompt that is either complementary to the user's current state where this is currently positive, thereby enhancing it, or intended to revert the user's current state to a better one, where it is currently negative. Hence it will be understood that where a feedback action is expected to alter a state of the user, in a restorative situation this may involve an action that changes the user's negative state to a new state, but conversely in a complimentary or supportive situation this may involve an action that maintains the user's current positive state when it may otherwise change adversely.

After a feedback action has occurred, such a user interface may similarly be used to provide additional feedback to the user after the user state altering action has been taken; this additional feedback may provide positive reinforcement of the intended user state after the action, or prompt the user to measure the effectiveness of the action on their state (for example for the purposes of training the feedback system, and/or self-evaluation to recognise/appreciate the effect of the feedback action).

Next, pharmaceutical feedback actions focus on pharmaceutical interventions to change the state of the user, and typically relate to interventions based on active ingredients, such as the amount or type, and when these are changed (for example reactively or pre-emptively, for example based upon correlations between current user factors and future user states or feedback actions), and the like. Such acts can also relate to selecting alternative modes of consumption, for example switching from vaping to snus or vice versa.

Consequently, the estimation processor may identify at least a first feedback action based upon one or more user factors as explained elsewhere herein, wherein the at least first feedback action relates to an amount or nature of an active ingredient delivered by the delivery device.

With regards to the amount, the identified feedback action may comprise a binary decision to not supply any amount of active ingredient (for example switch to a placebo output); this may occur if an active ingredient is known to have a certain physiological effect that may be considered adverse given the user's current physiological state as inferred from the available user factors. Hence for example an active ingredient that is likely to increase heart rate may be stopped if the user is detected to have elevated heart rate. For at least a short period of time, the placebo effect is likely to work as the user will still be consuming from the delivery device in the belief that they are receiving the active ingredient (or, in the case where consent has been sought, nonetheless engaging in an action with previous positive connotations).

It will be appreciated that stopping the provision of the active ingredient in this manner may be enacted either by stopping the delivery function of the delivery device altogether, or merely stopping the inclusion of the active ingredient within a delivery medium supplied by the delivery device. In the former case, a message may be provided to the user indicating that the feedback action has occurred, so that the user does not believe the delivery device is malfunctioning.

Alternatively to a binary decision, the identified feedback action may modify the concentration and/or amount of active ingredient provided to the user. For example, the concentration of active ingredient may increase or as appropriate decrease from a default level according to the estimation processor output, such as the extent of activation of a feedback action, or of an output value associated with the feedback action, as described elsewhere herein. Similarly the amount of active ingredient (for a given concentration) may be increased or decreased for example by modifying aerosol generation parameters

(such as heating temperature) to adjust the amount of aerosol generated per unit airflow through the delivery device.

Hence the feedback processor may respond to and/or with a quantitative value indicating the degree of change, or class of change, as appropriate to the identified feedback action.

It will be appreciated that the feedback action may modify the amount of active ingredient whether from none to a predetermined value, or along a sliding scale, for a single subsequent inhalation, or for inhalations taken during a following period of time, or equivalently for a predetermined number of inhalations, or equivalently for a predetermined sum volume of inhalation for example as estimated from delivery device airflow and inhalation duration measurements.

Alternatively or in addition to such modifications, the feedback system may also manage the distribution of delivery of an active ingredient over a period of time, for example to provide acute or chronic delivery regimes relatively independent of the user's pattern of consumption, providing the user consumes from the delivery device relatively frequently compared to the period of time. Hence for example if a given period of time typically encompasses 20 inhalation actions, then a feedback action may deliver the same overall amount of active ingredient for example in an average manner similar overall inhalations, or concentrate some, most or all of the active ingredient into a small number of inhalations within that 20, so as to provide either a chronic or acute delivery regime.

Such modifications to the delivery regime may be reactive, so that if the user appears particularly stressed, and acute delivery regime may be implemented, followed by a return to chronic delivery regime that may optionally be at a lower concentration and/or amount responsive to the concentration and/or amount during the acute phase. Conversely, such modifications to the delivery regime may be predictive, so that if the user shows mounting stress levels, or circumstantial or other user factors indicate an imminent stressful situation (for example about to start a driving test), then the feedback action may comprise an acute delivery regime to forestall the increase of stress, for example by providing an acute delivery regime in response to mounting stress, or immediately preceding the expected stressful situation. This may then optionally be followed by a correspondingly lower chronic delivery regime once user factors indicate that stress levels have dropped, and/or circumstantial user factors suggest the stressful situation is over.

The above discussions of changes to the amount of active ingredient, either on a per inhalation basis or over a predetermined period of time, assume a single active ingredient is available either on the same delivery device or a plurality delivery devices. However, two or more active ingredients may be available, and a feedback action may comprise switching from one active ingredient to another, or mixing or modifying a mix of two or more active ingredients.

Active ingredients can comprise any composition having a physiological effect on the user, for example changing heart rate, dopamine and/or cortisol levels, and/or having an effect on brain chemistry and/or subjective user experience, as described elsewhere herein.

The most common active ingredient is nicotine, but any suitable active ingredient may be considered.

Hence a feedback action where a delivery device already delivers active ingredient X to the user, it now introduces additional active ingredient Y during use. Meanwhile concentrations of active ingredient X may remain the same or may increase or decrease as appropriate, depending for example on whether a desired change to user state benefits from a transition from X to Y or the complementarity of X and Y.

Alternatively or in addition to the introduction of a second active ingredient either as a complementary composition or as part of a transition from one to the other, a feedback action may comprise simply switching from one active ingredient to another. This may be done within a single delivery device, for example by selective heating of a gel or other carrier medium for a respective active ingredient, or by recommending to the user that they change cartridge rather consumable with the same delivery device, or switch to another delivery device if they own more than one. In this latter case, typically the user will have registered their delivery devices with the feedback system, and optionally the delivery devices will have reported the type of payload they currently carry to the feedback system as a user factor.

A feedback action comprising switching from one active ingredient another may also include switching from one form of an active ingredient to a different form of an active ingredient.

For example, a feedback action may comprise selecting between providing protonated nicotine or not as the active ingredient, or adjusting the percentage of protonated nicotine delivered within an overall mix of nicotine or another active ingredient being supplied collectively as the active ingredient.

Protonated nicotine is absorbed more quickly by the lungs than non-protonated nicotine; this may be advantageous if for example the user factors indicate a sudden onset of stress.

The delivery device may comprise respective reservoirs, gels, or other delivery mechanisms for protonated and non-protonated nicotine, or may comprise means to protonated nicotine on demand.

Hence more generally a feedback action may comprise selecting a stronger version of the same active ingredient, or a more effective version of the same active ingredient from a pharmacokinetic / user response perspective.

It will be appreciated that where the strength, effectiveness, concentration, and/or amount of active ingredient can be altered, either by adjusting a mix of active and inactive aerosol ingredient, or a mix of two active ingredient within the aerosol, or substituting one active ingredient with another (whether a different active ingredient, or a different version of the same active ingredient), such alteration by way of a feedback action may be either reactive or proactive. For non-aerosol delivery devices, the adjustment of concentration may be by similarly mixing where appropriate, or by selecting different prepared blends. Meanwhile the adjustment of amount may be achieved by providing different amounts of an item, and/or different sizes of an item, as appropriate.

A reactive feedback action may for example be identified when user factors indicate physiological stress, and/or the presence of circumstantial and/or environmental stressors.

Meanwhile a proactive or anticipatory feedback action may for example be identified prior to the presence of the circumstantial and/or environmental stressor, for example based upon information in a user's calendar, texts and/or social media posts indicative of the stressful event such as a visit to the dentist or a driving test, or comments indicative of anticipated stress made by the user in texts or social media posts; similarly user factors relating to the user's location, or who they are with may become associated with elevated stress levels; for example the feedback system may learn a correspondence between physiological stress levels and other aspects of the user's situation such as their current circumstantial or environmental situation, so that if the user appears to be travelling towards a certain location, or is started to be surrounded by certain people or enter into some other circumstance which previously has been associated with high physiological stress, a proactive feedback action can alter the active ingredient as described above to be more effective against stress.

Hence for example providing an increased amount of nicotine and/or protonated nicotine within a given volume of aerosol prior to a user encountering a stressful situation will reduce the impact of the stressor on the user. Optionally the feedback system can use a pharmacokinetic model to determine whether a given puff preceding the anticipated stressful circumstance should be modified in this manner, so that the adjusted active ingredient (for example nicotine) has the desired effect at the anticipated time that the stressful event occurs.

It will be appreciated that whilst the above examples relate to mitigating stress, a similar approach may be used to promote or sustain positive states in the user.

Finally, non-consumption feedback actions typically relate to activating/controlling or simply recommending the use of devices not specifically related to the consumption of the active ingredient, such as aromatherapy systems/steamers, biofeedback devices, headphones (for example activating noise cancellation, or modifying volume or music selection), vehicle use (for example stress warnings, or route selection / reselection to longer but less congested or slower routes), and the like.

The selection notification sub-processor may be comprised of one or more real or virtual processors, and its functionality may be located or distributed within the server and/or one or more devices within the delivery ecosystem as appropriate.

Action implementation

The action implementation sub-processor may be optional; for example some devices may accept commands directly with no further interpretation or processing required. In this case the action implementation sub-processor may be either thought of as not required, or having its role implemented by the feedback processor / selection and notification sub-processor.

Meanwhile, in some cases the role of the action implementation sub-processor may in fact pre-exist within the device, which for example is capable of interpreting user interface commands to implement changes to the operation of the device; in this case, the commands from the feedback processor may optionally simply replicate such user interface commands.

In other cases, the action implementation sub-processor may be separately provided, for example by adapting a conventional processor according to suitable software instruction. Such an example may be an app on a user's mobile phone, operable to receive commands and modify one or more of aspects of the mobile phone and/or the app on the mobile phone, the delivery device, and/or one or more other devices in the delivery ecosystem. Similarly, a dock 200 for the delivery device may comprise such an action implementation sub-processor, as may some varieties of delivery device.

The action implementation sub-processor operates to implement the feedback action on the or each relevant device. Hence for example if a command relating to a feedback action describes changing heater temperature of the delivery device, then the action implementation sub-processor may change the power supply to the heater, and/or a duty cycle of the heater, to implement the specified change.

Similarly for example if a command relating to a feedback action describes reducing ambient noise levels for the user, then the action implementation sub-processor for a pair of noise cancelling headphones may activate the noise cancelling function; the action implementation sub-processor for the user's mobile phone may reduce the volume level of music being played into those headphones, and display a message to the user suggesting that they seek to avoid sources of noise in their environment.

The specific actions implemented by respective sub-processors may thus depend on the nature of the proposed feedback action and the nature of the device within the delivery ecosystem, but will typically represent a direct translation of the proposed feedback action into the mechanism by which it may be enacted within device.

As noted previously herein, feedback actions may be accompanied or followed up by requests or opportunities for the user to report on their efficacy. Alternatively or in addition, feedback actions may be accompanied or followed up by positive reinforcement of the expected state change, for example through a message on a UI, or a change of interface colour, haptic response or the like, or for example a positive goal being met in an app for a wearable. The reinforcement may be a simple message indicating that the feedback has occurred, or may be based upon measurements, for example to report that the user's heart rate has lowered, or to confirm that an action has worked well (by changing the user's state, typically as evidenced by a change to one or more user factors, or as self-reported by the user). The perception and/or expectation of a change in state engendered by such positive reinforcement can increase the effectiveness of at least some feedback actions.

The action implementation sub-processor may be comprised of one or more real or virtual processors, and its functionality may be located or distributed within the server and/or one or more devices within the delivery ecosystem as appropriate.

The autonomy of the action implementation sub-processor (and more generally the feedback processor, and/or feedback system) may be set globally, or vary according to the type of feedback action, or according to individual feedback actions. Here the autonomy means whether to what extent the action implementation sub- processor may proceed to implement a feedback action without notifying the user or a requesting their consent, whether as an initial permission, or each time a feedback action is performed.

Consequently one option is that the relevant device in the delivery ecosystem is arranged to automatically implement its part of a feedback action, for example by automatically selecting a flavour or adjusting a flavour concentration, automatically adjusting a vapour mass delivery rate (e.g. an amount of flavour), automatically providing feedback or text messages, or the like.

As a result, the feedback system automatically implements a feedback action expected to change a state of the user.

Such automatic adjustment may be applied globally, for example set at manufacturing for all feedback actions. Alternatively, such automatic adjustment may be applied only for certain feedback actions (for example feedback actions that are considered unlikely to prompt a refusal by the user) and/or such automatic adjustment may be applied only for certain user states (for example where it is considered likely that an automatic adjustment may be positively received). Alternatively, such automatic adjustment may be selected by the user, for example in an initial setup phase so that the user sets are initial preferences and does not need to be prompted again. Optionally in this case the user can revisit their preferences to change whether or not a particular feedback action is applied automatically.

Alternatively, an option is that the relevant device in the delivery ecosystem is arranged to prompt the user before taking any action that may adjust or influence the user state, and hence give the user control over whether the device implements its part of a feedback action.

In this case, depending on the device's user interface capabilities, the prompt may be a text or spoken prompt, or a haptic prompt, and audio prompt, or the activation of an LED or selection of a specific LED colour. Subsequently the user's response (at its simplest a yes/no response, or a yes by inaction or alternatively a no by inaction response) may similarly be determined by the device's user interface capabilities; for example on a touchscreen a user may an icon indicating permission or refusal, or may press a button indicative of permission or refusal. It will also be appreciated that for a predetermined period of time after a prompt has been given to the user, one or more buttons may be repurposed for the provision of consent; for example '+' and '-' buttons, used for example to change heater temperature or sound volume or any other the device may be temporarily repurposed so that '+' means consent and '-' means refusal. It will be appreciated that any suitable button may be repurposed in this way.

As with the case of automatic feedback action, prompts may be set to applied globally, or depending on the feedback action and/or depending on the user state to be modified from or to.

It will also be appreciated that the prompt may relate to the type of feedback action, or the type of change in user state anticipated as a consequence of the feedback action, or any mixture of the two. Hence for example the prompt may suggest to the user that they are in a certain mood, or have an elevated heart rate, or any other state discussed elsewhere herein, and ask if they want to change an aspect of the delivery process, or change their mood, or change their heart rate, as appropriate. Similarly for example a prompt may suggest that a certain feedback action will result in a different user state, or the maintenance of the current user state where it might otherwise change.

Hence alternatively or in addition to requesting consent for a particular feedback action, a prompt may offer the selection of a feedback action to the user; as described elsewhere herein, the feedback processor may automatically select from among a plurality of identified feedback actions, but alternatively this function may be adapted to involve the user. Optionally the feedback processor may preselect or shortlist identified feedback options, for example based upon currently available devices in the delivery ecosystem or elsewhere that can fulfil the identified feedback actions or their respective parts of these, but may then give the final choice to the user. Similarly the feedback processor may preselect or shortlist identified feedback options based on their frequency of use and/or selection either by the individual user or among a cohort of users, and/or the effectiveness of a feedback action, as reported by either the individual user or a cohort of users.

Typically as explained elsewhere herein the identified feedback actions have been identified in response to some or all of the user factors obtained by the user feedback system, and so are likely to be directed toward achieving similar effects. However they may do so in different ways, some of which are more preferable to the user than others. Accordingly user may be asked to choose one (or more) suggested feedback actions from among those identified. The feedback system may optionally modify any of these feedback actions, if implementing two or more would change the intended outcome, and/or similarly may dynamically grey out certain options if another incompatible option has been selected by the user.

Alternatively or in addition to asking the user to select amongst alternative feedback actions directed towards achieving similar effects, a device in the delivery ecosystem may propose different feedback actions relating to different user states that may be associated with the same user factors, or respective subsets of received user factors. Hence for example a user may be simultaneously calm, but subjectively feel either focused or lethargic; depending on the user factors available, these aspects of the user state may or may not be identified; subsequently if the user appears to be calm, then different feedback

actions relating to whether the user is focused, sleepy, or lethargic may be provided to the user so they can make their own choice. Hence for example a different flavour may be provided if the user is focused to when the user is sleepy, and/or a different heating profile may be used to modify the delivery of flavour and or ingredients within the inhalation process.

It will be appreciated that optionally where a user feedback system is initially based on average or cohort data for user behaviours, but is capable of learning about an individual user, then such a system may present more feedback actions to a user during initial and early use, but subsequently learn which feedback actions the user prefers and/or responds best to, and so reduce the number of choices. Hence also optionally, once a clear choice has been determined for the given user factors, the user feedback system may either only request confirmation from the user to proceed with that feedback action, or implement the feedback action automatically, as discussed elsewhere herein.

As a further alternative to either automatically implementing the identified feedback action or actions, or requesting permission to implement an identified feedback action or actions, or choosing from among suggested identified feedback actions, optionally a prompt may comprise instructions on how the user can implement the identified or actions feedback action themselves, for example a prompt to manually change the settings of a device within the delivery ecosystem, for example to increase heater temperature on the delivery device.

It will be appreciated that in any event optionally a prompt may be delivered on a device with a more capable user interface than the device upon which the feedback action is implemented.

Where consent or refusal is indicated, either explicitly or by inaction or by selection depending on the user interface chosen, this can be used to train the feedback system to better determine when to select feedback actions in future.

Processors

As noted previously, the obtaining processor, estimation processor, and feedback processor (and any sub processors) may comprise one or more real or virtual processors located within one or more servers and/or within the delivery ecosystem. Furthermore it will be appreciated that the demarcation of roles described herein is not fixed; for example the obtaining processor may receive information directly indicative of user state (for example by user self-reporting), and so the first step of a two step process by the estimation processor could be bypassed or supplemented by the obtaining processor; similarly in this case, the feedback processor may, for example, look up a corresponding proposed feedback action. Hence in this example, the role of the estimation processor is carried out by the obtaining processor and the feedback processor. Hence more generally these processors are representative of tasks that may be implemented by any processor under suitable software instruction, and can equivalently be considered to comprise a data gathering task, a feedback proposal task (whether or not based on an explicit estimation of the state of the user), and either a feedback training task or a feedback delivery task.

Summary Embodiments

In a summary embodiment of the present description, a user feedback system for a user of a delivery device (10) within a delivery ecosystem (1), comprises the following features. Firstly, an estimation processor (1020) is adapted to identify at least a first feedback action based upon one or more user factors, as described elsewhere herein. The at least first feedback action relates to an amount or nature of an active ingredient delivered by the delivery device, as described elsewhere herein, and is expected

to alter a state of the user (for example to modify a negative state, or maintain a positive state that might otherwise change) as indicated at least in part by the one or more user factors, as described elsewhere herein. The system also comprises a feedback processor (1030) adapted to select at least a first identified feedback action, and to cause a modification of one or more operations of at least a first device within the delivery ecosystem, according to the or each selected feedback action, as described elsewhere herein.

In an instance of the summary embodiment, the feedback action comprises adjusting a concentration and/or amount of the active ingredient provided by the delivery device, as described elsewhere herein.

In this instance, optionally the adjustment of the concentration and/or amount of the active ingredient is responsive to an output value generated by the estimation processor, in turn responsive to one or more user factors, as described elsewhere herein; for example the estimation processor may generate a dedicated output value, or the alpha value may be derived from the strength of activation of an identified feedback action, or from any other value having a correlation or correspondence with a beneficial degree of adjustment.

Similarly in this instance, optionally the feedback action comprises adjusting the concentration and/or amount of the active ingredient provided by the delivery device to zero, as described elsewhere herein. This may be performed by using a means capable of variable adjustment of the concentration and/or amount, or may be performed by disabling delivery of the active ingredient, for example by preventing its release from a reservoir, or where normally selectively heated, not selectively heating it whilst still heating other materials that contribute to the aerosol.

Similarly in this instance, optionally the concentration and/or amount of active ingredient is adjusted over a sequence of inhalation actions, as described elsewhere herein.

In an instance of the summary embodiment, the feedback action comprises stopping the delivery operation of the delivery device, as described elsewhere herein; for example rather than merely disabling delivery of the active ingredient, delivery of both the active ingredient and other materials contributing to the aerosol (or other mode of delivery or delivery medium), are prevented from being delivered to the user.

In an instance of the summary embodiment, the feedback action comprises changing the ratio of a mix of a first active ingredient and a second active ingredient, as described elsewhere herein.

In this instance, optionally the feedback action comprises switching from the first active ingredient to the second active ingredient, as described elsewhere herein; in this case will be appreciated that the mix ratio becomes 0% : 100%, or vice versa.

Similarly in this instance, optionally the first active ingredient is nicotine, and the second active ingredient is protonated nicotine, as described elsewhere herein, although it will be appreciated that any suitable combination of active ingredients may be envisaged.

In an instance of the summary embodiment, the feedback action is performed in anticipation of the user's state, as described elsewhere herein; hence a feedback action may be proactive or anticipatory.

In an instance of the summary embodiment, the feedback action comprises modifying an amount or nature of an active ingredient for a predetermined period of time, as described elsewhere herein.

In this instance, optionally the predetermined period of time is derived from one selected from the list consisting of a predetermined period of time, a predetermined number of inhalations, and a predetermined sum volume of inhalation.

In an instance of the summary embodiment, the selected feedback action also comprises causing the prompting of the user to provide feedback to the user feedback system in relation to the selected feedback action, as described elsewhere herein.

In an instance of the summary embodiment, the estimation processor is operable to identify one or more further proposed feedback actions relating to one or more selected from the list consisting of a behavioural feedback action for affecting at least a first behaviour of the use, and a non-consumption feedback action for affecting one or more non-consumption operations of the delivery ecosystem, as described elsewhere herein.

Instance of the summary embodiment, the feedback processor (1030) is adapted to select the at least first identified feedback action responsive to the current availability of respective devices for implementing feedback actions within the delivery ecosystem, as described elsewhere herein.

In an instance of the summary embodiment, the feedback processor is adapted to cause implementation of the at least first identified feedback action automatically, as described elsewhere herein.

In an instance of the summary embodiment, alternatively the feedback processor is adapted to prompt the user for consent to cause implementation of at least part of the at least first identified feedback action, and to only cause implementation of the at least part of the at least first identified feedback action, if consent is determined, as described elsewhere herein.

In an instance of the summary embodiment, the feedback processor is adapted to provide the user with a choice of identified feedback actions to select from, as described elsewhere herein.

In an instance of the summary embodiment, the feedback processor is adapted to provide the user with details of how to implement a selected identified feedback action themselves, as described elsewhere herein.

In an instance of the summary embodiment, the user feedback system comprises an obtaining processor (1010) adapted to obtain one or more user factors indicative of a state of the user, for use by the estimation processor, as described elsewhere herein.

In this instance, a respective one of the one or more user factors is based upon one selected from the list consisting of at least a first physical property associated with at least a first user inhalation action, at least a first physical property associated with user behaviour other than inhalation, at least a first physical property associated with user physiology other than in relation to inhalation, and at least a first aspect of the user's situation separate to their handling or operation of the delivery device, as described elsewhere herein.

In an instance of the summary embodiment, one or more user factors respectively relate to at least one class selected from the list consisting of historical data providing background information relating to the user, neurological data relating to the user, physiological data relating to the user, contextual data

relating to the user, environmental data relating to the user, deterministic data relating to the user, and usage data relating to usage of the delivery device by the user, as described elsewhere herein.

In an instance of the summary embodiment, the estimation processor does not generate an explicit estimation of user state as an interim step in the identification of the one or more proposed feedback actions, as described elsewhere herein.

In an instance of the summary embodiment, the delivery ecosystem comprises one or more selected from the list consisting of one or more delivery devices (10), one or more mobile terminals (100), one or more wearable devices (400), and one or more docking units (200) for the or each delivery device, as described elsewhere herein.

In an instance of the summary embodiment, functionality of one or more of the obtaining processor, estimation processor, and a feedback processor is provided at least in part by a remote server (1000), as described elsewhere herein.

In an instance of the summary embodiment, functionality of one or more of the obtaining processor, estimation processor, and a feedback processor is provided at least in part by one or more processors located within one or more devices (10, 100, 200, 300, 400) of the delivery ecosystem (1), as described elsewhere herein.

Turning now to Figure 7, in a summary embodiment of the present description, a user feedback method for a user of a delivery device (10) within a delivery ecosystem (1), comprises the following steps.

A first estimating step s710 comprises identifying at least a first feedback action based upon one or more user factors as described elsewhere herein. The at least first feedback action relates to an amount or nature of an active ingredient delivered by the delivery device, and the feedback action is expected to alter a state of the user as indicated at least in part by the one or more user factors, also as described elsewhere herein.

A second feedback step s720 then comprises the respective steps of selecting (s722) at least a first identified feedback action, and causing (s724) a modification of one or more operations of at least a first device within the delivery ecosystem according to the or each selected feedback action, as described elsewhere herein.

It will be apparent to a person skilled in the art that variations in the above method corresponding to operation of the various embodiments of the method and/or apparatus as described and claimed herein are considered within the scope of the present disclosure, including but not limited to that:

- the feedback action comprises adjusting a concentration and/or amount of the active ingredient provided by the delivery device, as described elsewhere herein;
    - in this instance, optionally the adjustment of the concentration and/or amount of the active ingredient is responsive to an output value generated by the estimation processor, in turn responsive to one or more user factors, as described elsewhere herein;
    - similarly in this instance, optionally the feedback action comprises adjusting the concentration and/or amount of the active ingredient provided by the delivery device to zero, as described elsewhere herein;

    - similarly in this instance, optionally the concentration and/or amount of active ingredient is adjusted over a sequence of inhalation actions, as described elsewhere herein;
- stopping the delivery operation of the delivery device, as described elsewhere herein;
- changing the ratio of a mix of a first active ingredient and a second active ingredient, as described elsewhere herein;
    - in this instance, optionally switching from the first active ingredient to the second active ingredient, as described elsewhere herein;
    - similarly in this instance, optionally the first active ingredient is nicotine, and the second active ingredient is protonated nicotine, as described elsewhere herein
- performing the feedback action in anticipation of the user's state, as described elsewhere herein;
- modifying an amount or nature of an active ingredient for a predetermined period of time;
    - in this instance, optionally the predetermined period of time is derived from one selected from the list consisting of a predetermined period of time, a predetermined number of inhalations, and a predetermined sum volume of inhalation;
- the selected feedback action comprises also causing the prompting of the user to provide feedback to the user feedback system in relation to the selected feedback action, as described elsewhere herein;
- the estimation step comprises identifying one or more further proposed feedback actions relating to one or more selected from the list consisting of a pharmaceutical feedback action for affecting the consumption of an active ingredient by the user, and a non-consumption feedback action for affecting one or more non-consumption operations of the delivery ecosystem, as described elsewhere herein;
- the feedback step comprises selecting the at least first identified feedback action responsive to the current availability of respective devices for implementing feedback actions within the delivery ecosystem, as described elsewhere herein;
- the feedback step comprises causing implementation of the at least first identified feedback action automatically, as described elsewhere herein;
- the feedback step alternatively prompting the user for consent to cause implementation of at least part of the at least first identified feedback action, and only causing implementation of the at least part of the at least first identified feedback action if consent is determined, as described elsewhere herein;
- the feedback step comprises providing the user with a choice of identified feedback actions to select from, as described elsewhere herein;
- the feedback step comprises providing the user with details of how to implement a selected identified feedback action themselves, as described elsewhere herein;
- an obtaining step comprises obtaining one or more user factors indicative of a state of the user, for use by the estimation processor, as described elsewhere herein;
    - In this instance, a respective one of the one or more user factors being based upon one selected from the list consists of at least a first physical property associated with at least a first user inhalation action, at least a first physical property associated with user behaviour other than inhalation, at least a first physical property associated with user physiology other than in relation to inhalation, and at least a first aspect of the user's situation separate to their handling or operation of the delivery device, as described elsewhere herein;

- one or more user factors respectively relate to at least one class selected from the list consisting of historical data providing background information relating to the user, neurological data relating to the user, physiological data relating to the user, contextual data relating to the user, environmental data relating to the user, deterministic data relating to the user, and usage data relating to usage of the delivery device by the user, as described elsewhere herein;
- the estimation step does not comprise the generation of an explicit estimation of user state as an interim step in the identification of the one or more proposed feedback actions, as described elsewhere herein;
- the estimation step comprises identifying one or more proposed feedback actions relating to a behavioural feedback action for affecting at least a first behaviour of the user, as described elsewhere herein;
- the delivery ecosystem comprises one or more selected from the list consisting of one or more delivery devices (10), one or more mobile terminals (100), one or more wearable devices (400), and one or more docking units (200) for the or each delivery device, as described elsewhere herein;
- functionality of one or more of the obtaining step, estimation step, and a feedback step is provided at least in part by a remote server (1000), as described elsewhere herein; and
- functionality of one or more of the obtaining step, estimation step, and a feedback step is provided at least in part by one or more processors located within one or more devices (10, 100, 200, 300, 400) of the delivery ecosystem (1), as described elsewhere herein.

It will be appreciated that the above methods may be carried out on conventional hardware (such as the obtaining processor, estimation processor feedback processor, on the server and/or any one of the devices of the delivery ecosystem) suitably adapted as applicable by software instruction or by the inclusion or substitution of dedicated hardware.

Thus the required adaptation to existing parts of a conventional equivalent device may be implemented in the form of a computer program product comprising processor implementable instructions stored on a non-transitory machine-readable medium such as a floppy disk, optical disk, hard disk, solid state disk, PROM, RAM, flash memory or any combination of these or other storage media, or realised in hardware as an ASIC (application specific integrated circuit) or an FPGA (field programmable gate array) or other configurable circuit suitable to use in adapting the conventional equivalent device. Separately, such a computer program may be transmitted via data signals on a network such as an Ethernet, a wireless network, the Internet, or any combination of these or other networks.

CLAIMS

1. A user feedback system for a user of a delivery device within a delivery ecosystem, comprising:

an estimation processor adapted to identify at least a first feedback action based upon one or more user factors,

the at least first feedback action relating to an amount or nature of an active ingredient delivered by the delivery device,

the feedback action being expected to alter a state of the user as indicated at least in part by the one or more user factors; and

a feedback processor adapted to select at least a first identified feedback action, and to cause a modification of one or more operations of at least a first device within the delivery ecosystem, according to the or each selected feedback action.

2. A user feedback system according to claim 1, in which the feedback action comprises adjusting a concentration and/or amount of the active ingredient provided by the delivery device.

3. A user feedback system according to claim 2, in which the adjustment of the concentration and/or amount of the active ingredient is responsive to an output value generated by the estimation processor, in turn responsive to one or more user factors.

4. A user feedback system according to claim 2 or claim 3, in which the feedback action comprises adjusting a concentration and/or amount of the active ingredient provided by the delivery device to zero.

5. A user feedback system according to any one of claims 2 to 4, in which the concentration and/or amount of the active ingredient is adjusted over a sequence of inhalation actions.

6. A user feedback system according to claim 1, in which the feedback action comprises stopping the delivery operation of the delivery device.

7. A user feedback system according to any proceeding claim, in which the feedback action comprises changing the ratio of a mix of a first active ingredient and a second active ingredient.

8. A user feedback system according to claim 7, in which the feedback action comprises switching from the first active ingredient to the second active ingredient.

9. A user feedback system according to claim 7 or claim 8, in which the first active ingredient is nicotine, and the second active ingredient is protonated nicotine.

10. A user feedback system according to any one of the preceding claims, in which the feedback action is performed in anticipation of a user's state.

11. A user feedback system according to any proceeding claim, in which the feedback action comprises modifying an amount or nature of an active ingredient for a predetermined period of time.

12. A user feedback system according to claim 11, in which the predetermined period of time is derived from one selected from the list consisting of:

i. a predetermined period of time;

ii. a predetermined number of inhalations; and

iii. a predetermined sum volume of inhalation.

13. A user feedback system according to any proceeding claim, in which the selected feedback action also comprises causing the prompting of the user to provide feedback to the user feedback system in relation to the selected feedback action.

14. A user feedback system according to any proceeding claim, in which the estimation processor is operable to identify one or more further proposed feedback actions relating to one or more selected from the list consisting of:

i. a behavioural feedback action for affecting at least a first behaviour of the use; and
ii. a non-consumption feedback action for affecting one or more non-consumption operations of the delivery ecosystem.

15. A user feedback system according to any preceding claim in which the feedback processor (1030) is adapted to select the at least first identified feedback action responsive to the current availability of respective devices for implementing feedback actions within the delivery ecosystem.

16. A user feedback system according to any proceeding claim in which the feedback processor is adapted to cause implementation of the at least first identified feedback action automatically.

17. A user feedback system according to any one of claims 1-15 in which the feedback processor is adapted to prompt the user for consent to cause implementation of at least part of the at least first identified feedback action, and to only cause implementation of the at least part of the at least first identified feedback action, if consent is determined.

18. A user feedback system according to any one of claims 1-17 in which the feedback processor is adapted to provide the user with a choice of identified feedback actions to select from.

19. A user feedback system according to any one of claims 1-18, in which the feedback processor is adapted to provide the user with details of how to implement a selected identified feedback action themselves.

20. A user feedback system according to any preceding claim, comprising:
an obtaining processor adapted to obtain one or more user factors indicative of a state of the user, for use by the estimation processor.

21. A user feedback system according to claim 20, in which:
a respective one of the one or more user factors is based upon one selected from the list consisting of:

i. at least a first physical property associated with at least a first user inhalation action;
ii. at least a first physical property associated with user behaviour other than inhalation;
iii. at least a first physical property associated with user physiology other than in relation to inhalation; and
iv. at least a first aspect of the user's situation separate to their handling or operation of the delivery device.

22. A user feedback system according to any proceeding claim, in which one or more user factors respectively relate to at least one class selected from the list consisting of:

i. historical data providing background information relating to the user;

ii. neurological data relating to the user;
iii. physiological data relating to the user;
iv. contextual data relating to the user;
v. environmental data relating to the user;

23. The user feedback system according any proceeding claim, in which the estimation processor does not generate an explicit estimation of user state as an interim step in the identification of the one or more proposed feedback actions.

24. The user feedback system according to any proceeding claim, in which the delivery ecosystem comprises one or more selected from the list consisting of:
i. one or more delivery devices;
ii. one or more mobile terminals;
iii. one or more wearable devices; and
iv. one or more docking units for the or each delivery device.

25. The user feedback system according to any proceeding claim, in which functionality of one or more of the obtaining processor, estimation processor, and a feedback processor is provided at least in part by a remote server.

26. The user feedback system according to any proceeding claim, in which functionality of one or more of the obtaining processor, estimation processor, and a feedback processor is provided at least in part by one or more processors located within one or more devices of the delivery ecosystem.

27. A user feedback method for a user of a delivery device within a delivery ecosystem, comprising:
an estimating step of identifying at least a first feedback action based upon one or more user factors,
the at least first feedback action relating to an amount or nature of an active ingredient delivered by the delivery device,
the feedback action being expected to alter a state of the user as indicated at least in part by the one or more user factors; and
a feedback step, comprising
selecting at least a first identified feedback action, and
causing a modification of one or more operations of at least a first device within the delivery ecosystem according to the or each selected feedback action.

28. A user feedback system according to claim 27, in which the feedback action comprises adjusting a concentration and/or amount of the active ingredient provided by the delivery device.

29. A user feedback system according to claim 28, in which the adjustment of the concentration and/or amount of the active ingredient is responsive to an output value generated by the estimation processor, in turn responsive to one or more user factors.

30. A user feedback system according to claim 27 or claim 28, in which the feedback action comprises adjusting the concentration and/or amount of the active ingredient provided by the delivery device to zero.

31. A user feedback system according to one of claims 27-30, in which the feedback action comprises changing the ratio of a mix of a first active ingredient and a second active ingredient.

32. A user feedback system according to claim 31, in which the feedback action comprises switching from the first active ingredient to the second active ingredient.

33. A user feedback system according to claim 31 or claim 32, in which the first active ingredient is nicotine, and the second active ingredient is protonated nicotine.

34. A user feedback system according to any one of claims 27-33, in which the feedback action is performed in anticipation of a user's state.

35. A user feedback system according to any one of claims 27-34, in which the feedback action comprises modifying an amount or nature of an active ingredient for a predetermined period of time.

36. A user feedback system according to claim 35, in which the predetermined period of time is derived from one selected from the list consisting of:

i. a predetermined period of time;
ii. a predetermined number of inhalations; and
iii. a predetermined sum volume of inhalation.

37. A user feedback method according to any one of claims 27-36, in which the estimating step comprises identifying one or more further proposed feedback actions relating to one or more selected from the list consisting of:

i. a behavioural feedback action for affecting at least a first behaviour of the use; and
ii. a non-consumption feedback action for affecting one or more non-consumption operations of the delivery ecosystem.

38. A user feedback method according to any one of claims 27-37 in which the feedback processor is adapted to cause implementation of the at least first identified feedback action automatically.

39. A user feedback method according to any one of claims 27-38 in which the feedback processor is adapted to prompt the user for consent to cause implementation of at least part of the at least first identified feedback action, and to only cause implementation of the at least part of the at least first identified feedback action, if consent is determined.

40. A user feedback method according to any one of claims 27-39 in which the feedback processor is adapted to provide the user with a choice of identified feedback actions to select from.

41. A computer program comprising computer executable instructions adapted to cause a computer system to perform the method of any one of claims 27-40.

42. A computer program product comprising the computer program of claim 41 stored on a non-transitory machine-readable medium.

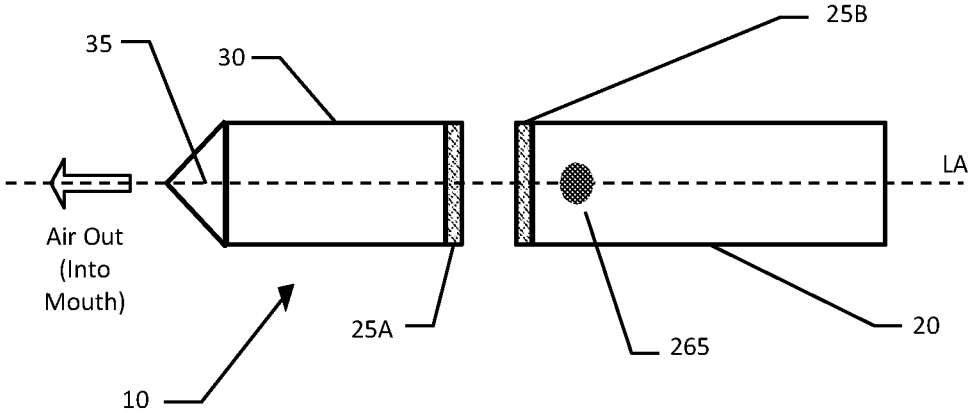

Figure 1

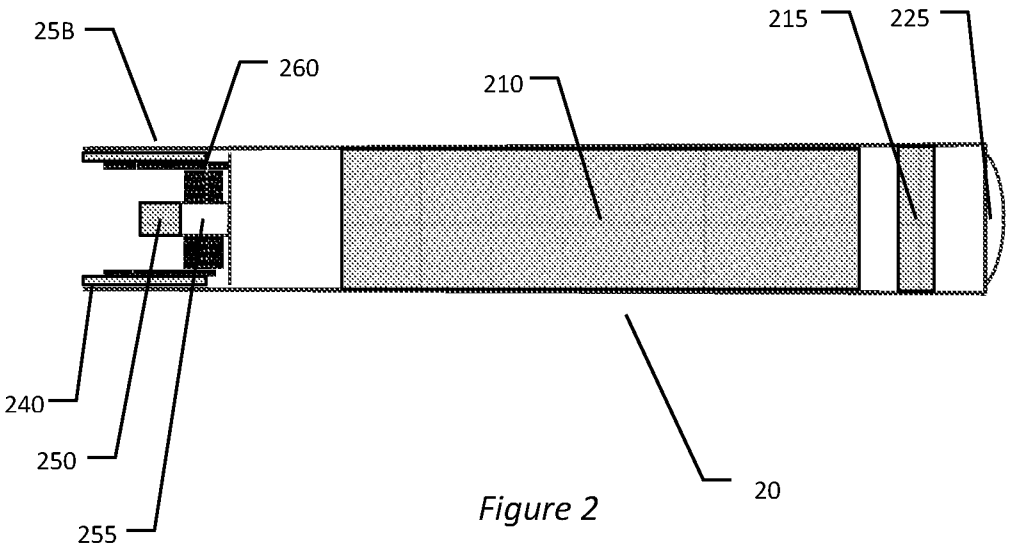

Figure 2

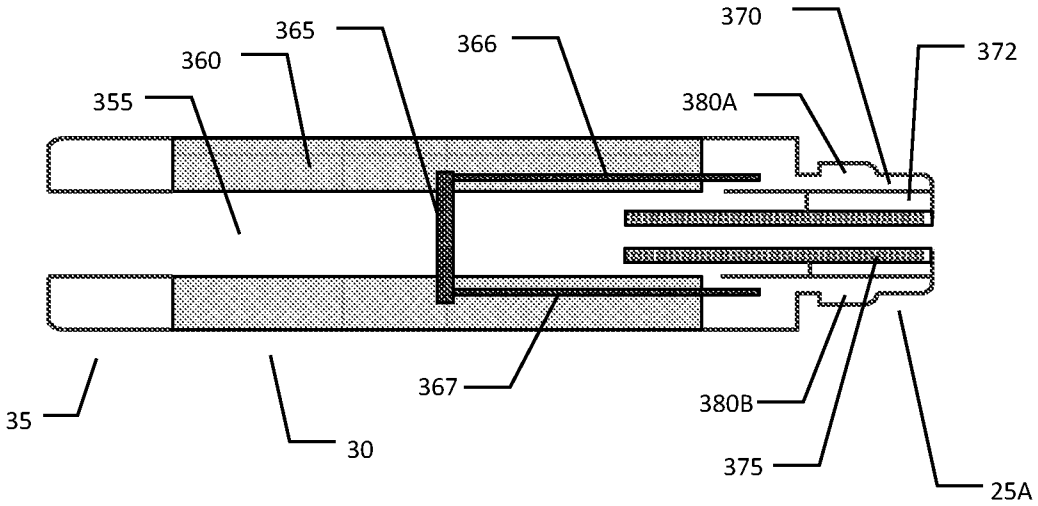

Figure 3

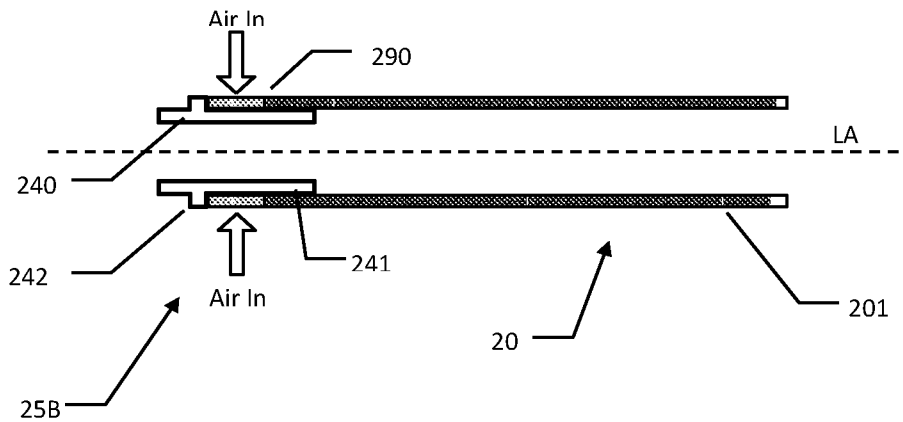

Figure 4

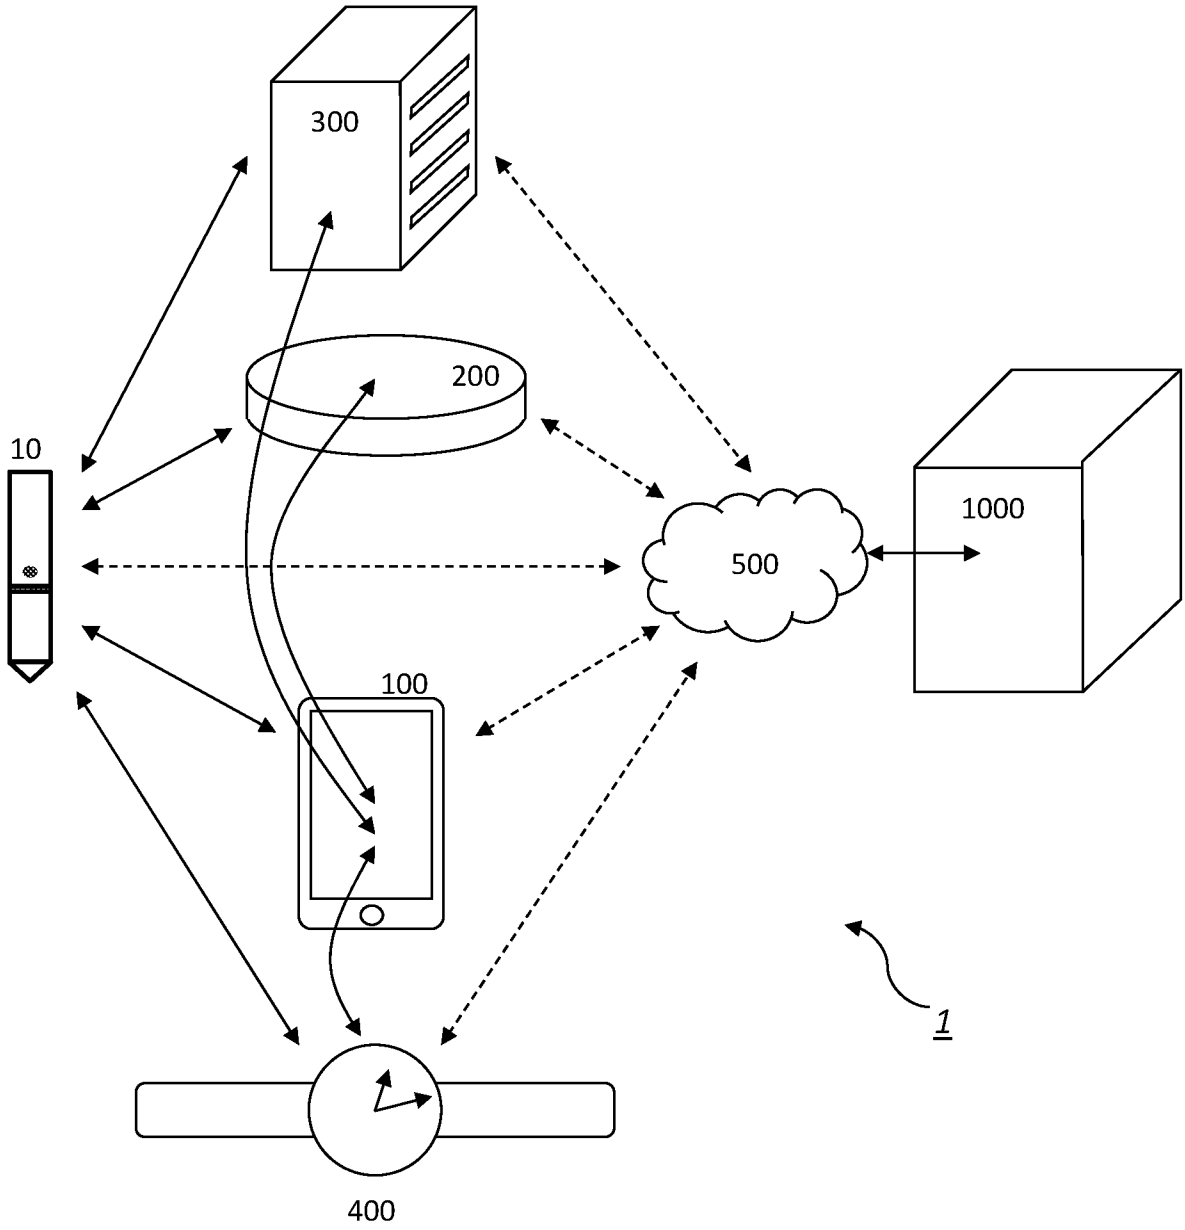

Figure 5

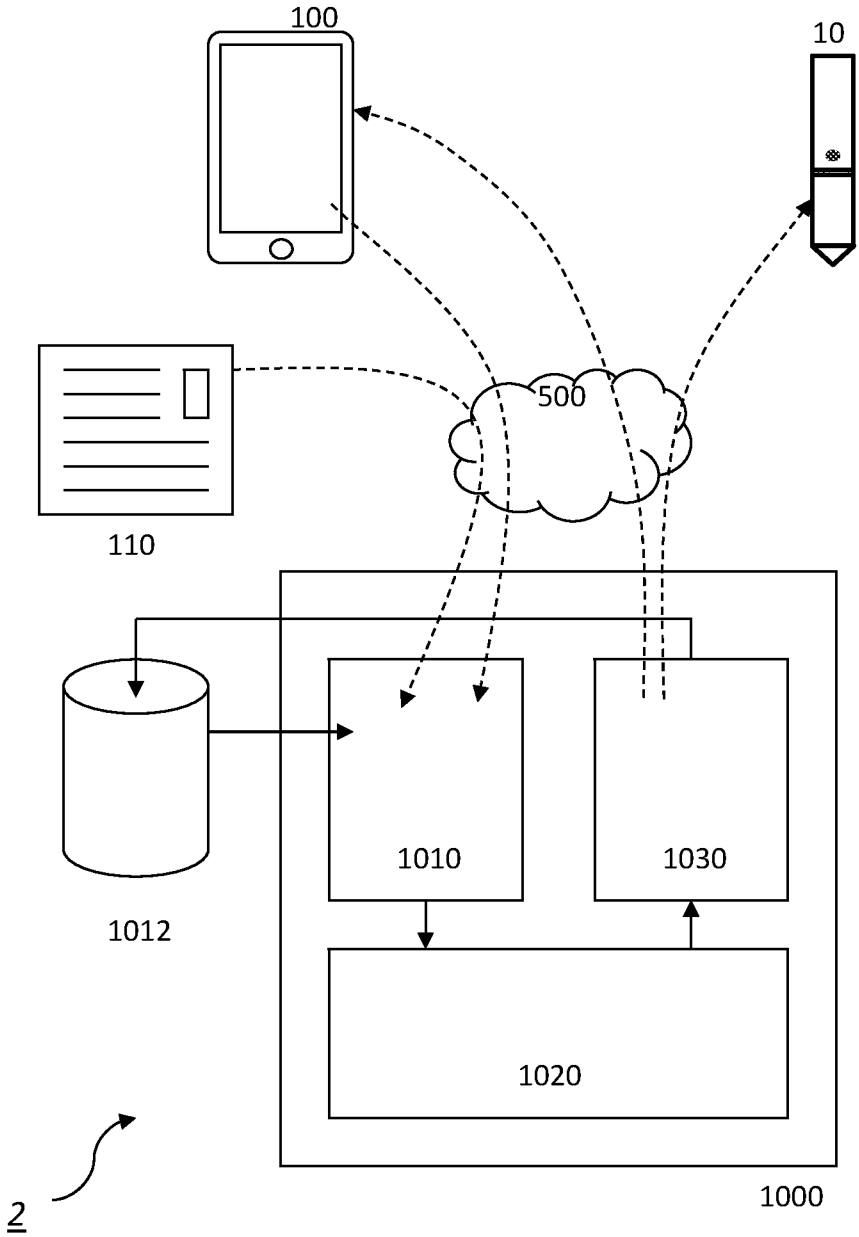

Figure 6

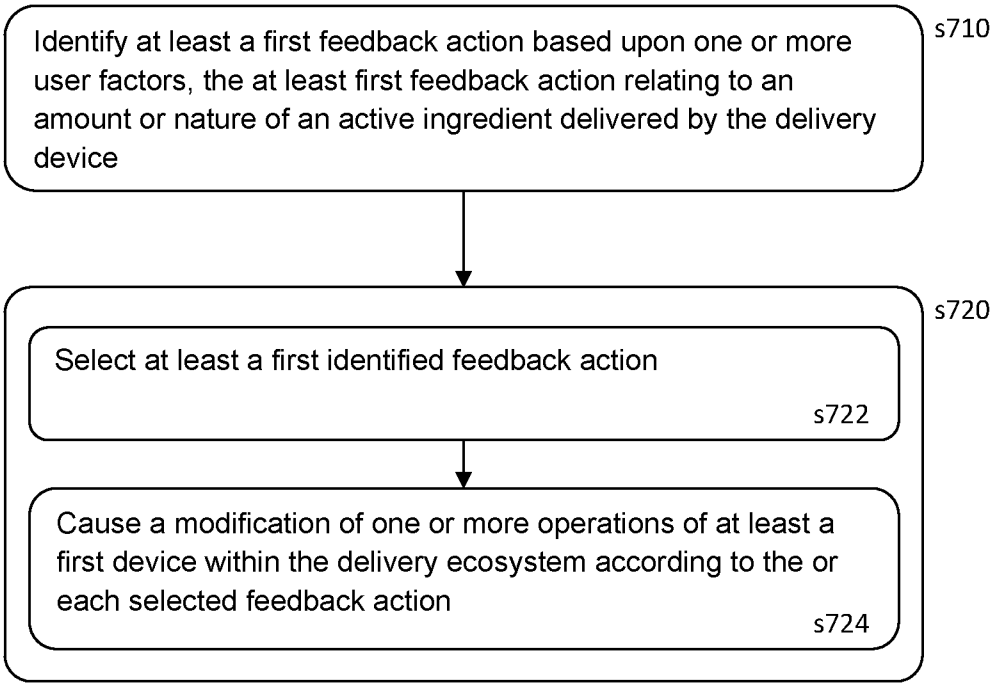
Identify at least a first feedback action based upon one or more user factors, the at least first feedback action relating to an amount or nature of an active ingredient delivered by the delivery device
s710
s720
Select at least a first identified feedback action
s722
Cause a modification of one or more operations of at least a first device within the delivery ecosystem according to the or each selected feedback action
s724

Figure 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2021/051523

A. CLASSIFICATION OF SUBJECT MATTER

INV. A24F40/00 A61M11/00 A61M15/06 G16H20/13 G16H40/63
G16H50/20
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

A24F G06F A61M G16H

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/291791 A1 (PRADEEP ANANTHA [US]) 22 November 2012 (2012-11-22) paragraph [0019] - paragraph [0028] figures 1, 2 -----	1-42
X	US 2019/387796 A1 (COHEN GAL A [US]) 26 December 2019 (2019-12-26) paragraph [0025] - paragraph [0064] paragraph [0069] - paragraph [0127] paragraph [0134] - paragraph [0144] paragraph [0152] - paragraph [0166] figures 1-9 -----	1-42

-/--

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
20 August 2021	30/08/2021

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Hauber, Jörg

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2021/051523

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"Method and system for smartwatch-based nicotine control in electronic cigarette - Project Specific invention ED - Darl Kuhn", IP.COM, IP.COM INC., WEST HENRIETTA, NY, US, 22 April 2019 (2019-04-22), XP013182960, ISSN: 1533-0001 page 1 figure 1 -----	1-42

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/GB2021/051523

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012291791 A1	22-11-2012	NONE	
US 2019387796 A1	26-12-2019	AU 2019295713 A1	21-01-2021
		CA 3104879 A1	02-01-2020
		EP 3813913 A1	05-05-2021
		US 2019387796 A1	26-12-2019
		WO 2020006109 A1	02-01-2020