nature

https://doi.org/10.1038/s41586-020-2442-2

Supplementary information

A mobile robotic chemist

In the format provided by the authors and unedited

Benjamin Burger, Phillip M. Maffettone, Vladimir V. Gusev, Catherine M. Aitchison, Yang Bai, Xiaoyan Wang, Xiaobo Li, Ben M. Alston, Buyi Li, Rob Clowes, Nicola Rankin, Brandon Harris, Reiner Sebastian Sprick & Andrew I. Cooper✉

Supplementary Information

A Mobile Robotic Chemist

Benjamin Burger, Phillip M. Maffettone, Vladimir V. Gusev, Catherine M. Aitchison, Yang Bai, Xiaoyan Wang, Xiaobo Li, Ben M. Alston, Buyi Li, Rob Clowes, Nicola Rankin, Brandon Harris, Reiner Sebastian Sprick and Andrew I. Cooper*

Leverhulme Centre for Functional Materials Design, Materials Innovation Factory and Department of Chemistry, University of Liverpool, 51 Oxford Street, Liverpool, L7 3NY, United Kingdom

Table of Contents

1. Autonomous Mobile Robot Platform

1.1 General Specifications.

The mobile robotic platform was supplied by KUKA and is shown in Extended Data Fig. 1. The platform consists of two robots. The first is a dexterous robotic arm with 7 independent degrees of freedom. The arm is a KUKA Leichtbauroboter ("built-for-light-weight robot") intelligent industrial work assistant, or LBR iiwa. The LBR iiwa has a load-bearing capacity of up to 14 kg and a reach of 820 mm.

The LBR iiwa arm was mounted on a Kuka Mobile Platform (KMP) base, which is an omnidirectional mobile base that is capable of moving the LBR iiwa throughout the laboratory. The size of the KMP is 1190 x 700 x 720 mm; it has a maximum speed of 3 km h^{-1} and can carry payloads of up to 200 kg. The omnidirectional wheels on this base allow the robot to move in any direction and it has a small turning circle, which is beneficial for use in a laboratory environment.

1.2 Laser Scanning.

The KMP uses a map for both navigation and localization, as shown in Fig. 1b in the main text. The robot is specified to navigate to within 1 mm precision under perfect, idealized conditions. However, in a real laboratory, there will be human operators in proximity to the robot while it is in operation—for example, when the stock of sample vials is being replenished—which may perturb the laser mapping. Also, laboratory benches and other large objects that form the basis of the map (Extended Data Fig. 1; Fig. S1) may move slightly over the course of the experiment; hence, the robot may not fulfill this nominal 1 mm precision specification. An experiment was therefore performed to assess how accurately the KMP is able to navigate within a real operational laboratory. In this experiment, the robot drove from each location on the map to every other location, measuring the position for a total of 563 steps. The results are shown in Figures S2–S5. These results show that the base is able to navigate to within 10 mm of the target location in a real laboratory environment (standard deviation 1.64 mm), with an orientation (angle) within 0.5 degrees of the target angle (standard deviation 0.070 degrees). This precision is more than high enough to navigate to the various stations in the experiment, but it is insufficient to carry out delicate operations such as loading solid-dispensing cartridges into the Quantos solid dispensing unit. To address this, a secondary positioning method was developed (section **1.3**).

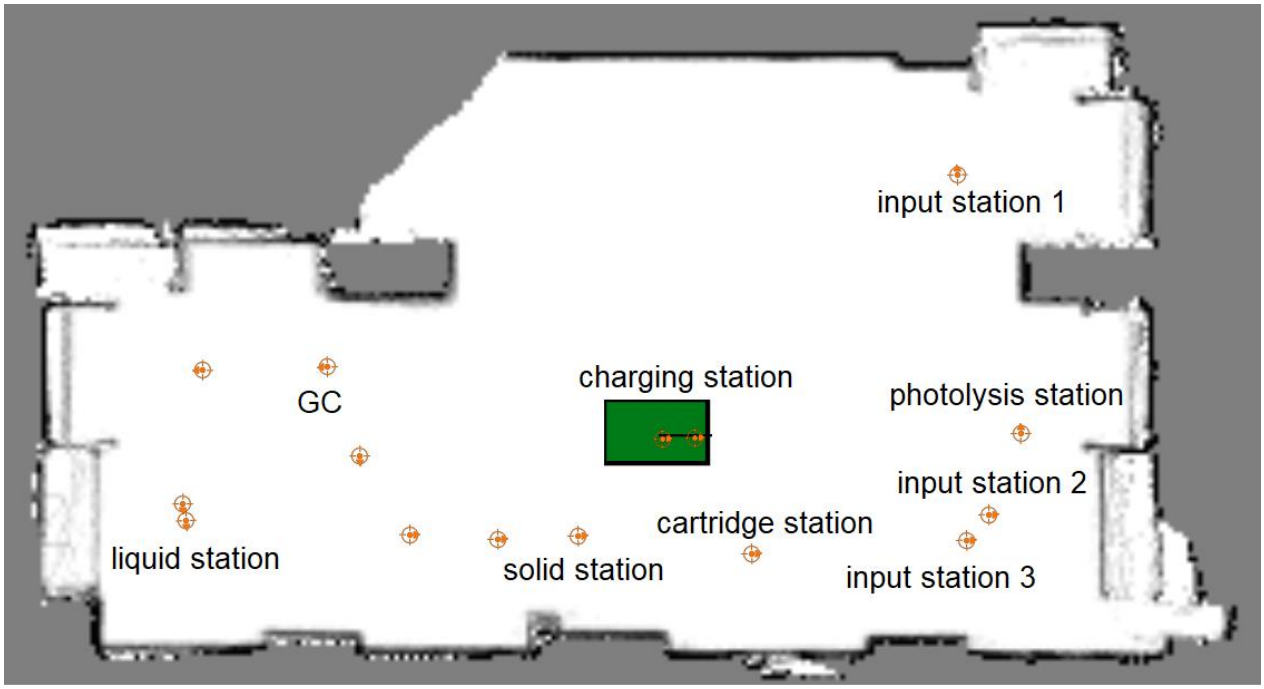

Figure S1. Laser-scanning map of the laboratory. The map is generated automatically by the laser scanners on the robot. The key locations for the workflow are labeled; the green rectangle represents the robot, in the same position as for the photograph, and the orange crosshairs on the map indicate recorded locations. Black location cubes are fixed to the benches at key location points (orange crosshairs; see also section **1.3**).

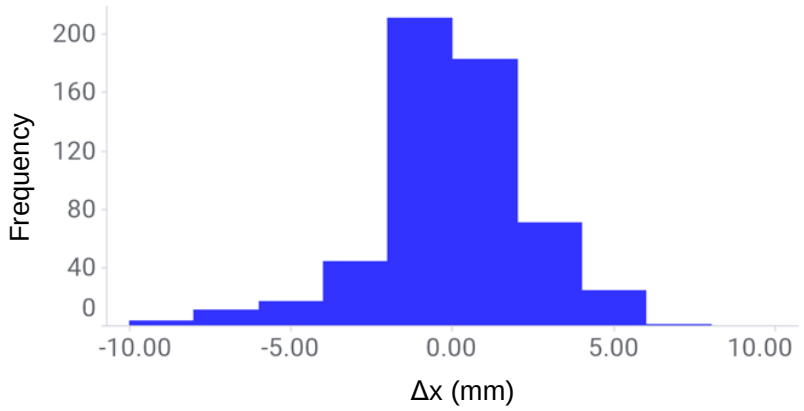

Figure S2. Histogram of Δx for the position of the KMP after laser navigation. $\Delta x = x_{cube} - x$, where x_{cube} = position of the calibration cube and x = actual position of KMP.

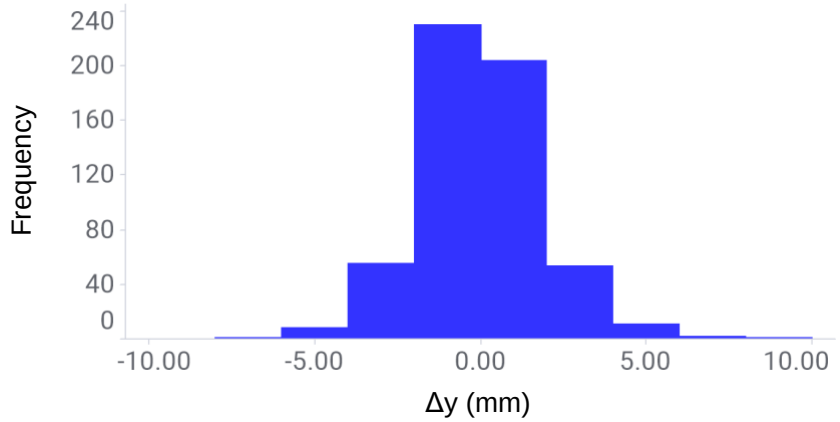

Figure S3. Histogram of Δy for the position of the KMP after laser navigation. $\Delta y = y_{cube} - y$, where y_{cube} = position of the calibration cube and y = actual position of KMP.

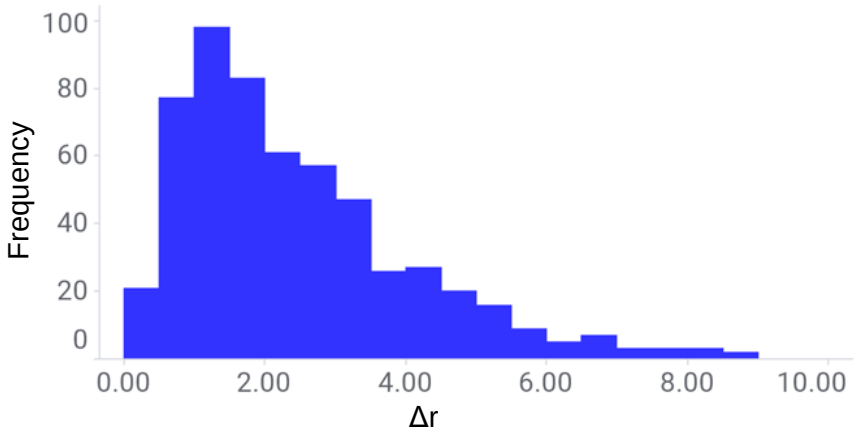

Figure S4. Histogram of r for the position of the KMP after laser navigation, $r_2 = \Delta x_2 + \Delta y_2$.

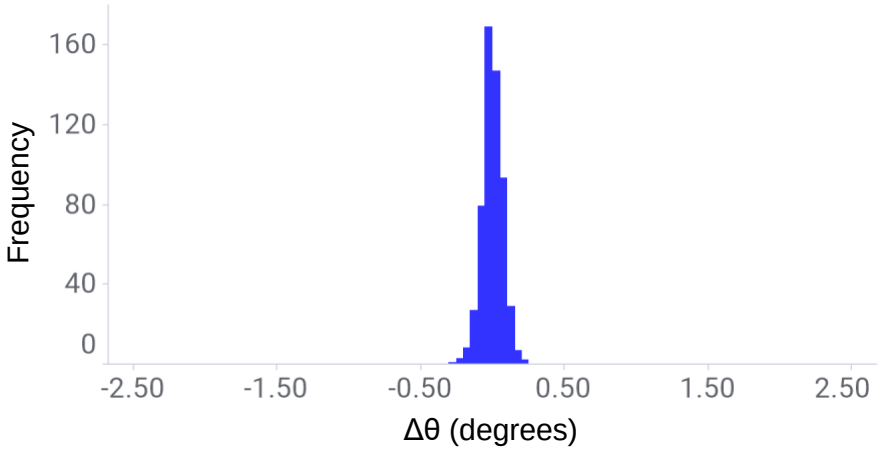

Figure S5. Histogram of the orientation of the KMP after laser navigation. $\Delta\theta = \theta_{location} - \theta$, where $\theta_{location}$ = target orientation at location and θ = actual orientation.

1.3 Six-Point Positioning

A force-feedback or 'touch' routine was used as a secondary positioning method to improve the location precision. When the robot approaches an experimental station, it measures its position with respect to a black cube (dimensions 100 x 100 x 50 mm) that is fixed in space with respect to the station (see *e.g.*, Extended Data Fig. 1). First, the robot touches three points on the top plane of the cube (Figure S6; t_1, t_2, t_3) that are detected by force feedback; these three points uniquely define plane x. The robot then performs two more touches (Fig. S6; t_4, t_5) on plane y that is perpendicular to plane x. The last touch (Fig. S6; t_6) is on the third plane, z that is perpendicular to the other two planes. The intersection of these planes defines the origin of the coordinate frame, and the normal vectors of the planes are used as axes.

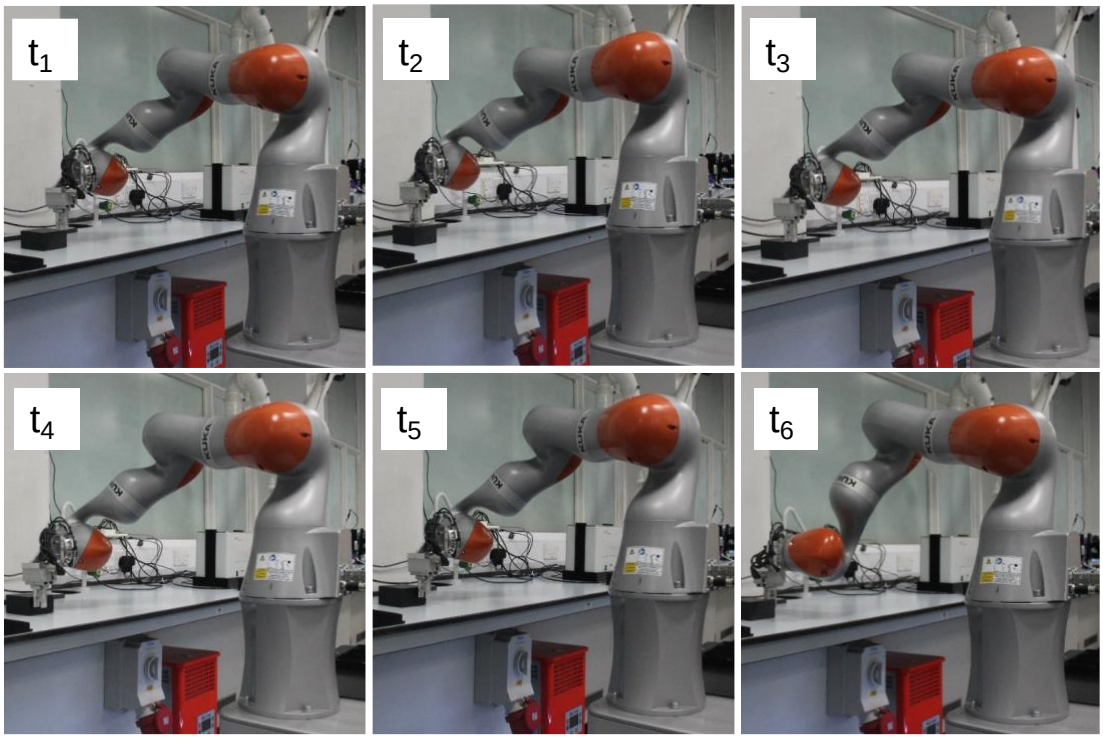

Figure S6. Robot touching a black calibration cube during 6-point calibration. Photographs t_{1-3} show the robot touching the top x-plane three times. Photographs t_4 and t_5 show the robot touching the perpendicular plane y on the side of the cube. Photograph t_6 shows the robot touching the third z plane.

To test the precision of this method, 100 independent 6-point calibrations were performed with the starting point—that is, the initial cube position estimate—being subjected to noise. Under this test, both the position (x, y) and the orientation (θ) of the robot may vary, as simulated using uniform noise in both x and y (up to ±15 mm) and θ ± (2.5 degrees). Note that the noise applied exceeds the measured positioning error for the KMP navigation (10 mm; Figs. S2–S5), and can therefore be considered as a worst-case scenario. The results of these tests are shown in Figures S7–S10. This showed that the cube position was found to within ±0.12 mm, and that any orientation offset is corrected by 6-point positioning to less than 0.005 degree.

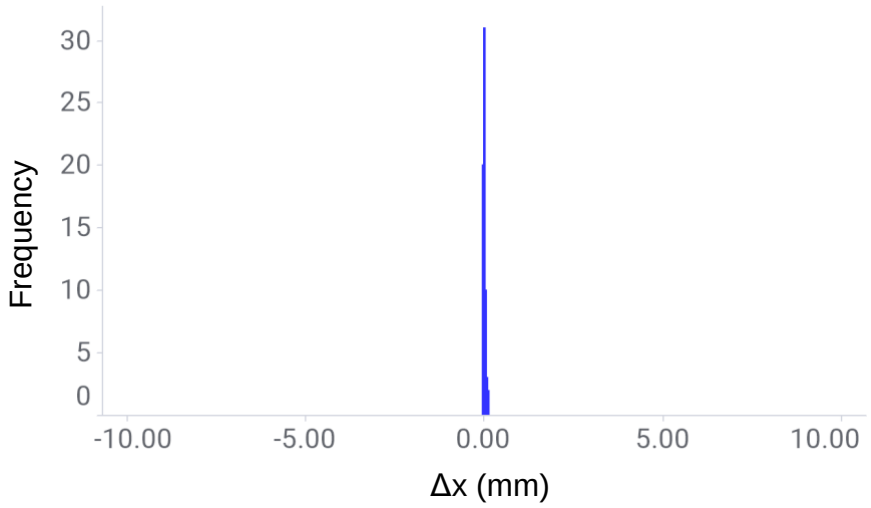

Figure S7. Histogram of Δx after 6-point positioning. $\Delta x = x_{cube} - x$, where xcube = position of the calibration cube and x = actual position of KMP (plotted on the same scale as Fig. S2).

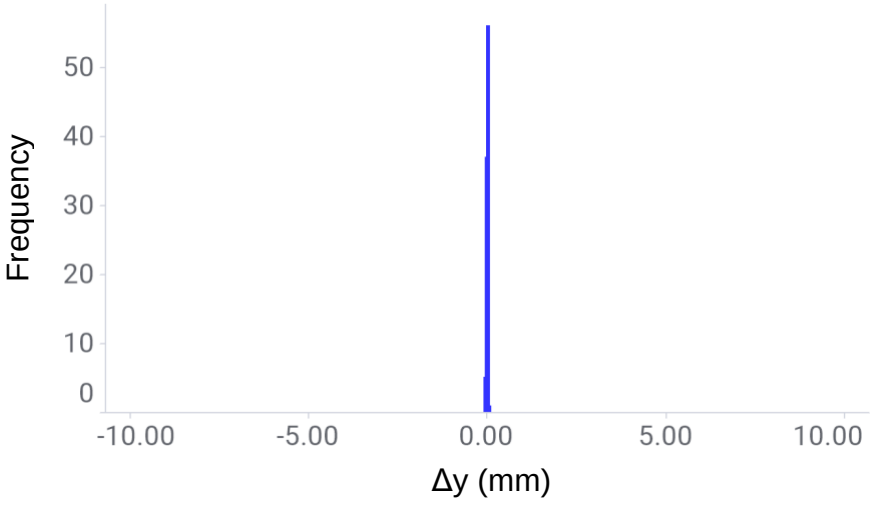

Figure S8. Histogram of Δy after 6-point positioning. $\Delta y = y_{cube} - y$, where ycube = position of the calibration cube and y = actual position of KMP (plotted on the same scale as Fig. S3).

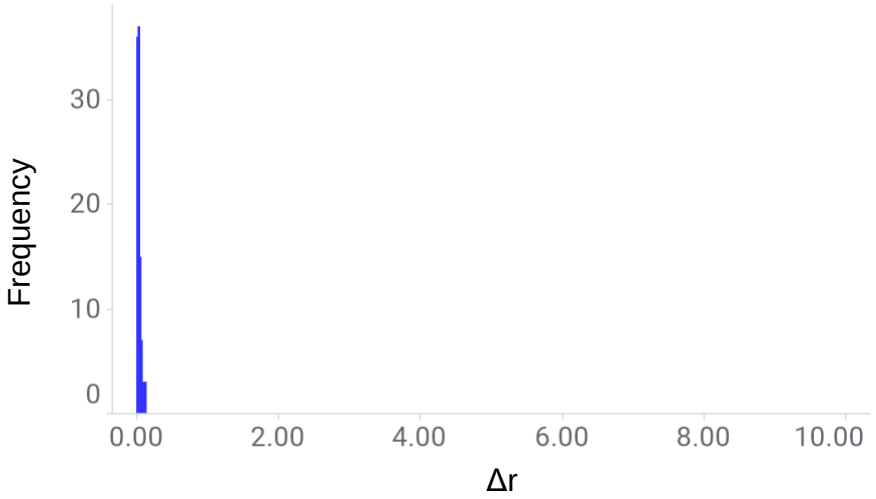

Figure S9. Histogram of r after 6-point positioning. $r_2 = \Delta x_2 + \Delta y_2$ (plotted on the same scale as Fig. S4)

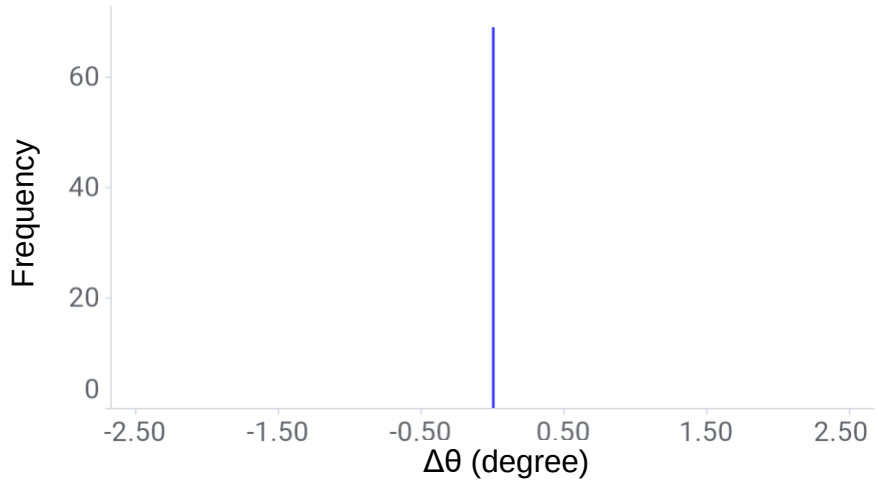

Figure S10. Histogram of the average angle after 6-point positioning. $\Delta\theta = \theta_{location} - \theta$, where $\theta_{location}$ = target orientation at location and θ = actual orientation (plotted on the same scale as Fig. S5).

Without the touch positioning routine, the localization precision for the robot was too low to perform the tasks required in the workflow. For example, reliable placement of vials (Extended Data Fig. 5c) and solid dispensing cartridges (Extended Data Fig. 5d) requires a repeatable precision of < 0.5 mm. As shown in Figure S38, the percentage error rate for the platform in the period May 2019 – July 2019 was reduced from 1.5 % to < 0.5 % (per sample processed). Of these errors, only a small subset can be attributed to the robot itself (0.07–0.1% per month; that is, less than one sample in 1000; grey bars in Fig. S38). This shows that the touch localization approach is robust. For the 3 robot errors that did occur in the period May–July 2019 (over a total of 5201 samples processed), we assign this to variance issues in the LIDAR map, rather than the touch localization itself.

1.4 Gripper and Sample Rack Design

The specialized metal gripper for this workflow is shown in Extended Data Fig. 5a. This parallel, multipurpose gripper was designed to accomplish a variety of tasks including 6-point calibration (section **1.3**), and object manipulations such as grasping solid-dispensing cartridges (section **3.1**), sample transport racks, and individual sample vials. The gripper features indents on the inside of the fingers (Extended Data Fig. 5a) that allow the gripper to grasp vials in two different ways. The gripper can grasp a vial top-down (Extended Data Fig. 5b) or sideways (Extended Data Fig. 5c). This allows for the placement of vials in racks (using the top grasp) as well as the positioning of vials into more constrained apparatus, such as the capping station (section **3.2**). The indents are also able to accommodate solid-dispensing cartridges (Extended Data Fig. 5d). The fingers of the gripper have protrusions on the outside (Extended Data Fig. 5a) that allow the gripper to lock into the sample racks by expanding outwards into two matching recesses, as shown in Extended Data Fig. 5e.

This gripper design was found to be extremely reliable: indeed, we have not yet encountered a grasping failure during operation. As such, the control software only checks for placement failures, rather than grasping failures, although in principle the hardware could be configured to do both in the future. This might become more important in workflows that involve grasping a wider variety of objects.

1.5 Safety Protocols

The mobile platform consists of two robots, the LBR (robotic arm) and the KMP200 (mobile base). Each of these robots has its own safety settings.

Mobile KMP200 base: The safety mechanism of the mobile KMP200 base uses a bounding box to detect obstacles by laser scanning. Once detected, the robot either slows down (decreasing the size of the bounding box) or stops completely if the obstacle does not move. We limited the speed of the mobile base to a maximum speed of 0.3 m s^{-1}. When the velocity is between 0.1 m s^{-1} and 0.3 m s^{-1}, then the bounding box extends 45 cm around the robot, when unexpected obstacles are detected, the robot decelerates to 0.1 m s^{-1}. When the speed is 0.1 m s^{-1} or lower, then the bounding box extends 30 cm around the robot and stops when objects are detected inside this box. In addition, the mobile base can only drive when the arm is tucked into the drive position; it cannot move with the arm extended beyond the mobile base. When samples are transported, they are always placed into the sample rack on the deck of the mobile base (see *e.g.*, the two sample racks in Fig. 1a); they are not carried by the robot arm itself. The lights surrounding the robot base flash just before it starts to move (see Supplementary Movies) as a visual warning. The mobile robot base also has a red emergency stop button.

LBR arm: The robot arm can only move while the base is stationary. The maximum speed of the arm was limited to 40 cm s^{-1} (at any point on the arm). During arm movement, there is a constant collision detection protocol that triggers a safety stop when the robot encounters unexpected torque above 35 Nm. For any placements, or when the robot is close to a fixture, the velocity of the robot arm is limited to 10 cm s^{-1}, and a lower force of 20 N triggers a safety stop.

With these safety settings applied, we assessed that it was safe for operators to work in the same area as the robot while it was carrying out experiments (for example, to reload consumables), but all such operators were given a full safety induction prior to doing this. The work area was also monitored 24/7 using CCTV cameras (section 9; Supplementary Fig. S39) that recorded activity at each station, thus allowing us to optimize performance and minimize any errors that could create a safety issue.

2. Experimental Stations

This workflow comprised eight separate experimental stations, distributed around the laboratory; a solid dispensing station (section **2.1**), a combined liquid dispensing and capping station (section **2.2**), a sonication station (section **2.3**), a photolysis station (section **2.4**), a gas chromatography (GC) analysis station (section **2.5**), and three separate input stations for consumables (sample vials, racks) and storage of completed racks (section **2.6**).

2.1 Solid Dispensing Station

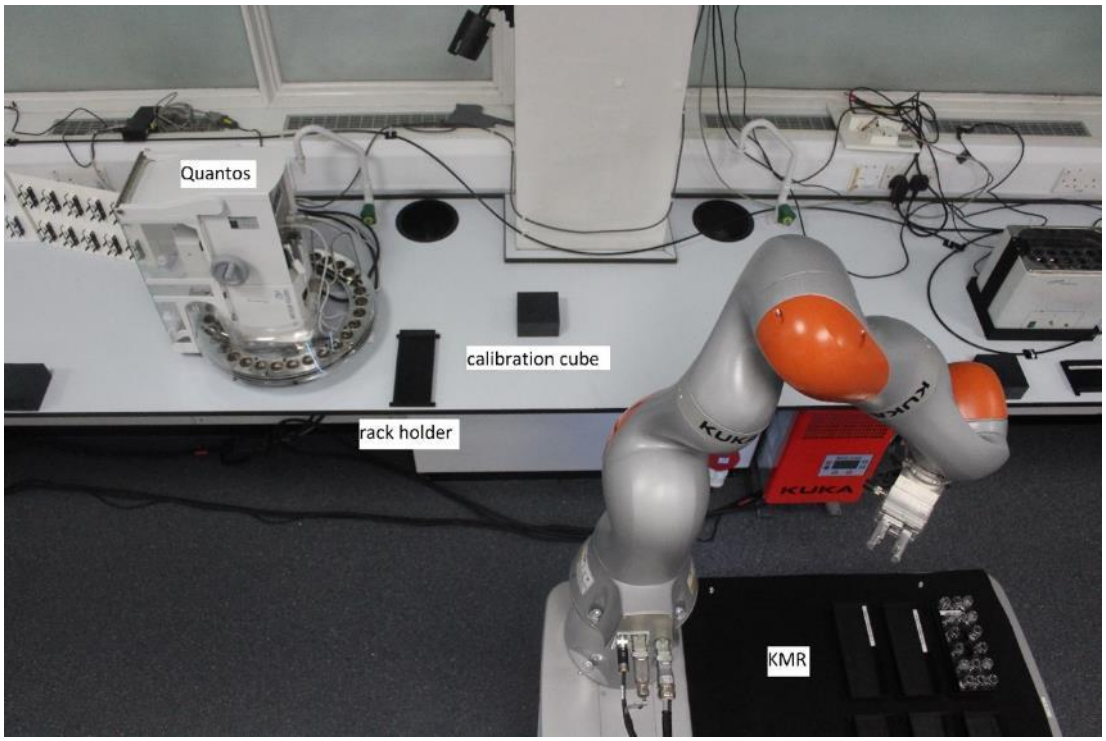

Figure S11. Photograph showing the robot at the solid-dispensing station, just prior to loading the Quantos solid dispenser with empty sample vials (see rack carried on the KMR).

The workflow used an automated Quantos (QS30) solid dispenser (Mettler-Toledo) to dispense solids. This instrument can dispense a range of organic and inorganic solids. The performance that we achieved for dispensing the polymer photocatalyst, P10, is shown in Figure S12 (6,608 dispenses). A similar performance was attained for a range of other organic solids over 6,600 separate dispenses, as shown in Figure S13. The mobile robot both loads empty sample vials into the Quantos carousel (Figure S11) prior to solid dispensing and loads/unloads the solid-dispensing cartridges into the Quantos instrument (cartridge station; Extended Data Fig. 3a), allowing up to 20 separate solids to be dispensed, or more if additional cartridge hotels are added.

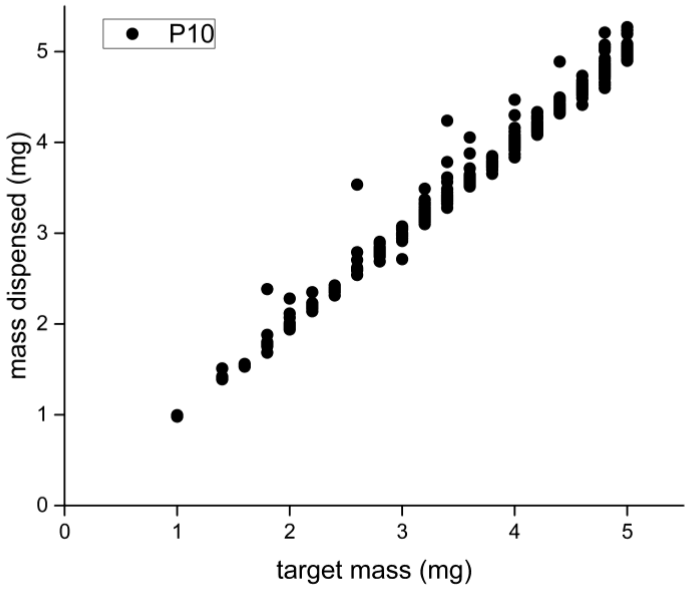

Figure S12. Plot showing solid dispensing performance for P10 during the autonomous search and earlier tests (total of 6608 dispenses).

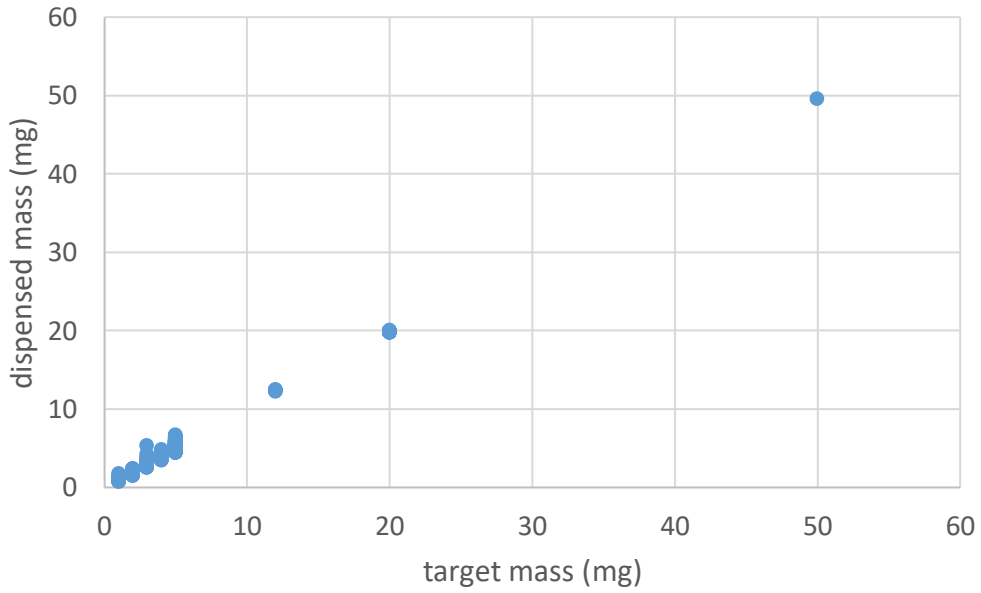

Figure S13. Solid dispensing performance over 6,600 dispenses for a range of organic solids (1–50 mg per dispense)

2.2 Liquid Dispensing and Capping Station

Inexpensive custom-built liquid dispensers were developed for this workflow, although it would also be straightforward to interface this mobile robot with a commercial liquid handler. A bulk industrial peristaltic pump with controller and a microcontroller was used to build two liquid stations, as shown in Figure S14. Peristaltic pumps from Watson and Marlow (200 series)[45] were used, which are essentially stepper motors combined with a peristaltic pump

head. These pumps employ a 0–5 V control voltage that was controlled here by using an Arduino micro-controller. Arduino controllers are limited to pulsed width modulation to generate such a signal, which is not suitable here because it is not continuous. We therefore used a Digital-to-Analog Converter (DAC) from Sparkfun, which converts the 12-bit signal from the Arduino into a 0–5V analog signal. This signal was then used to control the speed of the motors. The pumps are powered by 30 V DC, and the supply voltage to the pumps was controlled using mechanical relays. The wiring diagram for this is shown in Figure S15 for a single pump. This can be extended to multiple pumps by adding a relay for each pump required, assuming that only one pump is used at a given time. Direction and speed can be split for each pump, but an additional relay (marked as “Pump On”) was needed for each additional pump. In the set-up used here, there were seven pumps on each station allowing for up to 14 separate liquids to be dispensed. Because the robot is mobile and can move between multiple stations, this approach is scalable to a much larger number of liquids, if required, given sufficient bench space.

The amount of liquid dispensed was measured gravimetrically using two balances (Kern PCG 2500-1), which led to high dispensing accuracy for a range of aqueous solutions (Figs. S16 and S17). To achieve such reproducible and reliable dispensing, it was necessary to implement dynamic control the pump speed depending on how much liquid was left to dispense. That is, the dispensing speed needs to slow down toward the end of the dispensing operation. A traditional approach from the field of control theory was used to achieve proportional-integral-derivative (PID) control.

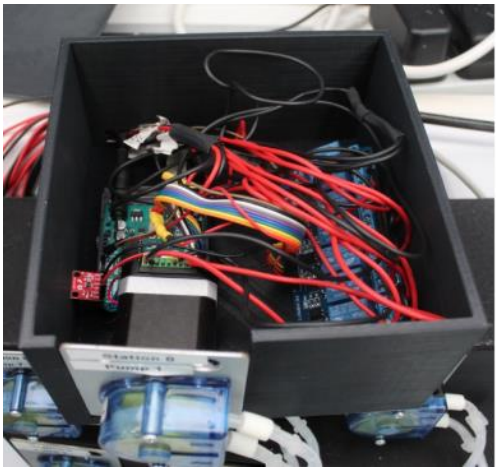

Figure S14. Custom-built peristaltic pump controlled by an Arduino microprocessor. The peristaltic pump and controller can be seen toward the bottom of this photograph; inside the box is the microcontroller and a digital-to-analogue-converter (DAC) to generate an analogue control signal for the pump.

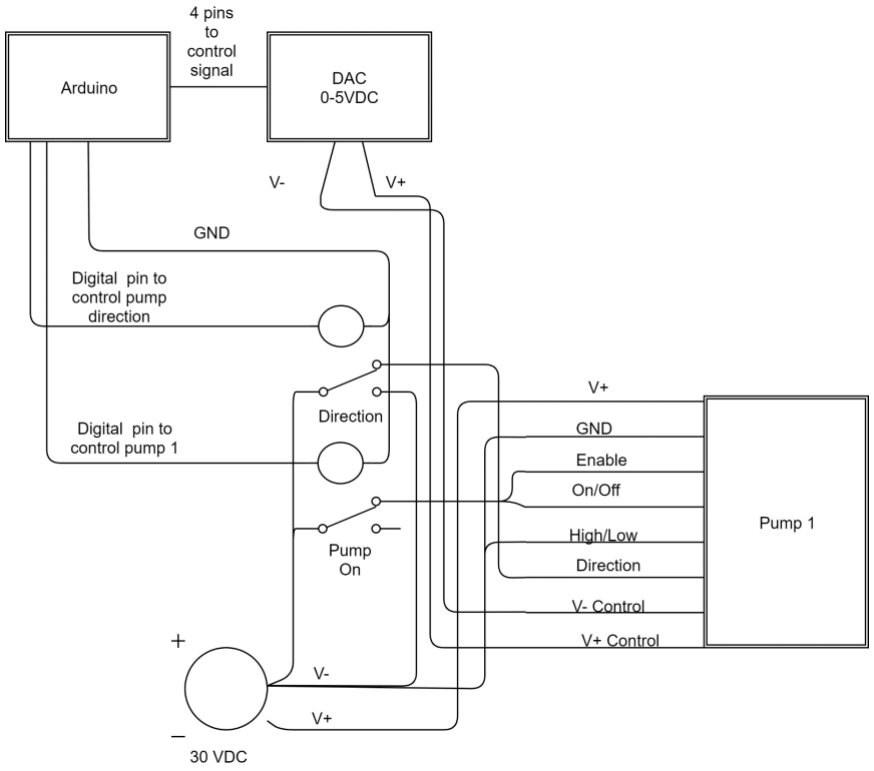

Figure S15. Wiring diagram for the liquid-dispensing station, showing how an Arduino microprocessor, a DAC, relays, and the pump were wired into a RS-232 controlled platform.

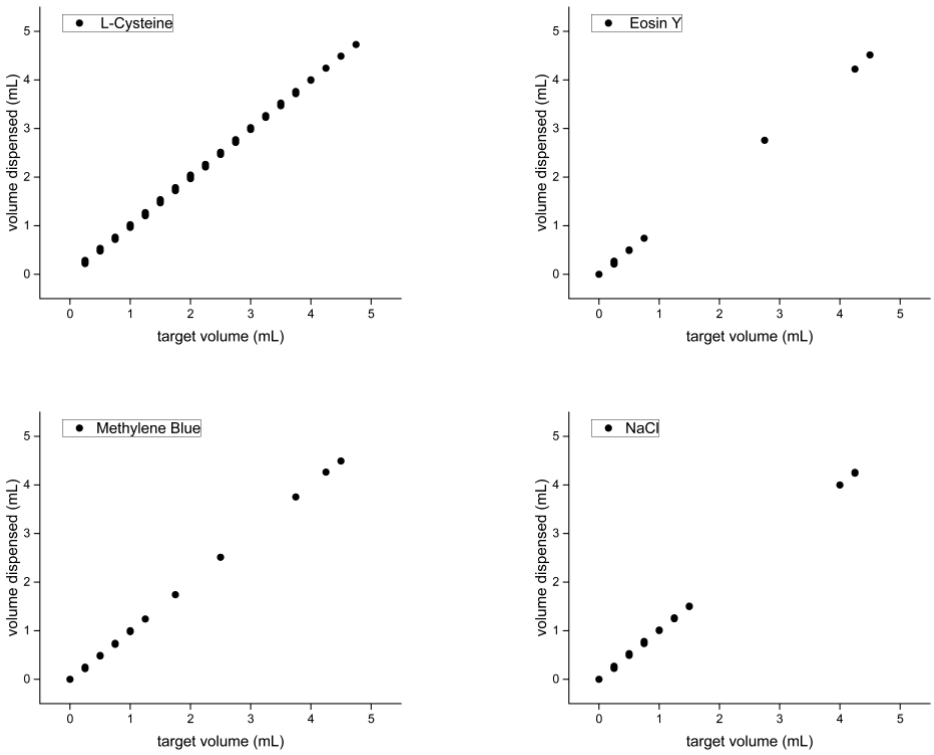

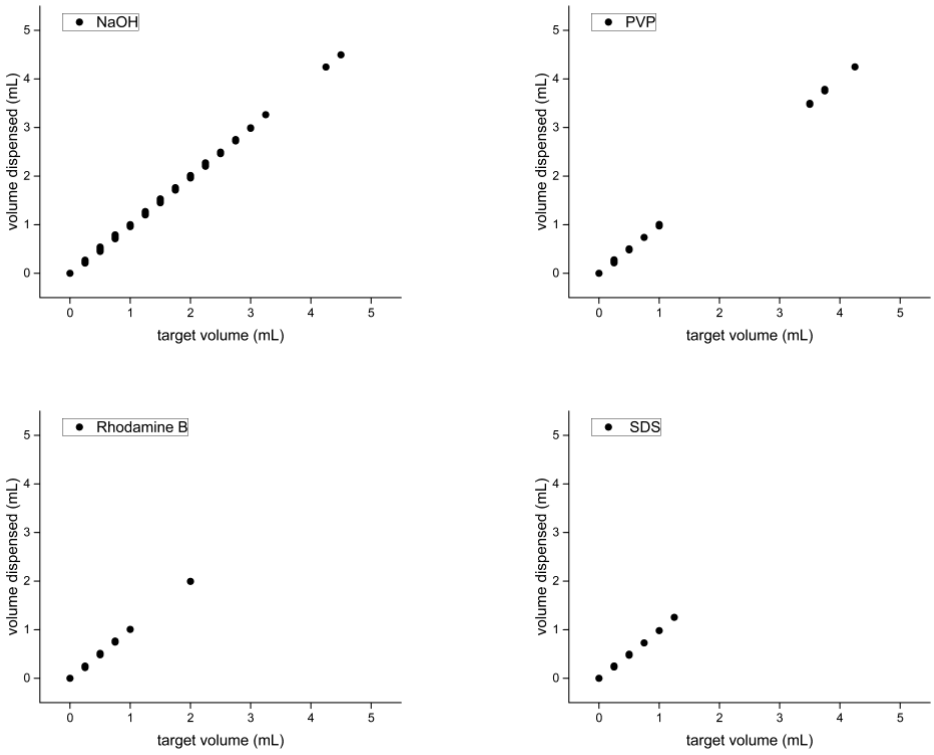

Figure S16 Liquid dispensing performance for aqueous solutions of L-cysteine, Eosin Y, Methylene Blue, NaOH, PVP, Rhodamine B, and SDS during the autonomous search.

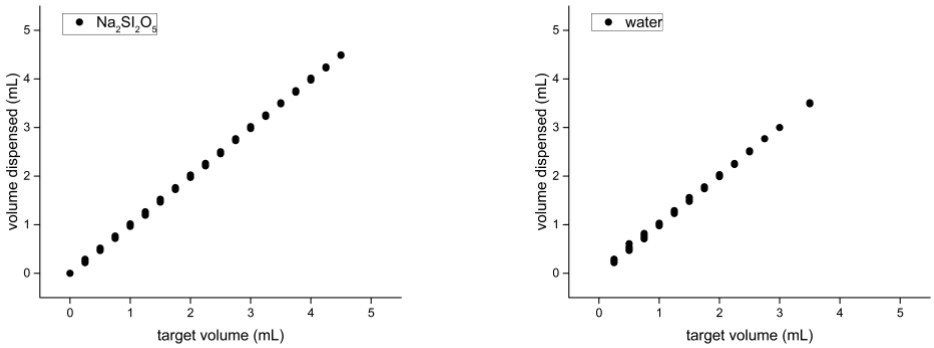

Figure S17 Liquid dispensing performance for an aqueous solution of $Na_2Si_2O_5$ and for pure water during the autonomous search.

The PID controller was used to control the speed at which the liquid was pumped into the sample vial. It did this by measuring how much liquid had been added into the vial to allow accurate and robust dispensing of small liquid volumes (0.25–5.0 mL). The PID control operated according to the following equation:

$$u(t) = K_p e(t) + K_i \int_0^t e(t)\,dt + K_d \frac{de(t)}{dt} \qquad \text{(S1)}$$

where *u(t)* is the input parameter for the system, *e(t)* is the error defined as the output measured by the system minus the target steady-state output, and K_p, K_i, K_d are constants for the proportional action, integral action and derivative action.

The first term of the equation is commonly referred to as proportional action (P-control) and is simply proportional to the control error; that is, this term effectively dispenses liquid at a speed that is linear to the error or the amount of liquid left to dispense. The second term, integral action (I-control), ensures that the process output agrees with the steady state. By using only P-control, the system will always have a (small) error, and with only P control this will lead to an either increasing or decreasing signal. The final term in Equation S1 is the derivative action (D), which improves the closed-loop stability. In a typical process, the dynamics will take some time before a change in the control variable is detectable in the output and will hence lag behind. The derivative term is used to predict the changes in the system and to improve the performance of the closed loop[46]. A snippet for the control code is shown in Figure S18.

In this implementation, the PID values (K_p = 20, K_i = 0.1, K_d = 0.8) were used; these values were optimized manually by trial and error. Based on these values, the voltage was controlled and used to regulate the speed of the pump. Validation results are shown in Figure S19 and Figure S20. Using PID control it is possible to dispense all liquids tested (including mixtures, organic solvents and viscous liquids) tested with the range of 0.5 mL to 5 mL accurately and with minimal error. Supplementary Movie S4 shows a liquid station dispensing a volume of 1 mL.

```
continue_dispensing = True
while(continue_dispensing) {
                rate = CalculateRate(current_mass_dispensed)
                pump.set(rate)
                wait(200)
                if target <= current_mass_dispensed {
                                continue_dispensing = False
        }
}
```

Figure S18. Code snippet for dispensing liquids under PID control.

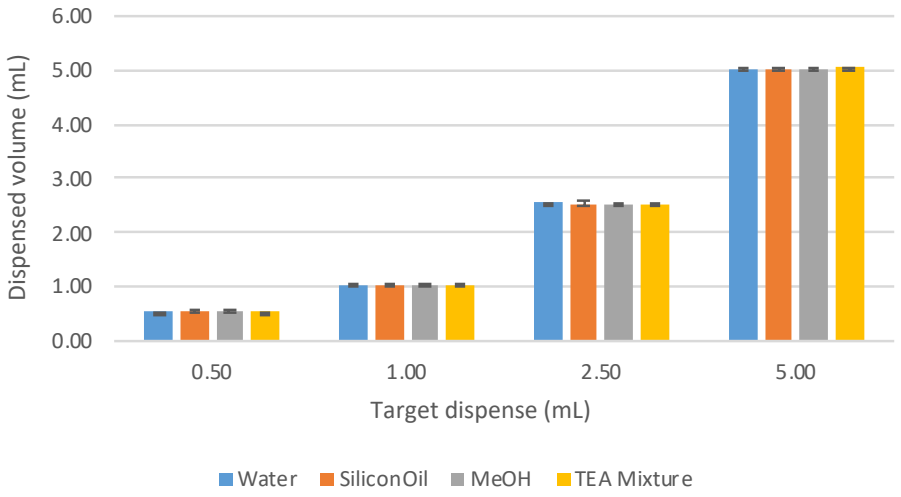

Figure S19. Bar chart showing accuracy for liquid dispensing using a feedback loop. Dispensing of these four liquids was accurate to within a margin of error of 0.05 mL. The error bars are for 50 repeat dispenses in each case.

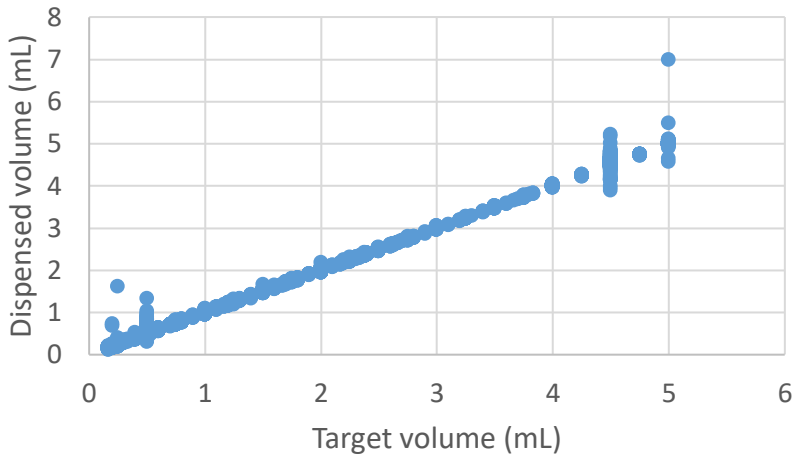

Figure S20. Liquid dispensing performance for more than 20,000 separate dispenses across a variety of liquids during the development of the robot workflow (2018/2019).

The liquid station was also specified to allow continual purging of the stock solutions with nitrogen over the course of the experiment. This was done to limit oxidation of the stock solutions and to ensure a consistent oxygen background.

After addition of both solids and liquids to the vials, the next step was to cap the vials under an inert atmosphere with a pierceable HPLC cap. Since no commercial instrument was available that could do this, we designed one specifically for this workflow; the Liverpool Inertization Capper-Crimper (LICC; built by Labman Automation). This was positioned adjacent to the two liquid stations, such that the liquid-filled vials could be capped immediately after filling, thus saving time and unnecessary transportation tasks (Extended Data Fig. 3c). The LICC is a custom-built capper-crimper that can purge the vials with nitrogen (or other gas) before capping (Figure S21). The instrument crimps 20 mm crimp caps onto the 10 mL headspace vials used in this workflow.

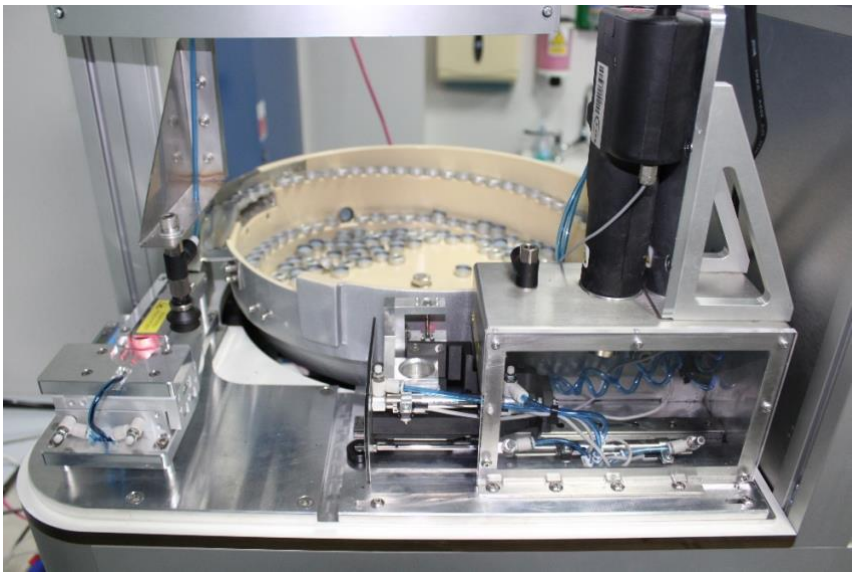

Figure S21. Photograph of the Liverpool Inertization Capper-Crimper (LICC). The bowl feeder that contains a stock of vial caps can be seen in the rear. The inertization chamber is in the lower right of the photograph and has a plastic window at the front to allow observation of the capping process in case of failure (see also section **9** on 24/7 monitoring). The crimper itself is the black, vertically-mounted instrument in the top right of the photograph.

The reproducibility of the inertization/capping process was tested by making four batches of 16 samples and measuring the residual oxygen background. Samples were prepared using 5 mL of N_2-purged water, dispensed using the liquid-dispensing stations, and then immediately inertised and cap-crimped. The results are shown in Figure S22, which shows good uniformity over the samples. The oxygen peak area is 419,664 pA (peak area) on average with standard deviation 10,897 pA, which corresponds to 1.40 ± 0.04 μmol.

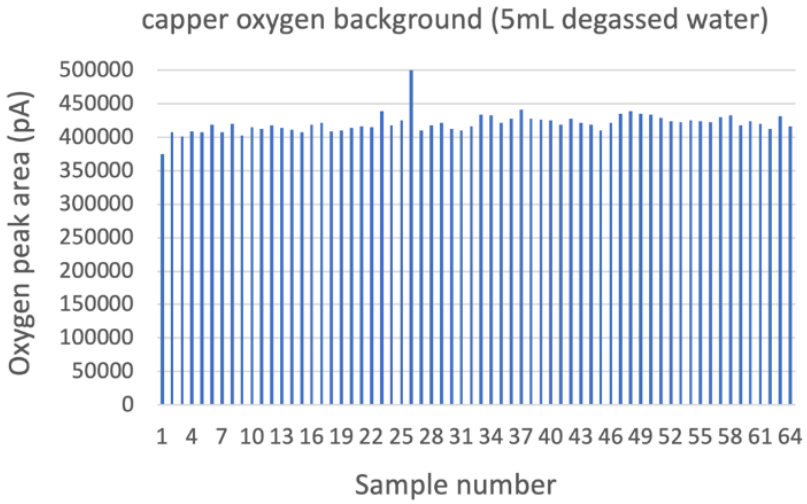

Figure S22. Oxygen background for repeat capping of inertized vials containing water. In this experiment, 64 vials were each purged for 60s, sonicated (10 min), and then submitted for GC analysis. The oxygen peak for sample (27) was off scale and failed the 'twist test' (*i.e.*, was easily twisted by hand) and was assumed to be improperly sealed.

2.3 Sonication Station

An ultrasonic bath was also integrated into the workflow to allow, where necessary, dispersal of solid catalysts into the aqueous mixture. The robot can load and unload the ultrasonic bath (Figure S23); we also fitted a vortex drying device to allow drying of the sample vials after immersed into water during sonication (Supplementary Movie S3).

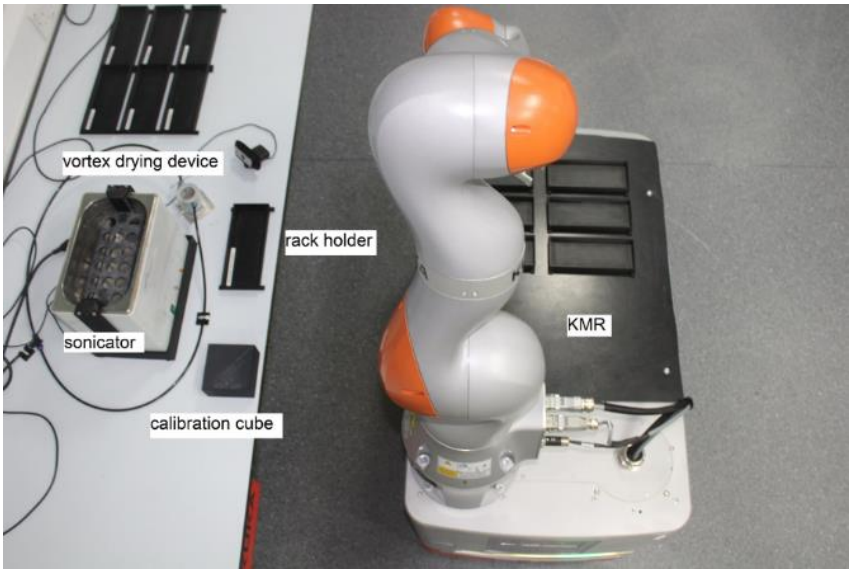

Figure S23. Photograph showing the robot at the sonication station. A vortex drying device was included to allow the robot to dry the sample vials (Supplementary Movie S3; 1 min 11 sec), because they are immersed in water during sonication.

2.4 Photolysis Station

To illuminate the samples after capping, we designed and built a custom photolysis station with a vibrating rack to suspend the photocatalyst during reaction (Extended Data Fig. 3b; Supplementary Movie S6). The system uses a vibratory motor (VIBTEC, M3/20) to vibrate the sample rack. Glass beads (4 mm diameter) were added into the vials to assist with sample mixing. To approximate solar light, the station uses a combined light source that is composed of two different types of lamps; LED panels (400–700 nm) and BL368 fluorescent tubes; the spectrum of the combined light source is shown in Figure S24. The samples are illuminated from below. Both the sample rack and the motor were cooled by a fan (12 W, axial,120 mm), mounted close to the vibratory motor. Additional fan cooling was also fitted to the aluminium illumination box itself (not visible in this photograph). The temperature of the sample rack was monitored using two TM (Papouch) thermocouples and was found to vary between 20.6–23.5 °C in this set of experiments.

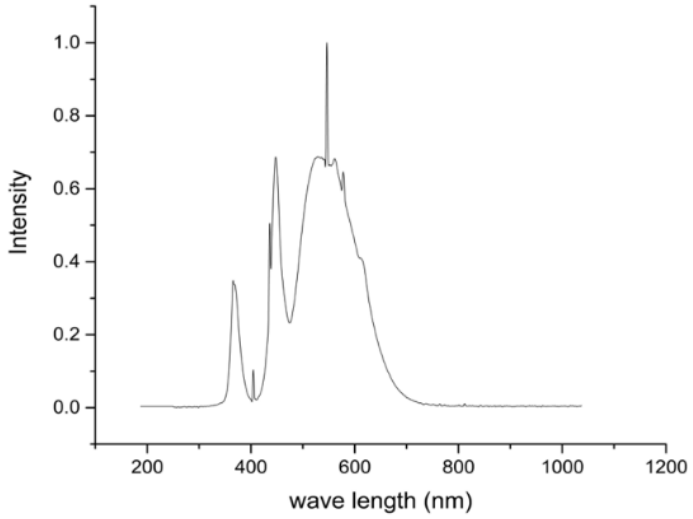

Figure S24. Combined spectral output of the photolysis station, achieved using a combination of LED panels and BL368 fluorescent tubes.

2.5 Headspace Gas Chromatography (GC) Station

The GC station comprised an Agilent GC (7890B) equipped with a Pulsed Discharge Detector from VICI and a headspace sampler (9697A, Extended Data Fig. 3d). Helium was used as the flow gas and the column oven was set to 90 °C. The Openlab software supplied by Agilent had no programmable interface and was therefore automated using simulated mouse-click and keyboard events using software supplied by CSols Ltd. Additionally, the robot was programmed to press a physical button on the headspace GC to park the gantry robot, thus using this interface in the same way that a human operator would (gantry robot is shown parked in the image in Fig. 30; see Supplementary Movie S3). This operation was required because the gantry robot blocks access to the sample rack holders on the head space sampler unless it is parked. The calibration curves for oxygen and hydrogen are shown in Figures S25 and S26.

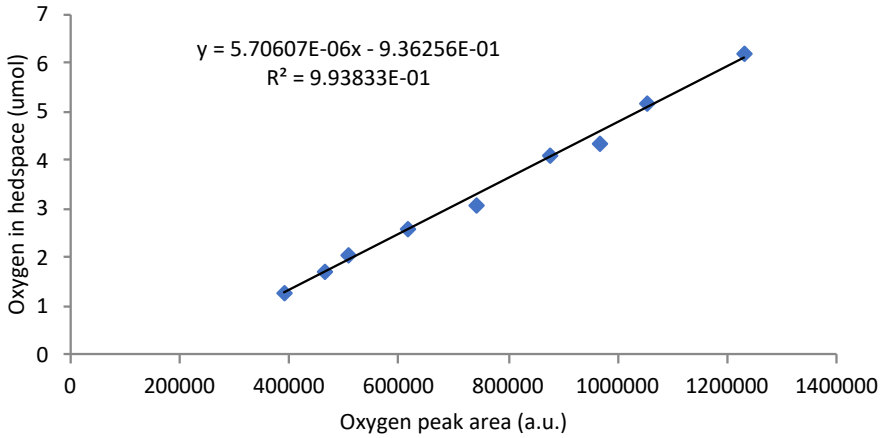

Figure S25. Oxygen calibration curve for headspace GC

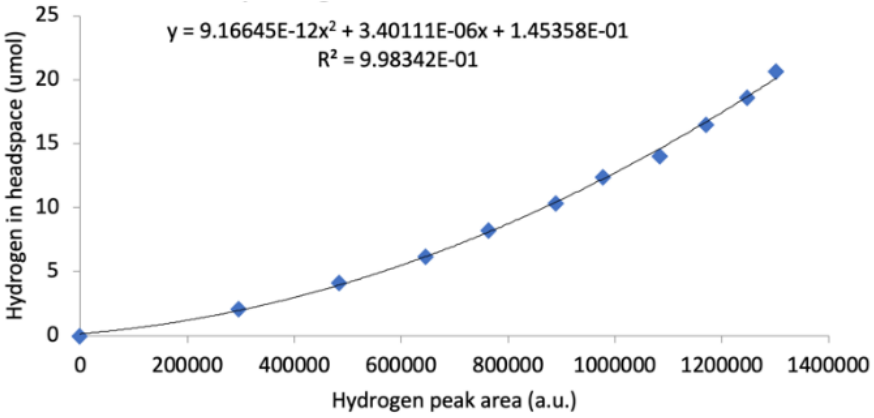

Figure S26. Hydrogen calibration curve for headspace GC

2.6 Input Stations

The workflow included three input stations (Figure S27) that allowed for a total of 39 racks (16 samples per rack) to be stored. These input stations were used both for supplying unloaded sample vials to the workflow and for storing completed racks after GC analysis.

The software system keeps track of the position of each rack holder / batch in the workflow. The system assigns the oldest empty rack in the system to new batches. The human operator is responsible for loading stock, including racks and empty vials, into and out of the workflow. Given the capacity of the input stations, the robot can operate autonomously for multiple days between loadings. The system uses the knowledge rack position and status to alert the operator by text message when the number of empty sample vials is running low, or if there is no more storage space for completed samples. A similar alert system is used to flag when the workflow runs out of liquids or solids, which are tracked automatically as they are used.

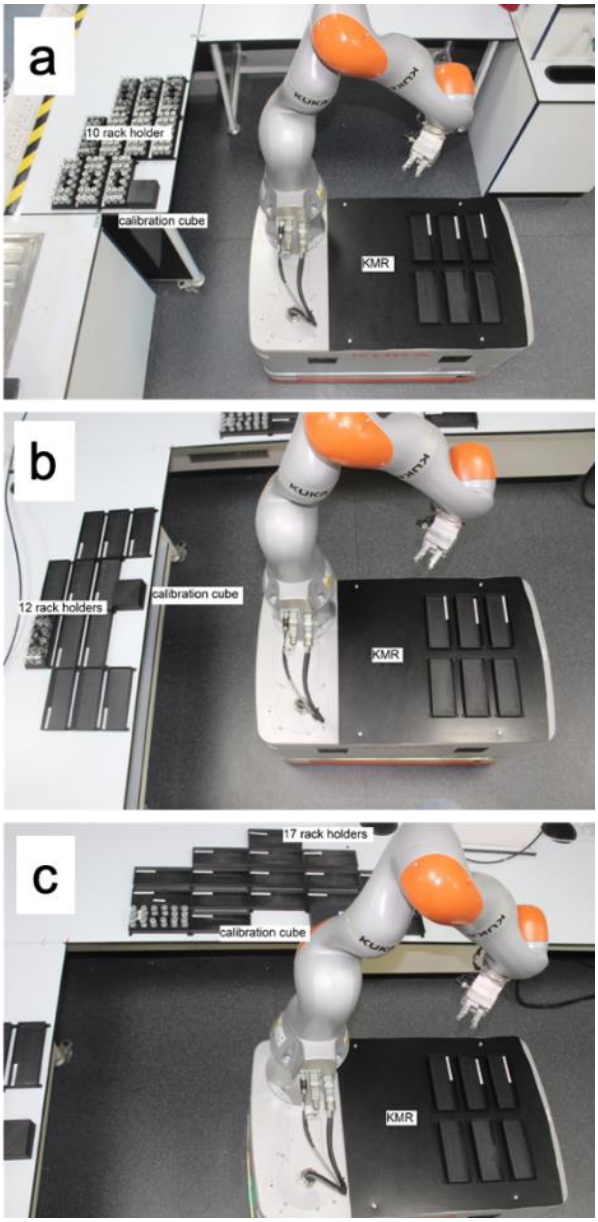

Figure S27. Photographs showing the robot parked at (a) Input Station 1; (b) Input Station 2, and; (c) Input Station 3.

2.7 Integration of Stations and Process Management System

The autonomous workflow comprised a process management system (PMS) module, an optimizer module, the robot, and the experimental stations described in the sections above. Each of these modules communicated with the server either directly or indirectly (Figure S28). At the center is the PMS software module, which contains all of the process logic for controlling the labware. The PMS communicates with all labware using various communication protocols, including the mobile robot (TCP/IP over WIFI), the GC (TCP/IP over LAN), various RS-232-controlled labware (temperature probe; Quantos solid dispenser; capping station; scale), and bespoke labware build using Arduino processors via RS-232 connections (photolysis station, sonicator and liquid pumps).

Initiation of an action for a piece of labware (*e.g.*, the Quantos solid dispenser) is started by sending a RS-232 command and finished when the labware sends a command that the action

has been completed. Solid dispensing takes a finite amount of time, but the same approach was also used for instantaneous actions, such as switching a relay. This was particularly important for solid dispensing, where the dispensing time can vary significantly (0.5–3 minutes per sample), and hence the use of a fixed-duration time delay is inefficient. The PMS sends commands to the Quantos dispenser to control both the sample carousel and the dispensing mechanism.

In the case of the liquid dispensing station, the PMS controls the both the pumps and scales (which weigh the liquids as they are dispensed). The PMS monitors the weight of the vial from the scale and controls the speed of the pumps (all using RS-232 commands), depending on the input received from the scales (see section 2.2 for more details on this process).

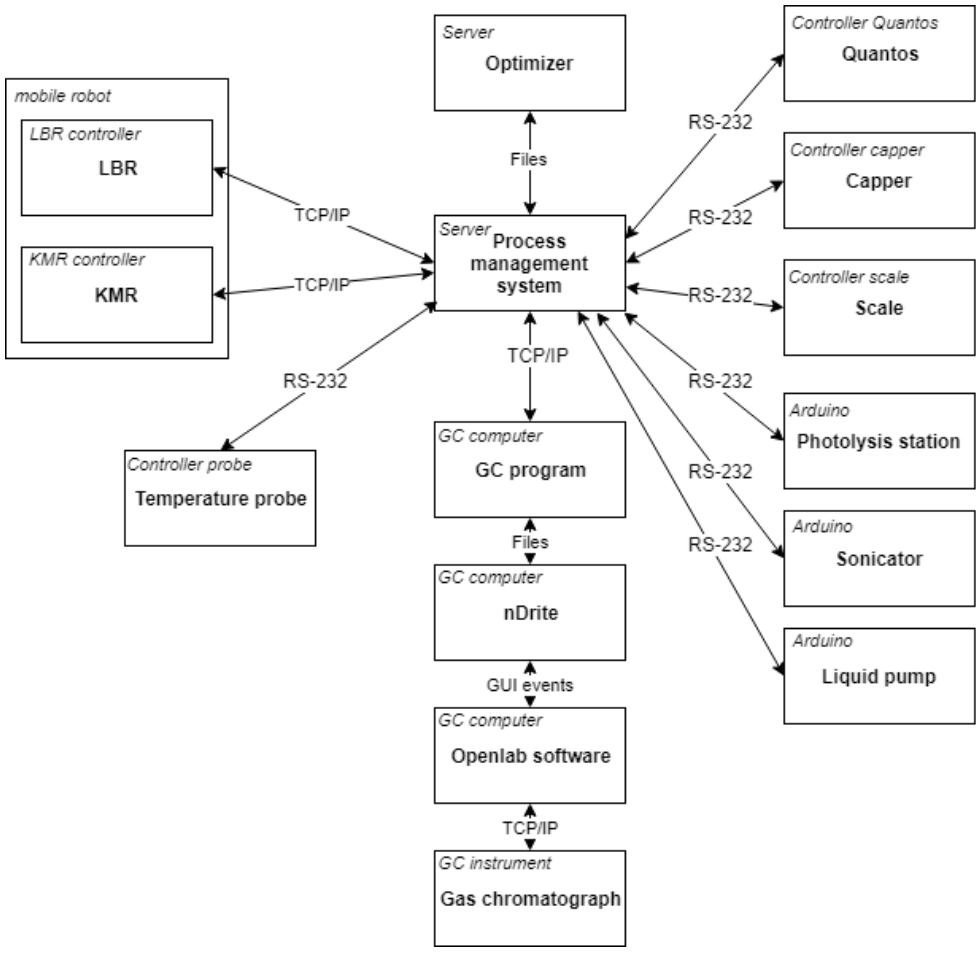

Figure S28. Overview of the hardware and software modules and communication protocols. Each node represents a module in the autonomous workflow; the text in italics (top-left corner of nodes) specifies the module controller. Interconnecting lines detail the communication protocols used between the various modules

3. Workflow Benchmarking

3.1 Baseline

To assess the baseline performance and sample-to-sample variability of the workflow, nine identical batches were run (16 samples per batch). First, six batches were ran in sequence over the course of 36 hours, and then one week later, another three batches were ran over the course of 24 hours. All samples were prepared with 5 mg of the photocatalyst, P10, in 5 % v/v TEOA, inertized for 60 seconds with N_2, and photolyzed for 60 minutes. The results are shown in Figure S29. Over the nine batches, an average hydrogen evolution of 17.84 µmol was observed with a standard deviation of 0.58 µmol and a coefficient of variation of 0.033. Similar standard deviations and coefficients of variation were observed within batches; these data are shown in Table S1.

Table S1. HER vs batch number for baseline experiment

Batch number	Average HER (µmol)	Standard deviation HER (µmol)	Coefficient of variation HER
1	18.25	0.69	0.038
2	17.91	0.49	0.027
3	17.98	0.54	0.03
4	17.94	0.65	0.036
5	18.04	0.45	0.025
6	17.57	0.46	0.026
7	17.7	0.6	0.034
8	17.24	0.4	0.023
9	17.87	0.39	0.022

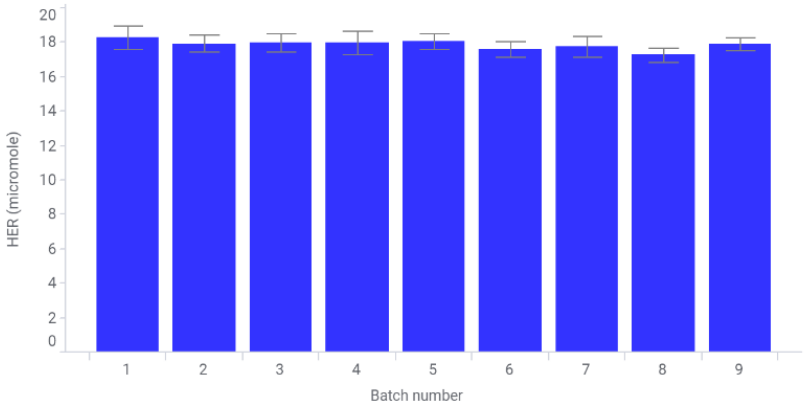

Figure S29. Average hydrogen evolution rate (HER) for repeat baseline experiments using P10 and TEOA as a hole scavenger, showing good batch-to-batch reproducibility.

3.2 Kinetic Tests

The robot can also carry out kinetic experiments, if required, by photolyzing identical samples for different amounts of time. To do this, the robot loads 16 identical samples into the photolysis station and removes one sample every 5 minutes. Kinetic experiments were performed using P10 (5 mg) in either TEOA 5% v/v (Figure S30, upper plot) or L-cysteine 10 g L^{-1} (Figure S30, lower plot). Linear kinetics were observed in both cases. These results also show that there is no background signal (*e.g.*, from photocatalysis that results from ambient light contamination), as indicated by the fact that the HER was zero at t=0 for all experiments. Here, this was achieved by painting the outside of the vials with an opaque black paint. Alternatively, it is also possible to use standard, clear glass vials and to run the experiment in the dark (Supplementary Movie S2).

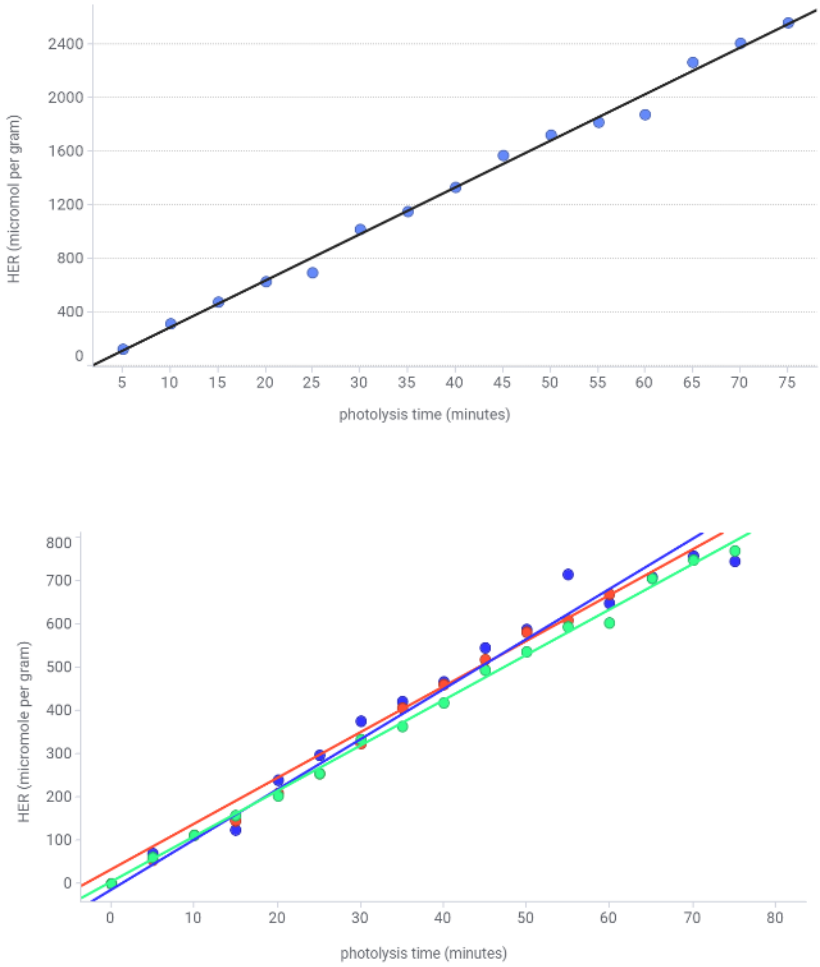

Figure S30. (Upper) Kinetic plot of hydrogen evolution versus photolysis time, as measured by the mobile robot (5 mg P10 in 5% v/v TEOA). The relationship between HER and photolysis time was fitted as a line with least squares. The equation of the line of best fit is given by y = -60.0 + 34.9x, with R2=0.994 and a mass-normalized rate (to the P10 catalyst) of 2092 µmol g^{-1} h^{-1}. (Lower) Kinetic plot of hydrogen evolution versus photolysis time, as measured by the mobile robot5 mg P10 in 5 mL 10 gL^{-1} L-cysteine). The relationship between the hydrogen evolution and the photolysis time was fitted with least squares. Batch 1 (red): y = 33.0 + 10.58x with R_2 = 0.970; mass-normalized rate = 635 µmol g^{-1} h^{-1}; Batch 2 (blue): y = -14.0 + 11.6x with R_2 = 0.995; mass-normalized rate = 695 µmol g^{-1} h^{-1}, and; Batch 3 (green): y = 4.53 + 10.5x with R_2 = 0.996; mass-normalized rate = 629 µmol g^{-1} h^{-1}.

3.3 Timescales for Steps in the Workflow

In this workflow, the robot prepares and analyzes samples in batches of 16. The modular nature of the workflow allows operations to be overlapped; for example, the robot can be dispensing solids for one batch while a second batch is being photolyzed and a third is being analyzed by GC. A summary of the times for the various operations, as averaged over 46 consecutive batches, is shown in Extended Data Fig. 6. The analysis time is rate limiting; it takes longer to analyze a batch of samples by GC (232 min) than the total time required for solid-dispensing, liquid-dispensing and photolysis of a batch (184 min). With this set-up, the robot can prepare and analyze 6.2 batches per day, on average.

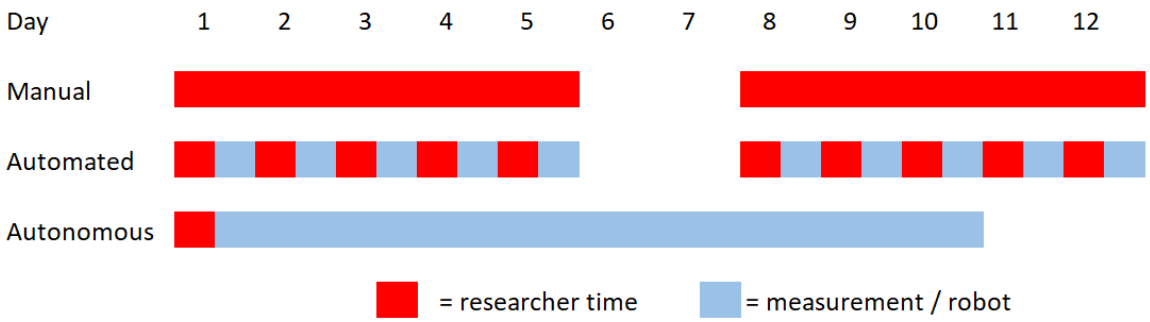

Figure S31. Comparison of the time needed to carry out 1000 photocatalysis experiments using different methodologies, assuming a 5-day working week. For the manual method, it is possible to carry out two experiments per day, requiring 500 days of researcher time (700 days total, assuming a 5-day work week). Using a semi-automated (but not autonomous) robotic method that we developed previously (ref. 33 in main text), we can carry out 100 experiments per day. This semi-automated method requires half a day of researcher time to set it up, so 1000 experiments requires 5 days of researcher time (12 days of experiment in total; see figure). The new autonomous approach developed here also requires a half-day set up, but then the experiment runs independently, 24/7, without any user input. As such, the autonomous workflow is 500–1000 times faster than manual approaches, and at least 10 times faster than semi-automated robotic methods based on dedicated researcher time. The true efficiency gains are larger than this because in the semi-automated method, the researcher would need to retrieve the results of the last batch of experiments and manually apply a process or algorithm to decide on the next batch. In the autonomous workflow, this is done (effectively) instantaneously. Furthermore, 100 experiments per batch is non-optimal from the perspective of batched Bayesian optimization, but a reduction in the batch size for the semi-automated experiment would reduce the throughput significantly. This is because the time saving in that method arises from batched photolysis and analysis of a large number of samples at one time (100), thus reducing the number of manual loading and unloading operations. In principle, we could operate the workflow illustrated in Supplementary Movie S3 at the same speed by using a human researcher to convey the samples from station to station, but we note that this would be an unappealing prospect (8 days of continuous work 24/7) and subject to user error (*e.g.*, the robot makes 4816 independent sample vial transfers during the 8-day experiment, all of which must occur in the correct sequence).

4. Sacrificial Hole Scavenger Screen

A range of candidate bioderived hole scavengers was tested, either at a concentration of 50 g L^{-1} or as saturated solutions for materials where the maximum solubility was lower than this (*i.e.*, L-asparagine, L-glutamine, L-histidine, L-methionine, L-threonine, leucine, and tyrosine). The robot was programmed to prepared samples with 5 mg of P10 in 5 mL of the stock solution; the samples were then photolyzed for 1 hour on the photolysis station and the hydrogen evolved was measured by GC. We applied conditional automation for any samples with a hydrogen concentration of more than 1 µmol, which were automatically retested 5 times by the robot. Most candidate donors produce little or no hydrogen (Extended Data Fig. 4), but L-cysteine emerges as a promising candidate, albeit with a hydrogen evolution rate that is lower than the benchmark petrochemical amine, triethanolamine (TEOA).

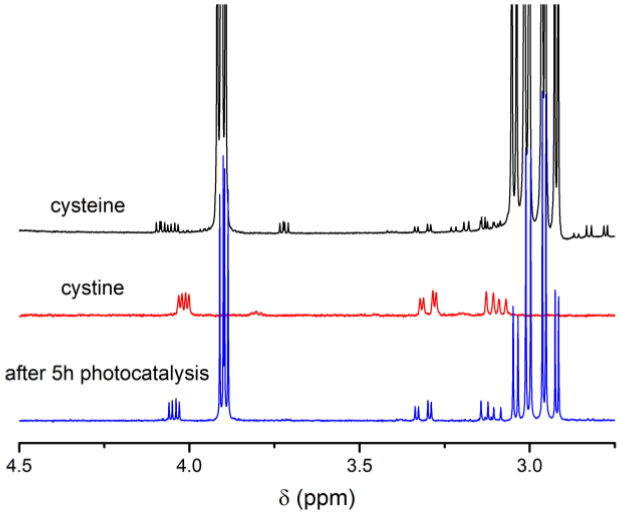

Figure S32. $_1$H NMR (D_2O) of L-cysteine, L-cystine and a P10-L-cysteine mixture after 5 h photolysis (pH = 7).

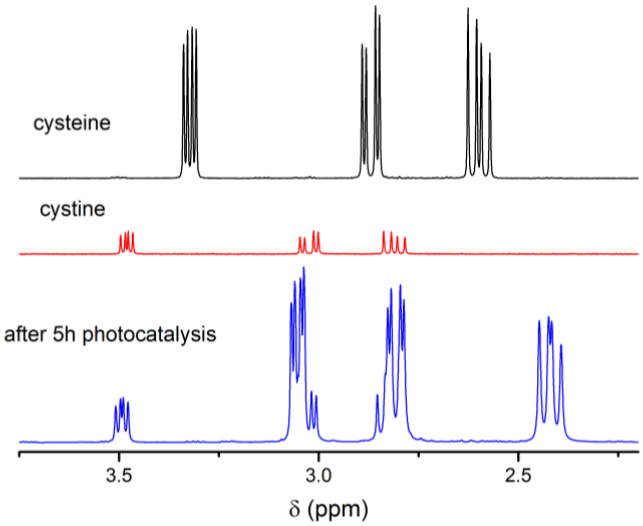

Figure S33. $_1$H NMR (D_2O) of L-cysteine, L-cystine and a P10-L-cysteine mixture after 5 h photolysis (pH = 11).

5. Baseline Control Samples

Two control experiments were included in each experimental batch to test the stability and reproducibility of the workflow because the optimization experiment would be running autonomously over multiple days. These control samples contained P10 (5 mg) and a 20 g L^{-1} solution of L-cysteine (5 mL). The average hydrogen evolution rate (HER) for these control experiments was 670 µmol $g^{-1}h^{-1}$ based on the P10 catalyst, with standard deviation 61 µmol g^{-1} h^{-1}. These data are normally distributed (Figures S34 & S35), as indicated by a Kolmogorov-Smirnov test (KS = 0.05331, p = 0.96225).

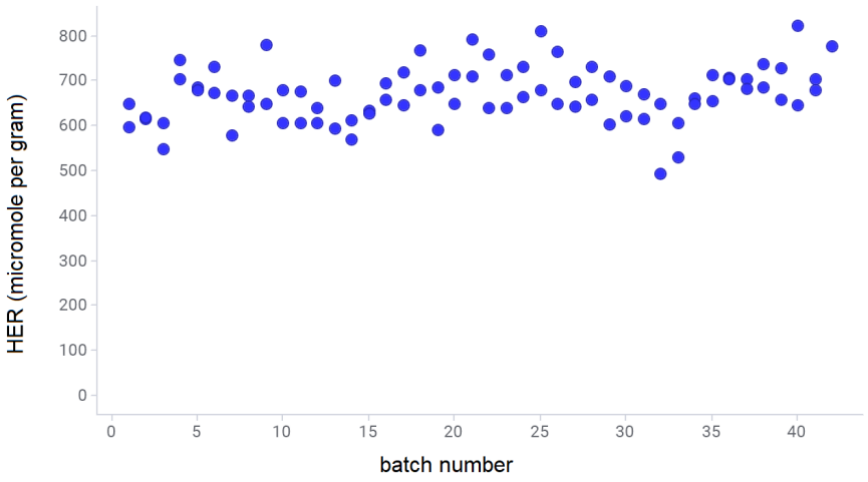

Figure S34. Hydrogen evolution rate (HER) for baseline control experiments performed throughout the autonomous search.

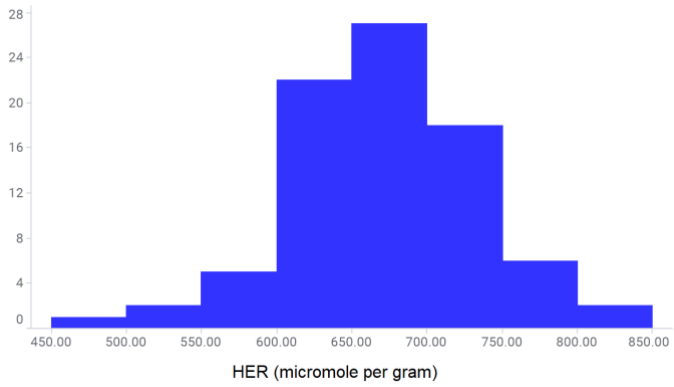

Figure S35. Histogram of HER for baseline control experiments (82 repeat measurements).

6. Control Experiments using other Light Sources

It is known that photocatalytic hydrogen evolution rates can be affected strongly by the choice of light source. To investigate the effectiveness of our home-built photolysis station, we carried out off-line manual tests using other light sources. Two quantitative measurements were carried out: one for a control solution of 25 mg of P10 dispersed in 25 mL 10 gL^{-1} L-cysteine and a second for the best solution arising from the autonomous search (25 mg P10, 30 mg NaOH, 1 g L-cysteine and 37.5 mg sodium disilicate in 25 mL water). For each experiment, a flask was charged with the P10 powder (25 mg) and 25 mL of the scavenger solution, and then sealed with a septum. The resulting suspension was ultrasonicated until the photocatalyst was dispersed before degassing by N_2 bubbling for 30 minutes. The reaction mixture was side illuminated with solar simulator irradiation (AM1.5G, classification ABA, ASTME927-10). Gas samples were taken with a gas-tight syringe and analysed using a Bruker 450-GC gas chromatograph equipped with a Molecular Sieve 13X 60-80 mesh 1.5 m × ⅛" × 2 mm ss column at 50 °C with an argon flow of 40.0 mL min^{-1}. The results are shown in Figure S36. The rate observed for the control experiment was 968 µmol g^{-1} and for the best, optimized solution it was 5448 µmol g^{-1} h^{-1}. The relative rate increase of the optimized solution versus

the control experiment, as measured using the solar simulator, was 5.6. This is consistent with the relative rate increase found during the autonomous robotic search (x 6.3).

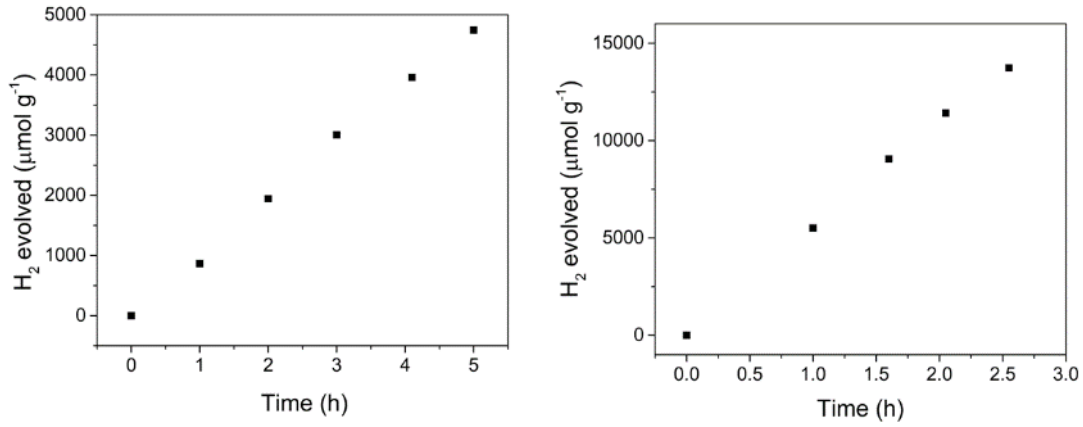

Figure S36. Photocatalytic hydrogen evolution for control experiment (left) and the best, optimized solution found during the autonomous search, as measured manually using a solar simulator for irradiation. This demonstrates that the findings from the autonomous workflow are transferable to a light source that more closely replicates natural sunlight.

7. *In Silico* Benchmarking using Virtual Experiments

To investigate the reliability and reproducibility of the optimization approach, a series of virtual experiments was performed on the calculated 10-dimensional hypersurface. Using the data collected during the autonomous experiment, we calculated a regression model using the same methodology as presented above. This model was then used to run virtual experiments by yielding computed hydrogen evolution rates to serve as the 'results' to feed the search algorithm. Each suggested point was perturbed to emulate instrumental precision by adding random noise to each result (between -0.5 and 0.5 μmol). We ran a virtual search using the same optimizer program, configuration, and parameterization as used in the physical robotic experiment (for results, see Figure 5 in main text). This virtual search was ran 100 times with different random starting conditions in each case. On average, 161 experiments were required to find solutions with 95% of the maximum HER. The benchmarking and regression model are available on GitHub.

8. Experimental Robustness

A high level of experimental robustness is required to carry out autonomous experimental searches that may run for multiple days. We evaluated the robustness of the workflow in detail as we were developing it over the period April–September 2019. To allow this, we implemented a protocol where all errors were logged in detail in a database, using CCTV footage to understand any errors that occurred while no operator was present (section 9; below). During the first 61 days of this period, 50 days were spent running experiments. In that 50-day period, a total of 3180 samples were measured. Over that 50-day period, the robot spent a total of 16 days carrying out manipulation tasks, 17 days idle (*i.e.*, not carrying out tasks, nor at the charging station), and 16.5 days either charging or waiting at the charging station. Overall, the robot spent 32.6% of its time in active operation over this 50-day period, and 4.5% in error. Figure S37 summarizes the errors that occurred during this 50-day period, shown as a pie chart.

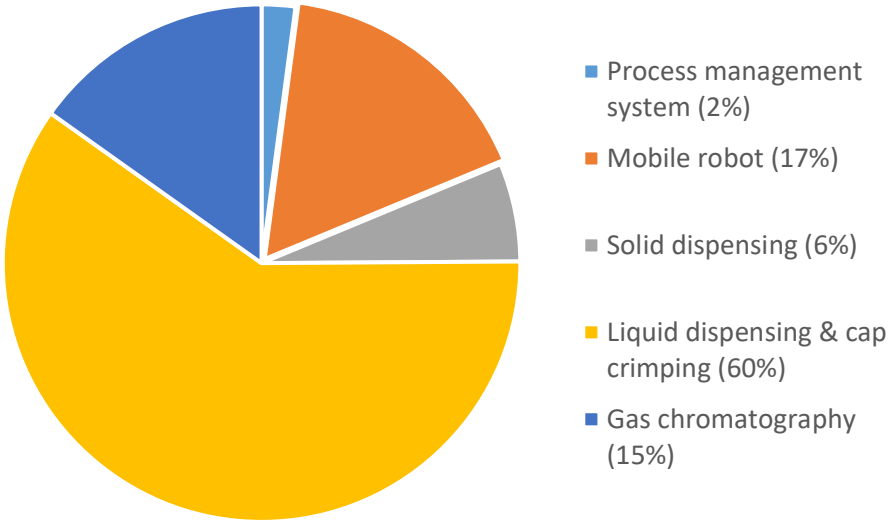

Figure S37 Percentage distribution of errors by time, based on running 3180 samples over the course of 2 months (April–May 2019). In this period, liquid dispensing and capping operations accounted for more than half of all errors, although this distribution evolved as the workflow (both hardware and software) was developed (Figure S37).

These error data, which were collected continuously, were used to make various improvements to the workflow, such as better positioning routines (*e.g.*, control of robot arm speed) and software improvements. We also introduced operator protocols, such as stricter pre-checks for faulty consumables (*e.g.*, malformed glass sample vials, etc.). This allowed the error rate to be steadily decreased over the course of 4 months of running the system, as shown in Figure S38. In July 2019, the average error rate was 0.5 %; effectively, one error per 200 samples. A significant number of these errors could be reset remotely by the operator without coming to the laboratory, since all stations were monitored by CCTV, thus allowing the operator to judge the nature of the failure, and whether it was safe to resume the experiment (see section **9** and Figure S39). As Fig. S38 shows, the single largest source of error in the July 2019 period was solid dispensing, which reflects its relative difficulty as an automated operation. The GC errors stem from limitations of the software control implementation and could be eliminated with further development. Based on this, we estimate that an error rate of 0.2 % (2 samples in 1000) might be achievable with additional workflow optimization, which could equate around 1–2 failures per week, on average, at the current level of sample throughput.

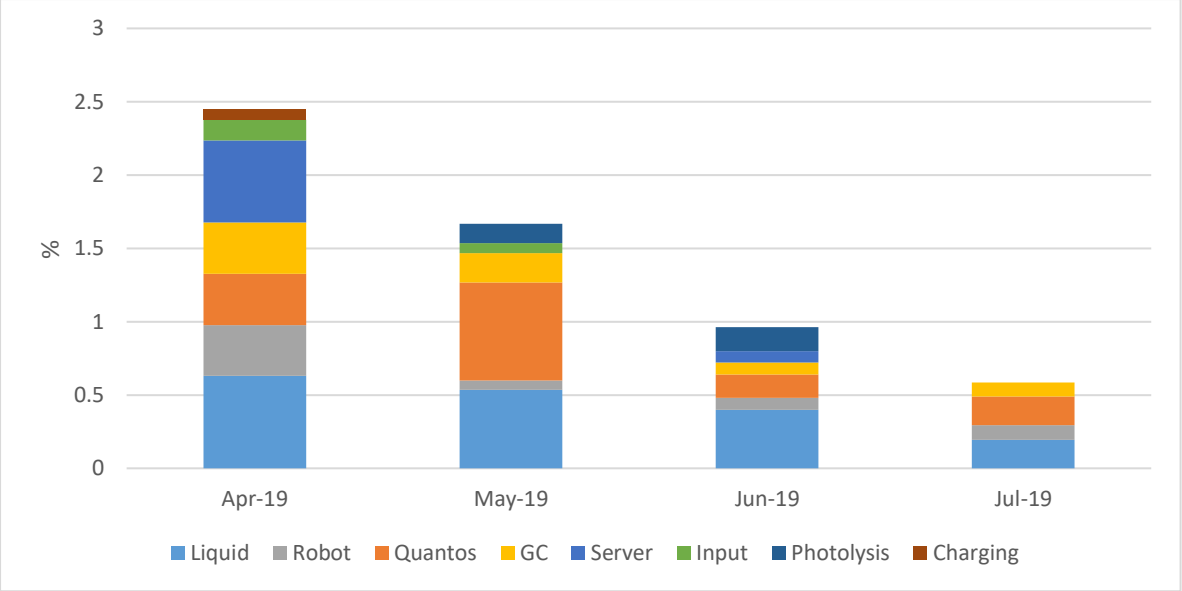

Figure S38. Percentage errors per sample in the workflow, categorized by error type, in the period April–July 2019. This plot is based on running 5201 samples over the course of these 4 months. Note that some errors, such as those connected to loading/unloading the input stations (green bars), server issues (pale blue bars), and the photolysis station (dark blue bars) were eliminated completely over this period. Some fraction of the remaining errors (*e.g.*, GC, solid dispensing) are inherent to the commercial instruments being used, rather than the robot or workflow. Note that the errors associated with the robot itself (grey bars) constitute a small fraction of the totals over this 4-month period.

9. 24/7 Monitoring

The workflow was monitored 24/7 using 8 CCTV cameras (Figure S39 shows examples). Footage was recorded continuously for all stations, allowing any errors to be analyzed, irrespective of when they occurred. In many cases, this allowed the workflow to be reset / restarted remotely without being in the laboratory, when it was deemed safe to do so. This also produced the detailed data required for to make continuous technical improvements to the workflow (section **8**).

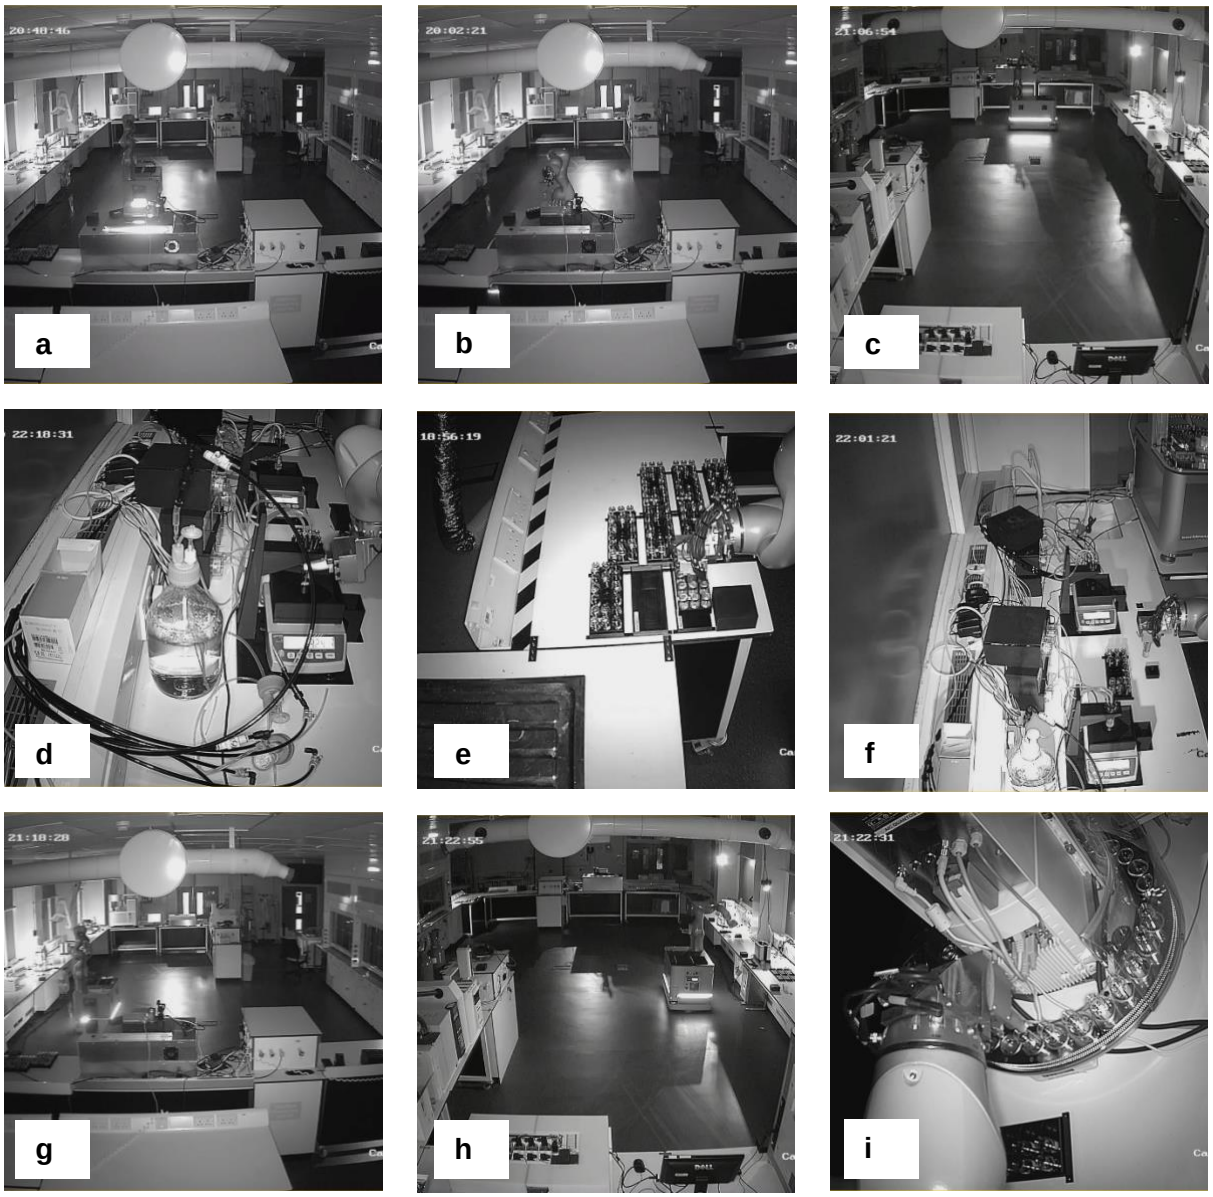

Figure S39. The workflow was monitored 24/7 using 8 CCTV cameras positioned around the laboratory. The images are screenshots from this CCTV monitoring system showing **a,** robot recharging its battery; **b**, robot loading the photolysis station; **c**, robot navigating away from the active photolysis station; **d**, robot loading the liquid station with a vial; **e**, robot placing measured samples into storage; **f**, robot preparing to load the capper/crimper; **g**, robot navigating across the laboratory; **h**, robot unloading the Quantos solid-dispenser; **i**, close-up of robot loading the Quantos solid-dispenser.

10. P10 Synthesis

P10 was synthesized in 20 mmol batches according to a modified literature procedure;[44] 3,7-dibromodibenzo[*b,d*]thiophene 5,5-dioxide (3.74 g, 10 mmol), 3,7-bis(4,4,5,5-tetramethyl-1,3,2-dioxaborolan-2-yl)dibenzo[*b,d*]thiophene 5,5-dioxide (4.68 g, 10 mmol), dimethylformamide (600 mL) and K_2CO_3 (aq) (2 M, 100 mL) were added to a dry flask and degassed by N_2 bubbling for 1 hour. Tetrakis(triphenylphosphine)palladium(0) (400 mg, 0.35 mmol, 3.5 mol %) was added and the mixture was degassed for a further 30 minutes. The reaction was refluxed at 140 °C for 48 hours under a nitrogen atmosphere. After cooling to room temperature, the mixture was poured into water (2 L) and stirred for 1 hour. The green solid was collected by filtration and washed with water (500 mL) and methanol (500 mL). The polymer was purified by Soxhlet extraction in methanol (1 day) and chloroform (2 days) before drying under vacuum. The polymer was ground before use to give a free-flowing green powder (3.99–4.24 g, 93–99 %).

11. Variable Correlation

The data from the 688-experiment autonomous search were analyzed for correlations between variables (Figure S40). There was a strong positive correlation between the catalyst, P10, and both NaOH and L-cysteine, and a weaker but significant correlation with sodium disilicate ($Na_2Si_2O_5$).

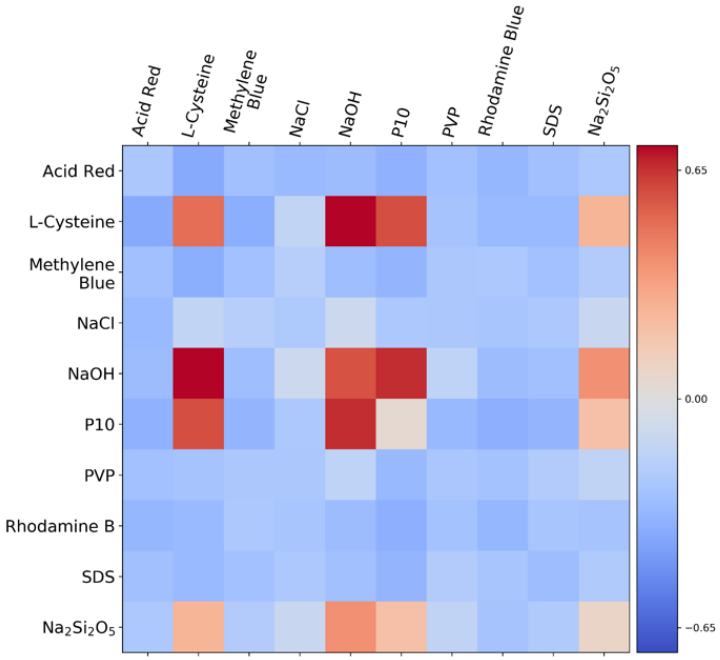

Figure S40 Matrix showing the Pearson correlation of the geometric mean of pairs of components with hydrogen evolution rate (HER).

12. Supplementary References

45 Williamson Stepper Motor Pumps, Product Code: 200-SMB-040-050, http://www.williamson-shop.co.uk/200-series-stepper-including-driver-for-bracket-mount-or-external-speed-control-7878-p.asp

45 Åström, K. J. & Hägglund, T. *PID controllers: theory, design, and tuning*. Vol. 2 (Instrument society of America Research Triangle Park, NC, 1995).