A closed loop discovery robot driven by a curiosity algorithm discovers proto-cells that show complex and emergent behaviours

Jonathan Grizou[1], Laurie J. Points[1], Abhishek Sharma[1] and Leroy Cronin[1]*

*[1]WestCHEM, School of Chemistry, University of Glasgow, Joseph Black Building, University Avenue, Glasgow G12 8QQ, U.K. *Corresponding author email: Lee.Cronin@glasgow.ac.uk*

We describe a chemical robotic discovery assistant equipped with a curiosity algorithm (CA) that can efficiently explore a complex chemical system in search of complex emergent phenomena exhibited by proto-cell droplets. The CA-robot is designed to explore proto-cell formulations in an open-ended way with no explicit discovery or optimization target. By applying the CA-robot to the study of multicomponent oil-in-water proto-cell droplets, we discovered an order of magnitude more instances of interesting behaviours than possible with a random parameter search. Among them, a formulation displaying a sudden and highly specific response to temperature was discovered. Six modes of proto-cell droplet motion were identified and classified using a time-temperature phase diagram and probed using a variety of techniques including NMR, which allowed the design of a payload release system triggered by temperature. This work shows how objective free search can lead to the discovery of useful and unexpected properties, with real-world applications in formulation chemistry.

The investigation of multicomponent chemical systems is difficult, as non-linear interactions originating at the molecular scale mean that an almost infinite number of experiments must be done to understand the system. Even an expert experimentalist requires enormous resources, often with the help of various design of experiments techniques, to make significant discoveries. Although lab automation has shown a remarkable increase in experimental throughput,[1] the rate of discovery has decreased as human intuition has been taken out of the feedback loop.[2,3] A viable solution would be to combine the curiosity and

knowledge inherent in scientific researchers together with the reliability and throughput of a robotic system. The automation of the laboratory experiments is important because it could dramatically improve the rate of scientific advances and new discoveries by designing experiments uninfluenced by preconceived ideas, bias or risk-averse tendencies.[4,5] Brute force automation has increased the number of experiments researchers could perform, but has not achieved the expected increase in new discoveries.[6] To address this, statistical methods were introduced to analyse the vast quantities of data being generated,[7,8] and recently machine learning methods have started to be integrated into laboratory equipment to better optimize predefined metrics.[9–11] However, these methods require prior knowledge of the specific desired phenomenon, and progress has stagnated because the autonomy of these systems remains limited.[2]

Herein, we describe an experimental method which implements state-of-the-art curiosity algorithms (CA) into a laboratory robot (CA-robot, Figure 1). CAs have been developed to replicate curiosity-driven learning in humans[12,13] and make use of knowledge acquired from developmental psychology, neuroscience, artificial intelligence and robotics.[14] CAs have been shown very efficient at exploring unknown systems[15–17] because they are designed to actively and autonomously select experiments that maximize the number of new and reproducible observations. Our CA, called random goal exploration,[15,18] decides on a target observation to try to generate a target goal rather than deciding directly experiment parameters to test (select different oil ratios). To understand this approach consider the analogy with to learning to play golf or a driving range for the first time, with no tuition. With each shot, you have a new target. You can vary how you hit the ball and with what club etc. Every time you play a shot, you learn from how it went, and apply that knowledge to your future shots. The 'targets' / holes are randomly positioned. Our system works in this way and starts by doing an experiment assuming uniform set of results (droplets of the same speed) will be observed in the physical system. The CA will generate a self-determined goal on the basis of the previous results observed (e.g. aiming for fast droplets

in large numbers). To attain that goal, the CA-robot must select experimental parameters to be tested. To do so, the CA-robot uses a dataset containing the results of the previous experiments observed to generate a picture or model of the process. The model is then used to infer the experimental parameters that are most likely to generate the CA-robot's self-determined goal. The selected experiment is then undertaken leading to a new observation which is added to the dataset of experimental observations. The CA-robot repeats this process for a given number of iterations, each time improving the quality of the model hence helping drive the exploration further.[15]

We tested our approach on dynamic oil-in-water droplets – promising protocell models[19,20] displaying an astonishing range of life-like behaviours, including movement, division, fusion and chemotaxis.[11,20–23] Although these droplets are thought to be driven by Marangoni instabilities originating from surface tension asymmetry,[24] to date, the understanding of even the most simple systems remains limited.[25,26] As such, oil-in-water droplets offer a great example of the challenges in studying complex and poorly understood systems where few components can lead to the emergence of a range of complex properties or behaviours, a topic of great relevance across many industries. Our CA-robot can perform droplet experiments, record and analyse the droplets' behaviours, and select the next experiments in full closed-loop autonomy (Figure 1, Supplementary Movie 1). Our aim was to explore the range of behaviours the droplets can exhibit with the minimum time and experimental budget possible. We first studied how our CA-robot compares with a random parameter search, a standard method used in high-throughput approaches. The droplet behaviours defining our observation space were chosen as the average droplet speed and the number of droplets, both selected due to their inherently interesting nature and similarity to the behaviours of simple lifeforms that can move and replicate. While these specific droplet behavioural metrics were relevant in this context, the methodology and principles applied herein are not specific and could apply to many other metrics or systems. Each exploration run consisted of 1000 droplet

experiments performed at 27.0 ±0.7 °C *(mean*±std*)*. Both the 'curious' and 'random' algorithms were tested three times.
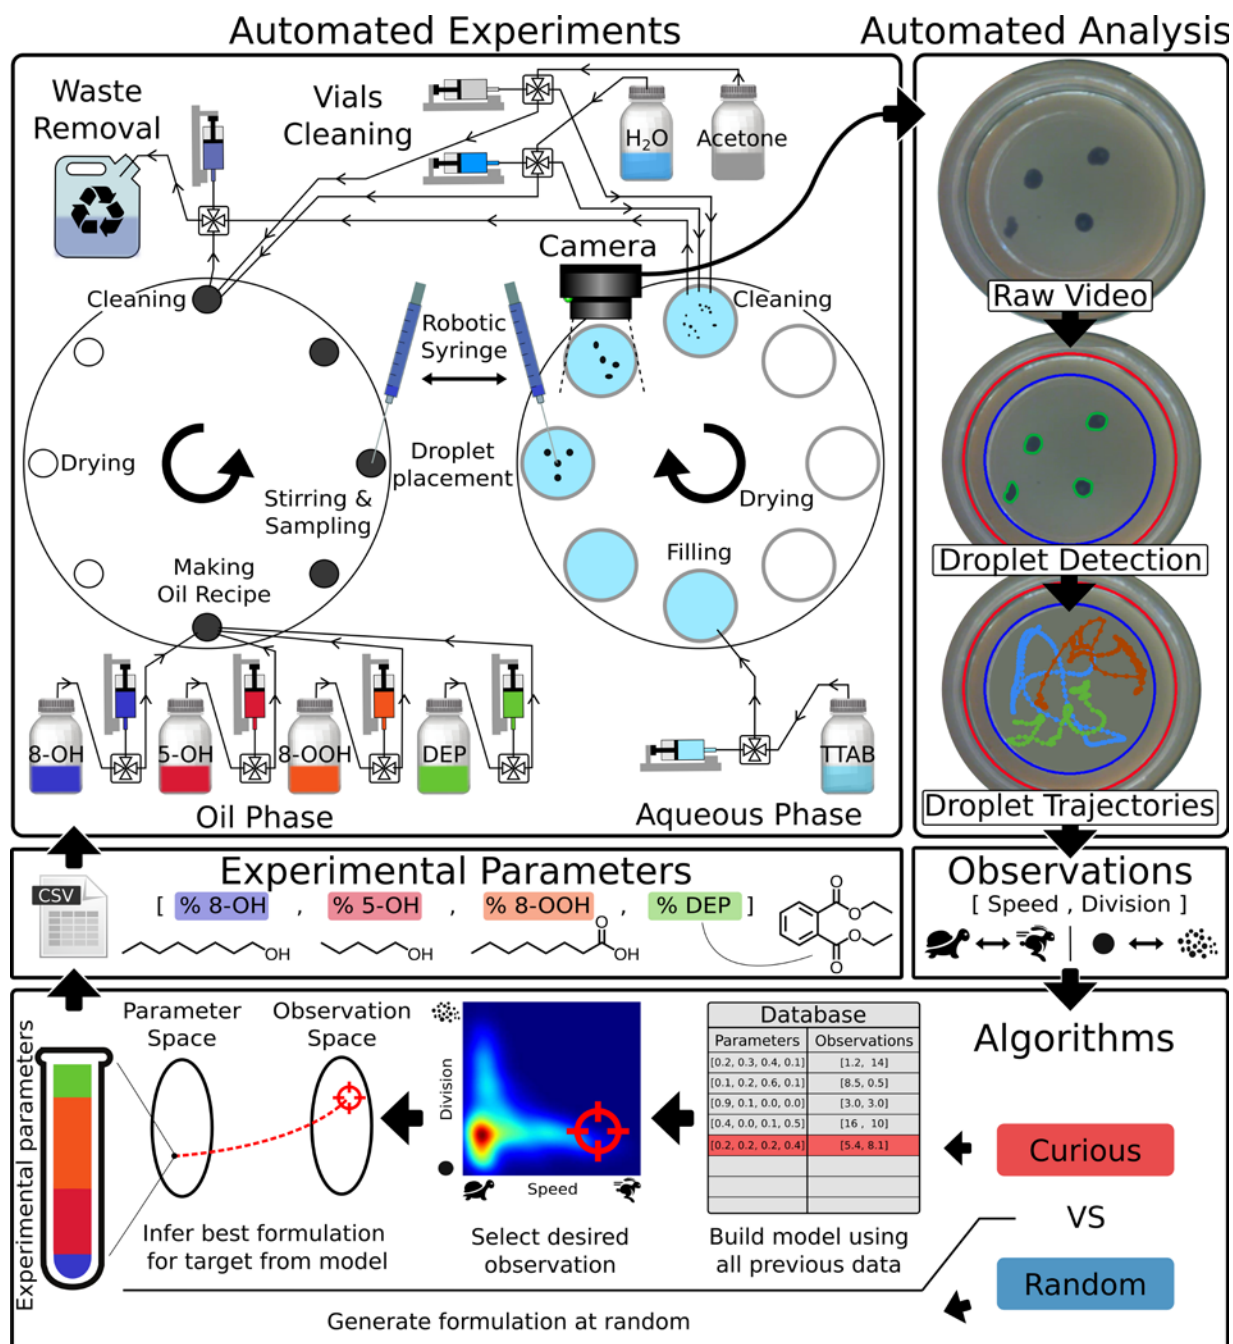

Figure 1. Diagram of the closed-loop workflow of the discovery robot driven by child-like curiosity. Top-Left: Schematic of new high-throughput droplet generating robot developed for this work. The robot runs the experiments by first mixing the oils accordingly, then prepares the aqueous phase and places droplets in the petri dish using a syringe. The motion of the droplets video is recorded and analysed. Once the experiment is completed, the platform cleans the entire system. Top-Right: Droplet contours and positions are extracted from the video data. Middle-Right: From the trajectories, the average speed and number of droplets generated per experiment was

determined. Bottom: The 'curious' algorithm collates all the observations and uses these to build a model and then selects a desired observation (a goal) and uses the built model to infer the experimental parameters most likely to achieve it. The 'random' algorithm generates experimental parameters uniformly at random. Middle-Left: Experimental parameters are the proportion of each oil composing our droplets, which are then used by the platform to perform the next experiment.

Given the same budget of 1000 experiments, the CA-robot generated significantly more varied droplet behaviours than the random parameter search (Figure 2 - notice the higher speed and division of droplets observed using the CA versus the random methodology). We quantified this exploration (Supplementary Information 2.2.3) and found that the CA enables us to observe 73.4 ±15.2% of the total observable space, *ca.* 3.3x more (p=0.039 - Welch's t-test) than a random parameter search (22.5 ±2.1%) within the same experimental budget. Strikingly the number of active droplet experiments observed (speed > 3mm.s^{-1}) is as low as 28.7 ±0.9 for random parameter search but jumps to 395.0 ±16.5 for the CA, a 14-fold improvement (p<0.001), without explicitly asking the robot to generate high speed experiments. Interestingly, after only 128 experiments the CA-robot already generated more varied experiments than random parameter search did in 1000 experiments (Figure 2A), a 7-fold efficiency gain in time and resources given the same hardware setup. Supplementary Movie 2 illustrates the exploration over time using both the CA and random; notice how even after as few as 50 experiments the CA driven exploration is already identifying more extreme cases of droplet behaviour, and this differentiation only increases as more experiments are undertaken. This is further visualised in Supplementary Movie 3, which shows videos of the 1st, 10th and 50th highest speed recipes from the two approaches. Furthermore, the CA-robot was able to, in a reproducible way, discover a delicate and never reported before temperature effect which was not observable using random parameter search given our experimental budget. Figure 2 compares the distribution of the speed of droplet experiments generated by both algorithms at 22.6 ±0.5°C and 27.0 ±0.7°C. The *ca.* 4.4°C temperature change has a significant impact on the observations made using the CA (395.0 ±16.5 vs 93 ±43.1 active droplets, p=0.005) whilst a negligible change is observed with

random parameter search (28.7 ±0.9 vs 19.3 ±7.6 active droplets, p=0.22). The vast differences seen with

small temperature changes are illustrated in Supplementary Movie 4.
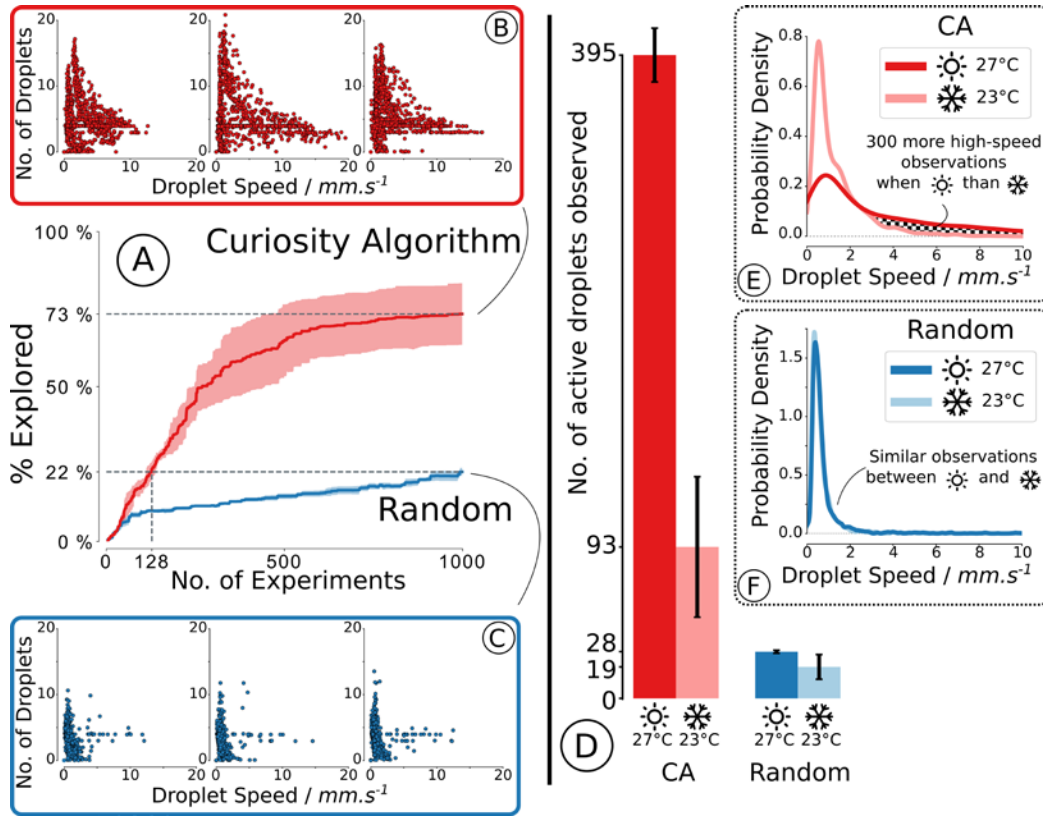

Figure 2. A summary of the results generated using our CA-robot, illustrating how the CA enables both significantly greater exploration of the behavioural space and the discovery of temperature sensitivity of the droplets. Left: Comparison of the observed droplet behaviours after 1000 individual experiments for CA and random – average of 3 repeats with shaded area showing 68% confidence interval. (A) Evolution of the percentage of the behaviour space explored between the two methods. CA explored 3.3 times more within the same experimental budget (73% vs 22%) and generated as diverse observations as random after only 128 experiments – a 7-fold reduction in time and financial cost for equivalent results. (B, C) Visualisation of the observations made by each method for each repeat; each scatter dot represents the average speed and number of droplets for a single 90 second droplet experiment. CA (B) leads to much more observations of rare and interesting droplets than random (C). Right: Effect of temperature (22.6 ±0.5°C vs 27.0 ±0.7°C) on the observations made using each algorithm. (D) Number of droplet experiments observed with a speed faster than 3 mms^{-1} for each method and temperature with error bar showing standard deviation. The CA-robot, by performing the same number of experiments, generated 14 times more interesting droplet recipes than random at 27.0°C (395 vs 28, p<0.001), and 5 times more at 22.6°C (93 vs 19, p=0.13). A change of only ca. 4.4°C led to a large and significant difference in the observed droplet behaviours when using the CA (395 vs 93, p=0.005). This difference in effect could not be significantly observed when using random (28 vs 19, p=0.22). This is confirmed by (E) and (F) which show the distribution of observation respectively for CA and random. (E) The distribution of observations has a strong tail indicating a wider exploration from the CA-robot, and there is a significant difference between observations made at 27.0°C and 22.6°C that is not observable with random (F). By focusing on the output space, the CA-robot provides a more accurate picture of the system for the same experimental budget, which allowed the discovery of this delicate temperature effect.

To study this newly observed effect in detail, we ran targeted droplet experiments within the range of temperatures accessible in the room (17-30°C). There were significant, unexpected and non-linear variations in the behaviour of the droplets of different compositions due to temperature (Supplementary Information 1.5). Such variations were highly reproducible as, for a given recipe, the observation of droplets' behaviour is enough to infer the room temperature with high accuracy (prediction error of 0.05 ± 0.66°C – Supplementary Information 1.6), a testament to both the reproducibility of the droplet behaviours and the existence of a delicate temperature effect. This is rather striking given the complexity of the system, the timescale of an experiment and the relative simplicity of our video-based analysis. One recipe of interest (composed of 1.9% octanoic acid, 47.9% DEP, 13.5% 1-octanol, and 36.7% 1-pentanol) was further analysed to probe the causes behind these observations via longer (15 minute) droplet experiments at a range of temperatures (Supplementary Information 1.8). Surprisingly, as shown in Figure 3A and Supplementary Movie 5, the droplets were seen to exhibit two peaks in their speed-time profile – they accelerate to achieve a first maximum speed, decelerate, and then accelerate again to reach a second maximum speed. The temperature effect on droplet motion can clearly be seen in the variation of their speed profile, with the peak speed timing and magnitude exhibiting clear trends with temperature, with the peaks occurring earlier and with a greater magnitude for hotter experiments.

Utilising droplet displacement data we identified six clear stages of droplet motion: initiation, fluctuation, irregular, deceleration, continuous and saturation, of which characteristic examples may be seen in Figure 3G (P1 to P6). During the initiation stage, the droplet vibrates around a point, showing little locomotion and low speeds. During fluctuation, these vibrations extend and the droplet speed increases before peaking during irregular motion, in which the droplet moves short distances in alternating directions. This is followed by a deceleration stage, during which the droplets slow down and display smoother motion, which then develops into continuous motion, during which concerted movement is seen and resulting in a more circular motion of the droplets around the dish. Eventually the saturation stage is

7

reached, in which the droplets slow down again and come to a halt. The peak speeds are observed for the irregular (purple) and continuous (orange) modes of motion, with the deceleration (green) period existing in between these two. A temperature-time phase diagram was derived showing the times at which each distinct phase of motion occurs at different temperatures (Figure 3C, Supplementary Information 1.9). The temperature-time phase diagram was created by calculating the intercept between cumulative distance travelled plots and linearly-fitted transition times (Figure 3B). The phase-transition times were each defined by characteristic points in the droplet acceleration-time plots. This phase diagram highlights the strong temperature dependence on the duration of each of the phases of motion and can be used to predict the mode of droplet motion observed at any time or temperature within the studied range.
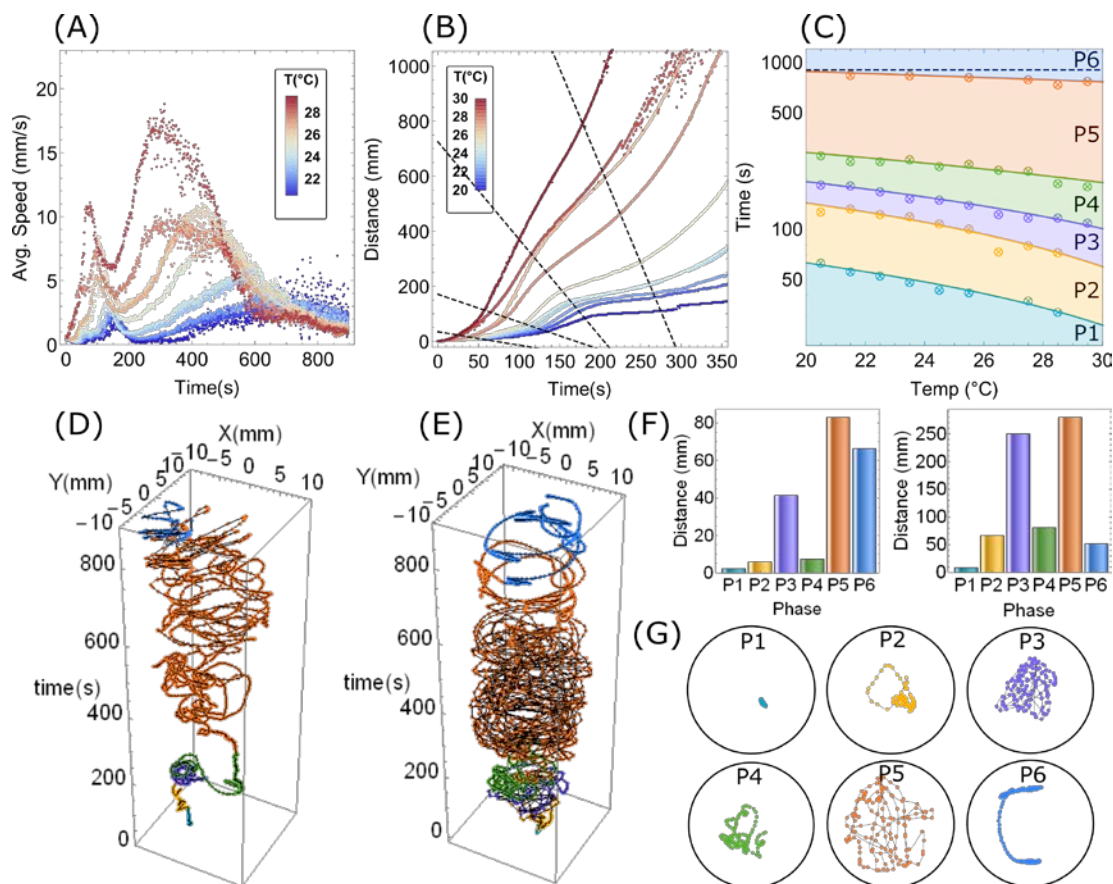
Figure 3. A summary of the analysis undertaken on a focus recipe, which resulted in the classification of 6 phases of droplet motion and the production of a time-temperature phase diagram. A) Temperature dependence of droplet speed vs. time. Each colour represents all experiments consisting of four droplets undertaken in a given temperature interval of 1°C. B) Temperature dependence of droplet cumulated distance

moved vs. time. The black dashed lines show the phase transitions in droplet motion that are used to estimate the phase diagram and are calculated by linear fitting of maxima and minima in the acceleration profile at each temperature interval. C) Temperature-Time Phase diagram of droplet motion showing different phases, initiation (P1), fluctuation (P2), irregular (P3), deceleration (P4), continuous (P5) and saturation (P6). The marked data points correspond to the intercepts shown in (B). D, E) The trajectory of a single droplet at 21.44°C (D) and 27.39°C (E), with different motion phases highlighted by colour. G) Exemplar 36s segments of each phase of motion, with each point showing the droplet location every 0.25 s at 27.39°C (E). Each example trajectory contains the same number of points to emphasise the differences in distance covered during the different phases, which is quantified in the cumulative distance per phase plots (F) for the droplet trajectories seen in (D - left) and (E -right).

Oil dissolution into the aqueous phase is hypothesised to play a major role in the observed droplet behaviours,[24,26] with oil dissolution impacting the interfacial tension, leading to droplet motion induced by Marangoni instabilities. We utilized a previously reported ^{1}H NMR spectroscopic method[22] to quantify the aqueous phase oil concentration during droplet motion at 22.4 ±0.2°C and 27.7 ±0.2°C (Supplementary Information 1.10). A 5°C temperature increase is seen to accelerate the dissolution of all oils (Figure 4).
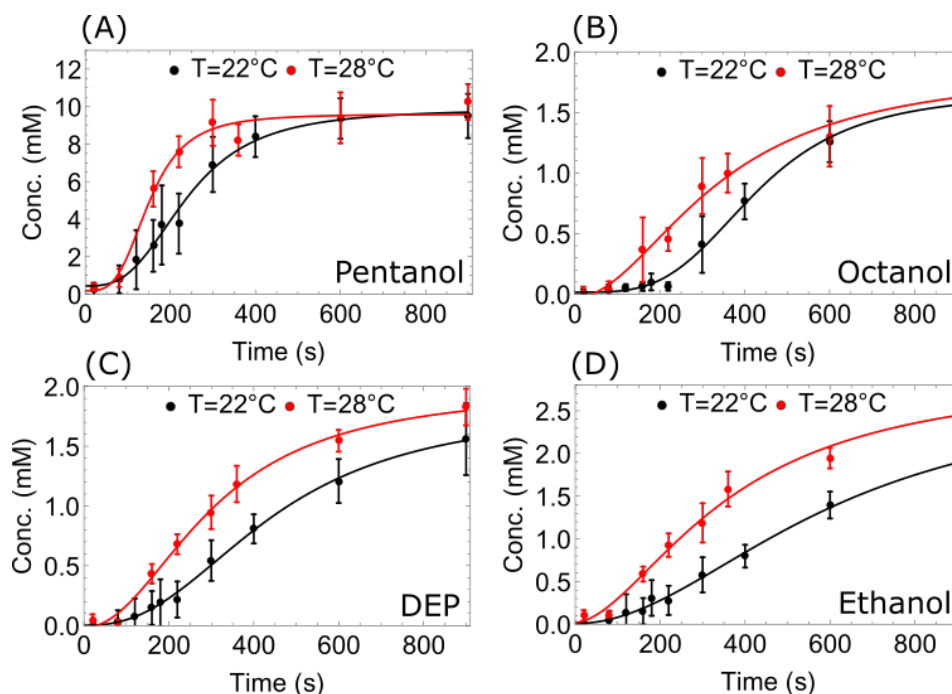

Figure 4. Oil concentration in the aqueous phase over time at 22°C (black) and 28°C (red), as quantified by ^{1}H NMR spectroscopy. Note how each oil dissolves faster at the higher temperature, whilst DEP and ethanol also dissolve to different final concentrations. Note differences in y-axis scale – pentanol dissolves around 5x more than the other oils. When regulated to a target of 22°C, temperature at the experimental location was 22.4 ±0.2°C. When regulated to a target of 28°C, temperature at the experimental location was 27.7 ±0.2°C.

Pentanol dissolves fastest and to the greatest level, as expected by its relative solubility. Octanoic acid dissolves to a fixed level early in the experiment and then stays constant; this is unsurprising due to its low concentration in the formulation and the fact it will rapidly deprotonate at high pH. As previously reported,[22] we note the presence of ethanol due to the base catalysed hydrolysis of DEP. Interestingly, DEP and ethanol have different final concentrations at the different temperatures, as temperature affects the equilibrium of the hydrolysis reaction, as opposed to only physical processes driving the other oils dissolution. Octanol, DEP and ethanol dissolution are delayed as compared to pentanol dissolution, suggesting that pentanol dissolution is the main contributor to the first peak of droplet motion.

To confirm this hypothesis, we compared the oil dissolution rates with the droplet motion data, as shown in Figure 5A and B and detailed in Supplementary Information 1.11. The rate of pentanol dissolution is seen to be rapidly increasing during the fluctuation and irregular phases, before rapidly decreasing during the deceleration phase. This indicates that pentanol dominates the early stages of droplet motion, and that its dissolution is the primary cause of the fluctuation and irregular forms of motion. As pentanol dissolves so fast in these early stages, it is not surprising that the motion is sporadic, as rapid dissolution in all directions (Figure 5C) prevents the initiation of structure, regular flows and a more continuous form of motion. Because pentanol dissolution has largely ceased by the time of the continuous phase of motion, whilst the other oils are still dissolving to significant levels, it appears that DEP/ethanol and/or octanol are the primary driving force of the continuous period of motion. We hypothesise (Figure 5D) that the more gradual rate of dissolution during the continuous phase of motion allows a positive feedback loop to be setup between oil motion, dissolution and Marangoni flows.[27,28] As the droplet moves in this phase, it advects 'fresh' surfactant solution onto its anterior face (via collision with empty micelles and free surfactant molecules) and leaves a trail of oil filled micelles in its wake (via oil dissolution). Thus, the interfacial tension is higher at the posterior face as there are more oil filled micelles and less free surfactants in this zone. As there is an interfacial tension differential between the anterior and posterior

faces of the moving droplet, a Marangoni flow is induced, supporting the forward direction of motion, providing a positive feedback loop for continued forward motion. This hypothesis is also supported by the observation that droplets often avoid following the recent path of other droplets. When the oil dissolution rates begin to saturate, the continuous motion slows and stops.
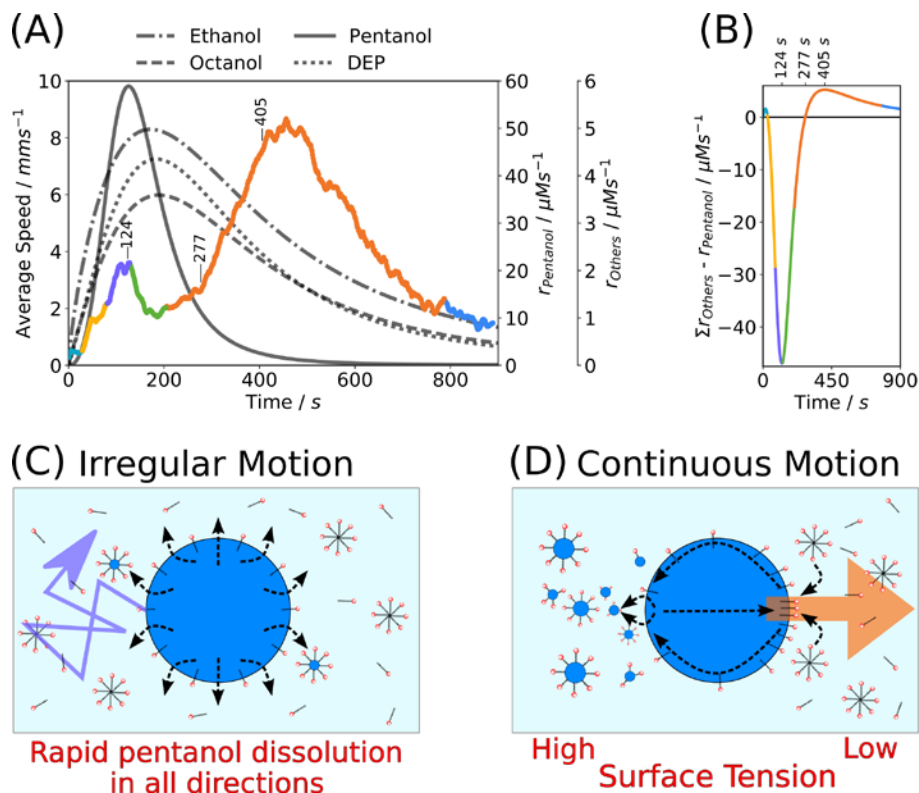
Figure 5. The correlation between oil dissolution and droplet behaviours and schematics illustrating the proposed mechanisms for the irregular and continuous phases of motion. A) The average oil droplet speed (coloured plot, left y axis) observed at 28°C (average across 8 experiments processed via a 10 seconds moving average), with the colour corresponding to the phase of motion (cyan-initiation, yellow-fluctuation, purple-irregular, green-deceleration, orange-continuous and blue-saturation). The grey lines illustrate the rates of oil dissolution (right hand y axes) from the fitted [1]H NMR spectroscopy dissolution data. B) The difference between the sum of the rates of DEP, ethanol and octanol dissolution and the rate of pentanol dissolution against time. Note the peak difference in favour of pentanol at 124 seconds, the point at which the rates are equal at 277 seconds and the peak difference in favour of DEP, ethanol and octanol at 405 seconds. These times are also marked in (A) and correlate closely with the irregular-deceleration transition, rapid acceleration in the continuous phase and the maximum droplet speeds in the continuous phase. C) Schematic illustrating the proposed mechanism for the fluctuation and irregular phases of motion. Rapid pentanol dissolution in all directions (black arrows) into a largely oil free aqueous phase containing many empty micelles and free surfactants leads to no concerted directional motion, but rather erratic motion in various directions (purple arrow). D) Schematic illustrating the proposed mechanism for the continuous phase of motion. At this time, total oil dissolution is slower. The front of the moving droplets contacts 'fresh' aqueous phase whilst the rear of the droplet leaves a trail of 'filled' micelles. Thus, the interfacial tension is lower at the front of the droplet leading to a positive feedback loop of forward motion via Marangoni flows.

11

We cannot ascertain from the previously discussed data whether it is DEP, ethanol and/or octanol dissolution that is the primary cause of the continuous phase of motion. To discriminate between these, we varied the pH of the surfactant containing aqueous phase, which had a significant impact on the oil droplet behaviour. As the pH and temperature are increased, DEP hydrolysis is significantly accelerated,[29] leading to an earlier and larger second continuous motion peak (Supplementary Figure 36). With increasing pH there is also a 10^6 fold increase in ionic strength, significantly reducing the aqueous solubility of alcohols,[30] thus lowering the dissolution of pentanol and reducing the irregular motion peak. These results together indicate that DEP hydrolysis is the primary cause of the second movement peak and continuous phase of motion. A range of experiments in which the pentanol-octanol ratio, the alcohol chain length and the number of droplets placed were varied further confirmed the links between pentanol and the first speed peak and DEP and the second speed peak (Supplementary Information 1.12).
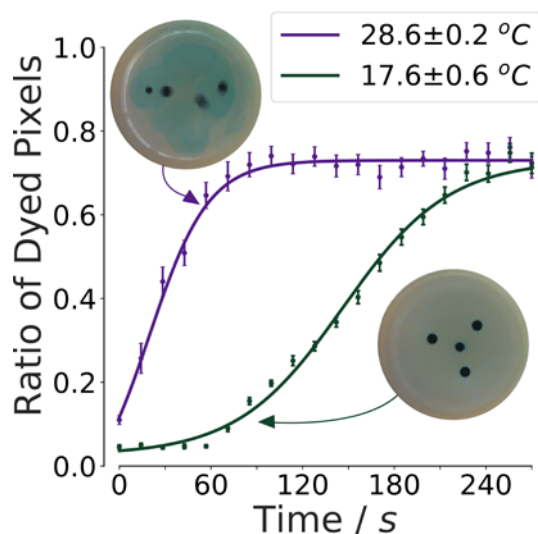
Figure 6. A plot illustrating how temperature sensitive oil droplets can be used as a targeted release system for the dye methylene-blue, as measured by the ratio of blue pixels during a droplet video. With a temperature increase of 11°C, complete dye release occurs ca. 2.5 times faster.

Knowing the chemical and physical mechanisms driving our droplet motion and their temperature dependence, we investigated a potential application of this discovery for temperature controlled chemical release. The dye methylene blue was found to be released approximately 2.5 times faster at 28.6 ±0.6°C

than at 17.6 ±0.2°C after 20 repeats using the studied recipe. (Figure 6, Supplementary Movie 6). This is a proof of concept that droplets such as these could be used as containers with temperature dependent release for active molecules such as dyes, drugs or catalysts.

By equipping a droplet generating robot with a curiosity algorithm (CA-robot), we were able to discover droplet systems with interesting properties, for example the temperature sensitivity of a specific oil-in-water droplet formulation. We showed that, given the same experimental budget, this temperature effect could not have been observed using a random parameter search. This illustrates that CA-robots can be of significant advantage to reveal properties of unknown systems in a fully automated fashion. Using physical and chemical analysis, we characterized the discovered effect and derived a phase diagram of droplet motion through time and temperature which links to the underlying oil dissolution processes. This chemical analysis revealed the astonishing complexity that underlies the dynamics of our 4-component oil-in-water droplet system. Finally, we harvested this property to design a chemical payload release system triggered by temperature. This is the first time a curiosity algorithm has been used for the exploration of a physical system, which differs from previous work because it is not an optimization or a hypothesis disambiguation problem, both of which require prior knowledge to select a metric to optimize or to generate hypotheses to be disambiguated. Future research will focus on constructing the observation dimensions autonomously from the droplet videos in an unsupervised way,[18] as in this work the observation space was designed by the authors which potentially introduces human bias that can limit possible discoveries.

Methods

Robotic Platform

We designed a high-throughput droplet-generating robot (Figure 1) that can execute and record a 90 s droplet experiment every 111 s, including mixing, syringe driven droplet placement, recording, cleaning

and drying. Such minimal overhead time was achieved by parallelizing all operations enabling our platform to routinely perform 300 droplet experiments per day in full autonomy. The platform and sequence of operations are fully described in Supplementary Information 2.1.

Droplet Chemistry

The oil-in-water system comprises four droplets composed of a mixture of four oils placed onto a surfactant containing aqueous phase in a petri dish [11]. An experiment consists of preparing a formulation of octanoic acid, diethyl phthalate (DEP), 1-octanol and 1-pentanol at a specific ratio determined by the algorithm and dyed with 0.5 mgmL^{-1} of Sudan Black B dye. The oil mixture is sampled by the robot using a 250μL syringe and delivered as 4 x 4μL droplets in a Y pattern from the center of a 32mm petri dish filled with 3.5mL of a 20mM cationic surfactant (myristyltrimethylammonium bromide, TTAB) solution raised to a high pH (ca. 13) using 8gL^{-1} NaOH. The droplet making procedures are fully described in Supplementary Information 2.3.1.

Image Analysis

The droplet activity is recorded at 20fps for 90 seconds and analyzed using computer vision. Droplet contours are extracted using a thresholding algorithm and tracked through frames using a proximity rule. The droplets' average speed and the average number of droplets in the dish (droplets can split, fuse or leave the tracking area) are quantified and used as the observation space. The droplet tracking procedures are fully described in Supplementary Information 2.1.5.

Algorithmic Implementation

Experimental parameters are generated as a 4-dimensional vector representing the ratio of each oil in the droplet mixture. Observations are represented as a 2-dimensional vector representing the average speed and average number of droplets in an experiment. For the random goal exploration algorithm, the forward model is built uniquely from previous observations using locally weighted linear regression and the

14

inverse model is solved for each goal using the CMA-ES algorithm on the learnt forward model. The CA implementation is fully described in Supplementary Information 2.2.

Data availability

Due to the large total size of the droplet videos (> 500GB of data), the experimental data used in this work are available upon request to the corresponding author at lee.cronin@glasgow.ac.uk.

Code availability

The code used to operate the robotic platform, generate and analyse results are available online in our group GitHub account at https://github.com/croningp and are fully described in the Supplementary Information.

ASSOCIATED CONTENT

Supplementary Information

The Supplementary Information Appendix contains further results and discussion including more detail on related work, an in-depth comparison of the algorithms and a detailed explanation of the physicochemical analysis undertaken, the modelling of droplet behaviour and the phase diagram preparation. Additional experiments are presented studying the sensitivity of our system to pH, proportion of each oil, chain length of alcohol used and the number of droplets place in the dish, as well as detail given on the dye release and droplets as temperature sensors experiments. The Supplementary Information Appendix also provides detailed information (and the relevant GitHub repositories) about the materials and methods including the full droplet robot design and code, the droplet tracking implementation, a formal description of the curious algorithm and its implementation, and the experimental procedure related to the chemical analysis. Finally, the Supplementary Movies are listed along with their explanatory captions.

AUTHOR INFORMATION

Corresponding Author

Correspondence should be addressed to LC (lee.cronin@glasgow.ac.uk).

Author Contributions

L.C. conceived the original idea and J.G. and L.C. together designed the project and the research plan; J.G. designed and built the robotic platform and implemented the computer vision and the algorithms. J.G. and L.J.P. performed droplet experiments. L.J.P. designed and ran the chemical analysis. J.G., L.J.P, A.S. respectively analyzed the algorithm, chemical and droplet motion data. A.S. extracted the phase diagram. L.J.P. designed and tested the chemical payload release system and J.G. quantified the release from the videos. J.G., L.J.P., A.S., and L.C wrote the paper.

REFERENCES

1.   MacArron, R. *et al.* Impact of high-throughput screening in biomedical research. *Nature Reviews*

Drug Discovery **10**, 188–195 (2011).

2.   Bleicher, K. H., Böhm, H. J., Müller, K. & Alanine, A. I. Hit and lead generation: Beyond high-throughput screening. *Nat. Rev. Drug Discov*. **2**, 369–378 (2003).

3.   Kodadek, T. The rise, fall and reinvention of combinatorial chemistry. *Chem. Commun*. **47**, 9757-9763, (2011).

4.   Häse, F., Roch, L. & Aspuru-Guzik, A. Chimera: Enabling Hierarchy Based Multi-Objective Optimization for Self-Driving Laboratories. (2018) Preprint: chemrxiv:6195176.

5.   King, R. D., Costa, V. S., Mellingwood, C. & Soldatova, L. N. Automating Sciences. *IEEE Technol. Soc. Mag.* **37**, 40–46 (2018).

6.   Greenaway, R. L. *et al.* High-throughput discovery of organic cages and catenanes using computational screening fused with robotic synthesis. *Nat. Commun.*, **9,** 2849 (2018).

7.   Buchanan, B., Sutherland, G. & Feigenbaum, E. A. Heuristic DENDRAL - A program for generating explanatory hypotheses in organic chemistry. *Mach. Intell.* **4**, 209–254 (1969).

8.   Goodacre, R., Vaidyanathan, S., Dunn, W. B., Harrigan, G. G. & Kell, D. B. Metabolomics by numbers: Acquiring and understanding global metabolite data. *Trends in Biotechnology* **22**, 245–252 (2004).

9.   Assion, A., Baumert, T., Bergt, M., Brixner, T., Kiefer, B., Seyfried, V., Strehle, M. & Gerber, G. Control of Chemical Reactions by Feedback-Optimized Phase-Shaped Femtosecond Laser Pulses. *Science* **282**, 919–922 (1998).

10.   Reizman, B. J., Wang, Y.-M., Buchwald, S. L. & Jensen, K. F. Suzuki–Miyaura cross-coupling optimization enabled by automated feedback. *React. Chem. Eng*. **1**, 658–666 (2016).

11.   Gutierrez, J. M. P., Hinkley, T., Taylor, J. W., Yanev, K. & Cronin, L. Evolution of oil droplets in a chemorobotic platform. *Nat. Commun*. **5**, 5571 (2014).

12.   Oudeyer, P.-Y., Kaplan, F. & Hafner, V. V. Intrinsic Motivation Systems for Autonomous Mental Development. *IEEE Trans. Evol. Comput*. **11**, 265–286, (2007).

13.   Gottlieb, J., Oudeyer, P.-Y., Lopes, M. & Baranes, A. Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends Cogn. Sci*. **17**, 585–593 (2013).

14.   Cangelosi A & Schlesinger M. *Developmental Robotics: From Babies to Robots*. (MIT Press, 2015).

15.   Baranes, A. & Oudeyer, P.-Y. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Rob. Auton. Syst*. **61**, 49–73 (2013).

16.   Kulkarni, T. D., Narasimhan, K. R., Saeedi, A. & Tenenbaum, J. B. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. in *Advances in neural information processing systems* **29** 3675–3683 (2016).

17.   Lehman, J. & Stanley, K. O. Abandoning Objectives: Evolution Through the Search for Novelty Alone. *Evol. Comput*. **19**, 189–223 (2011).

18. Péré, A., Forestier, S., Sigaud, O. & Oudeyer, P.-Y. Unsupervised learning of goal spaces for intrinsically motivated exploration. in *The Sixth International Conference on Learning Representations* (2018).

19. Dzieciol, A. J. & Mann, S. Designs for life: protocell models in the laboratory. *Chem. Soc. Rev.* **41**, 79–85 (2012).

20. Lach, S., Yoon, S. M. & Grzybowski, B. A. Tactic, reactive, and functional droplets outside of equilibrium. *Chem. Soc. Rev.* **45**, 4766–4796 (2016).

21. Toyota, T., Maru, N., Hanczyc, M. M., Ikegami, T. & Sugawara, T. Self-propelled oil droplets consuming 'Fuel' surfactant. *J. Am. Chem. Soc.* **131**, 5012–5013 (2009).

22. Points, L. J., Taylor, J. W., Grizou, J., Donkers, K. & Cronin, L. Artificial intelligence exploration of unstable protocells leads to predictable properties and discovery of collective behavior. *Proc. Natl. Acad. Sci.* **115**, 885–890 (2018).

23. Hanczyc, M. M., Toyota, T., Ikegami, T., Packard, N. & Sugawara, T. Fatty acid chemistry at the oil-water interface: Self-propelled oil droplets. *J. Am. Chem. Soc.* **129**, 9386–9391 (2007).

24. Herminghaus, S. et al. Interfacial mechanisms in active emulsions. *Soft Matter* **10**, 7008–7022 (2014).

25. Liu, H. & Zhang, Y. Phase-field modeling droplet dynamics with soluble surfactants. *J. Comput. Phys.* **229**, 9166–9187 (2010).

26. Schmitt, M. & Stark, H. Marangoni flow at droplet interfaces: Three-dimensional solution and applications. *Phys. Fluids* **28**, 012106 (2016).

27. Maass, C. C., Krüger, C., Herminghaus, S. & Bahr, C. Swimming Droplets. *Annu. Rev. Condens. Matter Phys.* **7**, 171–193 (2016).

28. Jin, C., Krüger, C. & Maass, C. C. Chemotaxis and autochemotaxis of self-propelling droplet swimmers. *Proc. Natl. Acad. Sci.* **114**, 5089–5094 (2017).

29. Huang, J., Nkrumah, P. N., Li, Y. & Appiah-sefah, G. in *Reviews of Environmental Contamination and Toxicology* (ed. Whitacre, D. M.) **224**, 39–52 (Springer-Verlag, 2013).

30. Bates, C. M., Stevens, F., Langford, S. C. & Dickinson, J. T. Motion and dissolution of drops of sparingly soluble alcohols on water. *Langmuir* **24**, 7193–7199 (2008).