

Genome-scale transcriptional dynamics and environmental biosensing

Garrett Graham^{a,1}, Nicholas Csicsery^{a,1}, Elizabeth Stasiowski^{a,1}, Gregoire Thouvenin^{a,1}, William H. Mather^b, Michael Ferry^b, Scott Cookson^b, and Jeff Hasty^{a,b,c,d,2}

^aDepartment of Bioengineering, University of California San Diego, La Jolla, CA 92093; ^bQuantitative BioSciences, Inc., San Diego, CA 92121; ^cMolecular Biology Section, Division of Biological Sciences, University of California San Diego, La Jolla, CA 92093; and ^dBioCircuits Institute, University of California San Diego, La Jolla, CA 92093

Edited by Charles R. Cantor, Retrotope, Inc., Del Mar, CA, and approved December 24, 2019 (received for review July 28, 2019)

Genome-scale technologies have enabled mapping of the complex molecular networks that govern cellular behavior. An emerging theme in the analyses of these networks is that cells use many layers of regulatory feedback to constantly assess and precisely react to their environment. The importance of complex feedback in controlling the real-time response to external stimuli has led to a need for the next generation of cell-based technologies that enable both the collection and analysis of high-throughput temporal data. Toward this end, we have developed a microfluidic platform capable of monitoring temporal gene expression from over 2,000 promoters. By coupling the “Dynamics” platform with deep neural network (DNN) and associated explainable artificial intelligence (XAI) algorithms, we show how machine learning can be harnessed to assess patterns in transcriptional data on a genome scale and identify which genes contribute to these patterns. Furthermore, we demonstrate the utility of the Dynamics platform as a field-deployable real-time biosensor through prediction of the presence of heavy metals in urban water and mine spill samples, based on the the dynamic transcription profiles of 1,807 unique *Escherichia coli* promoters.

high-throughput microfluidics | dynamics | *E. coli* transcriptomics | explainable AI | biosensor

In model organisms, studying the changing patterns of gene expression in reaction to experimentally induced environmental perturbations often elucidates the underlying signaling network (1–4). However, present techniques to measure genome-wide expression data are often destructive in nature and offer only snapshots of a cell’s state (5–11). Meanwhile, an increasing body of evidence points to dynamics as a key way biological systems encode information (12, 13).

Microfluidics, coupled to time-lapse fluorescence microscopy, has served as a means to measure gene-expression dynamics in precisely controlled environments (14–16). Recently, several studies have demonstrated how microfluidic parallelization permits the simultaneous tracking of hundreds to thousands of strains of *Saccharomyces cerevisiae* (17, 18), *Escherichia coli* (19, 20), or mammalian cell lines (21). The combination of microfluidics and genome-scale fluorescent reporter strain libraries has facilitated the study of genomic transcriptional dynamics. Existing approaches, though, have been hampered by short experimental lifespans, limited temporal resolution, static environmental conditions, and the use of single-purpose devices.

In addressing these needs, we have developed Dynamics, a straightforward and broadly applicable research platform that combines multiplexed microfluidics, fluorescence microscopy, and deep neural network (DNN) and explainable artificial intelligence (XAI) algorithms to better resolve transcriptional dynamics at the genome scale. The Dynamics platform enables continuous growth, precise environmental control, and optical monitoring of 2,176 microcolonies of unique GFP-reporter *E. coli* for up to 14 d. In addition to demonstrating the platform’s utility in studying transcriptional dynamics through time series

and fold-change data, we show the platform’s effective application as a continuous biosensor for heavy metals in water supplies at environmentally relevant concentrations and conditions, using DNN algorithms to predict the presence of heavy metals and XAI algorithms to identify which genes contribute to these predictions, unraveling the “black box problem” typical of machine learning (Fig. 14).

Results

Microfluidic Device Development. The Dynamics microfluidic device was designed for straightforward experimental setup, reliable trap filling and cell retention, and optimal fluorescent signal from each spotted microcolony (Fig. 1 *B–D*). The single media inlet–outlet device requires only two fluidic connections after cell spotting and chip bonding. The media inlet channel feeds a total of 2,176 4-μm-tall cell traps. Trap shape and spacing allow a 6,144 Society for Biomolecular Sciences (SBS) density pin pad to deposit cells into the back of the trap, where they grow toward the tapered opening interfacing with 50-μm-tall minor media channels. These minor channels branch off of a larger 230-μm-tall major media channel manifold, which eliminates the possibility of cell trap cross-contamination. Once spotted cells

Significance

New technologies are needed for the global analysis of intracellular signaling networks that encode information in the time domain. We developed a microfluidic platform capable of culturing over 2,000 bacterial strains simultaneously and subjecting them to dynamical perturbation. We used an explainable artificial intelligence classifier to reveal insights embedded in the temporal transcriptional response of *Escherichia coli* exposed to heavy metal stress. This enabled real-time predictions of the presence of heavy metals in complex environmental samples.

Author contributions: G.G., N.C., E.S., G.T., W.H.M., M.F., S.C., and J.H. designed research; G.G., N.C., E.S., and G.T. performed research; G.G. and W.H.M. contributed new reagents/analytic tools; G.G., N.C., E.S., G.T., W.H.M., M.F., S.C., and J.H. analyzed data; and G.G., N.C., E.S., G.T., W.H.M., and J.H. wrote the paper.

Competing interest statement: W.H.M., M.F., S.C., and J.H. have a financial interest in Quantitative BioSciences. Quantitative BioSciences has an exclusive license to IP stemming from this work, which is owned by the University of California San Diego.

This article is a PNAS Direct Submission.

Published under the PNAS license.

Data deposition: Preprocessed, labeled machine-learning features, the corresponding library strain position records, and the relevant metadata for our experiments are available on the University of California San Diego Biodynamics Laboratory website (<http://biodynamics.ucsd.edu/downloads>). The code used to process data from Dynamics experiments and train machine-learning models is available on GitHub at <https://github.com/GarrettCGraham/dynamics.public>.

¹G.G., N.C., E.S., and G.T. contributed equally to this work.

²To whom correspondence may be addressed. Email: hasty@bioeng.ucsd.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1913003117/-DCSupplemental>.




Fig. 1. The Dynamics platform. (A) Fluorescent strain libraries are loaded onto large-scale microfluidic devices that can be fully captured in a single image using custom optics. Parallel cultures of *E. coli* are subjected to multiple exposures of different stimuli with time series and fold changes used to quantify responsive strains. Machine-learning algorithms are trained on preprocessed data to enable real-time stimulus detection. (B) Design of the Dynamics 2,176-strain microfluidic device with cell traps in red and media channels in blue and yellow. (C) Detailed schematic of four cell traps with arrows showing direction of media flow. (D) View of the full Dynamics chip. (E) Mean fluorescence (solid blue) and SD (shaded blue) of the *E. coli* *zntA* promoter driving GFP to repeated cadmium inductions (gray bars) with periods increasing from left to right (30 min, 2 h, 4 h, and 8 h).

have reached confluence, inducer compounds can be pulsed in at user-specified frequencies with the dynamic response of each strain measured down to a 4-min temporal resolution (Fig. 1E).

Screening for Responsive Promoters to Heavy Metals. Using the Dynamics platform with a previously developed GFP *E. coli* promoter library (22), 1,807 unique *E. coli* promoters were screened against nine heavy metals (Cu(II), Zn(II), Fe(III), Pb(II), Cd(II), Cr(VI), Hg(II), As(III), Sb(III)) at environmentally relevant concentrations (SI Appendix, Table S3). Screening experiments lasted 7 to 14 d, with cells exposed to a different heavy metal every 24 h. Promoters responsive to each metal can be identified through a combination of clustering and fold-change analysis. A high-level view of the 1,807 promoter time traces (Fig. 2A) and subsequent clustering (Fig. 2B) reveals distinct classes of transcriptional responses to a single 4-h zinc exposure. In Fig. 2B, clusters 1 and 2 include promoters that are up- and down-regulated, respectively, in the presence of zinc, but return to

baseline expression levels within 15 h after zinc removal. Clusters 3 and 4 include promoters that are up- and down-regulated, respectively, but with slower dynamics. Gene ontology (GO) enrichment analysis (SI Appendix, Fig. S10) suggests that from these four clusters, genes associated with cellular stress are up-regulated (cellular detoxification, cellular response to toxic substance, and antibiotic catabolic processes) while genes involved in differing metabolic and biosynthesis can be either up- or down-regulated.

Individual responsive strains for each metal were identified, based on their fold-change response (Fig. 2D) to daily 4-h metal exposures (Fig. 2C). Fold-change measurements highlight the promoters displaying the strongest response to each metal. Subsequent investigation of the most responsive strains (Fig. 2E) quantitatively elucidates dynamical properties, such as




Fig. 2. Dynamics as a screening tool for heavy metal responsive promoters in *E. coli*. (A) Fluorescence response of an *E. coli* promoter library during a 4-h 50-ppb Zn induction (dashed window). Each row represents the promoter activity, normalized between 0 and 1, of a single strain, with 1,995 total strains represented. Four clusters from agglomerative clustering are labeled on the right. (B) Four clusters of strains calculated from agglomerative clustering from the data in A. The mean (dark blue line) $\pm$ 1 SD (dark blue shading) of all strains in each cluster is plotted. The dashed window denotes when zinc was present. (C) Responsive strains over the duration of a Dynamics experiment. Normalized fluorescence for two strains is plotted over the duration of one experiment, with 4-h heavy metal inductions (gray bars) occurring once daily. (D) Fold change for top responding strains to all metals. Log₂ of the average fold change is shown for the top responding strains to each heavy metal. * $P = 0.05$, ** $P = 0.01$, *** $P = 0.001$, respectively. (E) Significant single-strain normalized fluorescence response (blue line) $\pm$ 1 SD (blue shading) across all inductions for a given metal (dashed window).

amplitude, relaxation time, and response speed, all of which are important factors for their use in the study of gene expression regulation and continuous biosensing applications. While many of the identified sensing strains, such as *zntA* (23) or *cueO* (24), have well-documented metal interactions (*SI Appendix, Tables S5 and S6*), others are less studied or poorly annotated, particularly members of *E. coli* “y-ome” (25). Overall, these analyses demonstrate the utility of this platform as a screening tool for dynamic environmental-response phenotypes in a strain library. However, specific metal discrimination based on fold change alone is difficult to interpret due to promoter nonspecificity, cross-talk, noise, and low-amplitude responses.

Machine Learning. To better discriminate between *E. coli*'s responses to the heavy metals used in our screening, we trained and tested two types of machine-learning models on the Dynamics data (26). The first model, known as extreme gradient boosted trees (XGBoost), is a popular decision tree ensemble-based classifier known for its ability to learn nonlinear models (27). The second one, known as a long short-term memory recurrent neural network (LSTM-RNN), is a DNN (28) selected because of its ability to effectively utilize sample sequence history to classify time series data, a property not shared by XGBoost.

Both classification algorithms outperformed random guessing of the majority class (no toxin) on the standardized experiments' feature set, with the LSTM-RNN performing the best overall (*SI Appendix, Figs. S8 and S9*). As seen by examining the diagonal elements in the confusion matrix in Fig. 3A, the LSTM-RNN was able to distinguish both biotic and xenobiotic metal-spiked water from pure water with a high level of reliability.

The LSTM-RNN found iron and copper to be easily detectable biotic metals, which is not surprising given their importance to *E. coli* cellular function (24, 29). Cadmium was the most readily detected xenobiotic metal with the LSTM-RNN classifier, although it was sometimes confused with zinc. *E. coli* are known to use the same sensing and transport systems to capture and export excess amounts of these two metals, which possess the same number of valence electrons (23, 30). Most classification errors occurred during the 10 to 40 min at the start or end of the induction periods, when the LSTM-RNN occasionally had difficulty determining the exact time that each metal was added or removed from the media (Fig. 3B). This is most

pronounced with the prediction of lead, for which the classifier incorrectly predicted no toxin for 48% of time points where lead was present. This is largely due to the weak promoter responses induced by 0.03 ppm lead, which is only double the Environmental Protection Agency (EPA) maximum contaminant level. In lead exposures with poor prediction, time points at the start of the 4-h induction window are misclassified as no toxin, while lead is accurately predicted near the end of this window (*SI Appendix, Fig. S10*). While past studies have used machine-learning frameworks to assign cells to chronologically distinct phenotypes based on their transcriptomes (31), we believe this is a different instance of a multiclass classifier successfully leveraging genome-wide transcriptional dynamics in live cells to predict exposure of a biological organism to an environmental stressor.

Machine-Learning Introspection Using Explainable Artificial Intelligence. At present, a major obstacle to making scientific conclusions from machine-learning results is the black box problem: As an algorithm's ability to model complex phenomena grows, its decision-making processes become more and more obscured from its operators (32). Recently, explainable artificial intelligence techniques have been employed to explain the decision making of machine-learning algorithms in the life sciences (33–35), while contributions from coalitional game theory have led to the development of a mathematically consistent method for understanding the decision-making process of any AI classifier (36, 37).

Taking advantage of these recent advances, we trained a Shapley additive explanations (SHAP) learner on both our XGBoost and LSTM classifiers (36, 38). The SHAP algorithm scores a strain's impact on the classifier's predictions by calculating Shapley values from cooperative game theory. Shapley values are the mathematically unique way to divide game payout between players who have collaborated with each other to achieve a common goal, assuming basic rules of fairness (39). A major advantage of SHAP is that Lundberg and Lee (36) demonstrated that it is an umbrella method that mathematically unifies several commonly used feature attribution frameworks, including LIME, Layerwise Relevance Propagation, and DeepLIFT. Viewing both SHAP values (impact on classifier output) and feature values (data fed to the classifier) with respect to time offers




Fig. 3. Machine learning on heavy metal exposures. (A) Confusion matrix showing the recall (true positive rate) of the LSTM-RNN classifier in predicting six metals across all experimental data (14,332 time points). (B) LSTM-RNN classifier applied to time series data for all six detectable metals in two different experiments. Both experiments have a row for the true media condition and the predicted condition. In the case of correct classification, the color in the predicted row would match the color in the top row, whose color represents which metal was actually present at that time point. An easy way to tell if there has been a misclassification is by seeing if there are any regions flagged with red below the predicted row. Red indicates time points where the prediction does not match the ground truth. (C) Feature (blue) and SHAP (orange) time trajectories for individual promoters during metal exposures. Solid lines show the mean value over all inductions for that metal and the shaded region around lines represents SD. Dashed black lines represent metal exposure window. While some promoters are responsive to many different metals, additional information from other promoters helps the classifier to differentiate each metal. Many promoters with noisy and subtle metal responses also contribute to classifier performance.

insight into how the classifier operates in real time (Fig. 3C). The causes of misclassification are made clearer, as SHAP dynamics reveal that the predictive impact of a strain often varies within an induction window, particularly at its start and end. Furthermore, we see how some promoters, such as *zntA*, positively contribute to the detection of multiple metals, which causes the classifier to rely on promoters with less-pronounced responses and lower SHAP value magnitudes to distinguish the exposed metal, explaining some misclassification instances. Finally, promoters that may not have been identified as responsive using fold-change analysis because of subtle, low-amplitude, and noisy responses can be identified via XAI. While these responders may not serve as stand-alone biosensor strains, they provide promising targets for future sensor engineering efforts. These insights highlight the ability of the LSTM-RNN classifier to compile the influence of many strains, prominent and subtle, to make an accurate prediction of the metal exposure.

The SHAP algorithm also highlights similarities and differences between how the LSTM-RNN and XGBoost make decisions. Fig. 4A shows the 15 promoters with the highest mean impact on the model, plus the promoterless strain U139, which is included as a negative control. Both methods rely heavily on the metal-sensing promoter *zntA* for the detection and discrimi-

nation of multiple metals, especially cadmium and zinc. Beyond *zntA*, XGBoost relies heavily on single strains to detect single metals, in a manner comparable to human attention patterns. The LSTM-RNN, on the other hand, utilizes many strains of moderate influence in a combinatorial fashion; this tendency to find a different representation from that of the human visual system has been noted in other works (40). These trends are also seen when looking at the top 15 promoters for each individual metal class (SI Appendix, Fig. S11).

The ability of the explained classifiers to identify promoters involved in metal response serves as a valuable scientific tool, suggesting potential pathways and genes for further investigation. This value is highlighted by looking at a subset of the 10 most-impactful promoters individually for cadmium and iron inductions (Fig. 4B). These summary plots illustrate how the two classifiers make similar decisions through different methods. In the case of cadmium, *zntA* plays a significant role for both classifiers, while different sets of genes involved in ion transport or amino acid synthesis are identified for each. Most notably, the *metE* and *metB* promoters which are involved in methionine synthesis, an amino acid known to chelate cadmium (41), are identified by XGBoost, while the LSTM-RNN uses only the *metE* regulator, *metR*, for detection. Similarly with iron, we see XGBoost rely on members of the arginine synthesis, *argA* and *argC*, while the LSTM-RNN relies on different promoters that are involved in other metabolic or biosynthetic processes.

Biosensor Validation. Given the severe impact of heavy metals on human health (42) and the persistence of water quality issues in the United States (43), we sought to deploy the Dynamics platform as a real-time water quality biosensor. To verify that this device was functional on waters of varying ion compositions, we conducted experiments with media made from municipal water samples from San Diego, Seattle, Chicago, Miami-Dade, and New York City with added cadmium. Fig. 5A shows the LSTM-RNN classifier predictions for cadmium exposures on each city's water supply. While there is some misclassification of cadmium as zinc, there are few instances of incorrectly predicting the presence of a toxin versus water, even with largely different water compositions between cities. Additionally, the mean absolute SHAP values for each city correlated strongly with those for laboratory Milli-Q water ($R^2 = 0.853$), indicating that water composition did not affect gene response. *zntA* was the best predictor of cadmium presence across all water compositions (SI Appendix, Fig. S13).

The Dynamics device was also exposed to samples collected from the Gold King Mine spill in August 2015. Fig. 5B shows the predictions of the LSTM-RNN classifier on samples from the spill, collected from the San Juan River. The classifier predictions are output as multiclass, multilabel probability vectors. As the sample was introduced onto the device, the probability of uncontaminated water decreased significantly while the probabilities of the other metals increased. The metal with the highest probability, iron, was also the most abundant metal in the samples, as measured by inductively coupled plasma mass spectrometry (ICP-MS) (SI Appendix, Table S2). Despite the classifier not being trained on combinations of metals or at the concentrations present in these samples, the ability to reliably report the presence of the most prominent metal and, to a lesser degree, the less abundant metals suggests the broad applicability of this platform for heavy metal detection.

Discussion

In this work, we developed a high-throughput microfluidic platform to track the transcriptional dynamics of thousands of *E. coli* genes in parallel. The Dynamics platform offers a useful experimental approach through its high temporal resolution, degree of multiplexing, and precise experimental control. In




Fig. 4. XAI offers insights into the *E. coli* transcriptional dynamics contributing to metal classification. (A) Bar plot showing the cumulative contribution based on the SHAP values of 15 promoters and a negative control (promoterless strain U139) to the prediction of each metal for both XGBoost and LSTM-RNN classifiers. Colored bars for each metal represent the mean absolute SHAP value over all experimental time points. (B) SHAP values shown for 10 top promoters and a negative control (promoterless strain U139) for Cd(II) and Fe(III) for XGBoost and LSTM-RNN. Each point represents the feature value (normalized first derivative) at a given time point. Positive SHAP values suggest that a given metal is present while negative values suggest its absence. Up-regulated promoters (*zntA*, *codB*) give high SHAP values when feature values are high. Promoters are annotated with prominent gene ontology terms enriched between the two datasets.




Fig. 5. Dynamics and machine learning on environmental samples. (A) LSTM-RNN classification of cadmium contamination in five different urban water sources. Each city has a row for the true media condition, the predicted condition, and whether the time points are misclassified (red). The colors correspond to the metals in B, *Inset*. (B) Multiclass, multilabel classification of water samples from the San Juan River during the 2015 Gold King Mine waste water spill. Independent probabilities of each class are determined by the sigmoid activation function. The plot shows the sum of the classifier probabilities, averaged across triplicate sample exposures (addition and removal at vertical black lines). *Inset* bar chart shows the concentration of detectable metals in San Juan River samples as determined by ICP-MS. The colors of predicted toxins correspond to the metals plotted in *Inset*.

a high-throughput screen using Dynamics, we simultaneously exposed 1,807 strains of the promoter-based *E. coli* GFP library to nine different heavy metals. The fine-grained temporal gene expression data it produced highlighted the unique dynamics of stimuli-specific genes previously reported as heavy metal responsive (44) and identified gene clusters that shared similar response dynamics.

We illustrate our platform's potential for exploring the dynamics of transcriptional networks by applying machine-learning techniques to examine heavy metal stress responses in *E. coli*. Here we demonstrate that supervised machine learning can infer exposure to environmental stressors from real-time observation of transcriptional activity at the genome scale. Time series from 1,807 strains were used to differentiate between multiple biotic and xenobiotic heavy metals. We believe this study is an informative instance of dynamic mapping between transcriptomic changes captured in live microorganisms on the one hand and their surrounding environment on the other. These data, with genome-scale coverage and high sampling frequency, could be used in future studies to screen large strain libraries for common motifs, such as nonlinear interaction patterns and feedback loops, which are difficult to discern using static gene expression data (7). Furthermore, we use explainable AI techniques to gain insight into the features used by the predictive algorithms trained on our transcriptional data. The SHAP-XAI revealed that formally different algorithms rely on different biological features to classify transcriptomic adaptation to stress. While a decision tree-based model relied heavily on a small number of strains, a better-performing deep-learning algorithm based its prediction on many strains of moderate influence (Fig. 4). These findings reveal that there are different ways to segregate the high-dimensional space explored by an organism's transcriptome during sensory response.

Finally, we show the real-world applicability of our platform for the detection of heavy metals in both urban water sources and field samples from a recent environmental catastrophe. Compared to conventional methods of metal quantification, such as atomic absorption spectroscopy or ICP-MS, the Dynamics platform sacrifices detection sensitivity for the ability to report

continuous measurements, eliminating the need to take discrete samples. Although the Dynamics platform sometimes experiences a slight lag in the detection of metals when they are first introduced or removed, the platform is still a significant improvement over grab sampling. While previous approaches to microorganism-based heavy metal sensing have relied on engineering a small number of biosensors that are specific to one metal (45), here we use *E. coli*'s transcriptomic response at the genome scale to detect environmental stressors. Our biosensor was robust to the differences in ionic composition of five urban water sources and consistently detected cadmium in those samples. In addition, it was able to simultaneously detect multiple target metals in mine spill samples, despite not being trained to perform this type of multiclass, multilabel classification. This result suggests our approach may outperform single-purpose biosensors in accuracy and robustness and may be adaptable to more varied sensing tasks via optimization through testing combinations of metals and different concentrations of metals. In summary, combining high-throughput microfluidics and machine learning can produce insights into the coordination of cellular processes at a system level and this type of data can be leveraged for environmental monitoring.

Materials and Methods

Microfluidic Device Development and Fabrication. Our group has previously described the microfabrication techniques used to pattern SU-8 photoresist onto a silicon wafer to create the mold for our device (46). A polydimethylsiloxane (PDMS) device was made from the wafer by mixing 77 g of Sylgard 184 and pouring it on the wafer centered on a level 5'' × 5'' glass plate surrounded with an aluminum foil seal. The degassed wafer and PDMS were cured on a flat surface for 1 h at 95 °C.

Cell Preparation. The *E. coli* promoter library (22) was arrayed using the Singer ROTOR Stinger (Singer Instrument Co. Ltd.) attachment from 96-well density formatted agar plates onto four 1,536-density formatted agar plates to match the layout of the cell traps on the microfluidic device. At the time of experimental setup, the four 1,536-density agar plates were combined onto one 6,144-density agar plate using the Singer ROTOR and grown for 2 h at 37 °C before being transferred to the device.

Microfluidic Device Loading and Bonding. A PDMS device cleaned with 70% ethanol and adhesive tape was aligned to a custom fixture compatible with the Singer ROTOR. Both the fixture and a clean 4'' × 3'' glass slide sonicated with 2% Hellmanex III were exposed to oxygen plasma. Cells were spotted from the previously arrayed 6,144-density agar plate to the aligned PDMS device using the Singer ROTOR spotting robot. The device and glass slide were bonded together and cured at 37 °C for 2 h.

Experimental Protocol. Microfluidic experiments were performed on a custom optical assembly described in *SI Appendix*. Continuous imaging occurred every 10 min, imaging both the transmitted light and GFP fluorescence channels. Cells were grown in the device on LB media with kanamycin, 0.075% Tween-20, and 50 mM methyl α-D-mannopyranoside until traps were filled to confluence. The media were then switched to a heavy-metal-trace-free minimal media (HM9) minimal media described in *SI Appendix*, Table S1, which was based on a previous study (47) and optimized for microfluidic *E. coli* growth with minimal traces of metals. Cells were grown on HM9 for 48 h before inducing with heavy metals. Heavy metal inductions occurred once a day for 4 h with HM9 media flowing on chip for the remaining 20 h. Quintuplicate inductions of each metal were performed in a random order across multiple experiments, with each experiment lasting 7 to 14 d. A total of 2,176 time traces were collected from each experiment (Fig. 2 A–C). Extracted time traces were normalized to remove device background fluorescence and strain background fluorescence (*SI Appendix*). Detailed methods on experimental setup and data collection can be found in *SI Appendix*, Table S3.

Municipal Water Experimental Setup. Water samples were obtained from the Department of Water Management at the City of Chicago in Chicago, IL; the Alex Orr Water Treatment Plant in Miami, FL; the New York City Department of Environmental Protection and Bureau of Water Supply in Corona, NY; the Seattle Public Utilities Water Quality Laboratory in Seattle, WA; and

the Alvarado Water Treatment Plant in San Diego, CA. HM9 media for each city water experiment were prepared by diluting 5× HM9 concentrate made from Milli-Q water with the water obtained from each city. The microfluidic device was initially grown on LB media with kanamycin, 0.075% Tween-20, and 50 mM methyl α -D-mannopyranoside until traps were filled to confluence and then switched to HM9 made with city water for the remainder of the experiment. Cadmium diluted in the HM9 city water media was used to perform inductions as described in [SI Appendix](#).

Gold King Mine Spill Experimental Setup. Water was collected from Mexican Hat, UT in August 2015 when the Gold King Mine spill plume reached the collection point in the San Juan River. Samples were stored in 0.5% HCl acid until tested. HM9 media were prepared by diluting 5× HM9 concentrate made from Milli-Q water with filtered San Juan River samples. The pH was adjusted to 7.05. The metal concentrations of the HM9 San Juan River samples were tested by ICP-MS at the Environmental and Complex Analysis Laboratory (ECAL) at University of California, San Diego. Four-hour inductions were performed as described in [SI Appendix](#).

Machine-Learning and Data Analysis Methods. We transformed our 18 standardized experiments' time points into a first derivative-based feature for the training and testing feature sets. To optimize the classifiers, extensive

Bayesian optimization searches were used to find optimal hyperparameter combinations (48). Throughout our hyperparameter searches, we used leave-one-out cross-validation on a per-experiment basis and appropriate overfitting-prevention strategies to ensure that any resultant classifier would generalize to future datasets. All classifiers were evaluated using the F_1 -macro scoring metric. The F_1 -macro score, which is the per-class average of the harmonic mean of precision and recall, was especially well suited because of our dataset's large multiclass imbalances, with water making up ~86% of the final feature set (49). Finally, all generalization evaluations were performed by recording the results of using leave-one-out cross-validation with early stopping and then taking the mean prediction across the cross-validation's output.

Data Availability. Preprocessed, labeled machine-learning features, the corresponding library strain position records, and the relevant metadata and code for our experiments are available on the University of California San Diego Biodynamics Laboratory website (<http://biodynamics.ucsd.edu/downloads>).

ACKNOWLEDGMENTS. We thank Ryan Johnson and Patrick Mock (Quantitative BioSciences, Inc., San Diego, CA) for help designing hardware tools used in this work. This work was supported by the Defense Advanced Research Projects Agency.

- B. Kholodenko, M. B. Yaffe, W. Kolch, Computational approaches for analyzing information flow in biological networks. *Sci. Signal.* **5**, re1 (2012).
- R. Milo et al., Network motifs: Simple building blocks of complex networks. *Science* **298**, 824–827 (2002).
- F. Jacob, J. Monod, Genetic regulatory mechanisms in the synthesis of proteins. *J. Mol. Biol.* **3**, 318–356 (1961).
- T. S. Gardner, C. R. Cantor, J. J. Collins, Construction of a genetic toggle switch in *Escherichia coli*. *Nature* **403**, 339–342 (2000).
- M. Krupp et al., RNA-Seq Atlas-a reference database for gene expression profiling in normal tissue by next-generation sequencing. *Bioinformatics* **28**, 1184–1185 (2012).
- G. La Manno et al., RNA velocity of single cells. *Nature* **560**, 494–498 (2018).
- D. L. Shis, M. R. Bennett, O. A. Igoshin, Dynamics of bacterial gene regulatory networks. *Annu. Rev. Biophys.* **47**, 447–467 (2018).
- N. T. Ingolia, S. Ghaemmaghami, J. R. S. Newman, J. S. Weissman, Genome-wide analysis in vivo of translation with nucleotide resolution using ribosome profiling. *Science* **324**, 218–223 (2009).
- Y. Ho et al., Systematic identification of protein complexes in *Saccharomyces cerevisiae* by mass spectrometry. *Nature* **415**, 180–183 (2002).
- D. A. Lashkari et al., Yeast microarrays for genome wide parallel genetic and gene expression analysis. *Proc. Natl. Acad. Sci. U.S.A.* **24**, 13057–13062 (1997).
- M. J. Heller, DNA microarray technology: Devices, systems, and applications. *Annu. Rev. Biomed. Eng.* **4**, 129–153 (2002).
- N. Hao, B. A. Budnik, J. Gunawardena, E. K. O'Shea, Tunable signal processing through modular control of transcription factor translocation. *Science* **339**, 460–464 (2013).
- J. E. Purvis, G. Lahav, Encoding and decoding cellular information through signaling dynamics. *Cell* **152**, 945–956 (2013).
- M. R. Bennett et al., Metabolic gene regulation in a dynamically changing environment. *Nature* **454**, 1119–1122 (2008).
- J. Uhlenberg et al., Long-term model predictive control of gene expression at the population and single-cell levels. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 14271–14276 (2012).
- J. T. Mettetal, D. Muzzey, C. Gomez-Urbe, A. van Oudenaarden, The frequency dependence of osmo-adaptation in *Saccharomyces cerevisiae*. *Science* **319**, 482–484 (2008).
- N. Dénervaud et al., A chemostat array enables the spatio-temporal analysis of the yeast proteome. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 15842–15847 (2013).
- R. Zhang et al., High-throughput single-cell analysis for the proteomic dynamics study of the yeast osmotic stress response. *Sci. Rep.* **7**, 42200 (2017).
- Y. Taniguchi et al., Quantifying *E. coli* proteome and transcriptome with single-molecule sensitivity in single cells. *Science* **329**, 533–539 (2010).
- A. Prindle et al., A sensing array of radically coupled genetic 'biopixels'. *Nature* **481**, 39–44 (2012).
- C. Zhang et al., Ultra-multiplexed analysis of single-cell dynamics reveals logic rules in differentiation. *Sci. Adv.* **5**, eaav7959 (2019).
- A. Zaslaver et al., A comprehensive library of fluorescent transcriptional reporters for *Escherichia coli*. *Nat. Methods* **3**, 623–628 (2006).
- R. Sharma, C. Rensing, P. Rosen, B. Mitra, B. P. Rosen, The ATP hydrolytic activity of purified ZntA, a Pb(II)/Cd(II)/Zn(II)-translocating ATPase from *Escherichia coli*. *J. Biol. Chem.* **275**, 3873–3878 (2000).
- G. Grass, C. Rensing, CueO is a multi-copper oxidase that confers copper tolerance in *Escherichia coli*. *Biochem. Biophys. Res. Commun.* **286**, 902–908 (2001).
- S. Ghatak, Z. A. King, A. Sastry, B. O. Palsson, The y-ome defines the 35% of *Escherichia coli* genes that lack experimental evidence of function. *Nucleic Acids Res.* **47**, 2446–2454 (2019).
- G. Graham, N. Csicsery, E. Stasiowski, G. Thouvenin, Labeled data set for "Genome-scale transcriptional dynamics and environmental biosensing." <http://biodynamics.ucsd.edu/downloads>. Deposited 11 December 2019.
- T. Chen, C. Guestrin, XGBoost: A scalable tree boosting system. *ArXiv:1603.02754* (10 June 2016).
- S. Hochreiter, J. Schmidhuber, Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
- J. P. McHugh et al., Global iron-dependent gene regulation in *Escherichia coli*. *J. Biol. Chem.* **278**, 29478–29486 (2003).
- C. Rensing, B. Mitra, B. P. Rosen, The zntA gene of *Escherichia coli* encodes a Zn(II)-translocating P-type ATPase. *Biochemistry* **94**, 14326–14331 (1997).
- S. P. Singh et al., Machine learning based classification of cells into chronological stages using single-cell transcriptomics. *Sci. Rep.* **8**, 17156 (2018).
- D. Castelvécchi, Can we open the black box of AI? *Nat. News* **538**, 20–23 (2016).
- J. Ma et al., Using deep learning to model the hierarchical structure and function of a cell. *Nat. Methods* **15**, 290–298 (2018).
- J. H. Yang et al., A white-box machine learning approach for revealing antibiotic mechanisms of action. *Cell* **177**, 1649–1661.e9 (2019).
- J. Zhou, O. G. Troyanskaya, Predicting effects of noncoding variants with deep learning-based sequence model. *Nat. Methods* **12**, 931–934 (2015).
- S. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions. *ArXiv:1705.07874* (25 November 2017).
- S. M. Lundberg, G. G. Erion, S.-I. Lee, Consistent individualized feature attribution for tree ensembles. *ArXiv:1802.03888* (7 March 2019).
- S. M. Lundberg et al., Explainable AI for trees: From local explanations to global understanding. *ArXiv:1905.04610* (11 May 2019).
- L. S. Shapley, "A value for n-person games" in *Contributions to the Theory of Games*, H. W. Kuhn, A. W. Tucker, Eds. (Princeton University Press, 1953), vol. 2, pp. 307–317.
- S. Dodge, L. Karam, A study and comparison of human and deep learning recognition performance under visual distortions. <https://ieeexplore.ieee.org/abstract/document/8038465>. Accessed 25 May 2019.
- A. C. Esteves, J. Felcman, Study of the effect of the administration of Cd(II) cysteine, methionine, and Cd(II) together with cysteine or methionine on the conversion of xanthine dehydrogenase into xanthine oxidase. *Biol. Trace Elem. Res.* **76**, 19–30 (2000).
- P. B. Tchounwou, C. G. Yedjou, A. K. Patlolla, D. J. Sutton, "Heavy metal toxicity and the environment" in *Molecular, Clinical and Environmental Toxicology*, A. Luch, Ed. (Springer, Basel, 2012), pp. 133–164.
- M. Allaire, H. Wu, U. Lall, National trends in drinking water quality violations. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 2078–2083 (2018).
- R. A. LaVoie, A. O. Summers, Transcriptional responses of *Escherichia coli* during recovery from inorganic or organic mercury exposure. *BMC Genom.* **19**, 52 (2018).
- H. J. Kim, H. Jeong, S. J. Lee, Synthetic biology for microbial heavy metal biosensors. *Anal. Bioanal. Chem.* **410**, 1191–1203 (2018).
- M. S. Ferry, I. A. Razinkov, J. Hasty, Microfluidics for synthetic biology: From design to execution. *Methods Enzymol.* **497**, 295–372 (2011).
- R. A. LaRossa, D. R. Smulski, T. K. Van Dyk, Interaction of lead nitrate and cadmium chloride with *Escherichia coli* K-12 and *Salmonella typhimurium* global regulatory mutants. *J. Ind. Microbiol.* **14**, 252–258 (1995).
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, N. De Freitas, Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* **104**, 148–175 (2016).
- Z. C. Lipton, C. Elkan, B. Narayanaswamy, Thresholding classifiers to maximize F1 score. *ArXiv:1402.1892* (14 May 2014).