

## Learning predictive models of drug response that translate across biological contexts

Jianzhu Ma<sup>1,2</sup>, Samson H. Fong<sup>1,3</sup>, Yunan Luo<sup>4</sup>, Christopher J. Bakkenist<sup>5</sup>,

John Paul Shen<sup>6</sup>, Soufiane Mourragui<sup>7, 8</sup>, Lodewyk F. A. Wessels<sup>7, 8</sup>,

Marc Hafner<sup>9</sup>, Roded Sharan<sup>10</sup>, Jian Peng<sup>4</sup>, Trey Ideker<sup>1,3,\*</sup>

1. Department of Medicine, University of California San Diego, La Jolla, CA 92093, USA
2. Department of Computer Science, Purdue University, West Lafayette, IN 47907, USA
3. Department of Bioengineering, University of California San Diego, La Jolla, CA 92093, USA
4. Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
5. Department of Radiation Oncology, University of Pittsburgh School of Medicine, Pittsburgh, PA 15260, USA
6. Department of Gastrointestinal Medical Oncology, University of Texas MD Anderson Cancer Center, Houston, TX 77030, USA
7. Division of Molecular Carcinogenesis, Oncode Institute, The Netherlands Cancer Institute, Plesmanlaan 121, 1066 CX, Amsterdam, The Netherlands.
8. Faculty of EEMCS, Delft University of Technology, Delft, The Netherlands
9. Department of Bioinformatics and Computational Biology, Genentech, Inc, South San Francisco, CA 94080, USA
10. Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel

\* Correspondence to Trey Ideker at [tideker@ucsd.edu](mailto:tideker@ucsd.edu)

Cell-line screens create expansive datasets for learning predictive markers of drug response, but these models do not readily translate to the clinic with its diverse contexts and limited data. Here we apply a recently developed technique, few-shot machine learning, to train a versatile neural network model in cell lines that can be tuned to new contexts using few additional samples. The model quickly adapts when switching among different tissue types and in moving from cell-line models to clinical contexts, including patient-derived tumor cells and patient-derived xenografts. It can also be interpreted to identify the molecular features most important to a drug response, highlighting critical roles for *RB* and *SMAD4* in the response to CDK inhibition and *RNF8* and *CHD4* in the response to ATM inhibition. The new learning framework provides a bridge from the many samples surveyed in high-throughput screens (N-of-many) to the distinctive contexts of individual patients (N-of-one).

Translating biomarkers from basic research to clinical utility involves transfer of information across a series of contexts in which data are progressively harder to obtain. *In-vitro* platforms such as human cell culture are amenable to high-throughput screening, yielding large datasets characterizing the molecular profiles of thousands of cell lines and their responses to millions of chemical compounds, genetic interventions or environments<sup>1,2</sup>. Promising indications may progress to advanced culture systems and animal models<sup>3,4</sup>, few of which are further evaluated in human cohorts and, ultimately, used in diagnosis and treatment of individual patients.

It is well known that drug response predictions learned in cell-line or animal models do not always transfer to clinical contexts in a straightforward manner<sup>5-7</sup>. For example, dual inhibition of EGFR and VEGFR was found to induce sustained tumor regression in a mouse model of EGFR-mutant non-small cell lung cancer<sup>8</sup> (Naumov et al. 2009), whereas follow-up clinical studies failed to replicate such an effect<sup>9</sup> (Lee et al. 2012). Similarly, upregulation of IGF1R had been noted as a prominent marker of tamoxifen resistance in breast cancer cell lines<sup>10</sup> (Parisot et al. 1999), whereas opposite behavior – reduced IGF1R protein levels – was observed in tamoxifen-resistant patients<sup>11</sup> (Drury et al., 2011). It remains unclear, however, whether such failures are caused by fundamental irreconcilable differences between biological contexts or missed opportunities to identify the correct markers likely to translate. A key challenge in marker selection is that the common signal is easily overwhelmed by context-specific patterns, especially given the very limited amounts of data available in patients relative to cell lines.

To improve biomarker transfer across contexts, we formulated a neural network model, Translation of Cellular Response Prediction (TCRP), using the technique of few-shot learning<sup>12,13</sup>. Few-shot learning is an emerging method of transfer learning, a field that

postulates that prior knowledge acquired in one problem domain can be reused and applied to solve different but related problems<sup>14–16</sup>. Transfer learning has proven instrumental in fields such as linguistics, where people (and machines) can learn to speak a new language much more quickly if they have extensive prior knowledge of a related tongue which can be transferred efficiently to the new one<sup>17</sup>. Recent applications in biomedicine include an improved ability to identify chemical compounds with biological activity<sup>18</sup> or to classify tissue type and tumor grade in histopathological images<sup>19</sup>.

Few-shot learning aims to identify widely applicable input features by optimizing their transferability, rather than their overall prediction accuracy as in conventional learning approaches (**Online Methods**). In an initial “pre-training” phase (**Fig. 1 top**), the model is exposed to a variety of different pre-defined contexts, each of which is represented by numerous training samples. In a second “few-shot learning” phase (**Fig. 1 bottom**), the model is presented with a new context not seen previously, and further learning is performed on a small number of new samples. Neural networks trained by this two-phase design have been shown to learn surprisingly rapidly in the new context relative to models trained by conventional means<sup>20–23</sup>.

Here, we applied the few-shot learning paradigm to three context-transfer challenges of high interest in predictive medicine: (1) Transfer of a predictive model learned in one tissue type to the distinct contexts of other tissues; (2) Transfer of a predictive model learned in tumor cell lines to patient-derived tumor cell cultures *in vitro*; and (3) Transfer of a predictive model learned in tumor cell lines to the context of patient-derived tumor xenografts in mouse models *in vivo* (**Fig. 1, Supplementary Table 1**).

## Results

### Challenge 1: Transfer across tissue types

For the first challenge, we evaluated the ability of our TCRP model to predict the growth rates of tumor cell lines **from a target tissue for which very few samples were available for learning**. Data were taken from a recent survey of 335 human cell lines from 19 tissues, in which cell growth rates had been measured across a genome-wide panel of gene disruptions using CRISPR-Cas9<sup>1</sup> (**Online Methods**). For each cell line, this same survey summarized the binary genotype status of genes (0 = unmutated or synonymous mutation; 1 = non-synonymous mutation) and their quantitative mRNA abundance levels during nominal growth. For each CRISPR-Cas9 gene disruption (focusing on 469 genes with demonstrated tumor growth dependencies), we trained TCRP alongside a collection of conventional learning models to predict the growth responses of all cell lines. **During this process, 1 of the 19 tissues was designated as the target. A training set was then created that included all cell lines from the other 18 tissues but only a small number of cell lines from the target tissue; the remaining target cell lines constituted the test set. TCRP was trained in two phases, first on the large number of cell lines from the 18 tissues (pre-training phase), then on the small number of cell lines available from the target tissue (few-shot learning phase, Fig. 1, Online Methods).** Conventional models were trained using a standard one-phase training procedure, by pooling all samples designated as training, after which the model was evaluated on all samples designated **as test**. Key questions were how quickly a predictive model transfers to the new tissue, having been trained mainly on others, and to which tissues the model transfers worst/best.

Models displayed a range of prediction accuracies during pre-training, as assessed by five-fold cross validation, with conventional random forests performing best (**Supplementary Fig. 1, Online Methods**). However, when switching to the target tissue, no model performed better than random, demonstrating the difficulty posed by new contexts (**Fig. 2a**). We then entered into the few-shot learning phase. For conventional models, accuracy improved very slowly as samples from the new tissue were added to the training set. In contrast, TCRP improved rapidly, with an average gain of 829% in performance after examining only five additional samples (**Fig. 2a**). Tissues with the most improvement were the kidney, urinary tract and pancreas (**Fig. 2b**). For example, we observed a very high accuracy when predicting the response to CRISPR knockout of hepatocyte nuclear factor 1 beta (HNF1B), for which TCRP achieved a performance of 0.60 (Pearson correlation) in contrast to the second best approach (random forests, 0.19). The importance of HNF1B to tumor growth has been verified in multiple cancer types, including hepatocellular carcinoma, pancreatic carcinoma, renal cancer, ovarian cancer, endometrial cancer, and prostate cancer<sup>24</sup>.

We also conducted a related Challenge 1b, in which cell growth response data were drawn from a high-throughput pharmacogenomic screen of 255 anti-cancer drugs (including both FDA-approved and experimental compounds, **Online Methods**) administered to each of 990 cancer cell lines encompassing 30 tissues<sup>2</sup>. Similar to Challenge 1a, but for each drug, TCRP was trained alongside conventional learning models to predict the growth sensitivity of cell lines using their molecular markers. As before, TCRP learned rapidly when switching to the target tissue, with the largest improvements seen when learning from the first few cell-line samples (**Fig. 2c,d**). We found that the accuracy of drug predictions was highly correlated with the accuracy of CRISPR predictions across the tissues examined, with tissues like urinary tract

generating highly accurate predictions in both settings, and tissues like central nervous system, skin, and lung generating poor predictions (Spearman  $\rho = 0.73$ ,  $p = 0.012$ ).

## **Challenge 2: Transfer to patient-derived tumor cells**

Next, we studied whether models of drug response trained on cell lines could be transferred to a pre-clinical context (Challenge 2, **Fig. 3a**). For this challenge we used data on breast cancer patient-derived tumor cells (PDTCs) made available by Project Biobank<sup>4</sup> (**Online Methods**). In this prior study, tumors ( $n=83$ ) were biopsied, subjected to whole-exome and mRNA sequencing to generate molecular profiles, and implanted in immunodeficient mice. PDTCs were then isolated from the host mice and tested for drug responses *in vitro*. From these data we selected 50 drugs for which the protein targets were well characterized, with drugs administered to 15-19 PDTCs each. For each drug, TCRP was pre-trained using the GDSC cell-line drug response data from Challenge 1b before switching context to PDTCs.

As observed with earlier challenges, all models performed poorly when switching contexts, achieving accuracies near or below zero (**Supplementary Fig. 1**). However, once again we observed that TCRP improved substantially after exposure to each new patient sample: the average performance was  $r = 0.30$  at 5 samples, reaching  $r = 0.35$  at 10 samples versus  $r < 0.10$  for the runner up (**Fig. 3b,c, Supplementary Fig. 2a**). Nearly all drug predictions were improved by the few-shot paradigm. For example, the ATM inhibitor KU-55933 had the top performing drug response predictions, with a Pearson correlation of 0.56 between predicted and actual growth response measurements (top row of **Fig. 3c**, average performance over 5-10 samples). KU-55933 also represented the largest improvement over conventional approaches,

where the best performing conventional model, the random forest, obtained correlations of approximately 0.12.

### Challenge 3: Transfer to PDX in mice

Finally, in Challenge 3 we went a step further, moving from PDTs tested against drugs *in vitro* to patient-derived tumor xenografts (PDXs) tested against drugs in live mice (**Fig. 4a**, **Supplementary Fig. 2b**, **Supplementary Fig. 3**). For this purpose we obtained data for 228 PDX mouse models from the PDX Encyclopedia<sup>25</sup>, where each model was exposed to one of the five drugs on which TCRP had been trained in cell lines (Cetuximab, Erlotinib, Paclitaxel, Tamoxifen, and Trametinib; see Challenge 1b). Also provided were genotype and mRNA transcriptomes of each PDX, from which we obtained the molecular features used by TCRP to make drug response predictions. In cell lines, the predicted output from TCRP was the area under dose response curve (AUC); for PDXs, the analogous measurement was the percent change in tumor volume resulting from drug treatment *in vivo* ( $\Delta\text{vol}$ ). Therefore, these predicted and measured values were each normalized to a standard normal distribution to translate between the two (i.e. z-score; **Online Methods**).

Although TCRP models pre-trained on cell-line data initially performed poorly in predicting PDX responses, we observed significant improvements during training on the first few PDX samples (**Fig. 4a**). Such improvements were seen for all five drugs and led to a range of final prediction accuracies from  $r = 0.50$  for Erlotinib to  $r = 0.18$  for Paclitaxel (Spearman correlation between predicted and actual drug response after training on ten PDX samples, **Fig. 4a**, **Supplementary Fig. 2b**). We also explored the effect of translating the continuous-valued drug response predictions to discrete treatment outcomes, as are typically assigned in a clinical setting,



designating each response as either Progressive Disease (PD,  $\Delta\text{vol} \geq 30\%$ ) or Stable Disease or better (SD,  $\Delta\text{vol} < 30\%$ ). We found that these predicted binary classifications were strongly associated with the observed PD/SD outcomes, with a range of odds ratios from 3.0 (Cetuximab) to 10.5 (Tamoxifen) (**Fig. 4b,c**). For Cetuximab, Paclitaxel, Tamoxifen, and Trametinib but not Erlotinib, we found that the predicted PD/SD designations also showed significant differences in progression-free survival, depending on how many PDX samples had been used for few-shot learning (**Fig. 4d-g**).

### Interpreting the predictive models

A common critique of machine learning systems is that they produce ‘black boxes’ whose predictions are difficult to interpret<sup>26,27</sup>. Here, since we had focused on drugs with known specific targets, we found that model predictions were typically explainable by molecular markers within that target’s pathway (using models constrained to these features, **Online Methods**). For example, a top feature in predicting the response of PDTs to PD-0332991 (Palbociclib, **Fig. 5a,b**) was the expression of RB-Like-Factor RBL2, a cell-cycle transcriptional repressor inactivated by CDK4/6. RBL2 expression was associated with Palbociclib resistance (third from top in **Fig. 5c**,  $r = 0.47$ ), suggesting that high RBL2 activity masks upstream inhibition of CDK4/6 by the drug. Another important feature was somatic mutation of SMAD4, a transcriptional modulator repressing CDK4 transcription<sup>28</sup> (**Fig. 5d**). SMAD4 inactivation may release CDK4 to drive cell cycle<sup>29</sup>, with CDK4 inhibition counteracting this effect (**Fig. 5b**). While SMAD4 mutation was rare in PDTs (1/19 samples), it was much more common in cell lines (43/811 samples). The model thus learns to strongly rely on the SMAD4 mutation during pre-training, in which frequent SMAD4 mutations are strongly associated with drug response. When switching to the

PDTC data set, this prior information is combined with the prevalence and effect of SMAD4 mutation in the new dataset to jointly estimate its importance to the drug response.

As a second example, a top feature in the response to ATM inhibition (KU-55933, **Fig. 5e,f**) was the expression of RNF8, which is recruited to DNA double-stranded breaks (DSBs) following activation of ATM by DNA damage<sup>30–32</sup>. RNF8 expression was correlated with KU-55933 resistance (third from top in **Fig. 5g**,  $r = 0.54$ ), suggesting that, when RNF8 activity is high, ATM is not limiting for DSB repair. Also correlated with drug resistance was mutation of CHD4 (**Fig. 5h**), encoding the chromodomain-helicase-DNA-binding subunit of NuRD, a complex essential for chromatin relaxation at DSBs<sup>33</sup>. Disabled NuRD may interfere with DNA repair, masking the effects of ATM inhibition. Alternatively, it may dampen the impact of ATM on CHD4-dependent cell-cycle progression<sup>34</sup>.

A significant third example involved BRAF inhibition, which leads to drug sensitivity in the context of a BRAF activating mutation. It is well-established that some tissue types respond to BRAF inhibition more strongly than others; for instance, BRAF-mutant melanomas are generally responsive whereas colorectal tumors are not, for reasons that are not fully understood but are partially explained by expression of EGFR<sup>35</sup>. As expected from these prior observations, the TCRP model predicted significant sensitivity to the BRAF inhibitor Dabrafenib in BRAF-mutant cells, but not in wild-type cells, with a much more pronounced effect in melanoma than colorectal cancer (**Supplementary Fig. 4a**). Of note, the drug response predicted by TCRP was significantly more accurate than the response predicted solely based on BRAF mutation and EGFR expression status (**Supplementary Fig. 4b**), raising the question of which features TCRP had used to achieve higher accuracy. Further examination indicated that the model drew from a combination of novel features (**Supplementary Fig. 4c-f**). These included expression of mRAS,

which has been shown to function as a RAF phosphatase<sup>36</sup>, expression of 14-3-3 proteins YWHAE and YWHAH, which interact with RAF proteins in signal transduction<sup>37</sup>, and mutation of RAPGEF1, a human guanine nucleotide exchange factor central to activation of the Ras/Raf/MEK/ERK signal transduction pathway.

## Discussion

Recently an abundance of biological response data have been generated for targeted perturbations in numerous contexts. The usual way of analyzing these data is to pool all samples, under the assumption that accruing the maximal amount of data will result in a predictive model with greatest statistical power. Here, we have identified a more efficient means of building predictive models, using the technique of few-shot learning. The two-phase learning procedure overlays naturally on the process of translating observations from basic research *in vitro* to predictive markers in tumors (**Fig. 3a**): First, in a basic research phase, a general predictive model is pre-trained from extensive data generated in high-throughput cell-based screens. Second, in a pre-clinical or clinical phase, few-shot learning is used to tune the general model to make predictions for a specific type of human tumors, by testing drugs with high predicted sensitivity in settings such as patient-derived tumor cells and xenografts and, ultimately, patients. Thus far, few-shot learning shows encouraging performance in multiple datasets and translation scenarios where conventional learning fails. In all three challenges we examined, the initial pre-training phase was the same: optimizing the model for transfer across cell lines of different tissue types. Notably, this particular transfer task was sufficiently general to enable predictive models to transfer from cell lines to the settings of PDTs and PDXs.

Models like TCRP may have compelling applications in clinical contexts seeking to implement precision medicine, in which the task is to match a patient's specific molecular profile to an optimal course of therapy. For this purpose molecular tumor boards have been established in many cancer centers, where clinical experts must often make treatment decisions for a patient based on just a few precious cases with matching histopathology and molecular profiles. A second compelling application is in the pharmaceutical industry, in which a key goal is to select patients that are most likely to respond to a targeted agent. In both cases, classical predictive models have been significantly hampered by lack of access to large numbers of well-characterized clinical samples, i.e. for which molecular profiles have been coupled to precise information on treatment outcomes.

In this regard, an important question for future exploration concerns the degree to which an approach like TCRP is ready for use in clinical or pharmaceutical settings. There are many uncertainties when deciding on treatment, and how the predictive value of the models built here compare to other molecular and clinical markers, and their predictive values, will need to be determined for each disease setting. In terms of absolute predictive performance, we observed a range of accuracy across the drugs examined, with some drugs indeed yielding promising results. For example, in the PDX analysis of paclitaxel (**Figs. 4b,c,e**), a drug non-response was predicted for 23 tumors, of which 20 were in agreement with the actual observations of tumor growth in mice, a very high success rate by any standard ( $20/23 = 87\%$  correct predictions of progressive disease). As another example, non-response of PDX tumors to tamoxifen was correctly predicted in 23/24 of cases (96%). In these analyses, a non-response (progressive disease) was called if the change in tumor volume was  $\geq 30\%$ , the standard threshold implemented by the PDX Encyclopedia<sup>25,38</sup>. Given more data and a focused clinical study, one

could likely tune the prediction threshold to drive performance higher. For example, at a threshold value of >60%, TCRP predicts paclitaxel non-response with 100% accuracy given the current PDX dataset (14/14 patients). Future investigations with larger cohorts of PDX models or patients will be able to shed further light on the best uses of few-shot learning clinically.

In our analysis of both the PDTC (**Fig. 3b**) and PDX (**Fig. 4a**) datasets, we noted that the performance of few-shot learning improves quickly and then appears to saturate. Further inspection reveals that the reason for this phenomenon relates to the balance of training versus test samples during evaluation. Given a fixed number of tumor samples, as the number of few-shot training samples increases, the number of testing samples decreases proportionally. In turn, a fewer number of testing samples means that the statistical power used to evaluate the prediction performance gets weaker, with a concomitant increase in variance. For most drugs in the PDTC dataset, a total of 19 tumor samples were available to be split between training and validation. To evaluate performance for 1-shot learning, 18 of these samples were therefore available as a test set, whereas for 10-shot learning, only 9 samples were available for test.

We also observed that drug responses were better predicted in some tissues compared to others (**Fig. 2b,d**). Although the poor predictive power in some tissues is in need of further investigation, a potential factor relates to the significant molecular heterogeneity observed within some cancer tissue types. For example, cell lines of lung tumors have been organized into as many as nine subtypes based on their transcriptomic profiles, in contrast to pancreatic tumor cell lines which appear far more homogeneous<sup>39</sup>. These findings are superficially in agreement with those of our study, in that drug response predictions in lung cancer lines are less predictive than those of pancreas (**Fig. 2b,d**).

While the results demonstrated here were obtained with gene mutation and mRNA expression features, the TCRP framework is general with potential relevance to many other data types, such as copy number variants, features extracted from histopathological images, or data transferred from disease models in other species. Furthermore, while each perturbation by CRISPR (Challenge 1a) or drugs (all other challenges) was considered a separate machine learning task, a worthy future direction would be to explore the extent to which information can be transferred from one perturbation to another. If significant information is shared, one might pursue a single unified model with predictive capacity across many or all drugs rather than training models separately.

A final future direction is to better understand the relationship between the predictability of a drug and its pharmacological properties, including its number of recognized targets and off-target effects (i.e. polypharmacology). This relationship is difficult to study with the present TCRP, for which features are selected from the protein interaction network of each known target, yielding a tendency to include more features for drugs that have more known targets (**Online Methods**). On the other hand, our understanding of drug target genes and pathways is far from complete, and the protein network we used for feature selection is not cancer-specific. Future model configurations using the same numbers of biomarkers across drugs will potentially shed light on the complex interactions between drug response and polypharmacology.

## **Data Availability**

The software implementation of TCRP and dataset are available at <https://github.com/idekerlab/TCRP/>.

## **Acknowledgements**

We gratefully acknowledge support for this work provided by grants from the National Cancer Institute (U54CA209891 to TI, R01CA204173 to CB, K22CA234406 to JPS) and the National Institute of General Medical Sciences (P41GM103504 and R01HG009979 to TI). RS was supported by a research grant from the Israel Science Foundation (grant no. 715/18). JP was supported by a grant from the National Science Foundation (grant no. 1652815). LFAW and SM were supported by ZonMw TOP grant COMPUTE CANCER (40-00812-98-16012). JPS was supported by the Cancer Prevention and Research Institute of Texas (CPRIT RR180035).

## **Author Contributions**

JM and TI designed the study and developed the conceptual ideas. JM implemented the main algorithms. JM and SF collected all the input sources. JM, SM, LW, and MH developed the strategy for alignment of *in vitro* and *in vivo* drug responses. JM, CB and TI interpreted the results. JM, SF, RS, CB, JP, JPS and TI wrote the manuscript.

## **Competing Interests**

TI is co-founder of Data4Cure, Inc., is on the Scientific Advisory Board, and has an equity interest. TI is on the Scientific Advisory Board of Ideaya BioSciences, Inc., and has an equity interest. The terms of these arrangements have been reviewed and approved by the University of California San Diego in accordance with its conflict of interest policies. LW received project funding from Genmab BV.

## References

1. Meyers, R. M. *et al.* Computational correction of copy number effect improves specificity of CRISPR-Cas9 essentiality screens in cancer cells. *Nat. Genet.* **49**, 1779–1784 (2017).
2. Iorio, F. *et al.* A Landscape of Pharmacogenomic Interactions in Cancer. *Cell* **166**, 740–754 (2016).
3. Brabetz, S. *et al.* A biobank of patient-derived pediatric brain tumor models. *Nat. Med.* **24**, 1752–1761 (2018).
4. Bruna, A. *et al.* A Biobank of Breast Cancer Explants with Preserved Intra-tumor Heterogeneity to Screen Anticancer Compounds. *Cell* **167**, 260–274.e22 (2016).
5. Butler, D. Translational research: Crossing the valley of death. *Nature* vol. 453 840–842 (2008).
6. Lieu, C. H., Tan, A.-C., Leong, S., Diamond, J. R. & Eckhardt, S. G. From bench to bedside: lessons learned in translating preclinical studies in cancer drug development. *J. Natl. Cancer Inst.* **105**, 1441–1456 (2013).
7. Seyhan, A. A. Lost in translation: the valley of death across preclinical and clinical divide – identification of problems and overcoming obstacles. *Translational Medicine Communications* vol. 4 (2019).
8. Naumov, G. N. *et al.* Combined vascular endothelial growth factor receptor and epidermal growth factor receptor (EGFR) blockade inhibits tumor growth in xenograft models of EGFR inhibitor resistance. *Clin. Cancer Res.* **15**, 3484–3494 (2009).
9. Lee, J. S. *et al.* Vandetanib Versus placebo in patients with advanced non-small-cell lung cancer after prior therapy with an epidermal growth factor receptor tyrosine kinase inhibitor: a randomized, double-blind phase III trial (ZEPHYR). *J. Clin. Oncol.* **30**, 1114–1121 (2012).
10. Parisot, J. P., Hu, X. F., DeLuise, M. & Zalcberg, J. R. Altered expression of the IGF-1



- receptor in a tamoxifen-resistant human breast cancer cell line. *Br. J. Cancer* **79**, 693–700 (1999).
11. Drury, S. C. *et al.* Changes in breast cancer biomarkers in the IGF1R/PI3K pathway in recurrent breast cancer after tamoxifen treatment. *Endocr. Relat. Cancer* **18**, 565–577 (2011).
  12. Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science* **350**, 1332–1338 (2015).
  13. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D. & Lillicrap, T. Meta-Learning with Memory-Augmented Neural Networks. in *Proceedings of The 33rd International Conference on Machine Learning* (eds. Balcan, M. F. & Weinberger, K. Q.) vol. 48 1842–1850 (PMLR, 2016).
  14. Dai, W., Yang, Q., Xue, G.-R. & Yu, Y. Boosting for transfer learning. in *Proceedings of the 24th international conference on Machine learning* 193–200 (Association for Computing Machinery, 2007).
  15. Blitzer, J., McDonald, R. & Pereira, F. Domain adaptation with structural correspondence learning. in *Proceedings of the 2006 conference on empirical methods in natural language processing* 120–128 (2006).
  16. Argyriou, A., Evgeniou, T. & Pontil, M. Multi-Task Feature Learning. in *Advances in Neural Information Processing Systems 19* (eds. Schölkopf, B., Platt, J. C. & Hoffman, T.) 41–48 (MIT Press, 2007).
  17. Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. The Omniglot challenge: a 3-year progress report. *Current Opinion in Behavioral Sciences* **29**, 97–104 (2019).
  18. Altae-Tran, H., Ramsundar, B., Pappu, A. S. & Pande, V. Low Data Drug Discovery with One-Shot Learning. *ACS Cent Sci* **3**, 283–293 (2017).

19. Medela, A. *et al.* Few Shot Learning in Histopathological Images: Reducing the Need of Labeled Data on Biological Datasets. *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)* (2019) doi:10.1109/isbi.2019.8759182.
20. Snell, J., Swersky, K. & Zemel, R. Prototypical Networks for Few-shot Learning. in *Advances in Neural Information Processing Systems 30* (eds. Guyon, I. *et al.*) 4077–4087 (Curran Associates, Inc., 2017).
21. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K. & Wierstra, D. Matching Networks for One Shot Learning. in *Advances in Neural Information Processing Systems 29* (eds. Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I. & Garnett, R.) 3630–3638 (Curran Associates, Inc., 2016).
22. Finn, C., Abbeel, P. & Levine, S. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *arXiv [cs.LG]* (2017).
23. Preuer, K. *et al.* DeepSynergy: predicting anti-cancer drug synergy with Deep Learning. *Bioinformatics* **34**, 1538–1546 (2018).
24. Yu, D.-D., Guo, S.-W., Jing, Y.-Y., Dong, Y.-L. & Wei, L.-X. A review on hepatocyte nuclear factor-1beta and tumor. *Cell Biosci.* **5**, 58 (2015).
25. Gao, H. *et al.* High-throughput screening using patient-derived tumor xenografts to predict clinical trial drug response. *Nat. Med.* **21**, 1318–1325 (2015).
26. Lipton, Z. C. The Mythos of Model Interpretability. *arXiv [cs.LG]* (2016).
27. Ma, J. *et al.* Using deep learning to model the hierarchical structure and function of a cell. *Nat. Methods* (2018) doi:10.1038/nmeth.4627.
28. Liu, F. & Matsuura, I. Inhibition of Smad antiproliferative function by CDK phosphorylation. *Cell Cycle* **4**, 63–66 (2005).
29. Zhao, M., Mishra, L. & Deng, C.-X. The role of TGF- $\beta$ /SMAD4 signaling in cancer. *Int. J.*

*Biol. Sci.* **14**, 111–123 (2018).

30. Zhang, F., Bick, G., Park, J.-Y. & Andreassen, P. R. MDC1 and RNF8 function in a pathway that directs BRCA1-dependent localization of PALB2 required for homologous recombination. *J. Cell Sci.* **125**, 6049–6057 (2012).
31. Lu, C.-S. *et al.* The RING finger protein RNF8 ubiquitinates Nbs1 to promote DNA double-strand break repair by homologous recombination. *J. Biol. Chem.* **287**, 43984–43994 (2012).
32. Kobayashi, S. *et al.* Rad18 and Rnf8 facilitate homologous recombination by two distinct mechanisms, promoting Rad51 focus formation and suppressing the toxic effect of nonhomologous end joining. *Oncogene* **34**, 4403–4411 (2015).
33. Smith, R., Sellou, H., Chapuis, C., Huet, S. & Timinszky, G. CHD3 and CHD4 recruitment and chromatin remodeling activity at DNA breaks is promoted by early poly(ADP-ribose)-dependent chromatin relaxation. *Nucleic Acids Res.* **46**, 6087–6098 (2018).
34. Larsen, D. H. *et al.* The chromatin-remodeling factor CHD4 coordinates signaling and repair after DNA damage. *J. Cell Biol.* **190**, 731–740 (2010).
35. Prahallad, A. *et al.* Unresponsiveness of colon cancer to BRAF(V600E) inhibition through feedback activation of EGFR. *Nature* **483**, 100–103 (2012).
36. Young, L. C. *et al.* SHOC2–MRAS–PP1 complex positively regulates RAF activity and contributes to Noonan syndrome pathogenesis. *Proc. Natl. Acad. Sci. U. S. A.* **115**, E10576–E10585 (2018).
37. Tzivion, G., Luo, Z. & Avruch, J. A dimeric 14-3-3 protein is an essential cofactor for Raf kinase activity. *Nature* **394**, 88–92 (1998).
38. Schwartz, L. H. *et al.* RECIST 1.1—Update and clarification: From the RECIST committee.

*Eur. J. Cancer* **62**, 132–137 (2016).

39. Yu, K. *et al.* Comprehensive transcriptomic analysis of cell lines as models of primary tumors across 22 tumor types. *Nature Communications* vol. 10 (2019).

## Figure Legends

**Figure 1. Study design.** Three distinct translation challenges are considered. Each challenge involves a pre-training phase (top) based on cell-line response data across tissues, followed by a few-shot learning phase (bottom) in which data in the new context are presented for additional learning, one sample at a time. Challenge 1: Transfer of CRISPR (1a) or drug (1b) response model for prediction in the context of a new tissue. Challenge 2: Transfer of model to patient-derived tumor cells *in vitro*. Challenge 3: Transfer of model to patient-derived xenografts *in vivo*.

**Figure 2. Transfer of predictive models across tissue types.** (a) **Challenge 1a.** For each CRISPR gene knockout and target tissue, model accuracy is measured by the Pearson correlation between predicted and actual drug responses, considering only the test samples from the target. The plot shows the distribution of average model accuracy across CRISPR knockouts (y-axis, mean  $\pm$  95% CI) as a function of the number of cell lines from the target tissue provided to the model during training (x-axis). (b) Model accuracy (x-axis) is displayed separately for each tissue in Challenge 1a (y-axis). Accuracy is the average achieved when training includes 5-10 samples of the target tissue. The accuracy standard deviation is shown over all CRISPR gene knockouts (point size). (c) **Challenge 1b.** As for panel (a) for models trained on perturbations with targeted drugs. (d) As for panel (b) for models trained on perturbations with targeted drugs.

**Figure 3. Transfer of cell-line models to patient-derived tumor cell lines.** (a) Schema for translating a predictive model from cell lines to patients using few-shot learning. The model is trained over successive rounds of data, each with fewer samples but closer to the desired

clinical context. **(b) Challenge 2.** Predictive models were pre-trained using responses of breast cancer cell lines to targeted perturbations with a particular drug (**Supplementary Table 1**). Few-shot learning was then performed on 0-10 PDTC breast tumor samples exposed to that drug (x-axis), and model accuracy (Pearson correlation, y-axis) validated using the remaining held-out PDTC samples. Results averaged across 48 drugs. **(c)** Predictive accuracy (x-axis) is displayed separately for each drug model (y-axis). Colors as in previous figures.

**Figure 4. Transfer of cell-line models to patient-derived xenografts. (a) Challenge 3.** Predictive models were pre-trained using responses of cancer cell lines to targeted perturbations with drugs, one model per drug. Few-shot learning (x-axis, number of few-shot samples used) was performed using PDX samples exposed to one of five drugs (line colors), and the improved model used to predict the change in tumor volume ( $\Delta\text{vol}$ , **Online Methods**). Accuracy of this prediction was validated using the actual changes in volume of the remaining held-out PDX tumors (Pearson correlation, y-axis). **(b) Odds ratio.** We evaluated the odds of obtaining progressive disease to stable disease (PD:SD) outcomes when stratifying tumors into predicted responsive versus unresponsive subtypes (predicted  $\Delta\text{vol} < \text{or} \geq 30\%$ , respectively). Odds ratio (left) and corresponding contingency table (right) is shown for each drug. **(c)** Ranking of all PDX samples (x-axis) by the predicted  $\Delta\text{vol}$  (y-axis) for Trametinib, Paclitaxel and Erlotinib. Color indicates actual clinical outcome. **(d-g)** Kaplan–Meier survival plots when stratifying tumors into responsive versus unresponsive subtypes for **(d)** Cetuximab, **(e)** Paclitaxel, **(f)** Tamoxifen, and **(g)** Trametinib.

**Figure 5. Model interpretation to identify predictive markers. (a)** Measured versus predicted resistance to the CDK4/6 inhibitor Palbociclib after few-shot learning on five PDTC samples

treated with this drug. **(b)** Schematic of CDK pathway with Palbociclib targets and selected molecular markers. **(c)** Left: mRNA expression profiles for the top expression-based features of Palbociclib. Right: Pearson correlation of Palbociclib resistance and mRNA expression for the top expression-based features. **(d)** Left: Somatic mutation profiles for the top mutation-based features of Palbociclib. Right: Increase of Palbociclib resistance when comparing mutated and wild type samples for each top feature. **(e)** Same as panel **(a)** for the response to ATM inhibitor KU-55933. **(f)** Schematic of ATM pathway with selected predictive markers. **(g-h)** Same as panels c, d for the response to ATM inhibitor KU-55933. Numbered sample labels in panels (a, e) correspond to PDTC sample numbers in (c, d, g, h), in which molecular profiles for the 6 most sensitive and 6 most resistant samples are shown (PDTC1-6 and PDTC14-19, respectively).

Figures

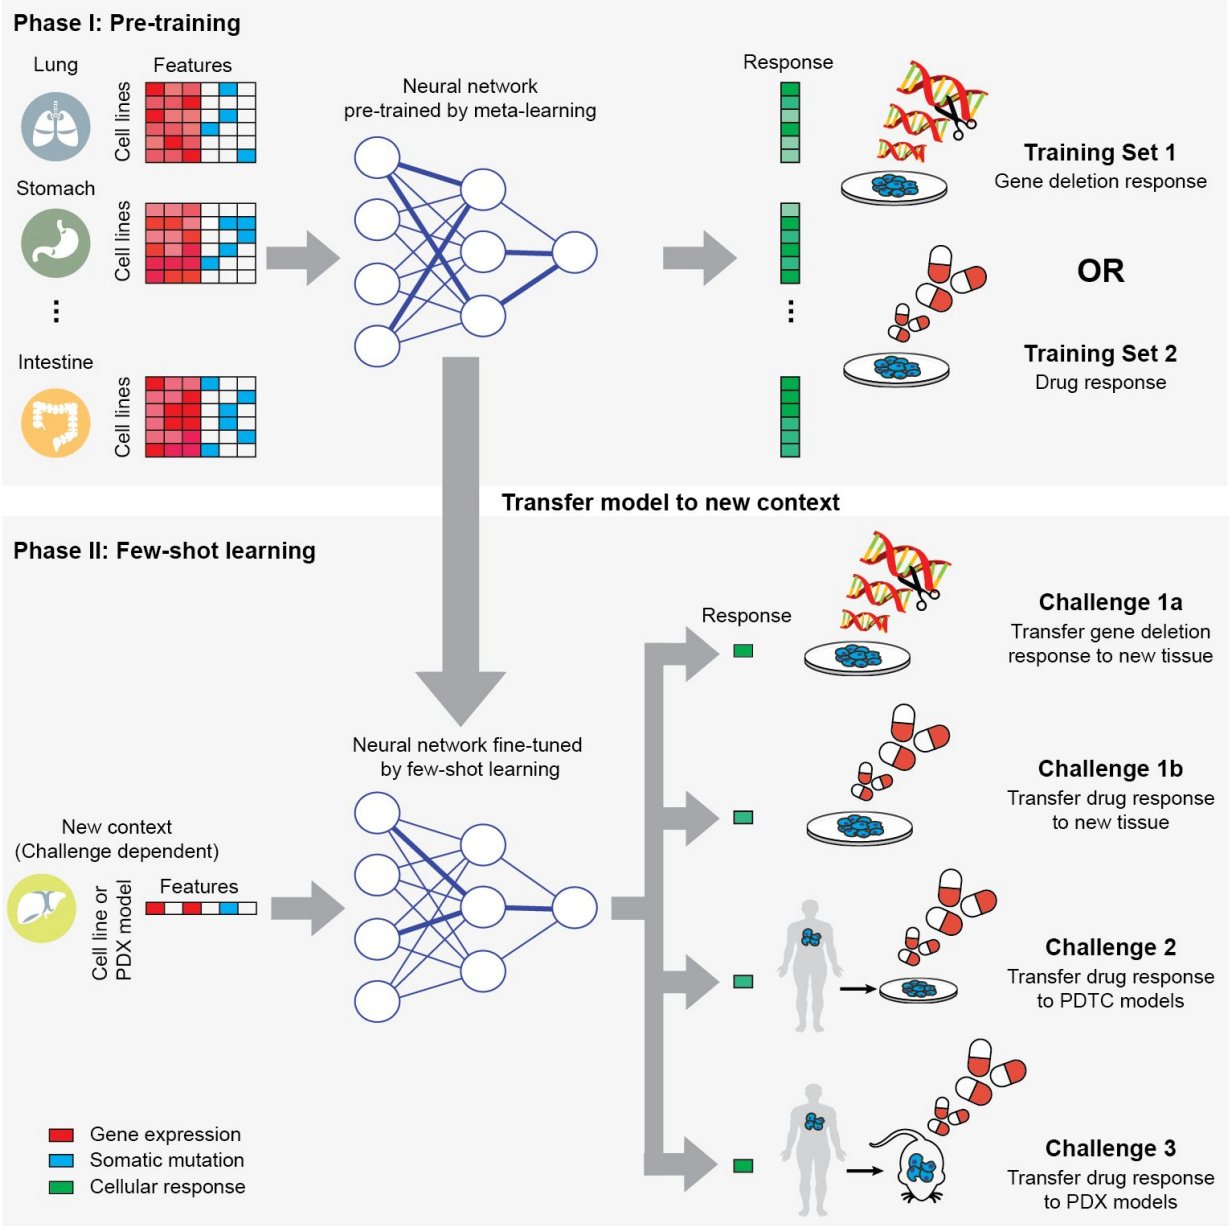


Figure 1



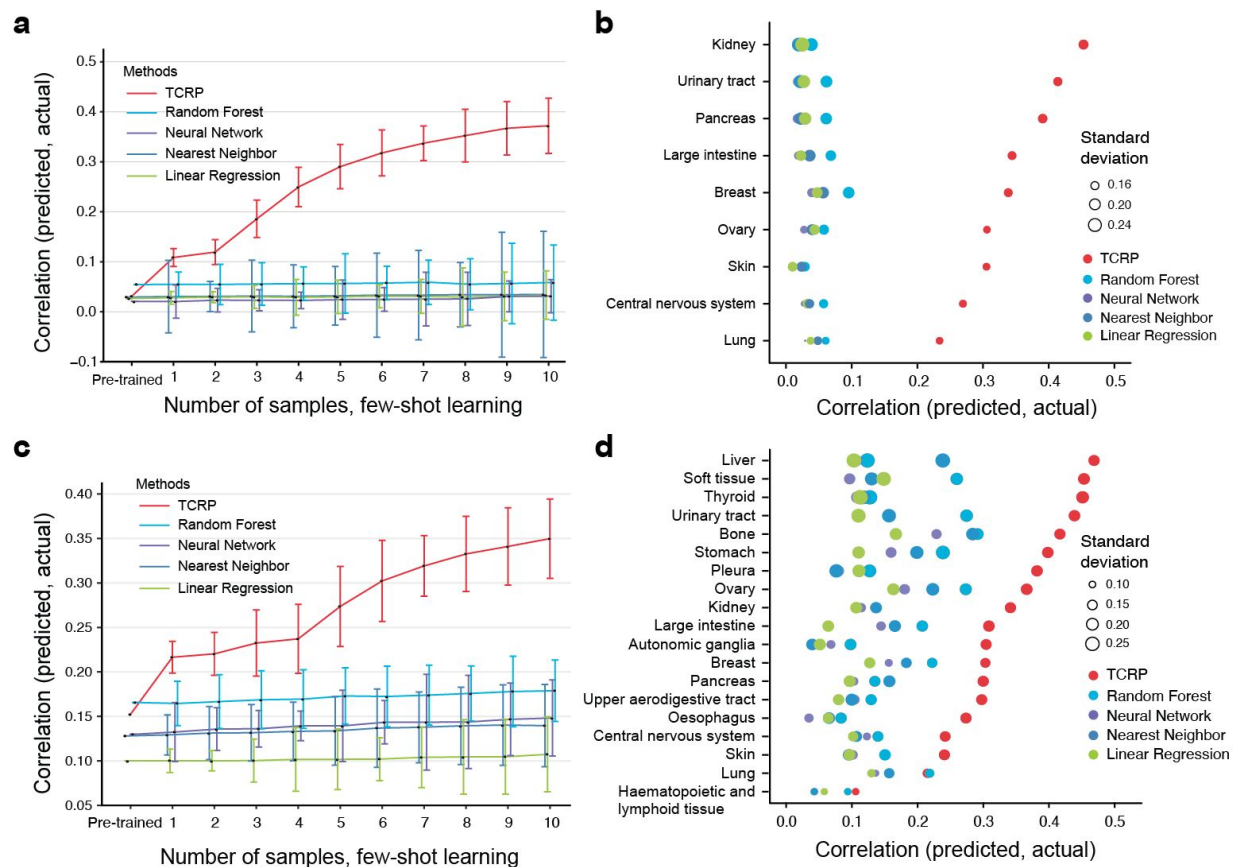


Figure 2

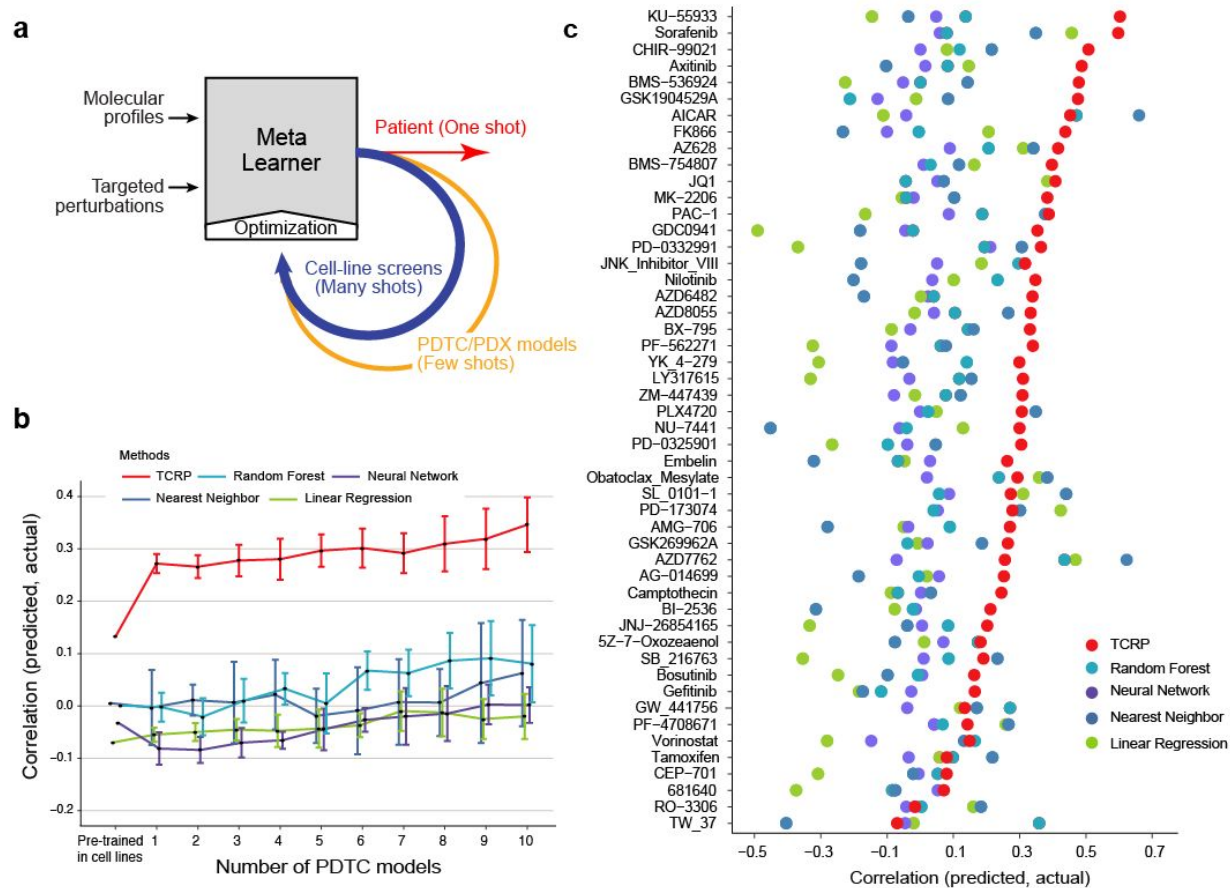


Figure 3

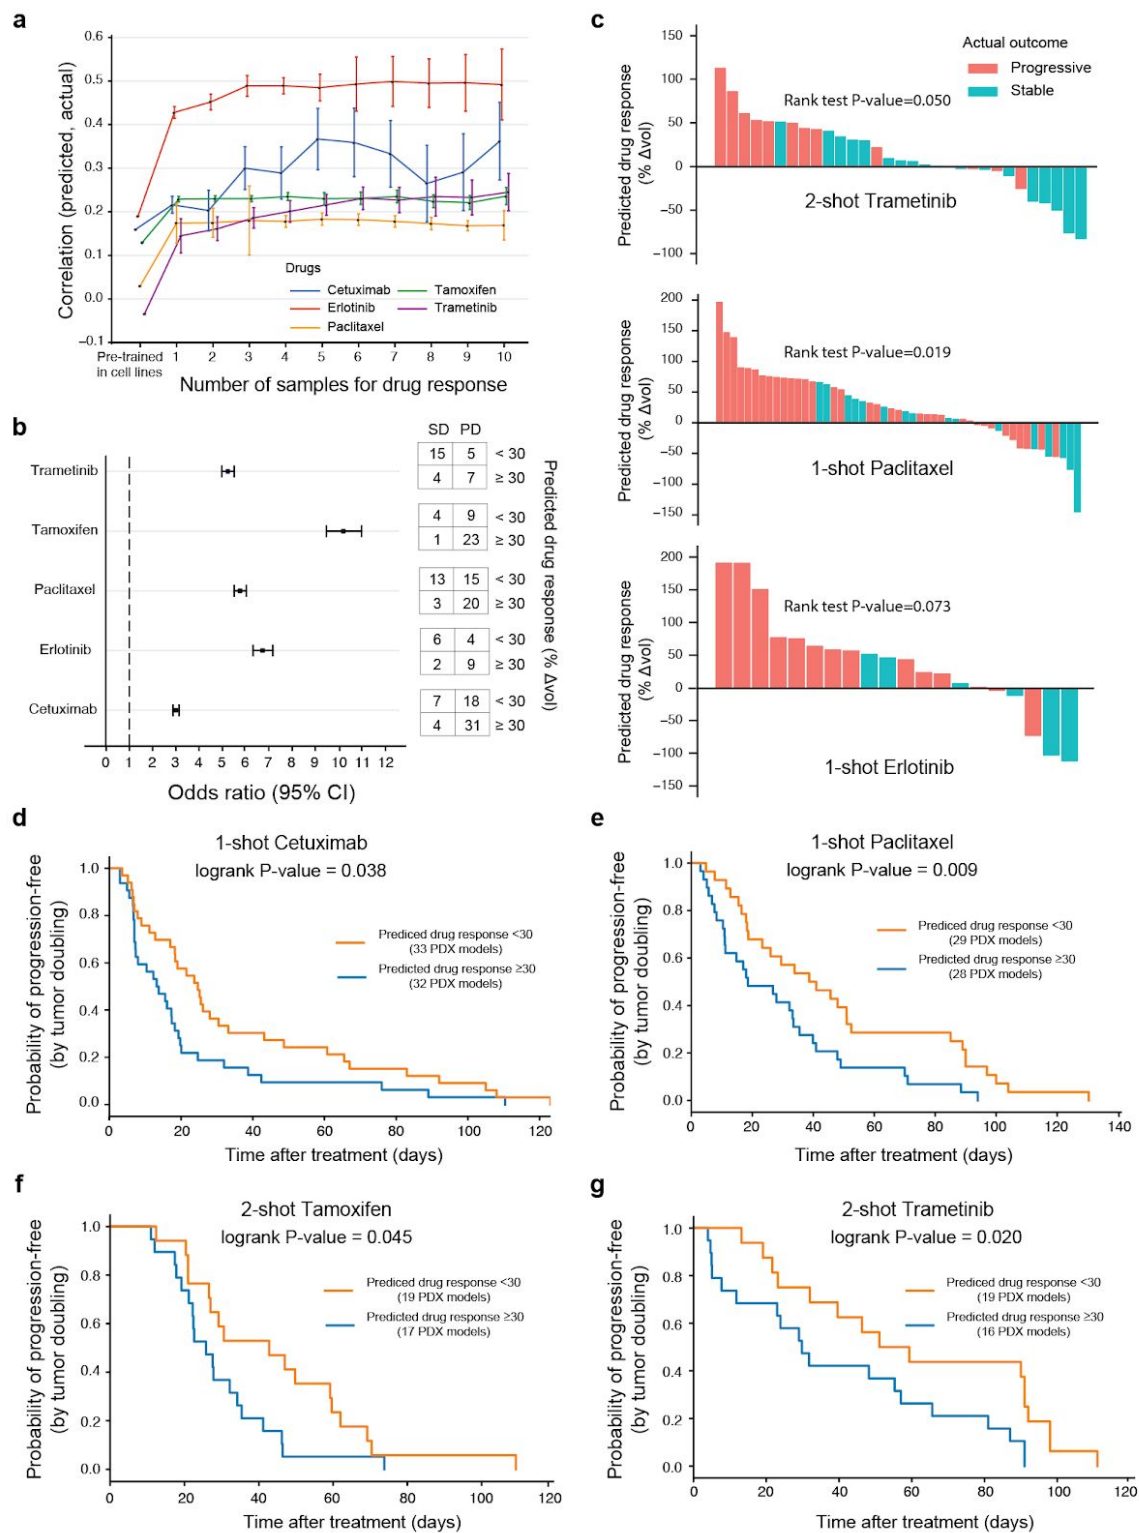


Figure 4

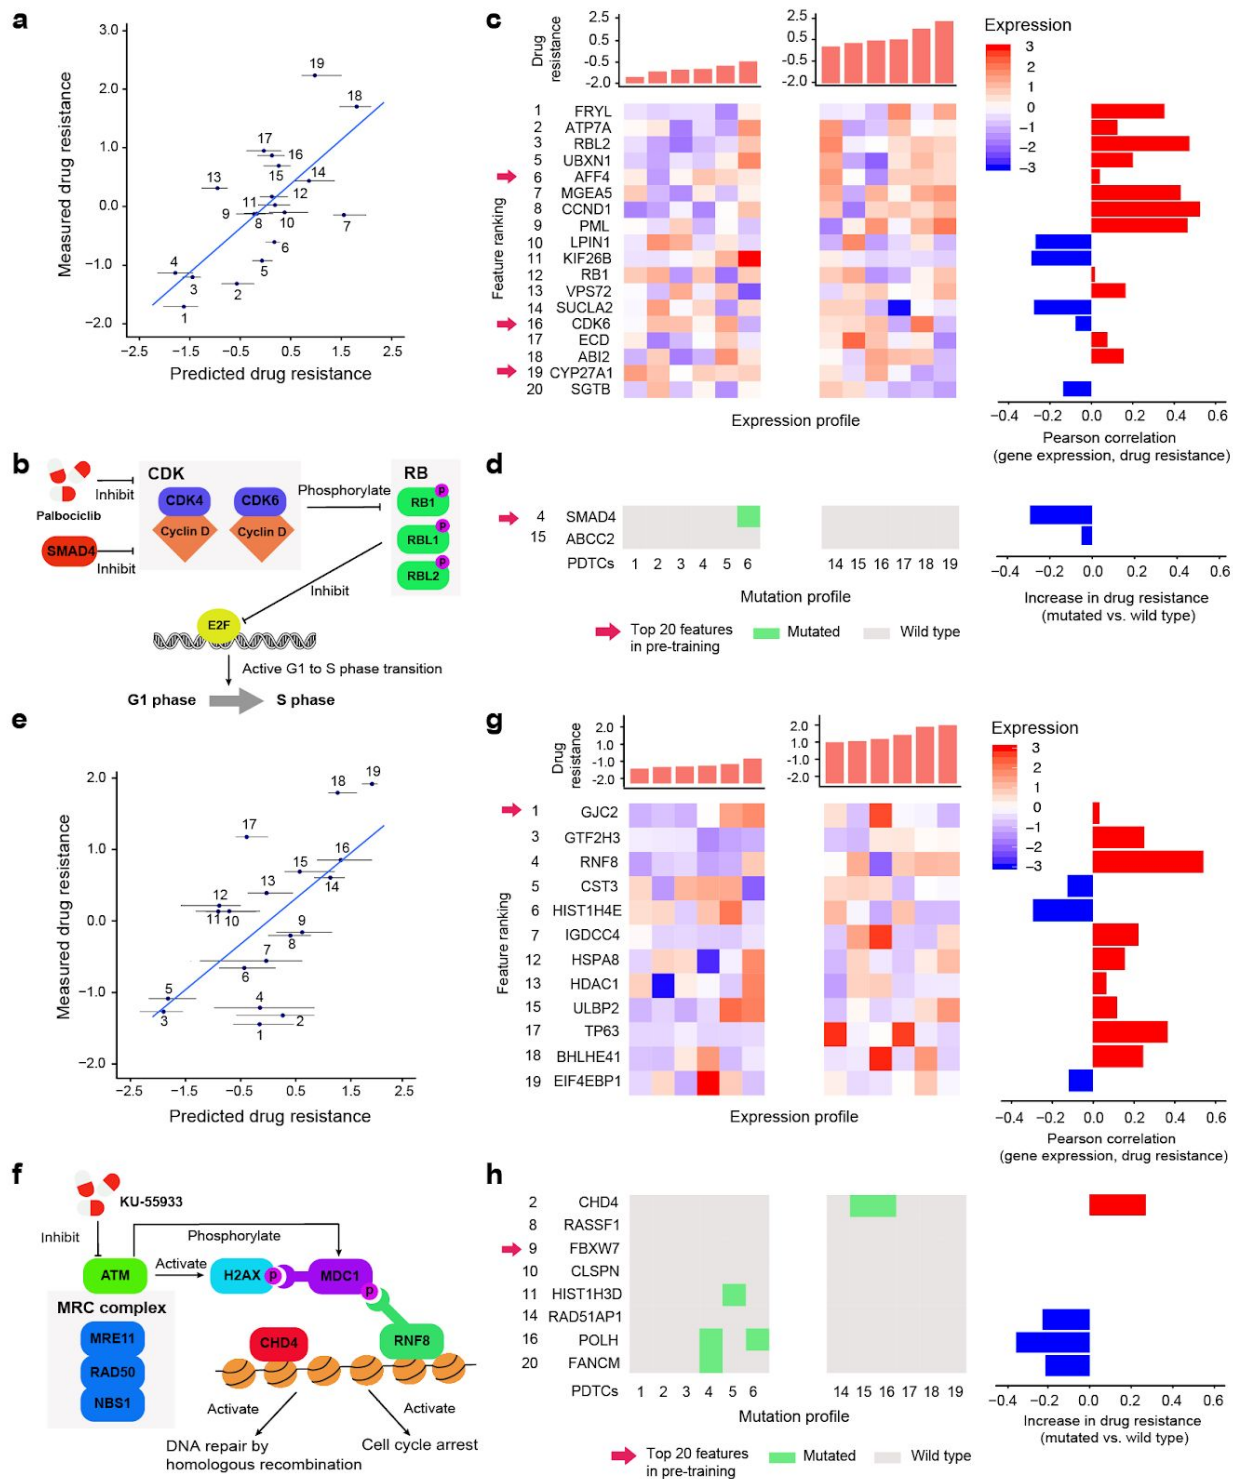


Figure 5

## Online Methods

### Challenge 1a

**Overview.** The first challenge was based on the Cancer Dependency Map (DepMap), which used CRISPR/Cas9 gene editing to disrupt nearly all (~17,700) human genes in each of 335 cancer cell lines (19 tissues), in each case measuring the relative cellular growth response<sup>1</sup>. The machine learning task was to use molecular features of each cell line to predict its growth response to the gene disruptions. Each gene disruption was considered as a separate learning task, in which cell lines represent learning samples. We studied 469 gene disruptions that had been reported by DepMap to have demonstrated ability to influence cellular growth, as evidenced by the presence of at least one cell line for which the response was at least six standard deviations away from the mean across cell lines<sup>1</sup>. **Even though there is a modest difference between the distribution of fitness values for all genes versus the selected genes (Supplementary Fig. 5a), we did not observe a strong relationship between the overall fitness effect of a gene knockout and model predictive performance (Supplementary Fig. 5b).**

**Task-specific features.** Features for learning were based on gene somatic mutations and expression levels for each cell line, as reported in the Cancer Cell Line Encyclopedia (CCLE) project<sup>2</sup> and downloaded from the DepMap website (<https://depmap.org/portal/download/>). For each learning task (CRISPR gene disruption, see above) we selected genes reported as having either a protein-protein interaction (PPI) or an mRNA co-expression relationship ( $|r| > 0.4$ ) with the disrupted gene. The PPI data were taken as the union of the InBioMap<sup>3</sup>, PathwayCommons<sup>4</sup> and CORUM<sup>5</sup> databases. The co-expression relationship is calculated over all the cell lines from the feature mRNA expression data. Such a feature selection strategy, based on the molecular network neighborhood of the disrupted gene, was similar to that adopted earlier by the DepMap project<sup>6</sup>. We further removed gene expression features for which the standard deviations fell into the lowest 10% over all genes and excluded genes with less than 10 somatic mutations across cell lines. The somatic mutations and mRNA expression levels of the remaining genes were applied to construct the input feature vector for each cell line.

**Labels.** Sample labels were taken as the growth response of a cell line to the CRISPR disruption of interest (see above) using the CERES-corrected single-gene disruption scores downloaded from DepMap (<https://depmap.org/portal/>). These scores are calculated by comparing the abundances of guide RNAs for the disrupted gene between the starting plasmid pool and the end of the CRISPR disruption experiment. The CERES method<sup>7</sup> then processes these scores by removing effects due to copy-number variation.

**Few-shot design.** For each gene disruption learning problem, the 19 tissues represented by DepMap cell lines were split such that 18 tissues were used in the pre-training phase and the remaining tissue was held out for the few-shot phase. To ensure sufficient samples for performance evaluation, this held-out tissue was selected from among the 9 tissues having  $\geq 15$  cell lines. In this few-shot phase, we randomly selected  $k$  cell lines as the few-shot samples to fine tune the model ( $k = [0..10]$ , plotted along the x axis of **Fig. 2a**) and used the remaining cell

lines as testing data. For each  $k$ , the selection of few-shot samples was random, thus we repeated this selection 20 times and reported the average and standard deviation of the prediction performance over these replicates (y axis of **Fig. 2a**).

### **Challenge 1b**

Overview. This challenge was based on the dataset collected by the Genomics of Drug Sensitivity in Cancer (GDSC) project<sup>8</sup>, which systematically tested the cellular growth responses elicited by a panel of 265 drugs applied to each of 1001 tumor cell lines (representing 30 tissues). The machine learning task was to use molecular features of each cell line to predict its growth response to a drug. Each drug was considered as a separate learning task, in which cell lines represent learning samples. We focused on 199 drugs for which the mechanism of action was at least partially characterized, i.e. with a documented protein target or pathway. Drug target and pathway information was obtained from Table S1G of the original GDSC 1000 paper<sup>8</sup>.

Task-specific features. Task-specific features were constructed for each drug by selecting genes having PPI or mRNA co-expression relationships ( $|r| > 0.4$ ) with the documented drug targets, with the PPI and mRNA co-expression networks defined as per Challenge 1a above. For drugs with multiple targets we included all PPI/co-expressed neighbors of these targets. As above, we further removed gene expression features for which the standard deviations fell into the lowest 10% over all genes and excluded genes with less than 10 somatic mutations across cell lines. Somatic mutations and mRNA expression levels of the remaining selected genes were applied to construct the input feature vector for each cell line.

Labels. Sample labels were taken as the growth response of a cell line to the drug of interest, using area under the dose-response curve (AUC) as the measure of drug response. All drug response data were downloaded from the GDSC 1000 website: [https://www.cancerrxgene.org/gdsc1000/GDSC1000\\_WebResources/](https://www.cancerrxgene.org/gdsc1000/GDSC1000_WebResources/).

Few-shot design. For each drug, the tissues were split such that one tissue was held out for the few-shot phase, and the remaining tissues were used in the pre-training phase. We required that the held-out tissue have data for  $\geq 15$  cell lines to provide sufficient samples for the few shot learning phase. A consequence of this requirement was that the number of held-out tissues differed from drug to drug, since drugs had a variable number of cell lines for which drug responses had been measured<sup>8</sup>. Similar to Challenge 1a, in the few-shot phase we randomly selected  $k$  cell lines from the held-out tissue as few-shot samples to fine tune the model ( $k = [0..10]$ , plotted along the x axis of **Fig. 2c**) and used the remaining cell lines as testing data. For each  $k$ , the selection of few-shot samples was random, thus we repeated this selection 20 times and reported average and standard deviation of prediction performance over all of these replicates (y axis of **Fig. 2c**).

### **Challenge 2**



**Overview.** In this second challenge, we pre-trained TCRP to predict drug responses in the GDSC dataset (see Challenge 1b) and then subjected this model to few-shot learning using a study of patient-derived tumor cells (PDTC)<sup>9</sup>. This study obtained 83 human breast tumor biopsies and, using mice as intermediary, established distinct human cell cultures from these tumors. Each of these human cell cultures was exposed to a panel of drugs, from which we considered the 50 drugs with known protein targets and for which cell-line responses had also been measured in the GDSC dataset. The machine learning task was to use the pre-trained model to predict the growth response of these PDTCs to each drug. Each drug was considered as a separate learning task, in which PDTCs represent learning samples.

**Features.** We considered gene expression and mutation features that had been characterized in both the PDTC and GDSC datasets. Both drug-specific features and mini cancer genome features were evaluated. Expression and somatic mutation data of the PDTC dataset were downloaded from the following URL:

[https://figshare.com/articles/Bruna\\_et\\_al\\_A\\_biobank\\_of\\_breast\\_cancer\\_explants\\_with\\_preserved\\_intra-tumor\\_heterogeneity\\_to\\_screen\\_anticancer\\_compounds\\_Cell\\_2016/2069274](https://figshare.com/articles/Bruna_et_al_A_biobank_of_breast_cancer_explants_with_preserved_intra-tumor_heterogeneity_to_screen_anticancer_compounds_Cell_2016/2069274).

**Labels.** For the PDTC responses we used AUC as the measure of drug response, similar to the GDSC dataset in Challenge 1b. These data were downloaded from:

[https://figshare.com/articles/Bruna\\_et\\_al\\_A\\_biobank\\_of\\_breast\\_cancer\\_explants\\_with\\_preserved\\_intra-tumor\\_heterogeneity\\_to\\_screen\\_anticancer\\_compounds\\_Cell\\_2016/2069274](https://figshare.com/articles/Bruna_et_al_A_biobank_of_breast_cancer_explants_with_preserved_intra-tumor_heterogeneity_to_screen_anticancer_compounds_Cell_2016/2069274).

**Few-shot design.** In the few-shot learning phase, we randomly selected  $k$  PDTCs as the few-shot samples to fine tune the model ( $k = [0..10]$ , plotted along the x axis of **Fig. 3b**) and used the remaining cell lines as testing data. For each  $k$ , the selection of few-shot samples was random, thus we repeated this selection 20 times and reported the average and standard deviation of prediction performance over all of these replicates (y axis of **Fig. 3b**).

### Challenge 3

In this third challenge, we pre-trained TCRP to predict drug responses in the GDSC dataset (see Challenge 1b) and then used few-shot learning to transfer it to make drug response predictions in a study of patient-derived tumor xenografts (PDX)<sup>10</sup>. This study created a large collection of mouse xenografts of human tumor biopsies, all characterized for tumor somatic mutations and mRNA expression levels. PDXs were exposed to a panel of drug treatments (one PDX per animal per treatment) during which *in-vivo* tumor growth was measured. The machine learning task was to use the pre-trained TCRP to predict tumor growth *in vivo*. In particular, we used data for 228 PDX mouse models, where each model was exposed to one of the five drugs on which TCRP had been trained in cell lines (Cetuximab, Erlotinib, Paclitaxel, Tamoxifen, and Trametinib).

**Mini cancer genome features.** Expression and somatic mutation data for all PDX samples were downloaded from Supplementary Table 1 of the original paper<sup>10</sup>. Most drugs in the PDX dataset do not have known drug targets, a requirement for feature selection in previous challenges (see

above). Therefore, we adopted an alternative means of selecting features which does not require knowledge of drug mechanism of action, as introduced in recent work<sup>11</sup>. These features were based on the “mini cancer genome panel”, a set of known cancer related genes collected by the Center for Personalized Cancer Treatment (CPCT, The Netherlands)<sup>12</sup>. From this panel, we first removed gene expression and mutation features that had not been characterized in both the PDX and GDSC datasets. Second, we removed gene expression features for which the standard deviations fell into the lowest 10% over all genes in GDSC, and we removed gene mutation features with less than 10 somatic mutations across GDSC cell lines. The somatic mutations and mRNA expression levels of the remaining selected genes were applied to construct the input feature vector for each cell line. In this scenario, all the learning tasks (drugs) shared the same feature set.

**Labels.** PDX drug response was measured by the minimum change in tumor volume in comparison to baseline, over the period from 10 days post-treatment until completion of the study ( $\Delta vol$  in the main text). This measure captures the speed, strength and durability of the *in vivo* response; all values were downloaded from Supplementary Table 1 of the original paper<sup>10</sup>. When comparing TCRP predictions to  $\Delta vol$  measurements, both were normalized to a standard normal distribution to translate between the two (i.e. z-score).

**Few-shot design.** In the few-shot learning phase, we randomly selected  $k$  PDXs as the few-shot samples to fine tune the model ( $k = [0..10]$ , plotted along the x axis of **Fig. 4a**) and used the remaining PDX samples as testing data. For each  $k$ , the selection of few-shot samples was random, thus we repeated this selection 20 times and reported average and standard deviation of prediction performance over all of these replicates (y axis of **Fig. 4a**).

### TCRP neural network model

We trained a multi-layer neural network model to predict the phenotype of a tumor sample using its molecular features. For each sample  $i$ , the output of the  $j + 1$  th layer  $h_i^{(j+1)}$  is defined as a non-linear function of the output of  $j$  th layer  $h_i^{(j)}$  as follows,

$$h_i^{(j+1)} = Relu ( Linear( h_i^{(j)} ) ) \quad (1)$$

where  $Linear( h_i^{(j)} )$  is a linear function of  $h_i^{(j)}$  defined as  $W^{(j)} * h_i^{(j)} + b^{(j)}$ .  $W^{(j)}$  is the weight matrix and  $b^{(j)}$  is the bias vector.  $Relu$  is the rectified linear activation function<sup>13</sup> which thresholds values less than 0 to exactly 0. The first layer  $h_i^{(1)}$  is the input molecular feature of sample  $i$  and the last layer  $h_i^{(N)}$  acts as its final prediction  $\hat{p}_i(\theta)$ , where  $\theta$  are parameters containing  $W^{(j)}$  and  $b^{(j)}$  from all the linear layers. For each machine learning task, we scan all combinations of layers={1,2} and hidden neurons={5,10,15,20} and determine the architecture of the neural network by cross validation. All parameters are trained by minimizing the mean square error function  $L$  which is a function of sample set  $C$  and parameters  $\theta$ :



$$L(C, \theta) = \frac{1}{M} \sum_{c_i \in C} (p_i - \hat{p}_i(\theta))^2 \quad (2)$$

where  $p_i$  is the measured label for sample  $i$  and  $M$  is the number of samples in  $C$ .

### Model pre-training phase

In the pre-training phase, the aim is to train a neural network model that can quickly adapt to a new learning task with only a few additional training samples. The rationale is to acquire prior knowledge from a set of related tasks where training samples are abundant. In this work, we adopted an established computational framework called the Model Agnostic Meta-Learning (MAML) algorithm<sup>14</sup>. Meta-learning approaches like MAML seek to identify universal knowledge across multiple conditions and then to transfer this knowledge to make robust predictions in a new condition. In recent studies, the MAML technique has shown superior performance in comparison to other meta-learning frameworks<sup>14</sup>, and it is flexible and model agnostic such that it can be applied to any gradient-based learning algorithm.

For each training iteration, we first sample a subset  $S_i$  of 12 tissue types from the pool  $S$  of all types available.  $S_i$  is then randomly partitioned into 2 non-overlapping sets of 6 cell lines  $T$  and 6 cell lines  $V$ . A loss function adapted from **Eq. (2)** is defined as follows with respect to  $S$ :

$$E_{S_i \in S} [ E_{<T, V> \in S_i} [ L(V, \theta - \alpha \frac{\partial L(T, \theta)}{\partial \theta}) ] ] \quad (3)$$

Here  $L$  is a mean square error function with respect to  $V$ . The second argument of the loss function is a one-step gradient descent which seeks a better regression loss for cell line set  $T$ . We then optimize **(3)** using the gradient descent algorithm ADAM<sup>15</sup>. Note that using gradient descent requires calculation of a second order gradient of loss function  $L$ . The intuition is that for each training iteration of minimizing **(3)**, we seek parameters  $\theta$  that can achieve a smaller regression loss on cell line set  $V$  after performing one iteration of gradient descent on a distinct cell-line set  $T$ . A total of 200 training iterations were performed, sampling different  $S_i$ , with each  $S_i$  including 20 partitions.

### Few-shot learning phase

In the second training phase, we observe a task  $Q$  with only a few training samples (e.g. cell lines, PDTs or PDX models). We perform only one iteration of gradient descent to achieve  $\theta_{few-shot}$  suitable for the new task (e.g. new tissue or mouse models):

$$\theta_{few-shot} = \theta_{pre-training} - \alpha \frac{\partial L(Q, \theta)}{\partial \theta} \Big|_{\theta=\theta_{pre-training}} \quad (4)$$

Here  $\theta_{pre-training}$  is the TCRP model trained in the pre-training phase. In theory, one can perform multiple iterations of gradient descent using **(4)** until convergence. However, one of the unsolved problems in the field of meta-learning is that the few-shot model can be easily over-fit

on a new task given its very few samples. Therefore, we chose to only update parameters once. Note that  $\alpha$  in (3) and (4) refers to the same hyperparameter. The structure of the neural network was defined as in Eq. (1).

### **Nested cross validation**

The appropriate architecture of a neural network is dependent on the particular problem and datasets. For drug prediction problems (Challenge 1b, 2, 3), all hyperparameters, including mini-batch size and the size of  $T$  and  $V$ , were determined by the technique of nested cross validation as previously described<sup>16</sup>. For Challenge 1a, we used regular cross-validation due to the greater number of prediction tasks.

### **Interpreting TCRP model predictions**

We used the framework of Local Interpretable Model-Agnostic Explanations (LIME)<sup>17</sup> to generate locally faithful explanations for the TCRP neural network model. LIME works by taking the feature vector of a query sample of interest and perturbing it randomly, resulting in many perturbed samples around this query. Subsequently, it trains a much simpler interpretable model on this perturbed neighborhood (**Supplementary Fig. 6**). In this way, LIME is able to select important features specific for sample  $i$ , which is the major difference from conventional feature selection methods which act globally over all samples, not locally to a sample of interest. More formally, for the molecular feature vector  $f_i$  of each sample  $i$ , we generated  $N$  ( $=10,000$ ) perturbed samples. Each of these perturbed samples  $j$  was created by adding to the original features independent Gaussian noise with mean 0 and standard deviation 1. For each perturbed sample, we made a prediction  $g_{ij}$  using the TCRP neural network. A second simpler model, regularized linear regression, was then trained to fit the perturbed samples to their corresponding neural network predictions  $\{g_{ij}\}$ . Empirically, we applied both Elastic Net<sup>18</sup> and Lasso<sup>19</sup> regularization methods with different sparsity parameters ( $=\{0.1, 0.01, 0.001, 0.0001\}$ ). The final ranking of features was averaged from the rankings produced by Elastic Net and Lasso over all sparsity parameters and over all tested samples.

LIME was chosen over alternative model interpretation techniques such as layer-wise relevance propagation<sup>20</sup>, as these other techniques do not generate sample-specific explanations. LIME is an approximation of gradient-based methods<sup>21</sup> and could be used interchangeably with those methods in our work.

### **Implementation details of competing methods**

We used the Python package “scikit-learn” (<http://scikit-learn.org/stable/index.html>) to implement four conventional machine learning methods: random forests, conventional neural networks, K nearest neighbors, and linear regression, as follows:

Random forests. For Random Forests, we chose the max depth for each of the learning tasks based on five-fold cross validation.

Conventional neural networks. Conventional neural network models were implemented using the PyTorch library (<https://pytorch.org/>), selecting the number of hidden neurons ( $=\{5, 10, 20, 30, 40, 50, 100\}$ ), layers ( $=\{1, 2\}$ ) and learning rates ( $=\{0.1, 0.01, 0.001\}$ ) based on five-fold nested cross validation. For each machine learning task (e.g. drugs and gene perturbations), there are approximately (or fewer than) 1000 cell line examples (plus <20 PDTC/PDX models in some cases); thus, the data do not support a very deep neural network architecture with many parameters. Therefore, we focused on exploration of small neural network architectures in this study. The number of hidden layers ( $=\{1, 2\}$ ) and the number of hidden neurons ( $=\{5, 10, 15, 20\}$ ) of the neural network were also determined by cross validation. We implemented the algorithm using the PyTorch library (<https://pytorch.org/>) running on Tesla K20 GPUs. The non-linear transformation was the same as **Eq. (1)** and optimized using ADAM<sup>15</sup>. Notice that both TCRP and this baseline method rely on a neural network model; however, the two models are trained in different ways and with potentially different network architectures (# hidden, layers) due to separate cross-validation processes.

K nearest neighbors (KNN). For the KNN algorithm, to evaluate the accuracy of a sample  $i$  in the training data, we ruled out sample  $i$  when making its prediction. Otherwise, KNN will achieve a zero prediction error on the training set. The best ‘K’ for KNN was selected using five-fold cross validation.

Linear regression. For the final conventional method, we implemented linear regression with the regular least squares loss function and without regularization.

## References (Online Methods)

1. Meyers, R. M. *et al.* Computational correction of copy number effect improves specificity of CRISPR-Cas9 essentiality screens in cancer cells. *Nat. Genet.* **49**, 1779–1784 (2017).
2. Ghandi, M. *et al.* Next-generation characterization of the Cancer Cell Line Encyclopedia. *Nature* **569**, 503–508 (2019).
3. Li, T. *et al.* A scored human protein-protein interaction network to catalyze genomic interpretation. *Nat. Methods* **14**, 61–64 (2017).
4. Cerami, E. G. *et al.* Pathway Commons, a web resource for biological pathway data. *Nucleic Acids Res.* **39**, D685–90 (2011).
5. Giurgiu, M. *et al.* CORUM: the comprehensive resource of mammalian protein complexes—2019. *Nucleic Acids Res.* **47**, D559–D563 (2019).
6. Tsherniak, A. *et al.* Defining a Cancer Dependency Map. *Cell* **170**, 564–576.e16 (2017).
7. Meyers, R. M. *et al.* Computational correction of copy-number effect improves specificity of CRISPR-Cas9 essentiality screens in cancer cells. doi:10.1101/160861.
8. Iorio, F. *et al.* A Landscape of Pharmacogenomic Interactions in Cancer. *Cell* **166**, 740–754 (2016).
9. Bruna, A. *et al.* A Biobank of Breast Cancer Explants with Preserved Intra-tumor Heterogeneity to Screen Anticancer Compounds. *Cell* **167**, 260–274.e22 (2016).
10. Gao, H. *et al.* High-throughput screening using patient-derived tumor xenografts to predict clinical trial drug response. *Nat. Med.* **21**, 1318–1325 (2015).
11. Kim, Y., Bismeyer, T., Zwart, W., Wessels, L. F. A. & Vis, D. J. Genomic data integration by

- WON-PARAFAC identifies interpretable factors for predicting drug-sensitivity in vivo. *Nat. Commun.* **10**, 5034 (2019).
12. Harakalova, M. *et al.* Multiplexed array-based and in-solution genomic enrichment for flexible and cost-effective targeted next-generation sequencing. *Nat. Protoc.* **6**, 1870–1886 (2011).
  13. Glorot, X., Bordes, A. & Bengio, Y. Deep Sparse Rectifier Neural Networks. *Aistats* (2011).
  14. Finn, C., Abbeel, P. & Levine, S. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *arXiv [cs.LG]* (2017).
  15. Kingma, D. & Ba, J. Adam: A Method for Stochastic Optimization. *arXiv [cs.LG]* (2014).
  16. Baumann, D. & Baumann, K. Reliable estimation of prediction errors for QSAR models under model uncertainty using double cross-validation. *J. Cheminform.* **6**, 47 (2014).
  17. Ribeiro, M. T., Singh, S. & Guestrin, C. Why should i trust you?: Explaining the predictions of any classifier. in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 1135–1144 (ACM, 2016).
  18. Zou, H. & Hastie, T. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Series B Stat. Methodol.* **67**, 301–320 (2005).
  19. Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *J. R. Stat. Soc. Series B Stat. Methodol.* **58**, 267–288 (1996).
  20. Binder, A., Montavon, G., Lapuschkin, S., Müller, K.-R. & Samek, W. Layer-Wise Relevance Propagation for Neural Networks with Local Renormalization Layers. in *Artificial Neural Networks and Machine Learning – ICANN 2016* 63–71 (Springer International Publishing, 2016).
  21. Simonyan, K., Vedaldi, A. & Zisserman, A. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *arXiv [cs.CV]* (2013).